

A data-driven regularization approach for template matching in spike sorting with high-density neural probes

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

Abstract—Spike sorting is the process of assigning neural spikes in an extracellular brain recording to their putative neurons. Optimal pre-whitened template matching filters that are used in spike sorting typically suffer from ill-conditioning. In this paper, we investigate the origin of this ill-conditioning and the way in which it influences the resulting filters. Two data-driven subspace regularization approaches are proposed, and those are shown to outperform a regularization approach used in recent literature. The comparison of the methods is based on ground truth data that are recorded in-vivo.

I. INTRODUCTION

Extracellular microelectrodes are often used for studying brain dynamics at the level of individual neurons. Such microelectrodes can measure, among other signals, the extracellular potential changes caused by action potentials, also known as spikes, from nearby neurons. Typically, multiple electrodes are used for a given experiment and their mutual spatial organisation is matched to the targeted brain region. A popular electrode configuration used in hippocampal studies is the wire tetrode [1], reflecting the region’s shallow cell layer. When studying cortical areas, probes [2] are more popular because they allow to measure along the entire cortical multilayer cell structure.

Single electrodes tend to pick up spikes from multiple neurons, while multiple electrodes pick up the spikes from a specific neuron. On top of that, multiple neurons can be simultaneously active. To study the dynamics of the *individual* neuron, one has to solve this neuronal cocktail party problem, i.e., unmixing the recorded spiking activity of the individual neurons that are embedded in the recording. A multitude of algorithmic pipelines have been proposed to achieve such unmixing. These pipelines are commonly referred to as spike sorting [3] algorithms.

A typical spike sorting algorithm consists of the following steps [4]: after *preprocessing* the recorded data a *spike detection* is performed, then *features* are extracted from those detected spikes, a *clustering* algorithm is applied to the extracted features, and finally a *template matching* phase

is conducted for classifying the spikes according to their putative neurons.

Template matching is used here for classification [5] because it has the ability to resolve overlapping spikes. Template matching is also interesting for real-time spike sorting, e.g., for a neural prosthesis based on single-cell activity [6]. Although template matching is the final step in the algorithmic pipeline, the template matching filters are applied to the preprocessed data directly. As such they are suitable for streaming data applications after an initial training phase.

In this paper we will study a failure mechanism in the template matching filter design, namely that of the ill-conditioned noise covariance matrix. Most spike sorting algorithms propose pragmatic solutions that seem to ignore the source of the ill-conditioning. The template matching filters obtained through such pragmatic approaches are suboptimal. We will propose a meaningful regularization approach, which is in line with the underlying failure mechanism.

The remainder of this paper is organized as follows: in Section II template matching is formally presented and the ill-conditioning failure mechanism is studied. Section III proposes two data-driven regularization approaches that can overcome the ill-conditioning. In Section IV the proposed methods are validated. Finally Section V concludes this paper with a brief discussion.

II. TEMPLATE MATCHING

Consider $\mathbf{x}_i[k] \in \mathbb{R}^N$ to be a selection of N channels of the high-pass filtered extracellular recording sampled at discrete time k , where typically only the channels are included that are relevant to isolate the spikes of a specific target neuron i . To represent the spatio-temporal L -tap delay line at time k the following stacked vector $\bar{\mathbf{x}}_i[k] \in \mathbb{R}^{NL}$ is constructed:

$$\bar{\mathbf{x}}_i[k] = \begin{bmatrix} \mathbf{x}_i[k]^T & \mathbf{x}_i[k-1]^T & \dots & \mathbf{x}_i[k-L+1]^T \end{bmatrix}^T, \quad (1)$$

where L should be at least as long as the spike length of the target neuron.

In this paper we focus only on the template matching filter design for spike sorting. Therefore, we assume that for every sortable neuron i we have some example spike times \mathcal{T}_i available. In practice these example spike times are obtained from the spike sorting clustering phase.

Without loss of generality, in the remainder of this paper we will focus on a single neuron, hence the neuron-related subscript is dropped. The spike template $\boldsymbol{\tau}$ for the neuron of interest can be calculated as follows:

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 the Research Foundation Flanders (FWO) project FWO G0D7516N. This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No 802895). 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

$$\tau = \text{median}_{l \in \mathcal{T}} \bar{x}[l], \quad (2)$$

where the median is applied in an element-wise fashion. The median operator was chosen because it provides robustness against outliers in the training data. The classical optimal template matching filter [7], [8] that can be used for the classification of spikes is the solution to the following optimization problem:

$$\max_{\mathbf{f}} \frac{\mathbf{f}^T \boldsymbol{\tau} \boldsymbol{\tau}^T \mathbf{f}}{\mathbf{f}^T \mathbf{R}_{\text{nn}} \mathbf{f}}, \quad (3)$$

with $\mathbf{R}_{\text{nn}} \in \mathbb{R}^{NL \times NL}$ the spatio-temporal noise-only covariance matrix, i.e., $\sum_{k \in \mathcal{N}} \bar{\mathbf{x}}[k] \bar{\mathbf{x}}[k]^T$ with \mathcal{N} containing all the sample times in which there is no spiking activity. The closed-form solution to (3) is given by [7], [8]:

$$\mathbf{f} = \mathbf{R}_{\text{nn}}^{-1} \boldsymbol{\tau}. \quad (4)$$

The actual spike classification is based on the filter output $\mathbf{f}^T \bar{\mathbf{x}}[k]$. This corresponds to a matched filter which also performs an implicit pre-whitening, thereby maximizing the signal-to-noise ratio at the target neuron spikes. Within the scope of this paper we will classify the spikes based on thresholding the squared template matching filter output. This approach is most in line with the objective function (3), and it gives a good indication about the quality of the obtained filters for use in spike sorting. However, in practice more advanced classification schemes based on, e.g., multiple thresholds and iterative filtering exist [9].

The filter given by (4) only exists if the spatio-temporal covariance matrix is invertible. The matrix is invertible if and only if its null space is trivial, i.e., $\text{Null}(\mathbf{R}_{\text{nn}}) = \{\mathbf{0}\}$ [10]. If the covariance matrix is non-invertible, then it has a null space different from the zero-vector in which case, the denominator in (3) can be made arbitrarily small by choosing \mathbf{f} “close” to that null space. Assuming that the template is not perfectly orthogonal to that null space, the objective function can be made arbitrarily large, without taking into account the template itself.

In practice most covariance matrices will be numerically invertible, if sufficient data are used for their estimation. This can be explained by the fact that the signals under investigation originate from experimental recordings. The recording electronics will typically introduce low power white noise which will lead to matrices that are full rank, therefore invertible. However, the power dominant spatio-temporal neural data typically have a low rank structure, especially when using modern high-density probes with spatial oversampling. Under the low power white noise assumptions, those matrices will often be ill-conditioned, i.e., they contain an eigenspace with corresponding eigenvalues that are several orders of magnitude smaller than the eigenvalues related to the neural data. As such the denominator in (3) can be made very small by pushing \mathbf{f} towards this eigenspace (not arbitrarily small as was the case for the non-invertible case), thereby maximizing the objective function. Here, again there is a high probability that the significance of the filter response to the template is not considered. Such a filter cannot emphasize the spiking

activity of its target neuron, hence it is useless in a spike sorting context.

III. REGULARIZATION

Fig. 1 shows the eigenvalues of \mathbf{R}_{nn} , i.e., the signal power across the different principal components in a principal component analysis (PCA), from two of the spatio-temporally expanded in-vivo recordings used in Section IV. Most signal power is restricted to a few PCA components, indicating that the data are ill-conditioned. As explained in the previous section, applying (4) without proper regularization will lead to filters that are practically useless for spike sorting purposes. Another need for regularization arises when only a limited number of samples are used for estimating the covariance matrix. In the remainder of this work we will expand on the topic of regularization within the context of template matching-based spike sorting.




Fig. 1. Power along every principal component (red) for the spatio-temporally expanded recording data of neuron 1 and neuron 2. The cumulative normalized power (dashed blue) is also shown and the β -fraction point (see Section III-B) is indicated by the black arrow.

In recent spike sorting literature [9], the ill-conditioning is practically tackled by only considering the spatial covariance for the design of the spatio-temporal matched filter. On top of that a diagonal loading of the covariance matrix is performed to further improve the condition. This approach can be summarized by the following optimization problem which has the same form as (3):

$$\max_{\mathbf{f}} \frac{\mathbf{f}^T \boldsymbol{\tau} \boldsymbol{\tau}^T \mathbf{f}}{\mathbf{f}^T \left(\begin{bmatrix} \mathbf{R}_{\text{nn}} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{\text{nn}} & & \vdots \\ \vdots & & \ddots & \mathbf{0} \\ \mathbf{0} & \dots & \mathbf{0} & \mathbf{R}_{\text{nn}} \end{bmatrix} + c\mathbf{I} \right) \mathbf{f}}, \quad (5)$$

where $\mathbf{R}_{\text{nn}} \in \mathbb{R}^{N \times N}$ is the spatial noise-only covariance matrix (based on $\mathbf{x}[k]$ instead of $\bar{\mathbf{x}}[k]$) and c is the diagonal loading constant which is a fixed parameter and data independent. By using this practically regularized filter design formulation the existing temporal covariance in the signal is ignored. The filters that result from this approach are suboptimal detectors from a spatio-temporal point of view. We will use this *spatial-only* regularization approach for benchmarking our proposed methods.

A. Use the signal+noise covariance matrix

From a matched filtering point of view there is no reason to exclude spiking activity samples from the noise covariance matrix estimate. One can show that including the template

from the neuron of interest in the noise covariance does not change the optimal filter [8]:

$$\arg \max_{\mathbf{f}} \frac{\mathbf{f}^T \boldsymbol{\tau} \boldsymbol{\tau}^T \mathbf{f}}{\mathbf{f}^T (\mathbf{R}_{\text{nn}} + \alpha \boldsymbol{\tau} \boldsymbol{\tau}^T) \mathbf{f}} = \arg \max_{\mathbf{f}} \frac{\mathbf{f}^T \boldsymbol{\tau} \boldsymbol{\tau}^T \mathbf{f}}{\mathbf{f}^T \mathbf{R}_{\text{nn}} \mathbf{f}}, \quad (6)$$

where α is an arbitrary weight factor. Therefore, we propose to use the signal+noise covariance matrix instead of the noise-only covariance matrix, i.e., we replace \mathbf{R}_{nn} with $\mathbf{R}_{\bar{\mathbf{x}\mathbf{x}}} = \sum_k \bar{\mathbf{x}}[k] \bar{\mathbf{x}}[k]^T$. This effectively increases the number of available samples used for the estimation of the covariance matrix, thereby making it invertible without regularization. This choice of covariance matrix also better describes the true underlying noise structure, which does include spikes from other (interfering) neurons¹. The optimal filter design objective function we will consider is then:

$$\max_{\mathbf{f}} \frac{\mathbf{f}^T \boldsymbol{\tau} \boldsymbol{\tau}^T \mathbf{f}}{\mathbf{f}^T \mathbf{R}_{\bar{\mathbf{x}\mathbf{x}}} \mathbf{f}}. \quad (7)$$

The objective function in (7) still suffers from ill-conditioning due to the low rank structure of the neural signal. To solve for this we propose a regularization approach that involves the design of the filter in a lower dimensional subspace.

B. Filter design in PCA subspace

We define a PCA subspace that accounts for a β -fraction of the total energy in the recorded (preprocessed) signal. A typical value would be $\beta = 0.95$. The subspace is readily available through the eigendecomposition of the covariance matrix:

$$\mathbf{R}_{\bar{\mathbf{x}\mathbf{x}}} = \mathbf{U} \boldsymbol{\Lambda} \mathbf{U}^T, \quad (8)$$

where \mathbf{U} contains the eigenvectors in its columns. Without loss of generality, we assume that the diagonal matrix $\boldsymbol{\Lambda}$ has all of its diagonal elements λ_i arranged in a descending order. The subspace is spanned by the columns of \mathbf{U}_β :

$$\mathbf{U}_\beta = [\mathbf{u}_1 \dots \mathbf{u}_k] \text{ with } k = \max \left\{ l \mid \frac{\sum_{i=1}^l \lambda_i}{\text{tr}(\boldsymbol{\Lambda})} < \beta \right\}, \quad (9)$$

where \mathbf{u}_j denotes the j^{th} column of \mathbf{U} . Note that the dimensionality of the subspace is defined by the underlying data, hence the term data-driven regularization. The template matching filter \mathbf{f}_β operating in this reduced subspace is the solution to the following optimization problem:

$$\max_{\mathbf{f}_\beta} \frac{\mathbf{f}_\beta^T \mathbf{U}_\beta^T \boldsymbol{\tau} \boldsymbol{\tau}^T \mathbf{U}_\beta \mathbf{f}_\beta}{\mathbf{f}_\beta^T \mathbf{U}_\beta^T \mathbf{R}_{\bar{\mathbf{x}\mathbf{x}}} \mathbf{U}_\beta \mathbf{f}_\beta}. \quad (10)$$

Transforming this optimal filter back to the original space results in the first subspace regularized filter:

$$\mathbf{f}_1 = \mathbf{U}_\beta (\mathbf{U}_\beta^T \mathbf{R}_{\bar{\mathbf{x}\mathbf{x}}} \mathbf{U}_\beta)^{-1} \mathbf{U}_\beta^T \boldsymbol{\tau}. \quad (11)$$

¹Recently, it has been shown that a more discriminative filter is obtained if \mathbf{R}_{nn} is replaced with an interference covariance matrix that is *only* built from segments containing interfering spikes [8]. In this case, the matrix may again become non-invertible, which is resolved by using the PCA preprocessing as explained in Sections III-B and III-C.

C. Filter design in template & PCA subspace

The above method assumes that the template is well-approximated in the filter design subspace. For $\beta = 0.95$ it is very likely that the template is indeed well-approximated in the PCA subspace, but there is no absolute guarantee. To ensure that the template is well-approximated within the subspace we propose to use a subspace spanned by the template itself, which is then extended with orthogonal basis vectors until a β -fraction of the variance is captured. Consider the template data direction:

$$\mathbf{u}_\tau = \boldsymbol{\tau} / \|\boldsymbol{\tau}\|_2. \quad (12)$$

The power that is associated with this direction is:

$$\lambda_\tau = \mathbf{u}_\tau^T \mathbf{R}_{\bar{\mathbf{x}\mathbf{x}}} \mathbf{u}_\tau. \quad (13)$$

Let $\text{Proj}_{\mathbf{u}_\tau} \bar{\mathbf{x}}[k] = \mathbf{u}_\tau \mathbf{u}_\tau^T \bar{\mathbf{x}}[k]$ be the projection of $\bar{\mathbf{x}}[k]$ onto the direction \mathbf{u}_τ . We define the new data vector $\bar{\mathbf{x}}_{-\tau}[k] = \bar{\mathbf{x}}[k] - \text{Proj}_{\mathbf{u}_\tau} \bar{\mathbf{x}}[k] = (\mathbf{I} - \mathbf{u}_\tau \mathbf{u}_\tau^T) \bar{\mathbf{x}}[k]$ which is by definition orthogonal to \mathbf{u}_τ , i.e., it does not contain any variance in the template direction. The covariance matrix associated with $\bar{\mathbf{x}}_{-\tau}[k]$ is then given by:

$$\mathbf{R}_{-\tau} = (\mathbf{I} - \mathbf{u}_\tau \mathbf{u}_\tau^T) \mathbf{R}_{\bar{\mathbf{x}\mathbf{x}}} (\mathbf{I} - \mathbf{u}_\tau \mathbf{u}_\tau^T)^T. \quad (14)$$

From the eigendecomposition (again assuming that the diagonal matrix $\boldsymbol{\Lambda}_{-\tau}$ has all of its diagonal elements $\lambda_{-\tau i}$ arranged in a descending order) of this covariance matrix:

$$\mathbf{R}_{-\tau} = \mathbf{U}_{-\tau} \boldsymbol{\Lambda}_{-\tau} \mathbf{U}_{-\tau}^T, \quad (15)$$

we can extract an orthogonal basis for the subspace that captures the most variance, while it is also guaranteed to contain the template:

$$\mathbf{T}_\beta = [\mathbf{u}_\tau \quad \mathbf{u}_{-\tau 1} \dots \mathbf{u}_{-\tau j}]. \quad (16)$$

The subspace dimensionality of $j + 1$ is then defined in a data-driven fashion by the following expression:

$$j = \max \left\{ l \mid \frac{\lambda_\tau + \sum_{i=1}^l \lambda_{-\tau i}}{\text{tr}(\boldsymbol{\Lambda})} < \beta \right\}. \quad (17)$$

The second subspace regularized filter is then given by:

$$\mathbf{f}_2 = \mathbf{T}_\beta (\mathbf{T}_\beta^T \mathbf{R}_{\bar{\mathbf{x}\mathbf{x}}} \mathbf{T}_\beta)^{-1} \mathbf{T}_\beta^T \boldsymbol{\tau}. \quad (18)$$

IV. EXPERIMENTS

We validate the two proposed regularization methods on 28 ground truth neurons. These ground truth neurons are sourced from publicly available paired recordings [11] and hybrid recordings [12]. The preprocessing consists only of a high-pass filtering with a cutoff frequency $f_c = 300$ Hz (note that no explicit prewhitening is performed here). For every neuron a spatial region with size N was chosen such that all electrodes within a 100 μm radius from the electrode with the maximum template signal deflection for that neuron are included in that spatial region. The discrete temporal spike window length is set to $L = 30$, which is equivalent to a 1 ms window. From every recording the first 3 min of data are used for the analysis. The training data \mathcal{T} include the ground truth spike times from the first 90 s. Validation is done on the remaining 90 s.

The filters that are used for benchmarking our two proposed regularized filter design methods are the solutions to the optimization problem given by (5). The free parameter was chosen to be $c = 10^{-18}$ as given by [9]. For our data this parameter was of very little use, because the spatial-only regularization was always dominant over the diagonal loading. As suggested before $\beta = 0.95$ was chosen for both our proposed regularized filter design methods given by (11) and (18), such that the low rank approximated recording does not deviate too much from the original recording.




Fig. 2. The F_1 score is shown as a measure of spike sorting performance for every ground truth neuron used. Three different regularization methods for calculating the template matching filters are compared: template & PCA subspace regularization (green), PCA subspace regularization (red), and spatial-only regularization (blue).

The three different regularized filter design methods are validated through their spike sorting capability, that is obtained from thresholding their respective squared filter outputs. For every combination of neuron and regularized filter design, the spike sorting capability is summarized by the maximal F_1 score across all thresholds. This score reflects the ability of a filter to extract only the activity of the neuron of interest and is depicted in Fig. 2. These results show an increased spike sorting performance in favour of both the proposed methods when compared to the spatial-only regularization. The subspace regularization that includes the template seems to slightly outperform the subspace regularization without explicit use of the template.

Using a Friedman test with $\alpha = 0.05$, there is a statistically significant difference in F_1 score between the different studied methods, $X^2(2) = 24.768$, $p \ll 0.05$. Post hoc

analysis with Wilcoxon signed-rank tests was conducted with a Bonferroni correction applied, resulting in a significance level set at $p < \frac{\alpha}{m} = 0.017$, with $m = 3$ indicating the number of comparisons. There are statistically significant differences in F_1 score between the spatial-only regularization and the PCA subspace regularization ($Z = -4.0145$, $p \ll 0.017$), and between the spatial-only and the template & PCA subspace regularization ($Z = -4.0145$, $p \ll 0.017$). However, there is no significant difference between the PCA subspace regularization and the template & PCA subspace regularization ($Z = -2.3541$, $p = 0.018$), despite a perceived superiority of the template & PCA subspace method. However, this may be due to the overly conservative choice for Bonferroni correction, as the p-value almost reaches the significance level of $\alpha = 0.017$.

V. DISCUSSION

A data-driven regularization approach for template matching filter design in spike sorting is proposed. The approach consists of designing the actual filter in a subspace for which the dimensionality is derived from the data. Subspace regularization can be interpreted as data denoising combined with a diagonal loading. This approach is shown to significantly outperform the spatial-only regularization approach in our simplified spike sorting framework based on a single thresholding. More spike sorting regularization approaches exist in literature, but all of them are based on prior fixed parameters, e.g., a fixed diagonal loading parameter, or a fixed size subspace. The integration of our new approach in a fully-featured spike sorting framework might lead to a better spike discrimination and therefore less time spent on manual curation.

REFERENCES

- [1] D. P. Nguyen *et al.*, “Micro-drive array for chronic in vivo recording: tetrode assembly,” *JoVE*, no. 26, 2009.
- [2] F. Michon *et al.*, “Integration of silicon-based neural probes and micro-drive arrays for chronic recording of large populations of neurons in behaving animals,” *J Neural Eng*, vol. 13, no. 4, p. 046018, 2016.
- [3] 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.
- [4] S. Gibson *et al.*, “Spike sorting: The first step in decoding the brain,” *IEEE Signal processing magazine*, vol. 29, no. 1, pp. 124–143, 2012.
- [5] O. Marre *et al.*, “Mapping a complete neural population in the retina,” *Journal of Neuroscience*, vol. 32, no. 43, pp. 14859–14873, 2012.
- [6] Y. Yang *et al.*, “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.
- [7] F. Franke *et al.*, “Bayes optimal template matching for spike sorting—combining fisher discriminant analysis with optimal filtering,” *Journal of computational neuroscience*, vol. 38, no. 3, pp. 439–459, 2015.
- [8] J. Wouters *et al.*, “Towards online spike sorting for high-density neural probes using discriminative template matching with suppression of interfering spikes,” *J Neural Eng*, vol. 15, no. 5, p. 056005, 2018.
- [9] P. Yger *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.
- [10] D. C. Lay, *Linear algebra and its applications*. Addison Wesley Publishing Company, 1994.
- [11] J. P. Neto *et al.*, “Validating silicon polytrodes with paired juxtacellular recordings: method and dataset,” *Journal of neurophysiology*, vol. 116, no. 2, pp. 892–903, 2016.
- [12] C. Rossant *et al.*, “Spike sorting for large, dense electrode arrays,” *Nature neuroscience*, vol. 19, no. 4, p. 634, 2016.