Computer Standards & Interfaces 92 (2025) 103897

Contents lists available at ScienceDirect

Computer Standards & Interfaces

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

Interpreting the latent space of a Convolutional Variational Autoencoder for semi-automated eye blink artefact detection in EEG signals

Sabatina Criscuolo [a,b,c], Andrea Apicella [a,c,d], Roberto Prevete [a,c,d,*], Luca Longo [b]

[a] University of Naples Federico II - Department of Information Technology and Electrical Engineering, Naples, Italy
[b] Artificial Intelligence and Cognitive Load Research Lab, Technological University Dublin, Dublin, Republic of Ireland
[c] Laboratory of Augmented Reality for Health Monitoring (ARHeMLab), Italy
[d] Laboratory of Artificial Intelligence, Privacy & Applications (AIPA Lab), Italy

ARTICLE INFO

Keywords:
Electroencephalography
Variational autoencoders
Convolution
Ocular artefacts detection
Latent space interpretation

ABSTRACT

Electroencephalography (EEG) allows the investigation of brain activity. However, neural signals often contain artefacts, hindering signal analysis. For example, eye-blink artefacts are particularly challenging due to their frequency overlap with neural signals. Artificial intelligence, particularly Variational Autoencoders (VAE), has shown promise in EEG artefact removal. This research explores the design and application of Convolutional VAEs for automatically detecting and removing eye blinks in EEG signals. The latent space of CVAE, trained on EEG topographic maps, is used to identify latent components that are selective for eye blinks. Receiver Operating Characteristic (ROC) curves and Area Under the Curve (AUC) are employed to evaluate the discriminative performance of each latent component. The most discriminative component, determined by the highest AUC, is modified to eliminate eye blinks. The evaluation of artefact removal involves visual inspection and Pearson correlation index assessment of the original EEG signal and the reconstructed clean version, focusing on the $Fp1$ and $Fp2$ channels most affected by eye-blink artefacts. Results indicate that the proposed method effectively removes eye blinks without significant loss of information related to the neural signal, demonstrating Pearson correlation values around 0.60 for each subject. The contribution to the knowledge offered by this research study is the design and application of a novel offline pipeline for automatically detecting and removing eye blinks from multi-variate EEG signals without human intervention.

1. Introduction

Electroencephalography, a non-invasive and cost-effective methodology for monitoring human brain electrical activity by placing electrodes on the scalp [1,2], has traditionally found its primary applications in clinical and laboratory settings. Nowadays, the integration of EEG into daily life applications has garnered increasing interest, driven by its wearability and portability [3]. However, the broader utilization of EEG faces challenges due to artefacts influencing cerebral signals, compromising signal analysis, and posing a significant hurdle in extracting meaningful information [4,5]. Artefacts can originate from non-physiological or physiological sources [6]. While non-physiological factors can often be mitigated through precise recording equipment and signal filtering techniques, physiological sources, such as eye movements, blinks, and muscle activity, present a more intricate challenge for removal or mitigation [7]. As a result, the recorded EEG signal is a complex, non-stationary, and non-linear mixture of pure EEG signal and artefacts. In this context, the EEG denoising stage plays an important role in enhancing the quality of the EEG data and making it easier to analyse and extract meaningful information about brain activity. Specifically, this phase removes artefacts and noise from EEG signals by improving the signal-to-noise ratio and preserving the underlying neural information.

Among the physiological artefacts, the eye-blink artefact is particularly relevant [8]. This activity generates spike-like-shaped signal waveforms with peaks reaching up to 800 μV, occurring approximately every 5 s. These movements can distort EEG signals in the frequency range from 0 Hz to 12 Hz, overlapping with neural activity [9,10]. Over the years, several techniques to mitigate eye-blink artefacts have been developed. These methods can be divided into four main categories: regression methods, filtering methods, blind source separation methods (BSS), and source decomposition methods [11,12]. However, these approaches often require additional reference channels such as electrooculography (EOG), threshold and parameter settings, and clean baseline signals. Recognizing the limitations of conventional methods, artificial intelligence (AI) and, in particular, deep learning (DL) architectures have been employed for more advanced and automated

* Corresponding author at: University of Naples Federico II - Department of Information Technology and Electrical Engineering, Naples, Italy.
E-mail address: roberto.prevete@unina.it (R. Prevete).

https://doi.org/10.1016/j.csi.2024.103897
Received 1 March 2024; Received in revised form 18 May 2024; Accepted 21 July 2024
Available online 23 July 2024

solutions [13–16]. Among DL models, autoencoders (AE) and their variations, particularly variational autoencoders (VAE), have shown promise in EEG denoising [17–20]. AEs are a class of deep learning models that consist of two main components: an encoder and a decoder. The encoder maps the input data into an internal (usually smaller) representation named *latent space*, while the decoder reconstructs the original data from this representation [21]. VAE introduces a probabilistic element to this process, allowing for the projection of the original data into a space capturing the inherent variability of the data distribution [22]. This internal representation is particularly interesting in EEG signal processing, encapsulating information about underlying neural activity and associated patterns [23,24].

To the best of our knowledge, the integration of VAE into the EEG signal denoising pipeline represents an innovative methodology. Leveraging the ability of VAEs to capture complex data distributions could help detect and remove the artefacts associated with eye blinks. The VAEs have independent latent components, and thus, by learning the underlying structure of EEG signals in an unsupervised manner, they could potentially disentangle blink-related artefacts from neural patterns, offering a promising way for more accurate EEG data analysis. Based on these considerations, the primary goal of this study is to investigate the application of VAEs in the context of automatic blink detection and removal from EEG signals. The primary contribution of this research lies in detecting and removing eye blinks from EEG signals through the interpretation of the latent space of a VAE. By analysing the latent space and assuming that information is gained into the latent components that capture the brain activity [20,23], the aim is to identify the significance of each latent component of the VAE trained with EEG topographic maps for eye blink detection. Specifically, this work presents a convolutional VAE-based (CVAE) model capable of representing the underlying structure of EEG data, focusing on blink-related patterns. In detail, the latent space of CVAE trained on topology-preserving EEG topographic maps of each subject is examined. After the EEG pre-processing stage and the training of such CVAE, the latent space is interpreted to identify components for distinguishing blink and no-blink events based on statistical measures. Once the most discriminative latent component/s is/are identified, its/their value/s in the topographic maps associated with a blink is updated to mitigate the identified artefacts. The evaluation of such an artefact removal pipeline involves the visual inspection and Pearson correlation assessment of the original EEG signal and the reconstructed, clean version, specifically focusing on the pre-frontal channels ($Fp1$ and $Fp2$) mostly affected by eye-blink artefacts. Results indicate that the proposed method effectively removes eye blinks without excessive loss of information related to the brain signal, demonstrating Pearson correlation values around 0.60 for each evaluated human subject. Moreover, the Pearson correlation of one of the electrodes placed on the midline sagittal plane of the skull (Pz) was less influenced by eye blinks, and it demonstrated a value close to 1. In other words, this means a negligible loss in the original EEG signal. Consequently, the proposed method achieved promising results in eye blink removal from EEG signals.

This research manuscript is organized as follows. Section 2 provides an overview of traditional and advanced methods for eye-blink detection beyond introducing VAE. Section 3 outlines the proposed methodology for eye-blink detection based on a latent space interpretation and analysis. Section 4 presents the experimental evaluation, whereas Section 5 shows the results and critically discusses them. Finally, Section 6 draws conclusions and final remarks.

2. Related works

The analysis of EEG signals aims to understand neural activity and cognitive processes [2]. However, EEG recordings are susceptible to various physiological artefacts that can compromise the integrity of the acquired data [4,7]. Eye blinks are a prominent challenge among these artefacts due to their overlapping with brain activity [10]. Consequently, these artefacts pose a significant difficulty in extracting meaningful information from EEG recordings. Over the years, several methods to detect and mitigate eye-blink artefacts have been developed [8,25]. A simple approach involves recording EEG while instructing the subject to avoid blinking consciously. However, given that eye blinks are involuntary activities, controlling blinking for extended periods is particularly challenging. Other techniques involve rejecting EEG signal segments containing the artefact, with the drawback of losing neural activity information [25]. Regression techniques and adaptive filtering have been traditionally involved in eye-blink artefact removal. However, they require a reference signal such as those from electroculography (EOG). Additionally, these algorithms assume that the neuronal activity within EEG and EOG signals is uncorrelated, leading to the inadvertent removal of shared neuronal activity from EEG signals [25,26]. One prevalent and extensively utilized method for detecting and mitigating eye-blink artefacts is Independent Component Analysis (ICA) [8]. Based on the BSS approach, this technique decomposes multichannel EEG data from different sources into independent components. Nonetheless, the identification of artefact components has relied on the manual inspection of topographic maps and time series of ICs. In recent years, the Artefact Subspace Reconstruction (ASR) technique was introduced [7]. However, this approach requires configuring several parameters to choose which components can be considered artefacts. ASR typically requires a further clean baseline signal to allow the algorithm to identify and characterize the artefact subspace. In contrast, Empirical Mode Decomposition (EMD) is a data-driven technique employed to break down time-domain signals into a series of intrinsic mode functions (IMFs) [8]. Following signal decomposition, the artefact is typically distinguished from informative EEG signals using an appropriate threshold or correlation function. The IMFs likely related to artefacts are then eliminated from the raw EEG signal. Subsequently, the signal is reconstructed through an inverse empirical mode decomposition, recombining IMFs to produce the artefact-free EEG signal. Despite EMD efficiency in reconstructing data by selectively removing artefactual IMFs, it is computationally complex [7,27]. Recognizing the limitations of conventional methods, artificial intelligence (AI) and, in particular, deep learning (DL) architectures have been employed for more advanced and automated solutions. Among the DL architectures aimed at denoising EEG signals, autoencoders (AEs) and their variants are attracting growing interest in this field [18,19,28]. An AE [29,30] is an artificial neural network (ANN) that consists of two main components: an encoder and a decoder. The encoder maps the input data into an internal (usually lower-dimension) space, the *latent space*, obtaining a representation of the data, named *latent code*. The decoder reconstructs the input from this representation, aiming for a reconstruction closely resembling the original.

In the context of EEG denoising, Yang et al. proposed a sparse AE to remove eye-blink artefacts [28]. Specifically, the designed AE autonomously acquires knowledge of the high-order statistical moments embedded in ocular artefacts using EOG signals. The trained AE model was then used to identify and extract preliminary ocular artefacts from a raw EEG signal. Thus, the AE was used only for artefact identification, whereas the subsequent interference removal involved an adaptive filter. A similar approach was followed in [17], where a sparse AE was trained on samples without eye-blink artefacts and then used to remove ocular artefacts from the EEG signals. However, the proposed approach needs clear EEG signals during the training phase, but achieving a perfectly clean signal is often impractical in most cases. Ghosh et al. developed a method for eye blink artefact removal with an AE [31]. Nevertheless, an initial phase involving a support vector machine for blink identification is necessary before the subsequent denoising using AE. Chaung et al. proposed a model combining ICA and AE, training the ANN with mixtures of independent EEG components [32]. However, this approach requires clean EEG signals to set the hyperparameters of the AE. Nguyen et al. presented a method based

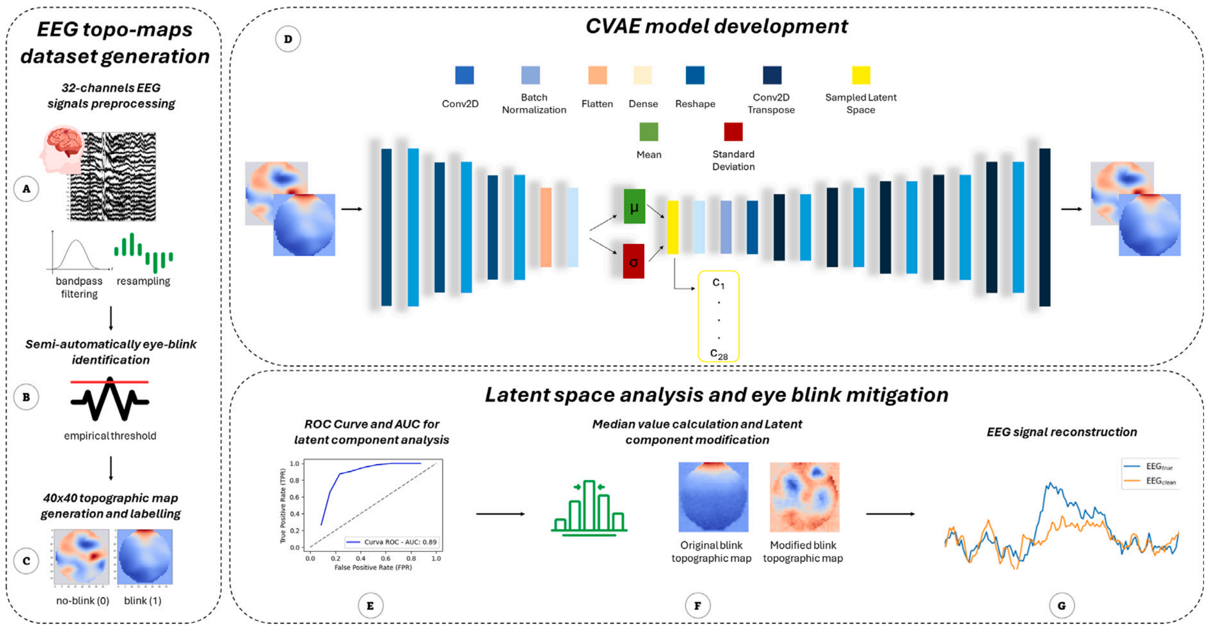

Fig. 1. Experimental assessment of the proposed method for eye-blink artefact detection and removal from the EEG signal. For each subject, 32-channel EEG signals were acquired. The EEG data was preprocessed by bandpass filtering and resampling (A). The eye blinks were manually identified using an empirical threshold (B). The EEG samples were transformed into topology-preserving topographic maps of 40 × 40 pixels and labelled (C) as 0 for no-blink maps, 1 for blink-associated maps, and 2 for transition maps between non-blink and blink states (and vice-versa). These maps were split into training, validation, and test sets. A CVAE was trained on each subject's topographic maps (D). Then, the latent space was analysed based on the ROC curve and AUC to determine the most discriminative component for blink (E). For each blink topographic map, the discriminative latent component is set to the median value of the training data without blinks (F), resulting in modified topographic maps. A clean version of the EEG signal is reconstructed (G).

on wavelet transform and sparse AE to suppress EOG artefacts [33]. Saba-Sadiya et al. proposed a deep AE to reconstruct clean EEG data after identifying EEG artefacts based on several extracted features and outlier detection algorithms [34]. Hwaidi et al. developed an unsupervised learning approach by using VAE to remove artefacts from EEG signals [19]. However, this approach relies on the availability of clean signals for training, posing a significant limitation in real-world applications where such reference signals may be unavailable. Over the years, exploration of artefact removal methodologies in EEG signals has revealed a spectrum of approaches. Traditional methods, including regression techniques and adaptive filtering, have been fundamental but face challenges preserving nuanced neural activity. ICA, ASR, and EMD have offered valuable alternatives, yet their applicability may require careful parameter tuning or baseline signals. Integrating DL architectures has emerged as a promising avenue for EEG denoising. In particular, the VAE, having independent latent components, appears more interpretable and could potentially disentangle artefacts from neural patterns.

Leveraging the limitations that emerged from the literature review, this study focuses on detecting and removing eye blinks from EEG signals by interpreting the latent space of a VAE. This approach comes from eXplainable Artificial Intelligence [35]. In this context, the aim is to understand the latent representation of a neural network-based architecture, in this case, a VAE, specifically by identifying the component associated with eye-blink artefacts. In detail, to mitigate an eye blink, the value of this latent component is modified while the values for the remaining components remain unchanged. This approach helps identify the components/s in the latent space related to eye blinks.

3. Proposed method

The technical problem that this study focuses on is the detection and removal of eye-blink artefacts within EEG signals by using Convolutional Variational Autoencoders (CVAEs). Specifically, the latent space of subject-specific CVAEs is interpreted to extract salient latent components for eye-blink artefact detection.

The proposed methodology follows several steps, as reported in Fig. 1.

(PHASE-A) For our experiments, for reproducible and comparable results, the public DEAP dataset [36], consisting in EEG acquisitions for an emotion recognition task from 32 subjects adopting a 32 channels device, was used. In particular, for each subject, 32-channel EEG signals are pre-processed by applying a band-pass filter (0.1 Hz to 45 Hz and resampling [37].

(PHASE-B) An initial identification of eye blinks is carried out semi-automatically by visual inspection and a threshold algorithm [38]. This process allowed us to obtain a reference for developing a system capable of automating blink detection.

(PHASE-C) The EEG samples were transformed into topology-preserving topographic head-maps from EEG data (topo-maps), as in [39,40].

(PHASE-D) The CVAE's reconstruction capacity is evaluated by using the Structural Similarity Index Measure (SSIM) and Mean-Squared Error (MSE). CVAE models for each subject were obtained at the end of each training process.

(PHASE-E) The distributions of each latent component are analysed to discern between blink and no-blink data. Assuming that blinks are anomalies in EEG signals [41], for each latent component, Receiver Operating Characteristic (ROC) curves were considered using different percentile ranges (5th–95th, 10th–90th, 15th–85th, 20th–80th, 25th–75th), as classification thresholds. Specifically, the values of the latent components falling within the percentile range are considered non-blink, while those outside are deemed blinks. The Area Under the Curve (AUC) is calculated, and the most discriminative latent component was identified by considering the one with the highest AUC.

(PHASE-F) The subsequent step involves modifying the values of the identified discriminative component in the test set. Several statistical parameters, like mean, median, and interquartile range, could mask this discriminative latent component for removing the eye-blink artefact. The median value is chosen based on theoretical considerations and empirical assessment. The use of the median proved superior to the mean and interquartile range in mitigating potential effects, primarily

due to its robustness to outliers. Thus, for each blink image, the discriminative latent component is set to the median value of the training data without blinks, resulting in modified reconstructed topo-maps.

(PHASE-G) Following [23], the final step involves signal reconstruction based on these modified topo-maps.

In the following sections, the various steps mentioned above are described in more detail.

4. Experimental assessment

The assessment of the automated denoising pipeline, as presented in the previous section, comprises four stages. The first phase involves preparing topographic maps from the raw EEG signals. In the second phase, these topo-maps train a subjet-specific CVAE model. Subsequently, the trained CVAE latent space is interpreted to identify the most discriminative component for eye blinking. Finally, a cleaned version of the signal is reconstructed.

4.1. EEG topology-preserving maps dataset generation

The dataset preparation phase consisted of generating labelled EEG topographic maps based on the presence of eye-blink artefacts. In this work, to transform the EEG signal into topographic maps, an initial scattered 3D topology-preserving representation is formed in a 3-dimensional space, following the coordinates of each electrode position on the scalp [39]. Azimuthal Equidistant Projection (polar) is subsequently used to transform this representation into a scattered 2D map, preserving the relative distance between adjacent electrodes [42]. Eventually, the Clough–Tocher method [43] is applied to fill the scattered 2D maps by estimating the values in-between the electrode over a new interpolated map, an image of 40 × 40. These maps were subsequently labelled based on the presence or absence of eye blinks, randomly split into training, validation, and test sets, and used to develop and evaluate a CVAE for each subject. To do this, semi-automatic blink detection was performed. This process ensured ground truth for evaluating the person-specific CVAEs. The different steps for generating the dataset are detailed below.

4.1.1. Raw EEG signals

EEG signals are part of a publicly available dataset, DEAP [36]. This contains EEG data collected from 32 subjects, recorded as each watched 40 one-minute-long excerpts of music videos. EEG data was acquired at a sampling rate of 512 Hz by using the Biosemi ActiveTwo system. For the recording, 32 EEG electrodes (‘Fp1’, ‘AF3’, ‘F7’, ‘F3’, ‘FC1’, ‘FC5’, ‘T7’, ‘C3’, ‘CP1’, ‘CP5’, ‘P7’, ‘P3’, ‘Pz’, ‘PO3’, ‘O1’, ‘Oz’, ‘O2’, ‘PO4’, ‘P4’, ‘P8’, ‘CP6’, ‘CP2’, ‘C4’, ‘T8’, ‘FC6’, ‘FC2’, ‘F4’, ‘F8’, ‘AF4’, ‘Fp2’, ‘Fz’, and ‘Cz’), placed according to the international 10–20 system, and 12 peripheral channels (such as EOG) were used. Moreover, for 22 of the 32 participants, a frontal face video was also recorded. Among these, in the context of this study, the first 7 subjects were considered. The rationale behind this decision is associated with the time-consuming nature of the subsequent phase involving the semi-automatic blink identification and labelling for each subject.

4.1.2. Eye-blink semi-automatically identification

The raw EEG signal of each of the 7 subjects is filtered with a bandpass frequency filter between 0.1 Hz to 45 Hz, reducing power line interference and focusing on the most significant frequency content of brain activity [37]. The signals were then resampled at 128 Hz (Fig. 1A). To create the labelled dataset for CVAE training, an initial analysis was performed to determine the specific instant times the eye-blink artefact occurred (Fig. 1B). Specifically, for each participant and trial, the $Fp1$ and $Fp2$ channels were picked, and a bandpass filter ranging from 1 Hz to 12 Hz was applied to isolate the influence of eye blinks. Subsequently, by leveraging on [38], an empirical subject-specific threshold was established to identify eye blinks by examining both EEG signals and face video recordings to optimize the detection of eye-blink artefacts. Then, 1 s EEG segments were selected, centred on the detected blink's identified peak. The number of such segments varies for each subject, depending on the number of blinks.

Table 1
Number of eyeblink and non-eyeblink samples for each subject.

	N° eyeblink	N° non-eyeblink
Subject #1	380	3620
Subject #2	439	3561
Subject #3	423	3577
Subject #4	450	3550
Subject #5	247	3753
Subject #6	399	3601
Subject #7	412	3588

4.1.3. EEG topographic map labelling

The related topo-maps were considered for each of 1 second EEG segments centred on eye blink. Based on the identified peaks, the topo-maps were labelled considering (i) no eye blink and (ii) eye blink. Fig. 1C reports topographic maps of a subject. Specifically, the figure displays the frontal area ($Fp1$ and $Fp2$) on the top. Since EEG topo-maps are generated from 1 s EEG segments, 128 maps were produced for each of them (since the EEG was resampled to 128 Hz). However, the corresponding number of EEG topo-maps associated with each participant differed due to variations in the total number of EEG segments extracted. From all the available data, 4000 maps for each subject were randomly selected. Table 1 reports the number of eyeblink and non-eyeblink samples for each subjects. This number was deemed reasonable given the computational capacity available during the training phase.

4.2. CVAE model development

Each subject's generated topology-preserving EEG topo maps were divided into training, validation, and test sets, with a percentage of 80%, 10%, and 10%, respectively. Possible data leakage phenomena were considered and avoided and the topographic maps associated to the test set are not used for training any more [44]. Additionally, all topographic maps were normalized with min–max normalization on a per-pixel basis, considering the minimum and maximum values derived from the entire training dataset to prevent data leakage [44].

Subject-specific CVAE architectures were designed to capture the complex patterns within EEG data (Fig. 1D). The optimal CVAE architecture is selected through a grid search by optimizing the learning rate and batch size as hyperparameters [45]. Specifically, the learning rates ranged between 10^{-7} and 10^{-5}, and batch sizes were 32, 64, and 128. The training epochs were settled at a maximum of 2500. Underfitting and overfitting are avoided with an early stopping mechanism on the validation set, with a patience value of 30. The number of components of latent space was fixed at 28, as optimally tuned in [40].

A 5-fold cross-validation strategy was adopted with a grid-search strategy to tune the model hyperparameters [45,46]. This approach is commonly employed in the literature owing to its capability for assessing the success rate of models used for classification [47,48]. The best hyperparameter configuration was selected considering the average of the Structural Similarity Index Measure (SSIM) and Mean-Squared Error (MSE) [49,50] among the five folds.

The SSIM and MSE allow the analysis of the similarity and accuracy of the reconstructed EEG topo-maps with the original topo-maps. The SSIM quantifies the image degradation by considering the change in structural information [50]. The SSIM index is calculated for various windows of an image. Formally, given two matrices (corresponding to patches of the topo-maps) x and y of size $N \times N$, the SSIM is defined as

$$\mathrm{SSIM}(x, y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)} \tag{1}$$

with μ_x the mean of x; μ_y the mean of y; σ_x^2 the variance of x; σ_y^2 the variance of y; σ_{xy} the covariance of x and y; $c_1 = (k_1 L)^2$ and $c_2 = (k_2 L)^2$ two variables to stabilize the division with weak denominator; L the dynamic range of the pixel-values (typically this is $2^{\#bits\ per\ pixel} - 1$); $k_1 = 0.01$ and $k_2 = 0.03$ by default.

Instead, the mean squared error (MSE) is calculated by considering the squared intensity differences of distorted and reference image pixels [50]. Therefore, given two images such as h(x,y) and $\hat{h}(x,y)$, the MSE is defined as

$$\text{MSE} = \frac{1}{MN} \sum_{n=0}^{M-1} \sum_{m=1}^{N} [\hat{h}(n,m) - h(n,m)]^2 \tag{2}$$

where M and N are the dimensions of the image.

4.3. Interpretability of CVAE's Latent space

Each latent component of a trained subject-specific CVAE is interpreted to identify its relation to eye-blink in EEG data. Subsequently, its value is modified to generate artefact-free topo-maps, which, in practice, means mitigating the artefacts in EEG data.

The 28 latent components of the training set were evaluated in detail by exploring their distribution. The Receiver Operating Characteristic (ROC) curves are considered for each latent component. A ROC curve is a method to show the performance of a classification model at different classification thresholds. The ROC curve identifies the relation between the True Positive Rate (TPR) and False Positive Rate (FPR). The TPR is defined by Eq. (3) and represents the proportion of actual positive instances correctly identified by a binary classification model out of the total positive instances. The FPR is defined by Eq. (4) and represents the proportion of negative cases erroneously classified as positive relative to the total true negatives.

$$\text{TPR} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}} \tag{3}$$

$$\text{FPR} = \frac{\text{False Positives}}{\text{False Positives} + \text{True Negatives}} \tag{4}$$

In this study, the ROC curves are constructed by referencing the less frequent class, the eyeblinks (anomalies), to focus on the class that is the subject of this investigation. Thus, assuming that blinks are anomalies in EEG signals [41], different percentile ranges, notably 5th–95th, 10th–90th, 15th–85th, 20th–80th, 25th–75th, are used as classification thresholds. The TPR indicates the proportion of eye-blink artefacts correctly identified, while the FPR is the proportion of no-blink classified as a blink (Fig. 1E). Then, the Area Under the Curve (AUC) was calculated to provide an aggregate performance measure across all possible classification thresholds. AUC ranges in value from 0 to 1, where 1 means perfect classification and 0 that all the elements are misclassified. In light of this, the discriminatory power of each latent component for blink identification is evaluated by leveraging AUC scores. The most discriminative latent component for eye blinking is identified by considering the one with the highest AUC.

4.4. Eye blink mitigation

After determining the most discriminative latent component responsible for capturing the information related to eye blinks in the training set, a systematic approach is employed to mitigate the impact of eye-blink artefacts on EEG signals. Specifically, the median value of the identified latent component in the no-blink training set is computed. Subsequently, considering eye-blink topo-maps in the test set, the current value of the relevant latent component is replaced with the calculated median value while leaving the other latent components unchanged (Fig. 1F). The median value is chosen after experimental assessment in which other statistical metrics, such as the mean value, were preliminarily tested. Moreover, the median value is a robust measure of central tendency, demonstrating insensitivity to extreme values or outliers within the dataset [51]. This characteristic makes it suitable for masking the latent component associated with blink artefacts that appear as outliers or extreme values in the EEG data. After modifying the topo-maps associated with eye blinks, the final step involves the reconstruction of a clean version of the EEG signal across time. Following [23], the conversion of EEG topographic maps into EEG signals is executed by extracting pixel values from the positions corresponding to the 32 EEG channels as originally attributed when topo-maps were generated from the original EEG data (Fig. 1G). To evaluate the similarity between the original EEG signal and the reconstructed clean version, the Pearson correlation between channels was calculated [52].

Table 2
SSIM and MSE scores for all subjects, averaged on 5-folds for each subject. SSIM is adimensional, while the MSE's unit is pixel intensity squared (I^2).

	SSIM (mean ± std)	MSE (mean ± std)
Subject #1	0.8556 ± 0.0128	$(1.000 \pm 0.005) \times 10^{-3}$
Subject #2	0.8373 ± 0.0115	$(1.062 \pm 0.054) \times 10^{-3}$
Subject #3	0.7223 ± 0.0135	$(4.184 \pm 0.544) \times 10^{-4}$
Subject #4	0.8705 ± 0.0217	$(3.767 \pm 0.398) \times 10^{-4}$
Subject #5	0.8013 ± 0.0249	$(1.186 \pm 0.129) \times 10^{-3}$
Subject #6	0.7202 ± 0.0157	$(8.951 \pm 0.312) \times 10^{-4}$
Subject #7	0.8282 ± 0.0084	$(6.986 \pm 0.495) \times 10^{-4}$

The Pearson correlation index is defined as:

$$r = \frac{cov(EEG_{true}, EEG_{clean})}{s_{EEG_{true}} \times s_{EEG_{clean}}} \tag{5}$$

where cov is the covariance between EEG_{true}, the original EEG signals, and EEG_{clean}, the reconstructed clean version, while $s_{EEG_{true}}$ and $s_{EEG_{clean}}$ are the standard deviation of the two signals. This metric ranges from −1 to 1, where 1 represents the highest level of correlation between the signals.

Finally, to evaluate the impact of blink artefact removal on the EEG spectral characteristics, the correlation of normalized spectral power between segments before and after artefact removal was calculated. The spectral power density for each segment was estimated using Welch's method [53]. The power spectral density was then normalized by the total power across all frequencies to obtain:

$$P_{\text{norm}}(f) = \frac{P(f)}{\sum_f P(f)} \tag{6}$$

where $P(f)$ is the power at frequency f and $\sum_f P(f)$ is the total power across all frequencies. The Pearson correlation coefficient between the normalized spectral powers of pre- and post-removal segments was computed to quantify the preservation of spectral information.

5. Results and discussion

The experimental investigation aimed at evaluating the effectiveness of the proposed subject-specific CVAE model in detecting and removing eye blinks within EEG signals. This section presents the obtained findings, showing the reconstruction performance of the CVAE in terms of SSIM and MSE and its ability to capture meaningful patterns associated with blink events.

5.1. Reconstruction performance of the CVAE

Table 2 summarizes the SSIM and MSE scores for all subjects, averaged on 5-folds for each subject. Fig. 2 illustrates the reconstruction capability of the CVAE model across the seven subjects considered by providing a visual representation of the variability in SSIM (a) and MSE (b). As can be observed, for some subjects, the variability in terms of SSIM and MSE appears to be higher (for example, in subject #5). This may be attributed to greater variability in the training data for the specific subject. Such variability in EEG patterns may be

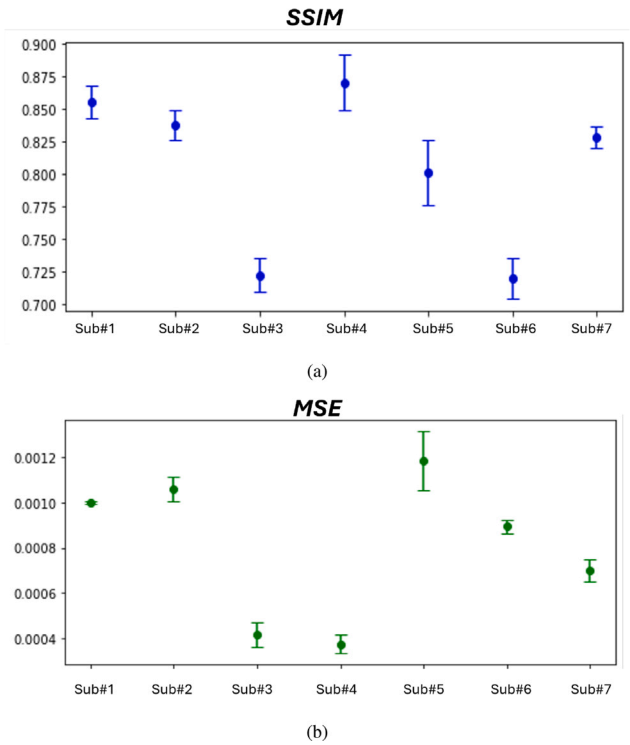

Fig. 2. CVAE's reconstruction capability through (a) boxplot of the Structural Similarity Index (SSIM) and (b) boxplot of Mean Squared Error (MSE). The boxplots highlight the distribution of variability in SSIM and MSE metrics across the 7 subjects considered. Both the metrics are dimensionless.

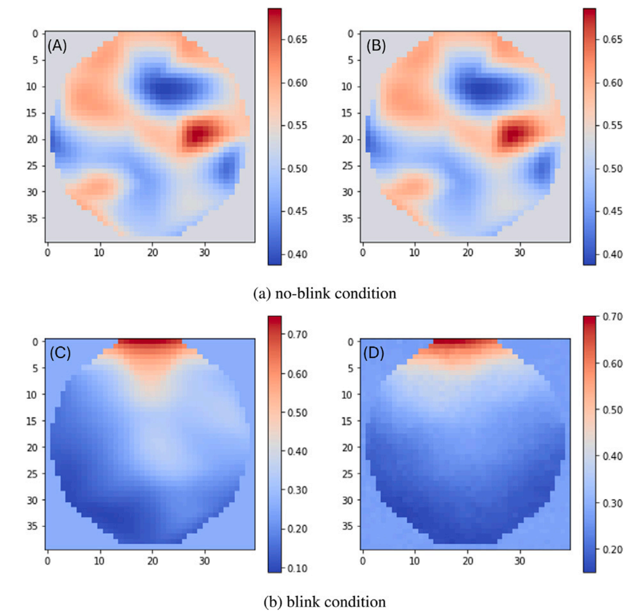

Fig. 3. EEG topomaps for one subject. No-blink condition: original (A) and reconstructed by the VAE (B). Blink condition: original (C) and reconstructed by the VAE (D). Note that the frontal channels are on the top.

Table 3
Hyperparameters and metrics (SSIM and MSE) for the best model for each subject.

	Best model	SSIM	MSE
Subject #1	learning rate: 10^{-5} batch_size: 32	0.8684	0.0011
Subject #2	learning rate: 10^{-5} batch size: 32	0.8522	0.0007
Subject #3	learning rate: 10^{-6} batch size: 32	0.7338	0.0004
Subject #4	learning rate: 10^{-5} batch size: 64	0.8911	0.0003
Subject #5	learning rate: 10^{-5} batch size: 32	0.8255	0.0012
Subject #6	learning rate: 10^{-6} batch size: 64	0.7355	0.0009
Subject #7	learning rate: 10^{-6} batch size: 64	0.8364	0.0007

associated with the subject's emotional state under analysis, as the EEG signals used were recorded. In contrast, subjects were engaged in a task (watching music videos). Instead, Fig. 3 shows an example of original EEG topographic maps and the reconstructed ones for one subject, providing a visual representation of the model performance. Table 3 shows the best model with associated SSIM and MSE for each subject.

5.2. Analysis of the latent space

Since the main goal of this study was to analyse the latent space to identify discriminative latent components for eye blinks, first, the 28 latent components on the training set were evaluated by exploring their distribution. Given the knowledge of whether each image is associated with a blink, a non-blink, or a transition, the histograms for each latent component on the training set were considered to evaluate how these events were distributed. Fig. 4 shows examples of the obtained histograms for some subjects and some latent components.

Subsequently, the ROC curves and AUC scores were analysed for each latent component of the training set to identify the most discriminative one for detecting eye blinks.

The ROC curve shows the TPR on the y-axis (refer to Eq. (3)), representing the proportion of correctly identified eye blinks, and the FPR on the x-axis (refer to Eq. (4)), which is the proportion of instances without blinks classified as blinks (further details are described in Section 4.3). By considering the AUC scores, the discriminatory capacity of each latent component for eye blinks was assessed, with the latent component most discriminative identified as having the highest AUC.

Table 4 reports the AUC scores for each component and subject, with the highest AUC value, indicated in bold. Subsequently, as described in Section 4.4, the median value of the identified latent component in only the topo maps without a blink in the training set was considered. In other words, the median value was calculated only on the topo-maps of the training set that are not associated with blinks. Considering only the topographic maps in the test set associated with eye blinks, the current value of the relevant latent component was replaced with the calculated median value while leaving the other latent components unchanged. Fig. 5 reports some topo-maps obtained after this procedure, also showing cases where the algorithm partially mitigates blink or fails in eye blink removal.

5.3. EEG signal reconstruction in time

A cleaned version of the EEG signal across time was reconstructed, such as in [40], to evaluate the blink mitigation (Fig. 1G).

Specifically, the focus was on the frontal channels $Fp1$ and $Fp2$, which are most affected by the eye-blink artefact. To evaluate the fidelity of reconstructing a clean version of the EEG signal while minimizing information loss, the Pearson correlation index (refer to Eq. (5)) was computed between the original and the reconstructed signals. The Fig. 6 shows the original signal EEG_{true} and the clean version EEG_{clean} for channels $Fp1$ (column#1) and $Fp2$ (column#2).

As can be seen, the eye blink results are mitigated for channels $Fp1$ and $Fp2$, and the correlation value is high, showing no excessive

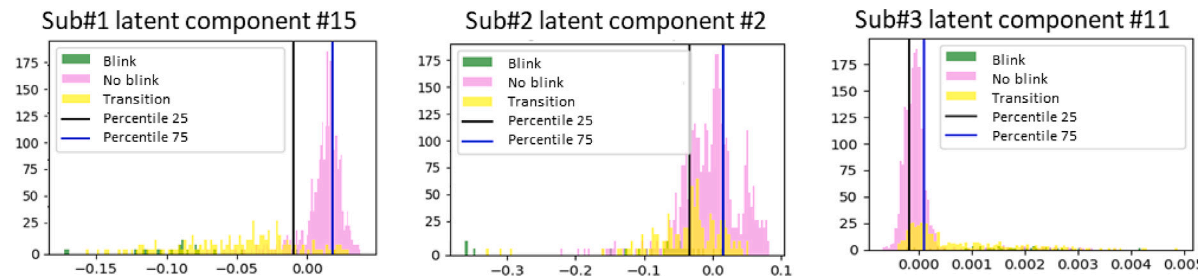

Fig. 4. Examples of histograms for some subjects and some latent components on the training set. The histograms show the distribution of blinks (green), no-blinks (pink), and transitions (yellow). The 25th–75th percentile range are also reported, noting that blinks (green) are usually out of this range. The y axis represents the number of elements for each bin of the histogram, while the x axis represents the value of latent components. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

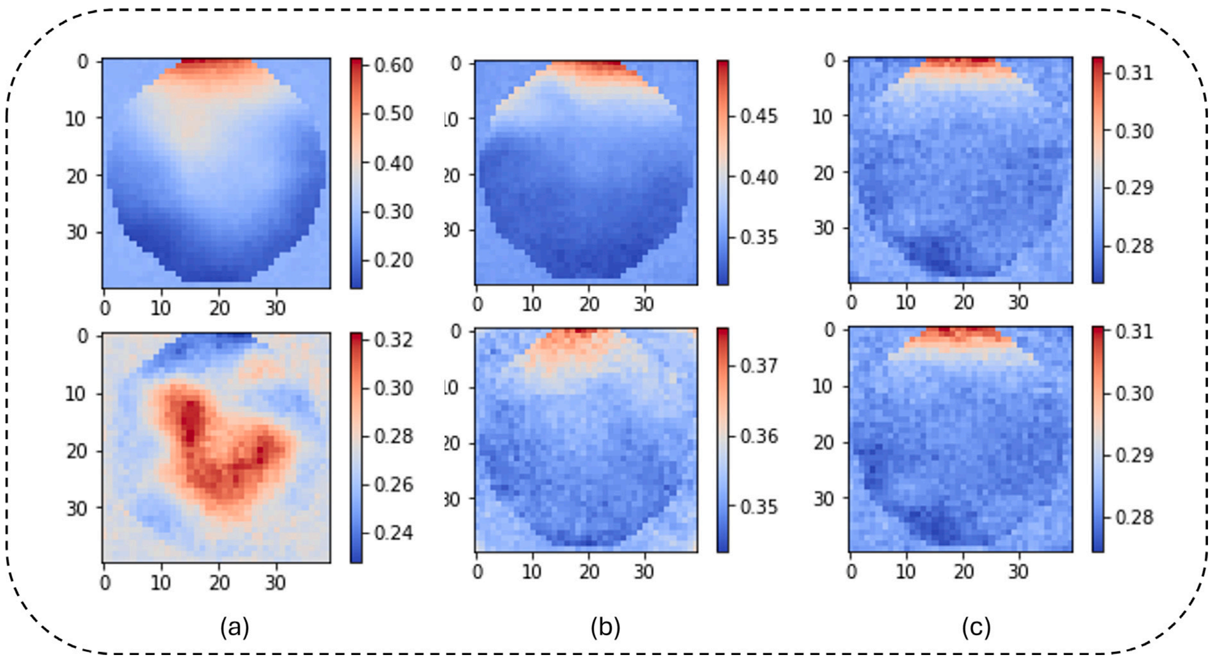

Fig. 5. Examples of topo-maps obtained after mitigating eye-blink with the proposed method: (a) blink is completely removed, (b) blink is partially mitigated, (c) the algorithm fails in eye blink removal.

loss of information related to the brain signal. Fig. 6 also includes the Pz channel (column#3) for comparative analysis. The selection of the central Pz channel is motivated by its lower susceptibility to eye-blink artefacts. Upon visual inspection, it is apparent that the reconstructed signal in the Pz channel is closely overlaid with the original signal. This visual assessment is further supported by a Pearson correlation value close to 1 for the Pz channel, signifying negligible loss of EEG signal. Table 5 shows the mean and standard deviation for Pearson correlation, calculated for $Fp1$, $Fp2$ and Pz channels for each subject. Specifically, the correlation value was averaged over all 1 s EEG segments around the identified blinks (see Section 4). In summary, it can be observed through both visual inspection and the Pearson correlation coefficients that the blink artefact is effectively mitigated in the $Fp1$ and $Fp2$ channels without an excessive loss of EEG signal information.

5.4. Analysis of frequency domain impact

To evaluate the impact of the proposed blink artefact removal method on EEG spectral characteristics, the correlation of normalized spectral power between segments before and after artefact removal were computed (refer to Eq. (6)). Specifically, also in this case, the focus was on the frontal channels $Fp1$ and $Fp2$, and Pz. As stated in Section 5.3, the frontal channels $Fp1$ and $Fp2$ are most affected by the eye-blink artefact. On the other hand, the selection of the central Pz channel is motivated by its lower susceptibility to eye-blink artefacts. Table 6 shows the mean and standard deviation for the calculated correlation. As can be observed, the correlations across all subjects and channels are generally high (all above 0.68), indicating that the blink removal process effectively preserves the frequency domain characteristics of the EEG. This indicates that spectral information in the EEG signals is retained following the removal of artefacts. The Pz channel consistently demonstrates high correlations (all above 0.81), indicating that the blink removal technique is highly effective at preserving spectral characteristics in the region of the scalp that is less susceptible to eye-blink artefacts.

The developed method for denoising EEG signals by using CVAE, introduces a novel approach to mitigate artefacts related to eye blinks. The methodology focuses on analysing the latent space, identifying the latent component most discriminative for eye blink, and adjusting its value accordingly, as described in Section 4. The obtained results showed that the proposed approach is promising in effectively reducing eye-blink artefacts, resulting in a improved quality of the reconstructed signal. The main advantage of this method is its automation. In contrast to conventional techniques such as ICA and ASR, the approach outlined in the present study necessitates minimal user intervention. As a matter of fact, the users do not manually select parameters, and the method

Table 4
AUC score for each latent component and each subject. The highest AUC values is reported in bold, identifying the most discriminative latent component for eye blink.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Subject #1	0.89	0.86	0.85	0.83	0.65	0.89	0.85	0.88	0.90	0.68	0.84	0.85	0.75	0.89
Subject #2	0.63	0.77	0.46	0.41	0.36	0.67	**0.98**	0.65	0.37	0.46	0.42	0.38	0.69	0.41
Subject #3	0.88	**0.96**	0.95	0.95	0.93	0.93	0.94	0.95	0.94	0.84	0.95	0.95	0.95	0.89
Subject #4	0.83	0.91	0.91	0.91	0.40	0.91	0.91	0.91	0.91	0.93	**0.95**	0.91	0.91	0.93
Subject #5	0.93	0.95	0.94	0.94	0.94	0.93	0.93	0.95	0.94	0.56	**0.98**	0.96	0.96	0.60
Subject #6	0.95	0.95	0.95	**0.96**	0.94	0.95	0.95	0.95	0.94	0.94	0.94	0.95	0.94	0.95
Subject #7	0.99	0.99	0.99	0.99	0.99	0.99	**1.00**	0.89	0.98	0.99	0.99	0.99	0.99	0.69
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
Subject #1	**0.93**	0.86	0.81	0.85	0.90	0.69	0.88	0.84	0.90	0.89	0.88	0.79	0.88	0.67
Subject #2	0.69	0.63	0.57	0.47	0.70	0.60	0.51	0.48	0.60	0.41	0.67	0.52	0.59	0.68
Subject #3	0.94	0.91	0.94	0.95	0.94	0.95	0.94	0.94	0.95	0.95	0.94	0.94	0.90	0.95
Subject #4	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91
Subject #5	0.95	0.84	0.96	0.96	0.64	0.95	0.58	0.62	0.95	0.94	0.95	0.91	0.96	0.95
Subject #6	0.95	0.94	0.94	0.95	0.94	0.94	0.94	0.60	0.95	0.94	0.95	0.94	0.95	0.94
Subject #7	0.99	0.56	0.63	0.99	0.99	0.40	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.88

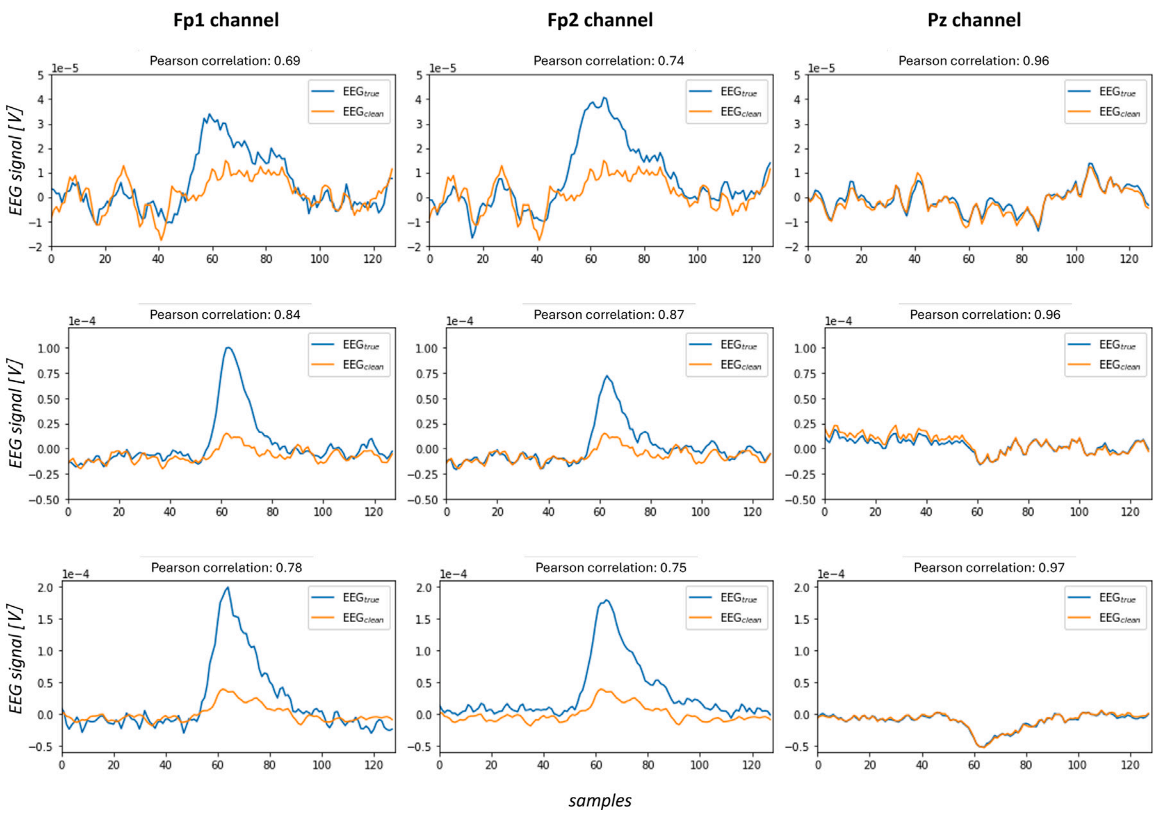

Fig. 6. Comparison between EEG_{true} (blue) and EEG_{clean} (orange) for channels $Fp1$ (first column), $Fp2$ (second column) and Pz (third column) for 1-s EEG segment for different subject (row). For each plot, on the x-axis and on y-axis, the samples of the EEG signals (128 samples) and the EEG signal value [μV] are reported. The Pearson correlation value is also shown. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 5
Pearson correlation value (mean ± std) for each subject. Specifically, the correlation value was averaged over all 1 s EEG segments around the identified blinks (see Section 4).

	Fp1	Fp2	Pz
Subject #1	0.64 ± 0.19	0.66 ± 0.19	0.91 ± 0.10
Subject #2	0.57 ± 0.17	0.50 ± 0.15	0.88 ± 0.07
Subject #3	0.69 ± 0.12	0.77 ± 0.19	0.90 ± 0.15
Subject #4	0.66 ± 0.20	0.65 ± 0.17	0.91 ± 0.09
Subject #5	0.75 ± 0.20	0.77 ± 0.19	0.85 ± 0.15
Subject #6	0.61 ± 0.20	0.54 ± 0.25	0.80 ± 0.20
Subject #7	0.65 ± 0.20	0.66 ± 0.22	0.88 ± 0.10

Table 6
Pearson correlation value (mean ± std) for normalized spectral power between segments before and after artefact removal for each subject. Specifically, the correlation value was averaged over all 1 s EEG segments around the identified blinks (see Section 4).

	Fp1	Fp2	Pz
Subject #1	0.81 ± 0.20	0.83 ± 0.15	0.95 ± 0.12
Subject #2	0.88 ± 0.19	0.89 ± 0.12	0.90 ± 0.10
Subject #3	0.80 ± 0.05	0.82 ± 0.09	0.97 ± 0.11
Subject #4	0.70 ± 0.32	0.76 ± 0.17	0.96 ± 0.04
Subject #5	0.85 ± 0.12	0.87 ± 0.10	0.95 ± 0.05
Subject #6	0.79 ± 0.21	0.68 ± 0.28	0.81 ± 0.19
Subject #7	0.91 ± 0.13	0.92 ± 0.06	0.94 ± 0.06

does not need a baseline clean signal for artefact removal. The proposed method allows an automatic selection of the most discriminative latent component for eye blink based on the ROC curve analysis. However, some limitations need to be taken into account. Notably, the proposed method operates offline, rendering it unsuitable for real-time applications, at present. Additionally, the neural network training phase introduces a time-consuming aspect to the methodology.

6. Conclusion

This study explored the adoption of Convolutional Variational Autoencoders (CVAEs) for automatically detecting and removing blink artefacts from EEG signals. Specifically, subject-specific CVAE models capable of representing the underlying structure of EEG data, with a focus on blink-related patterns, were developed. The investigation focused on analysing the latent space of such CVAEs trained on topology-preserving EEG maps. Starting with preprocessed signals from 32 EEG channels, including only band-pass filtering and raw data resampling, EEG topographic maps labelled based on eye blinks were generated for each subject. These maps were randomly divided into training, validation, and test sets to develop robust subject-specific CVAE models. Model reconstruction performance was assessed by considering SSIM and MSE. Subsequently, the latent space was explored to identify those components responsible for distinguishing between blink and non-blink events through statistical measures. Considering eye blinks as anomalies in EEG signals, ROC curves, taking into account different percentile ranges, were constructed for each latent component and the AUC scores were calculated. By leveraging AUC scores, the discriminatory power of each latent component responsible for eye blinks was evaluated. The most discriminative latent component for eye blinking was identified on the training set by considering the one with the highest AUC.

Subsequently, the value of this identified discriminative component was modified in the test set by setting the median value of the distribution of the same component associated with the portion of the training data without blinks, resulting in modified topographic maps. The final step involved reconstructing a clean version of the EEG signal based on these modified maps. Visual inspection of EEG signals and Pearson correlation values for the frontal channels $Fp1$ and $Fp2$ showed that the proposed methodology offers promising results in the automatic eye blink detection and removal with a minimal loss of EEG signal information.

This paper introduces a semi-automated, unsupervised, offline denoising pipeline for detecting and mitigating blinks in EEG signals, thereby making a significant contribution to the body of knowledge in neural signal processing.

However, some drawbacks exist to this.

For instance, this study was conducted by considering a limited sample size (7 subjects). This choice was driven by practical constraints, particularly the time-intensive nature of the labelling process. Future studies will involve expanded sample sizes, allowing better generalization of the results.

This investigation focused on employing the median value for modifying the latent component values in eye-blink topomaps. Future studies could explore alternatives to the median adjustment technique, by considering alternative statistical measures or other filtering methods to enhance denoising capabilities. Additionally, future directions might extend the evaluation of false positives and false negatives to gain a more profound insight into the algorithm's performance. This expansion could involve considering supplementary evaluation metrics and methodologies to comprehensively analyse classification errors.

This study employed the highest AUC score to identify the most discriminative latent component for eye blink detection. However, given that the highest AUC may not always indicate the most effective component when several scores are similar, future research will tackle the challenge of selecting among multiple components with comparable AUC scores. This will be achieved by integrating additional criteria, such as statistical significance, to ensure a more robust and accurate component selection process.

To enhance denoising performance, also the generation of the latent space through an automatic smart grid search across its size could be exploited.

While this study relied on visual inspection and the Pearson correlation index to quantify denoising effectiveness, forthcoming research could employ Signal-to-Noise Ratio (SNR) and Peak Signal-to-Noise Ratio (PSNR) metrics. This would allow for a meticulous assessment of the fidelity of the reconstructed EEG signal compared to the original, specifically within and outside blink regions.

The present study mainly considered the $Fp1$ and $Fp2$ channels as being influenced by eye blinks. Future work aims to include other channels in close proximity to $Fp1$ and $Fp2$. Moreover, future works could extend the methodology to other types of physiological artefacts and also to discriminate between involuntary and voluntary blink, widely used in reactive brain–computer interfaces [54–56] for specific commands.

Finally, in future studies, the potential of transfer learning and domain adaptation techniques [57] will be investigated to develop robust models able to perform across several subjects.

CRediT authorship contribution statement

Sabatina Criscuolo: Writing – original draft, Visualization, Validation, Supervision, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Andrea Apicella:** Writing – review & editing, Visualization, Validation, Supervision, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Roberto Prevete:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Luca Longo:** Writing – review & editing, Supervision, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

This work was supported by the European Union - FSE-REACT-EU, PON Research and Innovation 2014–2020, DM 1061/2021 contract number DOT19X7NYL-2, and DM 1062/2021 contract number 18-I-15350-2, Ministry of University and Research PRIN research project "BRIO – BIAS, RISK, OPACITY in AI: design, verification and development of Trustworthy AI.", Project no. 2020SSKZ7R, and by the PNRR MUR project PE0000013-FAIR (CUP: E63C22002150007).

The Authors would like to thank Taufique Ahmed, from the Artificial Intelligence and Cognitive Load Research Laboratory, School of Computer Science, Technological University Dublin, Ireland, for his support in the coding part of the implementation of the proposed method.

References

[1] J.C. Henry, Electroencephalography: basic principles, clinical applications, and related fields, Neurology 67 (11) (2006) 2092.

[2] W.O. Tatum IV, Handbook of EEG Interpretation, Springer Publishing Company, 2021.

[3] P. Arpaia, A. Esposito, L. Gargiulo, N. Moccaldi, Wearable Brain-Computer Interfaces: Prototyping EEG-Based Instruments for Monitoring and Control, CRC Press, 2023.

[4] J.A. Urigüen, B. Garcia-Zapirain, EEG artifact removal—state-of-the-art and guidelines, J. Neural Eng. 12 (3) (2015) 031001.

[5] X. Jiang, G.-B. Bian, Z. Tian, Removal of artifacts from EEG signals: a review, Sensors 19 (5) (2019) 987.

[6] W. Mumtaz, S. Rasheed, A. Irfan, Review of challenges associated with the EEG artifact removal methods, Biomed. Signal Process. Control 68 (2021) 102741.

[7] A. Cataldo, S. Criscuolo, E. De Benedetto, A. Masciullo, M. Pesola, R. Schiavoni, S. Invitto, A method for optimizing the artifact subspace reconstruction performance in low-density EEG, IEEE Sens. J. 22 (21) (2022) 21257–21265.
[8] R. Ranjan, B.C. Sahana, A.K. Bhandari, Ocular artifact elimination from electroencephalography signals: A systematic review, Biocybern. Biomed. Eng. 41 (3) (2021) 960–996.
[9] K.A. Robbins, J. Touryan, T. Mullen, C. Kothe, N. Bigdely-Shamlo, How sensitive are EEG results to preprocessing methods: a benchmarking study, IEEE Trans. Neural Syst. Rehabil. Eng. 28 (5) (2020) 1081–1090.
[10] M. Abo-Zahhad, S.M. Ahmed, S.N. Abbas, A new EEG acquisition protocol for biometric identification using eye blinking signals, Int. J. Intell. Syst. Appl. 7 (6) (2015) 48.
[11] P. Arpaia, E. De Bendetto, A. Esposito, A. Natalizio, M. Parvis, M. Pesola, Comparing artifact removal techniques for daily-life electroencephalography with few channels, in: 2022 IEEE International Symposium on Medical Measurements and Applications (MeMeA), IEEE, 2022, pp. 1–6.
[12] P. Arpaia, A. Esposito, A. Natalizio, M. Parvis, M. Pesola, Low-density EEG correction with multivariate decomposition and subspace reconstruction, IEEE Sens. J. (2023).
[13] W. Sun, Y. Su, X. Wu, X. Wu, A novel end-to-end 1D-ResCNN model to remove artifact from EEG signals, Neurocomputing 404 (2020) 108–121.
[14] H. Zhang, M. Zhao, C. Wei, D. Mantini, Z. Li, Q. Liu, EEGdenoiseNet: a benchmark dataset for deep learning solutions of EEG denoising, J. Neural Eng. 18 (5) (2021) 056057.
[15] P. Sawangjai, M. Trakulruangroj, C. Boonnag, M. Piriyajitakonkij, R.K. Tripathy, T. Sudhawiyangkul, T. Wilaiprasitporn, EEGANet: Removal of ocular artifacts from the EEG signal using generative adversarial networks, IEEE J. Biomed. Health Inf. 26 (10) (2021) 4913–4924.
[16] A.V. Chikkankod, L. Longo, On the dimensionality and utility of convolutional autoencoder's latent space trained with topology-preserving spectral EEG head-maps, Mach. Learn. Knowl. Extraction 4 (4) (2022) 1042–1064, http://dx.doi.org/10.3390/make4040053, URL https://www.mdpi.com/2504-4990/4/4/53.
[17] B. Yang, K. Duan, C. Fan, C. Hu, J. Wang, Automatic ocular artifacts removal in EEG using deep learning, Biomed. Signal Process. Control 43 (2018) 148–158.
[18] N.M.N. Leite, E.T. Pereira, E.C. Gurjao, L.R. Veloso, Deep convolutional autoencoder for EEG noise filtering, in: 2018 IEEE International Conference on Bioinformatics and Biomedicine, BIBM, IEEE, 2018, pp. 2605–2612.
[19] J.F. Hwaidi, T.M. Chen, A noise removal approach from eeg recordings based on variational autoencoders, in: 2021 13th International Conference on Computer and Automation Engineering, ICCAE, IEEE, 2021, pp. 19–23.
[20] T. Ahmed, L. Longo, Latent space interpretation and visualisation for understanding the decisions of convolutional variational autoencoders trained with eeg topographic maps, in: Joint Proceedings of the XAI-2023 Late-Breaking Work, Demos and Doctoral Consortium, Co-Located with the 1st World Conference on EXplainable Artificial Intelligence (XAI-2023), CEUR Workshop Proceedings, 2023, pp. 65–70.
[21] D. Bank, N. Koenigstein, R. Giryes, Autoencoders, in: Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook, Springer, 2023, pp. 353–374.
[22] D.M. Blei, A. Kucukelbir, J.D. McAuliffe, Variational inference: A review for statisticians, J. Am. Stat. Assoc. 112 (518) (2017) 859–877.
[23] T. Ahmed, L. Longo, Interpreting disentangled representations of person-specific convolutional variational autoencoders of spatially preserving eeg topographic maps via clustering and visual plausibility, Information 14 (9) (2023) 489.
[24] Q. Wang, J.M. Loh, X. He, Y. Wang, A latent state space model for estimating brain dynamics from electroencephalogram (EEG) data, Biometrics 79 (3) (2023) 2444–2457.
[25] M.M.N. Mannan, M.A. Kamran, M.Y. Jeong, Identification and removal of physiological artifacts from electroencephalogram signals: A review, Ieee Access 6 (2018) 30630–30652.
[26] S. Romero, M.A. Mañanas, M.J. Barbanoj, A comparative study of automatic techniques for ocular artifact reduction in spontaneous EEG signals based on clinical target variables: a simulation case, Comput. Biol. Med. 38 (3) (2008) 348–360.
[27] Y.-H. Wang, C.-H. Yeh, H.-W.V. Young, K. Hu, M.-T. Lo, On the computational complexity of the empirical mode decomposition algorithm, Phys. A 400 (2014) 159–167.
[28] B. Yang, K. Duan, T. Zhang, Removal of EOG artifacts from EEG using a cascade of sparse autoencoder and recursive least squares adaptive filter, Neurocomputing 214 (2016) 1053–1060.
[29] D.E. Rumelhart, G.E. Hinton, R.J. Williams, et al., Learning internal representations by error propagation, 1985.
[30] L. Rokach, O. Maimon, E. Shmueli, Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook, Springer Nature, 2023.
[31] R. Ghosh, N. Sinha, S.K. Biswas, Automated eye blink artefact removal from EEG using support vector machine and autoencoder, IET Signal Process. 13 (2) (2019) 141–148.
[32] C.-H. Chuang, K.-Y. Chang, C.-S. Huang, T.-P. Jung, IC-U-Net: a U-Net-based denoising autoencoder using mixtures of independent components for automatic EEG artifact removal, NeuroImage 263 (2022) 119586.
[33] H.-A.T. Nguyen, T.H. Le, T.D. Bui, A deep wavelet sparse autoencoder method for online and automatic electrooculographical artifact removal, Neural Comput. Appl. 32 (2020) 18255–18270.
[34] S. Saba-Sadiya, E. Chantland, T. Alhanai, T. Liu, M.M. Ghassemi, Unsupervised eeg artifact detection and correction, Front. Digit. Health 2 (2021) 608920.
[35] L. Longo, M. Brcic, F. Cabitza, J. Choi, R. Confalonieri, J.D. Ser, R. Guidotti, Y. Hayashi, F. Herrera, A. Holzinger, R. Jiang, H. Khosravi, F. Lecue, G. Malgieri, A. Páez, W. Samek, J. Schneider, T. Speith, S. Stumpf, Explainable artificial intelligence (XAI) 2.0: A manifesto of open challenges and interdisciplinary research directions, Inf. Fusion 106 (2024) 102301, http://dx.doi.org/10.1016/j.inffus.2024.102301.
[36] S. Koelstra, C. Muhl, M. Soleymani, J.-S. Lee, A. Yazdani, T. Ebrahimi, T. Pun, A. Nijholt, I. Patras, Deap: A database for emotion analysis; using physiological signals, IEEE Trans. Affect. Comput. 3 (1) (2011) 18–31.
[37] M.X. Cohen, Analyzing Neural Time Series Data: Theory and Practice, MIT Press, 2014.
[38] D.-K. Tran, T.-H. Nguyen, T.-N. Nguyen, Detection of EEG-based eye-blinks using a thresholding algorithm, Eur. J. Eng. Technol. Res. 6 (4) (2021) 6–12.
[39] L. Longo, Modeling cognitive load as a self-supervised brain rate with electroencephalography and deep learning, Brain Sci. 12 (10) (2022) http://dx.doi.org/10.3390/brainsci12101416, URL https://www.mdpi.com/2076-3425/12/10/1416.
[40] T. Ahmed, L. Longo, Examining the size of the latent space of convolutional variational autoencoders trained with spectral topographic maps of EEG frequency bands, IEEE Access 10 (2022) 107575–107586.
[41] S. Çınar, N. Acır, A novel system for automatic removal of ocular artefacts in EEG by using outlier detection methods and independent component analysis, Expert Syst. Appl. 68 (2017) 36–44.
[42] M. Sougranttika, J.K.D. Halim, S. Foong, H.L. Ng, C. Kang, S. Maryam, F. Chan, 3D human head parametric estimation and EEG localization with 2D azimuthal equidistant map projection, in: 2022 IEEE/ASME International Conference on Advanced Intelligent Mechatronics, AIM, IEEE, 2022, pp. 830–837.
[43] S. Mann, Cubic precision clough-tocher interpolation, Comput. Aided Geometr. Des. 16 (2) (1999) 85–88.
[44] A. Apicella, F. Isgrò, R. Prevete, Don't push the button! exploring data leakage risks in machine learning and transfer learning, 2024, arXiv preprint arXiv:2401.13796.
[45] C. Bishop, Neural networks for pattern recognition, Clarendon Press Google Schola 2 (1995) 223–228.
[46] S. Geisser, The predictive sample reuse method with applications, J. Am. Stat. Assoc. 70 (350) (1975) 320–328.
[47] C. Rao, Y. Wu, Linear model selection by cross-validation, J. Statist. Plann. Inference 128 (1) (2005) 231–240.
[48] I.K. Nti, O. Nyarko-Boateng, J. Aning, Performance of machine learning algorithms with different K values in K-fold cross-validation, J. Inf. Technol. Comput. Sci. 6 (2021) 61–71.
[49] Z. Wang, A.C. Bovik, H.R. Sheikh, E.P. Simoncelli, Image quality assessment: from error visibility to structural similarity, IEEE Trans. Image Process. 13 (4) (2004) 600–612.
[50] U. Sara, M. Akter, M.S. Uddin, Image quality assessment through FSIM, SSIM, MSE and PSNR—a comparative study, J. Comput. Commun. 7 (3) (2019) 8–18.
[51] D.C. Montgomery, G.C. Runger, Applied Statistics and Probability for Engineers, John wiley & sons, 2010.
[52] I. Cohen, Y. Huang, J. Chen, J. Benesty, J. Benesty, J. Chen, Y. Huang, I. Cohen, Pearson correlation coefficient, in: Noise Reduction in Speech Processing, Springer, 2009, pp. 1–4.
[53] L. Zhao, Y. He, Power spectrum estimation of the welch method based on imagery EEG, Appl. Mech. Mater. 278 (2013) 1260–1264.
[54] P. Arpaia, S. Criscuolo, E. De Benedetto, N. Donato, L. Duraccio, Evaluation of the effectiveness of a wearable, AR-based BCI for robot control in ADHD treatment, in: 2022 IEEE International Conference on Metrology for Extended Reality, Artificial Intelligence and Neural Engineering (MetroXRAINE), IEEE, 2022, pp. 630–634.
[55] L. Angrisani, P. Arpaia, A. Esposito, N. Moccaldi, A wearable brain–computer interface instrument for augmented reality-based inspection in industry 4.0, IEEE Trans. Instrum. Meas. 69 (4) (2019) 1530–1539.
[56] A. Apicella, P. Arpaia, E. De Benedetto, N. Donato, L. Duraccio, S. Giugliano, R. Prevete, Enhancement of SSVEPs classification in BCI-based wearable instrumentation through machine learning techniques, IEEE Sens. J. 22 (9) (2022) 9087–9094.
[57] A. Apicella, P. Arpaia, G. D'Errico, D. Marocco, G. Mastrati, N. Moccaldi, R. Prevete, Machine learning strategies to improve generalization in EEG-based emotion assessment: a systematic review, 2022, arXiv preprint arXiv:2212.08744.