

Dictionary Learning for Spontaneous Neural Activity Modeling

Eirini Troullinou^{*†}, Grigorios Tsagkatakis[†], Ganna Palagina^{‡§}
 Maria Papadopoulou^{*†}, Stelios Manolis Smirnakis^{‡§}, and Panagiotis Tsakalides^{*†}

^{*}Department of Computer Science, University of Crete, Heraklion, 70013, Greece

[†]Institute of Computer Science, Foundation for Research and Technology Hellas (FORTH), Heraklion, 70013, Greece

[‡]Department of Neurology, Brigham and Womens Hospital, Harvard Medical School, Boston MA 02115

[§]Boston VA Research Institute, Jamaica Plain Veterans Administration Hospital, Harvard Medical School, Boston, United States

Abstract—Modeling the activity of an ensemble of neurons can provide critical insights into the workings of the brain. In this work we examine if learning based signal modeling can contribute to a high quality modeling of neuronal signal data. To that end, we employ the sparse coding and dictionary learning schemes for capturing the behavior of neuronal responses into a small number of representative prototypical signals. Performance is measured by the reconstruction quality of clean and noisy test signals, which serves as an indicator of the generalization and discrimination capabilities of the learned dictionaries. To validate the merits of the proposed approach, a novel dataset of the actual recordings from 183 neurons from the primary visual cortex of a mouse in early postnatal development was developed and investigated. The results demonstrate that high quality modeling of testing data can be achieved from a small number of training examples and that the learned dictionaries exhibit significant specificity when introducing noise.

1. Introduction

Neurons are the elementary processing units in the central nervous system and are connected to each other in intricate patterns. Neuronal signals consist of short electrical pulses, termed action potentials or spikes, which form the elementary signals for transmission [1]. A chain of action potentials, generated by a single neuron is called a spike train and is a time sequence of firing events, which occur at regular or irregular intervals. Thus, these pulse-coded signals that represent the information encoded by a neuron, employ both binary and temporal coding mechanisms, which expand the signals into higher dimensional spaces [2].

Visual information relayed from the retina to the primary visual cortex (V1) is encoded in real time via the joint firing of multiple neurons. Although much has been learned about the properties of single neuronal units, the rules by which neurons coordinate their activity in cortical networks to represent information about the visual stimulus remains one of the fundamental unanswered questions in neuroscience [3]. To understand why, one needs only to consider that responses of single units are both noisy and ambiguous,

that is responses to the same stimulus vary considerably and responses to multiple different stimuli can be the same.

The high dimensionality of the observations introduces a challenge in terms of signal analysis. The curse of dimensionality is an indicative example of such challenge. In order to capture only the significant information encoded in the high-dimensional spaces, dimensionality reduction methods can be applied [4]. Despite the benefits of traditional approaches like PCA, the majority of methods are unable to include prior knowledge through learning of the underlying signals statistics. The objective of this work is to explore to what extent it is possible to reduce the ambient dimensionality of multiple neuron activation patterns to the effective dimension of the underlying process without significant penalty in the subsequent data analysis.




Figure 1: The proposed dictionary learning framework for neuronal signal modeling. (1) Data is acquired. (2) Action potentials are transformed to binary values. (3) Data is represented in matrix format. (4) Dictionary learning is applied on the training data. (5) Testing examples are represented by a linear combination of a few dictionary elements. (6) Reconstruction error is quantified.

To address this challenge, we propose modeling the neural activation data as an instance of a sparse coding and dictionary learning problem [5], [6]. A visual illustration of the data analysis chain is shown in Figure 1, which depicts how dictionaries that learned directly from the observations

can be employed to represent neuronal binary signals as linear combination of a small number of prototypical signals called atoms. Although sparsity seeking dictionary learning methods have achieved great success in modeling complex signals from images [7] to wireless sensor network [8] data, no such approach has been considered for handling the intricacy of the multiple neuron activation patterns modeling considered in this work.

This work includes the following innovative aspects:

- The application of sparsity modeling and dictionary learning on neural network data.
- The evaluation of the generalization capacity of sparsifying dictionary learning.
- The evaluation on sensitivity of the learned dictionary with respect to noise.

2. Proposed modeling approach

The working hypothesis in this work is that neuron activation patterns can be efficiently represented as linear combinations of a small number of prototypical patterns of activity called atoms. Each data point, corresponding to the activation pattern of a set of neurons during a small temporal time unit, can thus be represented through a sparse vector of coefficients, (i.e. sparse coding process). Furthermore, the collection of atoms, called dictionary, can be either explicitly modeled or, in case when no such models exist, it can be learned from a set of training data through a dictionary learning process. Here, we apply the K-SVD dictionary learning algorithm [9] on a real dataset to quantify the quality of reconstruction of sparse low-dimensional representation.

Formally, according to the sparse representation framework, given a dictionary $\mathbf{D}$ that contains K prototype signal-atoms for columns and an input signal $\mathbf{y} \in \mathbb{R}^N$, we search for a vector $\mathbf{x}$ that optimizes a certain sparsity level. An approach to this problem is the minimization of the following l_0 norm problem:

$$\min_{\mathbf{x}_i} \|\mathbf{y}_i - \mathbf{D}\mathbf{x}_i\|_2^2 \quad \text{subject to} \quad \|\mathbf{x}_i\|_0 \leq T_0 \quad (1)$$

where $\|\mathbf{x}_i\|_0$ is the l_0 pseudo-norm which counts the number of non-zeros elements. T_0 is the sparsity level which denotes the number of nonzero elements for every $\mathbf{x}_i$, namely for every column i of sparse coefficient vector $\mathbf{x}$.

An important issue regarding the formulation in Eq.(1) is that the l_0 minimization is an NP-hard problem and therefore inefficient to solve for even moderate sized problems. To address this issue, greedy approaches, such as the OMP [10] algorithm, have been proposed. OMP greedily tries to identify the elements that contain most of the signal energy by iteratively selecting the dictionary element that best matches the signal by projecting the input signal to the linear span of the selected elements and estimating the residual error until an acceptable approximation limit or a maximum number of iterations is reached.

Eq.(1) assumes that a dictionary $\mathbf{D}$ is available, which is the case when assumptions regarding the characteristics of the signals are made. To handle the intricacy of

modeling multiple neural activities, we employ K-SVD, a dictionary learning algorithm [9] designed for maximizing signal reconstruction quality subject to sparsity constraints. In the sparse representation problem, each input signal is represented by a linear combination of a small number of dictionary elements. Formally, let $\mathbf{Y} \subset \mathbb{R}^M$ denote the set of training signals of interest, e.g. a neuronal signal. Given $[\mathbf{y}_1, \dots, \mathbf{y}_N] \in \mathbf{Y}$, the first stage of dictionary learning is to learn a dictionary $\mathbf{D} \subset \mathbb{R}^N$, by finding a set of signals, the atoms, $\mathbf{D} = [\mathbf{d}_1, \dots, \mathbf{d}_K]$, that form the building blocks of $\mathbf{Y}$. An input signal $\mathbf{y} \in \mathbf{Y}$ can then be represented by a linear combination of a small number of atoms, i.e. $\mathbf{y}_i = \mathbf{D}\mathbf{x}_i + \mathbf{E}$ where $\mathbf{E}$ captures the contribution of noise due to modeling.

The training phase of K-SVD involves searching for the best dictionary that will support the sparse representation of the testing set $\mathbf{Y}$ by minimizing the error E . Depending on the particular application, desired accuracy and nature of the signals, dictionary learning may take different forms. Yet is often formulated as a least squares optimization of the form:

$$\min_{\mathbf{D}, \mathbf{X}} \|\mathbf{Y} - \mathbf{D}\mathbf{X}\|_F^2 \quad \text{subject to} \quad \|\mathbf{x}_i\|_0 \leq T_0 \quad \forall i. \quad (2)$$

where $\|\cdot\|_F$ denotes the frobenius norm of a matrix. The coefficient matrix $\mathbf{X}$ has thus T_0 nonzero entries in every column i and these can have arbitrary values.

In K-SVD, the expression in (2) is iteratively minimized and each iteration includes two steps. First, the dictionary $\mathbf{D}$ is assumed to be fixed and the goal is to find the optimal coefficient matrix $\mathbf{X}$. As finding the optimal $\mathbf{X}$ is an NP-hard problem, the OMP method is used for the sparse coding. Once the sparse coding task is completed, the dictionary is updated one column at a time. Specifically to update column $\mathbf{d}_k$, all other columns in $\mathbf{D}$ are fixed and the atom is updated such that it minimizes the representation error, through a singular value decomposition. The algorithm repeats these two steps until reaching the stopping conditions based on a threshold in the number of iterations or a representation error. To account for the fact that the reconstruction can produce real-valued signals, a hard thresholding to either 0 or 1 is performed in this work.

3. Evaluation

3.1. Dataset Collection

To evaluate the merits of the proposed modeling approach, a novel dataset consisting of true measurements is employed. Data was collected using two-photon calcium imaging in the neocortex of a 9-day old mouse (C57BL/6). Simultaneously 183 layer 2/3 neurons were imaged using calcium indicator OGB-1 (imaging depth 130 microns from pia). 29 minutes of spontaneous activity were recorded, comprised of 11K, i.e. 11000 frames, each of 0.1451 sec duration. The raw fluorescence movie was motion-corrected to remove slow xy-plane drift. After motion correction, we used ImageJ software [11] to draw the ROIs of cells around cell body centers, staying 12 pixels from the margin of

a cell to avoid contamination with neuropil signals. We then averaged the signals of cell ROI pixels and converted them to dF/F [12]. To determine the onsets of spontaneous calcium responses, the dF/F timecourse for each cell was thresholded, using the noise portion of the data, to 3 standard deviations above noise. To make a binary eventogramme of the responses, for each cell the frames containing the onsets for this particular cell were assigned the value 1, and all other frames were assigned the value 0. The resulting binary eventogramme of all cells was used in subsequent analysis.

3.2. Evaluation metrics

In this section we report the performance of the dictionary learning algorithm on the modeling of neuronal signals. To assess the performance of K-SVD in neuronal signal reconstruction, we explore the impact of the following parameters: (i) dictionary size, i.e., the number of elements considered in the dictionary, (ii) the sparsity level, i.e., the number of atoms used for representation, and (iii) the training size used for dictionary learning.

The analysis aims to examine if a good reconstruction can be achieved via a trained dictionary. For that, it selects a random signal of size 2K, 3K, and 4K (out of the entire 11K) for training (training time instances) and employs 5K testing signals. The random selection of the signals in the training and test sets, allows us to introduce invariance with respect to temporal correlations and focus on the synchronicity of the neurons activity. For a test set consisting of 5K instances, the total number of events for the 183 neurons is 915K.

For the following figures each point corresponds to mean performance over ten realizations randomly splitting the dataset into training and testing sets, while the error bar demonstrates the standard deviation. Without loss of generality we make the following simplification: The reconstructed events are mainly $\{0, 1\}$ but because of the fact that we deal with a reconstruction problem, sometimes arbitrary reconstructed values that are neither 0 nor 1 appear in small numbers. Thus, to make the outputs of the dictionary modeling process binary, values greater than 0.5 are considered activations, while the rest as not. Future work will explore potential modification of sparse coding and dictionary learning to the purely binary regime.

3.3. Evaluation of signal modeling

Fig. 2 illustrates the performance of K-SVD as a function of the dictionary size for three different training set sizes. In this figure where the sparsity level is fixed to 4 (i.e. the sparse coefficient matrix $\mathbf{X}$ has at most 4 non-zero entries in every column), we observe that the total number of misclassified events comes at maximum up to 2K, which is dramatically smaller compared to the 915K total events, in the order of 0.2%.

The results demonstrate that increasing the number of examples has a positive effect on the system's learning ability until a dictionary size of 250 atoms is used. Increasing the dictionary size up to 250, the reconstruction




Figure 2: Total number of misclassified events with respect to dictionary size for sparsity level 4

error becomes significantly smaller, which is an expected behavior as there is a wider variety of dictionary atoms and the system selects those ones that will better approach the original test signal. However, increasing the dictionary size more than 250 dictionary elements causes the reconstruction error to increase. This can be attributed to overfitting of the system due to the increased dictionary size, in combination with the hard sparsity constraints.

For sufficient sparsity however, one expects that increasing the dictionary size would have a positive effect on the reconstruction quality. This is indeed confirmed by Fig. 3, where sparsity level is increased to 20. The results demonstrate that the overall number of misclassified events is smaller compared to lower sparsity levels, while the error is monotonically decreasing with increased dictionary size.




Figure 3: Number of misclassified events with respect to dictionary size for sparsity level 20

The behavior demonstrates that by increasing the sparsity level, namely the non-zero entries in every column of coefficient matrix $\mathbf{X}$, the system has higher flexibility in modeling complicated high dimensional signals and achieves better generalization capacity. Similar to the case in Fig 2, increasing the number of training signals leads to better performance.

A natural question to ask is whether the number of misclassified events will keep decreasing if we increase the sparsity level. To answer this question, the sparsity level was increased to 50 and the results are shown in Fig. 4.

The results suggest that the number of misclassified events increased compared to Fig. 3. One possible explanation for the behavior is that by increasing the sparsity




Figure 4: Number of misclassified events with respect to dictionary size for sparsity level 50

level, the system exhibits signs of overfitting, which limits its generalization ability for modeling new test signals. Additionally, we notice that the number of training examples does not affect the reconstruction error. A key observation from these results is that it is possible to achieve a good reconstruction quality, and thus modeling capability, with a relatively small number of training examples.

Given the fact that firing events occupy only the 0.36% of the whole dataset and the rest of the events are zeros, a critical question is if our system can detect and reconstruct the firing events. In order to examine this, we report a confusion matrix in Table 1, where we examined the case of training the system with 2K and 4K signals for sparsity levels 4, 20 and 50, while performance is reported on 5K testing signals.

Table 1. Confusion Matrix of Reconstructed Events

Training Examples		2000		4000	
Sparsity	Predicted \ Actual	0	1	0	1
4	0	911523	162	911685	0
	1	1953	1362	1546	1769
20	0	911663	22	911674	11
	1	652	2663	613	2702
50	0	911526	159	911668	17
	1	1082	2233	1751	1564

The worst performance in terms of reconstructed events occurs when we use 2K (the smallest) number of training examples combined to a hard sparsity constraint, namely with sparsity level 4. In contrast to this, the best results are obtained when the number of training examples is increased to 4K with sparsity level 20, where we observe an important improvement in the reconstruction of aces. Finally, for sparsity level 50 the system exhibits signs of poor generalization indicated by worse performance compared to sparsity level 20, justifying to an extent the applicability of sparsity as a modeling constraint.

3.4. Evaluation of sensitivity in modeling noisy test signals

In this experiment, we explore the ability of the learned dictionary to discriminate true patterns from noisy signals. Similar to the previous section, the dictionaries are trained using noise-free examples, while the performance in reconstruction is evaluated when noisy examples are sparsely represented in the learned dictionaries. For all of the following experiments the training examples are set to 3K and the test examples to 5K, while the sparsity level is set to 20 based on the observed behavior in the previous experiments.

We examine three cases in the number of noisy neuronal signals by adding noise to 10, 80 and 183 neuronal signals out of the 183 in total. To create noisy neuronal signals, we randomly select a subset of neurons and randomly change a number of events for each neuron from firing to non-firing, namely 0 values are turned into 1 values and vice versa. This type of noise is quite intense, since the independence of the noisy signals completely changes the characteristics of the noise-free test signals. The objective of this experiment, is to quantify the extent to which the learned dictionaries truly encode neuronal activities or simply model random binary noise. In order to verify the learning ability of the K-SVD, we expect to see *significant increases* in reconstruction error when noise starts dominating the clear signal.

To evaluate the reconstruction error in the experiments that correspond to Fig. 5 and 6, we compare the *noise-free* test signals with the *reconstructed* test signals for dictionary sizes 200 and 400 respectively, in terms of number of noise-afflicted neurons. The 0 value in the horizontal axes corresponds to the noise-free case, namely the reconstruction is based on clean test signal (0 flipped events)




Figure 5: Number of misclassified events as a function of the number of flipped events, for dictionary size 200.

In Fig. 5 one can observe that there is a monotonic increase in the number of misclassified events with increasing number of erroneous events. This behavior demonstrates that the learned dictionary was able to capture the underlying statistics of the true signals and does not simply model random noise. Furthermore, introducing noise in more neurons has a direct impact on the ability of the learned dictionary to represent such signal ensembles.

Similar behavior is observed for larger dictionaries as shown in Fig. 6. An increase in dictionary size favors the




Figure 6: Number of misclassified events as a function of the number of flipped events, for dictionary size 400.

learning capacity of the system, and since the underlying distribution changes due to the noise, it becomes more challenging for the system to reconstruct the observations. This is manifested by the increased reconstruction error.

Fig. 7 demonstrates the performance of K-SVD in terms of the number of flipped events but in this case, we compare the *noisy* test signals with the *reconstructed* test signal. Specifically, while in Fig. 5 we explore if the modeling process can separate the noise from the signal, in Fig. 7 we examine the degree to which the trained dictionary is specific to representing the original (neural) data, as opposed to the randomly degraded noisy patterns we introduced.




Figure 7: Differences in events between noisy test signals and reconstructions, for dictionary size 200.

Similar to the results in Fig. 5, increasing either the number of noisy neurons or the number of erroneous events leads to higher reconstruction error. Furthermore, we observe that as we increase the amount of noise, this leads to an exponential increase in the error, which doesn't occur in the case of Fig. 5, in which we observe a more robust estimation of the noise-free signal in the presence of increased noise. The behavior demonstrated in these experimental results suggest that increasing the amount of noise has a profound impact on the reconstruction quality, thus the learned dictionary effectively captures the underlying signal characteristics.

4. Conclusion

In this work we have investigated the possibility of representing neuronal signals in low-dimensional subspaces

via dictionary learning. In order to achieve this, we used the dictionary learning algorithm K-SVD. The fundamental idea is that a set of neuronal signals can be represented as a linear combination of a few basic elements learned directly from the data. Extensive experimental results show the efficacy of the algorithm in representing such signals in low-dimensional subspaces, maintaining their structure (e.g. synchronicity of firing events) and simultaneously preserving only the necessary information needed for research analysis. The experimental results over neuronal ensemble signal reconstruction were very encouraging, which suggests that further research is necessary.

Acknowledgment

This work is funded by the DEDALE project, contract no. 665044, within the H2020 Framework Program of the European Commission. Maria Papadopoulou was funded by a Fulbright Visiting Faculty Award 2017. Stelios Smirnakis was supported through a Simons Foundation Research Award and R21 #NS096640 from NINDS.

References

- [1] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.
- [2] J. P. Cunningham and M. Y. Byron, "Dimensionality reduction for large-scale neural recordings," *Nature neuroscience*, vol. 17, no. 11, pp. 1500–1509, 2014.
- [3] G. Palagina, J. F. Meyer, and S. M. Smirnakis, "Complex visual motion representation in mouse area v1," *Journal of Neuroscience*, vol. 37, no. 1, pp. 164–183, 2017.
- [4] R. Pang, B. J. Lansdell, and A. L. Fairhall, "Dimensionality reduction in neuroscience," *Current Biology*, vol. 26, no. 14, pp. R656–R660, 2016.
- [5] I. Tosic and P. Frossard, "Dictionary learning," *IEEE Signal Processing Magazine*, vol. 28, no. 2, pp. 27–38, 2011.
- [6] R. G. Baraniuk, V. Cevher, and M. B. Wakin, "Low-dimensional models for dimensionality reduction and signal recovery: A geometric perspective," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 959–971, 2010.
- [7] K. Fotiadou, G. Tsagkatakis, and P. Tsakalides, "Low light image enhancement via sparse representations," in *International Conference Image Analysis and Recognition*. Springer International Publishing, 2014, pp. 84–93.
- [8] G. Tsagkatakis and P. Tsakalides, "Dictionary based reconstruction and classification of randomly sampled sensor network data," in *Proc. 7th IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM 12)*, 2012.
- [9] M. Aharon, M. Elad, and A. Bruckstein, "K-svd: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Transactions on signal processing*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [10] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *Information Theory, IEEE Transactions on*, vol. 53, no. 12, pp. 4655–4666, 2007.
- [11] M. D. Abràmoff, P. J. Magalhães, and S. J. Ram, "Image processing with imagej," *Biophotonics international*, vol. 11, no. 7, pp. 36–42, 2004.
- [12] C. Stosiek, O. Garaschuk, K. Holthoff, and A. Konnerth, "In vivo two-photon calcium imaging of neuronal networks," *Proceedings of the National Academy of Sciences*, vol. 100, no. 12, pp. 7319–7324, 2003.