A Particle Swarm Optimised Independence Estimator for Blind Source Separation of Neurophysiological Time Series

Agnese Grison, Alexander Kenneth Clarke, Silvia Muceli, *Senior Member, IEEE,* Jaime Ibáñez, Aritra Kundu, *Senior Member, IEEE,* Dario Farina, *Fellow, IEEE*

***Abstract*—The decomposition of neurophysiological recordings into their constituent neural sources is of major importance to a diverse range of neuroscientific fields and neuroengineering applications. The advent of high density electrode probes and arrays has driven a major need for novel semi-automated and automated blind source separation methodologies that take advantage of the increased spatial resolution and coverage these new devices offer. Independent component analysis (ICA) offers a principled theoretical framework for such algorithms, but implementation inefficiencies often drive poor performance in practice, particularly for sparse sources. Here we observe that the use of a single non-linear optimization function to identify spiking sources with ICA often has a detrimental effect that precludes the recovery and correct separation of all spiking sources in the signal. We go on to propose a projection-pursuit ICA algorithm designed specifically for spiking sources, which uses a particle swarm methodology to adaptively traverse a polynomial family of non-linearities approximating the asymmetric cumulants of the sources. We robustly prove state-of-the-art decomposition performance on recordings from high density intramuscular probes and demonstrate how the particle swarm quickly finds optimal contrast non-linearities across a range of neurophysiological datasets.**

***Index Terms*—Independent component analysis, particle swarm optimisation, blind source separation, intramuscular electromyography, intracortical recording.**

I. Introduction

In vivo neurophysiological signals recorded by invasive probes, such as intramuscular electromyography (iEMG) and intracortical probe recordings, generally represent the summed electrical activity of many individual spiking neural sources [1]. The automated and semi-automated extraction of these sources from the multi-unit time series without any *a priori* knowledge of their spike timings or the shapes of their action potential (AP) waveforms is the domain of blind source separation (BSS) algorithms [2]. Such neurophysiological signal decomposition, also known as spike sorting, has been a core component of many neuroscientific research workflows for decades, in many cases supplanting tedious manual sorting [3]. In recent years, there has been a consistent trend towards multivariance in neurophysiological time series, driven by dramatic advances in the miniaturisation of recording hardware [4], [5].

Modern high density (HD) invasive probes can have hundreds of sensing electrodes, arranged in high density grids to achieve impressive spatial resolutions [6]. Unprecedented volumes of data have resulted in correspondingly unprecedented yields of decomposed sources [7], but have also introduced new challenges in error correction and reproducibility [8], whilst worsening known problems related to the effects of signal non-stationarity [9]. In many settings, manual spike sorting continues to outperform automated methods in both source yield and accuracy [10], adding an additional workload to studies that is generally unmanageable in the HD setting.

Spike sorting BSS algorithms can be broadly divided by whether they use lower-order or higher-order statistics to separate sources [11]. Lower-order methods find correlations between segmented AP waveforms and follow a pipeline of detection and sorting [3]. The timestamp of a waveform is first found through thresholding and temporal alignment, then the timestamp is assigned to a source by clustering [7], usually after a step of linear feature extraction, such as principal components analysis or wavelet convolution [9], [12]. This order of operations is reversed in higher-order methods; metrics of statistical independence are first used to find the separated source time series [13], which are then thresholded to obtain the spike timestamps [14], [15]. Higher-order statistical methods, such as independent component analysis (ICA), both manage AP waveform superpositions better and are generally more resistant to non-stationarities in the data [16], [14], as they do not require AP alignment and consider the entire time series simultaneously. Higher-order statistical methods have been largely ignored in the spike sorting literature. The exception is recording modalities such as surface EMG, which require high-order statistical methods for successful BSS. This is most likely due to their comparatively heavy computational loads. In recent years, this has been mitigated by the widespread adoption of parallel

AG, AKC, SM, JI, AK and DF are with the Department of Bioengineering, Imperial College London, London, UK. SM is also with the Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden. JI is also with BSICoS, IIS Aragón, Universidad de Zaragoza, Zaragoza, Spain. Please direct all correspondence to DF (d.farina@imperial.ac.uk). Code and experimental data available upon reasonable request.

AG was supported in part by UK Research and Innovation [UKRI Centre for Doctoral Training in AI for Healthcare grant number EP/S023283/1] and in part by Huawei Technologies Research & Development (UK) Limited. AKC was supported by the EPSRC Centre for Neurotechnology (EP/L016737/1). SM was supported by HybridNeuro (HORIZON-WIDERA-2021-ACCESS-03 - 101079392). JI was supported by project ECHOES (ERC Starting Grant 101077693), and by a Ramón y Cajal grant (RYC2021-031905-I) funded by MCIN/AEI/10.13039/501100011033 and by the NextGeneration EU/PRTR fund. AK was supported by NISNEM (EPSRC EP/T020970/1). DF was supported by NaturalBionicS (ERC Synergy 810346) and NISNEM (EPSRC EP/T020970/1).

computing architectures such as graphics processing units.

Another potential reason for the poor adoption of higher-order BSS algorithms in the invasive neurophysiological recording domain is that current methodologies are not specialised towards solving the BSS problem in spiking distributions, which are highly sparse and skewed. In particular, the nature of these distributions suggests a different optimal form for the contrast function, the non-linearity used as a metric of statistical independence. In the successful FastICA algorithm, on which most higher order BSS algorithms for spike sorting are based [17], [18], [14], [19], this contrast function is designed to be robust to outliers in the source. However, for spiking sources it is precisely these outliers that are interesting. This motivates the design of new contrast functions, in particular adaptive forms that can balance the requirements of fine source separation and optimisation stability. In this study, we propose to find the best form of this contrast function for each source using a particle swarm optimisation (PSO) methodology.

In summary, in this study we contribute the following:

1) We explore in detail how the design of the non-linearity in projection pursuit ICA affects the successful separation of invasive neurophysiological recordings, whose sources often have very similar AP waveforms.
2) Using these insights, we propose an adaptive contrast function which can be modified for each source to maximise the chances of a successful separation vector optimisation. We incorporate this into swarm contrastive decomposition (SCD), a novel automated BSS algorithm which uses particle swarm optimisation to find the best non-linearity for each source.
3) We demonstrate state-of-the-art performance of SCD on a range of experimentally-collected intramuscular EMG datasets, before going on to show potential utility in the intracortical domain.

II. Theory and Algorithm

A. Blind Source Separation in Neurophysiological Signal

The signal detected during neurophysiological recordings is a volume conducted electrical action potential from the activation of one or more cells [20]. The volume conductor has an attenuating effect on the magnitude of the potential recorded, meaning that sensing electrodes that are further away from the source will read a lower voltage than one closer [21]. In the context of a HD microelectrode array, and assuming each sensing electrode has the same active area, the detected activity from a source will be much more prominent in some channels than others. Thus, if we disregard non-stationarities by assuming a fixed anatomical geometry and no movement in the recording array, each source will contribute an associated spatiotemporal AP waveform to the signal on every activation [22]. For EMG, the waveform is not the output of a single cell, but instead represents the synchronised activity of the set of muscle fibers controlled by a single motor neuron [23], called a motor unit (MU), and is called a MU action potential (MUAP). The objective of BSS in spiking time series is to identify the timestamps of all source activations and to assign each timestamp to the correct source.

We can model the AP waveform as the result of a finite impulse response (FIR) filter of time support L, with each time instant l being specified by a $M \times N$ dimension matrix $\boldsymbol{H}(l)$, where M is the number of sensing electrodes in the array and N is the number of spiking sources. The temporal support of the spikes in the source is assumed to be much less than L, so they are approximated as a binary time series, with the activity at each time instant t being specified by the $N \times 1$ vector $\boldsymbol{s}(t)$. The $M \times 1$ observation vector $\boldsymbol{x}(t)$ of activity sensed by an array of M electrodes at time t is consequently a convolutive mixture of the FIR filters being driven by the binary time series, with an $M \times 1$ additive noise vector $\boldsymbol{\xi(t)}$:

$$\boldsymbol{x}(t) = \sum_{l=0}^{L-1} \boldsymbol{H}(l)\boldsymbol{s}(t-l) + \boldsymbol{\xi}(t) \tag{1}$$

The convolutive mixture in equation 1 can be written in an instantaneous form if we extend $\boldsymbol{x}(t)$ and $\boldsymbol{s}(t)$ with their L past values. In practice, L is not known, so we instead choose a value G, which is a hyperparameter choice made on the basis of the stability of the inversion procedure being used to approximate $\boldsymbol{s}$:

$$\tilde{\boldsymbol{x}}(t) = \tilde{\boldsymbol{H}}\tilde{\boldsymbol{s}}(t) + \tilde{\boldsymbol{\xi}}(t) \tag{2}$$

where $\tilde{\boldsymbol{x}}$ is the $M(G-1) \times 1$ extended observation vector:

$$\begin{aligned} \tilde{\boldsymbol{x}}(t) &= [\tilde{\boldsymbol{x}}_1(t), ..., \tilde{\boldsymbol{x}}_m(t) ..., \tilde{\boldsymbol{x}}_M(t)]^T \\ \tilde{\boldsymbol{x}}_m(t) &= [x_m(t), ..., x_m(t-g), ..., x_m(t-G)] \end{aligned} \tag{3}$$

$\tilde{\boldsymbol{s}}$ is the $N(G-1) \times 1$ extended source vector:

$$\begin{aligned} \tilde{\boldsymbol{s}}(t) &= [\tilde{\boldsymbol{s}}_1(t), ..., \tilde{\boldsymbol{s}}_n(t) ..., \tilde{\boldsymbol{s}}_N(t)]^T \\ \tilde{\boldsymbol{s}}_n(t) &= [s_n(t), ..., s_n(t-g), ..., s_n(t-G)] \end{aligned} \tag{4}$$

and $\tilde{\boldsymbol{H}}$ is the $M(G-1) \times N(G-1)$ extended filter matrix, also called a mixing matrix:

$$\tilde{\boldsymbol{H}} = \begin{bmatrix} \tilde{\boldsymbol{h}}_{11} & \cdots & \tilde{\boldsymbol{h}}_{1n} & \cdots & \tilde{\boldsymbol{h}}_{1N} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \tilde{\boldsymbol{h}}_{m1} & \cdots & \tilde{\boldsymbol{h}}_{mn} & \cdots & \tilde{\boldsymbol{h}}_{mN} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \tilde{\boldsymbol{h}}_{M1} & \cdots & \tilde{\boldsymbol{h}}_{Mn} & \cdots & \tilde{\boldsymbol{h}}_{MN} \end{bmatrix} \tag{5}$$

$$\tilde{\boldsymbol{h}}_{\boldsymbol{mn}} = \begin{bmatrix} h_{mn}(0) & \cdots & h_{mn}(G-1) & 0 & \cdots & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & h_{mn}(0) & \cdots & h_{mn}(G-1) \end{bmatrix}$$

The objective of BSS is thus to stably find the separation matrix $\boldsymbol{W}$, which is the approximate inverse of $\tilde{\boldsymbol{H}}$. The columns of $\boldsymbol{W}$ are the separation vectors $\boldsymbol{w}_n$, recovering an estimate of the nth source when applied to the extended observation matrix.

B. Projection Pursuit Independent Component Analysis

Independent component analysis is a method for solving the BSS problem where the separation matrix $\boldsymbol{W}$ is optimised such that it maximises a metric of statistical independence [13]. Statistical independence can be defined in a number of ways, but broadly, two sources are independent if they remain uncorrelated after any non-linear transformation of either or both of the sources [24]. The non-linearity used to estimate independence is often called a contrast function, and is an important hyperparameter. By the central limit theorem, when sources are mixed the resultant distribution tends to be more Gaussian than that of the original sources, meaning that the independent sources can be discovered by finding projections which are maximally independent from a Gaussian distribution.

To improve optimisation stability and speed, the extended observation matrix is centred and then whitened prior to an ICA run. The advantage of pre-whitening the data is that the eventual separation matrix $\boldsymbol{W}$ will be orthogonal, meaning that far fewer parameters need to be estimated during optimisation [13]. This operation involves the estimation of a transformation matrix $\boldsymbol{A}$ for $\tilde{\boldsymbol{x}}$ which leaves the channels of the $M(G-1)\times 1$ whitened observation vector $\tilde{\boldsymbol{z}}$ linearly decorrelated:

$$\begin{aligned} \tilde{\boldsymbol{z}} &= \boldsymbol{A}\tilde{\boldsymbol{x}} \\ E[\tilde{\boldsymbol{z}}\tilde{\boldsymbol{z}}^T] &= \boldsymbol{I} \end{aligned} \tag{6}$$

where E is the expectation and $\boldsymbol{I}$ is the identity matrix. The matrix $\boldsymbol{A}$ can always be found, for example by the zero-component analysis (ZCA) method [25], which uses singular value decomposition of the centred observation covariance:

$$\begin{aligned} A &= \boldsymbol{U}\boldsymbol{\Sigma}^{-\frac{1}{2}}\boldsymbol{U}^T \\ E[\tilde{\boldsymbol{x}}\tilde{\boldsymbol{x}}^T] &= \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T \end{aligned} \tag{7}$$

where $\boldsymbol{\Sigma}$ is a diagonal matrix.

In projection pursuit formulations of ICA, such as the FastICA algorithm [26], sources are found one-by-one using an optimisation method such as gradient descent or fixed point iteration. Each $\boldsymbol{w}_n$ is progressively updated such that it maximises the contrast function. To prevent the same source being found multiple times, the FastICA algorithm uses source deflation, where on each iteration the updated $\boldsymbol{w}_n$ is projected into the orthogonal space of the previously found separation vectors. However, for spike sorting it is usually quite simple to estimate the contribution of a source to the signal by convolving the spike triggered average with the source [16], which can then be subtracted from the signal, a process called peel-off.

C. An Adaptable Contrast Function

The selection of the contrast function is important for the stability of numerical optimisation methods. One family of non-linearities which has an intuitive theoretical appeal are the cumulants, where the ith order cumulant κ_i is defined by the ith derivative of the cumulant generating function $K(\gamma)$ evaluated at $\gamma = 0$:

$$\begin{aligned} \kappa_i &= \left.\frac{d^i K(\gamma)}{dt^i}\right|_{\gamma=0} \\ K(\gamma) &= \log E[e^{\gamma \boldsymbol{s}_n}] \end{aligned} \tag{8}$$

If the n-th estimated source is first Z-score standardised, i.e.:

$$y_n = \frac{\boldsymbol{s}_n - E[\boldsymbol{s}_n]}{(E[\boldsymbol{s}_n^2] - E[\boldsymbol{s}_n]^2)^{\frac{1}{2}}} \tag{9}$$

Then the first six cumulants are as follows:

$$\begin{aligned} \kappa_1(\boldsymbol{y}_n) &= 0 \\ \kappa_2(\boldsymbol{y}_n) &= 1 \\ \kappa_3(\boldsymbol{y}_n) &= E[\boldsymbol{y}_n^3] \\ \kappa_4(\boldsymbol{y}_n) &= E[\boldsymbol{y}_n^4] - 3 \\ \kappa_5(\boldsymbol{y}_n) &= E[\boldsymbol{y}_n^5] \\ \kappa_6(\boldsymbol{y}_n) &= E[\boldsymbol{y}_n^6] - 15E[\boldsymbol{y}_n^4] + 30 \end{aligned} \tag{10}$$

Gaussian distributions have no higher order cumulants as they are entirely specified by κ_1 and κ_2, respectively the mean and variance. So by optimising the separation vector to maximise a higher order cumulant, most commonly the kurtosis κ_4, nongaussianity is also maximised and therefore the likelihood that the source is independent. In practice, as the cumulant order increases, so does its sensitivity to outliers [24]. This can be a problem in many signal domains, motivating the use of more robust cumulant estimators such as the hyperbolic tangent. However, for spiking sources, outliers are precisely the samples of interest. The cumulant order can therefore be seen as a method of tuning the selectivity of $\boldsymbol{w}_n$ to the spikes in a signal and so is extremely important for accurate decomposition.

We propose to adaptively find the contrast function that has the optimal level of selectivity. Cumulants have the disadvantage of being specified only by the integer order, so instead we propose a continuous and strictly asymmetric polynomial approximate:

$$f(\boldsymbol{y}_n) = E[sign(\boldsymbol{y}_n)|\boldsymbol{y}_n|^e] \tag{11}$$

where $sign(.)$ is a function that operates on each element in a vector, returning 1 when an element is positive and -1 when negative. An asymmetric function is chosen to match the distribution asymmetry of spiking sources, whose samples are mostly close to zero with rare excursions into highly positive values during an activation.

As an example for why tuning the contrast function is important in spike sorting in particular, consider an observation matrix consisting of additive noise and two active sources with very similar AP waveforms. When the value of e in equation 11 is low, its value will be dominated by all spikes in the signal, and $\boldsymbol{w}_n$ will therefore optimise to enhance all spikes, failing to separate the two sources. When e is optimal, it will be dominated by a smaller pool of outliers, and $\boldsymbol{w}_n$ will enhance one source and suppress the other. At higher values

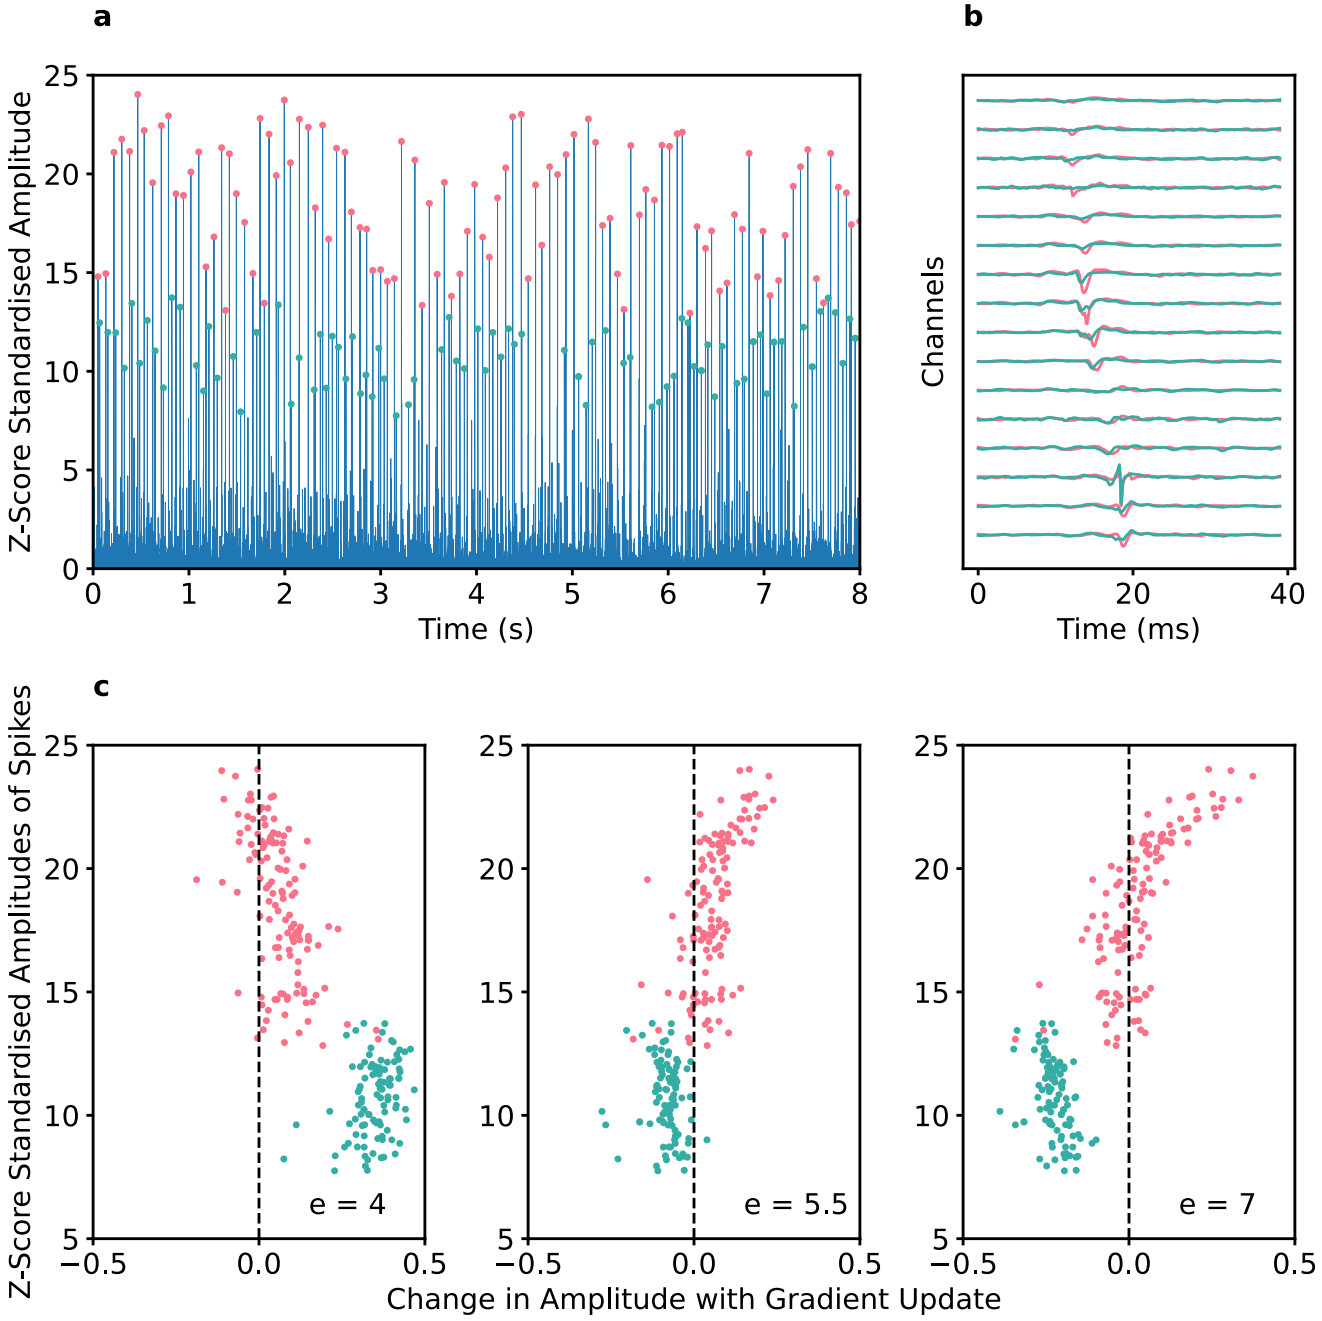

Fig. 1. **a** The Z-score standardised source activity estimated by applying a partially-optimised separation vector to whitened intramuscular EMG signal. The activity features two sources (pink and green) which have very similar action potential waveforms and hence are both being highlighted by the separation vector. **b** Superimposed spike triggered averages of the two sources. **c** The change in spike amplitude after a gradient step versus the current spike amplitude for three different values of e in $f(\boldsymbol{y}_n) = E[sign(\boldsymbol{y}_n)|\boldsymbol{y}_n|^e]$. When e is too small at 4.0, between-source variance is suppressed by the separation vector. When e is too high at 7.0, in-source variance is enhanced by the separation vector and the optimisation begins to collapse to a single spike. Stable source separation begins only when e is at an optimal value at 5.5, suppressing the spikes from the green source, whilst maintaining those from the pink source.

of e, the outlier pool will decrease further, and the within-source variance in the enhanced source's waveforms caused by the additive noise will start to dominate, until eventually $\boldsymbol{w}_n$ collapses down to enhancing only a single spike. In short, there exists an optimal value of e for each source, where $\boldsymbol{w}_n$ is enhancing between-source waveform variance, but not enhancing within-source waveform variance.

Figure 1 uses experimental iEMG to illustrate this effect. For lower values of e, between-source variance for sources with similar AP waveforms (refer to panel **b** to visualise the similarity between the sources) is suppressed and sources are incorrectly merged, whereas for higher values of e the in-source variance is enhanced and the source begins to collapse. The main issue is that the optimal value of e is unknown *a priori* and likely to vary between sources and data types. However, in the spike sorting literature there exists a number of efficient methods for validating and comparing the quality of decomposition automatically, i.e., the experiments are expensive, but the validation is cheap. This motivates the use of an efficient search algorithm that quickly finds the optimal e to build the most selective separation vector [27], [28].

D. Particle Swarm Optimisation of Contrast Functions

Particle swarm optimisation is a population-based stochastic optimisation methodology [29]. A number of candidate solution particles are evaluated, with the particle positions then updated on each step based on their personal best solution, the global best solution across particles and an inertia term which encourages exploration [30]. We propose to use this methodology to optimise the exponent e of equation 11, using as the objective function a metric of source quality $Q(\boldsymbol{s}_n)$ after a projection pursuit ICA run. J particles are created, each represented by a duplicate of a randomly initialised separation vector and an initial value of e (refer to Section III-D). On each step an ICA run is conducted and then the j-th particle's current value of e, e_j^0 is updated to its new value e_j^1 by:

$$e_j^1 = e_j^0 + v_j^1 \tag{12}$$

where v_j^1 is the particle's new velocity and is calculated from the old velocity v_j^0 by:

$$v_j^1 = wv_j^0 + c_1 r_1 (pBest_j - e_j^0) + c_2 r_2 (gBest - e_j^0) \tag{13}$$

with $pBest_j$ being the personal best of the j-th particle, i.e. the value of e which returns the highest $Q(\boldsymbol{s}_j)$. Likewise, the global best, $gBest$, is the value of e that returns the best quality metric value so far. The inertia coefficient w, cognitive coefficient c_1 and social coefficient c_2 are hyperparameters which tune the balance of the particle update towards its current trajectory, its personal best and the global best, determining exploration-exploitation trade-offs. Finally, r_1 and r_2 are samples from a normal distribution of mean 0 and standard deviations σ_1 and σ_2.

An important feature of the algorithm, which we call Swarm-Contrastive Decomposition (SCD), is that on each PSO step, the separation vector for all particles is reinitialised to be the spike triggered average of the spike timestamps from the best source from that round of ICA, as estimated by a quality metric Q. This helps counter the fact that small changes in e can lead to different sources being found by their respective ICA runs, without heavily constraining the optimisation. SCD is run until the global best stops improving for a set number of steps, at which point a final review of the best source is conducted, with the time stamps of the acceptable sources saved and the contribution of the source to the observation matrix estimated and peeled off.

E. Source Quality Metrics

A number of source quality metrics exist for automatically assessing the success of a decomposition operation. Selecting a sensitive and accurate metric is vitally important for the success of the SCD algorithm, with different metrics being the most appropriate for different data types. In this study, we use a robust coefficient of variance calculated on the interspike intervals for the iEMG data and a pseudo-silhouette score calculated post-thresholding for the intracortical data. A peak finding algorithm is first used to filter the source samples, followed by a two class k-medoid clustering algorithm to select the peaks which are likely spikes and add them to a timestamp set.

MU firing rates are quite regular in case of sustained contractions at a fixed target force [31], as in our experiments, meaning a good metric of source quality is the coefficient of variation (CoV) of the interspike intervals. The interspike intervals are first calculated on the timestamp set, with intervals five times larger than the median discarded. Rather than then estimating the coefficient of variation directly, we instead opted for a bootstrap version which penalises sources with less spikes, which can bias the CoV. We create 1000 timestamp sets by sampling from the original timestamp set with replacement, and then calculate the CoV on each set, the standard deviation of the intervals divided by the mean. We then take the 75th quantile of the 1000 CoV values, which we will refer to as the Q_{COV}.

In intracortical data, there is no similar expectation of regularity in the firing statistics. Instead, we use the values of the two k-medoids and the sample assignments to calculate a pseudo-silhouette score:

$$Q_{SIL} = \frac{1}{D}\sum_{d=1}^{D} \frac{|\alpha_d - \bar{\mu}_{out}| - |\alpha_d - \bar{\mu}_{in}|}{max(|\alpha_d - \bar{\mu}_{out}|, |\alpha_d - \bar{\mu}_{in}|)} \tag{14}$$

where D is the number of timestamps, α_d the source's value at the d-th timestamp. $\bar{\mu}_{in}$ and $\bar{\mu}_{out}$ are respectively the values of the medoid that the d-th timestamp is and is not assigned to.

III. Materials and Methods

A. Experimental Datasets

The effectiveness and reliability of the proposed SCD algorithm were validated using a diverse range of invasive HD iEMG and intracortical recordings.

The iEMG signals were recorded from either the Tibialis Anterior (TA) or the forearm muscles. In all iEMG experiments, multi-channel intramuscular electrodes designed for acute recordings were used [17], [32]. The EMG signals were recorded with a multi-channel amplifier (OT-Bioelettronica, Torino, Italy), sampled at 10240 Hz, high pass filtered at 100 Hz, and finally A/D converted with 12-bit resolution for the TA sets and with 16-bit resolution for the forearm data. The EMG signals were acquired in unipolar derivation. A reference electrode was placed on the ankle (TA) or the wrist (forearm). In all experiments, the participants were instructed to sustain an isometric contraction (ankle dorsiflexion for the TA, and extension of the little finger for the forearm). The relative forces were determined in all cases as percentages of the maximal voluntary contraction (MVC). A visual feedback of the exerted force and the target was provided to the subject. All experimental procedures adhered to the ethical guidelines set by Imperial College London (FM) and by the Ethical Committee of the University Medical Centre of Göttingen (TA1 and TA2). The studies were performed according to the Declaration of Helsinki, with an informed consent form signed by all participants before each experiment.

The TA recordings were made using two thin-film intramuscular electrodes, with 16 platinum detection points (140 $\mu m\times$ 40 μm) placed on a linear array with 1 mm interelectrode distance [32], inserted in the TA muscle of two participants (TA1 and TA2), approximately 2 cm apart in the longitudinal direction. The subjects were seated in a Biodex System 3 (Biodex Medical Systems Inc., NY, USA) with the right leg and foot constrained to the equipment. The subjects were instructed to perform three 30 second isometric contractions at 10, 20, and 30 %MVC. A central segment of 10 s was then used for algorithm validation. To study the robustness of the algorithm to different noise conditions, we generated signals with [30, 25, 20, 15, 10, 5, 0] dB of Additive White Gaussian Noise (AWGN) to the original reconstructed signals [33].

For the forearm muscles, data were recorded from three thin-film arrays with 40 platinum electrodes (140 $\mu m\times$ 40 μm) linearly distributed over a length of 2 cm, with 0.5 mm interelectrode distance [17], inserted in the forearm muscles of one participant (FM). To capture the precise activity of the fingers in extension, the muscles targeted were: the extensor pollicis longus (EPL), the extensor digitorum (ED), and the flexor digitorum superficialis (FDS). The experimental protocol was designed following MRI scans of the forearm, and the insertions were performed with the guidance of a portable ultrasound (Butterfly iQ+) (Figure 2 **a**). The forearm and hand of the participant were constrained in an instrumented platform to measure isometric forces from each finger individually, in flexion and in extension (Figure 2 **b**). The data acquisition framework featured ten calibrated load cells (TAL 220, 10 kg) mounted on a platform and tailored to accommodate the individual size of the participant. The force data were concurrently recorded with the EMG data through 10 auxiliary connections with the Quattrocento to ensure perfect synchronization between the physiological signals and the behavior. Within this framework, the participant was asked to follow a trapezoidal force trajectory at 15% MVC whilst extending the little finger. As with the TA recordings, a central segment of 10 s was then used for algorithm validation.

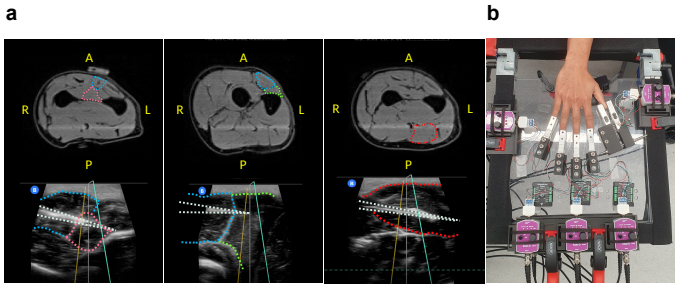

Fig. 2. **a** MRI scans of the forearm. The boundaries of the muscles targeted are highlighted in color. In the first MRI image, a marker is visible, used to determine the precise relative positioning of the slice image within the forearm's anatomical structure. Ultrasound images of the muscles during the insertion of the thin-film electrodes. The boundary of the guiding needle is highlighted in white. The boundaries of the relevant muscles are in color (EPL: Pink, ED: Blue, Extensor Carpi Radialis (ECR): green, FDS: red). **b** Platform built for the forearm experimental protocol.

In addition to iEMG signals, we tested the proposed algorithm on a publicly-available dataset of invasive intracortical recordings [34]. The dataset contains signals recorded from two or three levels of the neocortex of 20 anaesthetised rats using a single-shank silicon probe with a 32x4 dense array. The signals were sampled at 20 kHz and filtered with a 250 Hz high-pass cut-off. Each recording had also been decomposed into units using the Kilosort2 MATLAB package [35], followed by further manual curation, which gave a validated ground truth that we could compare SCD against. We selected the first 30 s of the recordings from the 12 rats with data from three anatomical levels, giving 36 time series to decompose.

B. Validation Methodology

We initially planned to validate the intramuscular against an existing set of spike timestamps [17], obtained by a skilled manual operator using the EMGLAB software [36]. However, we found that SCD discovered a number of units that were not in this set. These units were not detected by manual decomposition of the expert operator but, when their spike trains were shown to the operator, they were confirmed as valid sources. These sources were likely initially missed due to two reasons. Firstly, manual decomposition of signal from HD probes places an extremely high burden on the manual operator, which increases the risk of missing units. Secondly, some of the units found only by SCD were extremely small and missed by the automatic thresholding method in EMGLAB, despite having spiking temporal statistics and a clear spike triggered average (PTP 0.12 mV), with propagation of the AP across channels (Fig. 3). To compensate for this deficit in the ground truth, we elected to present the results by SCD in two modes. In SCDa, the algorithm was run in a fully automated mode, with no *post hoc* correction. In SCDm, we went on to load the output timestamps from SCDa into EMGLAB and perform a spike-by-spike validation of each unit, before searching for any units in the residual signal that may have been missed, as in normal EMGLAB operation. This

means that whenever SCDm has a higher yield than SCDa, it is because SCDa missed some units which were then found manually by an operator in EMGLAB.

We also compared SCDa to the current state-of-the-art method for automated iEMG decomposition [15], i.e. convolutive blind source separation (cBSS). This algorithm performs FastICA with a fixed contrast function set to $e = 3$ to find a source, before conducting an additional stage of optimisation, where the separation vector is recalculated using the spike triggered average and then reapplied to the signal in a loop until the CoV stops decreasing [37].

C. Assessment Metrics

For source acceptance, after a full set of PSO iterations, the Q_{SIL} is first calculated. For both iEMG and intracortical data, a source was rejected if the Q_{SIL} was below a threshold τ_{SIL} of 0.85. Of the sources that pass this check, iEMG sources are accepted if the Q_{COV} is below a threshold τ_{COV} and the firing rate Q_{FR} is below a threshold $\tau_{FR<}$, whilst for intracortical sources only the firing rate is calculated, which must be more than $\tau_{FR>}$ (refer to Section III-D).

We quantitatively assessed the accuracy of the automatic decomposition methods by determining their Rate of Agreement (RoA) with the ground truth [17], [38]. This evaluation was undertaken for each unit (MU), wherein the RoA was computed as the fraction of correctly identified discharges relative to the sum total of all discharges, both correctly and incorrectly identified. This relation can be formulated as:

$$RoA = \frac{O}{O + U_1 + U_2} \times 100 \tag{15}$$

Here, O refers to the number of matched predicted activations within a deviation margin of ±0.5 ms, whilst U_1 and U_2 represent the counts of unmatched predicted activations due to missed spikes and to additional, incorrect spikes, respectively.

The unit yield of decomposition was defined as the number of units that matched with a ground truth unit to an RoA of at least 30%, which was the case for all units.

D. Hyperparameters and Training

The hardware used for running the decomposition was an AMD Rome processor with an Nvidia RTX 6000 GPU. The *PyTorch* machine learning package was used to build the algorithm [39]. The optimal hyperparameters for the particle swarm update and source acceptance were determined by grid search, and were as follows: a standard deviation of 0.1 for both σ_1 and σ_2; respectively a value of 0.3 and 0.15 for c_1 and c_2 and finally an inertia of 1 which is then decayed by 0.1 with each PSO step. The number of particles J was set to 6 and associated starting contrast function e exponents set to [2, 3, 4, 5, 6, 7]. The same methodology was used to set the source acceptance thresholds τ_{SIL} and τ_{COV}, set to 0.85 and 0.4 respectively, whilst $\tau_{FR<}$ and $\tau_{FR>}$ was set to 35 Hz and 1 Hz respectively based on the literature [15], [40]. For the optimisation of the separation vector, we employed stochastic gradient descent with a learning rate of 0.1 and a momentum of 0.9. Within the main decomposition loop, the maximum number of iterations was limited to 1000. If no new acceptable sources were found consecutively for 30 iterations, the termination criterion would be triggered. A source was considered a repeat if it had a maximum rate of agreement of 30% with the previously found sources. In such cases, the timestamps of the repeated source would be used in the peel-off process again, but not then saved to the library of timestamp sets. The computational time taken for the decomposition of 10 s of the 30 %MVC signal of TA1 was 128 hours for the "Manual Only" approach and 32 minutes for SCDa.

IV. Results

A. Intramuscular Recording Decomposition

The proposed algorithm, SCDa, significantly outperformed the current state-of-the-art automated BSS algorithm, cBSS [15], for iEMG decomposition in both source yield and timestamp accuracy as measured by the rate of agreement with the manual decomposition (Table I). In fact, in a first for iEMG we are aware of in the literature, SCDa actually outperforms manual decomposition in terms of MU yield on four of the six recordings for which manual decomposition timestamps were available. The bulk of the sources found by SCDa and not manual decomposition had very small peak-to-peak amplitudes and were not detected by EMGLAB's automated thresholding procedure, as can be seen in Figure 3. The semi-automated approach SCDm, when a manual operator instead decomposes the iEMG using SCDa-generated timestamps as a reference, gives a higher number of sources than both the automated and manual methods. The median and IQR of the Q_{COV} and Q_{FR} for the 231 manually decomposed units were 12.3 (IQR 10.6-14.3) and 12.3 (IQR 10.8-13.6) respectively. For the 47 extra units found from SCDa and successively cleaned in EMGLAB we report a median Q_{COV} of 12.7 (IQR 11.2-14.5) and median Q_{FR} of 13.6 (IQR 12.4-14.9), further confirming their physiological characteristics. Consequently, we use the SCDm yield and timestamps as gold standards with which to compare the automated methods.

Across the iEMG recordings TA1, TA2 and FM, SCDa found 85.8% of the total sources found by SCDm, whilst cBSS found 28.1%. All 93 sources found using cBSS were also found by SCDa, with a further 191 sources found by SCDa. None of the sources found by either SCDa or cBSS were subsequently found to not be sources on manual inspection. Per-source decomposition accuracy of the automated algorithms was also assessed, by calculating the rate of agreement between their output timestamps and the manually-curated timestamps from SCDm. The median RoA of SCDa across all iEMG recordings was 98.8% (IQR 96.1%-100.0%), whilst for only those sources matched with cBSS it was 100.0% (IQR 99.5%-100.0%) versus 98.8% (IQR 98.5%-99.6%) for cBSS, a highly significant difference as measured by Wilcoxon signed-rank test ($p < 0.0001$).

B. Intracortical Recording Decomposition

In the mouse intracortical recordings SCDa found 1062 sources, 449 of which were matched with the manually-curated Kilosort2 decomposition (Fig. 4), with a median RoA

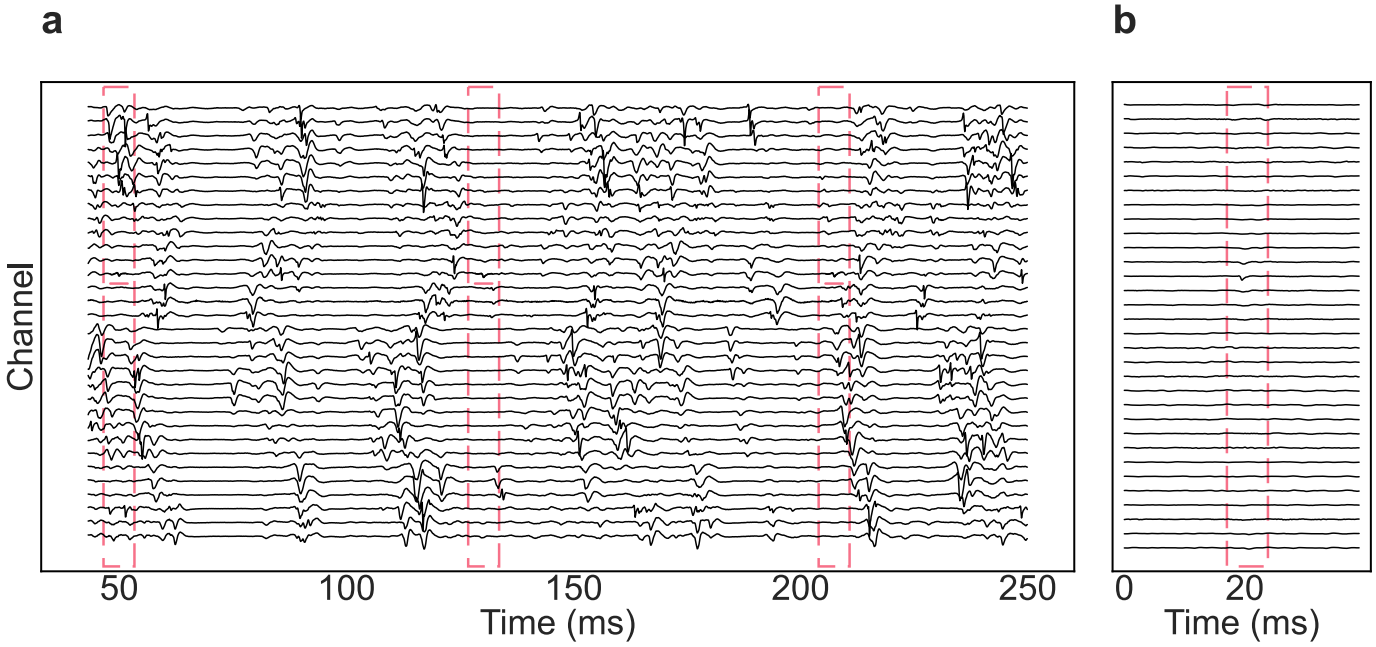

Fig. 3. **a** A section of HD iEMG signal from the tibialis anterior of subject 1. Highlighted in the red boxes are the spikes from a unit found by SCD, but missed by EMGLAB's automated thresholding and by the manual operator, due to a low peak-to-peak amplitude (0.12 mV). **b** Spike triggered average of the same unit. The waveform is clearly visible, which, combined with the clearly repeating nature of the timestamps, can allow us to be confident that this is indeed a real unit.

TABLE I
YIELD AND PER-SOURCE RATE OF AGREEMENT (MEDIAN AND INTERQUARTILE RANGES) RESULTS FOR THE INTRAMUSCULAR EMG DECOMPOSITIONS.

			Source Yield				RoAs (%)		
ID	Muscle	Force Level (%MVC)	cBSS [15]	SCDa	Manual Only (EMGLAB [36])	SCDm (SCDa→EMGLAB)	cBSS (all units / common with SCDa)	SCDa (common with cBSS)	SCDa (all units)
TA1	TA	10	10	48	53	57	95.5 (95.1-96.6)	95.6 (95.3-99.6)	96.1 (94.4-99.3)
	TA	20	13	66	57	71	98.4 (97.0-99.2)	100.0 (99.2-100.0)	99.1 (96.6-100.0)
	TA	30	11	64	60	82	98.5 (96.7-100.0)	99.7 (98.7-100.0)	95.9 (91.6-99.2)
TA2	TA	10	11	17	14	17	100.0 (100.0-100.0)	100.0 (100.0-100.0)	100.0 (99.3-100.0)
	TA	20	13	16	20	21	99.0 (97.3-100.0)	100.0 (98.5-100.0)	100.0 (98.5-100.0)
	TA	30	12	28	27	30	99.6 (99.1-100.0)	100.0 (99.0-100.0)	98.8 (96.8.0-100.0)
FM	ED	15	4	10	-	13	98.6 (95.2-100.0)	99.0 (96.6-99.5)	95.5 (92.8-98.4)
	EPL	15	8	10	-	10	98.8 (97.8-99.4)	100.0 (99.1-100.0)	100.0 (99.1-100.0)
	FD	15	11	25	-	30	100.0 (99.6-100.0)	100.0 (99.6-100.0)	98.8 (95.3-100.0)

of 96.7% (IQR 92.8%-98.5%). Figure 5 shows an example raster from an SCDa decomposition from depth two of rat six, alongside some waveform examples and principal component projections. SCDa was able to find even very low peak-to-peak amplitude AP waveforms. Another advantage is that the separation vector acts to align the AP waveforms temporally, meaning no further alignment is needed after thresholding, as is needed in lower-order spike sorting methods like KiloSort [7].

C. Effect of Particle Swarm Optimisation

Once it was confirmed that SCDa was performing well across the data types included in this study, we further examined the effect of using PSO to discover the best ICA contrast function for each source. Figure 6 shows how the PSO algorithm was able to quickly converge on an optimal contrast exponent e for each source, which varied significantly at the population level between data types (one-way ANOVA, $p < 0.0001$).

A further ablation study was performed using TA1 and TA2, were the PSO component was removed and the cubic contrast

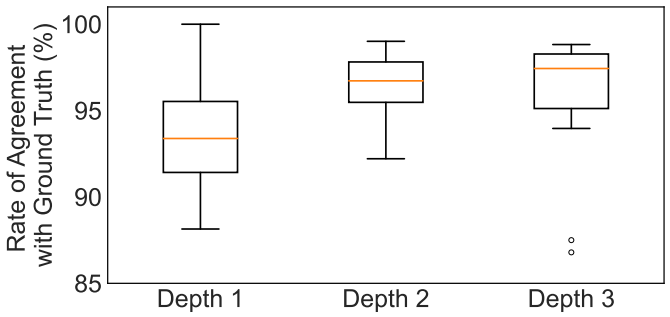

Fig. 4. Distribution box plot of the median rates of agreement calculated between SCDa and KiloSort2 across all intracortical recordings.

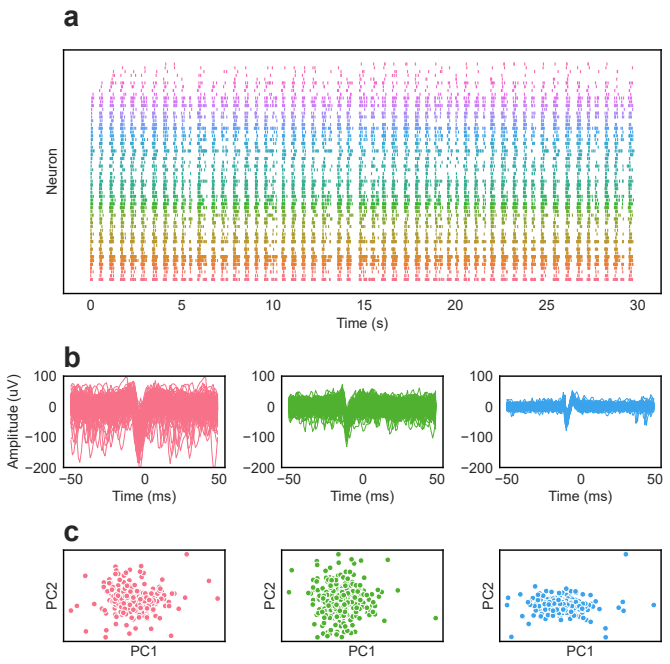

Fig. 5. **a** Example raster plot of all sources after running SCDa. **b** Spike-triggered waveforms of three found sources with low peak-to-peak amplitudes ($\leq$ 0.30 mV). **c** Projection of the action potential waveforms in principal component space for each of the above sources. One cluster is visible, confirming the correct feature was extracted.

function of $E[\boldsymbol{y}_n^3]$ used instead [15], the results of which are shown in table II. SCDa had a significantly higher yield than the fixed projection pursuit ICA as assessed by paired t-test (p = 0.01), proving its utility in the decomposition of invasive neurophysiological data.

D. Noise Sensitivity Analysis

Finally, in Table III, we report the source yield for the dataset when a range of additive normally distributed noise was added to the reconstructed signals to give specific signal to noise ratios (SNR) as calculated by Welch's method. The yield is relatively stable until 15 dB, when it decreases quickly, eventually collapsing at 0 dB. The RoA is consistently high across different levels of additive noise, most likely because these are only the largest amplitude AP waveforms that can be found.

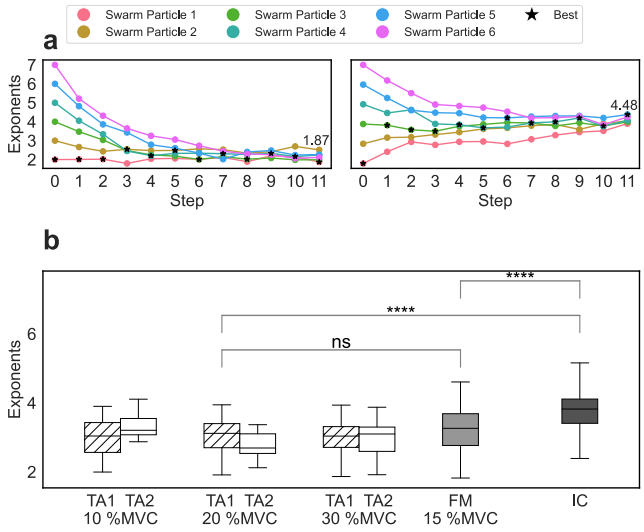

Fig. 6. **a**: Convergence of swarm particles for two sources during projection pursuit ICA on iEMG data. Particle swarm optimisation quickly finds the optimal exponent value e in the contrast function $f(\boldsymbol{y}_n) = E[sign(\boldsymbol{y}_n)|\boldsymbol{y}_n|^e]$ for each source. **b** Box and whisker plots of the final best value of e across the different data types, showing the optimal e can be significantly different.

TABLE II
YIELDS AND *MEDIAN RATE OF AGREEMENTS (IQR)* FOR TA1 AND TA2 OF A FIXED CUBIC CONTRAST (AS IN [15]) VERSUS SCDA.

ID	Force Level (%MVC)	Fixed Cubic	SCDa
TA1	10	**42**, *95.9 (93.8-97.9)*	**48**, *96.1 (94.4-99.3)*
TA1	20	**52**, *98.8 (97.2-99.5)*	**66**, *99.1 (96.6-100.0)*
TA1	30	**53**, *97.0 (92.8-98.5)*	**64**, *95.9 (91.6-99.2)*
TA2	10	**15**, *100.0 (99.3-100.0)*	**17**, *100.0 (99.3-100.0)*
TA2	20	**9**, *100.0 (99.2-100.0)*	**16**, *100.0 (98.5-100.0)*
TA2	30	**26**, *99.1 (96.4-100.0)*	**28**, *98.8 (96.8.0-100.0)*

TABLE III
THE EFFECT OF ADDITIVE WHITE NOISE ON SCDA **YIELDS** AND *MEDIAN RATE OF AGREEMENTS (IQR)* COMPARED TO THE GROUND TRUTH.

SNR (dB)	10 (%MVC)	20 (%MVC)	30 (%MVC)
30	**44**, *100.0 (99.0-100.0)*	**44**, *99.3 (98.5-100.0)*	**47**, *100.0 (98.8-100.0)*
25	**42**, *100.0 (97.4-100.0)*	**49**, *99.3 (98.5-100.0)*	**45**, *100.0 (97.9-100.0)*
20	**36**, *100.0 (96.6-100.0)*	**45**, *99.3 (97.1-100)*	**46**, *99.1 (94.3-100.0)*
15	**31**, *100.0 (96.5-100.0)*	**33**, *99.2 (95.4-100.0)*	**38**, *98.7 (93.7-100.0)*
10	**25**, *98.9 (96.6-100.0)*	**23**, *97.1 (92.6-99.3)*	**19**, *97.9 (91.1-99.7)*
5	**11**, *97.7 (96.3-100.0)*	**11**, *97.0 (94.4-99.6)*	**9**, *97.2 (94.5-99.3)*
0	**1**, *96.9*	**3**, *97.0 (95.3-98.0)*	**1**, *95.8*

V. CONCLUSION

This study presents a re-evaluation of the non-linear contrast function used to estimate independence of a source in projec-

tion pursuit ICA, specifically when those sources are the sparse time series that make up multichannel invasive recordings. We show that there exists an optimal level of non-linearity when optimising a separation vector: too low and the vector will fail to separate similar sources; too high and the optimisation collapses as the within-source variance begins to dominate the selectivity of the vector. We then demonstrate that a particle swarm optimisation method can quickly find this optimal non-linearity. The resultant BSS algorithm, swarm contrastive decomposition, extracts large volumes of neural sources in a variety of invasive data, including state-of-the-art performance in the intramuscular EMG domain, even outperforming manual source separation in some recordings.

A high-accuracy and fully automated decomposition pipeline offers a realistic method by which high-density intramuscular recordings can be processed. Consequently, it paves the way for population-level muscle data investigations, refining neural interfaces and analytical methodologies for a deeper understanding of neural movement control. Regarding potential limitations, like most BSS algorithms, the *post hoc* assessment of sources is limited by the accuracy of the quality metric. This motivates planned work integrating better methods of automatically assessing neurophysiological sources. Whilst it was not a focus of the study, SCD showed promise in intracortical data as well; we hope to validate this more robustly in a future study.

VI. Acknowledgments

The authors express their gratitude for the valuable contributions made by Irene Mendez Guerra. We would like to thank Dr. Angela Kedgley for her technical guidance on the development of the platform to record data from the forearm muscles. This study utilised the expertise and equipment at the Imperial College Advanced Hackspace.

References

[1] E. Stark and M. Abeles, "Predicting movement from multiunit activity," *Journal of Neuroscience*, vol. 27, pp. 8387–8394, Aug. 2007.

[2] H. Sawada, N. Ono, H. Kameoka, D. Kitamura, and H. Saruwatari, "A review of blind source separation methods: two converging routes to ilrma originating from ica and nmf," *APSIPA Transactions on Signal and Information Processing*, vol. 8, p. e12, 2019.

[3] H. G. Rey, C. Pedreira, and R. Q. Quiroga, "Past, present and future of spike sorting techniques," *Brain research bulletin*, vol. 119, pp. 106–117, 2015.

[4] J. J. Jun, N. A. Steinmetz, J. H. Siegle, D. J. Denman, M. Bauza, B. Barbarits, A. K. Lee, C. A. Anastassiou, A. Andrei, Ç. Aydın, *et al.*, "Fully integrated silicon probes for high-density recording of neural activity," *Nature*, vol. 551, no. 7679, pp. 232–236, 2017.

[5] J. Putzeys, B. C. Raducanu, A. Carton, J. De Ceulaer, B. Karsh, J. H. Siegle, N. Van Helleputte, T. D. Harris, B. Dutta, S. Musa, *et al.*, "Neuropixels data-acquisition system: a scalable platform for parallel recording of 10 000+ electrophysiological signals," *IEEE transactions on biomedical circuits and systems*, vol. 13, no. 6, pp. 1635–1644, 2019.

[6] J. E. Chung, K. K. Sellers, M. K. Leonard, L. Gwilliams, D. Xu, M. E. Dougherty, V. Kharazia, S. L. Metzger, M. Welkenhuysen, B. Dutta, *et al.*, "High-density single-unit human cortical recordings using the neuropixels probe," *Neuron*, vol. 110, no. 15, pp. 2409–2421, 2022.

[7] M. Pachitariu, S. Sridhar, and C. Stringer, "Solving the spike sorting problem with kilosort," *bioRxiv*, pp. 2023–01, 2023.

[8] J. Magland, J. J. Jun, E. Lovero, A. J. Morley, C. L. Hurwitz, A. P. Buccino, S. Garcia, and A. H. Barnett, "Spikeforest, reproducible web-facing ground-truth validation of automated neural spike sorters," *Elife*, vol. 9, p. e55167, 2020.

[9] A. P. Buccino, S. Garcia, and P. Yger, "Spike sorting: new trends and challenges of the era of high-density probes," *Progress in Biomedical Engineering*, vol. 4, no. 2, p. 022005, 2022.

[10] K. M. Mayer, A. Del Vecchio, B. M. Eskofier, and D. Farina, "Unsupervised neural decoding of signals recorded by thin-film electrode arrays implanted in muscles using autoencoding with a physiologically derived optimisation criterion," *Biomedical Signal Processing and Control*, vol. 86, p. 105178, 2023.

[11] S. Takahashi, Y. Anzai, and Y. Sakurai, "Automatic sorting for multi-neuronal activity recorded with tetrodes in the presence of overlapping spikes," *Journal of neurophysiology*, vol. 89, no. 4, pp. 2245–2258, 2003.

[12] J. Lee, C. Mitelut, H. Shokri, I. Kinsella, N. Dethe, S. Wu, K. Li, E. B. Reyes, D. Turcu, E. Batty, *et al.*, "Yass: Yet another spike sorter applied to large-scale multi-electrode array recordings in primate retina," *BioRxiv*, pp. 2020–03, 2020.

[13] A. Hyvärinen and E. Oja, "Independent component analysis: algorithms and applications," *Neural networks*, vol. 13, no. 4-5, pp. 411–430, 2000.

[14] C. Leibig, T. Wachtler, and G. Zeck, "Unsupervised neural spike sorting for high-density microelectrode arrays with convolutive independent component analysis," *Journal of neuroscience methods*, vol. 271, pp. 1–13, 2016.

[15] F. Negro, S. Muceli, A. M. Castronovo, A. Holobar, and D. Farina, "Multi-channel intramuscular and surface EMG decomposition by convolutive blind source separation," *Journal of neural engineering*, vol. 13, no. 2, p. 026027, 2016.

[16] D. Farina and A. Holobar, "Characterization of human motor units from surface EMG decomposition," *Proceedings of the IEEE*, vol. 104, no. 2, pp. 353–373, 2016.

[17] S. Muceli, W. Poppendieck, A. Holobar, S. Gandevia, D. Liebetanz, and D. Farina, "Blind identification of the spinal cord output in humans with high-density electrode arrays implanted in muscles," *Science advances*, vol. 8, no. 46, p. eabo5040, 2022.

[18] D. Jäckel, U. Frey, M. Fiscella, F. Franke, and A. Hierlemann, "Applicability of independent component analysis on high-density microelectrode array recordings," *Journal of neurophysiology*, vol. 108, no. 1, pp. 334–348, 2012.

[19] A. P. Buccino, E. Hagen, G. T. Einevoll, P. D. Häfliger, and G. Cauwenberghs, "Independent component analysis for fully automated multi-electrode array spike sorting," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 2627–2630, IEEE, 2018.

[20] S. B. Rutkove, "Introduction to volume conduction," *The clinical neurophysiology primer*, pp. 43–53, 2007.

[21] D. F. Stegeman, D. Dumitru, J. C. King, and K. Roeleveld, "Near- and far-fields: source characteristics and the conducting medium in neurophysiology," *Journal of clinical neurophysiology*, vol. 14, no. 5, pp. 429–442, 1997.

[22] R. Merletti and D. Farina, *Surface electromyography: physiology, engineering, and applications*. John Wiley & Sons, 2016.

[23] R. Merletti and S. Muceli, "Tutorial. surface EMG detection in space and time: Best practices," *Journal of Electromyography and Kinesiology*, vol. 49, p. 102363, 2019.

[24] A. Hyvärinen, J. Karhunen, and E. Oja, *Introduction*. John Wiley Sons, Ltd, 2001.

[25] A. J. Bell and T. J. Sejnowski, "The "independent components" of natural scenes are edge filters," *Vision research*, vol. 37, no. 23, pp. 3327–3338, 1997.

[26] A. Hyvarinen, "Fast and robust fixed-point algorithms for independent component analysis," *IEEE transactions on Neural Networks*, vol. 10, no. 3, pp. 626–634, 1999.

[27] J. Huang, B. Xue, Y. Sun, M. Zhang, and G. G. Yen, "Particle swarm optimization for compact neural architecture search for image classification," *IEEE Transactions on Evolutionary Computation*, vol. 27, no. 5, pp. 1298–1312, 2023.

[28] G. Yuan, B. Wang, B. Xue, and M. Zhang, "Particle swarm optimization for efficiently evolving deep convolutional neural networks using an autoencoder-based encoding strategy," *IEEE Transactions on Evolutionary Computation*, pp. 1–1, 2023.

[29] J. Kennedy and R. Eberhart, "Particle swarm optimization," vol. 4, pp. 1942–1948, 1995.

[30] Y. Shi and R. Eberhart, "A modified particle swarm optimizer," in *1998 IEEE international conference on evolutionary computation proceedings. IEEE world congress on computational intelligence (Cat. No. 98TH8360)*, pp. 69–73, IEEE, 1998.

[31] E. Martinez-Valdes, F. Negro, C. Laine, D. Falla, F. Mayer, and D. Farina, "Tracking motor units longitudinally across experimental sessions

with high-density surface electromyography," *The Journal of physiology*, vol. 595, no. 5, pp. 1479–1496, 2017.

[32] S. Muceli, W. Poppendieck, F. Negro, K. Yoshida, K. P. Hoffmann, J. E. Butler, S. C. Gandevia, and D. Farina, "Accurate and representative decoding of the neural drive to muscles in humans with multi-channel intramuscular thin-film electrodes," *The Journal of Physiology*, vol. 593, no. 17, pp. 3789–3804, 2015.

[33] A. Mishra, V. Bhateja, A. Gupta, and A. Mishra, "Noise removal in eeg signals using swt–ica combinational approach," in *Smart Intelligent Computing and Applications*, pp. 217–224, Springer Singapore, 2019.

[34] C. Horváth, L. F. Tóth, I. Ulbert, and R. Fiáth, "Dataset of cortical activity recorded with high spatial resolution from anesthetized rats," 2015.

[35] M. Pachitariu, N. A. Steinmetz, S. N. Kadir, M. Carandini, and K. D. Harris, "Fast and accurate spike sorting of high-channel count probes with kilosort," *Advances in neural information processing systems*, vol. 29, 2016.

[36] K. C. McGill, Z. C. Lateva, and H. R. Marateb, "EMGLAB: An interactive EMG decomposition program," *Journal of Neuroscience Methods*, vol. 149, no. 2, pp. 121–133, 2005.

[37] A. Holobar and D. Zazula, "Multichannel blind source separation using convolution kernel compensation," *IEEE Transactions on Signal Processing*, vol. 55, no. 9, pp. 4487–4496, 2007.

[38] A. Holobar, M. A. Minetto, A. Botter, F. Negro, and D. Farina, "Experimental analysis of accuracy in the identification of motor unit spike trains from high-density surface EMG," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 18, no. 3, pp. 221–229, 2010.

[39] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, pp. 8024–8035, 2019.

[40] C. M. Laine, E. Martinez-Valdes, D. Falla, F. Mayer, and D. Farina, "Motor neuron pools of synergistic thigh muscles share most of their synaptic input," *Journal of Neuroscience*, vol. 35, pp. 12207–12216, 9 2015.