Intrinsic Dimensionality of Extracellular Action Potentials

Kathryn Scannell[1], Agnieszka F. Szymanska[2], and Zoran Nenadic[2,3]

Abstract— Linear approaches to low-dimensional feature extraction may not be appropriate when statistical data are generated by a nonlinear interaction of parameters. Equally inadequate are linear methods for determining the dimension of the feature space. This article estimates the intrinsic dimension of extracellular action potentials (EAPs), which can be viewed as the minimum number of nonlinearly interacting parameters sufficient to describe the data. When combined with nonlinear feature extraction methods, this information may lead to a more faithful, low-dimensional EAP representation. These points are demonstrated using EAPs recorded experimentally by a multi-sensor electrode.

I. INTRODUCTION

Extracellular recording of action potentials represents a core activity in experimental neuroscience [1]. It enables behavior and electrophysiology to be linked at a single-neuron resolution. Before extracellular action potentials (EAPs) can be interpreted, they must be detected in recordings that contain biological and non-biological noise [2]. As a single extracellular electrode often senses signals from multiple neurons, EAPs must also be classified according to their neuron of origin [3]. This is especially important in multi-sensor electrodes [4], capable of simultaneous recording from a dozen or more neurons [5].

Classification of EAPs is typically preceded by extraction of low-dimensional data features, which facilitates their more accurate statistical description and classification [6]. A classical approach to this problem assumes that n-dimensional data are confined to a d-dimensional subspace ($d \ll n$) which can be recovered by a suitable linear transformation, followed by the truncation of $n - d$ noise features. In the context of EAP classification, the best known representatives of this approach are principal component analysis (PCA) [7], [8] and the wavelet transform (WT) [9], [10]. The dimension of the signal subspace d is typically determined by discarding the principal eigenvalues (wavelet coefficients) which fall below a suitably chosen threshold.

When n-dimensional statistical data depend nonlinearly on d parameters, the above linear approaches may not be appropriate as the data lie on a d-dimensional hypersurface (manifold). This observation has inspired the proliferation of nonlinear feature extraction techniques such as locally linear embedding [11], Isomap [12], diffusion maps [13], and many more. Unlike linear approaches that utilize global data properties, nonlinear methods extract features by preserving the local neighborhood of each data point. The approach to estimating the number of underlying parameters, d, must then be modified accordingly.

This article presents an algorithm to estimate *intrinsic dimensionality* of EAPs. This step may be viewed as determining the dimension of a possibly nonlinear feature "space" in which EAP classification is to be performed. When combined with nonlinear feature extraction techniques and classifiers that exploit local data properties (e.g. a nearest neighbor classifier), this approach may hold advantage over traditional linear techniques. These points will be elaborated upon below, and the algorithm will be illustrated on EAPs experimentally recorded by a 4-sensor electrode (tetrode).

II. INTRINSIC DIMENSIONALITY

A. Motivation

To motivate the concept of an intrinsic dimension, let us consider a model of an EAP given by:

$$s(t) = e^{-\frac{t}{\tau}} \sin\left(\frac{2\pi t}{T}\right) + \eta(t) \tag{1}$$

where τ is a decay time constant, T is a period of oscillations, and η is temporally correlated noise. Note that s depends nonlinearly on τ and T. A collection of EAPs can be generated by drawing τ and T from a uniform distribution with a mean of 0.5 ms and 1.5 ms, respectively, and a standard deviation of 10% of the mean. The standard deviation of η is set at 2.5% of the peak value of s. Fig. 1 shows EAPs simulated in this manner, as well as their intrinsic dimension estimated by the algorithm (details in Section II-B). Since the waveforms are parametrized by τ and T, the intrinsic dimension, $d_I = 2$, is correctly estimated when 13 or more nearest neighbors are used.
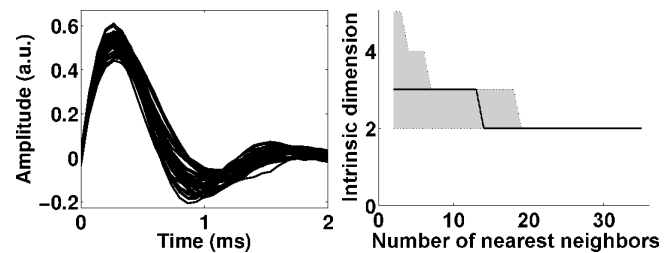

Fig. 1. (Left) 40 EAPs simulated according to (1) with a sampling rate of 15 kHz. (Right) Estimated intrinsic dimension as a function of the number of nearest neighbors. The solid line represent the median over 10 simulations, rounded to the nearest whole number. The shaded region marks the range of intrinsic dimensions across 10 simulations.

B. Intrinsic Dimension Calculation

Our methodology follows closely the original algorithm of Pettis et al. [14]. The main idea behind their approach is that the number of data points in a d-dimensional hyperball of radius R is proportional to R^d. A similar idea has led to the concept of a correlation dimension in nonlinear dynamics [15].

Assume that $s_j \in \mathbb{R}^n$ ($j = 1, 2, \cdots, N$) is a collection of EAPs, where n is the dimension and N is the number of EAPs. In the first step, the Euclidean distances between s_j and its K nearest neighbors are found, where $K \geq 2$ is a conveniently chosen number of neighbors. Denote these distances by r_{k,s_j} ($k = 1, 2, \cdots, K$; $j = 1, 2, \cdots, N$). By linking the probability density function of $r_{k,s}$ and the number of neighbors of s in a d-dimensional space, Pettis et al. arrive at the following relationship [14]:

$$\log(\bar{r}_k G_{k,d}) = \frac{1}{d}\log(k) + \log(C_n), \quad \forall k = 1, 2, \cdots, K \quad (2)$$

where

$$G_{k,d} = k^{\frac{1}{d}} \frac{\Gamma(k)}{\Gamma\left(k + \frac{1}{d}\right)}, \quad (3)$$

$\bar{r}_k = \frac{1}{N}\sum_{j=1}^{N} r_{k,s_j}$, and C_n is a term that is independent of k. Note that $\Gamma(\cdot)$ in (3) denotes the gamma function. The algorithm proceeds by fitting a regression line in (2), with $\log(1), \log(2), \cdots, \log(K)$ as independent variables, and $\log(\bar{r}_1 G_{1,d}), \log(\bar{r}_2 G_{2,d}), \cdots, \log(\bar{r}_K G_{K,d})$ as dependent ones. The intrinsic dimension, d_I, is then found as the inverse of the regression slope. Note that the intercept term, $\log(C_n)$, can also be estimated in this process, although it is not used in d_I calculations. Also note that since $G_{k,d}$ depends on d, the regression slope must be found iteratively. The initial estimate, $\hat{d}_0$, can be found from (2) by assuming $\log(G_{k,d}) = 0$ and solving for d as done in [14], or as $\hat{d}_0 = \frac{1}{K-1}\sum_{k=1}^{K-1} \hat{d}_{0,k}$, where based on [16], we have:

$$\hat{d}_{0,k} \approx \frac{1}{k} \frac{\bar{r}_k}{\bar{r}_{k+1} - \bar{r}_k}, \quad \forall k = 1, 2, \cdots, K-1$$

The latter approach was taken in the present study, although the choice of $\hat{d}_0$ did not significantly affect the performance. The iteration process was terminated once $|\hat{d}_i - \hat{d}_{i-1}| < d_{\text{tol}}$ was satisfied, where d_{tol} is a suitably chosen tolerance (0.01 in the present study). The algorithm can be summarized as follows:

(a) **begin** set: $K \in \mathbb{N}$, $d_{\text{tol}} \in \mathbb{R}$, $i_{\max} \in \mathbb{N}$
(b) get EAPs: $s_1, s_2, \cdots, s_N \in \mathbb{R}^n$
(c) compute: r_{k,s_j}, $k = 1, 2, \cdots, K$, $j = 1, 2, \cdots, N$
(d) compute: $\bar{r}_k = \text{median}_{1 \leq j \leq N}\{r_{k,s_j}\}$, $\hat{d}_0$
(e) **initialize** $i \leftarrow 0$, $\Delta d = d_{\text{tol}}$, $\hat{d} = \hat{d}_0$
(f) **while** ($|\Delta d| \geq d_{\text{tol}}$ and $i \leq i_{\max}$)
(g)   compute: $\log(G_{k,\hat{d}})$, $\forall k$ [see (3)]
(h)   compute: $S = \text{regress}(\log(\bar{r}_k G_{k,\hat{d}}), \log(k))$ [see (2)]
(i)   update: $d_{\text{old}} = \hat{d}$, $\hat{d} = \frac{1}{S}$, $\Delta d = \hat{d} - d_{\text{old}}$, $i \leftarrow i + 1$
(j) **end**
(k) **return** $d_I(K) = \hat{d}$
(l) **end**

To increase the robustness of the algorithm, Pettis et al. [14] discarded the samples of r_{k,s_j} that were one standard deviation above the mean, and then re-estimated the mean $\bar{r}_k$. The problem with this approach is that the outlier threshold itself is susceptible to outliers, hence a robust estimate of the mean, $\bar{r}_k$, was obtained by taking the median [step (d)].

The algorithm is typically run for a range of nearest neighbors $K \in [2, K_{\max}]$ (e.g. Fig. 1). To obtain confidence intervals on these estimates, the algorithm may be run multiple times (with novel data), and $d_I(K)$ can be found by averaging. Since the intrinsic dimension is expected to be a whole number, the values of $d_I(K)$ should be rounded. For a sufficiently large number of samples, $d_I(K)$ tends to remain stable over a wide range of nearest neighbors (Fig. 1), and the overall intrinsic dimension can be taken as the most frequent value across K.

III. RESULTS

A. Experimental Data

Data consist of recordings from the antennal lobe of an adult locust [17]. The data were collected with a planar 4-sensor silicon electrode (tetrode) and their 20-s-long segment is publicly available from [18]. The signals were amplified, band-pass filtered (300-5,000 Hz), and sampled at 15 kHz.
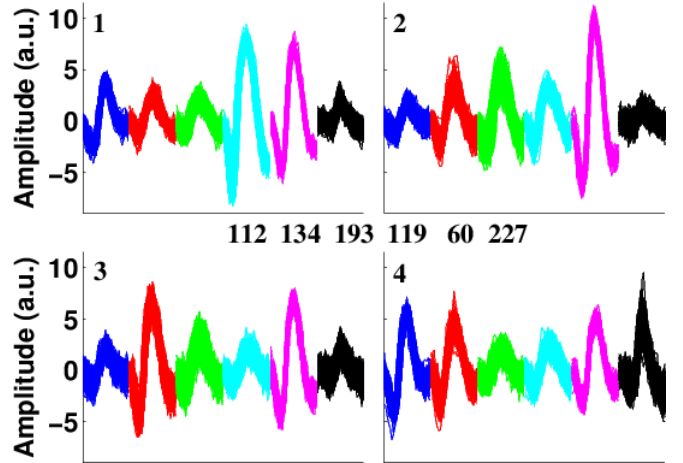

Fig. 2. Six EAP classes shown in different colors with 4 panels representing different tetrode sensors. Each trace is 2 ms long. The number of EAPs in each class is shown in the middle. Note that 185 outliers have been removed.

EAPs were detected in an unsupervised manner using the continuous WT method described in [2]. This resulted in detecting 1030 EAPs (result comparable to a supervised detection method in [19]). The EAPs were aligned [10], and 2 ms of data centered at the EAP peak were extracted. To facilitate further analysis, the EAPs were represented in the wavelet domain using bior1.3 wavelet class [2], [20], and classified using the top 15 wavelet coefficients. Our prior study [19] has revealed 6 EAP classes in these data, and so k-means clustering ($k = 6$) was performed in the 15-D wavelet space. The results of classification are shown in Fig. 2. Note that the classification results will *not* be used in intrinsic dimension calculation—their only role is to help with the interpretation of results.

B. Intrinsic Dimension of Extracellular Action Potentials

Upon outlier removal, the remaining 845 EAPs (Fig. 2) were reshaped into a vector form and used to estimate d_I as follows. For each $K \in [2, K_{\max}]$, the EAPs were randomly split into 10 groups (with 84–85 EAPs per group), and $d_I(K)$ was estimated as in Section II-B based on EAPs from 9 groups (one group withheld). Note that this limits the number of nearest neighbors to $K_{\max} < 760$. The above procedure was repeated each time withholding a different group, and the overall $d_I(K)$ was found as the median over these 10 values rounded to the nearest whole number. To remove randomization bias, the EAP group assignment was re-randomized for each K. The results are shown if Fig. 3, where $d_I(K)$ drops steadily with K, ultimately reaching a plateau, $d_I = 4$, in the range $K \in [298, 715]$. The histogram of $d_I(K)$ also shows that $d_I = 4$ is the most frequent value, and so it can be concluded that the intrinsic dimensionality of the set of EAPs shown in Fig. 2 is 4.
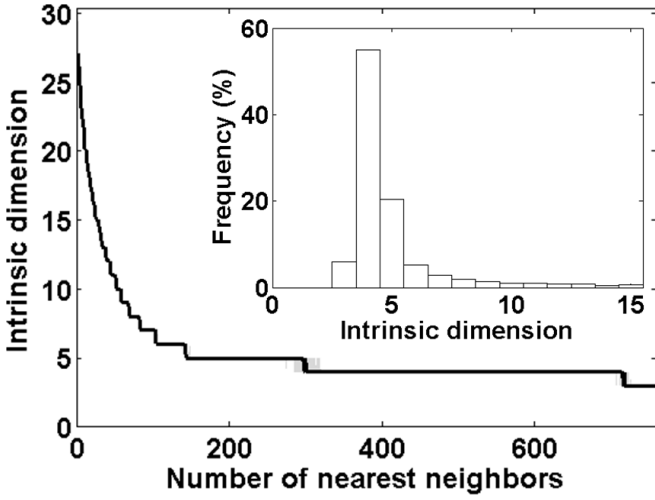

Fig. 3. Intrinsic dimension, d_I, as a function of the number of nearest neighbors $K \in [2, 759]$. Shaded regions represent confidence intervals estimated in the same way as in Fig. 1. (Inset) The histogram of d_I (truncated at $d_I = 15$), showing a dominant mode at $d_I = 4$.

The tight d_I confidence intervals (Fig. 3) suggest that d_I can be reliably estimated without subsampling and group randomization. This was confirmed by calculating d_I using all available data (single run), which resulted in a distribution of d_I nearly identical to the one in Fig. 3. Similarly, the results did not change when the calculations were repeated with 5%, 10%, 25% and 50% of EAPs (data not shown in the interest of space), as long as all 6 groups of EAPs were sufficiently represented in the sample. Therefore, the intrinsic dimensionality of this dataset is largely independent of the sample size.

It may be tempting to conclude that the intrinsic dimension of this dataset is somehow constrained by the number of sensors used. To show that this is not the case, we repeated the above analysis with various combinations of sensors (see Table I). It can be seen that 3-sensor combinations yielded similar d_I estimates, although the histograms of d_I (not

TABLE I
d_I FOR A DIFFERENT COMBINATION OF SENSORS.

sensors:	1,2,3	1,2,4	1,3,4	2,3,4	1	2	3	4
d_I:	4	4	5	4	5	3	8	4

shown) were much more broadly distributed, with $d_I = 5$ being the close second choice. An exception was the sensor combination (1,3,4) for which $d_I = 5$ (frequency: $\sim$40%), with $d_I = 4$ ($\sim$32%) as the second best candidate. When single-sensor data were used, the distribution of d_I was even broader, with the d_I values of 3, 4, and 5, often being nearly tied. A notable exception is the sensor 3 for which $d_I = 8$. This could be potentially explained by a higher noise level at the sensor 3, as established previously in [21] (see also Fig. 2). Namely, the presence of a strong noise in data may cause the algorithm to treat noise variances and correlations as additional data parameters [14].

C. Feature Extraction and Classification

Since d_I represents the intrinsic or topological dimension of an EAP hypersurface, nonlinear feature extraction methods are preferred in order to "flatten" the manifold into a d_I-dimensional Euclidean subspace. While many techniques can be used to accomplish this goal, we chose the diffusion maps (DM) approach [13] because of its simplicity and flexibility. The method works by forming a graph connectivity matrix $G = G^T \in \mathbb{R}^{N \times N}$, whose (i,j)-th element is:

$$ g(i,j) = e^{-\frac{\|s_i - s_j\|^2}{2D^2}}, \quad \forall i \geq j = 1, 2, \cdots, N $$

and D is a parameter that controls the strength of local interactions. The connectivity matrix is then converted into the transition probability matrix, $P \in \mathbb{R}^{N \times N}$, of an underlying Markov chain by: $p(i,j) := g(i,j)/\sum_{j=1}^{N} g(i,j)$. Note that $\sum_j p(i,j) = 1$. The set of N EAPs can then be represented by a feature matrix $F \in \mathbb{R}^{d_I \times N}$. Its rows are diffusion coordinates defined by [13]: $f_i := \lambda_i \psi_i^T$ $(i = 1, 2, \cdots, d_I)$, where $1 = \lambda_0 > |\lambda_1| \geq |\lambda_2| \geq \cdots$ are the eigenvalues of P, and $\psi_i \in \mathbb{R}^N$ are the corresponding eigenvectors.

The representation of the 845 EAPs with $d_I = 4$ diffusion coordinates is shown in Fig. 4. The interaction parameter was $D = 3.5\,\sigma$, where σ is a robust (median-based) estimate of the standard deviation of $\|s_i - s_j\|$ ($\forall i \geq j$). Several projections, e.g. f_1–f_2, f_1–f_4, and f_2–f_4, reveal the underlying 6-class data structure. In contrast, PCA (see Fig. 5), which is a linear technique, reveals good clustering in f_1–f_2, and to some extent in f_2–f_4 projection. Another linear method, WT (not shown), produced multiple projections with 6 clusters, however the cluster overlaps were significant.

To assess the class separability, we performed the k nearest neighbor (kNN) classification in the respective 4-D subspaces, while using the classification results in Section III-A as the ground truth. For all 3 methods (DM, PCA, and WT), the classification was performed using a leave-one-out approach, i.e. a single EAP was classified based on the class membership of its kNN ($k = 1, 3, \cdots, 23$) according to the
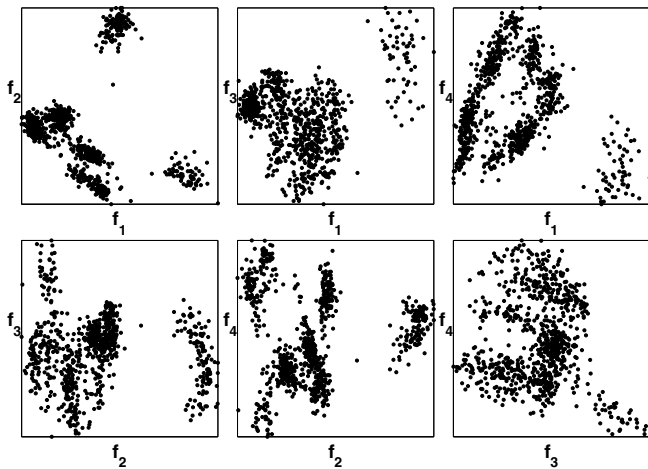
Fig. 4. Various 2-D projections of scatter plots showing the 845 EAPs. The features f_1, f_2, f_3, and f_4, were extracted using diffusion maps [13].
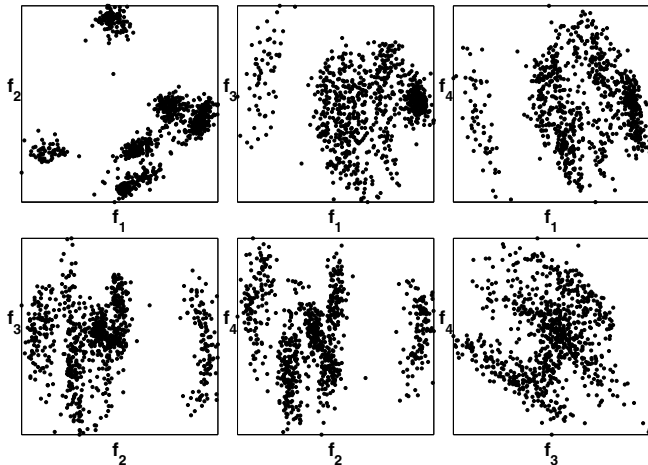

Fig. 5. Equivalent to Fig. 4, with f_1, f_2, f_3, and f_4, extracted by PCA.

majority rule. The predicted class label was then compared to the class label of the EAP in question, as determined in Section III-A. This procedure was repeated by looping over all EAPs, and the total number of misclassified EAPs was reported in Table II. In general, the classification results

TABLE II

COMPARISON OF CLASSIFICATION ERRORS FOR DM, PCA, AND WT.

k	1	3	5	7	9	11	13	15	17	19	21	23	mean
DM	3	3	1	1	1	1	1	3	3	4	4	4	2.42
PCA	5	4	3	4	3	4	4	4	3	4	4	4	3.83
WT	14	12	11	11	12	12	13	14	13	13	14	13	12.67

based on DM features were superior to those of PCA and WT over a range of k, reducing the error by $> 66\%$. Similar results were achieved with both $D = 4.0\,\sigma$ and $D = 4.5\,\sigma$.

IV. CONCLUSIONS

The algorithm to estimate the intrinsic dimension of statistical data was applied to EAPs recorded by a tetrode. It produced stable estimates of the intrinsic dimension, $d_I = 4$, across a wide range of samples, N, and nearest neighbors, K. Combined with the diffusion maps approach, this information led to a faithful representation of EAPs in a 4-D space, unambiguously revealing 6 distinct EAP clusters. When compared to its linear counterparts, such as PCA and WT, it also yielded a substantial reduction in classification errors, as judged by the kNN classifier. Finding a faithful low-dimensional data representation may also be useful in other neural signal processing applications. Our future work will be directed at analysis of additional EAP data as well as other types of electrophysiological signals.

REFERENCES

[1] D.R. Humphrey. *Electrophysiological Techniques*. Society for Neuroscience, Atlanta, 1979.
[2] Z. Nenadic and J.W. Burdick. Spike detection using the continuous wavelet transform. *IEEE Trans. Biomed. Eng.*, 52(1):74–87, 2005.
[3] M.S. Lewicki. Bayesian modeling and classification of neural signals. *Neural Computation*, 6(5):1005–1030, 1994.
[4] C.M. Gray, P.E. Maldonado, M. Wilson, and B. McNaughton. Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J Neurosci Methods*, 63(1-2):43–54, 1995.
[5] M.A. Wilson and B.L. McNaughton. Dynamics of the hippocampal ensemble code for space. *Science*, 261(5124):1055–1058, 1993.
[6] Z. Nenadic. Information discriminant analysis: Feature extraction with an information-theoretic objective. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(8):1394–1407, 2007.
[7] M. Abeles and M.H. Goldstein. Multispike train analysis. *Proc. IEEE*, 65(5):762–773, 1977.
[8] K.D. Harris, D.A. Henze, J. Csicsvari, H. Hirase, and G. Buzsáki. Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J. Neurophysiol.*, 84(1):401–414, 2000.
[9] R.Q. Quiroga, Z. Nadasdy, and Y. Ben-Shaul. Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput*, 16(8):1661–1687, 2004.
[10] Z. Nenadic and J.W. Burdick. A control algorithm for autonomous optimization of extracellular recordings. *IEEE Trans. Biomed. Eng.*, 53(5):941–955, 2006.
[11] S.T. Roweis and L.K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.
[12] J.B. Tenenbaum, V. de Silva, and J.C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
[13] R.R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.
[14] K.W. Pettis, T.A. Bailey, A.K. Jain, and R.C. Dubes. An intrinsic dimensionality estimator from near-neighbor information. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1(1):25–37, 1979.
[15] P. Grassberger and I. Procaccia. Measuring the strangeness of strange attractors. *Physica D: Nonlinear Phenomena*, 9(1-2):189–208, 1983.
[16] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Morgan Kaufmann, San Francisco, CA, 2nd edition, 1990.
[17] C. Pouzat, O. Mazor, and G. Laurent. Using noise signature to optimize spike-sorting and to assess neuronal classification quality. *J. Neurosci. Methods*, 122(1):43–57, 2002.
[18] C. Pouzat and G. Laurent. Locust data. Available online, http://www.biomedicale.univ-paris5.fr/SpikeOMatic/Data.html 2005.
[19] A.A. Szymanska, A. Hajirasooliha, and Z. Nenadic. Source location as a feature for the classification of multi-sensor extracellular action potentials. In *Proc. 6th Ann. Int. IEEE EMBS Conf. Neural Eng.*, pages 235–238, 2013.
[20] R. Benitez and Z. Nenadic. Robust unsupervised detection of action potentials with probabilistic models. *IEEE Trans. Biomed. Eng.*, 55(4):1344–1354, 2008.
[21] A.A. Szymanska and Z. Nenadic. Wavelet-approximated generalized matched filter for the detection of multisensor extracellular action potentials. In *Proc. 6th Ann. Int. IEEE EMBS Conf. Neural Eng.*, pages 810–814, 2013.