POLITECNICO DI TORINO
Repository ISTITUZIONALE

AdS Carroll Chern-Simons supergravity in 2+1 dimensions and its flat limit

Original
AdS Carroll Chern-Simons supergravity in 2+1 dimensions and its flat limit / Ravera, L.. - In: PHYSICS LETTERS. SECTION B. - ISSN 0370-2693. - ELETTRONICO. - 795:(2019), pp. 331-338. [10.1016/j.physletb.2019.06.026]

Availability:
This version is available at: 11583/2876430 since: 2022-01-12T13:19:35Z

Publisher:
Elsevier B.V.

Published
DOI:10.1016/j.physletb.2019.06.026

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright
Elsevier postprint/Author's Accepted Manuscript

© 2019. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
http://creativecommons.org/licenses/by-nc-nd/4.0/.The final authenticated version is available online at:
http://dx.doi.org/10.1016/j.physletb.2019.06.026

(Article begins on next page)

27 November 2026

Biomedical Signal Processing and Control 126 (2026) 110855

Contents lists available at ScienceDirect

Biomedical Signal Processing and Control

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

Automatic wave recurrences extraction for single-channel decomposition

Luca Mesin

Mathematical Biology and Physiology, Dept. Electronics and Telecommunications, Politecnico di Torino, Turin, Italy

ARTICLE INFO

Keywords:
Recurrent waveform detection
Time-series decomposition
Blind source separation
Template-based signal reconstruction
Unsupervised signal analysis
Adaptive signal processing

ABSTRACT

Recurring waveforms in time-series reflect the underlying system dynamics and are common in physiological signals. Identifying these recurrences enables the signal decomposition into meaningful components. Automatic wave recurrences extraction (AWaRE) is a novel low-cost decomposition method for single time-series pre-processing, based on recurrent waveform detection. AWaRE follows a three-step pipeline: (1) detecting candidate recurrent waveforms via matching signal fragments to the trace, (2) resolving overlapping occurrences using multi-kernel deconvolution, and (3) iterative refining of residuals to identify additional waveform onsets. AWaRE was evaluated on simulated data with varying waveform counts, recurrence rates and noise levels, and compared against two methods with similar computational cost: ensemble empirical mode decomposition with independent component analysis (EEMD-ICA) and a spike sorting approach. Mean absolute reconstruction errors in percentage of the mean absolute of the original signal (given as median ± inter-quartile range) were 85.3 ± 30.8% (EEMD-ICA), 65.5 ± 15.0% (spike sorting) and 24.3 ± 35.9% (AWaRE); waveform train estimation errors were 111.0 ± 34.5%, 103.2 ± 33.4% and 32.6 ± 47.1%; processing times for 10 s of data were 2.05 ± 3.13 s, 0.94 ± 0.27 s and 0.78 ± 2.00 s, respectively. All differences were statistically significant in the paired comparisons. Then, examples of applications in fetal ECG detection, intramuscular EMG decomposition, EEG artifact removal and neural spike sorting are provided, showing that the automatically extracted waveforms have clear physiological meaning. In conclusion, AWaRE enables adaptive time-series decomposition by extracting recurrent waveforms from the signal. It is an efficient and robust tool for biomedical signal analysis, facilitating the study of the activities of different sources through repeated waveform pattern identification.

1. Introduction

The repetition of specific waveforms in a time-series is a clear manifestation of determinism in the system generating it. Indeed, the analysis of the recurrences of patterns is the basis of theories providing quantitative measurements of determinism (recurrence quantification analysis [1]) and chaos [2], and predicting the dynamics of nonlinear systems [3]. Moreover, many biological and physiological signals reflect the repetitive activation of different sources, e.g., the ECG manifesting the repetition of heartbeats [4,5], the firings of neurons in micro-electrode recordings (MER) [6], or the activation of motor units (MU) measured by the electromyogram (EMG) [7].

When many sensors are available to detect the activities of the same sources, different statistical methods are available for their discrimination, e.g., blind source separation (BSS) by independent components analysis (ICA) and its variants [8,9]. However, sometimes the mixture of different sources is measured by a single-channel, and separating their contributions becomes a great challenge. Different approaches have been introduced to address this decomposition problem, sometimes adapted to a specific application (e.g., spike sorting of neural data [10] or intramuscular EMG decomposition [11]). Significant efforts have been made to validate decomposition methods, e.g., for MERs [12]. In fact, neuro-scientific and clinical applications increasingly rely on the precise analysis of single-channel recordings, e.g., to study the firings of neurons at a high resolution and to discriminate their coordinated activity [13]. These signals can represent a mixture of neural or MU activities (in MER and intramuscular EMG, respectively), and single-channel decomposition techniques attempt to separate them [14,15].

In this work, a novel approach is introduced to represent a single time-series as the sum of distinct components. It is based on the premise that the time-series is produced by a deterministic system, which inherently produces recurring waveforms. These recurrent waveforms are identified and exploited to decompose the time-series and to gain insight into the deterministic system that generates it. The process is called automatic wave recurrences extraction (AWaRE). By identifying these recurrent patterns, different contributions can be effectively separated without imposing additional assumptions regarding

E-mail address: luca.mesin@polito.it.

https://doi.org/10.1016/j.bspc.2026.110855
Received 24 November 2025; Received in revised form 2 May 2026; Accepted 13 June 2026
Available online 20 June 2026
1746-8094/© 2026 The Author. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Abbreviations

ANOVA	Analysis of variance
AWaRE	Automatic Wave Recurrences Extraction
BSS	Blind source separation
CDL	Convolutional dictionary learning
CWT	Continuous wavelet transform
DL	Deep learning
DTW	Dynamic time warping
ECG	Electrocardiogram
EEG	Electroencephalogram
EEMD	Ensemble empirical mode decomposition
EMD	Empirical mode decomposition
EMG	Electromyogram
FECG	Fetal ECG
GMM	Gaussian Mixture Models
HMM	Hidden Markov model
ICA	Independent component analysis
IMF	Intrinsic mode function
MER	Micro-electrode recording
ML	Machine learning
MU	Motor unit
MUAP	Motor unit action potential
PCA	Principal component analysis
PSD	Power spectral density
RICA	Robust independent component analysis
RMS	Root mean squared
SCICA	single-channel independent component analysis
SNR	Signal-to-noise ratio
STFT	Short time Fourier transform
TFR	Time-frequency representation
VMD	Variational mode decomposition
WICA	Wavelet independent component analysis

their temporal shapes or bandwidths. An efficient method for detecting recurrences in a time-series is proposed, which offers a general framework for isolating and distinguishing the waveforms that compose the studied signal. This approach is based on cross-correlating the signal with specific templates, an operation that can be performed at a low computational cost (i.e., logarithmic with respect to the number of samples of the signal, instead of the quadratic cost needed to compare all pairs of possible near neighbors). This idea has already been exploited in other methods, such as template matching [16] or wavelet transform [17], but here the prototype waveforms are different portions of the signal itself. Alternative approaches for identifying waveforms within a signal include methods based on thresholding (to detect portions exceeding a given level) and clustering (such as the aforementioned approaches for spike sorting [10] or EMG decomposition [11]). AWaRE differs from these methods in its waveform selection criterion, which is based on recurrence rather than amplitude alone, thereby enabling the detection of even low-amplitude but consistently recurring contributions. The only prerequisite is that the sources generate contributions that repeat over time and occasionally appear in isolation, thereby enabling the identification of their recurrence. In addition, a method was used to address and resolve residual overlaps between waveforms originating from different sources. Another relevant approach for decomposing a single time series resolving overlaps is Convolutional Dictionary Learning (CDL), which formulates the problem as a computationally intensive optimization to jointly estimate both waveforms and their activation times. Compared to CDL, AWaRE is characterized by a significantly lower computational cost.

The remainder of this paper is organized as follows: Section 2 discusses the proposed methodology and the testing signals (used to benchmark it against two alternative methods), the results are shown in Section 3, Section 4 discusses the findings and limitations (including a discussion in relation to existing literature), and Section 5 concludes the study with key insights. Finally, two appendices are included: the first reviews additional methods from the literature, and serves as background for readers unfamiliar with the approaches referenced in the main text; the second presents theoretical considerations on AWaRE.

2. Methods

AWaRE is based on three parts: identification of recurrences (implicitly interpreted as the result of the activity of different sources), separation of overlaps and final refinement.

2.1. Identification of recurrences

The innovative algorithm is based on a set of template matching operations, where the templates are portions of the signal used to look for recurrences. Specifically, starting from the beginning of the signal, a group of consecutive samples is selected to form a waveform that is used as a template for comparison with the entire signal. The length of the template (called the embedding dimension in the theory of nonlinear time-series) is a parameter, indicated by ED, which can also be a vector that includes different values to be tested. The template is centered with respect to its centroid. If the template has a small amplitude, it is skipped (the amplitude threshold and specific norm to measure it are selected by the user, indicated as A_{th} and $Norm$, respectively).

Portions of the data already identified as recurrences of previously used templates are omitted. Matches are identified if the root mean squared (RMS) difference with the template is smaller than the threshold M_{th} (which is another parameter selected by the user).

The estimation of recurrences, which would generally incur a quadratic cost, is made efficient by leveraging the following observation. The RMS difference is computed using three terms: the energy of the template (which is fast to compute), the energy of the considered portion of the signal and the cross-correlation between the template and the signal. Note that the last two terms can be computed using numerical convolutions [18], which can be obtained efficiently in the Fourier transform domain.

After these steps, recurrences of different patterns are found. However, the results could be affected by several possible problems:

- the templates are portions of raw data, that are directly affected by noise, so that some matches with other portions of the signal could be missed;
- some recurrences could be missed due to noise or artifacts superimposed to the signal or the overlap of different waveforms;
- some waveforms could be found as recurrences of more templates considered different as they do not satisfy the similarity criterion;
- some portions of the signal could be similar to the template only due to chance which makes two noisy portions of the signal similar (it is reasonable to expect that they are found only a few times).

To compensate for those problems (at least in part), templates with some common recurrences were merged together. Moreover, recurrent waveforms were aligned (in the Fourier domain to get maximal time resolution) and averaged. The obtained cleaned templates were then used again to find recurrences. Finally, the waveforms with less than a minimum number of recurrences (indicated by the selectable parameter N_{WF}) were removed.

2.2. Solving waveforms overlaps

Given the templates of recurrent waveforms from the previous part, the corresponding firing patterns were refined to resolve possible overlaps (which is a problem frequently found in intramuscular EMG decomposition or spike sorting [19]). A modified version of the multi-kernel deconvolution method proposed in [20] was used. The processed signal is approximated using the identified recurrent waveforms (considered as the multiple kernels of the deconvolution method). The best solution is obtained by solving an optimization problem that minimizes the reconstruction error functional. As noticed in [21], the mean squared error is affected by outliers and too tolerant to small values. Using the L_1 instead of the L_2 norm, the solution becomes more stable to outliers and is sparse [22], better resembling a firing pattern. Thus, L_1 deconvolution with regularization was implemented [23], considering the recurrent waveforms as kernels $\{K_i(t)\}$ convoluted with their firing patterns $\{f_i(t)\}$

$$\arg\min_{\hat{f}_i(t)} \left\| s(t) - \sum_{i=1}^{N_K} K_i(t) * \hat{f}(t) \right\|_1 + \alpha \left\| \hat{f}(t) \right\|_1 \tag{1}$$

where α is the regularization parameter chosen as detailed below. The problem was discretized as in [21,24], writing convolution as the multiplication of the unknown firing patterns with a matrix A, with $AX \approx b$, where A is the matrix collecting the samples of the waveforms, $b = \{b_i\} = \{s(t_i)\}$ is the vector of recorded data samples and X the vector of unknown firing patterns (all concatenated):

$$A = \begin{bmatrix} K_1 & K_2 & \cdots & K_N \end{bmatrix}$$
$$\text{where} \quad K_i = \begin{bmatrix} K_i[1] & 0 & \cdots & 0 \\ K_i[2] & K_i[1] & \ddots & \vdots \\ \vdots & \vdots & \ddots & 0 \\ K_i[M] & \cdots & K_i[2] & K_i[1] \end{bmatrix} \tag{2}$$

and

$$X = \begin{bmatrix} f_1(t) \\ f_2(t) \\ \vdots \\ f_N(t) \end{bmatrix} \tag{3}$$

where the case of N kernels is considered.

The iteratively reweighted least squares (IRLS) method was employed as a computational step within the proposed framework to solve the L_1 problem (using 10 iterations). IRLS is a well-established approach for solving optimization problems involving sparsity or robustness constraints. It has been successfully applied in various domains, including medical image reconstruction [25], large-scale statistical estimation [26], and sparse signal recovery [23]. The following error functional was minimized

$$\|W_e(AX - b)\|^2 + \alpha \|W_x X\|^2 = \left\| \begin{bmatrix} W_e A \\ \sqrt{\alpha} W_x \end{bmatrix} X - \begin{bmatrix} W_e b \\ 0 \end{bmatrix} \right\|^2 \tag{4}$$

where W_e and W_x are diagonal matrices including the reciprocal of the L_1 norm of the error $AX - b$ and of the solution X, respectively. The previous expression can be written as

$$\|BX - c\|^2 \tag{5}$$

where

$$B = \begin{bmatrix} W_e A \\ \sqrt{\alpha} W_x \end{bmatrix}, \qquad c = \begin{bmatrix} W_e b \\ 0 \end{bmatrix} \tag{6}$$

The solution of the regularized problem can be obtained analytically as

$$X = (B^T B)^{-1} B^T c = (A^T W_e^T W_e A + \alpha W_x^T W_x)^{-1} A^T W_e^T W_e b \tag{7}$$

The regularization parameter α was chosen as the 1% of the maximum eigenvalue of $A^T A$, which allowed to stabilize the inversion operation. A fine tuning on a few signals have been performed, showing stable results across a wide range (from about 2% to 0.2% of the maximum eigenvalue of $A^T A$).

This solution was iteratively computed to converge toward a weighted L_2 norm that approximates the L_1 norm. At each iteration, the result was passed through a sigmoid function defined as

$$\sigma(x) = \frac{1}{1 + \exp[-10(x - 0.5)]}, \tag{8}$$

which maps values into the interval $[0, 1]$ and accentuates the tendency toward the extremes of this range. After convergence, a binarization step was finally applied: values greater than 0.9 were set to 1, and all others to 0. This binarized output defines the presence or absence of the recurrence of a waveform, enabling reconstruction of the original signal using only the predefined kernels (without altering their amplitudes).

Notice that the problem has a quadratic computational cost with respect to the number of kernels and length of the processed signal segments. As indicated in [21], shortening the processed epochs (including an overlap and then building the final solution) allows to reduce the computational cost. In this study, the processed epoch was twice the maximal ED and the overlap was 50%.

2.3. Final refinement

Once the recurrent waveforms have been identified and any overlaps resolved, the signal can be reconstructed accordingly. Nevertheless, some instances of these waveforms may still have been missed. This can occur because of unresolved overlaps involving multiple waveforms, or due to factors such as noise or slight temporal misalignment (e.g., a half-sample shift), which may cause the RMS error to marginally exceed the detection threshold.

To address this issue, the residual, obtained after subtracting the estimated signal, was re-examined by scanning (by template matching) for one identified recurrent waveform (kernel) at a time. The patterns were processed in descending order of amplitude. To increase the sensitivity, the matching threshold M_{th} was increased by 20%, allowing for the detection of additional occurrences that may have been missed previously. This adjustment also accounts for possible distortions introduced into the residual by the subtraction of the estimated and raw signals. Any newly identified recurrence was incorporated into the existing sequence for the corresponding pattern and removed from the residual before analyzing the next pattern.

2.4. Summary of the method

AWaRe is based on three main parts, i.e., identification of recurrent waveforms, separation of overlaps and final refinement, with the first further split into different steps. Figs. 1 and 2 show these steps in an example, i.e., a simulated signal, including 8 waveforms (detailed in the following), each with 30 random occurrences in a time-series of 10 s and additional noise. The steps are:

1. preliminary identification of recurrences;
2. removal of waveforms with few recurrences (i.e., less than N_{WF}, possibly given by noisy portions, overlaps of more waveforms or not recurrent patterns) and robust estimation of their shape by averaging;
3. estimation of recurrences of the new estimated waveforms;
4. removal of similar waveforms;
5. solving overlaps;
6. refinement on residual.

Notice that the first 4 steps constitute the first part of the method (described in Section 2.1), which outputs a set of recurrent waveforms. They are shown in Fig. 1. The second and third parts of AWaRE (Sections 2.2 and 2.3, respectively) are in Fig. 2, applied to the same signal. Notice that the second part, devoted to overlap treatment, is the

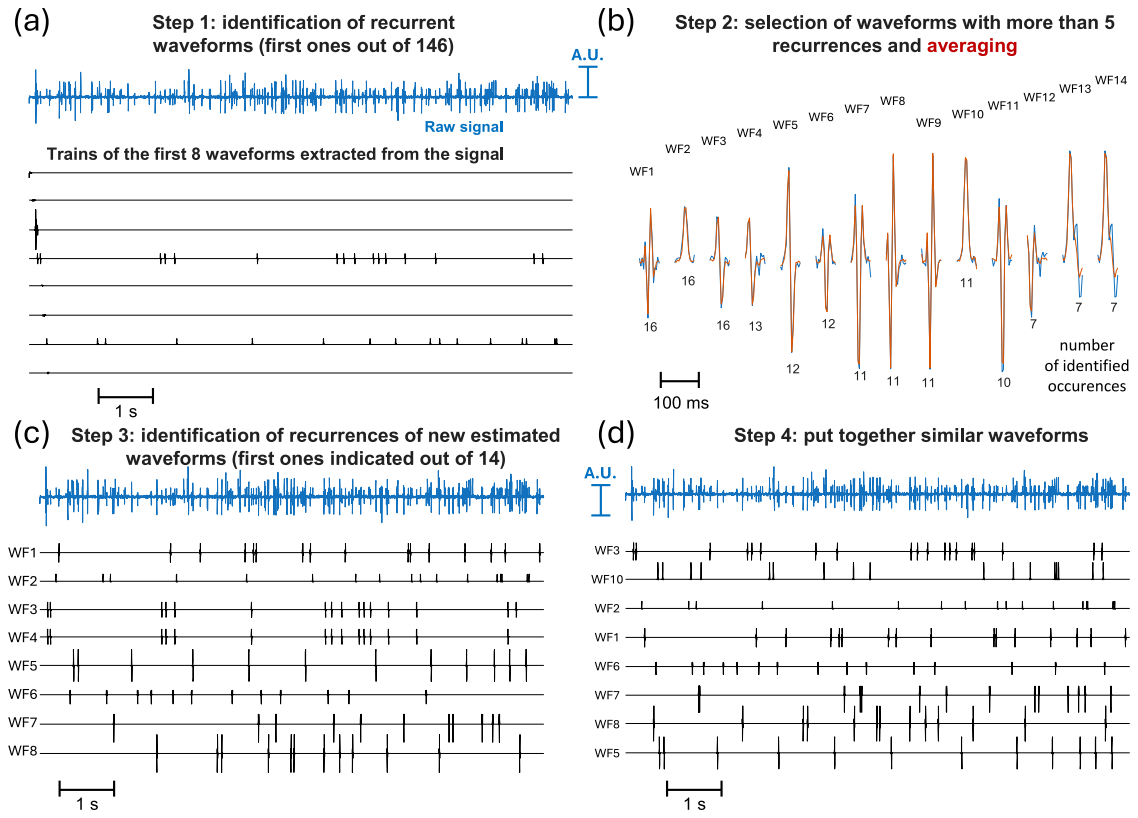


Fig. 1. First part of the steps of the algorithm AWaRE. It is applied to a simulated signal of 10 s, sampled at 500 Hz, including 8 different waveforms, appearing 30 times, to which a white Gaussian noise was superimposed with SNR = 15 dB. The following parameters were used (see Table 1 for the definition of these parameters): $A_{th} = 0.025$, $Norm = \overline{|s(t)|}$, $ED = [15\ 31]$ (vector of these two elements), $M_{th} = 0.15$, $N_{WF} = 5$. (a) Identification of recurrences of all patterns, considering all centered portions starting from the beginning of the signal. The raw signal is shown in blue on the top of the panel and the first waveforms trains are in black, below the raw signal (the same format is used also for panels (c), (d), and for (e) of Fig. 2). Most of the patterns are found only once. Processing time: 757 ms. (b) Patterns with few recurrences are removed, obtaining 14 waveforms (out of the 146 found in the first step; waveforms are ordered according to the number of identified occurrences and called WF followed by a number). Residual waveforms are cleaned by averaging. Both original selected waveforms and cleaned versions are shown, ordered in terms of the number of identified firings. Processing time: 69 ms. (c) Estimation by template matching of the recurrences of the clean waveforms obtained in the previous step. The names of the shown waveforms are indicated. Processing time: 59 ms. (d) Integration of similar waveforms, cleaning by averaging and identification by template matching of their onsets in the signal. Notice that some waveforms are removed (being integrated with others) and their order is changed (according to the new number of occurrences). Processing time: 6 ms.

one requiring the longest processing time in the shown example (the duration is strongly dependent on the number of identified recurrent waveforms).

Some theoretical considerations on the method are also provided in Appendix B.

Another example of application is shown in Fig. 3. A simple simulated signal is considered, including 8 waveforms (as in Figs. 1 and 2), each with 20 random occurrences in a time-series of 10 s, out of which a 2 s epoch is shown. The identified waveforms are presented in the order of detection. Waveform trains are displayed either superimposed or individually. Once a waveform is identified along with its recurrences, new subsequent parts of the signal are investigated, avoiding all portions already matched. Each newly investigated segment of the signal is centered to identify a well-balanced waveform.

The list of the parameters of the methods is provided in Table 1. Notice that the ranges of feasible values strongly depend on the specific application. A sensitivity analysis is provided in Appendix B. In summary, if reasonable ranges are selected, results show that the reconstruction error is only marginally affected by parameter variations (with mean effect sizes below 5%), indicating good robustness of the method. The main impact concerns computational cost, which increases particularly for lower and higher values of N_{WF} and M_{th}, respectively. Notably, unlike threshold-based approaches, AWaRE is largely insensitive to the amplitude threshold, as waveform selection is primarily driven by recurrence rather than amplitude.

2.5. Test of performances

Simulations were used to test the innovative algorithm in comparison with two other approaches. Subsequently, a few different example applications on experimental data have been prepared.

2.5.1. Simulations

Eight different waveforms were simulated and then used to build single-channel time-series. The waveforms were a Gaussian function with standard deviation $\sigma = 2.5$ ms truncated between ± 10 ms and its first, second and third derivatives, rescaled in order to have either RMS amplitude equal to 1 or to 2. Each waveform was convoluted with a firing pattern (i.e., a random permutation of a binary series with a fixed number of zeros and ones, the latter indicating the firing instants), thus obtaining a train of delayed repetitions for each. These waveform trains were then summed up to get the time-series to be processed.

The numbers of waveforms used and firings were varied across different simulations. White Gaussian noise with a specific signal to noise ratio (SNR) was finally added, with different values of SNR across simulations.

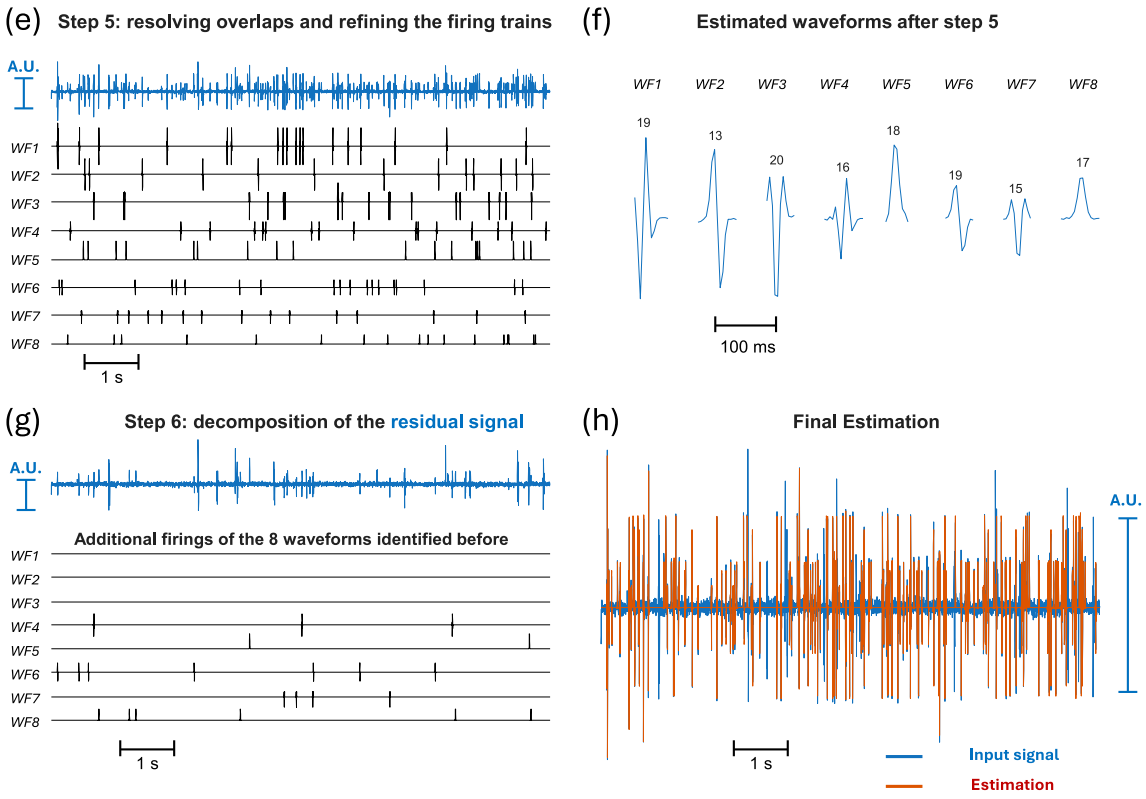


Fig. 2. Second and third parts of the algorithm AWaRE, applied to the same signal shown in Fig. 1 and with the same parameters. (e) Resolution of overlaps by deconvolution and updating of the firing patterns. New firings are found and accordingly new waveforms are obtained by averaging, now renamed as in (f). Processing time: 3805 ms. (f) New waveforms obtained after solving overlaps, now ordered according to their RMS amplitude. (g) Processing of the residual of the signal (shown on top, in blue) after subtracting the identified waveform trains. Each waveform is searched again in the residual (by template matching), starting from the largest, and removed before considering the subsequent waveform. Additional firings are shown. Processing time: 53 ms. (h) Final result: original (blue) and reconstructed (dark red) signals are shown.

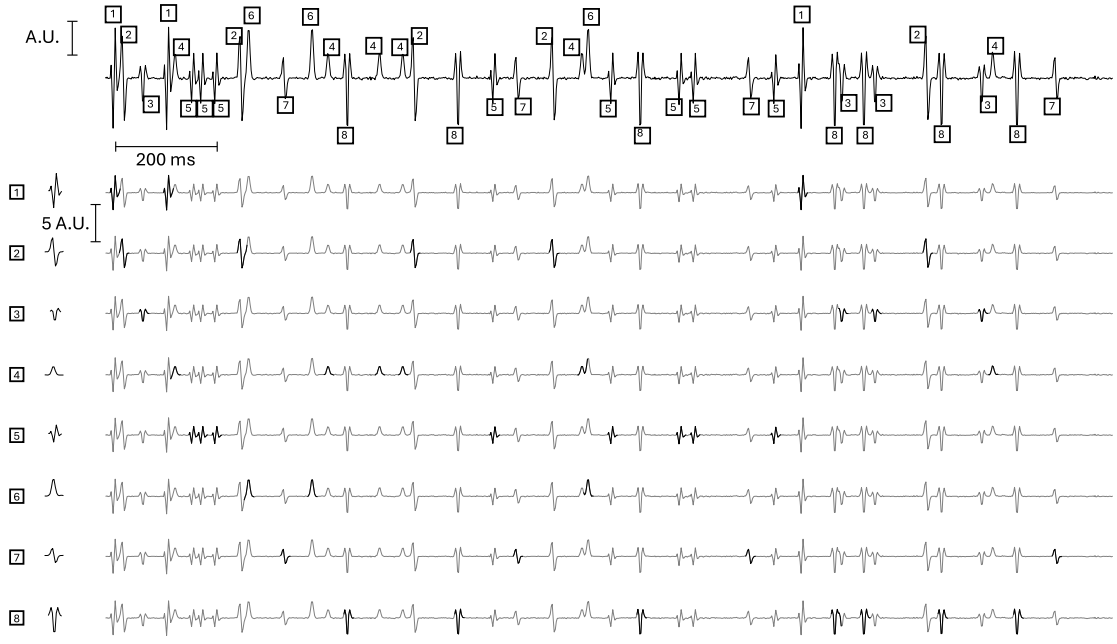


Fig. 3. Example of application of AWaRE to a simulated signal (including 8 waveforms, repeated 20 times in 10 s and including white Gaussian noise with SNR of 25 dB). The same parameters of AWaRE were used as in the previous Figures. Two seconds of the raw signal and the identified labeled waveforms are shown above, in the order of appearance (and of identification); single waveform trains are shown below for each of the waveforms (the 8 simulated waveforms are indicated on the left of each train).

Table 1
List of AWaRE parameters - ARV: average rectified value; RMS: root mean squared value; mean value of a time series $x(t)$ indicated by $\overline{x(t)}$.

Parameter	Description	Range of reasonable choices
$Norm$	Norm to measure the amplitude threshold	$ARV = \overline{\|s(t)\|}$, $RMS = \sqrt{\overline{\|s(t)\|^2}}$, $Range\|s(t)\|$
A_{th}	Amplitude threshold above which a waveform is considered	0.01–0.05 (normalized with respect to the norm of the signal)
ED	Embedding dimension (duration of the portions of the signal tested for recurrences)	10–50
M_{th}	RMS difference between a tested portion of the signal and the template waveform to indicate a recurrence	0.1–0.25 (normalized with respect to the RMS of the signal)
N_{WF}	Minimum number of recurrences of a waveform	3–10 (depending on the duration of the signal)

The signals were sampled at 500 Hz and the duration of each simulation was 10 s.

2.5.2. Alternative approaches for comparison

AWaRE was compared to two other methods, representing different methodological approaches. Simple methods have been selected, with similar computational costs as AWaRE.

- Ensemble empirical mode decomposition with ICA (EEMD-ICA; see Appendix A for details). The EEMD [27] was implemented, considering 230 ensembles with Gaussian white noise with a standard deviation equal to 1% of that of the signal to be processed and a maximum of 110 intrinsic mode functions (IMF; the parameters were chosen by Bayesian optimization minimizing the RMS error of the estimation of the waveform trains on a signal including 8 waveforms, with 40 firings in 10 s, SNR of 5 dB). FastICA was then applied to the modes extracted by the EEMD [28]. The component accounting for maximal energy of the original signal was selected and its peaks were identified, obtaining a firing pattern. The corresponding waveform was then extracted using spike-triggered averaging and convolved with the firing pattern to reconstruct the waveform train. This was then subtracted from the signal and the process was iterated to extract another component until reaching either the maximum number of 8 iterations or 75% of signal energy recovered.
- Spike sorting pipeline. The unsupervised method introduced in [29] was used. Spikes were detected using a median-based threshold (i.e., a threshold factor equal to 3 multiplied by the noise estimate, which was the median of the absolute value of the signal divided by 0.6745 [30]). They were then aligned and standardized using z-score normalization. Principal component analysis (PCA) was then applied to reduce dimensionality, keeping components explaining 95% of the variance, with a limit of 15 components. The peak amplitude and the principal components were used as features for clustering that separated different waveforms (k-means, with the Gap Statistic [31] to determine the optimal number of clusters).

2.5.3. Simulated dataset and statistical analysis

Different simulations were performed using the following parameters:

- Number of waveforms included in the signal $\{2, 4, 6, 8\}$;
- Number of occurrences per waveform $\{10, 20, 30, 40\}$;
- SNR levels $\{5, 10, 15, 20\}$ dB;
- Trial repetitions 10.

Different performance indices were estimated from AWaRE, EEMD-ICA and spike sorting:

- RMS error in reconstructing the original signal;
- mean RMS error of estimating the simulated waveform trains (after getting the best match between estimated and correct waveform trains with the Munkres Assign Algorithm [32] applied to the matrix of cross-correlation of simulated and estimated trains);
- computational cost (i.e., mean time to process 10 s of data).

The algorithm was implemented in MATLAB R2024b and computations were performed on a fixed workstation featuring an AMD Ryzen 7 9700X CPU (8 physical cores, 16 logical threads, 3.8 GHz) and 64 GB of RAM. A dedicated GPU was not used.

The data were organized in matrices, each with dimensions $4 \times 4 \times 4 \times 10$ (i.e., the dimensions of the 4 parameters indicated above). Paired comparisons were performed to compute the difference between the indices obtained using either AWaRE or one of the others (EEMD-ICA or spike sorting). The normality of the distributions was assessed using the Lilliefors, Anderson–Darlingand and Jarque–Bera tests. Given the systematic violation of the normality assumption in all tests ($p<0.05$), a monotonic transformation was applied as proposed in [33], resulting in positive normality indications across all tests. A three-way analysis of variance (ANOVA) with quadratic interactions was then performed to test the possible significant effects of the following factors:

- Number of waveforms (4 levels),
- Number of spikes (4 levels),
- SNR (4 levels).

When significant effects were disclosed, the Wilcoxon signed rank test was applied to the raw data to look for statistical differences between AWaRE and each of the two reference methods.

2.5.4. Examples of experimental data

Tests on experimental data have been prepared based on public datasets. They are simply intended as example of possible applications, in which exploring repetitions of arbitrary waveforms can be useful to decompose the data into components, which could possibly help interpretation. The following time-series were considered.

- An abdominal ECG recording from the PhysioNet dataset [34] was used to test the possibility of separating maternal and fetal ECG [35]. A single abdominal recording (channel number 4 of file `r08.edf`) was considered and compared to the direct fetal ECG (in contrast to the original paper [35], in which a Generalized Singular Value Decomposition was used to integrate the information from abdominal recordings to facilitate the separation of the mother and fetal ECGs).
- Intramuscular EMGs were taken from the "Examples of Electromyograms" dataset, available on PhysioNet [34,36]. They were recorded using a 25 mm concentric needle electrode placed in the tibialis anterior muscle, during a gentle dorsiflexion of the foot against resistance. Data were obtained from a healthy subject and two patients, with myopathy or neuropathy, respectively.

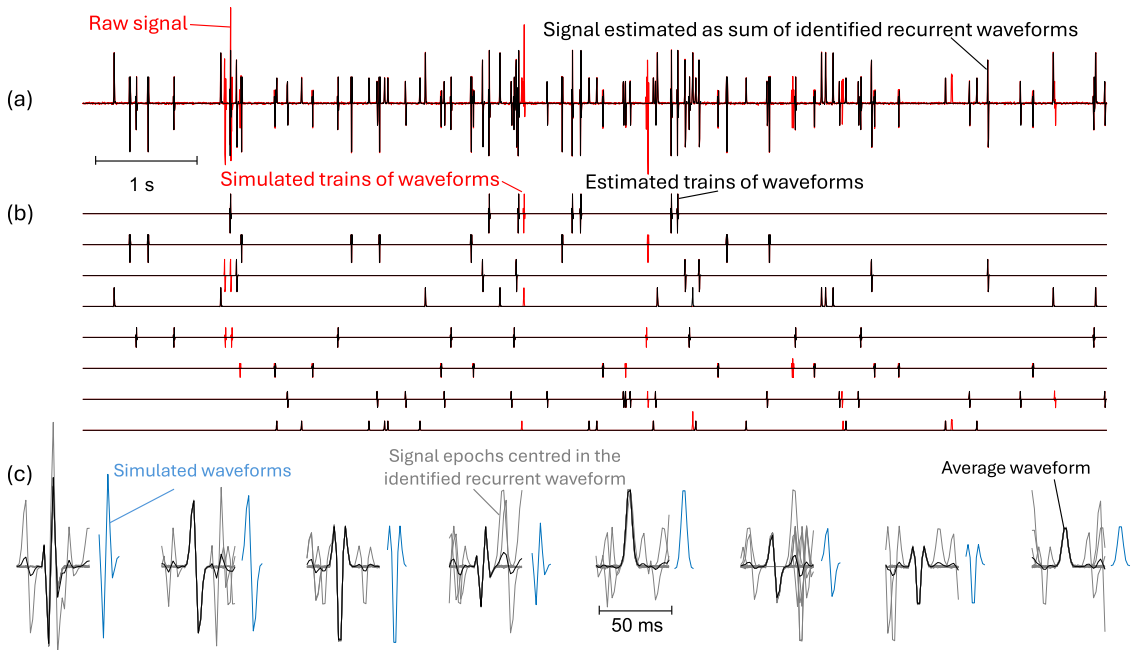


Fig. 4. Example of application of AWaRE to a simulated signal (including 8 waveforms, repeated 10 times in 10 s and including white Gaussian noise with SNR of 25 dB). The following parameters were used: $A_{th} = 0.025$, $Norm = \overline{|s(t)|}$, $ED = [15\ 31]$ (vector of these two elements), $M_{th} = 0.15$, $N_{WF} = 3$. (a) Raw signal (red) and reconstruction (black) from the estimated repeated waveforms. (b) Simulated and estimated waveform trains. (c) Simulated waveforms and estimation by averaging, with all aligned epochs in gray and average in black.

- An EEG signal with ocular artifacts was constructed by summing clean data epochs with artifact-contaminated segments, each concatenated into continuous sequences. From the EEGdenoiseNet dataset [37], 10 artifacts with a cross-correlation larger than 92.5% were selected. EEG and the train of ocular artifacts were normalized and summed after scaling the amplitude of the latter by a factor of 2.5. The objective was to identify the artifact, which was the only clear repeated wave.
- A simulation of MER. The file `C_Difficult1_noise01.mat` was used, taken from the *SYNTH_MONOTRODE* dataset, generated using the model [38], which is included as part of the SpikeForest platform [12]. The first 10 s of the signal were processed.

3. Results

Fig. 4 shows an example of the application of AWaRE to a simulated signal. The entire 10 s signal is shown, indicating the simulated and estimated waveform trains. Some waveforms, possibly due to superposition with others, have not been correctly identified, either leaving them unresolved or misclassified. Epochs including the identified onset of a specific waveform are also shown aligned in Fig. 4(c): some superpositions can be clearly seen.

Fig. 5 shows an example of comparison of AWaRE with EEMD-ICA and the spike sorting approach described in the Methods section. Only 4 waveforms were simulated. The alternative methods have problems in separating correctly the different waveform trains: EEMD-ICA is not designed specifically for the identification of spiking sources and finds only 2 waveforms; the spike sorting approach does not include a method to solve overlaps and finds some spurious contributions. AWaRE shows good performance in estimating the waveform trains for this specific simulation, confirming the results anticipated by the previous figures.

Fig. 6 displays the results of many tests on simulations similar to the specific one considered in Fig. 5. The parameters varied to generate more simulations were the number of different waveforms included in the signal, the number of occurrences in the 10 s long time-series and the SNR of the superimposed additive white Gaussian noise. As performance indices, the reconstruction errors for both raw data and waveform trains, along with the processing time, were considered. AWaRE was compared with EEMD-ICA and spike sorting, considering the difference between the performance indices. As the data were not normally distributed, they were transformed monotonically to obtain Gaussian residuals. ANOVA was then applied to the transformed data and it was assessed that all factors had significant effects. The Wilcoxon signed rank test was applied to check for possible significant effects when comparing AWaRE to the two reference methods. Most comparisons showed significant differences; only the exceptions are highlighted in the figure. The reconstruction errors (of both the raw signal and waveform trains) are always lower when using AWaRE. The median processing time is lower than that of the other methods, but with large variability, as the resolution of overlaps can be computationally expensive when many waveforms are identified.

The final figures show a few sample experimental applications. Fig. 7 provides FECG estimation. The processed data, shown in (b), includes two distinct waveforms. The parameters of the method were tuned to ensure that two waveforms were identified. Then the peaks of the direct FECG, shown in (a), were compared with those identified by the train of the fetal waveform estimated from the processed data. Some false negatives and positives were obtained, usually missing an overlap (as in (d) on the left) or misinterpreting a deformation of the mother ECG with the superposition of a FECG (as in (d) on the right). Note that no a-priori assumptions about the regularity of heartbeats or the minimum interval between two beats have been imposed; those constraints could help prevent missing a spike or detecting two firings that are too close together. For this specific test, the sensitivity of the method is 99.1% (close to the average obtained in [35] considering the entire dataset, i.e., about 99.0%, where information from more abdominal recordings were integrated to facilitate the separation of the mother and fetal ECGs); the positive predictive value was 94.4% (lower than the 99.0% reached in [35], so that AWaRE makes more false positives, that could be potentially reduced by imposing constraints, if it was to be fit to this specific application).

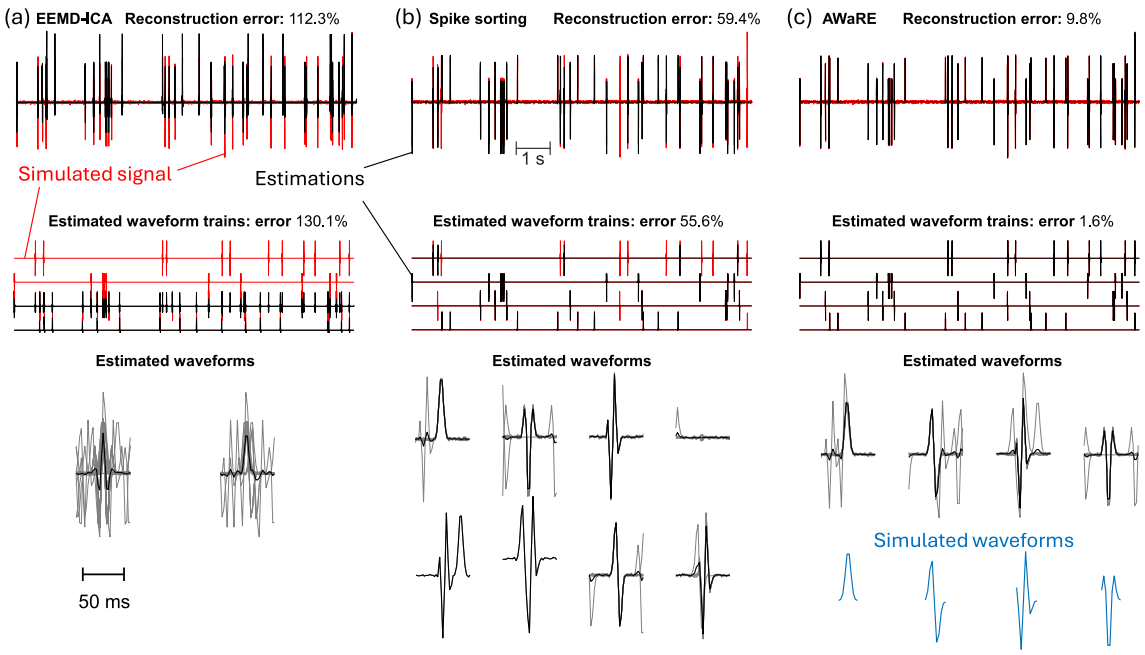


Fig. 5. Comparison of different methods for single-channel decomposition on a specific simulated signal, including 4 waveforms, repeated 10 times in 10 s and including white Gaussian noise with SNR of 20 dB. The parameters used for EEMD-ICA and spike sorting are indicated in Section 2.5.2. The following parameters were used for AWaRE (same as for the previous figure): A_{th} = 0.025, $Norm = \overline{|s(t)|}$, ED = [15 31] (vector of these two elements), M_{th} = 0.15, N_{WF} = 3. Processing time for 10 s of signal were 2.01 s for EEMD-ICA, 1.45 s for Spike Sorting and 0.93 s for AWaRE. Simulated and estimated signal and waveform trains are in red and black color, respectively.

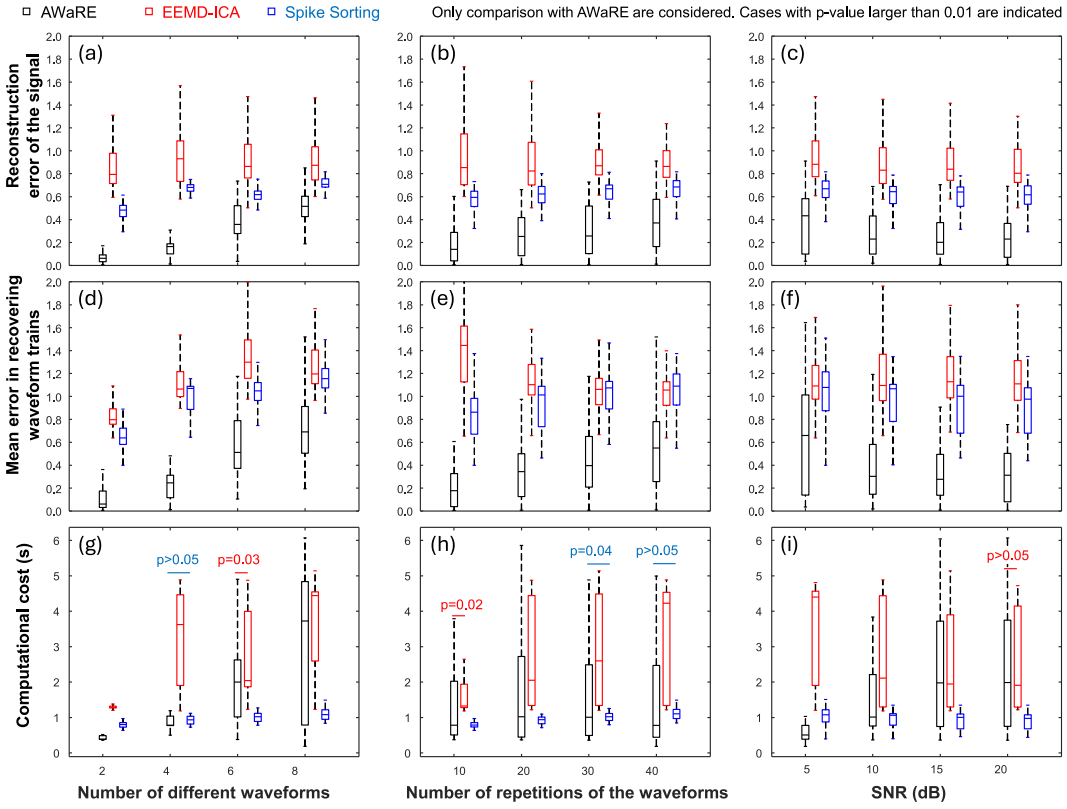


Fig. 6. Results of many tests in simulation to compare AWaRE with EEMD-ICA and a specific spike sorting approach (same parameters were used for the 3 methods as for Fig. 5). Performance indices are the RMS error in reconstructing the original signal (a–c), the mean RMS error of estimating the simulated waveform trains (d–f) and the mean time to process the 10 s simulated signals (g–i). Results are aggregated by factors, which are the number of waveforms (a, d, g), the number of occurrences per waveform (b, e, h) and the SNR level (c, f, i). For each condition, 10 random trials were simulated, obtaining matrices of results with dimensions $4 \times 4 \times 4 \times 10$. ANOVA tests of monotonically transformed data indicated that all factors provided statistically significant effects on all performance indices in paired comparisons of AWaRE with each of the alternatives. Wilcoxon signed rank test was used as post-hoc: all comparisons provided $p < 0.01$ with the exceptions indicated.

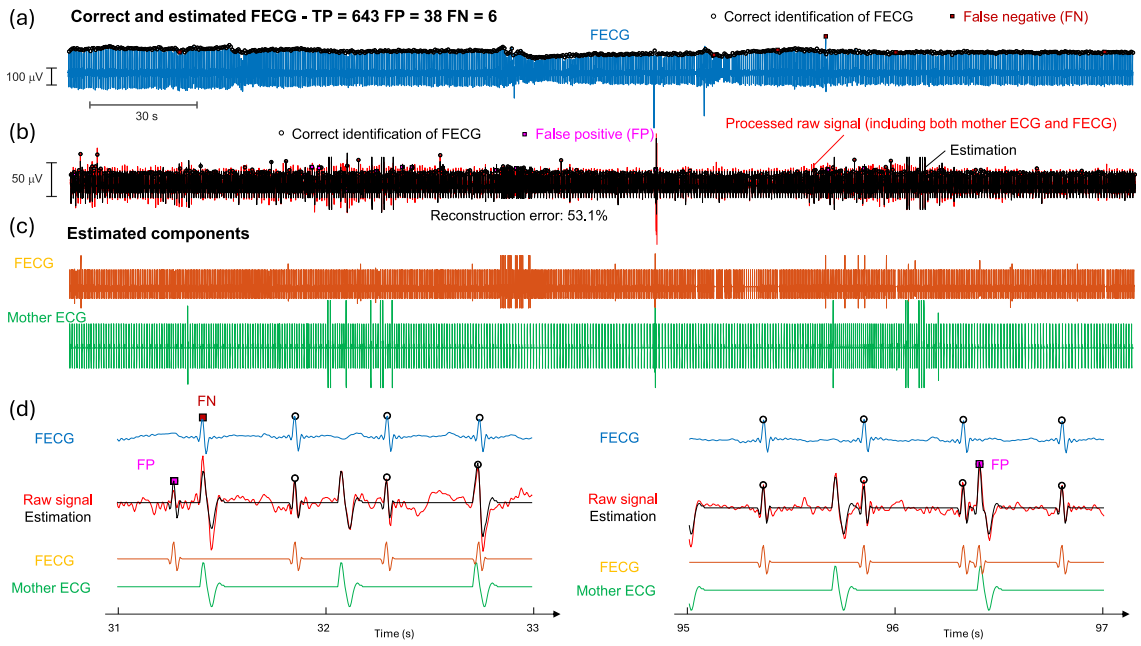


Fig. 7. Example of application of AWaRE to Fetal ECG (FECG). The data (originally sampled at 1 kHz) were bandpass filtered between 2 and 40 Hz and downsampled at 200 Hz. The following parameters were used for AWaRE: $A_{th} = 0.13$, $Norm$ = 98th percentile of $|s(t)|$, $ED = 25$, $M_{th} = 0.32$, $N_{WF} = 3$. Processing time was 5.5 s per 30 s of data. (a) Direct FECG with indication of identified peaks and false negatives obtained by processing the signal in (b), in which the FECG is superimposed to the ECG of the mother (and true and false positives are indicated). (c) Estimated trains of FECG and mother ECG. (d) Zoom of two portions of data in which some false positives or negatives are found.

Fig. 8 shows three EMGs recorded from different patients. The parameters of AWaRE were selected by using Bayesian optimization, minimizing the sum of the reconstruction error and the number of identified waveforms divided by 12. This error functional allowed to detect exactly 3 waveforms for each signal (changing parameters, more waveforms could be identified, e.g., even more than 10 on the signal from the healthy subject, obtaining a very small reconstruction error). This preliminary decomposition could provide some insight into the regularity of the firings of different MUs. By changing the parameters, e.g., reducing the amplitude thresholds A_{th}, smaller waveforms can be detected and their inclusion can reduce the reconstruction error.

Fig. 9 shows two example applications to neural data: (a) an EEG corrupted by ocular artifact and (b) a simulated MER. In the first case, only the blink artifact is identified as recurrent waveform and can be removed to clean the EEG. Note that only a single recurrent waveform was identified and subtracted from the raw signal. However, the ten blinks differed from one another (as indicated in Section 2.5.4), making the subtraction of their mean waveform suboptimal. Consequently, although the improvement compared to the raw signal is evident, the estimation of the slowest EEG rhythms remains substantially affected by residual artifacts, with only the beta band reliably recovered.

In the case of MER, three waveforms were correctly identified, but they are very similar; as a consequence, in the comparison of true and estimated waveform trains, there is some confusion in discriminating the correct class.

Overall, the applications to experimental data highlight the flexibility of AWaRE in handling different types of signals characterized by the presence of recurrent waveforms. While some limitations remain (e.g., residual artifacts or imperfect separation of contributions, whose assessment remains inherently uncertain in any case, in the absence of ground truth), the method consistently provides a decomposition into recurrent waveform components that supports data interpretation. Therefore, AWaRE is shown to be a low-cost pre-processing tool for extracting meaningful structure from data, either for exploratory analysis or as a preliminary step for downstream tasks such as classification or clustering.

4. Discussion

4.1. Motivation

Many methods have been proposed to decompose a signal as the sum of more components: sinusoids or complex exponentials for the Fourier transform, scaled versions of a waveform for the wavelet transform, modes adapted to the signal for EMD or variational mode decomposition (VMD), etc. Decomposition can provide new keys for understanding the original signal, e.g., decoding the activity of different sources. Moreover, it can help to clean the raw data, e.g., when high frequency noise or slow trends/artifacts are separated from the signal of interest in the Fourier transformed domain. Furthermore, specific patterns may emerge that facilitate interpretation, as in time-frequency representations (TFR), offering a 2-dimensional (redundant) view of a time-series, in which different components can show specific signatures that are more easily discriminable.

This paper introduces a new method (called AWaRE) for the analysis of a time-series, which decomposes it into separate components extracted from the signal itself. It leverages various pieces of knowledge.

- Template matching (by cross-correlation) is optimal in the least squared sense for identifying specific waveforms with random noise superimposed (equivalent to matched filtering). However, usually only a specific template is used and it is selected by the user depending on which predefined waveform should be found.
- Continuous wavelet transform (CWT) processes a signal by cross-correlating it with a set of templates. Best compression or localization of the estimates is obtained if such templates are similar to the waveforms that constitute the signal. However, they are all scaled versions of the mother wavelet, that is a predefined function, which could not be representative of the various waves that constitute the signal of interest.
- Recurrences are often found in a biological signal and are a clear indication of the subtended determinism of the system generating it. Specifically, recurrences of subsequent samples are searched in the theory of time delay embedding.

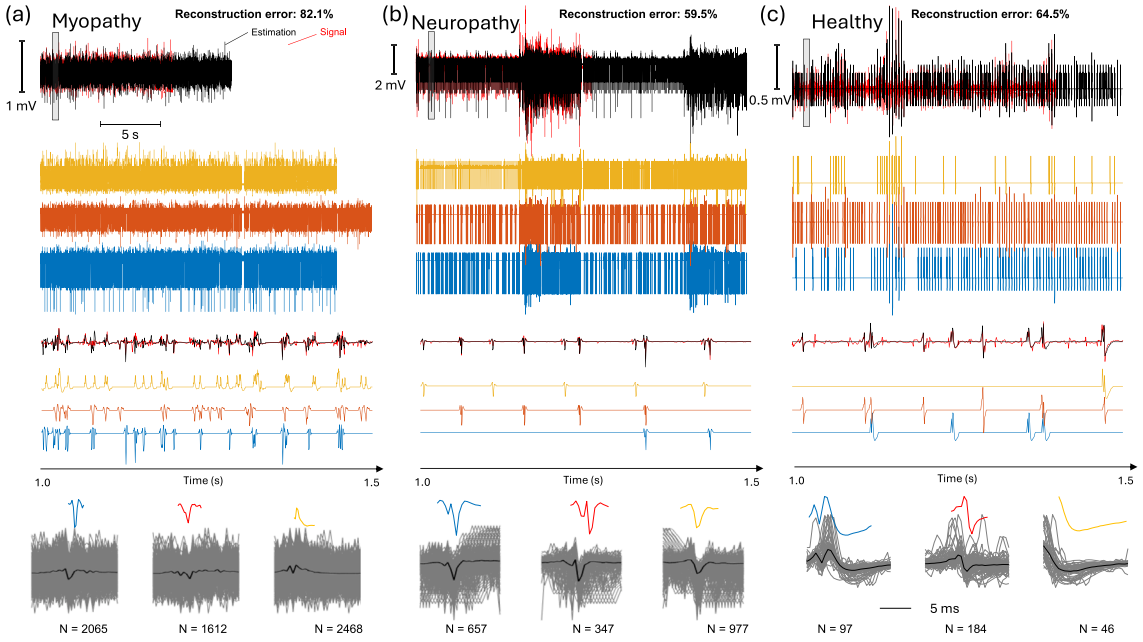


Fig. 8. Examples of intramuscular EMGs sampled at 4 kHz. They were filtered with bandwidth [25–1000] Hz and with a notch at 50 Hz, before downsampling at 2 kHz. For all processed signals $Norm = \max|s(t)|$. (a) Signal from a patient with myopathy: parameters of the method $A_{th} = 0.29$, $ED = 27$, $M_{th} = 0.4$, $N_{WF} = 2$ firings per second; processing time 7.0 s per 27.6 s of signal. (b) Signal from a patient with neuropathy: parameters of the method $A_{th} = 0.23$, , $ED = 25$, $M_{th} = 0.358$, $N_{WF} = 2$ firings per second; processing time 8.4 s per 37 s of signal. (c) Signal from a healthy subject: parameters of the method $A_{th} = 0.081$, $ED = 27$, $M_{th} = 0.251$, $N_{WF} = 2$ firings per second; processing time 2.6 s per 12.7 s of signal. In all cases, the raw signal is shown on top together with the estimation; then, moving downward, there are the three estimated waveform trains, a zoom of the original and estimated decomposition, the estimated waveforms and the epochs centered in the firing instants of each of them and their average (notice that the estimated waveform can be different from the average, as the latter is obtained considering also the epochs with overlaps).

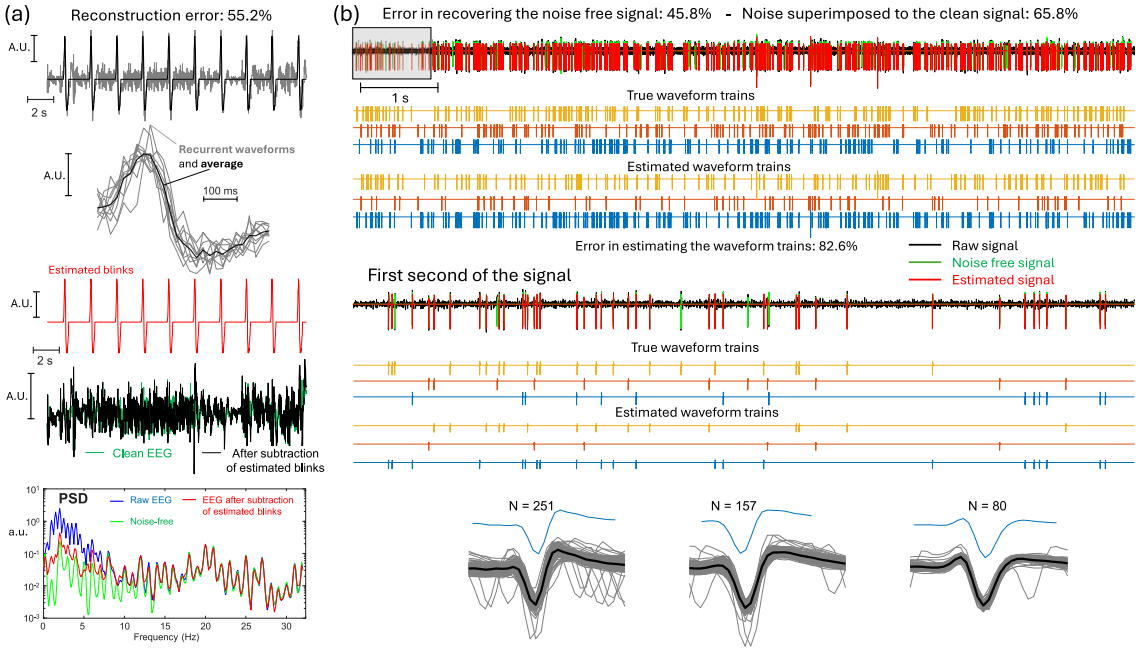


Fig. 9. (a) Application to epochs of EEG corrupted by ocular artifact, sampled at 256 Hz and downsampled at 64 Hz after a bandpass filter with bandwidth [0.1–28] Hz. From top to bottom: raw signal with the artifact highlighted in black; estimated recurrent waveforms after alignment and averaging; estimated artifact; artifact-free EEG and signal obtained after subtraction of the estimated artifact from the raw data; Power Spectral Density (PSD; Welch method, with Hamming window, 8 sub-epochs, 50% of overlap) of the raw signal, artifact-free signal, and residual signal after subtraction of the estimated blinks. Parameters of the method: $A_{th} = 0.025$, $Norm = \overline{|s(t)|}$, $ED = 25$, $M_{th} = 0.15$, $N_{WF} = 4$ firings; processing time 295 ms per 20 s of signal. (b) Application to a simulated MicroElectrode Recording (MER) with sampling frequency 24 kHz. Data were bandpass filtered in the range [300-6000] Hz. The parameters of the method were tuned to get a good reconstruction: $A_{th} = 5$ times the noise estimate (i.e., the median of $|s(t)|$ divided by 0.6745), $Norm = \max|s(t)|$, $ED = 27$, $M_{th} = 0.078$, $N_{WF} = 2$ firings per second; processing time 30 s per 10 s of signal.

These considerations motivated the development of a method that searches for recurrent waveforms, based on the assumption that such recurrences reflect an underlying determinism in the signal and may provide insight into its interpretation. Recurrence detection is efficiently implemented via template matching operations, using segments of the signal as templates. Consequently, the decomposition employs signal-derived basis functions, ensuring that the representation is inherently tailored to the structure of the specific time-series. The waveforms already identified as recurrences are not further used as templates for matching (as shown in Fig. 3), thus speeding up the detection.

Therefore, more matching operations are used, considering different templates, as in the case of CWT; however, those templates are not based on a predefined prototype function, but they are waves automatically extracted from the signal itself; moreover, they are not scaled versions of the same shape, but they can be all different. It is not even required that those waveforms satisfy specific constraints (e.g., that they have separate bandwidths, as in single-channel ICA or in the discrete wavelet transform), but they can be any portion of the processed signal.

4.2. Comments of the results

Recurring waveforms can be considered as successive activations of specific sources. Fig. 4 shows how AWaRE can be applied to recover simulated waveform trains superimposed in the single processed signal. The problem of separating different sources from a single-channel recording has been extensively studied in the literature, adapting to different applications: for example, estimating heartbeats from ECG [4,39], sorting spiking neurons [6,30], decomposing intramuscular EMG [40], identifying blink artifacts in EEG [41]. Two very different approaches have been compared to AWaRE: EEMD-ICA, which splits the single signal into IMFs and then uses ICA to extract the activity of different sources; a spiking neuron pipeline, using thresholding to identify spikes that are then clustered to be separated into different trains. Figs. 5 and 6 show a comparison of the three methods on simulated signals including spiking waveforms. The alternative methods have problems in managing the specific simulations used for the tests. This can be explained by the fact that EEMD-ICA is not specifically designed for the identification of spiking sources and the considered spike sorting approach did not solve overlaps. AWaRE overcame those alternatives in terms of performances and has also a low computational cost (with a median processing time lower than the others, with increasing cost when many waveforms are found, as the multi-kernel deconvolution for overlap resolution has a quadratically increasing cost with the number of kernels).

However, the primary goal of this work is not to surpass the performance of existing methods in any particular application, but to propose a novel framework for time-series analysis, based on decomposing the signal into a sum of repeated patterns inherently embedded in the data. The identified recurrent waveforms can be considered as the basic bricks that build the signal. Knowing these constituent bricks could open new keys for interpreting the data. The mere recurrence of waveforms indicates a deterministic structure, as mentioned before. Each different waveform could be associated to a source. The possible regularity or complexity of the onsets of a specific waveform could provide important information (as for the heartbeats and their variability, or the firing rates of neurons). The separate waveform trains could be further processed and studied, whereas they could be difficult to be interpreted when mixed with other sources.

As an example application, a FECG was extracted from a mixture of mother and fetal ECG, showing good identification of the fetal heartbeats (Fig. 7). Note that, for this specific application, including some a-priori information on the fetal heartbeat could be beneficial to improve the accuracy (e.g., the expected maximal firing rate could force to avoid the identification of beats too close to each other; on the other hand, the identification of too distant waves could suggest finding another in between).

MU action potential (MUAP) trains identification in intramuscular EMG is another example in which the recurrence of the onset of the same waveform can guide the decomposition of a single-channel. As shown in Fig. 8, EMGs from three different patients were investigated and the parameters were automatically selected to get the same number of trains. The total number and regularity of the firings and the MUAP shapes could provide useful information to compare the patients.

Even the residual signal could be more easily interpreted after removing the estimated waveform trains: this is the case of the EEG signal of Fig. 9(a) after removing the recurrent ocular artifact. As a further example involving neural data, Fig. 9(b) illustrates the use of AWaRE on a simulated MER, addressing a challenging spike sorting task where the waveforms were highly similar. The method identified most of the spikes, but made many misclassifications: probably, the mean squared difference was not enough stable to discriminate those similar spikes perturbed by noise.

In summary, representative examples are provided, drawn from biomedical applications involving regularly firing sources. For each application, physiological knowledge can guide to impose constraints to improve the identification and discrimination of different waveforms. However, AWaRE was here applied without any additional information, being blind about the application. In fact, it is proposed as a general decomposition method, aimed to identify recurrent waveforms. Thus, the examples provided simply support the significance of recurrent waveforms in a physiological time-series. The domain-agnostic formulation makes AWaRE applicable to structurally different contexts, providing a useful pre-processing. However, as AWaRE is not conceived as a task-specific approach for either decomposition or denoising, other approaches excelling in those specific tasks are expected to provide better focused results. Therefore, AWaRE should be interpreted as complementary to advanced techniques, providing a preliminary exploration of the signal, before applying more focused methods, if needed. For example, recent advances in EEG denoising have led to specialized adaptive decomposition approaches. Methods such as the Fourier–Bessel series expansion empirical wavelet transform combined with local polynomial approximation total variation filtering [42] and the dyadic boundary point empirical wavelet transform with second-order gradient-based total variation filtering [43] decompose a single time-series into data-driven sub-bands and selectively suppress artifact-dominated components through nonlinear variation-based regularization.

A strength of AWaRE is its moderate computational cost, which is comparable to that of the alternative lightweight methods included in our comparison. Another approach that addresses the same problem of separating a single time series into the sum of waveforms trains is CDL [44], which adopts a fully blind strategy in which both the dictionary (i.e., the set of waveform templates in our application) and the corresponding activation maps must be inferred directly from the signal. This requires solving large-scale sparse-coding and dictionary-update sub-problems, resulting in a considerably higher computational burden; for this reason, CDL was not included in our benchmark. In contrast, AWaRE explicitly identifies recurring waveform patterns by efficiently comparing each template candidate with all portions of the signal through convolution operations. Templates are therefore extracted directly from the data and retained on the basis of their recurrences. Importantly, amplitude plays a discriminative role in AWaRE: templates with identical shapes but different amplitudes are treated as distinct patterns, whereas CDL enforces unit-norm filters and thus collapses such cases into a single atom. This distinction is crucial in applications such as neural spike detection in MERs or MUAP identification in EMGs, where waveforms of similar shape but different amplitude correspond to different physiological events.

4.3. Limitations

AWaRE is specifically designed to identify localized waveforms. It assumes that at least some waveform occurrences appear in isolation, enabling reliable estimation through recurrence. In the presence of moderate overlap, recurrent patterns can still be identified, whereas in the extreme case of highly variable and consistent overlaps that prevent recurrences from being reliably detected, the method is expected to degrade, as templates cannot be extracted. This assumption of AWaRE is generally satisfied in many bioelectrical signals (e.g., MER, EMG, ECG), while signals lacking recurrent patterns may require alternative approaches. For example, AWaRE is not suitable for completely random sources, where statistical properties such as uncorrelation or independence can be more effectively exploited for source separation (e.g., via PCA- or ICA-based approaches, respectively). Moreover, it is not directly applicable to sources that are temporally distributed, such as those effectively separated by wavelet ICA (WICA) and EEMD-ICA in [27]. To address this limitation, one possible strategy involves initially applying the method to selectively activated sources, allowing the extraction of their constituent basis waveforms. These learned components could then be used to separate mixtures of such sources. Nonetheless, because the method is based on the morphological structure of specific signal segments (rather than on statistical properties like independence) its performance in this context strongly depends on the distinctiveness of the underlying basis waveforms. This is particularly critical given that phase cancellations in mixtures may significantly distort the signal morphology.

Another limitation is the presence of many parameters, e.g., the minimum amplitude of waveforms and the distance threshold under which a recurrence is found. By changing these parameters, different signal decompositions can be obtained. This renders the method flexible, but also sensitive to these parameters (refer to Appendix B for a sensitivity analysis on a simulated dataset). The automatic selection of parameters to achieve a specific aim (e.g., the reduction of the reconstruction error with a fixed number of waveforms) could be used in focused applications.

5. Conclusion

Recurrence of specific waveforms within a signal clearly reflects the deterministic nature of the underlying system. It also implies the presence of multiple contributing sources, providing valuable information for distinguishing their respective activities. This study presents a novel method for decomposing a single time-series into a sum of repeated waveforms, that are automatically extracted from the signal itself. These waveforms are initially identified as segments of the time-series and then refined by averaging. Their temporal localization is further enhanced by resolving the overlaps caused by the superposition of other signal components. After thorough validation on simulated data, the method is demonstrated across diverse applications, including identification of fetal heartbeats from ECG recordings contaminated by maternal signals, intramuscular EMG decomposition, EEG artifact removal, and spike sorting in neural recordings.

Funding

This research received no external funding.

Declaration of competing interest

The author declares no conflict of interest.

Appendix A. Single channel decomposition methods in the literature

Different methods have been proposed to decompose a single signal into the sum of different contributions. Obviously, the problem is under-determined, so infinite choices are possible to select a plurality of signals with a sum equal to a given one. Thus, different methods make different assumptions for selecting the solution of interest. A list of the methods related to source separation is provided below. Some of them are suitable for specific applications, whereas others are more general. The integration of more approaches can also be useful for building pipelines aimed at solving this problem. Because of the very different ideas, the methods are clustered here, by grouping those that come from similar techniques.

Blind Source Separation (BSS) approaches

The following methods are useful when more time-series are available and can be used in combination with others (listed below) to split the single-channel signal into more components.

- Principal Component Analysis (PCA)
  PCA reduces the dimensionality of multiple data, by imposing non-correlated sources. In single-channel decomposition, it can be applied to project spikes identified by thresholding (see below) onto principal components to facilitate clustering [14].
- Independent Component Analysis (ICA)
  ICA identifies statistically independent sources, making it useful for decomposing non-Gaussian mixed signals [45]. This technique can outperform PCA when dealing with complex signals, but is computationally more demanding. In order to reduce the computational cost, it has been implemented in several efficient algorithms, e.g., the FastICA [28] (which relies on a fixed-point iteration scheme for rapid maximization of an approximation of the neg-entropy), the Infomax algorithm [46] (based on information maximization principles), and the RobustICA approach [47] (which optimizes a kurtosis contrast with an analytically determined step size).
- Robust ICA (RICA) is a variant of ICA that is more resilient to outliers [48]. RICA is valuable for single-channel biomedical data with noise spikes, separating the underlying sources while being less sensitive to the non-Gaussian assumption [49]. It should not be confused with RobustICA [47] (mentioned above), which is instead an efficient method for multi-channel ICA implementation.

Splitting the signal into more components

The following methods are useful to split a single-channel into more components.

- Time frequency representations (TFR)
  They decompose a single time-series in a two-dimensional space (time-frequency) using different approaches such as the short-time Fourier transform (STFT, affected by a time-frequency resolution tradeoff) or quadratic approaches (with better resolution than STFT, but affected by cross-terms). This redundancy can be exploited for source separation [50].
- Wavelet Transform
  Wavelet transform decomposes signals into a time-scale representation, allowing multi-scale feature extraction. This approach is robust to noise and is adaptable to varying spike widths. A crucial step is to select an appropriate mother wavelet [17]. It is a way to split a single-channel into more time-series, which can then be processed by methods imposing uncorrelation (PCA) or independence (ICA) to separate sources (e.g., in the wavelet ICA method, WICA [51]).

- Empirical-Mode Decomposition (EMD)
  It splits the signal into Intrinsic Mode Functions (IMF), i.e., simple oscillatory modes are iteratively extracted from it. Ensemble EMD (EEMD) is a more stable version that repeats EMD on more versions of the signal corrupted by noise and then averages the results. When combined with ICA (EMD-ICA or EEMD-ICA), it can decompose a single-channel into different sources, as shown for EMG and EEG [27].
- Variational Mode Decomposition (VMD)
  It separates the signal into a finite set of modes, each with a compact frequency support around a center frequency. It has been used to denoise and remove artifacts from ECG [52], EMG [53,54] and EEG [55].
- Phase space reconstruction by time delay embedding.
  The single-channel is embedded in a multidimensional space including delayed versions, before applying a BSS approach for source separation, e.g., ICA [56].
- Single-channel ICA (SCICA)
  A time delay representation is used. SCICA can separate sources, but only if they have different frequency bands [57].
- Convolutional Dictionary Learning (CDL)
  Recent advances in CDL [44] provide a powerful framework for representing a signal as a sum of convolutions between learned filters and sparse activation maps. CDL jointly optimizes both filters and coefficient maps through the minimization of a global non-convex functional, imposing good reconstruction and constraints of sparseness and unit-norm on the filters.

Identification and classification

Classification and machine learning methods can be integrated with identification algorithms to separate the contributions of interest.

- Thresholding
  Thresholding is a fundamental technique for detecting spikes by identifying signal amplitudes that exceed a predetermined value (i.e., the threshold). It is computationally efficient and straightforward to implement, but may produce false positives in noisy data and miss low-amplitude spikes [14].
- Clustering techniques
  Clustering algorithms such as k-means, Gaussian Mixture Models (GMM), and other optimization methods (e.g., super-paramagnetic clustering) [58,59] are flexible and can handle complex spike shapes, but they can be sensitive to initial conditions and parameter selection.
- Supervised machine learning (ML) methods
  Supervised learning algorithms, including support vector machines and neural networks, classify spikes based on labeled training data [60]. Although they achieve high accuracy, they require a labeled dataset, which is often difficult to obtain.
- Deep learning (DL) approaches
  DL methods have been proposed to decompose a signal into different sources, e.g., to remove artifacts in EEG [61] or to decompose intramuscular EMG [62]. These techniques typically require a large amount of labeled data for training. An alternative is to use methods that simply learn the signal (e.g., autoencoders) and extract features that can be useful for identifying different sources [63].

Matching with prototypes

If a prototype waveform is known, the following methods can be used to search it in the signal.

- Template Matching
  When specific prototypes are known, template matching searches for the waveforms closest to those templates, often using cross-correlation to measure similarity [18]. In spike sorting applications, templates are automatically identified (e.g., by thresholding for identification and clustering for separation of different spiky waveforms), before matching them with the signal [16].
- Dynamic Time Warping (DTW) for Template Matching
  DTW aligns spikes with templates by allowing flexible time shifts and effectively handles shape variations [64]. Although robust, DTW incurs a high computational cost.
- Hidden Markov Models (HMM)
  HMMs represent spike sequences as transitions between states that are effective for waveforms with consistent firing patterns. A method that uses information on both the spike shape and inter-spike interval was introduced for sorting [65]. HMMs require careful initialization and training, which increases their computational cost.

Integration of methods

Depending on the application, different approaches can be integrated.

- In the case of spike sorting, sharp waveforms with amplitude emerging from the background noise are sought. A typical pipeline [29] identifies the spikes (e.g., by thresholding) and clusters similar shapes (e.g., after extracting features by PCA and using a clustering algorithm on them). This approach can be followed by a method to solve the overlap of different spikes.
- When sources are not localized spikes, but are distributed on the entire signal (e.g., two independent wide-sense stationary processes superimposed on the same single-channel), the pipeline for separation could involve using a decomposition approach that maps the single-channel into more time-series (e.g., by wavelet transform, EMD, VMD), and applying a multidimensional BSS method (e.g., PCA or ICA) to identify different sources.

Final overview

In summary, each decomposition method has its unique advantages and drawbacks. Thresholding is computationally simple, but lacks robustness in noisy environments. The wavelet transform and clustering techniques adapt well to complex spike shapes. HMM and supervised learning methods excel in labeled datasets, although their computational demands are high. The method selection depends on factors such as noise levels, spike shape variability, and available computational resources. PCA and ICA enable dimensionality reduction, with ICA being particularly effective for non-Gaussian sources and can be used when more channels are available, if more epochs can be extracted from the same signal (as in the case of evoked potentials or when there are spikes that can be identified by thresholding and windows can be taken including them), or if the single-channel can be decomposed into different time-series (as in EMD, wavelet transform, VMD, and phase space representation). Single-channel variants of ICA, such as SCICA and RICA, are feasible under specific hypotheses (e.g., different bandwidths of the contribution to be separated). The advantages and disadvantages of these methods are presented in Table A.2.

Appendix B. Theoretical considerations

The theoretical basis and the domain of validity of the proposed method are here discussed. Although the algorithm consists of multiple stages (template detection, clustering of recurrent waveforms, template refinement via averaging, re-detection, overlap resolution, and

Table A.2
Comparison of different methods for signal decomposition.

Method	Advantages	Main Disadvantages
PCA	Reduces dimensionality	Assumes linear separability; suitable with more channels recording the same sources
ICA	Can separate non-Gaussian sources	Assumes linear separability; computationally demanding; suitable with more channels recording the same sources
RICA	More robust to outliers than ICA; suitable also for under-determined problems	May require optimization of parameters
TFR	Represents the scalar signal into a two-dimensional space	Time–frequency resolution tradeoff (for STFT) or cross-terms (quadratic approaches)
Wavelet transform	Multi-scale analysis	Sensitive to the choice of the mother wavelet
EMD	Adaptive and feasible for nonlinear and nonstationary signals	Sensitive to noise (EEMD could provide a countermeasure)
VMD	Adaptive, robust to noise, control over number of modes	Computationally demanding and requires parameter tuning
Phase space reconstruction	No prior knowledge of the number of sources, effective for nonlinear and nonstationary data	Sensitive to parameter selection (optimal delay and embedding dimension), high computational cost
SCICA	Effective for source separation if signals have different frequency bands; captures temporal dependencies and delays	Can only separate sources with different frequency bands; high computational complexity; may require selection of optimal number of components
CDL	Very general model in which the estimation of both templates (constituting the dictionary) and their onsets are optimized	High computational cost; physiological interpretation of the templates is not guaranteed
Thresholding	Fast, simple	Sensitive to noise; may miss low spikes
Clustering	Handles complex shapes	Sensitive to initial conditions and parameters selection
Supervised ML and DL	High accuracy with labels	Needs labeled data; high computational cost
Template matching	Accurate with stable shapes	The template should be known
DTW	Robust to shape variation	Computationally intense
HMM	Can impose regularity of the firing patterns	Requires good initialization

refinement on residual) we can still outline the underlying assumptions, the conditions for validity, and the fundamental properties inherited from the classical methods upon which it is built (e.g., template matching and signal averaging). These considerations provide a theoretical framework that complements the simulation-based validation presented in the Results Section.

Assumptions of the model

We assume that the observed signal can be written as

$$x(t) = \sum_{k=1}^{K} \sum_{i=1}^{N_k} a_k(t - \tau_{k,i}) + n(t), \tag{B.1}$$

where $a_k(t)$ denotes the deterministic waveform generated by source k, $\{\tau_{k,i}\}$ are the occurrence times of that waveform (which has N_k onsets), and $n(t)$ is additive noise. This model describes a broad class of physiological signals composed of localized and recurrent waveforms, such as MUAPs in intramuscular EMG, action potentials in neural recordings, or cardiac beats in ECG.

The method relies on the following structural assumptions.

1. *Temporal localization:* each waveform $a_k(t)$ is time-limited or effectively supported within a short window. Some occurrences should not overlap significantly with waveforms from other sources, so that the templates could be reliably identified.
2. *Recurrence:* each source produces multiple occurrences of similar shape. A minimum number of occurrences must be present for each waveform:

   $$N_k \geq N_{\min}, \tag{B.2}$$

   where $N_{\min}$ depends on the noise power and needed reduction (by averaging). Waveforms occurring too rarely cannot be distinguished from noise. Moreover, waveforms from the same source are expected to retain a consistent morphology; notice that this assumption can be undermined, e.g., in applications on EMG, by gradual shape variations arising from contact instability or fatigue-induced changes of MUAPs.
3. *Sparsity of overlap:* although overlaps between sources may occur, a non-negligible fraction of occurrences must appear in isolation (as indicated above). Then, once estimated the different template waveforms, they can be used as multiple kernels for sparse deconvolution estimation.
4. *Distinctiveness:* waveforms from different sources must differ sufficiently to be distinguished by template matching. This means that each pair of templates, e.g., a_k and a_ℓ, must satisfy a cross-correlation bound of the form

   $$\langle a_k, a_\ell \rangle < \rho_{\text{crit}}, \tag{B.3}$$

   for some threshold ρ_{crit}. If two sources generate nearly identical waveforms, a reliable separation of them is not possible.

These assumptions define the domain of applicability of the method, i.e., signals containing sparse and recurrent localized events (e.g., spiking neurons, MUAPs in intramuscular EMG, heartbeats in ECG). Signals lacking temporal localization or recurrence (e.g., broadband stochastic processes or continuous oscillatory activity without isolated events, or superimposed waveforms which are non-recurrent or fully overlapped) cannot be decomposed by the proposed procedure.

Foundational properties of the algorithmic components

The waveforms identification is based on classical template matching, which corresponds to the least-squares optimal estimator for detecting a known waveform in additive white Gaussian noise.

Once multiple occurrences of a candidate waveform are identified, they are aligned and averaged to refine the template. If occurrences have the form $a_k(t - \tau_{k,i}) + n_i(t)$, where $n_i(t)$ are independent noise

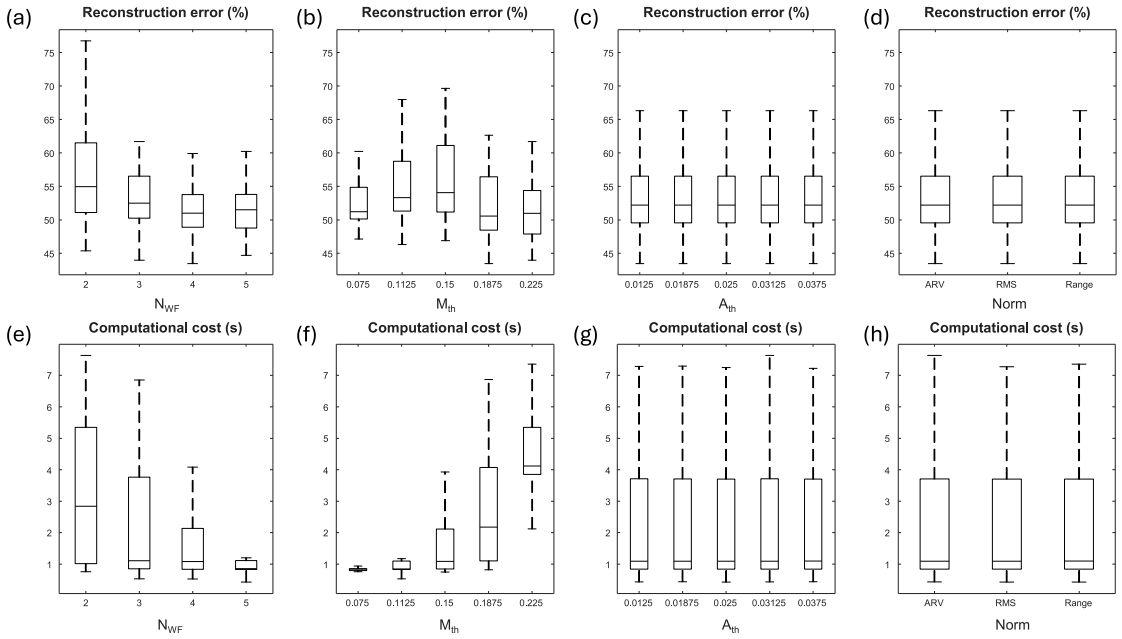


Fig. B.10. Sensitivity analysis of the main parameters of the method, applied to the same simulated data considered in Fig. 6, fixing the SNR to 10 dB, the number of different waveforms to 8 and their recurrences to 30. Ten trials of 10 s duration were tested for each condition. The values of the parameters used for Fig. 6 (i.e., $A_{th} = 0.025$, $Norm = \text{ARV} = \overline{|s(t)|}$, $ED = [15\ 31]$, $M_{th} = 0.15$, $N_{WF} = 3$) were used as default and 4 of them were changed one at a time. (a) Error in reconstructing the clean signal as a function of N_{WF}. (b) Error in reconstructing the clean signal as a function of M_{th}. (c) Error in reconstructing the clean signal as a function of A_{th}. (d) Error in reconstructing the clean signal as a function of $Norm$. (e) Computation time as a function of N_{WF}. (f) Computation time as a function of M_{th}. (g) Computation time as a function of A_{th}. (d) Computation time as a function of $Norm$.

realizations, then the average over N aligned occurrences converges to the true waveform reducing noise power at a rate proportional to $1/N$.

Once obtained robust and clean estimations of the recurrent waveforms, deconvolution with multiple kernels is used to remove overlaps.

Limitations of a closed-form performance analysis

A complete analytical characterization of the algorithm's performance as a function of SNR, waveform similarity, and overlap rate is not feasible due to the multi-stage and nonlinear nature of the method. Each stage depends on the accuracy of the previous one; for example, errors in initial template detection may be corrected during refinement, and the degree of refinement depends on the (unknown) recurrence structure of the data. Obtaining closed-form bounds would require adopting a highly restrictive generative model that does not reflect the variability of real physiological signals. For this reason, the method was tested on controlled simulations that systematically investigated the impact on practical performance of the number of different waveforms, their rate of recurrence (which, when higher, partly improves the estimation but also raises the overlap fraction) and noise level.

Sensitivity analysis

AWaRE requires to choose many parameters. As mentioned in Section 4.3 on Limitations, this provides flexibility to the method, but also some sensitivity to the choice of the parameters. Fig. B.10 shows two performance indexes, namely the reconstruction error of the clean signal and computation time, when deviating a single parameter at a time from a default model. Specifically, the dataset is equivalent to the simulated data considered in Fig. 6, fixing specific values of noise level and waveforms recurrences. Then parameters were changed within reasonable ranges for the specific application. Notice that all parameters of the method (summarized in Table 1) are considered, with the exception of ED, as it is specific for the range of durations of the waveforms that the user expects to find in the signal; moreover, a set of possible values can be provided (in this case, 15 and 31, to consider both short and long waveforms). Results indicate that the reconstruction error is not largely affected: using the median of paired absolute differences as effect size, the mean across all possible pairs with pooled data were 2.6%, 4.1%, 0.0% and 0.0% for N_{WF}, M_{th}, A_{th} and $Norm$, respectively. Considering the computational cost, the mean effect sizes are 0.6 s for N_{WF}, 1.6 s for M_{th}, and about 0.01 s for A_{th} and $Norm$. Notice that the computation time is increased when N_{WF} is reduced, since more distinct waveforms are included in the initial stages of the algorithm, some of which could exhibit only a few recurrences. Moreover, there is a tendency for the computational cost to grow with larger values of M_{th}, as a larger tolerance increases the number of waveforms considered as recurrent, thereby raising the overhead associated with their selection, discrimination, and cleaning in the last steps of the algorithm. In the specific tests considered, negligible effects were observed for the choice of the amplitude threshold A_{th} (which is instead a critical parameter for spike sorting methods) and the norm to measure it (i.e., $Norm$): in fact, even if a waveform is above threshold, it is selected only if it is repeated more times in the signal.

Data availability

Data will be made available on request.

References

[1] C.L. Webber, J.P. Zbilut, Recurrence quantification analysis of nonlinear dynamical systems, Phys. Lett. A 171 (3) (1994) 199–203, http://dx.doi.org/10.1016/0375-9601(94)90282-8.

[2] N. Marwan, M. Carmen Romano, M. Thiel, J. Kurths, Recurrence plots for the analysis of complex systems, Phys. Rep. 438 (5–6) (2007) 237–329, http://dx.doi.org/10.1016/j.physrep.2006.11.001.

[3] F. Takens, Detecting strange attractors in turbulence, in: Dynamical Systems and Turbulence, in: Lecture Notes in Mathematics, vol. 898, Springer, 1981, pp. 366–381, http://dx.doi.org/10.1007/BFb0091924.

[4] J. Pan, W.J. Tompkins, A real-time QRS detection algorithm, IEEE Trans. Biomed. Eng. 32 (3) (1985) 230–236, http://dx.doi.org/10.1109/TBME.1985.325532.
[5] L. Mesin, Heartbeat monitoring from adaptively down-sampled electrocardiogram, Comput. Biol. Med. 84 (2017) 217–225, http://dx.doi.org/10.1016/j.compbiomed.2017.03.023.
[6] R. Quian Quiroga, Z. Nadasdy, Y. Ben-Shaul, Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering, Neural Comput. 16 (8) (2004) 1661–1687, http://dx.doi.org/10.1162/089976604774201631.
[7] C.J. De Luca, A. Adam, R. Wotiz, L.D. Gilmore, S.H. Nawab, Electromyography and motor control: Detecting motor unit action potentials, Clin. Neurophysiol. 117 (3) (2006) 535–551, http://dx.doi.org/10.1016/j.clinph.2005.10.009.
[8] A. Hyvärinen, J. Karhunen, E. Oja, Independent Component Analysis, Wiley, 2001, http://dx.doi.org/10.1002/0471221311.
[9] A. Holobar, D. Zazula, Multichannel blind source separation using convolution kernel compensation, IEEE Trans. Signal Process. 55 (9) (2007) 4487–4496, http://dx.doi.org/10.1109/TSP.2007.894265.
[10] R. Quian Quiroga, Spike sorting, Scholarpedia 2 (12) (2007) 3583, http://dx.doi.org/10.4249/scholarpedia.3583.
[11] T. Yu, K. Akhmadeev, E. Le Carpentier, Y. Aoustin, R. Gross, Y. Péréon, D. Farina, Recursive decomposition of electromyographic signals with a varying number of active sources: Bayesian modeling and filtering, IEEE Trans. Biomed. Eng. 67 (2) (2020) 428–440, http://dx.doi.org/10.1109/TBME.2019.2914966.
[12] J. Magland, J.J. Jun, E. Lovero, A.J. Morley, C. Lincoln Hurwitz, A.P. Buccino, S. Garcia, A.H. Barnett, SpikeForest, reproducible web-facing ground-truth validation of automated neural spike sorters, ELife 9 (2020) e55167, http://dx.doi.org/10.7554/eLife.55167.
[13] M.R. Cohen, A. Kohn, Measuring and interpreting neuronal correlations, Nature Neurosci. 17 (7) (2014) 858–864, http://dx.doi.org/10.1038/nn.2842.
[14] M.S. Lewicki, A review of methods for spike sorting: the detection and classification of neural action potentials, Netw., Comput. Neural Syst. 9 (4) (1998) 53–78, http://dx.doi.org/10.1088/0954-898X/9/4/001.
[15] H.G. Rey, C. Pedreira, R. Quian Quiroga, Past, present and future of spike sorting techniques, Brain Res. Bull. 119 (Pt B) (2015) 106–117, http://dx.doi.org/10.1016/j.brainresbull.2015.04.007.
[16] K.J. Laboy-Juárez, S. Ahn, D.E. Feldman, A normalized template matching method for improving spike detection in extracellular voltage recordings, Sci. Rep. 9 (12087) (2019) http://dx.doi.org/10.1038/s41598-019-48456-y.
[17] R.Q. Quiroga, Z. Nadasdy, Y. Ben-Shaul, Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering, Neural Comput. 16 (8) (2004) 1661–1687, http://dx.doi.org/10.1162/089976604774201631.
[18] L. Mesin, M. Valerio, A. Beaumanoir, G. Capizzi, Automatic identification of slow biphasic complexes in EEG: An effective tool to detect encephalitis, Biomed. Phys. Eng. Express 5 (2019) 045006, http://dx.doi.org/10.1088/2057-1976/ab2086.
[19] G. Chiarion, L. Mesin, Resolution of spike overlapping by biogeography-based optimization, Electronics 10 (12) (2021) 1469, http://dx.doi.org/10.3390/electronics10121469, URL https://www.mdpi.com/2079-9292/10/12/1469.
[20] L. Mesin, Motor unit discharges from multi-kernel deconvolution of single channel surface electromyogram, Electronics 10 (16) (2021) 2022, http://dx.doi.org/10.3390/electronics10162022, URL https://www.mdpi.com/2079-9292/10/16/2022.
[21] L. Mesin, Single channel surface electromyogram deconvolution to explore motor unit discharges, Med. Biol. Eng. Comput. 57 (9) (2019) 2045–2054, http://dx.doi.org/10.1007/s11517-019-02010-0.
[22] D. Craven, B. McGinley, L. Kilmartin, M. Glavin, E. Jones, Compressed sensing for bioelectric signals: A review, IEEE J. Biomed. Health Inf. 19 (2015) 529–540, http://dx.doi.org/10.1109/JBHI.2014.2327194.
[23] I. Daubechies, M. Defrise, C. De Mol, An iterative thresholding algorithm for linear inverse problems with a sparsity constraint, Comm. Pure Appl. Math. 63 (1) (2010) 1–38, http://dx.doi.org/10.48550/arXiv.math/0307152.
[24] L. Mesin, D. Cocito, A new method for the estimation of motor nerve conduction block, Clin. Neurophysiol. 118 (4) (2007) 730–740, http://dx.doi.org/10.1016/j.clinph.2006.11.015.
[25] C. Chen, L. He, H. Li, J. Huang, Fast iteratively reweighted least squares algorithms for analysis-based sparse reconstruction, Med. Image Anal. 49 (2018) 141–152, http://dx.doi.org/10.1016/j.media.2018.08.002.
[26] C. Chen, Distributed iteratively reweighted least squares and applications, Stat. Interface 6 (2013) 585–593, http://dx.doi.org/10.4310/SII.2013.v6.n4.a15.
[27] B. Mijović, M. De Vos, I. Gligorijević, J. Taelman, S. Van Huffel, Source separation from single-channel recordings by combining empirical-mode decomposition and independent component analysis, IEEE Trans. Biomed. Eng. 57 (9) (2012) 2188–2196, http://dx.doi.org/10.1109/TBME.2010.2051440.
[28] A. Hyvärinen, Fast and robust fixed-point algorithms for independent component analysis, IEEE Trans. Neural Netw. 10 (3) (1999) 626–634, http://dx.doi.org/10.1109/72.761722.
[29] V. Forestiere, M. Lazzeri, M. Lenge, F. Giordano, L. Mesin, Analysis of intra-operative MER signals in pediatric patients with dystonia undergoing DBS of the globus pallidus internus, in: Proceedings of GNB 2025, Palermo, Italy, 2025, June 16th–18th.
[30] H.G. Rey, C. Pedreira, R. Quian Quiroga, Past, present and future of spike sorting techniques, Brain Res. Bull. 119(Pt B) (2015) 106–117, http://dx.doi.org/10.1016/j.brainresbull.2015.04.007.
[31] R. Tibshirani, G. Walther, T. Hastie, Estimating the number of clusters in a data set via the gap statistic, J. R. Stat. Soc. Ser. B Stat. Methodol. 63 (2) (2001) 411–423, http://dx.doi.org/10.1111/1467-9868.00293.
[32] J.K. Wong, A new implementation of an algorithm for the optimal assignment problem: An improved version of Munkres' algorithm, BIT Numer. Math. 19 (4) (1979) 418–424, http://dx.doi.org/10.1007/BF01930994.
[33] L. Mesin, Nonlinear spatio-temporal filter to reduce crosstalk in bipolar electromyogram, J. Neural Eng. 21 (2) (2024) 026001, http://dx.doi.org/10.1088/1741-2552/ad1f5a.
[34] A.L. Goldberger, L.A. Amaral, L. Glass, J.M. Hausdorff, P.C. Ivanov, R.G. Mark, J.E. Mietus, G.B. Moody, C.K. Peng, H.E. Stanley, PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals, Circ. [Online] 101 (23) (2000) e215–e220, http://dx.doi.org/10.1161/01.cir.101.23.e215.
[35] J. Jezewski, A. Matonia, T. Kupka, D. Roj, R. Czabanski, Determination of the fetal heart rate from abdominal signals: evaluation of beat-to-beat accuracy in relation to the direct fetal electrocardiogram, Biomed. Eng./Biomed. Tech. 57 (5) (2012) 383–394, http://dx.doi.org/10.1515/bmt-2011-0130.
[36] S. Rutkove, Examples of electromyograms, 2009, URL https://physionet.org/content/emgdb/1.0.0/, (Accessed 30 May 2025).
[37] 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) http://dx.doi.org/10.1088/1741-2552/ac2bf8.
[38] L.A. Camuñas-Mesa, R.Q. Quiroga, A detailed and fast model of extracellular recordings, Neural Comput. 25 (5) (2013) 1191–1212, http://dx.doi.org/10.1162/NECO_a_00433.
[39] F. Sana, T. Ballal, M. Shadaydeh, I. Hoteit, T.Y. Al-Naffouri, Fetal ECG extraction exploiting joint sparse supports in a dual dictionary framework, Biomed. Signal Process. Control. 48 (2019) 46–60, http://dx.doi.org/10.1016/j.bspc.2018.08.023.
[40] H.R. Marateb, S. Muceli, K.C. McGill, R. Merletti, D. Farina, Robust decomposition of single-channel intramuscular EMG signals at low force levels, J. Neural Eng. 8 (6) (2011) 066015, http://dx.doi.org/10.1088/1741-2560/8/6/066015.
[41] M. Wang, X. Cui, T. Wang, T. Jiang, F. Gao, J. Caos, Eye blink artifact detection based on multi-dimensional feature fusion and PSO optimization, Biomed. Signal Process. Control. 83 (2023) 104657, http://dx.doi.org/10.1016/j.bspc.2023.104657.
[42] P. Gajbhiye, R.K. Tripathy, R.B. Pachori, Elimination of ocular artifacts from single-channel EEG signals using FBSE-EWT based rhythms, IEEE Sensors J. 20 (13) (2020) 7317–7327, http://dx.doi.org/10.1109/JSEN.2020.2979276.
[43] S. Dash, P. Gajbhiye, P. Sawangjai, G. Sawadwuthikul, T. Sudhawiyangkul, R.K. Tripathy, R.B. Pachori, Ocular artifact removal from EEG using DBPEWT and second-order gradient total variation filtering, Biomed. Signal Process. Control. 85 (2023) 104996, http://dx.doi.org/10.1016/j.bspc.2023.104996.
[44] C. Garcia-Cardona, B. Wohlberg, Convolutional dictionary learning: A comparative review and new algorithms, IEEE Trans. Comput. Imaging 4 (3) (2018) 366–381, http://dx.doi.org/10.1109/TCI.2018.2840334.
[45] G.D. Brown, S. Yamada, T.J. Sejnowski, Independent component analysis at the neural cocktail party, Trends Neurosci. 24 (1) (2001) 54–63, http://dx.doi.org/10.1016/s0166-2236(00)01683-0.
[46] S. Makeig, A.J. Bell, T.-P. Jung, T.J. Sejnowski, Independent component analysis of electroencephalographic data, in: Advances in Neural Information Processing Systems 8, NIPS 1995, 1995, pp. 145–151, http://dx.doi.org/10.1016/j.neunet.2003.08.003.
[47] V. Zarzoso, P. Comon, Robust independent component analysis by iterative maximization of the kurtosis contrast with algebraic optimal step size, IEEE Trans. Neural Netw. 21 (2) (2010) 248–261, http://dx.doi.org/10.1109/TNN.2009.2035920.
[48] V. Zarzoso, P. Comon, Robust independent component analysis by iterative maximization of the kurtosis contrast with algebraic optimal step size, IEEE Trans. Neural Netw. 21 (2) (2010) 248–261, http://dx.doi.org/10.1109/TNN.2009.2035920.
[49] Y.O. Li, T. Adali, V.D. Calhoun, Estimating the number of independent components for functional magnetic resonance imaging data, Hum. Brain Mapp. 22 (3) (2015) 238–251, http://dx.doi.org/10.1002/hbm.20359.
[50] F. Abrard, Y. Deville, A time-frequency blind signal separation method applicable to underdetermined mixtures of dependent sources, Signal Process. 85 (7) (2005) 1389–1403, http://dx.doi.org/10.1016/j.sigpro.2005.02.010.
[51] J. Lin, A. Zhang, Fault feature separation using wavelet-ICA filter, Nondestruct. Test. Eval. Int. 38 (6) (2005) 421–427, http://dx.doi.org/10.1016/j.ndteint.2004.11.005.
[52] X. Hu, Q. Yu, H. Yu, An ECG denoising method combining variational modal decomposition and wavelet soft threshold, Concurr. Comput. Pr. Exper. e7048 (2022) http://dx.doi.org/10.1002/cpe.7048.
[53] F. Xiao, D. Yang, X. Guo, Y. Wang, VMD-based denoising methods for surface electromyography signals, J. Neural. Eng. 16 (5) (2019) 056017, http://dx.doi.org/10.1088/1741-2552/ab33e4.
[54] H. Ashraf, U. Shafiq, Q. Sajjad, A. Waris, O. Gilani, M. Boutaayamou, O. Brüls, Variational mode decomposition for surface and intramuscular EMG signal denoising, Biomed. Signal Process. Control. 82 (2023) 104560, http://dx.doi.org/10.1016/j.bspc.2022.104560.

[55] C. Dora, P.K. Biswal, Correlation-based ECG artifact correction from single channel EEG using modified variational mode decomposition, Comput. Methods Programs Biomed. 183 (105092) (2020) http://dx.doi.org/10.1016/j.cmpb.2019.105092.

[56] H. Pengju, S. Tingting, L. Wenhui, Y. Weibiao, Single channel blind source separation on the instantaneous mixed signal of multiple dynamic sources, Mech. Syst. Signal Process. 113 (3) (2018) 22–35, http://dx.doi.org/10.1016/j.ymssp.2017.04.004.

[57] M.E. Davies, C.J. James, Source separation using single channel ICA, Signal Process. 87 (8) (2007) 1819–1832, http://dx.doi.org/10.1016/j.sigpro.2007.01.011.

[58] R. Quian Quiroga, Z. Nadasdy, Y. Ben-Shaul, Unsupervised spike detection and sorting with wavelets and super-paramagnetic clustering, Neural Comput. 16 (2004) 1661–1687, http://dx.doi.org/10.1162/089976604774201631.

[59] C. Pedreira, J. Martinez, M.J. Ison, R. Quian Quiroga, How many neurons can we see with current spike sorting algorithms? J. Neurosci. Methods 211 (1) (2012) 58–65, http://dx.doi.org/10.1016/j.jneumeth.2012.07.010.

[60] C. Rossant, S.N. Kadir, D.F.M. Goodman, J. Schulman, M.L.D. Hunter, A.B. Saleem, A. Grosmark, M. Belluscio, G.H. Denfield, A.S. Ecker, A.S. Tolias, S. Solomon, G. Buzsaki, M. Carandini, K.D. Harris, Spike sorting for large, dense electrode arrays, Nature Neurosci. 19 (4) (2016) 634–641, http://dx.doi.org/10.1038/nn.4268.

[61] B. Kalita, N. Deb, D. Das, AnEEG: Leveraging deep learning for effective artifact removal in EEG data, Sci. Rep. 14 (1) (2024) 24234, http://dx.doi.org/10.1038/s41598-024-75091-z.

[62] W. Sun, R. Tang, Y. Lang, J. He, H. Qiang, Decomposing single-channel intramuscular electromyography signal sampled at a low frequency into its motor unit action potential trains with a generative adversarial network, J. Electromyogr. Kinesiol. 48 (2019) 187–196, http://dx.doi.org/10.1016/j.jelekin.2019.07.015.

[63] L. Xing, A.J. Casson, Deep autoencoder for real-time single-channel EEG cleaning and its smartphone implementation using TensorFlow lite with hardware/software acceleration, IEEE Trans. Biomed. Eng. 71 (11) (2024) 3111–3122, http://dx.doi.org/10.1109/TBME.2024.3408331.

[64] Y. Cao, N. Rakhilin, P.H. Gordon, X. Shen, E.C. Kan, A real-time spike classification method based on dynamic time warping for extracellular enteric neural recording with large waveform variability, J. Neurosci. Methods 261 (2016) 97–109, http://dx.doi.org/10.1016/j.jneumeth.2015.12.006.

[65] S. Samiee, M.B. Shamsollahi, V. Vigneron, A new Markovian approach towards neural spike sorting, in: 2011 8th International Conference on Information, Communications & Signal Processing, IEEE, 2011, pp. 1–5, http://dx.doi.org/10.1109/ICICS.2011.6173566.