

Uncovering low-dimensional structure in high-dimensional representations of long-term recordings in people with epilepsy

Joaquín Rapela¹, Timothée Proix^{1,2}, Dmitrii Todorov¹, Wilson Truccolo^{1,2}

Abstract—Effective representations of recordings of epileptic activity for seizure prediction are high-dimensional, which prevents their visualization. Here we introduce and evaluate methods to find low-dimensional (2D or 3D) descriptors of these high-dimensional representations, which are amenable for visualization. Once low-dimensional descriptors are found, it is useful to identify structure in them. We evaluate clustering algorithms to automatically identify this structure. In addition, typical recordings of epileptic activity are long, extending for several days or weeks. We present and assess extensions of the previous methods to handle large datasets.

I. INTRODUCTION

Building good representations of data is essential for the performance of machine learning methods in general [2], and for the detection and prediction of epileptic seizures in particular [1], [7]. We previously developed novel data representations [6], [7] for the unique epilepsy recordings, that we analyze in this manuscript. These are recordings from layers two and three of human cortex obtained from a 10×10 ($4 \times 4 \text{ mm}^2$) microelectrode array intracortically implanted in epileptic patients. A challenge with these representations is their high dimensionality, which prevents their display and the visual recognition of patterns in the data. Various approaches have been proposed to visualize high-dimensional data. Some of these approaches display more than three dimensions, e.g., parallel coordinates [4] and pixel-based visualizations [5]. Other approaches convert high-dimensional data into two or three dimensions that can be easily displayed, e.g., t-distributed stochastic neighbor embedding (t-SNE, [11]). Here we introduce and evaluate methods to learn low-dimensional descriptors of high-dimensional representations of our microelectrode array recordings. We then use clustering algorithms to automatically discover structure in the learned low-dimensional descriptors.

Recent work have established the feasibility of long-term recordings for epileptic patients, extending from several days to years [3]. Hence, there is a need for new methods to extract low-dimensional descriptors and to discover structure in these very-long recordings of epileptic activity.

*This research was supported by the National Institute of Neurological Disorders and Stroke (NINDS), grant R01NS079533 (WT); the U.S. Department of Veterans Affairs, Merit Review Award I01RX000668 (WT); the Pablo J. Salame '88 Goldman Sachs endowed Assistant Professorship of Computational Neuroscience at Brown University (WT). The contents do not represent the views of the the United States Government or the U.S. Department of Veterans Affairs.

¹Department of Neuroscience & Carney Institute for Brain Science, Brown University, Providence RI, USA; ²U.S. Department of Veterans Affairs, Center for Neurorestoration and Neurotechnology, Providence RI, USA. Email: {joaquin_rapela,timothee_proix, dmitrii_todorov, wilson_truccolo}@brown.edu

In Section II-A we present and evaluate methods to find low-dimensional descriptors of high-dimensional representations of recordings of epileptic activity, and in Section II-B we describe and assess clustering methods to automatically discover structure in these low-dimensional descriptors. Each of these sections first describes and evaluates methods for smaller recordings and then extends these methods to handle very-long datasets.

II. RESULTS

A. Low-dimensional descriptors of high-dimensional representations of recordings of epileptic activity

1) *Short recordings: non-parametric t-SNE*: Figure 1a shows a scatter plot of the t-SNE two-dimensional descriptors (Section IV-B) of a 2,102-dimensional representation (Section IV-A) of a recording block containing three epileptic seizures. t-SNE separates well inter-ictal (black and gray), pre-ictal (green), ictal (red) and post-ictal (blue) samples. In contrast, the PCA two-dimensional descriptors (Section IV-C) failed to separate non-ictal features (Figure 1b).

We quantified the quality of the low-dimensional descriptors by classifying points as inter-ictal vs. non-inter-ictal, pre-ictal vs. non-pre-ictal, ictal vs. non-ictal and post-ictal vs. non-post-ictal based on their low-dimensional coordinates (multinomial logistic regression; function `multinom` of package `nnet` of [8]). We reasoned that a relevant low-dimensional descriptor should separate well different seizure stages. Therefore, the better a low-dimensional descriptor of recordings of epileptic activity is, the better we should be able to classify the seizure stage of a sample point based on its low-dimensional descriptor. Table I shows the area under the receiver operating characteristic curve (AUC; computed using function `roc` of package `pROC` of [8]) for classifying seizure states using t-SNE and 2D-PCA as low-dimensional descriptors and 487D-PCA as high-dimensional descriptors. For all seizure stages, excluding the ictal stage, t-SNE yields better classifications than PCA. The difference in AUC between the t-SNE low-dimensional descriptors and the 487D-PCA high-dimensional descriptors is not large, suggesting that t-SNE captures most of the relevant information of the high-dimensional representation of seizure recordings.

2) *Long-term recordings: parametric t-SNE*: In the previous section we showed that non-parametric t-SNE produced relevant low-dimensional (2-D) descriptors of high-dimensional (2,102-D) representations of recordings of epileptic activity. A limitation of non-parametric t-SNE is that its memory requirements grows as $O(N^2)$, where N is the number of input data points. Therefore, it cannot

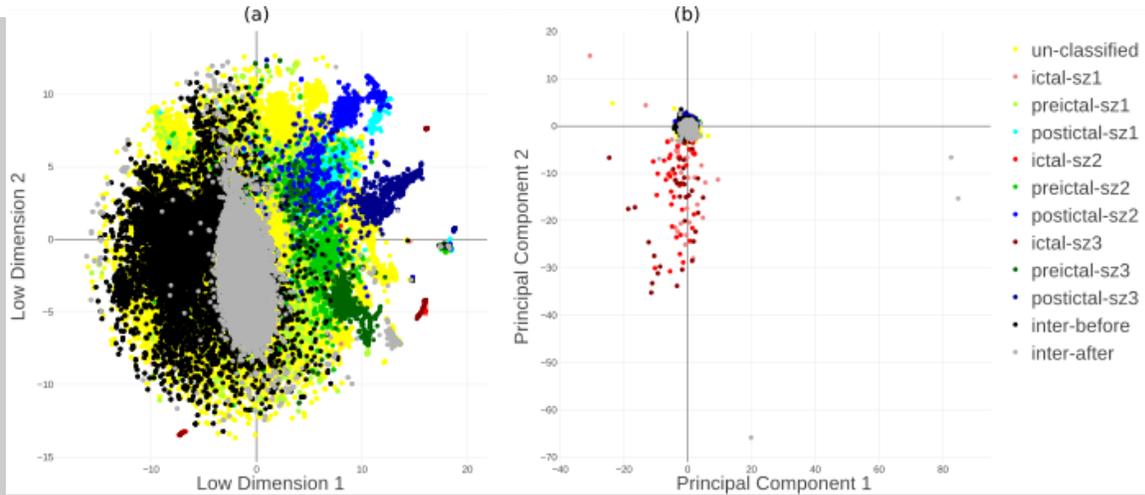


Fig. 1. Scatter-plot for t-SNE (a) and PCA (b) two-dimensional descriptors. Colors indicate seizure stages (black and gray: inter-ictal, green: pre-ictal, red: ictal, blue: post-ictal) and shades indicate seizure (light: seizure one, medium: seizure two, dark: seizure three). Black and gray indicate inter-ictal samples at the beginning and end, respectively, of the recording block. t-SNE, but not PCA, separates well inter-ictal, pre-ictal, ictal and post-ictal samples.

	t-SNE	2D-PCA	487D-PCA
inter-ictal	0.95	0.59	1.00
pre-ictal	0.90	0.58	0.99
ictal	0.71	0.88	0.86
post-ictal	0.99	0.68	0.99

TABLE I

AREA UNDER THE RECEIVER OPERATING CHARACTERISTIC CURVE FOR THE CLASSIFICATION OF SAMPLES AS BELONGING OR NOT TO DIFFERENT SEIZURE STATES (ROWS) FROM DIFFERENT DIMENSIONALITY REDUCTION METHODS (COLUMNS).

dimensional descriptors (Figure 2)c are more similar to the t-SNE low-dimensional descriptors (Figure 2)a than the linear-regression low-dimensional descriptors (Figure 2e). This superiority is reflected in a smaller MAE for the XGBoost (MAE=0.98) than for the linear-regression (MAE=4.01) descriptors. The superiority of XGBoost over the linear-regression model is also observed, although to a lesser degree, in the test dataset (compare Figures 2b, 2d and 2f; MAE XGBoost 3.61, linear regression 4.24).

We partitioned the complete dataset into training and testing subsets, and repeated the operation 100 times. For each partition we estimated XGBoost and linear-regression models with the training samples and evaluated their prediction power using the training and test subsets. Figures 3ab show histograms of MAEs for training and testing samples, respectively, for XGBoost (blue) and linear-regression (pink) models. For both training and testing samples MAEs from XGBoost models are significantly lower than those by linear-regression models ($p < 1e-4$), confirming the superiority of XGBoost over linear regression as a parametric t-SNE model.

Although not perfect, the XGBoost and linear-regression models provide reasonable approximations of t-SNE low-dimensional descriptor of high-dimensional inputs. Additionally, XGBoost performs significantly better than linear regression. Thus XGBoost is a practical parametric t-SNE model to obtain low-dimensional descriptors of very large number of samples, e.g. for long-term recordings in epilepsy.

B. Discovering structure in low-dimensional descriptors

In the previous section we described a methodology to build low-dimensional descriptors of high-dimensional descriptions of electrophysiological recordings from epilepsy. We observed that the obtained low-dimensional descriptors were separating well inter-ictal, pre-ictal, ictal and post-ictal data points. However, as shown in Figure 1a, this low-dimensional descriptor has a finer structure than the

represent very long-duration recordings extending for weeks (several million samples), typical of recent recordings for epileptic patients.

To overcome this limitation, we trained parametric models (XGBoost and linear-regression models, Section IV-D) to map high-dimensional representations to their t-SNE low-dimensional descriptors. Once a parametric model has been trained, we can input an arbitrarily large number of high-dimensional feature vectors and obtain approximations of their t-SNE low-dimensional descriptors.

Figure 2a shows the t-SNE low-dimensional descriptors used to train parametric t-SNE models and Figure 2b shows the low-dimensional descriptors used to test their accuracy. Figure 2c and 2d shows the low-dimensional descriptors predicted by XGBoost models for the train and test data, respectively. Each figure's title reports the Maximum Absolute Error (MAE) achieved by the model (XGBoost MAE train=0.98, test=3.61). Figures 2f and 2g show the t-SNE low-dimensional descriptors for the train and test dataset, respectively, predicted by the linear-regression model and the corresponding MAEs (linear regression MAE train=4.01, test=4.24).

We observe that for the training dataset the XGBoost low-

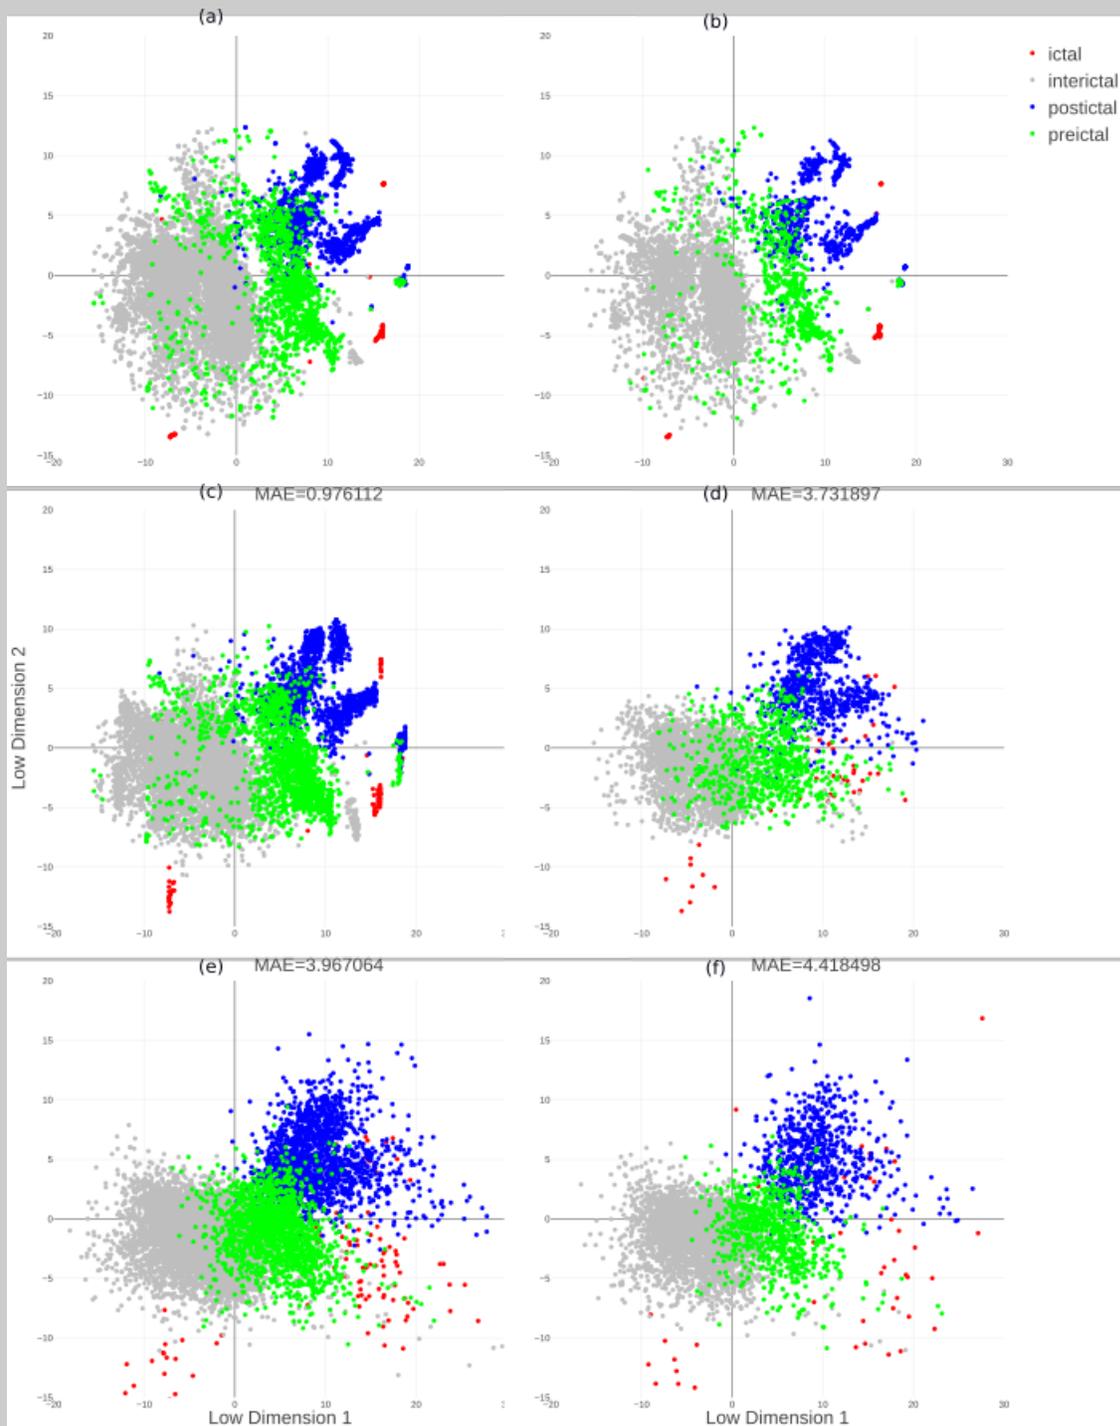


Fig. 2. XGBoost and linear regression as parametric t-SNE descriptors. (a,b) t-SNE low-dimensional descriptor used to train and test the parametric models in c-f. (c,d) XGBoost predictions of t-SNE low-dimensional descriptors from the train and test datasets. (e,f) as (c,d) but for linear-regression parametric t-SNE model. Color indicate seizure stage (gray: inter-ictal, green: pre-ictal, red: ictal and blue: post-ictal). Titles in c-f show maximum absolute error (MAE) between parametric model predictions and corresponding non-parametric t-SNE descriptors (e.g., the title in (c) shows the MAE between the XGBoost predictions for train samples in (c) and the t-SNE low-dimensional descriptor of train data in (a)). For both train and test datasets, MAEs of t-SNE are smaller than those of linear regression (see also resampling analysis in Figure 3).

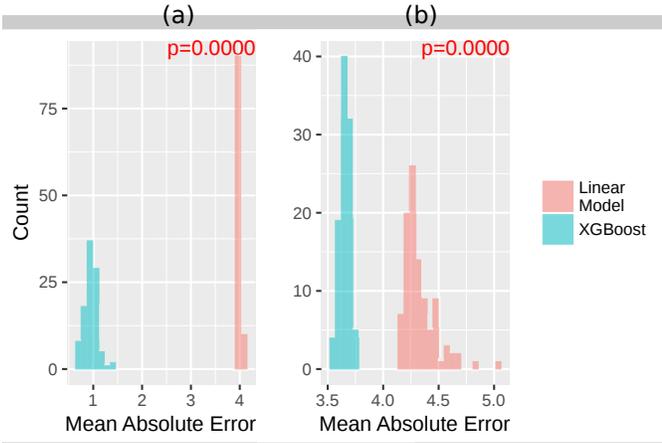


Fig. 3. Maximum absolute deviations (MAEs) between low-dimensional descriptors of train (a) and test (b) samples predicted by XGBoost models (blue) and by linear-regression models (pink) for different random data partitions into train and test samples. p-value in top-right is the result of a paired t-test for equality of the mean of XGBoost and linear-regression predictions. Predictions by the XGBoost model are significantly better than those by the linear-regression model, for both train and test samples.

one given by the previous seizures stages. In this section we introduced methods to automatically discover this finer structure, first in shorter-duration recordings and next in longer-duration ones.

1) *Short recordings: parametric t-SNE*: To automatically discover structure in t-SNE low-dimensional descriptors of recordings of epileptic activity we clustered these descriptors using the algorithm described in [9]. The obtained clusters are shown in Figure 4a and seem to capture more visually salient structure than those obtained using k-means (Section IV-E) shown in Figure 4b.

2) *Longer-duration recordings: non-parametric t-SNE*: A challenge of most clustering algorithms is that their memory requirement grows as $O(N^2)$, where N is the number of points to be clustered. Therefore, most clustering algorithms cannot be used to group very large number of samples, typical of recordings in epilepsy.

To cluster a very large number of low-dimensional t-SNE descriptors, we propose the following procedure:

- 1) Use non-parametric t-SNE to obtain low-dimensional descriptors of a smaller (representative) training dataset of high-dimensional representations of recordings of epileptic activity (as in Figure 1a),
- 2) Use the clustering algorithm developed by [9] to cluster this low-dimensional descriptors of the train dataset (as in Figure 4a),
- 3) Use a parametric version of t-SNE (e.g., XGBoost or linear regression) to derive low-dimensional descriptors of (a possibly very large number of) new samples, not included in the train data (as in Figure 2df).
- 4) Assign the low-dimensional descriptor of the new samples to a cluster, from step two. Note: the method of assignment of new samples to existing clusters varies between clustering algorithms (e.g., in k-means

a new sample should be assigned to the cluster with closest centroid, but in the algorithm by [9] it should be assigned to the same cluster as its nearest neighbor of higher local density).

For a dataset of reasonable size, we can simply run non-parametric t-SNE on the whole dataset (as in Figure 1a), and then cluster the output of t-SNE using the algorithm by [9] (as in Figure 4a). The cluster labels obtained in this way are the gold standard that the above procedure should achieve when applied to this not too large dataset. To test the accuracy of the above procedure we used a smaller dataset, for which we could compute the above gold standard. We then partitioned the dataset into train and test subsets. If this procedure worked perfectly on this smaller dataset, the cluster assigned to a new sample in step four should be equal to the cluster assigned to it when non-parametric t-SNE and the algorithm by [9] are run on the whole dataset. Thus, to quantify the accuracy of the above procedure, we measure percentage of correct assignments (i.e., cases where the cluster label assigned to a sample, when t-SNE and the clustering algorithm are run on the whole dataset, is the same as the cluster label assigned using the above procedure).

We partitioned the complete dataset into train and test samples 100 times, and for each partition we computed the percentage of correct assignments using XGBoost and linear regression parametric t-SNE models. The obtained percentage of correct assignments for train and test samples are shown in Figures 5a and 5b, respectively.

XGBoost assigned new samples to their correct cluster around 80% of the time, and significantly outperformed linear-regression. This shows that the procedure introduced at the beginning of this section, using XGBoost as parametric t-SNE model, is practical to discover low-dimensional structure in high-dimensional representations of long-duration recordings of epileptic activity.

III. DISCUSSION

We found that t-SNE yields low-dimensional descriptors of high-dimensional representations of recordings of epileptic activity that separates well seizure stages, and that are superior to low-dimensional descriptors derived from PCA (Figure 1). We proposed and evaluated the use of XGBoost and linear regression as parametric t-SNE models, showed that XGBoost yields accurate predictions of low-dimensional t-SNE descriptors, and demonstrated that XGBoost is better than linear regression as a parametric t-SNE model (Figure 2 and 3). We showed that the clustering algorithm described in [9] identifies visually salient clusters in t-SNE low-dimensional descriptors, which are superior to clusters found by PCA (Figure 4). We proposed a methodology to discover structure in high-dimensional representations of long-duration recordings from epileptic subjects. We provided preliminary evidence for its feasibility by showing that, with XGBoost as parametric t-SNE model, this methodology assigned correct structural labels to 80% of new samples.

The assessment of the methods introduced here is preliminary. We evaluated our methods in a 23-hour and 10-minute

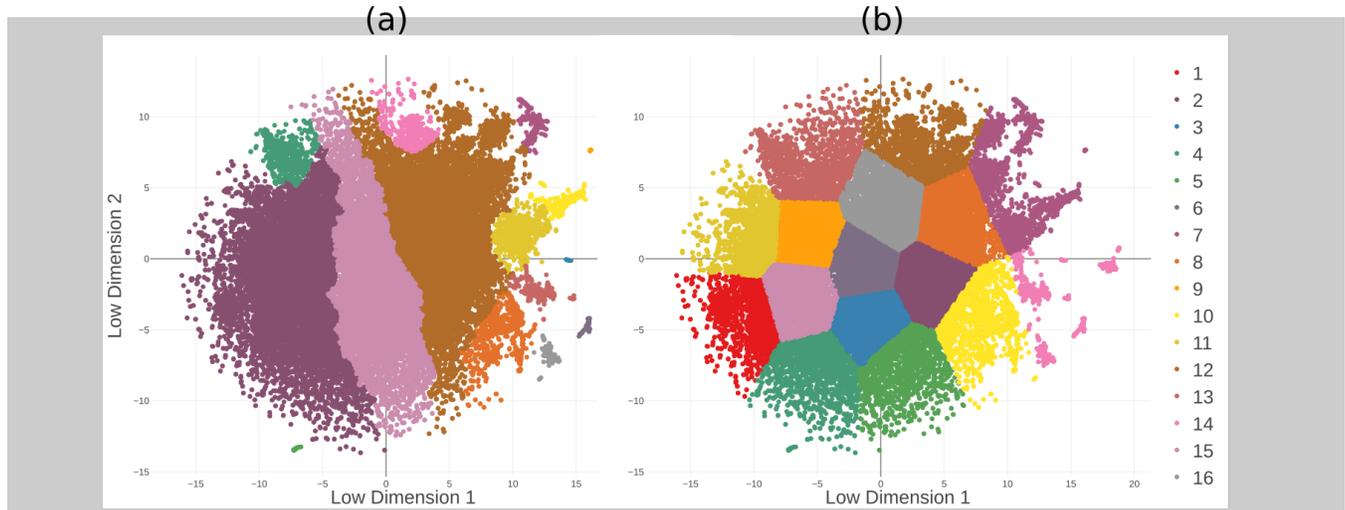


Fig. 4. Scatter-plot of t-SNE two-dimensional descriptor of all samples colored with cluster assignments from the [9] (a) and the k-means (b) algorithm. The former algorithm, but not the latter, identifies visually-salient clusters. Here we use the [9] algorithm to discover structure in t-SNE low-dimensional descriptors.

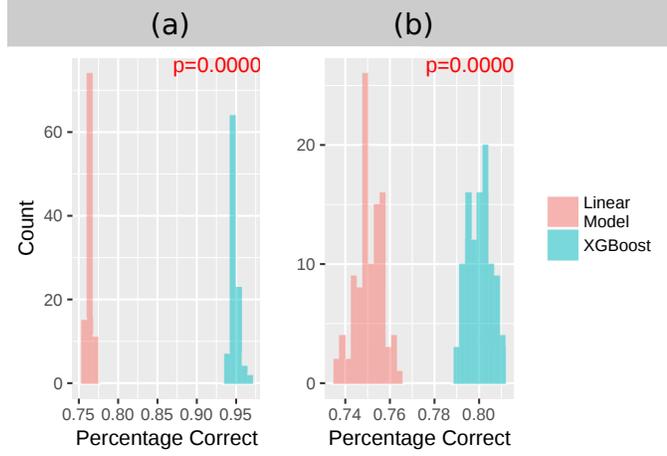


Fig. 5. Accuracy in the assignment of parametric t-SNE (XGBoost and linear-regression) predictions to [9] clusters of non-parametric t-SNE samples. Percentage of parametric t-SNE predictions assigned to the same cluster as the corresponding non-parametric t-SNE sample, for 100 parametric t-SNE models estimated and evaluated with different random train and test subsamples of the complete dataset. XGBoost predictions (blue) are significantly more often assigned to the same cluster as the corresponding non-parametric t-SNE sample than linear-regression predictions (pink). The 80% mean correct assignment of XGBoost predictions to clusters of non-parametric t-SNE samples suggests that the procedure using XGBoost, described at the beginning of Section II-B.2, is a good parametric t-SNE method, demonstrating that it is practical to find structure in t-SNE low-dimensional descriptors of large datasets of epilepsy recordings.

block from one participant with epilepsy containing three spike-and-wave seizures. Future research will test the new methods in more subjects and other types of seizures. We will also study the performance of the proposed methods on epilepsy neural recordings extending for several days and weeks. The clusters estimated by the [9] algorithm are appealing because they are visually salient. Next, we will study if these clusters are physiologically relevant. We hope

that some of these clusters will correspond to neural activity anticipating the occurrence of epileptic seizures, allowing the use of the methods introduced here for forecasting the occurrence of seizures.

IV. METHODS

A. High-dimensional representations of epilepsy recordings

We recorded broadband intracortical field potentials (0.3 Hz - 7.5 kHz; sampling at 30 kHz) in a 23-hour and 10-minute block period (42,896 samples) using a 10×10 (4×4 mm²) microelectrode array (96 recording electrodes plus 4 references) implanted in a person with focal epileptic seizures [10], [12]. Recordings were obtained according to an approve IRB protocol at MGH and with patient consent. We did not include recordings from 6 electrodes due to poor signal quality.

We measured local field potentials (LFPs) and multi-unit activity (MUA) counts. LFPs were calculated by low-pass filtering the broadband field potentials (Butterworth filter, order nine, cutoff frequency 500 Hz) and downsampling the result at a frequency of 2 kHz. MUA counts we computed by high-pass filtering the broadband field potentials (Butterworth filter, order nine, cutoff frequency 250 Hz), calculating the standard deviation of this high-passed signal, sd , identifying MUA spikes as fluctuations in the high-passed signal below -3 times sd , and calculating MUA counts by summing the number of MUA spikes in 0.5 millisecond windows (i.e., 2 kHz frequency).

Field potentials were segmented in four-second (overlapping) time windows, extracted every two seconds. This procedure yielded 42,896 sample times, and 90 four-second sample windows per sampling time (one sample window for each valid recording electrode). For each sample time, we represented the 90 sample windows as 2,102-dimensional vectors by extracting features from local-field potential (LFP)

and multi-unit activity (MUA) derived from the recorded voltages. These features included LFP power spectrum (multitaper) in 10 frequency bands, eigenvector centrality, leading eigenvalue, mean and variance of LFP coherence matrix for 10 frequency bands, MUA counts for each recording electrode, eigenvector centrality and largest eigenvalue of the correlation matrix of the MUA count, and eigenvector centrality and largest eigenvalue of the correlation matrix of the envelope of the MUA count. Please refer to [6], [7] for further details on feature extraction from LFP and MEA extracted from our microelectrode array recordings.

B. t-SNE low-dimensional descriptor

The input for t-SNE was a matrix A containing the high-dimensional representation of the epilepsy recordings; $A \in \mathbb{R}^{n_{\text{Features}}=2,102 \times n_{\text{Samples}}=42,896}$. We first reduced the dimensionality of this matrix using PCA (Section IV-C) keeping $n_{\text{PComp}}=487$ principal components capturing 90% of the variance of the high-dimensional representations, giving a matrix $B \in \mathbb{R}^{n_{\text{PComp}} \times n_{\text{Samples}}}$. Next, we computed L_1 distances between all the dimensionality reduced sample points (i.e., between all pairs of columns in matrix B) yielding a matrix $C \in \mathbb{R}^{n_{\text{Samples}} \times n_{\text{Samples}}}$ where $C[i, j] = L_1\text{-distance}(B[:, i], B[:, j])$. This distance matrix C was the input to the Barnes-Hut implementation of t-SNE (function `Rtsne` from package `Rtsne` of [8]) with default parameters. We used the default parameter value `perplexity=30` and verified that the results presented here were qualitatively similar for `perplexity` $\in [10, 50]$.

C. PCA low-dimensional descriptor

We used the function `prcomp` from [8] with parameters `center=scale=TRUE` to z-score the high-dimensional features across time) and keeping a pre-specified number of principal components ($n_{\text{PComp}} = 2$ for Figure 1b and for column 2D-PCA of Table I; $n_{\text{PComp}} = 487$ for t-SNE pre-processing, Section IV-B and for column 487D-PCA of Table I).

D. Parametric t-SNE models

XGBoost: we used the function `xgboost` from package `xgboost` of [8] with regression trees (maximum depth=10, learning rate=0.1). The optimal number of regression trees was estimated by 5-fold cross validation using the function `xgb.cv` from package `xgboost` of [8].

linear regression: we used the function `lm` from package `stats` of [8].

E. Clustering t-SNE low-dimensional descriptors

Rodriguez & Laio (2014): we used the Matlab implementation of the algorithm reference in their paper. As input to this implementation we provided an L_1 distance matrix between all pairs of t-SNE samples. As recommended in the article, we set the parameter percentage of neighbors to 2%. We manually selected $\rho_{\text{min}} = -0.1$ (minimum local density) and $\delta_{\text{min}} = 2.0$ (minimum distance to a point of higher density) in order to get a relatively large number of clusters ($n_{\text{Clusters}}=16$).

k-means: we used the same number of clusters as for the [9] algorithm (function `kmeans` of package `stats` of [8]).

REFERENCES

- [1] Mehdi Aghagolzadeh, Leigh R Hochberg, Sydney S Cash, and Wilson Truccolo. Predicting seizures from local field potentials recorded via intracortical microelectrode arrays. In *Engineering in Medicine and Biology Society (EMBC), 2016 IEEE 38th Annual International Conference of the*, pages 6353–6356. IEEE, 2016.
- [2] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [3] Mark J Cook, Terence J O'Brien, Samuel F Berkovic, Michael Murphy, Andrew Morokoff, Gavin Fabinyi, Wendyl D'Souza, Raju Yerra, John Archer, Lucas Litewka, Sean Hosking, Paul Lightfoot, Vanessa Ruedebusch, W Douglas Sheffield, David Snyder, Kent Leyde, and David Himes. Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: a first-in-man study. *The Lancet Neurology*, 12(6):563–571, June 2013.
- [4] Alfred Inselberg and Bernard Dimsdale. Parallel coordinates for visualizing multi-dimensional geometry. In *Computer Graphics 1987*, pages 25–44. Springer, 1987.
- [5] Dominik Jäckle, Fabian Fischer, Tobias Schreck, and Daniel A Keim. Temporal mds plots for analysis of multivariate data. *IEEE transactions on visualization and computer graphics*, 22(1):141–150, 2016.
- [6] Yun S Park, Leigh R Hochberg, Emad N Eskandar, Sydney S Cash, and Wilson Truccolo. Early detection of human epileptic seizures based on intracortical local field potentials. In *Neural Engineering (NER), 2013 6th International IEEE/EMBS Conference on*, pages 323–326. IEEE, 2013.
- [7] Yun S Park, Leigh R Hochberg, Emad N Eskandar, Sydney S Cash, and Wilson Truccolo. Early detection of human focal seizures based on cortical multiunit activity. In *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE*, pages 5796–5799. IEEE, 2014.
- [8] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2012. ISBN 3-900051-07-0.
- [9] Alex Rodriguez and Alessandro Laio. Clustering by fast search and find of density peaks. *Science*, 344(6191):1492–1496, 2014.
- [10] Wilson Truccolo, Jacob A Donoghue, Leigh R Hochberg, Emad N Eskandar, Joseph R Madsen, William S Anderson, Emery N Brown, Eric Halgren, and Sydney S Cash. Single-neuron dynamics in human focal epilepsy. *Nature neuroscience*, 14(5):635, 2011.
- [11] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [12] Fabien B Wagner, Emad N Eskandar, G Rees Cosgrove, Joseph R Madsen, Andrew S Blum, N Stevenson Potter, Leigh R Hochberg, Sydney S Cash, and Wilson Truccolo. Microscale spatiotemporal dynamics during neocortical propagation of human focal seizures. *Neuroimage*, 122:114–130, 2015.