

Signal-to-peak-interference ratio maximization with automatic interference weighting for threshold-based spike sorting of high-density neural probe data

Jasper Wouters¹, Fabian Kloosterman^{2,3,4} and Alexander Bertrand¹

Abstract—An innovative filter design method is proposed for threshold-based spike sorting of high-density neural recordings. Threshold-based spike sorting is the process of assigning each detected spike in an extracellular recording to its putative neuron, using only linear filters and simple thresholding operations. The low computational complexity of threshold-based spike sorting makes it interesting for real-time (hardware) implementations with potential applications in the field of brain-machine interfaces. The proposed method extends our earlier work on discriminative template matching and avoids the need for a prior heuristic definition of an interference covariance matrix. A new optimal filter design objective function is proposed, which automatically selects interference-dominated signal segments based on the output signal of the filter under design. This new method leads to filters that are signal-to-peak-interference ratio (SPIR) optimal. The method is validated on ground truth data recorded in-vivo.

I. INTRODUCTION

The spiking activity of a population of neurons is widely believed to reflect inter-neuronal communication. This spiking activity can be measured and recorded using extracellular electrodes placed in close proximity to those neurons. The spikes measured on a certain electrode typically originate from multiple neurons, and the spikes from a specific neuron are often picked up by multiple electrodes. Resolving such a multi-electrode recording into the spike times of its individual contributing neurons is known as spike sorting [1].

A multitude of methods have been proposed to solve the problem of spike sorting [2], [3], [4]. The development of spike sorting methods has been an ongoing process of exploiting progress in both measurement devices and computational techniques. Recently, the availability of high-density neural probes [5] has resulted in several new algorithms capable of processing the high-dimensional data collected by such probes [6], [7].

Many spike sorting algorithms contain the following steps in their algorithm pipeline: 1.) a spike detection step is performed, followed by 2.) a feature extraction phase, then those

This work was carried out at the ESAT Laboratory of KU Leuven, in the frame of KU Leuven Special Research Fund projects C14/16/057 and CoE PFV/10/002 (OPTEC), and the Research Foundation Flanders (FWO) project FWO G0D7516N (Distributed signal processing algorithms for spike sorting in nextgeneration high-density neuroprobes). The scientific responsibility is assumed by its authors.

¹KU Leuven, Electrical Engineering Dept. (ESAT), Stadius Center for Dynamical Systems, Signal Processing, and Data Analytics, Leuven, Belgium

²Neuro-Electronics Research Flanders (NERF), Leuven, Belgium

³KU Leuven, Brain & Cognition Research Unit, Leuven, Belgium

⁴VIB, Leuven, Belgium

extracted features are 3.) clustered, and finally 4.) template matching is performed to resolve overlapping spikes. The first three steps lead to the identification of the different recorded neurons that are active, also referred to as the recorded *units*. Unfortunately, the recording device also picks up spatio-temporally overlapping spikes from different neurons. Detected waveforms consisting of overlapping spikes can not be classified in the feature space [2], preventing a reliable determination of the *single-unit activity*.

To resolve those overlapping spikes, the algorithmic pipeline is augmented with the template matching phase. For every found cluster a template can be estimated. Care must be taken not to use non-unit clusters, e.g., clusters originating from correlated firing of multiple neurons. The templates are used to construct matched filters. Those matched filters are then used to determine the activity of the individual units [8].

Once the matched filters have been derived from training data, they can be applied to new data in a streaming fashion. Hence, template matching is also useful for real-time spike sorting. However, obtaining single-unit activity from the matched filter outputs requires iterative template subtraction to reliably sort the data [7], [8]. Such subtractive iteration increases the algorithmic complexity and introduces an additional undeterministic delay, which is undesirable for real-time applications.

To avoid the use of iterative template subtraction, discriminative template matching filters were proposed [9], [10], [11], [12]. The idea behind such discriminative filters is that they only have a high output variance when the corresponding neuron, also referred to as their *target* neuron, is active while suppressing interfering spikes from other neurons. As such the filter outputs can be converted into single-unit activity by the use of a simple thresholding operation. Because of the computational simplicity of performing threshold-based spike sorting, it is an interesting technique for real-time spike sorting applications. One possible application is a single-unit activity controlled brain-machine interface (BMI). Although the use of single-unit activity in BMIs have been questioned [13], recent work on BMIs still involve single-unit activity [14], [15]. Moreover, other recent work [16] suggests that single-unit activity can lead to insights that are beneficial for the long-term stability of BMIs.

However, such discriminative template matching filters were shown to be difficult to construct, when using conventional recording devices with low electrode densities [17].

Recently, discriminative template matching was applied to high-density neural probe data, and was shown to enable threshold-based spike sorting [12]. Also, a novel data-driven filter design method was proposed in [12] based on signal-to-peak-interference (SPIR) optimality. The data-drivenness removes the necessity of knowing all the templates of the interfering neurons compared to the earlier discriminative template matching filter design methods [9], [10], [11].

The filter design method in [12] made use of a prior heuristic procedure to define the interference-dominated segments in the data, leading to filter outputs that were not always truly SPIR-optimal. In some cases that filter design method (not the actual filtering) had to be re-applied several times in an iterative fashion to result in useful filters. Moreover, there was no clear convergence criterion available using the heuristic approach. In this work a formal approach on data-driven discriminative template matching filter design is presented, which optimizes the SPIR of the actual filter output. The proposed method includes an automatic interference weighting to avoid the need for a heuristic approach to define the interference covariance matrix.

The outline of this paper is as follows. In Section II the SPIR-optimal filter design method is derived. In Section III the method is validated on in-vivo recorded data. Finally, conclusions are drawn in Section IV.

II. METHOD

The method is derived here for a single target neuron. In practice, this method has to be applied in parallel (both filter design and actual filtering) for each neuron from which the single-unit activity needs to be determined.

A high-density probe recording is represented by $\mathbf{x}[k] \in \mathbb{R}^N$ with N the number of relevant electrodes on the probe from which a sample is observed at every discrete time point k . Only a local neighbourhood of N electrodes is considered around a center electrode which has the highest absolute voltage deflection for target spikes. This local neighbourhood is described by a radius R around that center electrode. All other channels are omitted from $\mathbf{x}[k]$. It is assumed that $\mathbf{x}[k] = \mathbf{s}_T[k] + \mathbf{s}_{NT}[k] + \mathbf{n}[k]$, where $\mathbf{s}_T[k]$ and $\mathbf{s}_{NT}[k]$ are the spiking activity of the target neuron and non-target neurons respectively, and $\mathbf{n}[k]$ contains both physiological and electrical noise. Consider $\bar{\mathbf{x}}[k] = [\mathbf{x}[k]^T \dots \mathbf{x}[k-L+1]^T]^T \in \mathbb{R}^{NL}$ containing an L -taps delay line for each electrode signal. The application of a linear spatio-temporal finite impulse response (FIR) filter with its coefficients described by \mathbf{f} , can then be written as $y[k] = \mathbf{f}^T \bar{\mathbf{x}}[k]$, with $y[k]$ the corresponding filter output. In this work we will focus on the design of such a spatio-temporal FIR filter \mathbf{f} , such that a simple thresholding operation *Thr* on the filter output power can be used to retrieve the single-unit activity spike train $u[k]$ of the target neuron:

$$u[k] = \begin{cases} 1 & \text{(target spike)} & \text{if } y[k]^2 > Thr \\ 0 & \text{(no target spike)} & \text{otherwise.} \end{cases} \quad (1)$$

Spike sorting is based on the assumption that neuronal spikes originating from the same neuron are measured as similar spatio-temporal voltage deflections on the probe. Clustering all detected spikes typically reveals different active units in the recording. As discussed in the introduction, to resolve overlapping spikes a template matching step can be performed. A template $\boldsymbol{\tau}$ is usually estimated by taking the median across the spike waveform instances $\bar{\mathbf{x}}[k]$ that belong to the cluster corresponding to the target neuron after proper temporal alignment [8].

Here, the goal is to design an optimal filter for threshold-based spike sorting. The filter design is powered by the following regularized SPIR optimization problem (note that the SPIR maximization is reformulated here as a constrained minimization problem):

$$\begin{aligned} \min_{\mathbf{f}} \frac{1}{|\mathcal{T}|} \sum_{k \in \mathcal{T}} w(\mathbf{f}, \bar{\mathbf{x}}[k]) (\mathbf{f}^T \bar{\mathbf{x}}[k])^2 + \lambda \mathbf{f}^T \mathbf{f} \\ \text{s.t. } (\mathbf{f}^T \boldsymbol{\tau})^2 = K, \end{aligned} \quad (2)$$

with \mathcal{T} a set containing some training data $\bar{\mathbf{x}}$, $w(\mathbf{f}, \bar{\mathbf{x}})$ the *activation weighting function* which selects interfering data samples, λ is a regularization constant, and K is an arbitrary positive constant that determines the filter's output power in response to a target spike instance. For conciseness the time indices in the mathematical notation will be omitted from here on. The activation weighting function is a shifted sigmoid as a function of the filter output power, described by:

$$w(\mathbf{f}, \bar{\mathbf{x}}) = \frac{1}{1 + e^{-((\mathbf{f}^T \bar{\mathbf{x}})^2 - \beta K)}}, \quad (3)$$

where β is used for setting the interference threshold. Here, a sigmoid was chosen as the activation weighting function, since the sigmoid is a differentiable approximation of a step function (at least when K is chosen sufficiently large, see also Section III). The activation weighting function is shown in Figure 1, where it is plotted as a function of the filter output power $(\mathbf{f}^T \bar{\mathbf{x}})^2$.




Fig. 1. Shifted sigmoid activation weighting function.

The solution of (2) thus tries to minimize the filter's output power for training data samples $\bar{\mathbf{x}}$ that generate a filter output power that is bigger than a β -fraction of the template's filter output power, while maintaining a fixed filter output power for that template. The regularization term was added to

prevent overfitting. In [12] the activation weighting function was implicitly defined a-priori using a heuristic method on the input data. The filter design method presented here uses an activation weighting function which depends on the actual filter output and which automatically activates relevant signal segments.

To solve the constrained non-linear non-convex optimization problem (2), SLSQP [18] is used, which is a gradient descent-based solver. Because the numerical evaluation of the objective function gradient is computationally intensive, the gradient is analytically derived and passed to the solver. The gradient of the objective function is given below for completeness:

$$\frac{1}{|\mathcal{T}|} \sum_{\mathcal{T}} \left(\nabla_{\mathbf{f}} w(\mathbf{f}, \bar{\mathbf{x}}) (\mathbf{f}^T \bar{\mathbf{x}})^2 + 2w(\mathbf{f}, \bar{\mathbf{x}}) \mathbf{f}^T \bar{\mathbf{x}} \bar{\mathbf{x}} \right) + 2\lambda \mathbf{f}, \quad (4)$$

with

$$\nabla_{\mathbf{f}} w(\mathbf{f}, \bar{\mathbf{x}}) = - \frac{e^{-((\mathbf{f}^T \bar{\mathbf{x}})^2 - \beta K)} (-2\mathbf{f}^T \bar{\mathbf{x}} \bar{\mathbf{x}} + 2\beta \mathbf{f}^T \boldsymbol{\tau} \boldsymbol{\tau})}{\left(1 + e^{-((\mathbf{f}^T \bar{\mathbf{x}})^2 - \beta K)}\right)^2}. \quad (5)$$

III. EXPERIMENTS

The proposed filter design method is validated on three ground truth datasets. Two in-vivo recorded ground truth datasets [19] and one hybrid dataset [4] are used. The performance of our algorithm is compared to the performance of the previously proposed discriminative template matching filter design method [12], which will be referred to as the heuristic method from here on.

From each dataset the first five minutes of data are used. The first 150 s of data are used for training the filters, whereas the last 150 s are used for testing. Prior to the filter design, the data is high-pass filtered with a cut-off frequency of 300 Hz. The radius describing a target spike neighbourhood is set to $R = 100 \mu\text{m}$.

The weighting function parameter β in (3) is chosen to be 0.1, as to try to attenuate the interfering peaks in the output to a SPIR of at least 10 dB. The desired template filter output power K should be chosen sufficiently high for the sigmoid to properly suppress interference. Assume we would choose $K = 1$, the value of the weighting function for data with a zero power output response would still be weighted by a factor 0.475. Therefore, we choose $K = 1000$, which allows for a wide range of output powers which are weighted by a factor close to 0, such that unimportant samples can be ignored during the filter design. The regularization parameter λ in (2) is chosen for every dataset separately in a grid-search fashion. The regularization parameter that maximizes the threshold-based spike sorting performance on the training data is retained. The values for λ are shown in Table I.

The template $\boldsymbol{\tau}$ is estimated using the ground truth timing information from the training data. In practice, this timing information can be obtained from clustering the training data, as explained in the introduction. It is also this template that

is used to initialise the numerical optimization solver with $\mathbf{f}^{\text{init}} = \alpha \boldsymbol{\tau}$, where α is chosen such that the constraint in (2) is satisfied.

The hybrid dataset (dataset 3) is a well-chosen dataset for which the heuristic method [12] does not result in a filter which is useable for threshold-based spike sorting in a straightforward way. The heuristic method has to be applied in an iterative fashion to resolve this problem, but because of its heuristic nature, no clear convergence criterion is available. In the heuristic case the goodness of the filter depends on an expert, who manually has to stop the iteration. Although the heuristic method was shown to perform well most of the time, two mechanisms could lead to suboptimal filters: firstly, if the interfering peaks fall into the so called target safe zone [12], they are ignored in the estimation of the interference covariance in which case they won't be suppressed in the output. Secondly, the heuristic filter projects the data onto a subspace in which there is no guarantee that new interferers are amplified which were not visible in the original input data. The newly proposed method does not have those problems.

Figure 2A shows that using the SPIR-optimal method vs. the heuristic without expert leads to different solutions. As can be seen from Figure 2B, the SPIR-optimal filter output response to a non-target spike (left) is largely reduced with respect to the target spike output response (right). This difference in response was not present in the heuristic design case. Such a difference in filter output response enables threshold-based spike sorting.

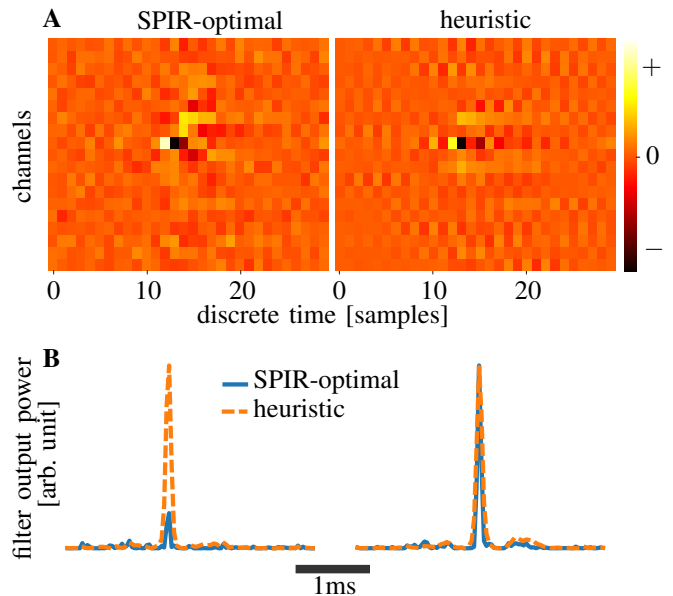


Fig. 2. **A** Color map representation of the filter coefficients in \mathbf{f} designed using the proposed method (left) and the heuristic method (right). **B** Squared filter output for a non-target spike (left) and target spike (right) using both the SPIR-optimal (solid blue) and heuristic (dashed orange) filter.

To quantify the threshold-based spike sorting performance of both filter design methods, the sensitivity and precision are calculated. The sensitivity (sens) and precision (prec) are defined as follows:

$$\text{sens} = \frac{\text{TP}}{\text{P}} \quad \text{and} \quad \text{prec} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad (6)$$

where the ground truth labels can be used to determine the true positives (TP), false positives (FP) and all ground truth occurrences (P). These performance metrics are defined for a fixed optimal threshold Thr , which is determined making use of the the ground truth labels (again, in practice clustering information can be used). The optimal threshold is determined by swiping the threshold over the filter output and retaining the threshold that optimizes the sum of sensitivity and precision. In case the resulting precision is lower than 0.9 the threshold which maximizes the precision is retained instead.

TABLE I

THRESHOLD-BASED SPIKE SORTING SENSITIVITY (SENS) AND PRECISION (PREC) FOR BOTH THE PROPOSED SPIR-OPTIMAL FILTER DESIGN METHOD AND THE HEURISTIC FILTER DESIGN METHOD IN [12]. THE LAST COLUMN CONTAINS THE DATASET SPECIFIC VALUE FOR λ AS USED IN (2).

dataset	SPIR sens/prec	heuristic sens/prec	λ
1	1.0/0.982	1.0/1.0	10
2	0.955/0.901	0.921/0.911	100
3	0.990/0.990	0.912/0.263	10

The quantitative threshold-based spike sorting results can be found in Table I. The in-vivo recorded datasets 1 and 2 on which the heuristic method performed fine, were used to compare the performance of the newly proposed method. For dataset 1 and 2 the performance using either the newly proposed method and the heuristic method is on par. For dataset 1 the difference in precision is due to two false positive detections. For dataset 2 the SPIR-optimal filter has a higher sensitivity, but a slightly lower precision. For dataset 3 there is a clear increase in both the precision and sensitivity using the newly proposed method.

IV. DISCUSSION AND CONCLUSION

A new filter design method is proposed for the use in threshold-based spike sorting. The filters are (local) minima of the proposed SPIR objective functions. It was shown that the filter design method outperforms a heuristic design method proposed earlier. Although the actual spike sorting is fast and computationally lightweight, the filter design can be computationally heavy. Other numerical solvers might be considered to speed up the filter design process. Of potential interest are stochastic gradient descent based methods, which take into account only a random subset of samples at every iteration step. Another possible improvement, which might potentially speed up the filter design process, is the use of another informed initialisation point, e.g., based on the heuristic method. Since the objective function is non-convex, it is possible that a sub-optimal local minima is given as a solution. Therefore, it is important to always check the threshold-based spike sorting performance on the training set. For all experimental data used in this work, the filter solutions were practically useable.

REFERENCES

- [1] M. S. Lewicki, "A review of methods for spike sorting: the detection and classification of neural action potentials," *Network: Computation in Neural Systems*, vol. 9, no. 4, pp. R53–R78, 1998.
- [2] S. Gibson, J. W. Judy, and D. Marković, "Spike sorting: The first step in decoding the brain," *IEEE Signal processing magazine*, vol. 29, no. 1, pp. 124–143, 2012.
- [3] S. N. Kadir, D. F. Goodman, and K. D. Harris, "High-dimensional cluster analysis with the masked EM algorithm," *Neural computation*, vol. 26, no. 11, pp. 2379–2394, 2014.
- [4] C. Rossant, S. N. Kadir, D. F. Goodman, J. Schulman, M. L. Hunter, A. B. Saleem, A. Grosmark, M. Belluscio, G. H. Denfield, A. S. Ecker, *et al.*, "Spike sorting for large, dense electrode arrays," *Nature neuroscience*, vol. 19, no. 4, p. 634, 2016.
- [5] 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, p. 232, 2017.
- [6] J. E. Chung, J. F. Magland, A. H. Barnett, V. M. Tolosa, A. C. Tooker, K. Y. Lee, K. G. Shah, S. H. Felix, L. M. Frank, and L. F. Greengard, "A fully automated approach to spike sorting," *Neuron*, vol. 95, no. 6, pp. 1381–1394, 2017.
- [7] P. Yger, G. L. Spampinato, E. Esposito, B. Lefebvre, S. Deny, C. Gardella, M. Stimberg, F. Jetter, G. Zeck, S. Picaud, *et al.*, "A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo," *ELife*, vol. 7, p. e34518, 2018.
- [8] O. Marre, D. Amodè, N. Deshmukh, K. Sadeghi, F. Soo, T. E. Holy, and M. J. Berry, "Mapping a complete neural population in the retina," *Journal of Neuroscience*, vol. 32, no. 43, pp. 14 859–14 873, 2012.
- [9] W. M. Roberts and D. K. Hartline, "Separation of multi-unit nerve impulse trains by a multi-channel linear filter algorithm," *Brain research*, vol. 94, no. 1, pp. 141–149, 1975.
- [10] R. Vollgraf and K. Obermayer, "Improved optimal linear filters for the discrimination of multichannel waveform templates for spike-sorting applications," *IEEE Signal Processing Letters*, vol. 13, no. 3, pp. 121–124, 2006.
- [11] F. Franke, M. Natora, C. Boucsein, M. H. Munk, and K. Obermayer, "An online spike detection and spike classification algorithm capable of instantaneous resolution of overlapping spikes," *Journal of computational neuroscience*, vol. 29, no. 1-2, pp. 127–148, 2010.
- [12] J. Wouters, F. Kloosterman, and A. Bertrand, "Towards online spike sorting for high-density neural probes using discriminative template matching with suppression of interfering spikes," *Journal of neural engineering*, 2018.
- [13] R. D. Flint, Z. A. Wright, M. R. Scheid, and M. W. Slutzky, "Long term, stable brain machine interface performance using local field potentials and multiunit spikes," *Journal of neural engineering*, vol. 10, no. 5, p. 056005, 2013.
- [14] A. Mendez, A. Belghith, and M. Sawan, "A DSP for sensing the bladder volume through afferent neural pathways," *IEEE transactions on biomedical circuits and systems*, vol. 8, no. 4, pp. 552–564, 2014.
- [15] Y. Yang, S. Boling, and A. J. Mason, "A hardware-efficient scalable spike sorting neural signal processor module for implantable high-channel-count brain machine interfaces," *IEEE transactions on biomedical circuits and systems*, vol. 11, no. 4, pp. 743–754, 2017.
- [16] J. E. Downey, N. Schwed, S. M. Chase, A. B. Schwartz, and J. L. Collinger, "Intracortical recording stability in human brain–computer interface users," *Journal of neural engineering*, vol. 15, no. 4, p. 046016, 2018.
- [17] F. Franke, D. Jäckel, J. Dragas, J. Müller, M. Radivojevic, D. Bakkum, and A. Hierlemann, "High-density microelectrode array recordings and real-time spike sorting for closed-loop experiments: an emerging technology to study neural plasticity," *Frontiers in neural circuits*, vol. 6, p. 105, 2012.
- [18] D. Kraft, "A software package for sequential quadratic programming," *Forschungsbericht- Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt*, 1988.
- [19] J. P. Neto, G. Lopes, J. Frazão, J. Nogueira, P. Lacerda, P. Baião, A. Aarts, A. Andrei, S. Musa, E. Fortunato, *et al.*, "Validating silicon polytrodes with paired juxtacellular recordings: method and dataset," *Journal of neurophysiology*, vol. 116, no. 2, pp. 892–903, 2016.