Toward a Direct Measure of Video Quality Perception Using EEG

Simon Scholler, Sebastian Bosse, Matthias Sebastian Treder, Benjamin Blankertz, Gabriel Curio,
Klaus-Robert Müller, and Thomas Wiegand, *Fellow, IEEE*

Abstract—An approach to the direct measurement of perception of video quality change using electroencephalography (EEG) is presented. Subjects viewed 8-s video clips while their brain activity was registered using EEG. The video signal was either uncompressed at full length or changed from uncompressed to a lower quality level at a random time point. The distortions were introduced by a hybrid video codec. Subjects had to indicate whether they had perceived a quality change. In response to a quality change, a positive voltage change in EEG (the so-called P3 component) was observed at latency of about 400–600 ms for all subjects. The voltage change positively correlated with the magnitude of the video quality change, substantiating the P3 component as a graded neural index of the perception of video quality change within the presented paradigm. By applying machine learning techniques, we could classify on a single-trial basis whether a subject perceived a quality change. Interestingly, some video clips wherein changes were missed (i.e., not reported) by the subject were classified as quality changes, suggesting that the brain detected a change, although the subject did not press a button. In conclusion, abrupt changes of video quality give rise to specific components in the EEG that can be detected on a single-trial basis. Potentially, a neurotechnological approach to video assessment could lead to a more objective quantification of quality change detection, overcoming the limitations of subjective approaches (such as subjective bias and the requirement of an overt response). Furthermore, it allows for real-time applications wherein the brain response to a video clip is monitored while it is being viewed.

Index Terms—Electroencephalography (EEG), perception, video coding, video quality.

S. Scholler, M. S. Treder, and B. Blankertz are with the Machine Learning Laboratory, Berlin Institute of Technology, 10587 Berlin, Germany (e-mail: simon.scholler@googlemail.com; matthias.treder@tu-berlin.de; benjamin.blankertz@tu-berlin.de).

K.-R. Müller is with the Machine Learning Laboratory, Berlin Institute of Technology, 10587 Berlin, Germany, and also with the Department of Brain and Cognitive Engineering, Korea University, Anam-dong, Seongbuk-gu, Seoul 136-713, Korea (e-mail: klaus-robert.mueller@tu-berlin.de).

S. Bosse is with the Fraunhofer Heinrich Hertz Institute (Fraunhofer HHI), 10587 Berlin, Germany (e-mail: sebastian.bosse@hhi.fraunhofer.de).

G. Curio is with the Department of Neurology and Clinical Neurophysiology, Charité, 10117 Berlin, Germany (e-mail: gabriel.curio@charite.de).

T. Wiegand is with the Department of Image Processing, Fraunhofer HHI, and with the Image Communication Laboratory, Berlin Institute of Technology, 10587 Berlin, Germany (e-mail: twiegand@ieee.org).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIP.2012.2187672

I. INTRODUCTION

VIDEO signals are typically intended to be viewed by humans. For their transmission at bit rates that is suitable for today's channels or storage devices, these signals are digitized and potentially compressed. With an increasing reduction in bit rates, the compression algorithm starts to introduce distortions that are visible to humans. The measurement of these distortions is essential in most video transmission tasks, e.g., for controlling the tradeoff between bit rate and distortion (R–D) or for assessing the visual quality of a transmitted video signal.

One approach toward measuring subjective distortion has been the modeling of the human visual system [1], [2]. The basic idea for obtaining such a model is to measure transfer functions given various temporal and/or spatial stimuli and then to combine these measurements into a more complete model. Such approaches led to a profound understanding of the limitations of the human visual system, e.g., the spatiotemporal limits of perception.

However, the question at hand of how to quantify visual distortion remained unsatisfyingly answered by these approaches. One reason is that various top-down mechanisms (which are difficult to model) influence the sensitivity of humans to distortions. The world, as we see it, is based on a number of inferences that are related to the visual input such as motion or depth, or content- and context-related inferences. These inferences are a major part of (active) perception and constitute the way we perceive. A model that does not dwell on these top-down processes will remain incomplete.

Hence, a precise model for the subjective perception of distortion is not available. The most common current approach to quantifying subjective distortion still is a judgment experiment: a human observer is presented a stimulus and gives an overt response. Such subjective testing for visual quality assessment has been formalized in [3] for television applications and in [4] for multimedia applications. The typical procedure in any of these recommendations is that the subject has to rank the quality of a set of test videos.

This may be done with or without showing a reference video. These subjective tests are widely used in practice and deliver quality assessments for video signals when averaged over many subjects. They share the drawback that ratings are highly variable across subjects. Further, these ratings given by the human observer are the result of a conscious process, which may be inferred by various aspects and which is prone to be affected by subjective factors (e.g., bias, expectation, and strategies).

A potential solution, proposed here, is to directly monitor brain activity during the observation of video clips using electroencephalography (EEG). For the first time, measurements of

brain activity are used to quantify the perception of a human observer when being shown a change in video quality. Our approach capitalizes on the P3 component (or P300 component), which is a large positivity that is usually observed 300–500 ms after a rare and/or significant event. Its amplitude peaks over the central–parietal brain regions [5]. The oddball paradigm is the classical paradigm in which a P3 response is elicited: Frequent and infrequent stimuli are shown, and a P3 component is found in response to infrequent stimuli. The P3 reflects cognitive processing and is observed independent of the sensory modality. In contrast to a visually evoked potential (VEP), it is not directly linked to sensory processing. Despite the high amount of noise in the EEG,[1] single-trial detection of the P3 component has been demonstrated in applications such as visual [7], [8] and auditory [9], [10] brain–computer interfaces (BCIs) and in audio quality assessment [11]. For reviews on BCIs and BCI technology, refer to [12]–[14].

This paper is organized as follows. In Section II, the viewing experiment, including the stimuli and EEG, is described. The measured EEG signals and their analysis are described in Section II-C. Results are presented in Section III followed by a discussion on future directions.

II. VIEWING EXPERIMENT

Our incentive was to investigate whether the (possibly subconscious) perception of video codec artifacts can be measured with EEG. To this end, an experiment was conducted in which subjects watched short video clips, some of which featured a sudden quality change from high quality to a lower quality level during the video (see Section II-B). In order to make the measurement independent of the image statistics at the actual gaze position at the time of the quality change, the stimulus material that was chosen is spatially and temporally roughly homogeneous.

A. Stimulus Material

In order to obtain more control over the properties of the video stimuli than would be possible with real-world video sequences, the stimuli were synthetically generated. By using video material without semantically meaningful content and the absence of salient objects, influences due to high-level image understanding were minimized. Furthermore, by using homogeneous stimuli, we assumed the exact position of the eye gaze to be of little effect on the experimental data. On the other hand, the video material should not be too abstract but contain simple real-world textures and motion.

In order to meet these requirements, video sequences were generated based on a synthetic image of a textured checkerboard, as shown in Fig. 1. The image was deformed over time by simulating a swaying water surface on top of the checkerboard. The deformation was calculated by solving the 2-D wave function $(1/c^2)(\partial^2 u/\partial t^2) = (\partial^2 u/\partial x^2) + (\partial^2 u/\partial y^2) - p(\partial u/\partial t)$ (where v is speed of wave, p is the decay of wave, t is the time, and x and y are the image coordinates) iteratively using a finite-difference approach [15]. The deformation $u(x, y, t)$ was applied to the image by convolution, resulting in a smooth


Fig. 1.   Time course of the video sequence. For 1 s, a fixation point is shown. For at least 2 s, the undistorted video is presented. At a random time point (uniformly distributed between 2 and 6 s), the video quality drops instantaneously. On the right, undistorted and distorted (highest quality loss QC10) frames are shown exemplarily.

and modest movement. Reflections have not been taken into account.

The sequence was generated in a resolution of 832 by 480 pixels, with a frame rate of 60 fps and a duration of 8 s.

From this undistorted sequence, 500 new sequences have been generated by introducing 1 out of 10 magnitude values of quality drops in one out of 50 time instances into the undistorted sequence. The time instances of the quality change were randomly chosen and uniformly distributed between 2 and 6 s, i.e., between the 121th and 360th frames. This assured that the subject cannot predict the time point of a quality change. The magnitude of the quality change was controlled by the quantization parameter of the video coder (see the following). These 500 sequences with one quality change each and the undistorted sequences served as stimuli. The different stimuli have been labeled with QC1–QC10 and QC0, where QC1 denotes the most subtle quality change, QC10 denotes the strongest quality change, and QC0 is the undistorted sequence.

Prior to each presentation of a stimulus, a central fixation point was shown for 1 s. Fig. 1 shows the time course of the stimuli and gives an example of the appearance of a stimulus over time.

The quality loss considered in the experiments was induced by lossy compression of the synthesized video sequence. The used video coder is a state-of-the-art hybrid motion-compensated block-based coder [16]. It is architecturally similar to the emerging HEVC standard [17], offering a flexible quadtree structure for prediction and transform. Statistical redundancies are exploited by blockwise temporal and spatial prediction. The residual signal is transformed blockwise, and coefficients are quantized in the transform domain. Coding artifacts, which are perceived by the human observer as a loss of visual quality, are introduced by the quantization.

The encoder decides about the best representation of the video to be coded by minimizing the Lagrangian R–D cost functional $J = D + \lambda R$ on block basis with D being the distortion and R being the frame rate.

We configured the coder with a maximum prediction block size of 64×64 and a maximum quadtree depth of 4, equivalent

---

[1]In the simulations in [6], only half of the surface scalp potential comes from sources within a 3-cm radius around the electrode.

Fig. 2. Experimental setup.

TABLE I
INDIVIDUAL STIMULUS LEVELS FOR ALL SUBJECTS. FOR THREE SUBJECTS, STIMULUS LEVEL QC-I WAS OMITTED DUE TO TIME CONSTRAINTS

Subject	undistorted	QC-I	QC-II	QC-III	QC-IV	QCmax
S1	QC0	QC4	QC5	QC6	QC7	QC10
S2	QC0	QC4	QC5	QC6	QC7	QC10
S3	QC0	-	QC5	QC6	QC7	QC10
S4	QC0	-	QC5	QC6	QC7	QC10
S5	QC0	QC4	QC5	QC6	QC7	QC10
S6	QC0	QC3	QC4	QC5	QC6	QC10
S7	QC0	QC3	QC4	QC5	QC6	QC10
S8	QC0	-	QC5	QC6	QC7	QC10
S9	QC0	QC4	QC5	QC6	QC7	QC10

to a minimum prediction. We used an IPPP prediction structure and an intra period longer than the coded sequence so that only the first frame is coded intra-only.

In order to obtain sequences of different visual qualities, we changed the quantization parameter QP $\in \{32, 34, 36, 38, 40, 42, 44, 46, 48, 51\}$. The quantization parameter QP corresponds to the quantization step size $\Delta_{i,j}$ with $\Delta_{i,j} = 2^{\text{QP}/6} \times m_{i,j}$, where $m_{i,j}$ is a multiplier depending on the transform matrix and the position of the coefficient. However, the perceived quality loss in the generation of the stimuli is controlled by and is monotone with the quantization parameter QP.

B. Experimental Design

Videos were shown on a 30″ screen (Dell UltraSharp 3008WFP) with a native resolution of 2560 × 1600 pixels (see Fig. 2). The screen was normalized according to the specifications in [3]. The video resolution was 832 × 480 pixels or 20.8 × 12 cm, which corresponds to 24 × 14° of visual angle. The viewing distance was 48 cm (four times the video height) in compliance with the specifications in [3].

Nine subjects (three females and six males in the age group of 20–32) participated in the experiment. They were naive with respect to the purpose of the experiment and they had not participated in a video assessment study before. All had normal or corrected-to-normal vision. Subjects sat in front of the screen in a dimly lit room.

Following a visual acuity test and a general introduction to the experiment, a pretest of 100 trials with ten different stimulus levels was conducted. In each trial, the subjects first had to fixate a central fixation point for 1 s. Subsequently, the video was shown for a duration of 8 s. In 83% of the video sequences, a quality change occurred in the interval of 2–6 s; in the other 17% of sequences, no change occurred (QC0). The stimulus levels of the experiment (QC1–QC10) are the magnitude of the quality change (c.f. Section II-A). At the end of the video, subjects indicated via a button press whether they perceived a change in quality at any point during the video (yes/no task). Based on the behavioral data from the pretest, four subject-specific stimulus levels were chosen that targeted the slope of the psychometric functions, i.e., stimulus levels around the threshold of perception. If QCx is the stimulus level closest to the perception threshold, then QCx−2, QCx−1, QCx, QCx+1 were the

selected stimulus levels.[2] In addition, the undistorted condition without a quality change (QC0) and a maximal quality change condition (QCmax) were included in order to have clearly perceived and not perceived trials, respectively, as a reference (see Table I).

The main experiment consisted of 600 trials (100 trials per stimulus level that are randomly shuffled). These trials were subdivided into eight blocks of 75 trials each. A block lasted about 15 min followed by a few minutes break. Including cap preparation, the experiment lasted about 4 h.

C. EEG Data

Brain activity was recorded using a 64-channel actiCAP active electrodes setup and BrainAmp amplifiers (Brain Products, Munich, Germany). The following electrode sites were used: AF3-4; Fp1-2; Fz,1-10; FCz,1-8; Cz,1-8; CPz,1-8; CP,1-9; Pz,1-10; POz,1-6; and Oz,1-2. Data were recorded at 1000 Hz with impedance kept below 20 kΩ. For offline analysis, data were downsampled to 200 Hz and low-pass filtered at 40 Hz in order to attenuate line noise. Trials in which subjects responded before the end of the video were omitted. No artifact rejection was performed.

D. Data Analysis

1) Behavioral Data: The psychometric function characterizes the performance of an observer in a detection task as a function of a physical quantity, i.e., in this case, the change in video quality. The performance is given by the detection rate, i.e., the fraction of trials where subjects are reported to have detected the stimulus divided by the total number of trials for the stimulus level. To obtain the psychometric curve, a logistic function was fitted to the detection rates of the stimulus levels with the *psignifit* toolbox [18] using bootstrapping to determine the confidence intervals of the fit.

2) Neurophysiology: Event-related potentials (ERPs) are obtained by aligning EEG data from a number of trials according to a predefined time point and averaging over trials. The time instance of alignment was chosen to be the time instance of the quality drop. ERP waveforms reflect only neuronal activity that is phase locked to the stimulus because activity that is not phase locked to the quality change averages out. Averaging also helps

[2]Stimulus levels with Latin numbers denote the pooled quality change levels over subjects. Since the stimulus levels are selected individually for each subject, the pooled stimulus levels might have different physical characteristics but are most similar perceptually.
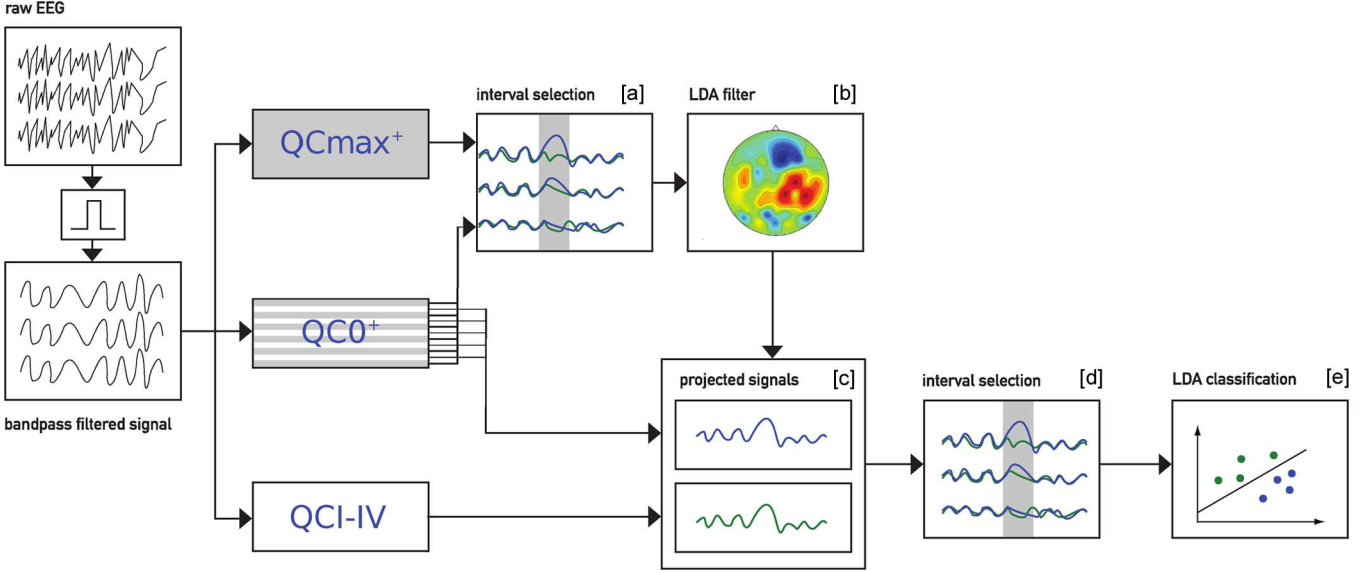
Fig. 3. EEG single-trial classification scheme. First, the raw EEG signal is bandpass filtered. Then, QC0$^+$ trials (where the superscript "+" indicates the subset of trials that were correctly labeled by the subject, i.e., correct rejections) are split into two equisized sets in an even–odd manner, as indicated by the gray-and-white horizontal bars. One of these two sets and the QCmax$^+$ trials (again, "+" indicates correctly labeled trials, in this case hits) are used to train the LDA filter. To this end, the sgn-r^2 between QCmax$^+$ and QC0$^+$ is used to select a discriminative interval [a] on a basis of which an LDA filter is determined [b]. Then, the sets of quality change trials QC-I, QC-II, QC-III, and QC-IV (depicted as QCI–IV) and the second set of QC0$^+$ trials are projected to virtual channels using the LDA filter [c]. For each QC level and for hits and misses separately, two intervals that discriminate best between quality changes and undistorted trials are selected [d]. This yields 2-D features that are finally used to classify single trials using LDA [e].

to increase the signal-to-noise ratio of the transient ERP waveform. Common parameters to characterize an ERP are the amplitude and the latency (with respect to stimulus onset) of the peaks in the waveform.

3) LDA: The classification of EEG data in the temporal domain is normally done by comparing the ERP data of trials from different stimulus levels. For each trial, features are derived from each channel at different time points, which are then used for classification. Usually, a single feature corresponds to the voltage averaged in a certain time window for a particular EEG channel.

For a linear classifier, the separating hyperplane is defined by $\mathbf{w}^T\mathbf{x} + b = 0$, where $\mathbf{w}$ is the projection vector, $\mathbf{x}$ is the data point on the separating hyperplane, and b is the bias term (classification threshold). The classification of data point $\mathbf{x_i}$ is then given by $y_i = sgn(\mathbf{w}^T\mathbf{x_i} + b)$.

Linear discriminant analysis (LDA) is a linear classifier that aims at finding a projection that maximizes the ratio of the class mean distance (of two or more classes) and the within-class variances. The projection vector of LDA is defined as

$$\mathbf{w} = \hat{\boldsymbol{\Sigma}}_c^{-1}(\hat{\boldsymbol{\mu}}_2 - \hat{\boldsymbol{\mu}}_1)$$

where $\hat{\boldsymbol{\mu}}_i$ is the estimated mean of class i and $\hat{\boldsymbol{\Sigma}}_c = 1/2(\hat{\boldsymbol{\Sigma}}_1 + \hat{\boldsymbol{\Sigma}}_2)$ is the estimated covariance matrix (the matrix of covariances of all EEG channels), i.e., the average of the classwise empirical covariance matrices $\hat{\boldsymbol{\Sigma}}_i$.

A linear classifier trained on temporal EEG features can be regarded as a spatial filter. Thus, the linear classifier may be interpreted as a "backward model" to recover the signal of discriminative sources. The weight vector $\mathbf{w}$ of the classifier can be visualized as a scalp map (c.f. Section II-D-4).

Let $\mathbf{w} \in \mathbb{R}^C$ be the LDA projection vector and $\mathbf{D} \in \mathbb{R}^{C \times T}$ be a matrix of EEG signals with C channels and T samples per channel. Then

$$\mathbf{D}_f := \mathbf{w}^\top\mathbf{D} \quad \in \mathbb{R}^{1 \times T}$$

is the result of spatial filtering: each EEG channel in $\mathbf{D}$ gets weighted with the corresponding component of $\mathbf{w}$ and summed up to yield a single virtual channel. In other words, each virtual channel is a linear combination of the original EEG channels.

For Gaussian distributions with equal covariance matrices for both classes, LDA is the optimal classifier, i.e., the risk of misclassification for samples drawn from the same distributions is minimized [19]. Since ERP signals are approximately Gaussian distributed and the covariances are dominated by the background EEG signal rather than by class-specific covariations, LDA is perfectly suitable for classifying ERP signals.

However, the stated optimality of LDA relates to known parameters μ_1, μ_2, and $\boldsymbol{\Sigma}_c$. For real applications, the dimensionality of the feature space is often high, whereas the number of observations is comparatively low. For EEG data, the ratio between the number of observations and the number of channels is low, which leads to a systematic misestimation of the covariance matrix [20]. Due to that, the estimated distribution parameters, in particular $\hat{\boldsymbol{\Sigma}}_c$, are error prone, which render classification suboptimal. Accordingly, classification was done by LDA with automatic regularization of the estimated covariance matrix using shrinkage [20], which outperforms the standard LDA when classifying ERPs on a single-trial level. For detailed overviews on single-trial classification of EEG data, see [21] and [22].

4) Classification: Classification is done subjectwise in a two-step procedure (c.f. Fig. 3). First, discriminative time intervals between correctly reported QC0 trials (correct rejections,

Fig. 4. Scalp plots depicting spatial distribution as a top view on the head, with nose pointing upwards and crosses marking the electrode positions. Left: sgn-r^2 values for QCmax$^+$ against QC0$^+$ for subject S1. The spatial distribution is similar to the P3 component, suggesting that class differences are mostly due to the P3. Right: LDA filter trained on the two classes. If the channels were uncorrelated, the LDA filter would be proportional to the difference of the class means. If the noise is not substantial (as it is the case for our narrowband filtered data), the spatial distribution roughly consists of dipoles along the gradient of the sgn-r^2 scalp plot.
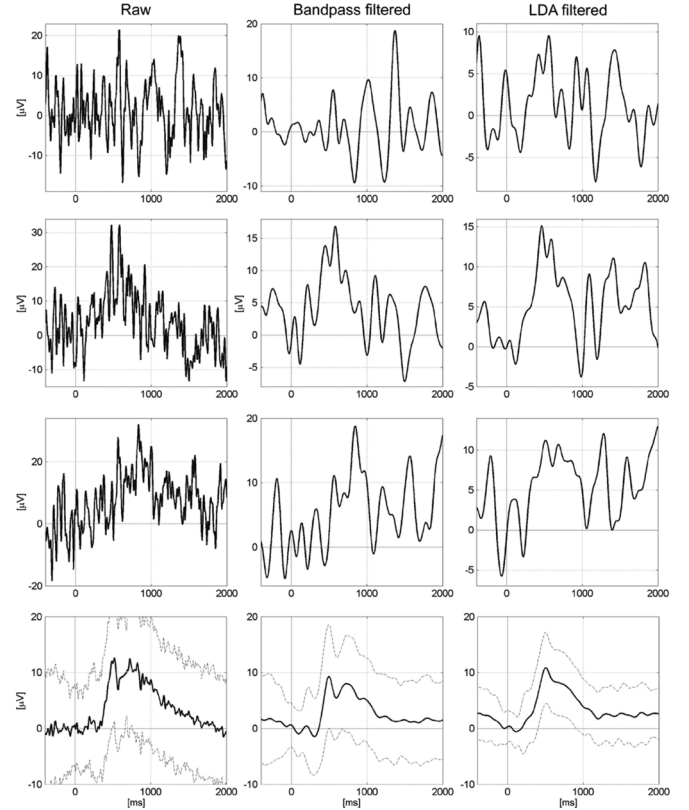

Fig. 5. EEG data from hit trials of stimulus-level QC-IV of subject S1. First column: Raw EEG data at channel CPz. Second column: Bandpass-filtered EEG data at channel CPz. Third column: LDA-projected EEG data. Rows 1–3: Single-trial data from three hit trials of subject S1. Last row: ERP and standard deviation over all hit trials of QC-IV for subject S1. The ERP peak is similar for all three data types, but the standard deviation is substantially reduced by each processing step. Further, the processing clearly increases the prominence of the P3 component in single trials. This illustrates that the filtering steps are beneficial for the enhancement the P3 component and, therefore, classification performance.

denoted as QC0$^+$) and correctly detected trials with highest quality change (hits, denoted as QCmax$^+$) are computed [see Fig. 3(a)].

To measure class discrimination for each channel and time point (over trials), we use the signed squared biserial correlation coefficient sgn-r^2 := sgn$(r) \cdot r^2$. The biserial correlation coefficient is defined as

$$r := \frac{\sqrt{N_1 \cdot N_2}}{N_1 + N_2} \frac{\mu_1 - \mu_2}{\sigma}$$

where N_1 and N_2 are the number of samples in the two classes, μ_i is the mean over samples (EEG voltage levels in this case) of class i (stimulus levels), and σ is the standard deviation over all samples. Roughly, the r^2 measures how much of the total variance for all samples can be explained by class membership.

Since classification is done on the P3, which is a single positive component, it is sufficient and neurophysiologically reasonable to choose only one time interval.[3] For the channelwise mean of this interval, the LDA projection vector is computed [see Fig. 3(b)]. Fig. 4 shows the projection vector obtained for one subject.

The data for the different stimulus levels are projected on the LDA filter (see Fig. 3(c); c.f. Section II-D-3). The LDA filter is computed on QCmax$^+$ versus QC0$^+$ since we expect the P3 component to be most prominent for this stimulus level (this is verified in Fig. 7, top left). By projecting EEG data of the other stimulus levels onto the LDA filter, we ensure that the classification is done on the P3 component instead of artifacts such as eye blinks that might also be discriminative to some degree. Furthermore, as illustrated in Fig. 5, LDA prefiltering reduces the variability across trials and thus enhances class separation on a single-trial level. In addition, the dimensionality (i.e., the number of channels) is thereby largely reduced. This is particularly important as LDA relies on an estimation of the covariance matrix, which is systematically skewed if

---

[3]The VEP is not used for classification since, in this paper, its amplitude was considerably smaller than the P3 amplitude and was of negligible class discriminability.

the ratio between the number of features and the number of observations is high.

In the second step, the LDA-prefiltered data are classified stimulus-levelwise (QC I–IV) against EEG data of correctly rejected undistorted trials. Again, the sgn-r^2 heuristic is used to extract the most discriminative time intervals [21] [see Fig. 3(d)]. Simulations have shown that two intervals suffice. By searching for the most discriminative subject-specific time intervals in the LDA-projected data, the classification becomes invariant with respect to the exact position of the P3 component relative to the stimulus onset. The trialwise means within the selected time intervals are the features (2-D) for the classification with LDA [see Fig. 3(e)]. Thus, this analysis is based on the assumption that a P3 component for lower quality changes has a similar spatial distribution to the P3 component of QCmax, although amplitude and latency might differ. Note that two disjunct sets of undistorted trials from QC0, created by alternating trials chronologically, are used for step 1 and 2.

For each subject, hits (true positives) and misses (false negatives) of each stimulus level are cross-validated separately against correctly reported undistorted trials (QC0$^+$) in a leave-one-out fashion. Classification is only performed on

Fig. 6. Psychometric function fits from the psychophysical data (small circles), each function represents one subject. Horizontal lines depict the 95% confidence intervals of the fit.
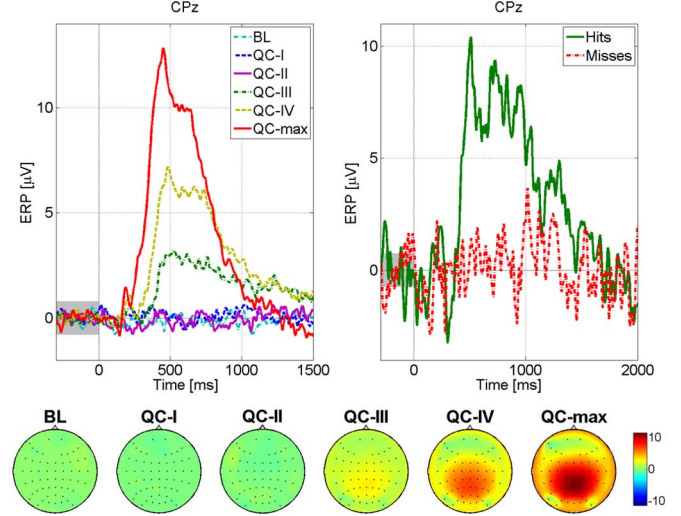

Fig. 7. Grand average ERP plots for the different stimulus levels. Top left: ERP for undistorted trials and the different quality changes at channel CPz. Top right: ERP for a selected stimulus level (QC-III) for subject S1, subdivided in hits (wherein the quality change was perceived) and misses (wherein the quality change was not perceived). Bottom: Scalp topographies for all channels. Each circle depicts a top view of the head, with the noise pointing upwards. Colors code the mean voltage for the time interval from 400–700 ms after quality change. ERP plots for single subjects can be found in the supplementary material.

stimulus levels with sufficient numbers of hit or miss trials ($n \geq 20$). Classifying hits against $QC0^+$ serves as a proof of concept. It shows that the features derived from EEG indeed reflect the subject's perception of the quality change and that EEG-based classification on a single-trial basis is possible. By classifying misses against $QC0^+$ trials, we investigate whether the EEG signals of these two classes differ, although their behavioral response is the same. A classification performance above chance level would substantiate that EEG offers a sensitivity that goes beyond what can be inferred from the behavioral data.

Classification performance is measured by the AUC, the area under the curve of the receiver operating characteristic (ROC) [23]. The ROC curve displays the true positive rate and the false positive rate of a 2-class problem as the discrimination criterion is varied. It is frequently used in machine learning as a measure of classification separability that is invariant to the number of observations per class. The AUC value equals the probability that a randomly chosen instance of class 1 has a higher value than one of class 2. An AUC of 1 (or 0) thus reflects perfect class separation, and an AUC of 0.5 reflects classification by chance.

For classification, data are bandpass filtered with a butterworth filter between 0.2 and 7 Hz to attenuate the influence of slow wave and alpha activity.

III. RESULTS

A. Behavioral Data

Fig. 6 depicts the psychometric functions fitted to the detection rate of the subjects at different stimulus levels.

The stimulus levels selected in the pretest were identical for the majority of subjects (see Table I). In addition, the corresponding psychometric functions are roughly consistent across subjects.

B. ERPs

The ERP waveform is dominated by the P3 component that is evident from 400–600 ms following the quality change (see Fig. 7). Its peak is broadly distributed over the scalp with a


Fig. 8. Relationship between neurophysiological and behavioral measures. Left: Amplitude and latency of the P3 component for QCmax shows a significant linear correlation across subjects ($r = -0.72$, $p < 0.05$, left plot). Right: Within subjects and across stimulus levels, amplitude and detection rate are positively correlated ($r = 0.84 \pm 0.15$; right plot).

center at central–parietal electrode sites (channel CPz). For QCmax, the P3 was present for all subjects. We found that P3 amplitude increases with the magnitude of quality change, which implies that it is a good index of stimulus intensity. This is supported by the fact that there is a strong correlation ($r = 0.89$ on average) between the detection rate of the subjects for the different stimulus levels and the corresponding P3 amplitude (see Fig. 8). In other words, the average neuronal response directly reflects task difficulty: If the quality change is easy to detect, P3 amplitude is high [24], [25]. For quality changes near the threshold of perception, the amplitude is low. For stimulus levels below the perceptual threshold (QC-I/QC-II), the grand average shows no P3 component.
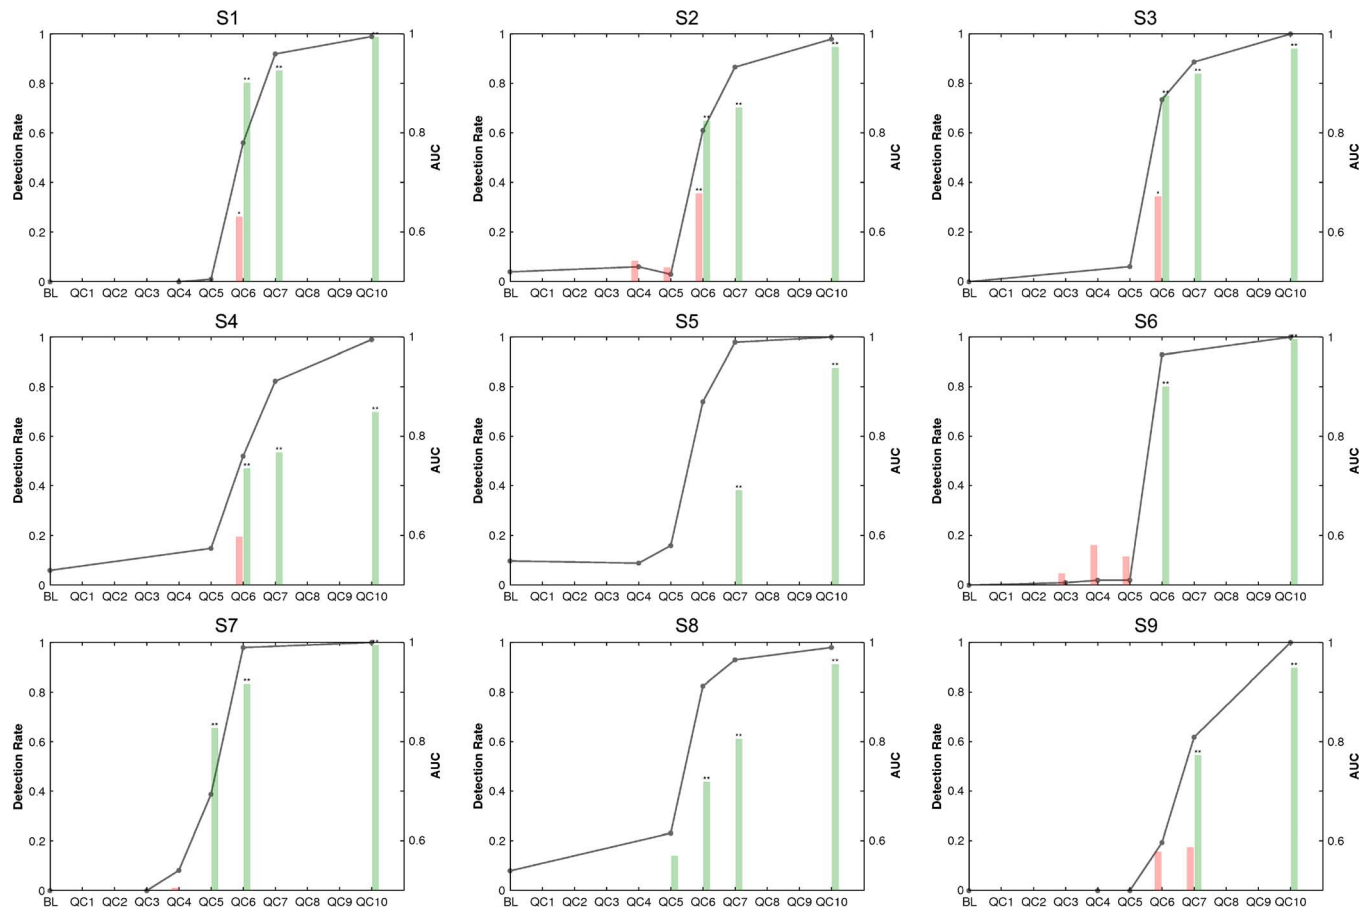
Fig. 9. Classification results for all subjects (S1–S8). Green bars show the classification performance (AUC value) of hits against $QC0^+$; red bars depict misses against $QC0^+$. One or two asterisks denote the significance level of the classification outcome in a Wilcoxon rank-sum test ($p < 0.05$ or $p < 0.01$, respectively). The gray curve depicts the detection rate over stimulus levels. Note that the detection rate and the classification performance have no direct connection as the classification is done separately on hit and miss trials.

Separate ERPs over reported and unreported quality changes differ considerably for all subjects, even within the same stimulus class (see Fig. 7, top right), which further indicates that the P3 reflects neuronal processing of the quality change.

C. Classification

Classification results for all subjects are depicted in Fig. 9. Single-trial classification performance of hits against undistorted trials is highly significant ($p < 0.01$) for all subjects, with AUC values close to 1 for QCmax in most cases. For lower stimulus levels, the AUC drops slightly to 0.8–0.9 for most subjects. An increase in classification accuracy with higher stimulus levels was observable for all subjects. Since ERP amplitude increases with the stimulus level, classification performance is tightly coupled to ERP amplitude. Across subjects, a comparably small ERP amplitude for QCmax also tends to lead to lower performance (c.f. Fig. 9; subjects S4, S5, and S9) and a high amplitude for QCmax leads to a high performance irrespective of the stimulus level (subjects S1, S3, and S7).

Classifying misses against undistorted trials yielded significant results for three subjects, with AUC values around 0.65. For all other subjects, classification did not exceed chance level.

In line with the hits versus undistorted classification, there was a clear tendency for classification performance to increase with the stimulus level (e.g., all three statistically significant classifications were measured for the highest stimulus class with a sufficient number of misses).

However, classification performance alone is not a sufficient measure when the objective is to detect conservative behavior, i.e., trials that were reported as "not perceived," although the subject faintly did. Since the fraction of these trials within all miss trials within a stimulus level may be small and may differ across subjects, classification performance does not reflect how well these trials have been detected.

An alternative approach can be motivated as follows. If a quality change was perceived residually but not reported, this should be reflected by a (possibly low amplitude) P3 component in the EEG. If not, the ERP waveform should be indistinguishable from the undistorted trials. Thus, if the subjects answered conservatively, a P3 component should be detectable by averaging over trials that were not labeled as undistorted trials by the classifier ("classifier hits"). Generally, the P3 itself is characterized by its spatial and its temporal form. However, for LDA-prefiltered classifier hits, the spatial distribution is roughly fixed by the LDA filter, i.e., the classifier will only label those trials as positive that have a P3-like spatial distribution in the classification time interval. However, as a P3-like spatial distribution
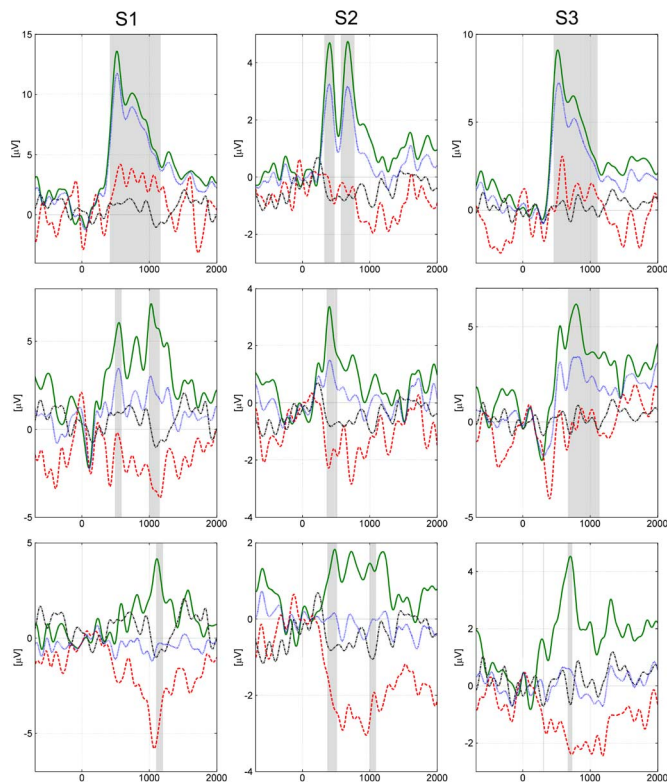
Fig. 10. ERPs of LDA-prefiltered data for subjects S1, S2, and S3. The blue dotted line shows the ERP over all trials of the stimulus level, and the black dash–dotted line shows the ERP of the $QC0^+$ undistorted trials. The green and red lines give the ERP over the trials of the stimulus level, subdivided into trials that are classified as hits (green) and misses (red) by the classifier. The gray-shaded areas depict the time intervals from which classification features are computed. Rows are the results for different classification runs. Top row: Hits versus $QC0^+$ for QC-III. Middle row: Misses versus $QC0^+$ for QC-III. These are the three classifications on miss trials that are statistically significant. Bottom row: Misses versus $QC0^+$ for the lowest QC-level of the subject. Note the different scaling of the plots.

might also arise due to random fluctuations in the EEG, the temporal time course is important to tell apart a random distortion from a true P3.

In Fig. 10, three different classification tasks were performed corresponding to the three rows. In each column, the results of the same three example subjects are shown. The curves correspond to the ERP averages of the EEG data after it has been projected onto a single channel using the LDA filter. For the classification of hits against undistorted trials (top row), a clear P3-like component is present for all three example subjects. This component is visible for the QC-III$^+$ trials, i.e., the distorted QC-III trials that have been detected by the classifier (green line), but it is also visible when one averages over all QC-III trials, including those that have not been detected by the classifier (blue dotted line). For QC-III$^-$ trials (undetected distorted trials) and undistorted trials, no P3 component is visible. For misses against undistorted trials (middle row), a P3 component with similar latency but smaller amplitude can be seen in at least two subjects (S2 and S3). Note that this positive peak is not only visible for trials classified as hits but also for the average over all miss trials. For lower stimulus levels in which the classification was not significant (bottom row), there are no apparent differences in the average waveform between miss trials and undistorted trials. For these cases, some trials (classifier hits) still seem to have a P3-like shape, but as classifier misses have a negative peak during the classification interval, this effect can be attributed to random fluctuations within the EEG signal.

IV. DISCUSSION

A. General Pattern of Activation

A P3 component related to the quality change was detected in all subjects. This component shows a graded response, i.e., its amplitude scales with the magnitude of the quality distortions for all subjects (mean correlation: $r = 0.89$). The P3 has long been known to vary with stimulus probability and with stimulus intensity in both the auditory and visual modality [26]. In this experiment, the probability of a large change as in QCmax is low (17%) as all other quality changes are clearly more subtle, which is reflected neuronally by a very large P3 amplitude for QCmax compared with the other stimulus levels.

Across subjects, there is also a large variability in P3 amplitude (7–23 μV). However, amplitude differences of the P3 between subjects have been recognized for decades and are understood to depend on a variety of psychological and biological factors [27] rather than on different processing of the stimuli.

For stimulus levels below the perceptual threshold (QC-I/QC-II) for which detection rates are below 0.15, we did not find a difference between the ERPs.

B. Classification

Psychophysical experiments suffer from lapses of subjects, in which clearly perceivable stimuli are not reported (e.g., due to inattentiveness or a wrong button press) or stimuli far below the perceptual threshold are reported as perceived. For fitting a psychometric function to the behavioral data, methods have been developed to deal with these lapses [28]. However, in a yes/no task, methods purely based on the behavioral data cannot detect conservative response behavior, i.e., trials in which subjects reported to have seen no change in quality although they faintly did. We showed that the EEG classification has the potential to detect these "mislabeled" trials. For those cases where the classification is significant, the ERP of miss trials classified as hits shows a P3-like temporal (see Fig. 10) and spatial topography (positive peak centered over central–parietal areas, unpublished data), indicating that the classification is also reasonable neurophysiologically. That is, we can assume that the subject either consciously perceived the change or that the brain processed it to some degree without conscious awareness.

In some cases, single $QC0^+$ trials were classified as hits. This is possibly due to random fluctuations in the EEG that have a similar spatial distribution as the P3 component. Although these events are rare, these fluctuations impede an error-free classification, and further investigations have to be made on how to reduce their influence. Overall, for subjects with high P3 amplitude, the classification tends to be better and more stable over stimulus levels. The reason is that, on single-trial level, low-amplitude P3 components are hard to detect since the background EEG noise is high, particularly in the low frequency range from which the P3 originates.[4] Subjects with high P3 amplitude are

---

[4]The power spectrum of EEG background noise has approximately a $1/f$ distribution.

thus less prone to random fluctuations in the EEG and therefore more suited for our approach than subjects with low P3 amplitude.

Subjects did not have to press the button instantly when they perceived a quality change since this would lead to neuronal motor activity that might interfere with classification. As a result, we cannot determine whether subjects perceived the quality distortion immediately or with a time lag. In particular, for stimuli around the perceptual threshold, subjects might not notice the change exactly at its onset. In the ERP waveform, lower amplitude and slower decay can be observed for more subtle quality changes (see Fig. 7), which seems at least partially to be due to a temporal jitter of the P3 component. Although LDA prefiltering is spatial and therefore time invariant, the discriminative intervals for classification are selected using all trials in the training set. If the P3 components of these all have similar latency, this will lead to a peak at this latency for the sgn-r^2 discrimination function. If the latency is variable, this discriminative peak will be smoothed out, and an interval will be chosen that may not be optimal for all trials.

EEG data are highly noisy, nonstationary, and easily affected by psychological factors such as attentiveness, motivation, or fatigue. As a result, EEG varies considerably over the time course of an experiment, and methods that make the signals more robust against such fluctuations are important for a good single-trial classification performance. For instance, adaptive LDA classifiers [29] and stationary subspace analysis (SSA) [30] are two methods that enhance single-trial classification in BCIs [31], [32]. Future work has to examine to what extent these and other methods prove useful in video quality experiments.

C. Toward Neurotechnology in Video Quality Assessment

A full assessment of the human perception of video quality is beyond the reach of conventional experimental methods. The purpose of this paper is to pave the way for a neurotechnology-based approach to video quality assessment by giving a proof-of-concept. While we do not present a full-fleshed solution for objective, robust, and reliable assessment of video quality perception, we believe that our method shows that neurotechnology can be a useful complement to, and an extension of, established behavioral methods. To substantiate this, it seems feasible to enumerate its potential merits in quality assessment.

- Overt response. The direct monitoring of brain activity releases one from the necessity to record overt responses such as button presses. Overt responses not only interrupt the experimental flow but they can also interfere with the subject's evaluation of a stimulus.
- Real-time monitoring. Brain activity can be monitored and processed in real time, potentially giving an impression of the user's assessment of video material while it is being viewed.
- Objectivity. Behavioral methods often suffer from response bias, i.e., some subjects are more inclined than others to give a particular response or rating. Tapping the brain response directly promises a more objective account on the perception and assessment of a stimulus than behavioral methods.

- Sensitivity. EEG is sensitive to stimuli that are at or even below the threshold of conscious perception. In this paper, we have some indications for that from subject S6 (see Fig. 9), who showed a brain response to low-distortion stimuli, although he did not report perceiving them. More robust evidence stems from a recent study on visual flicker, which, by using machine learning, showed that the brain can respond to flicker even when it is not reported by the subject [33].

D. Caveats

Before an out-of-the-box solution to video quality assessment is in reach, several more hurdles have to be overcome. This paper is limited in the following respects.

First, the coding artifacts arising in our stimulus video do not cover the full range of possible artifacts, and those artifacts appear in a sudden change of quality. Thus, the quality characteristics of the used stimuli are similar to the quality deviation in the special case of packet loss in fidelity-scalable video coding [34]. For a full assessment, several experiments have to be conducted to examine the different types of codec artifacts and typical artifact combinations.

Second, the EEG approach yielded significant results only with a subset of subjects. In particular, our approach at present requires a sufficiently large P3 amplitude of the subject. At present, two routes are possible to remedy this limitation. One route is to increase the sensitivity and robustness of EEG using advanced signal processing methods such as those mentioned earlier. Another route is to have a method to quickly identify the subjects that are suitable for an EEG-based approach. Such screening methods have already been explored in the context of BCIs [35], [36]. In the present context, one might envisage a short measurement using a classical oddball paradigm, based on which the P3 amplitude can be estimated. P3 amplitude could serve as a good predictor of classification performance due to its strong correlation.

Third, it seems that the extra preparation time required by an EEG setup ($>1/2$ h) may counteract the benefits of EEG measurements. However, a new generation of EEG caps using dry electrodes [37], [38] eliminates the time-expensive preparation of the EEG cap. This could substantially increase the practical applicability of our approach in quality assessment experiments.

Fourth, the present proof-of-concept study used artificially generated movie clips featuring a homogeneous 10-s video stream. Thus, one might wonder how this approach would work for real videos, particularly given the regular occurrence of scene changes, which could be perceived as unexpected changes and might thus trigger themselves a P3 component. In this respect, it is important to note that scene changes may be considered not an interfering nuisance, rather they could serve as a kind of calibration normal for each subject under study: Scene changes (which could be automatically detected from the video stream) define instances where the individual spatiotemporal profile of the P3 voltage maps can be obtained, against which the P3 triggered by artificially inserted quality changes can be compared. Thereby, one could obtain a natural metric for characterizing the extent of quality-change-related EEG data across subjects.

V. Conclusion

We proposed a novel approach based on neural correlates from EEG. This approach holds the promise to provide a less biased and more objective account of quality perception than obtained with behavioral methods.

Clearly, the method presented here needs further improvement, but we conjecture that it can be a valuable complement in psychophysical experiments wherein the number of trials is limited by relabeling mislabeled trials. Furthermore, since its focus is the perception of quality *change*, not quality perception per se, it forms only the first step of a neurotechnology-based approach to video quality assessment. Future work will aim toward a direct measure of quality perception using tonic EEG features such as oscillatory components (e.g., alpha rhythm). A direct neural index would allow for the real-time assessment of perceived image quality during the observations of videos, and it would relinquish the need for an overt response by the subject.

Acknowledgment

References

[1] K. Seshadrinathan and A. C. Bovik, "Motion tuned spatio-temporal quality assessment of natural videos," *IEEE Trans. Image Process.*, vol. 19, no. 2, pp. 335–350, Feb. 2010.

[2] A. B. Watson and J. Malo, "Video quality measures based on the standard spatial observer," in *Proc. ICIP*, Rochester, NY, 2002, pp. III-41–III-44.

[3] *Methodology for the Subjective Assessment of the Quality of Television Pictures*, Rec. ITU-R BT.500-11, 2002.

[4] *Subjective Video Quality Assessment Methods for Multimedia Applications*, Rec. ITU-T P.910, 2008.

[5] J. Polich, "Updating P300: An integrative theory of P3a and P3b," *Clin. Neurophysiol.*, vol. 118, no. 10, pp. 2128–2148, Oct. 2007.

[6] P. Nunez, R. Srinivasan, A. Westdorp, R. Wijesinghe, D. Tucker, R. Silberstein, and P. Cadusch, "EEG coherency: I: Statistics, reference electrode, volume conduction, Laplacians, cortical imaging, and interpretation at multiple scales," *Electroencephalogr. Clin. Neurophysiol.*, vol. 103, no. 5, pp. 499–515, Nov. 1997.

[7] L. Farwell and E. Donchin, "Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials," *Electroencephalogr. Clin. Neurophysiol.*, vol. 70, no. 6, pp. 510–523, Dec. 1988.

[8] M. S. Treder, N. M. Schmidt, and B. Blankertz, "Gaze-independent brain–computer interfaces based on covert attention and feature attention," *J. Neural Eng.* vol. 8, no. 6, p. 066003, Oct. 2011 [Online]. Available: http://dx.doi.org/10.1088/1741-2560/8/6/066003

[9] M. Schreuder, B. Blankertz, and M. Tangermann, "A new auditory multi-class brain–computer interface paradigm: Spatial hearing as an informative cue," *PLoS One*, vol. 5, no. 4, p. e9813, Apr. 2010.

[10] J. Höhne, M. Schreuder, B. Blankertz, and M. Tangermann, "Two-dimensional auditory P300 speller with predictive text system," in *Proc. IEEE Annu. Int. Conf. EMBC*, 2010, pp. 4185–4188.

[11] A. Porbadnigk, J. Antons, B. Blankertz, M. S. Treder, R. Schleicher, S. Möller, and G. Curio, "Using ERPs for assessing the (sub) conscious perception of noise," in *Proc. IEEE Annu. Int. Conf. EMBC*, 2010, pp. 2690–2693.

[12] B. Blankertz, G. Dornhege, M. Krauledat, K.-R. Müller, and G. Curio, "The non-invasive Berlin brain–computer interface: Fast acquisition of effective performance in untrained subjects," *Neuroimage* vol. 37, no. 2, pp. 539–550, Aug. 2007 [Online]. Available: http://dx.doi.org/10.1016/j.neuroimage.2007.01.051

[13] *Toward Brain–Computer Interfacing*, G. Dornhege, J. del R. Millán, T. Hinterberger, D. McFarland, and K.-R. Müller, Eds. Cambridge, MA: MIT Press, 2007.

[14] B. Blankertz, R. Tomioka, S. Lemm, M. Kawanabe, and K.-R. Müller, "Optimizing spatial filters for robust EEG single-trial analysis," *IEEE Signal Process. Mag.* vol. 25, no. 1, pp. 41–56, Jan. 2008 [Online]. Available: http://dx.doi.org/10.1109/MSP.2008.4408441

[15] C. Gerald and P. Wheatley, *Applied Numerical Analysis*. Reading, MA: Addison-Wesley, 1989.

[16] D. Marpe, H. Schwarz, S. Bosse, B. Bross, P. Helle, T. Hinz, H. Kirchhoffer, H. Lakshman, T. Nguyen, S. Oudin, M. Siekmann, K. Suhring, M. Winken, and T. Wiegand, "Video compression using nested quadtree structures, leaf merging, and improved techniques for motion representation and entropy coding," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 20, no. 12, pp. 1676–1687, Dec. 2010.

[17] T. Wiegand, W.-J. Han, B. Bross, J.-R. Ohm, and G. J. Sullivan, JCTVC, WD3: Working Draft 3 of High-Efficiency Video Coding Mar. 2011.

[18] I. Fründ, N. Haenel, and F. Wichmann, "Inference for psychometric functions in the presence of nonstationary behavior," *J. Vis.*, vol. 11, no. 6, p. 16, May 2011.

[19] R. Duda, P. Hart, and D. Stork, *Pattern Classification*. New York: Wiley, 2001.

[20] J. Schäfer and K. Strimmer, "A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics," *Stat. Appl. Genetics Mol. Biol.*, vol. 4, no. 1, p. 32, 2005.

[21] B. Blankertz, S. Lemm, M. S. Treder, S. Haufe, and K.-R. Müller, "Single-trial analysis and classification of ERP components-a tutorial," *Neuroimage*, vol. 56, no. 2, pp. 814–825, May 2011.

[22] K.-R. Müller, C. Anderson, and G. Birch, "Linear and nonlinear methods for brain–computer interfaces," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 11, no. 2, pp. 165–169, Jun. 2003.

[23] D. Green and J. Swets, *Signal Detection Theory and Psychophysics*. Huntington, NY: Krieger, 1974.

[24] G. Hagen, J. Gatherwright, B. Lopez, and J. Polich, "P3a from visual stimuli: Task difficulty effects," *Int. J. Psychophysiol.*, vol. 59, no. 1, pp. 8–14, Jan. 2006.

[25] K. Kim, J. Kim, J. Yoon, and K. Jung, "Influence of task difficulty on the features of event-related potential during visual oddball task," *Neurosci. Lett.*, vol. 445, no. 2, pp. 179–183, Nov. 2008.

[26] J. Polich, P. Ellerson, and J. Cohen, "P300, stimulus intensity, modality, and probability," *Int. J. Psychophysiol.*, vol. 23, no. 1/2, pp. 55–62, Aug./Sep. 1996.

[27] J. Polich and A. Kok, "Cognitive and biological determinants of P300: An integrative review," *Biol. Psychol.*, vol. 41, no. 2, pp. 103–146, Oct. 1995.

[28] F. Wichmann and N. Hill, "The psychometric function: I. Fitting, sampling, and goodness of fit," *Perception Psychophys.*, vol. 63, no. 8, pp. 1293–1313, Nov. 2001.

[29] C. Vidaurre, C. Sannelli, K.-R. Müller, and B. Blankertz, "Machine-learning-based coadaptive calibration for brain–computer interfaces," *Neural Comput.*, vol. 23, no. 3, pp. 1–28, Mar. 2011.

[30] P. von Bünau, F. Meinecke, F. Király, and K.-R. Müller, "Finding stationary subspaces in multivariate time series," *Phys. Rev. Lett.*, vol. 103, no. 21, pp. 214 101-1–214 101-4, Nov. 2009.

[31] P. von Bünau, F. Meinecke, S. Scholler, and K.-R. Müller, "Finding stationary brain sources in EEG data," in *Proc. IEEE Annu. Int. Conf. EMBC*, 2010, pp. 2810–2813.

[32] C. Vidaurre and B. Blankertz, "Towards a cure for BCI illiteracy," *Brain Topography*, vol. 23, no. 2, pp. 194–198, Jun. 2010.

[33] A. K. Porbadnigk, S. Scholler, B. Blankertz, A. Ritz, M. Born, R. Scholl, K.-R. Müller, G. Curio, and M. S. Treder, "Revealing the neural response to imperceptible peripheral flicker with machine learning," in *Proc. IEEE Eng. Med. Biol. Soc. Conf.*, 2011, vol. 2011, pp. 3692–3695.

[34] H. Schwarz, D. Marpe, and T. Wiegand, "Overview of the scalable video coding extension of the H.264/AVC standard," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 17, no. 9, pp. 1103–1120, Sep. 2007.

[35] B. Blankertz, C. Sannelli, S. Halder, E. Hammer, A. Kübler, K.-R. Müller, G. Curio, and T. Dickhaus, "Neurophysiological predictor of SMR-based BCI performance," *NeuroImage*, vol. 51, no. 4, pp. 1303–1309, Jul. 2010.

[36] M. S. Treder, A. Bahramisharif, N. M. Schmidt, M. van Gerven, and B. Blankertz, "Brain–computer interfacing using modulations of alpha activity induced by covert shifts of attention," *J. Neuroeng. Rehabil.* vol. 8, p. 24, May 2011 [Online]. Available: http://www.jneuroengrehab.com/content/8/1/24/abstract

[37] F. Popescu, S. Fazli, Y. Badower, B. Blankertz, and K.-R. Müller, "Single trial classification of motor imagination using 6 dry EEG electrodes," *PLoS one*, vol. 2, no. 7, p. e637, Jul. 2007.

[38] C. Grozea, C. Voinescu, and S. Fazli, "Bristle-sensors—Low-cost flexible passive dry EEG electrodes for neurofeedback and BCI applications," *J. Neural Eng.*, vol. 8, no. 2, p. 025008, Mar. 2011.

Simon Scholler received the B.S. degree in cognitive science from the University of Osnabrück, Osnabrück, Germany, in 2007 and the M.S. degree in computational neuroscience from the Bernstein Center for Computational Neuroscience Berlin, Berlin, Germany, in 2011.

From 2007 to 2011, as a Student Assistant, he has done research on brain–computer interfaces with the Intelligent Data Analysis Group, Fraunhofer FIRST, Berlin. Then, he worked on models for audio recognition. He is currently with the Machine Learning Group, Berlin Institute of Technology, Berlin, as a Research Associate.

Sebastian Bosse received the Diploma degree from RWTH Aachen University, Aachen, Germany, in 2008.

Since 2009, he has been with the Image and Video Coding Group, Fraunhofer Heinrich Hertz Institute (Fraunhofer HHI), Berlin, Germany, as a Research Associate. His research interests include image and video coding, computer vision, and human visual perception.

Matthias Sebastian Treder He received the M.Sc. degree in cognitive psychology and the Ph.D. degree in human visual from Radboud University, Nijmegen, The Netherlands, in 2004 and 2010, respectively.

Since 2009, he has been a Postdoctoral Researcher with Berlin Brain–Computer Interface Group, Berlin Institute of Technology, Berlin, Germany. His research interests include brain–computer interfacing, neuroergonomics, and visual perception using electroencephalography, and near-infrared spectroscopy.

Benjamin Blankertz received the Diploma degree in mathematics and the Ph.D. degree in mathematical logic from the University of Münster, Münster, Germany, in 1994 and 1997, respectively.

Since 2000, he has been with the Intelligent Data Analysis Group, Fraunhofer FIRST, Berlin, Germany. Since 2007, he has been with Berlin Institute of Technology, Berlin. He is the Head of the Berlin Brain–Computer Interface Project and is the Principal Investigator in the Bernstein Focus: Neurotechnology Berlin. His scientific interests include machine learning, analysis of neuronal data, and psychoacoustics.

Gabriel Curio received the Dr. Med. degree with a thesis about attentional influences on smooth pursuit eye movements from Freie Universität Berlin (FU Berlin), Berlin, Germany, and Erasmus Universiteit Rotterdam, Rotterdam, The Netherlands.

Since 1991, he has been leading the Neurophysics Group, FU Berlin. He is currently a Professor of neurology and a Deputy Director with the Department of Neurology and Clinical Neurophysiology, Campus Benjamin Franklin of the Charité–University Medicine Berlin, Berlin. He is a Founding Codirector of the Bernstein Center for Computational Neuroscience Berlin and the Berlin NeuroImaging Center, a Founding Member of the Bernstein Focus Neurotechnology Berlin, and a Faculty member of the Berlin Excellence School "Mind of Brain." He holds Board specializations in neurology and psychiatry. His current research interests include integrating the neurophysics of noninvasive electromagnetic brain monitoring with both, basic and clinical neuroscience concepts, such as spikelike activities in somatosensory evoked brain responses, neuromagnetic detection of injury currents, magnetoneurography, the comparison of cortical processing of phonemes versus musical chords, speech–hearing interactions, neurophysiological underpinnings of face recognition, single-trial electroencephalography/magnetoencephalography analysis, and brain–computer interfacing.

Dr. Curio is member of the Technical Commission of the German Society for Clinical Neurophysiology since 1998. He served as Section Editor for the International Federation of Clinical Neurophysiology journal *Clinical Neurophysiology* from 2003 to 2008.

Klaus-Robert Müller received the Diploma degree in mathematical physics and the Ph.D. degree in theoretical computer science both from the University of Karlsruhe, Karlsruhe, Germany, in 1989 and 1992, respectively.

From 1992 to 1994, he was a Postdoctoral Researcher with Fraunhofer FIRST, Berlin, Germany. From 1994 to 1995, he was a European Community Research Fellow with the University of Tokyo, Tokyo, Japan. In 1995, he founded the Intelligent Data Analysis (IDA) Group, Fraunhofer FIRST, and directed it until 2008. From 1999 to 2006, he was a Professor of computer science with the University of Potsdam, Potsdam, Germany. Since 2006, he has been a Professor of computer science with Berlin Institute of Technology, Berlin and also the Director of the Bernstein Focus on Neurotechnology Berlin, Berlin. His current research interests include IDA, machine learning, statistical signal processing and statistical learning theory with the application foci computational finance, computational chemistry, computational neuroscience, genomic data analysis, and, to study the interface between brain and machine, noninvasive electroencephalography-based brain–computer interfacing.

Dr. Müller was the recipient of the Olympus Prize by the German Pattern Recognition Society in 1999 and the Alcatel SEL Communication Award in 2006.

Thomas Wiegand (M'05–SM'08–F'11) received the Dipl.-Ing. degree in electrical engineering from the Technical University of Hamburg, Germany, in 1995 and the Dr.-Ing. degree from the University of Erlangen–Nuremberg, Erlangen–Nuremberg, Germany, in 2000.

From 1993 to 1994, he was a Visiting Researcher with Kobe University, Kobe, Japan. In 1995, he was a Visiting Scholar with the University of California at Santa Barbara, Santa Barbara. From 1997 to 1998, he was a Visiting Researcher with Stanford University, Palo Alto, CA, and served as a consultant to 8×8, Inc., Santa Clara, CA. In 2000, he joined the Department of Image Processing, Fraunhofer Heinrich Hertz Institute as the Head of the Image Communication Group. From 2006 to 2008, he was a Consultant with Stream Processors, Inc., Sunnyvale, CA. From 2007 to 2009, he was a Consultant with Skyfire, Inc., Mountain View, CA. He is currently a Professor with the Department of Electrical Engineering and Computer Science, Berlin Institute of Technology, Berlin, Germany, and the Chair of the Image Communication Laboratory, and is jointly heading the Department of Image Processing, Fraunhofer Heinrich Hertz Institute, Berlin. His research interests include video processing and coding, multimedia transmission, and computer vision and graphics.

Dr. T. Wiegand has been an active participant in standardization for multimedia with successful submissions to ITU-T VCEG, ISO/IEC MPEG, 3GPP, DVB, and IETF since 1995. He has been a member of the Technical Advisory Board of Vidyo, Inc., Hackensack, NJ, since 2006. He was appointed as the Associated Rapporteur of ITU-T VCEG in October 2000. He was appointed as the Associated Rapporteur/Co-Chair of the JVT in December 2001. He was appointed as the Editor of the H.264/MPEG-4 AVC video coding standard and its extensions (FRExt and SVC) in February 2002. From 2005 to 2009, he was the Co-Chair of MPEG Video. He was a Guest Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS for Video Technology for its *Special Issue on the H.264/AVC Video Coding Standard* in July 2003, its *Special Issue on Scalable Video Coding-Standardization and Beyond* in September 2007, and its *Special Section on the Joint Call for Proposals on High-Efficiency Video Coding Standardization*. He has been an Associate Editor OF IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY since January 2006. He was the recipient of the SPIE VCIP Best Student Paper Award in 1998; the Fraunhofer Award and the ITG Award of the German Society for Information Technology in 2004; the 2008 ATAS Primetime Emmy Engineering Award and a pair of NATAS Technology and Engineering Emmy Awards for the the projects that he co-chaired for the development of the H.264/AVC standard; the Innovations Award of the Vodafone Foundation, the EURASIP Group Technical Achievement Award, and the Best Paper Award of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY in 2009; the Eduard Rhein Technology Award in 2010; the Best Paper Award of EURASIP and the Karl Heinz Beckurts Award in 2011; and the IEEE Masaru Ibuka Technical Field Award in 2012.