



HAL
open science

Learning the Morphology of Brain Signals Using Alpha-Stable Convolutional Sparse Coding

Mainak Jas, Tom Dupré La Tour, Umut Şimşekli, Alexandre Gramfort

► **To cite this version:**

Mainak Jas, Tom Dupré La Tour, Umut Şimşekli, Alexandre Gramfort. Learning the Morphology of Brain Signals Using Alpha-Stable Convolutional Sparse Coding. *Advances in neural information processing systems*, Dec 2017, Long Beach, United States. hal-01590988

HAL Id: hal-01590988

<https://hal.science/hal-01590988v1>

Submitted on 20 Sep 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning the Morphology of Brain Signals Using Alpha-Stable Convolutional Sparse Coding

Mainak Jas^{*1}, Tom Dupré La Tour¹, Umut Şimşekli¹ and Alexandre Gramfort²

¹LTCI, Télécom ParisTech, Université Paris-Saclay, Paris, France

²INRIA, Université Paris Saclay, Saclay, France

Abstract

Neural time-series data contain a wide variety of prototypical signal waveforms (atoms) that are of significant importance in clinical and cognitive research. One of the goals for analyzing such data is hence to extract such ‘shift-invariant’ atoms. Even though some success has been reported with existing algorithms, they are limited in applicability due to their heuristic nature. Moreover, they are often vulnerable to artifacts and impulsive noise, which are typically present in raw neural recordings. In this study, we address these issues and propose a novel probabilistic convolutional sparse coding (CSC) model for learning shift-invariant atoms from raw neural signals containing potentially severe artifacts. In the core of our model, which we call α CSC, lies a family of heavy-tailed distributions called α -stable distributions. We develop a novel, computationally efficient Monte Carlo expectation-maximization algorithm for inference. The maximization step boils down to a weighted CSC problem, for which we develop a computationally efficient optimization algorithm. Our results show that the proposed algorithm achieves state-of-the-art convergence speeds. Besides, α CSC is significantly more robust to artifacts when compared to three competing algorithms: it can extract spike bursts, oscillations, and even reveal more subtle phenomena such as cross-frequency coupling when applied to noisy neural time series.

1 Introduction

Neural time series data, either non-invasive such as electroencephalography (EEG) or invasive such as electrocorticography (ECoG) and local field potentials (LFP), are fundamental to modern experimental neuroscience. Such recordings contain a wide variety of ‘prototypical signals’ that range from beta rhythms (12–30 Hz) in motor imagery tasks and alpha oscillations (8–12 Hz) involved in attention mechanisms, to spindles in sleep studies, and the classical P300 event related potential, a biomarker for surprise. These prototypical waveforms are considered critical in clinical and cognitive research [1], thereby motivating the development of computational tools for learning such signals from data.

Despite the underlying complexity in the morphology of neural signals, the majority of the computational tools in the community are based on representing the signals with rather simple, predefined bases, such as the Fourier or wavelet bases [2]. While such bases lead to computationally efficient algorithms, they often fall short at capturing the precise morphology of signal waveforms, as demonstrated by a number of recent studies [3, 4]. An example of such a failure is the disambiguation of the alpha rhythm from the mu rhythm [5], both of which have a component around 10 Hz but with different morphologies that cannot be captured by Fourier- or wavelet-based representations.

Recently, there have been several attempts for extracting more realistic and precise morphologies directly from unfiltered electrophysiology signals, via dictionary learning approaches [6–9]. These methods all aim to extract certain *shift-invariant* prototypical waveforms (called ‘atoms’ in this context) to better capture the temporal structure of the signals. As opposed to using generic bases that have predefined

*mainak.jas@telecom-paristech.fr

shapes, such as the Fourier or the wavelet bases, these atoms provide a more meaningful representation of the data and are not restricted to narrow frequency bands.

In this line of research, Jost et al. [6] proposed the MoTIF algorithm, which uses an iterative strategy based on generalized eigenvalue decompositions, where the atoms are assumed to be orthogonal to each other and learnt one by one in a greedy way. More recently, the ‘sliding window matching’ (SWM) algorithm [9] was proposed for learning time-varying atoms by using a correlation-based approach that aims to identify the recurring patterns. Even though some success has been reported with these algorithms, they have several limitations: SWM uses a slow stochastic search inspired by simulated annealing and MoTIF poorly handles correlated atoms, simultaneously activated, or having varying amplitudes; some cases which often occur in practical applications.

A natural way to cast the problem of learning a dictionary of shift-invariant atoms into an optimization problem is a convolutional sparse coding (CSC) approach [10]. This approach has gained popularity in computer vision [11–15], biomedical imaging [16] and audio signal processing [10, 17], due to its ability to obtain compact representations of the signals and to incorporate the temporal structure of the signals via convolution. In the neuroscience context, Barthélemy et al. [18] used an extension of the K-SVD algorithm using convolutions on EEG data. In a similar spirit, Brockmeier and Príncipe [7] used the matching pursuit algorithm combined with a rather heuristic dictionary update, which is similar to the MoTIF algorithm. In a very recent study, Hitziger et al. [8] proposed the AWL algorithm, which presents a mathematically more principled CSC approach for modeling neural signals. Yet, as opposed to classical CSC approaches, the AWL algorithm imposes additional combinatorial constraints, which limit its scope to certain data that contain spike-like atoms. Also, since these constraints increase the complexity of the optimization problem, the authors had to resort to dataset-specific initializations and many heuristics in their inference procedure.

While the current state-of-the-art CSC methods have a strong potential for modeling neural signals, they might also be limited as they consider an ℓ_2 reconstruction error, which corresponds to assuming an additive Gaussian noise distribution. While this assumption could be reasonable for several signal processing tasks, it turns out to be very restrictive for neural signals, which often contain heavy noise bursts and have low signal-to-noise ratio.

In this study, we aim to address the aforementioned concerns and propose a novel probabilistic CSC model called α CSC, which is better-suited for neural signals. α CSC is based on a family of *heavy-tailed* distributions called α -stable distributions [19] whose rich structure covers a broad range of noise distributions. The heavy-tailed nature of the α -stable distributions renders our model robust to impulsive observations. We develop a Monte Carlo expectation maximization (MCEM) algorithm for inference, with a weighted CSC model for the maximization step. We propose efficient optimization strategies that are specifically designed for neural time series. We illustrate the benefits of the proposed approach on both synthetic and real datasets.

2 Preliminaries

Notation: For a vector $v \in \mathbb{R}^n$ we denote the ℓ_p norm by $\|v\|_p = (\sum_i |v_i|^p)^{1/p}$. The convolution of two vectors $v_1 \in \mathbb{R}^N$ and $v_2 \in \mathbb{R}^M$ is denoted by $v_1 * v_2 \in \mathbb{R}^{N+M-1}$. We denote by x the observed signals, d the temporal atoms, and z the sparse vector of *activations*. The symbols \mathcal{U} , \mathcal{E} , \mathcal{N} , \mathcal{S} denote the univariate uniform, exponential, Gaussian, and α -stable distributions, respectively.

Convolutional sparse coding: The CSC problem formulation adopted in this work follows the Shift Invariant Sparse Coding (SISC) model from [10]. It is defined as follows:

$$\min_{d,z} \sum_{n=1}^N \left(\frac{1}{2} \|x_n - \sum_{k=1}^K d^k * z_n^k\|_2^2 + \lambda \sum_{k=1}^K \|z_n^k\|_1 \right), \quad \text{s.t.} \quad \|d^k\|_2^2 \leq 1 \text{ and } z_n^k \geq 0, \forall n, k, \quad (1)$$

where $x_n \in \mathbb{R}^T$ denotes one of the N observed segments of signals, also referred to as a *trials* in this paper. We denote by T as the length of a trial, and K the number of atoms. The aim in this model is to approximate the signals x_n by the convolution of certain *atoms* and their respective *activations*, which are

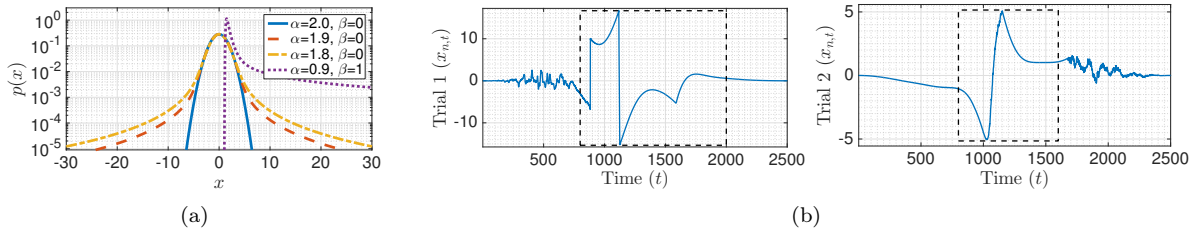


Figure 1: (a) PDFs of α -stable distributions. (b) Illustration of two trials from the striatal LFP data, which contain severe artifacts. The artifacts are illustrated with dashed rectangles.

sparse. Here, $d^k \in \mathbb{R}^L$ denotes the k th atom of the *dictionary* $d \equiv \{d^k\}_k$, and $z_n^k \in \mathbb{R}_+^{T-L+1}$ denotes the activation of the k th atom in the n th trial. We denote by $z \equiv \{z_n^k\}_{n,k}$.

The objective function (1) has two terms, an ℓ_2 data fitting term that corresponds to assuming an additive Gaussian noise model, and a regularization term that promotes sparsity with an ℓ_1 norm. The regularization parameter is called $\lambda > 0$. Two constraints are also imposed. First, we ensure that d^k lies within the unit sphere, which prevents the scale ambiguity between d and z . Second, a positivity constraint on z is imposed to be able to obtain physically meaningful activations and to avoid sign ambiguities between d and z . This positivity constraint is not present in the original SISC model [10].

α -Stable distributions: The α -stable distributions have become increasingly popular in modeling signals that might incur large variations [20–22] and have a particular importance in statistics since they appear as the limiting distributions in the generalized central limit theorem [19]. They are characterized by four parameters: α , β , σ , and μ : (i) $\alpha \in (0, 2]$ is the *characteristic exponent* and determines the tail thickness of the distribution: the distribution will be heavier-tailed as α gets smaller. (ii) $\beta \in [-1, 1]$ is the *skewness* parameter. If $\beta = 0$, the distribution is symmetric. (iii) $\sigma \in (0, \infty)$ is the *scale* parameter and measures the spread of the random variable around its mode. Finally, (iv) $\mu \in (-\infty, \infty)$ is the *location* parameter.

The probability density function of an α -stable distribution cannot be written in closed-form except for certain special cases; however, the characteristic function can be written as follows:

$$x \sim \mathcal{S}(\alpha, \beta, \sigma, \mu) \iff \mathbb{E}[\exp(i\omega x)] = \exp(-|\sigma\omega|^\alpha [1 + i \operatorname{sign}(\omega)\beta\psi_\alpha(\omega)] + i\mu\omega) ,$$

where $\psi_\alpha(\omega) = \log|\omega|$ for $\alpha = 1$, $\psi_\alpha(\omega) = \tan(\pi\alpha/2)$ for $\alpha \neq 1$, and $i = \sqrt{-1}$. As an important special case of the α -stable distributions, we obtain the Gaussian distribution when $\alpha = 2$ and $\beta = 0$, *i.e.* $\mathcal{S}(2, 0, \sigma, \mu) = \mathcal{N}(\mu, 2\sigma^2)$. In Fig. 1(a), we illustrate the (approximately computed) probability density functions (PDF) of the α -stable distribution for different values of α and β . The distribution becomes heavier-tailed as we decrease α , whereas the tails vanish quickly when $\alpha = 2$.

The moments of the α -stable distributions can only be defined up to the order α , *i.e.* $\mathbb{E}[|x|^p] < \infty$ if and only if $p < \alpha$, which implies the distribution has infinite variance when $\alpha < 2$. Furthermore, despite the fact that the PDFs of α -stable distributions do not admit an analytical form, it is straightforward to draw random samples from α -stable distributions [23].

3 Alpha-Stable Convolutional Sparse Coding

3.1 The Model

From a probabilistic perspective, the CSC problem can be also formulated as a maximum a-posteriori (MAP) estimation problem on the following probabilistic generative model:

$$z_{n,t}^k \sim \mathcal{E}(\lambda), \quad x_{n,t}|z, d \sim \mathcal{N}(\hat{x}_{n,t}, 1), \quad \text{where,} \quad \hat{x}_n \triangleq \sum_{k=1}^K d^k * z_n^k . \quad (2)$$

Here, $z_{n,t}^k$ denotes the t th element of z_n^k . We use the same notations for $x_{n,t}$ and $\hat{x}_{n,t}$. It is easy to verify that the MAP estimate for this probabilistic model, *i.e.* $\max_{d,z} \log p(d, z|x)$, is identical to the original optimization problem defined in (1).

It has been long known that, due to their light-tailed nature, Gaussian models often fail at handling noisy high amplitude observations or outliers [24]. As a result, the ‘vanilla’ CSC model turns out to be highly sensitive to outliers and impulsive noise that frequently occur in electrophysiological recordings, as illustrated in Fig. 1(b). Possible origins of such artifacts are movement, muscle contractions, ocular blinks or electrode contact losses.

In this study, we aim at developing a probabilistic CSC model that would be capable of modeling challenging electrophysiological signals. We propose an extension of the original CSC model defined in (2) by replacing the light-tailed Gaussian likelihood with heavy-tailed α -stable distributions. We define the proposed probabilistic model (α CSC) as follows:

$$z_{n,t}^k \sim \mathcal{E}(\lambda), \quad x_{n,t}|z, d \sim \mathcal{S}(\alpha, 0, 1/\sqrt{2}, \hat{x}_{n,t}), \quad (3)$$

where \mathcal{S} denotes the α -stable distribution. While still being able to capture the temporal structure of the observed signals via convolution, the proposed model has a richer structure and would allow large variations and outliers, thanks to the heavy-tailed α -stable distributions. Note that the vanilla CSC defined in (2) appears as a special case of α CSC, as the α -stable distribution coincides with the Gaussian distribution when $\alpha = 2$.

3.2 Maximum A-Posteriori Inference

Given the observed signals x , we are interested in the MAP estimates, defined as follows:

$$(d^*, z^*) = \arg \max_{d,z} \sum_{n,t} \left(\log p(x_{n,t}|d, z) + \sum_k \log p(z_{n,t}^k) \right). \quad (4)$$

As opposed to the Gaussian case, unfortunately, this optimization problem is not amenable to classical optimization tools, since the PDF of the α -stable distributions does not admit an analytical expression. As a remedy, we use the product property of the symmetric α -stable densities [19, 25] and re-express the α CSC model as conditionally Gaussian. It leads to:

$$z_{n,t}^k \sim \mathcal{E}(\lambda), \quad \phi_{n,t} \sim \mathcal{S}\left(\frac{\alpha}{2}, 1, 2(\cos \frac{\pi\alpha}{4})^{2/\alpha}, 0\right), \quad x_{n,t}|z, d, \phi \sim \mathcal{N}\left(\hat{x}_{n,t}, \frac{1}{2}\phi_{n,t}\right), \quad (5)$$

where ϕ is called the *impulse* variable that is drawn from a *positive* α -stable distribution (*i.e.* $\beta = 1$), whose PDF is illustrated in Fig. 1(a). It can be easily shown that both formulations of the α CSC model are identical by marginalizing the joint distribution $p(x, d, z, \phi)$ over ϕ .

The impulsive structure of the α CSC model becomes more prominent in this formulation: the variances of the Gaussian observations are modulated by stable random variables with infinite variance, where the impulsiveness depends on the value of α . It is also worth noting that when $\alpha = 2$, $\phi_{n,t}$ becomes deterministic and we can again verify that α CSC coincides with the vanilla CSC.

The conditionally Gaussian structure of the augmented model has a crucial practical implication: if the impulse variable ϕ were to be known, then the MAP estimation problem over d and z in this model would turn into a ‘weighted’ CSC problem, which is a much easier task compared to the original problem. In order to be able to exploit this property, we propose an expectation-maximization (EM) algorithm, which iteratively maximizes a lower bound of the log-posterior $\log p(d, z|x)$, and algorithmically boils down to computing the following steps in an iterative manner:

$$\text{E-Step:} \quad \mathcal{B}^{(i)}(d, z) = \mathbb{E}[\log p(x, \phi, z|d)]_{p(\phi|x, z^{(i)}, d^{(i)})}, \quad (6)$$

$$\text{M-Step:} \quad (d^{(i+1)}, z^{(i+1)}) = \arg \max_{d,z} \mathcal{B}^{(i)}(d, z). \quad (7)$$

where $\mathbb{E}[f(x)]_{q(x)}$ denotes the expectation of a function f under the distribution q , i denotes the iterations, and $\mathcal{B}^{(i)}$ is a lower bound to $\log p(d, z|x)$ and it is tight at the current iterates $z^{(i)}, d^{(i)}$.

The E-Step: In the first step of our algorithm, we need to compute the EM lower bound \mathcal{B} that has the following form:

$$\mathcal{B}^{(i)}(d, z) = + - \sum_{n=1}^N \left(\|\sqrt{w_n^{(i)}} \odot (x_n - \sum_{k=1}^K d^k * z_n^k)\|_2^2 + \lambda \sum_{k=1}^K \|z_n^k\|_1 \right), \quad (8)$$

where $=^+$ denotes equality up to additive constants, \odot denotes the Hadamard (element-wise) product, and the square-root operator is also defined element-wise. Here, $w_n^{(i)} \in \mathbb{R}_+^T$ are the *weights* that are defined as follows: $w_{n,t}^{(i)} \triangleq \mathbb{E}[1/\phi_{n,t}]_{p(\phi|x,z^{(i)},d^{(i)})}$. As the variables $\phi_{n,t}$ are expected to be large when $\hat{x}_{n,t}$ cannot explain the observation $x_{n,t}$ – typically due to a corruption or a high noise – the weights will accordingly suppress the importance of the particular point $x_{n,t}$. Therefore, the overall approach will be more robust to corrupted data than the Gaussian models where all weights would be deterministic and equal to 0.5.

Unfortunately, the weights $w^{(i)}$ cannot be computed analytically, therefore we need to resort to approximate methods. In this study, we develop a Markov chain Monte Carlo (MCMC) method to approximately compute the weights, where we approximate the intractable expectations with a finite sample average, given as follows: $w_{n,t}^{(i)} \approx (1/J) \sum_{j=1}^J 1/\phi_{n,t}^{(i,j)}$, where $\phi_{n,t}^{(i,j)}$ are some samples that are ideally drawn from the posterior distribution $p(\phi|x, z^{(i)}, d^{(i)})$. Unfortunately, directly drawing samples from the posterior distribution of ϕ is not tractable either, and therefore, we develop a *Metropolis-Hastings* algorithm [26], that asymptotically generates samples from the *target* distribution $p(\phi|\cdot)$ in two steps. In the j -th iteration of this algorithm, we first draw a random sample for each n and t from the prior distribution (cf. (5)), *i.e.*, $\phi'_{n,t} \sim p(\phi_{n,t})$.

We then compute an acceptance probability for each $\phi'_{n,t}$ that is defined as follows:

$$\text{acc}(\phi_{n,t}^{(i,j)} \rightarrow \phi'_{n,t}) \triangleq \min \left\{ 1, p(x_{n,t}|d^{(i)}, z^{(i)}, \phi'_{n,t})/p(x_{n,t}|d^{(i)}, z^{(i)}, \phi_{n,t}^{(i,j)}) \right\} \quad (9)$$

where j denotes the iteration number of the MCMC algorithm. Finally, we draw a uniform random number $u_{n,t} \sim \mathcal{U}([0, 1])$ for each n and t . If $u_{n,t} < \text{acc}(\phi_{n,t}^{(i,j)} \rightarrow \phi'_{n,t})$, we accept the sample and set $\phi_{n,t}^{(i+1)} = \phi'_{n,t}$; otherwise we reject the sample and set $\phi_{n,t}^{(i+1)} = \phi_{n,t}^{(i)}$. This procedure forms a Markov chain that leaves the target distribution $p(\phi|\cdot)$ invariant, where under mild ergodicity conditions, it can be shown that the finite-sample averages converge to their true values when J goes to infinity [27]. More detailed explanation of this procedure is given in the supplementary document.

The M-Step: Given the weights w_n that are estimated during the E-step, the objective of the M-step (7) is to solve a weighted CSC problem, which is much easier when compared to our original problem. This objective function is not jointly convex in d and z , yet it is convex if one fix either d or z . Here, similarly to the vanilla CSC approaches [9, 10], we develop a *block coordinate descent* strategy, where we solve the problem in (7) for either d or z , by keeping respectively z and d fixed. We first focus on solving the problem for z while keeping d fixed, given as follows:

$$\min_z \sum_{n=1}^N \left(\|\sqrt{w_n} \odot (x_n - \sum_{k=1}^K D^k \bar{z}_n^k)\|_2^2 + \lambda \sum_k \|z_n^k\|_1 \right) \quad \text{s.t. } z_n^k \geq 0, \forall n, k. \quad (10)$$

Here, we expressed the convolution of d^k and z_n^k as the inner product of the zero-padded activations $\bar{z}_n^k \triangleq [(z_n^k)^\top, 0 \dots 0]^\top \in \mathbb{R}_+^T$, with a Toeplitz matrix $D^k \in \mathbb{R}^{T \times T}$, that is constructed from d^k . The

Algorithm 1 α -stable Convolutional Sparse Coding

Require: Regularization: $\lambda \in \mathbb{R}_+$, Num. atoms: K ,
Atom length: L , Num. iterations: I, J, M

- 1: **for** $i = 1$ to I **do**
- 2: /* E-step: */
- 3: **for** $j = 1$ to J **do**
- 4: Draw $\phi_{n,t}^{(i,j)}$ via MCMC (9)
- 5: **end for**
- 6: $w_{n,t}^{(i)} \approx (1/J) \sum_{j=1}^J 1/\phi_{n,t}^{(i,j)}$
- 7: /* M-step: */
- 8: **for** $m = 1$ to M **do**
- 9: $z^{(i)} = \text{L-BFGS-B}$ on (10)
- 10: $d^{(i)} = \text{L-BFGS-B}$ on the dual of (11)
- 11: **end for**
- 12: **end for**
- 13: **return** $w^{(I)}, d^{(I)}, z^{(I)}$

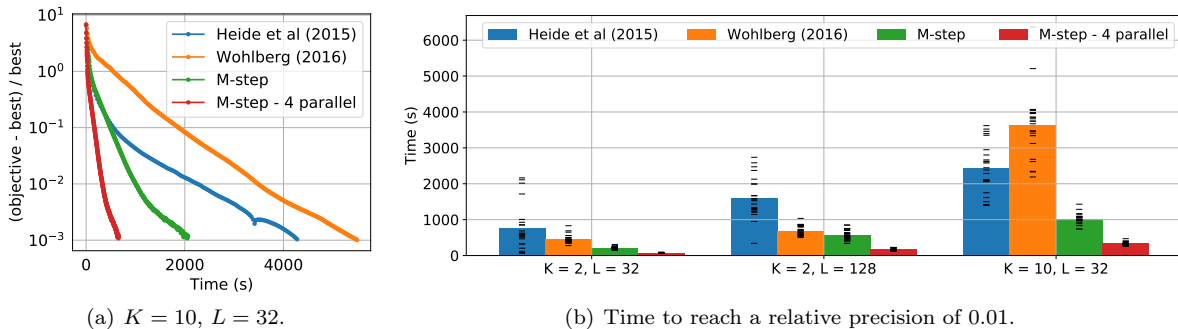


Figure 2: Comparison of state-of-the-art methods with our approach. (a) Convergence plot with the objective function relative to the obtained minimum, as a function of computational time. (b) Time taken to reach a relative precision of 10^{-2} , for different settings of K and L .

matrices D^k are never constructed in practice, and all operations are carried out using convolutions. This problem can be solved by various constrained optimization algorithms. Here, we choose the quasi-Newton L-BFGS-B algorithm [28] with a box constraint: $0 \leq z_{n,t}^k \leq \infty$. This approach only requires the simple computation of the gradient of the objective function with respect to z (*cf.* supplementary material). Note that, since each trial is independent from each other, we can solve this problem for each z_n in parallel.

We then solve the problem for the atoms d while keeping z fixed. This optimization problem turns out to be a constrained weighted least-squares problem. In the non-weighted case, this problem can be solved either in the time domain or in the Fourier domain [10–12]. The Fourier transform simplifies the convolutions that appear in least-squares problem, but it also induces several difficulties, such as that the atom d_k have to be in a finite support L , an important issue ignored in the seminal work of [10] and addressed with an ADMM solver in [11, 12]. In the weighted case, it is not clear how to solve this problem in the Fourier domain. We thus perform all the computations in the time domain.

Following the traditional filter identification approach [29], we need to embed the one-dimensional signals z_n^k into a matrix of delayed signals $Z_n^k \in \mathbb{R}^{T \times L}$, where $(Z_n^k)_{i,j} = z_{n,i+j-L+1}^k$ if $L-1 \leq i+j < T$ and 0 elsewhere. Equation (1) then becomes:

$$\min_d \sum_{n=1}^N \|\sqrt{w_n} \odot (x_n - \sum_{k=1}^K Z_n^k d^k)\|_2^2, \quad \text{s.t. } \|d^k\|_2^2 \leq 1. \quad (11)$$

Due to the constraint, we must resort to an iterative approach. The options are to use (accelerated) projected gradient methods such as FISTA [30] applied to (11), or to solve a dual problem as done in [10]. The dual is also a smooth constraint problem yet with a simpler positivity box constraint (*cf.* supplementary material). The dual can therefore be optimized with L-BFGS-B. Using such a quasi-Newton solver turned out to be more efficient than any accelerated first order method in either the primal or the dual (*cf.* benchmarks in supplementary material).

Our entire EM approach can be summarized in the Algorithm 1. Note that during the alternating minimization, thanks to convexity we can warm start the d update and the z update using the solution from the previous update. This significantly speeds up the convergence of the L-BFGS-B algorithm, particularly in the later iterations of the overall algorithm.

4 Experiments

In order to evaluate our approach, we conduct several experiments on both synthetic and real data. First, we show that our proposed optimization scheme for the M-step provides significant improvements in terms of convergence speed over the state-of-the-art CSC methods. Then, we provide empirical evidence that our algorithm is more robust to artifacts and outliers than three competing CSC methods [6, 7, 12].

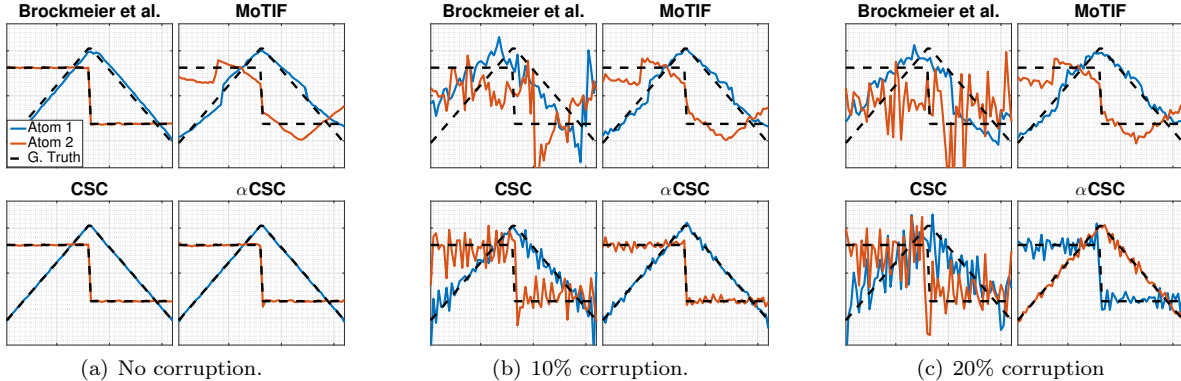


Figure 3: Simulation to compare state-of-the-art methods against α CSC.

Finally, we consider LFP data, where we illustrate that our algorithm can reveal interesting properties in electrophysiological signals without supervision, even in the presence of severe artifacts.

Synthetic simulation setup: In our synthetic data experiments, we simulate N trials of length T by first generating K zero mean and unit norm atoms of length L . The activation instants are integers drawn from a uniform distribution in $\llbracket 0, T - L \rrbracket$. The amplitude of the activations are drawn from a uniform distribution in $[0, 1]$. Atoms are activated only once per trial and are allowed to overlap. The activations are then convolved with the generated atoms and summed up as in (1).

M-step performance: In our first set of synthetic experiments, we illustrate the benefits of our M-step optimization approach over state-of-the-art CSC solvers. We set $N = 100$, $T = 2000$ and $\lambda = 1$, and use different values for K and L . To be comparable, we set $\alpha = 2$ and add Gaussian noise to the synthesized signals, where the standard deviation is set to 0.01. In this setting, we have $w_{n,t} = 1/2$ for all n, t , which reduces the problem to a standard CSC setup. We monitor the convergence of ADMM-based methods by Heide et al. [11] and Wohlberg [12] against our M-step algorithm, using both a single-threaded and a parallel version for the z -update. As the problem is non-convex, even if two algorithms start from the same point, they are not guaranteed to reach the same local minimum. Hence, for a fair comparison, we use a multiple restart strategy with averaging across 24 random seeds.

During our experiments we have observed that the ADMM-based methods do not guarantee the feasibility of the iterates. In other words, the norms of the estimated atoms might be greater than 1 during the iterations. To keep the algorithms comparable, when computing the objective value, we project the atoms to the unit ball and scale the activations accordingly. To be strictly comparable, we also imposed a positivity constraint on these algorithms. This is easily done by modifying the soft-thresholding operator to be a rectified linear function. In the benchmarks, all algorithms use a single thread, except “M-step - 4 parallel” which uses 4 threads during the z update.

In Fig. 2, we illustrate the convergence behaviors of the different methods. Note that the y-axis is the precision relative to the objective value obtained upon convergence. In other words, each curve is relative to its own local minimum (see supplementary document for details). In the right subplot, we show how long it takes for the algorithms to reach a relative precision of 0.01 for different settings (*cf.* supplementary material for more benchmarks). Our method consistently performs better and the difference is even more striking for more challenging setups. This speed improvement on the M-step is crucial for us as this step will be repeatedly executed.

Robustness to corrupted data: In our second synthetic data experiment, we illustrate the robustness of α CSC in the presence of corrupted observations. In order to simulate the likely presence of high amplitude artifacts, one way would be to directly simulate the generative model in (3). However, this would give us an unfair advantage, since α CSC is specifically designed for such data. Here, we take an alternative approach, where we corrupt a randomly chosen fraction of the trials (10% or 20%) with strong Gaussian noise of standard deviation 0.1, *i.e.* one order of magnitude higher than in a regular trial. We used a regularization parameter of $\lambda = 0.1$. In these experiments, by CSC we refer to α CSC with $\alpha = 2$, that resembles using only the M-step of our algorithm with deterministic weights $w_{n,t} = 1/2$ for all n, t . We used a simpler setup where we set $N = 100$, $T = 512$, and $L = 64$. We used $K = 2$ atoms, as shown in

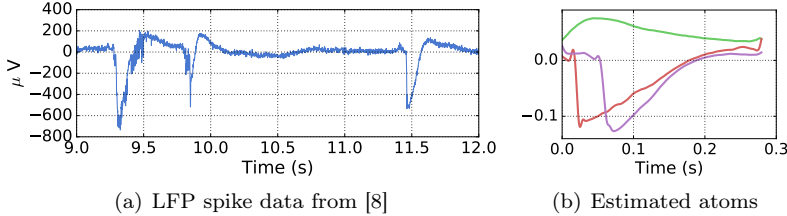


Figure 4: Atoms learnt by α CSC on LFP data containing epileptiform spikes with $\alpha = 2$.

dashed lines in Fig. 3.

For α CSC, we set the number of outer iterations $I = 5$, the number of iterations of the M-step to $M = 50$, and the number of iterations of the MCMC algorithm to $J = 10$. We discard the first 5 samples of the MCMC algorithm as burn-in. To enable a fair comparison, we run the standard CSC algorithm for $I \times M$ iterations, i.e. the *total* number of M-step iterations in α CSC. We also compared α CSC against competing state-of-art methods previously applied to neural time series: Brockmeier and Príncipe [7] and MoTIF [6]. Starting from multiple random initializations, the estimated atoms with the smallest ℓ_2 distance with the true atoms are shown in Fig. 3.

In the artifact-free scenario, all algorithms perform equally well, except for MoTIF that suffers from the presence of activations with varying amplitudes. This is because it aligns the data using correlations before performing the eigenvalue decomposition, without taking into account the strength of activations in each trial. The performance of Brockmeier and Príncipe [7] and CSC degrades as the level of corruption increases. On the other hand, α CSC is clearly more robust to the increasing level of corruption and recovers reasonable atoms even when 20% of the trials are corrupted.

Results on LFP data In our last set of experiments, we consider real neural data from two different datasets. We first applied α CSC on an LFP dataset previously used in [8] and containing epileptiform spikes as shown in Fig. 4(a). The data was recorded in the rat cortex, and is free of artifact. Therefore, we used the standard CSC with our optimization scheme, (i.e. α CSC with $\alpha = 2$). As a standard preprocessing procedure, we applied a high-pass filter at 1 Hz in order to remove drifts in the signal, and then applied a tapered cosine window to down-weight the samples near the edges. We set $\lambda = 6$, $N = 300$, $T = 2500$, $L = 350$, and $K = 3$. The recovered atoms by our algorithm are shown in Fig. 4(b). We can observe that the estimated atoms resemble the spikes in Fig. 4(a). These results show that, without using any heuristics, our approach can recover similar atoms to the ones reported in [8], even though it does not make any assumptions on the shapes of the waveforms, or initializes the atoms with template spikes in order to ease the optimization.

The second dataset is an LFP channel in a rodent striatum from [31]. We segmented the data into 70 trials of length 2500 samples, windowed each trial with a tapered cosine function, and detrended the data with a high-pass filter at 1 Hz. We set $\lambda = 10$, initialized the weights w_n to the inverse of the variance of the trial x_n . Atoms are in all experiments initialized with Gaussian white noise.

As opposed to the first LFP dataset, this dataset contains strong artifacts, as shown in Fig. 1(b). In order to be able to illustrate the potential of CSC on this data, we first *manually* identified and removed the trials that were corrupted by artifacts. In Fig. 5(a), we illustrate the estimated atoms with CSC on the manually-cleaned data. We observe that the estimated atoms correspond to canonical waveforms found in the signal. In particular, the high frequency oscillations around 80 Hz are modulated in amplitude by the low-frequency oscillation around 3 Hz, a phenomenon known as cross-frequency coupling (CFC) [32]. We can observe this by computing a comodulogram [33] on the entire signal (Fig. 5(b)). This measures the correlation between the amplitude of the high frequency band and the phase of the low frequency band.

Even though CSC is able to provide these excellent results on the cleaned data set, its performance heavily relies on the manual removal of the artifacts. Finally, we repeated the previous experiment on the full data, without removing the artifacts and compared CSC with α CSC, where we set $\alpha = 1.2$. The results are shown in the middle and the right sub-figures of Fig. 5(a). It can be observed that in the presence of strong artifacts, CSC is not able to recover the atoms anymore. On the contrary, we observe that α CSC can still recover atoms as observed in the artifact-free regime. In particular, the cross-frequency coupling phenomenon is still visible.

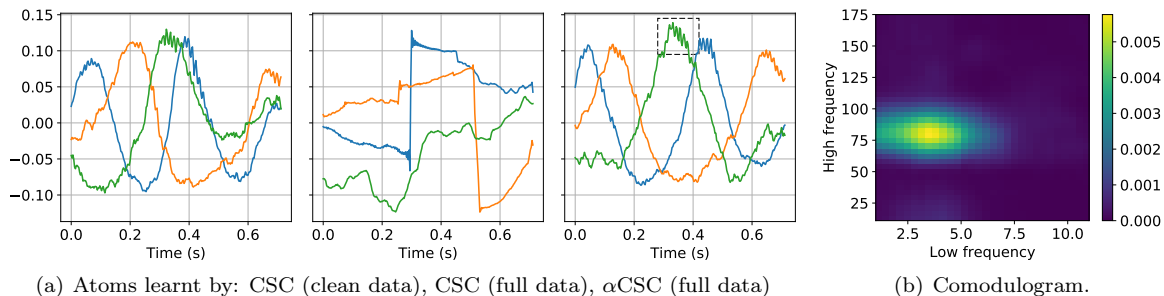


Figure 5: (a) Three atoms learnt from a rodent striatal LFP channel, using CSC on cleaned data, and both CSC and α CSC on the full data. The atoms capture the cross-frequency coupling of the data (dashed rectangle). (b) Comodulogram presents the cross-frequency coupling intensity computed between pairs of frequency bands on the entire cleaned signal, following [33].

5 Conclusion

This work addresses the present need in the neuroscience community to better capture the complex morphology of brain waves [1, 9]. Our data-driven approach to this problem is a probabilistic formulation of a convolutional sparse coding model [10]. We propose an inference strategy based on Monte Carlo EM to deal efficiently with heavy tailed noise and take into account the polarity of neural activations with a positivity constraint. Our problem formulation allows the use of fast quasi-Newton methods for the M-step which outperform previously proposed state-of-the-art ADMM-based algorithms [11, 12, 34, 35], even when not making use of our parallel implementation. Results on LFP data demonstrate that such algorithms can be robust to the presence of transient artifacts in data and reveal insights on neural time-series without supervision. We will make our code publicly available.

6 Acknowledgement

The work was supported by the French National Research Agency grants ANR-14-NEUC-0002-01 and ANR-16-CE23-0014 (FBIMATRIX), as well as the ERC Starting Grant SLAB ERC-YSstG-676943.

References

- [1] S. R. Cole and B. Voytek. Brain oscillations and the importance of waveform shape. *Trends Cogn. Sci.*, 2017.
- [2] M. X. Cohen. *Analyzing neural time series data: Theory and practice*. MIT Press, 2014. ISBN 9780262319560.
- [3] S. R. Jones. When brain rhythms aren’t ‘rhythmic’: implication for their mechanisms and meaning. *Curr. Opin. Neurobiol.*, 40:72–80, 2016.
- [4] A. Mazaheri and O. Jensen. Asymmetric amplitude modulations of brain oscillations generate slow evoked responses. *The Journal of Neuroscience*, 28(31):7781–7787, 2008.
- [5] R. Hari and A. Puce. *MEG-EEG Primer*. Oxford University Press, 2017.
- [6] P. Jost, P. Vandergheynst, S. Lesage, and R. Gribonval. MoTIF: an efficient algorithm for learning translation invariant dictionaries. In *Acoustics, Speech and Signal Processing, ICASSP*, volume 5. IEEE, 2006.
- [7] A. J. Brockmeier and J. C. Príncipe. Learning recurrent waveforms within EEGs. *IEEE Transactions on Biomedical Engineering*, 63(1):43–54, 2016.

- [8] S. Hitziger, M. Clerc, S. SAILLET, C. Benar, and T. Papadopoulo. Adaptive Waveform Learning: A Framework for Modeling Variability in Neurophysiological Signals. *IEEE Transactions on Signal Processing*, 2017.
- [9] B. Gips, A. Bahramisharif, E. Lowet, M. Roberts, P. de Weerd, O. Jensen, and J. van der Eerden. Discovering recurring patterns in electrophysiological recordings. *J. Neurosci. Methods*, 275:66–79, 2017.
- [10] R. Grosse, R. Raina, H. Kwong, and A. Y. Ng. Shift-invariant sparse coding for audio classification. In *23rd Conference on Uncertainty in Artificial Intelligence*, UAI’07, pages 149–158. AUAI Press, 2007. ISBN 0-9749039-3-0.
- [11] F. Heide, W. Heidrich, and G. Wetzstein. Fast and flexible convolutional sparse coding. In *Computer Vision and Pattern Recognition (CVPR)*, pages 5135–5143. IEEE, 2015.
- [12] B. Wohlberg. Efficient algorithms for convolutional sparse representations. *Image Processing, IEEE Transactions on*, 25(1):301–315, 2016.
- [13] M. D. Zeiler, D. Krishnan, G.W. Taylor, and R. Fergus. Deconvolutional networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2528–2535. IEEE, 2010.
- [14] M. Šorel and F. Šroubek. Fast convolutional sparse coding using matrix inversion lemma. *Digital Signal Processing*, 2016.
- [15] K. Kavukcuoglu, P. Sermanet, Y-L. Boureau, K. Gregor, M. Mathieu, and Y. Cun. Learning convolutional feature hierarchies for visual recognition. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1090–1098, 2010.
- [16] M. Pachitariu, A. M Packer, N. Pettit, H. Dalgleish, M. Hausser, and M. Sahani. Extracting regions of interest from biological images with convolutional sparse block coding. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1745–1753, 2013.
- [17] B. Mailhé, S. Lesage, R. Gribonval, F. Bimbot, and P. Vandergheynst. Shift-invariant dictionary learning for sparse representations: extending K-SVD. In *16th Eur. Signal Process. Conf.*, pages 1–5. IEEE, 2008.
- [18] Q. Barthélemy, C. Gouy-Pailler, Y. Isaac, A. Souloumiac, A. Larue, and J. I. Mars. Multivariate temporal dictionary learning for EEG. *J. Neurosci. Methods*, 215(1):19–28, 2013.
- [19] G. Samorodnitsky and M. S. Taqqu. *Stable non-Gaussian random processes: stochastic models with infinite variance*, volume 1. CRC press, 1994.
- [20] E. E. Kuruoglu. *Signal processing in α -stable noise environments: a least L_p -norm approach*. PhD thesis, University of Cambridge, 1999.
- [21] B. B. Mandelbrot. *Fractals and scaling in finance: Discontinuity, concentration, risk. Selecta volume E*. Springer Science & Business Media, 2013.
- [22] Y. Wang, Y. Qi, Y. Wang, Z. Lei, X. Zheng, and G. Pan. Delving into α -stable distribution in noise suppression for seizure detection from scalp EEG. *J. Neural. Eng.*, 13(5):056009, 2016.
- [23] J. M. Chambers, C. L. Mallows, and B. W. Stuck. A method for simulating stable random variables. *Journal of the american statistical association*, 71(354):340–344, 1976.
- [24] P. J. Huber. *Robust Statistics*. Wiley, 1981.
- [25] S. Godsill and E. Kuruoglu. Bayesian inference for time series with heavy-tailed symmetric α -stable noise processes. *Proc. Applications of heavy tailed distributions in economics, eng. and stat.*, 1999.

- [26] S. Chib and E. Greenberg. Understanding the Metropolis-Hastings algorithm. *The American Statistician*, 49(4):327–335, 1995.
- [27] J.S. Liu. *Monte Carlo strategies in scientific computing*. Springer, 2008.
- [28] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995.
- [29] E. Moulines, P. Duhamel, J-F. Cardoso, and S. Mayrargue. Subspace methods for the blind identification of multichannel FIR filters. *IEEE Transactions on signal processing*, 43(2):516–525, 1995.
- [30] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- [31] G. Dallérac, M. Graupner, J. Knippenberg, R. C. R. Martinez, T. F. Tavares, L. Tallot, N. El Massioui, A. Verschueren, S. Höhn, J.B. Bertolus, et al. Updating temporal expectancy of an aversive event engages striatal plasticity under amygdala control. *Nature Communications*, 8:13920, 2017.
- [32] O. Jensen and L. L. Colgin. Cross-frequency coupling between neuronal oscillations. *Trends in cognitive sciences*, 11(7):267–269, 2007.
- [33] A. BL. Tort, R. Komorowski, H. Eichenbaum, and N. Kopell. Measuring phase-amplitude coupling between neuronal oscillations of different frequencies. *J. Neurophysiol.*, 104(2):1195–1210, 2010.
- [34] B. Wohlberg. Efficient convolutional sparse coding. In *Acoustics, Speech and Signal Processing, ICASSP*, pages 7173–7177. IEEE, 2014.
- [35] H. Bristow, A. Eriksson, and S. Lucey. Fast convolutional sparse coding. In *Computer Vision and Pattern Recognition (CVPR)*, pages 391–398, 2013.

A Details of the E-Step

Computing the weights that are required in the M-step requires us to compute the expectation of $\frac{1}{\phi_{n,t}}$ under the posterior distribution $p(\phi_{n,t}|x, d, z)$, which is not analytically available.

Monte Carlo methods are numerical techniques that can be used to approximately compute the expectations of the form:

$$\mathbb{E}[f(\phi_{n,t})] = \int f(\phi_{n,t})\pi(\phi_{n,t})d\phi_{n,t} \approx \frac{1}{J} \sum_{j=1}^J f(\phi_{n,t}^{(j)}) \quad (12)$$

where $\phi_{n,t}^{(j)}$ are some samples drawn from $\pi(\phi_{n,t}) \triangleq p(\phi_{n,t}|x, d, z)$ and $f(\phi) = 1/\phi$ in our case. However, in our case, sampling directly from $\pi(\phi_{n,t})$ is also unfortunately intractable.

MCMC methods generate samples from the target distribution $\pi(\phi_{n,t})$ by forming a Markov chain, whose stationary distribution is $\pi(\phi_{n,t})$, so that $\pi(\phi_{n,t}) = \int \mathcal{T}(\phi_{n,t}|\phi'_{n,t})p(\phi'_{n,t})d\phi'_{n,t}$, where \mathcal{T} denotes the transition kernel of the Markov chain.

In this study, we develop a Metropolis-Hastings (MH) algorithm, that implicitly forms a transition kernel. The MH algorithm generates samples from a target distribution $\pi(\phi_{n,t})$ in two steps. First, it generates a random sample $\phi'_{n,t}$ from a *proposal* distribution $\phi'_{n,t} \sim q(\phi'_{n,t}|\phi_{n,t}^{(j)})$, then computes an acceptance probability $\text{acc}(\phi_{n,t}^{(j)} \rightarrow \phi'_{n,t})$ and draws a uniform random number $u \sim \mathcal{U}([0, 1])$. If $u < \text{acc}(\phi_{n,t}^{(j)} \rightarrow \phi'_{n,t})$, it accepts the sample and sets $\phi_{n,t}^{(j+1)} = \phi'_{n,t}$; otherwise it rejects the sample and sets $\phi_{n,t}^{(j+1)} = \phi_{n,t}^{(j)}$. The acceptance probability is given as follows

$$\text{acc}(\phi_{n,t} \rightarrow \phi'_{n,t}) = \min\left\{1, \frac{q(\phi_{n,t}|\phi'_{n,t})\pi(\phi'_{n,t})}{q(\phi'_{n,t}|\phi_{n,t})\pi(\phi_{n,t})}\right\} = \min\left\{1, \frac{q(\phi_{n,t}|\phi'_{n,t})p(x_{n,t}|\phi'_{n,t}, d, z)p(\phi'_{n,t})}{q(\phi'_{n,t}|\phi_{n,t})p(x_{n,t}|\phi_{n,t}, d, z)p(\phi_{n,t})}\right\} \quad (13)$$

where the last equality is obtained by applying the Bayes rule on π .

The acceptance probability requires the prior distribution of ϕ to be evaluated. Unfortunately, this is intractable in our case since this prior distribution is chosen to be a positive α -stable distribution whose PDF does not have an analytical form. As a remedy, we choose the prior distribution of $\phi_{n,t}$ as the proposal distribution, such $q(\phi_{n,t}|\phi'_{n,t}) = p(\phi_{n,t})$. This enables us to simplify the acceptance probability. Accordingly, for each $\phi_{n,t}$, we have the following acceptance probability:

$$\text{acc}(\phi_{n,t}^{(i,j)} \rightarrow \phi'_{n,t}) \triangleq \min\left\{1, \exp(\log \phi_{n,t}^{(i,j)} - \log \phi'_{n,t})/2 + (x_{n,t} - \hat{x}_{n,t}^{(i)})^2(1/\phi_{n,t}^{(i,j)} - 1/\phi'_{n,t})\right\}. \quad (14)$$

Thanks to the simplification, this probability is tractable and can be easily computed.

B Details of the M-Step

B.1 Solving for the activations

In the M-step, we optimize (10) to find the activations $z_n^{(i)}$ of each trial n independently. To keep the notation simple, we will drop the index for the iteration number i of the EM algorithm.

First, this equation can be rewritten by concatenating the Toeplitz matrices for the K atoms into a big matrix $D = [D^1, D^2, \dots, D^K] \in \mathbb{R}^{T \times KT}$ and the activations for different atoms into a single vector $\bar{z}_n = [(\bar{z}_n^1)^\top, (\bar{z}_n^2)^\top, \dots, (\bar{z}_n^K)^\top]^\top \in \mathbb{R}_+^{KT}$ where $(\cdot)^\top$ denotes the transposition operation. Recall that \bar{z}_n^k is a zero-padded version of z_n^k . This leads to a simpler formulation and the objective function $\mathcal{L}(d, z)$:

$$\mathcal{L}(d, z) = \sum_{n=1}^N \frac{1}{2} \|\sqrt{w_n} \odot (x_n - D\bar{z}_n)\|_2^2 + \lambda \mathbb{1}^\top \bar{z}_n, \quad (15)$$

where $\mathbb{1} \in \mathbb{R}^{KT}$ is a vector of ones.

The derivative w.r.t. z_n now reads:

$$\frac{\partial \mathcal{L}(d, z)}{\partial \bar{z}_n} = D^\top (w_n \odot (x_n - D\bar{z}_n)) + \lambda \mathbf{1}^\top . \quad (16)$$

In practice, this big matrix D is never assembled and all operations are carried out using convolutions. Note also that we do not update the zeros from the padding in \bar{z}_n^k . Now that we have the gradient, the activations can be estimated using an efficient quasi-Newton solver such as L-BFGS-B, taking into account the box positivity constraint $0 \leq z_n \leq \infty$.

For each trial, one iteration costs $\mathcal{O}(LKT)$.

B.2 Solving for the atoms

In the M-step, we optimize (11) to find the atoms d^k . As when solving for the activations z_n , we can remove the summation over the atoms by concatenating the delayed matrices into $Z_n = [Z_n^1, Z_n^2, \dots, Z_n^K] \in \mathbb{R}^{T \times KL}$ and $d = [(d^1)^\top, (d^2)^\top, \dots, (d^K)^\top]^\top \in \mathbb{R}^{KL}$. This leads to the simpler formulation:

$$\min_d \sum_{n=1}^N \frac{1}{2} \|\sqrt{w_n} \odot (x_n - Z_n d)\|_2^2, \quad \text{s.t. } \|d^k\|_2^2 \leq 1 . \quad (17)$$

The Lagrangian of this problem is given by:

$$g(d, \beta) = \sum_{n=1}^N \frac{1}{2} \|\sqrt{w_n} \odot (x_n - \sum_{k=1}^K Z_n^k d^k)\|_2^2 + \sum_k \beta^k (\|d^k\|_2^2 - 1) \quad \text{s.t. } \beta^k \geq 0 , \quad (18)$$

where $\beta = (\beta^1, \beta^2, \dots, \beta^K)$ are the dual variables. Therefore, the dual problem is:

$$\min_d g(d, \beta) = g(d^*, \beta) \quad (19)$$

where d^* , the primal optimal, is given by:

$$d^* = \left(\sum_{n=1}^N Z_n^\top (w_n \odot Z_n) + \bar{\beta} \right)^{-1} \sum_{n=1}^N (w_n \odot Z_n)^\top x_n \quad (20)$$

with $\bar{\beta} = \text{diag}([\mathbf{1}\beta^1, \mathbf{1}\beta^2, \dots, \mathbf{1}\beta^K]) \in \mathbb{R}^{KL}$ with $\mathbf{1} \in \mathbb{R}^L$. The gradient for the dual variable β^k is given by:

$$\frac{\partial g(d^*, \beta)}{\partial \beta^k} = \|d^{*k}\|_2^2 - 1, \quad (21)$$

with d^{*k} computed from (20). We can solve this iteratively using again L-BFGS-B taking into account the positivity constraint $\beta^k \geq 0$ for all k . What we have described so far solves for all the atoms simultaneously. However, it is also possible to estimate the atoms sequentially one at a time using a block coordinate descent (BCD) approach, as in the work of [36]. In each iteration of the BCD algorithm, a residual r_n^k is computed as given by:

$$r_n^k = x_n - \sum_{k' \neq k} Z_n^{k'} d^{k'} \quad (22)$$

and correspondingly subproblem 17 becomes:

$$\min_{d^k} \sum_{n=1}^N \frac{1}{2} \|\sqrt{w_n} \odot (r_n^k - Z_n^k d^k)\|_2^2, \quad \text{s.t. } \|d^k\|_2^2 \leq 1, . \quad (23)$$

which is solved in the same way as subproblem 17. Now, in the simultaneous case, we construct one linear problem in $\mathcal{O}(L^2 K^2 TN)$ and one iteration costs $\mathcal{O}(L^3 K^3)$. However, in the BCD strategy, we construct K linear problems in $\mathcal{O}(L^2 TN)$ and one iteration costs only $\mathcal{O}(L^3)$. Interestingly, when the weights w_n are all identical, we can use the fact that for one atom k , the matrix $\sum_{i=1}^N (Z_i^k)^T Z_i^k$ is Toeplitz. In this case, we can construct K linear problems in only $\mathcal{O}(LTN)$ and one iteration costs only $\mathcal{O}(L^2)$.

For the benefit of the reader, we summarize the complexity of the M-step in Table 1. We note p and q the number of iterations in the L-BFGS-B methods for the activations update and atoms update.

Method	Complexity
Solving activations z	$p \min(L, \log(T))KTN$
Solving atoms d	$L^2K^2TN + qL^3K^3$
Solving atoms d (sequential)	$LKTN + qL^2K$

Table 1: Complexity analysis of the M-step, where p and q are the number of iterations in the L-BFGS-B solvers for the activations and atoms updates.

C Additional Experiments: M-step speed benchmark

C.1 Comparison with state of the art

Here, we compare convergence plots of our algorithm against a number of state-of-art methods. The details of the experimental setup are described in Section 4. Fig. 6 demonstrates on a variety of setups the computational advantage of our quasi-Newton approach to solve the M-step. Note that Fig. 2b is in fact a summary of Fig. 6. Indeed, we can verify that ADMM methods converge quickly to a modest accuracy, but take much longer to converge to a high accuracy.[37]

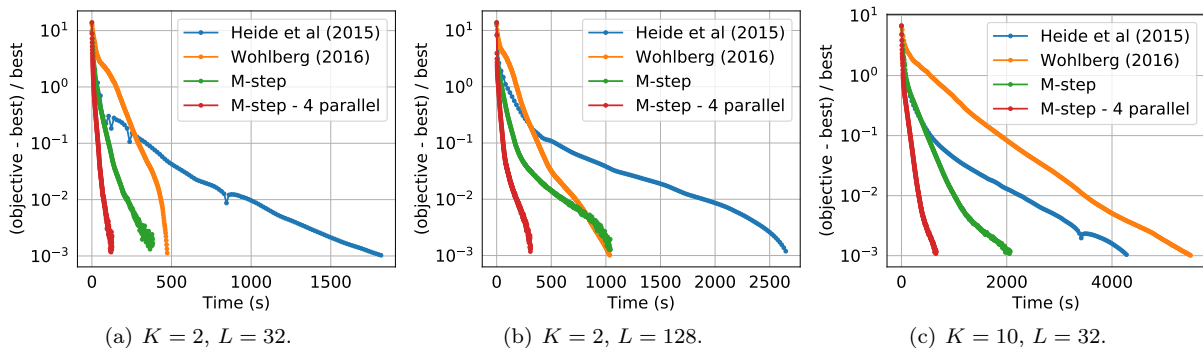


Figure 6: Convergence speed of the relative objective function. The y-axis shows the objective function relative to the obtained minimum for each run: $(f(x) - f(x^*)) / f(x^*)$. Each curve is the geometrical mean over 24 different random initializations.

Next, in Fig. 7, we show more traditional convergence plots. In contrast to Fig. 2 or 6 where the relative objective function is shown, here we plot the absolute value of the objective function. We can now verify that each of the methods have indeed converged to their respective local minimum. Of course, owing to the non-convex nature of the problem, they do not necessarily converge to the same local minimum.

C.2 Comparison of solver for the activations subproblem

Finally, we compare convergence plots of our algorithm using different solvers for the z -update: ISTA, FISTA, and L-BFGS-B. The rationale for choosing a quasi-Newton solver for the z -update becomes clear in Fig. 8 as the L-BFGS-B solver turns out to be computationally advantageous on a variety of setups.

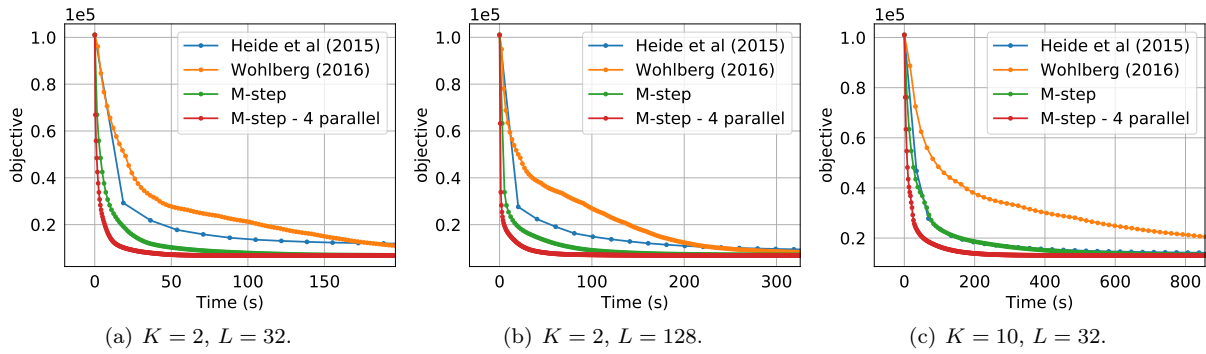


Figure 7: Convergence of the objective function as a function of time. The y-axis shows the absolute objective function $f(x)$. Each curve is the mean over 24 different random initializations.

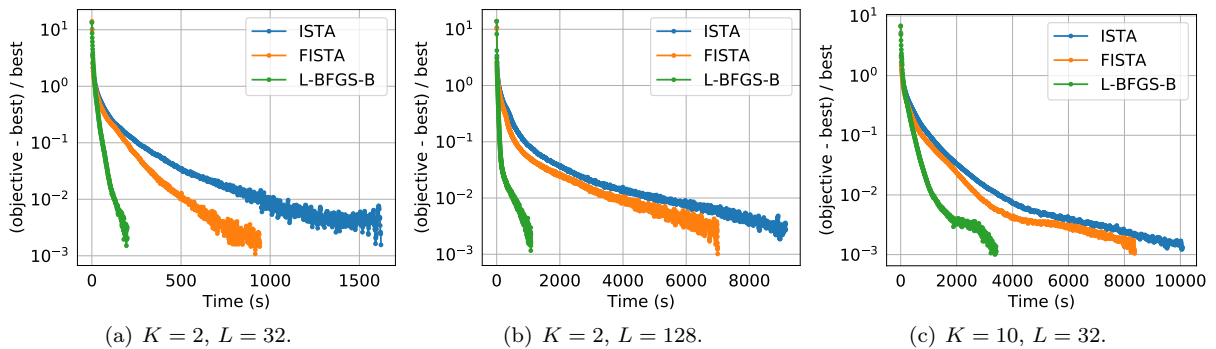


Figure 8: Convergence speed of the relative objective function. The y-axis shows the objective function relative to the obtained minimum for each run: $(f(x) - f(x^*)) / f(x^*)$. Each curve is the geometrical mean over 24 different random initializations.

References

- [36] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*, 11(Jan):19–60, 2010.
- [37] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.