



HAL
open science

Tensor Convolutional Sparse Coding with Low-Rank activations, an application to EEG analysis

Pierre Humbert, Laurent Oudre, Nivolas Vayatis, Julien Audiffren

► **To cite this version:**

Pierre Humbert, Laurent Oudre, Nivolas Vayatis, Julien Audiffren. Tensor Convolutional Sparse Coding with Low-Rank activations, an application to EEG analysis. 2020. hal-02888698v1

HAL Id: hal-02888698

<https://hal.science/hal-02888698v1>

Preprint submitted on 3 Jul 2020 (v1), last revised 9 Jul 2020 (v2)

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.

Tensor Convolutional Sparse Coding with Low-Rank activations, an application to EEG analysis

Pierre Humbert¹ Laurent Oudre² Nivolas Vayatis¹ Julien Audiffren³

¹ Université Paris-Saclay, CNRS, ENS Paris-Saclay, Centre Borelli,
F-91190 Gif-sur-Yvette, France

² Université Sorbonne Paris Nord, L2TI, UR 4443,
F-93430 Villetaneuse, France

³ eXascale Infolab, University of Fribourg, Fribourg, Switzerland

July 3, 2020

Abstract

Recently, there has been growing interest in the analysis of spectrograms of ElectroEncephaloGram (EEG), particularly to study the neural correlates of (un)-consciousness during General Anesthesia (GA). Indeed, it has been shown that order three tensors (channels \times frequencies \times times) are a natural and useful representation of these signals. However this encoding entails significant difficulties, especially for convolutional sparse coding (CSC) as existing methods do not take advantage of the particularities of tensor representation, such as rank structures, and are vulnerable to the high level of noise and perturbations that are inherent to EEG during medical acts. To address this issue, in this paper we introduce a new CSC model, named Kruskal CSC (K-CSC), that uses the Kruskal decomposition of the activation tensors to leverage the intrinsic low rank nature of these representations in order to extract relevant and interpretable encodings. Our main contribution, TC-FISTA, uses multiple tools to efficiently solve the resulting optimization problem despite the increasing complexity induced by the tensor representation. We then evaluate TC-FISTA on both synthetic dataset and real EEG recorded during GA. The results show that TC-FISTA is robust to noise and perturbations, resulting in accurate, sparse and interpretable encoding of the signals.

1 Introduction

General anesthesia (GA) consists in a medically induced state of unconsciousness, through the use of different drugs such as inhalational hypnotic anesthetics. It is a cornerstone of modern medicine, and is crucial for the realization of many medical and surgical procedures [Purdon et al., 2013]. However, anesthesia may carry some important risks (e.g. cognitive dysfunction [Punjasawadwong et al., 2018], postoperative delirium [Fritz et al., 2016]). Consequently, a sustained and careful monitoring of the level of consciousness of the patient – also referred to as the Depth of Anesthesia (DoA) – is required. As a direct measurement of the electrical brain activity, the Electroencephalograms (EEG) recordings remain the gold-standard to assess the DoA and have extensively been used to study the phenomenons occurring during anesthesia-induced loss of consciousness [Liu and Rinehart, 2016]. For instance, Purdon et al. [2013] has shown that the evolution of α and δ -waves (respectively in the 8-13 Hz and 1-3Hz ranges) are promising predictors of the recovery of consciousness after GA that have been induced with propofol.

However, there is still no consensus on the existence and nature of universal markers that would track DoA, especially during transition phases such as Loss and Recovery of Consciousness (LOC, and ROC respectively). This is partially due to the heavily influence of the choice of the anesthetic drug over the behavior of EEG [Purdon et al., 2013], and intrinsic technical limitations. For instance, when data are recorded during real surgeries, the EEG signals are prone to low signal to noise ratio, impulsive noise due to sensor malfunctions, and artifacts caused by e.g. electro-surgical devices that are used to cut and cauterize tissue [Dubost et al., 2019]. Nevertheless, there has been growing interest in the automatic analysis of EEG patterns in such settings [Becker et al., 2015]. In particular, we can distinguish two families of statistical learning methods that are used in this domain.

In tensor based approaches (see e.g. [Becker et al., 2010; Zhao et al., 2011]), EEG are frequently analyzed by computing a short-time Fourier transform (STFT) for each channel of the EEG, resulting in a tensor of order 3 encoding a Space-Time-Frequency representation. The resulting tensor is then studied through the prism of the CANDECOMP/PARAFAC (CP) decomposition [Becker et al., 2014], in order to exploit the interactions among multiple modes. This approach is frequently paired with the addition of a low CP-rank constraint on the tensor decomposition, resulting in representations that are more robust to noise and easier to interpret [Zhou et al., 2013; Zhao et al., 2011; Cong et al., 2015; Rabusseau and Kadri, 2016].

Similarly, there have been several attempts to extract new representations of unfiltered electrophysiological signals using shift-invariant dictionary learning approaches such as Convolutional Sparse Coding (CSC) (see e.g. [Jas et al., 2017; La Tour et al., 2018]). While these methods exhibit interesting results, their models are centered around atoms of order one, i.e. atoms that vary along a single mode. Indeed, CSC methods are mainly focused on resolution for univariate signals or images [Garcia-Cardona and Wohlberg, 2018a], and therefore do not take into account the possible interaction between all the different modes of multivariate data – a key property of the EEG spectrogram representation [Cong et al., 2015]. Moreover, these methods are frequently vulnerable to noise [Jas et al., 2017; Simon and Elad, 2019] and structured perturbations such as impulsive noise [Wang et al., 2020].

Following these remarks, in this work we aim at combining the two previous approaches by introducing a Tensor CSC model, where both activations and atoms can be high order tensors. Inspired by recent works from Phan et al. [2015] and Humbert et al. [2020], we constrain the solutions of this model to have low CP-rank activations, with the aim of obtaining accurate and interpretable decompositions of multidimensional data.

Contributions. In this article, we develop a low rank tensor based CSC model, named Kruskal CSC (or K-CSC for short), that is particularly adapted to analyse noisy tensor signals such as EEG spectrograms recorded during GA (Section 2). We propose an efficient optimization strategy to learn the different components of this model, named TC-FISTA, that entails a low complexity and computation time (Section 3). Using synthetic data, we show that this approach, where the rank constraint is enforced on activations instead of atoms, is significantly more robust to noise than its unconstrained counterparts. Finally, we evaluate our method on a dataset of EEG recording during GA. We show that TC-FISTA successfully isolates the complex noise and artifacts of the signal and produces a sparse, accurate and interpretable decomposition (Section 4).

Notation. In the rest of this paper, tensors (resp. matrices, vectors) are represented with capital, bold calligraphic letters (resp. capital bold letters, bold lower letters). $\mathcal{Y}_1, \dots, \mathcal{Y}_N$ denote N tensor-valued signals in $\mathbb{Y} \triangleq \mathbb{R}^{n_1 \times \dots \times n_p}$, $\{\mathcal{D}_k\}_{1 \leq k \leq K}$ are the K multivariate atoms in $\mathbb{D} \triangleq \mathbb{R}^{w_1 \times \dots \times w_p}$, with $(w_1 \leq n_1, \dots, w_p \leq n_p)$. $\{\mathcal{Z}_{n,k}\}_{1 \leq k \leq K, 1 \leq n \leq N} \in \mathbb{Y}$ are (multidimensional) activation maps. Regarding tensor operations, \otimes denotes the (circular) convolutional operator [Dudgeon, 1983], $\hat{\cdot}$ denotes the frequency representation of a signal via the Discrete Fourier Transform (DFT), \circ represents the outer product, \otimes is the Kronecker product, \odot is the Khatri-Rao product, and vec denotes the vectorization transformation. Finally, we will use the word *rank* to refer to the notion of *Canonical Polyadic rank* (CP-rank) of tensors [Zhou et al., 2013].

2 Low rank Tensor CSC

In this section, we briefly define the multivariate CSC setting, before introducing K-CSC, the model used in this work. Then we discuss the ideas behind this model and its relation with the existing literature.

Multivariate CSC. Formally, given a finite set of N multivariate signals $\mathcal{Y}_1, \dots, \mathcal{Y}_N \in \mathbb{Y}$, tensors inputs of order $p > 0$ and $\alpha > 0$ a scalar, the multivariate CSC problem is defined as

$$\min_{\{\mathcal{D}_k \in \mathbb{D}, \mathcal{Z}_{n,k} \in \mathbb{Y}\}} \frac{1}{2} \sum_{n=1}^N \left(\left\| \mathcal{Y}_n - \sum_{k=1}^K \mathcal{D}_k \otimes \mathcal{Z}_{n,k} \right\|_F^2 + \alpha \sum_{k=1}^K \|\mathcal{Z}_{n,k}\|_1 \right) \quad \text{s.t.} \quad \|\mathcal{D}_k\|_F \leq 1 \quad \forall k. \quad (1)$$

In this formulation, the $\mathcal{Z}_{n,k}$ are activation maps which specify where the atoms \mathcal{D}_k are placed in the input signals. Note that both $\mathcal{Z}_{n,k}$ and \mathcal{D}_k are tensors of order p . The constraint on \mathcal{D}_k prevents the

scaling indeterminacy problem between atoms and activations, as in the standard CSC (see e.g. [Bristow et al., 2013]). The convolutional sparse coding of multi-channel signals has received only limited attention in the past [Barthélemy et al., 2012; Wohlberg, 2016; Garcia-Cardona and Wohlberg, 2018a], and only few existing work consider the application of convolutional sparse representations to signals with more than three channels [Garcia-Cardona and Wohlberg, 2018b].

Low rank Tensor CSC. The key idea behind K-CSC is the use of a low rank constraint on the *activation maps* in order to take advantage of the tensor structure. Indeed, previous works have shown that adding a low rank constraint to tensor learning problems significantly improves the performance of many methods (see e.g. [Zhou et al., 2013; Zhao et al., 2011; Rabusseau and Kadri, 2016]). To enforce this constraint in CSC, we used the Kruskal operator, whose definition is recalled below

Definition 1. (Kruskal operator) – Let $R > 0$ and $p > 0$. Then the Kruskal operator of rank R and of order p denoted $\llbracket \cdot \rrbracket_{R,p}$ is defined as

$$\llbracket \cdot \rrbracket_{R,p} : \mathbb{R}^{n_1 \times R} \times \dots \times \mathbb{R}^{n_p \times R} \mapsto \mathbb{Y}, \quad \llbracket \mathbf{X}^{(1)}, \dots, \mathbf{X}^{(p)} \rrbracket_{R,p} \rightarrow \sum_{r=1}^R \mathbf{x}_r^{(1)} \circ \dots \circ \mathbf{x}_r^{(p)},$$

where $\forall 1 \leq i \leq p, \forall 1 \leq r \leq R, \mathbf{x}_r^{(i)} \in \mathbb{R}^{n_i}$ is the r -th column of $\mathbf{X}^{(i)}$.

In the following, and to make notation easier, we use $\llbracket \cdot \rrbracket$ to denote $\llbracket \cdot \rrbracket_{R,p}$ when the values of R, p are clear from the context. Note that $\llbracket \cdot \rrbracket_{R,p}$ takes value in the subspace of tensors of rank lower than R , and conversely any tensor in \mathbb{Y} of rank lower than R can be written using the Kruskal operator (see e.g. [Kruskal, 1989]). However this decomposition is not necessarily unique.

Using the Kruskal operator, the Low rank Tensor CSC optimization problem can be rewritten as:

$$\min_{\{\mathcal{D}_k, \mathcal{Z}_{n,k}\}} \frac{1}{2} \sum_{n=1}^N \left(\left\| \mathcal{Y}_n - \sum_{k=1}^K \mathcal{D}_k \otimes \mathcal{Z}_{n,k} \right\|_F^2 + \underbrace{\sum_{q=1}^p \alpha_q \sum_{k=1}^K \|\mathcal{Z}_{n,k}^{(q)}\|_1}_{\text{Mode Sparsity}} + \underbrace{\sum_{q=1}^p \beta_q \sum_{k=1}^K \|\mathcal{Z}_{n,k}^{(q)}\|_F}_{\text{Identifiability Regularization}} \right) \quad (2)$$

$$\text{s.t.} \quad \begin{cases} \mathcal{Z}_{n,k} = \llbracket \mathcal{Z}_{n,k}^{(1)}, \dots, \mathcal{Z}_{n,k}^{(p)} \rrbracket_{R,p} & \forall n, k \quad (\text{Rank Constraint}) \\ \mathcal{D}_k \in \mathbb{D}, \|\mathcal{D}_k\|_F \leq 1 & \forall k. \end{cases}$$

Compared to (1), this optimization problem presents the following crucial differences.

- **Rank constraint on the activations.** Due to this constraint, the rank of $\mathcal{Z}_{n,k}$ is lower than R . This aims at controlling the linear relation between the different modes of the activations maps and thus taking into account the natural structure of the data. While using the Kruskal decomposition to enforce a rank constraint on the dictionary has been studied (see e.g. [Rigamonti et al., 2013]), applying this constraint on the activations has received little attention [Humbert et al., 2020]. The interests of this approach and its relation with the literature are discussed below.
- **Identifiability and the Frobenius regularization.** The Kruskal decomposition, i.e. the tensor decomposition using the Kruskal operator, is prone to scaling indeterminacy¹, resulting in a continuous manifold of equivalent solutions. The Frobenius regularization in (2) is known to address this problem (see [Paatero, 1997] and [Acar et al., 2011] for more details).
- **Mode sparsity regularization.** The ℓ_1 -norms regularization on the modes $\mathcal{Z}_{n,k}^{(q)}$ of the activation tensors enforces sparse solutions. However, contrary to the usual sparsity constraint in (1), we take advantage of the Kruskal decomposition and enforce different degree of sparsity for each mode of the activation tensors independently. The sparsity in each mode is therefore controlled independently, allowing a higher flexibility – particularly in applications where modes have different meaning, such as in EEG analysis where the sparsity constraint on channels is significantly lower than the one on time (see Section 4).

¹Equally valid Kruskal decompositions of a tensor can be obtained by multiplying two modes by respectively γ and $1/\gamma$ for any $\gamma > 0$.

Low-rank constraints on sparse activations. The idea of enforcing low-rank constraints for CDL is not novel: Rigamonti et al. [2013] and Sironi et al. [2014] used the idea of separable filters for learning low-rank atoms in order to improve computational runtime. More recent publications including Quesada et al. [2018]; Silva et al. [2018]; Quesada et al. [2019] have also successfully used low-rank constraints on dictionary to introduce efficient solvers but only in 2-D. Importantly, in all the aforementioned works, the low-rank constraints is enforced on the atoms and not on the activations. Nevertheless, we claim that constraining the activations to be low rank brings two majors advantages. First, in multiple application contexts the low-rank structure naturally appears in the activations rather than in the atoms/dictionary (see [Humbert et al., 2020], and Section 4). Our model allows to leverage this prior knowledge. Second, low-rank constraints on activation entail a better robustness with respect to noise, which is one of the main weakness of the activation learning part of CDL [Simon and Elad, 2019]. Finally, the sparsity of the activations is supported by recent results on neuroscience which postulate that neural activity consists more of transient bursts of isolated events rather than rhythmically sustained oscillations van Ede et al. [2018]; Karvat et al. [2020]. Hence, such activities could be described not only by their frequency and amplitude but also by their rate, duration, and shape suggesting that multivariate CDL is well-adapted to analyze them.

3 Resolution of the problem: TC-FISTA

We now introduce TC-FISTA, an algorithm that solves the K-CDL problem by iterating over a dedicated FISTA solver. TC-FISTA minimizes the objective function (2) by using a bloc-coordinate strategy approach [Hildreth, 1957; Nikolova and Tan, 2017]. Indeed, TC-FISTA splits the main non-convex problem into several convex subproblems, by either 1) freezing all the atoms \mathcal{D}_k and all the modes of \mathcal{Z} except one (referred as \mathcal{Z} -step, see Section 3) or 2) freezing \mathcal{Z} referred as \mathcal{D} -step. It is important to note that while TC-FISTA is not guaranteed to converge toward the global minimum (similarly to other CDL solvers, see [Bristow et al., 2013]), in our empirical evaluations the algorithm almost always converges to a near-optimal local minimum. Since the optimization problem pertaining to \mathcal{D} -step is identical to the \mathcal{D} -step of Multivariate CDL (1), it can be solved using any existing methods (see [Garcia-Cardona and Wohlberg, 2018a] for a review). Following the recent result of [Wohlberg, 2015], we chose to use ADMM with iterative application of the Sherman-Morrison equation for the \mathcal{D} -step of TC-FISTA. In our setting, this is by far the fastest method as reported in [Wohlberg, 2015] and [Garcia-Cardona and Wohlberg, 2018a]. The rest of this section is devoted to the \mathcal{Z} -step.

\mathcal{Z} -step : Activations update

In the remainder of the section, we consider that the atoms \mathcal{D}_k are fixed. Additionally, and to alleviate the notation, we only consider the K-CSC problem with $N = 1$ signal, and denote $\mathcal{Z}_q \doteq \mathcal{Z}_{1,q}$ and $\mathbf{Z}_{1,k}^{(q)} \doteq \mathbf{Z}_k^{(q)}$. All the results and methods can be immediately extended to any value of $N > 1$ as the activation maps are independent across the signals $\mathcal{Y}_1, \dots, \mathcal{Y}_N$.

Block-Coordinate Descent. One of the main difficulties of the K-CDL model is that because of the low rank constraint, (2) is no longer convex with respect to \mathcal{Z} . While it is theoretically possible to solve the optimization problem by rewriting it as a regression problem and using methods proposed by [Zhou et al., 2013; He et al., 2018], this approach would require the construction and manipulation of a circulant tensor of size $\approx K \times (n_1 \times \dots \times n_p)^2$ which is not tractable in practice. To address this problem, we consider the modes of \mathcal{Z} , defined as $(\mathbf{Z}_1^{(q)}, \dots, \mathbf{Z}_K^{(q)})$, for any value of $q \in [1, \dots, p]$. Indeed, (2) is convex with respect to each of the modes of \mathcal{Z} individually, i.e. when all the other modes are fixed. Furthermore, the two regularization terms of (2) are separable with respect to the $\mathbf{Z}_k^{(q)}$. Building on this remark, TC-FISTA does not directly solve (2) but instead uses a block-coordinate strategy, minimizing p different sub-problems where all *modes* except the q -th one ($1 \leq q \leq p$) of each activation tensor are fixed. The resulting optimization sub-problems can be rewritten as

$$\min_{\mathbf{Z}_1^{(q)}, \dots, \mathbf{Z}_K^{(q)}} \frac{1}{2} \left\| \underbrace{\mathcal{Y} - \sum_{k=1}^K \mathcal{D}_k \otimes_{1, \dots, p} [\mathbf{Z}_k^{(1)}, \dots, \mathbf{Z}_k^{(p)}]}_{f(\{\mathbf{Z}_k^{(r)}\}_{k=1, r=1}^{K,p})} \right\|_F^2 + \underbrace{\alpha_q \sum_{k=1}^K \|\mathbf{Z}_k^{(q)}\|_1 + \beta_q \sum_{k=1}^K \|\mathbf{Z}_k^{(q)}\|_F^2}_{g(\{\mathbf{Z}_k^{(q)}\}_{k=1}^K)}, \quad (3)$$

where f is the *fidelity term* which controls the difference between the input and its reconstruction, and g is the sum of the regularization terms.

Solving the block-coordinate problem. In standard CSC, two solvers dominate the literature to learn the activation maps: Alternating Direction Method of Multipliers (ADMM) [Glowinski and Marroco, 1975] and Fast Iterative Soft Thresholding Algorithm (FISTA) [Beck and Teboulle, 2009]. Both are competitive and recent comparative reviews [Garcia-Cardona and Wohlberg, 2018a] indicate a very wide range of performances across existing methods. In this article, we chose to focus on FISTA to solve (3), but our approach can also be used with the ADMM solver.

Solving (3) with FISTA requires repeated A) computation of the gradient of f and B) evaluation of the proximal operator of g . For the later, as g is fully separable, its proximal is given by the weighted soft-thresholding operator for each $\mathbf{Z}_{k,q}$ [Parikh et al., 2014]. Notice that any g with known proximal can be used with TC-FISTA. For instance, in Section 4, we add a non-negative constant on the activations when using spectrograms derived from EEG, resulting in a proximal given by the soft non-negative thresholding operator. Regarding A), calculating the gradient of f is computationally demanding, particularly in the tensor setting, and the rest of this section is devoted to introduce several important optimizations.

Accelerating gradient computation using Fourier transform. Recent works [Wohlberg, 2016; Garcia-Cardona and Wohlberg, 2018a] have shown that projecting the computation of the gradient of f in the Fourier domain drastically decreases its complexity, resulting in a significantly faster optimization. This method has proven to be important to perform CSC quickly and efficiently on complex and large signals, and is crucial to the tensor-based CSC presented in this paper due to the exponential increase of the size of signals and atoms. However, the generalization of this technique to our tensor setting (3) is not straightforward. In the following, we provide a new formulation of the gradient in the Fourier domain, that is easy to compute efficiently.

Theorem 1. (Compact vectorized formulation) – Let $\widehat{\mathbf{y}}^{(q)}$ (resp. $\widehat{\mathbf{d}}_k^{(q)}$) be the vectorization of the folding of $\widehat{\mathbf{Y}}$ (resp. $\widehat{\mathbf{D}}_k$) along the dimension q , $\widehat{\mathbf{Z}}_k^{(q)} = [\widehat{\mathbf{Z}}_k^{(q)}(:, 1) \mid \dots \mid \widehat{\mathbf{Z}}_k^{(q)}(:, R)]$ be the DFT of $\mathbf{Z}_k^{(q)}$ along its columns and $\widehat{\mathbf{z}}_k^{(q)} = \text{vec}(\widehat{\mathbf{Z}}_k^{(q)})$. Then

$$f\left(\{\mathbf{Z}_k^{(r)}\}_{k=1, r=1}^{K,p}\right) \propto \frac{1}{2} \left\| \widehat{\mathbf{y}}^{(q)} - \widehat{\mathbf{\Gamma}}^{(q)} (\widehat{\mathbf{A}}^{(q)} \otimes \mathbf{I}^{(q)}) \widehat{\mathbf{z}}^{(q)} \right\|_F^2, \quad (4)$$

where $\widehat{\mathbf{z}}^{(q)} = [\widehat{\mathbf{z}}_1^{(q)\top}, \dots, \widehat{\mathbf{z}}_K^{(q)\top}]^\top \in \mathbb{C}^{KRn_q}$, $\widehat{\mathbf{\Gamma}}^{(q)} = [\text{diag}(\widehat{\mathbf{d}}_1^{(q)}), \dots, \text{diag}(\widehat{\mathbf{d}}_K^{(q)})] \in \mathbb{C}^{n_1 \dots n_p \times Kn_1 \dots n_p}$, $\mathbf{I}^{(q)} \in \mathbb{R}^{n_q \times n_q}$ is the identity matrix and

$$\widehat{\mathbf{A}}^{(q)} = \begin{pmatrix} \widehat{\mathbf{B}}_1^{(q)} & & 0 \\ & \ddots & \\ 0 & & \widehat{\mathbf{B}}_K^{(q)} \end{pmatrix} \in \mathbb{C}^{K \prod_{1, i \neq q} n_i \times KR} \quad \text{with} \quad \widehat{\mathbf{B}}_k^{(q)} = \left(\overset{\leftarrow p}{\odot}_{i=1, i \neq q} \widehat{\mathbf{Z}}_k^{(i)} \right).$$

Proof. Equation (4) is derived by using Plancherel theorem and successive tensor manipulations. The full proof can be found in the appendix. \square

From Theorem 1, we can derive the gradient of f in the Fourier domain.

Corollary 1. (Gradient of f) – With the notation of Theorem 1, the gradient of f with respect to $\mathbf{z}^{(q)} = [\text{vec}(\mathbf{Z}_1^{(q)})^\top, \dots, \text{vec}(\mathbf{Z}_K^{(q)})^\top]^\top$ is given by

$$\nabla_{\mathbf{z}^{(q)}} f\left(\{\mathbf{Z}_k^{(r)}\}_{k=1, r=1}^{K,p}\right) = \text{IDFT} \left[\left((\widehat{\mathbf{A}}^{(q)} \otimes \mathbf{I}) \widehat{\mathbf{\Gamma}}^{(q)} \right)^H \left(\widehat{\mathbf{\Gamma}}^{(q)} (\widehat{\mathbf{A}}^{(q)} \otimes \mathbf{I}) \widehat{\mathbf{z}}^{(q)} - \widehat{\mathbf{y}}^{(q)} \right) \right], \quad (5)$$

where $\text{IDFT}[\cdot]$ stands for the Inverse Discrete Fourier Transform.

Equation (5) provides a closed-form equation for computing the values of the partial derivative of f . Crucially, some of its terms can be precomputed, further improving the rapidity of the \mathbf{Z} -step. This point is discussed in the next paragraph. In the following, we drop the upper-script $\cdot^{(q)}$ of $\widehat{\mathbf{A}}$ and $\widehat{\mathbf{\Gamma}}$ for clarity.

Table 1: Time and Space complexities of one iteration of the \mathcal{Z} -step for TC-FISTA, **FCSC-ShM** with Sherman-Morrison iterations [Bristow et al., 2013; Wohlberg, 2015] and **ConvFISTA-FD** in the Fourier domain [Wohlberg, 2015; Garcia-Cardona and Wohlberg, 2018a]. K denotes the number of atoms, n_i is the size of the signal along order i , $M = \prod_{i=1}^p n_i$ is the size of the signal and R is the maximal rank.

| Algorithm | Time complexity (\mathcal{Z} -step) | Dominant term | Space |
|---------------------|---|--|--------------------------------------|
| FCSC-ShM | $\underbrace{KM}_{\text{Linear systems}} + \underbrace{KM \log(M)}_{\text{FFTs}} + \underbrace{KM}_{\text{Shrinkage}}$ | $K \left(\prod_{q=1}^p n_q \right) \log \left(\prod_{q=1}^p n_q \right)$ | $K \left(\prod_{q=1}^p n_q \right)$ |
| ConvFISTA-FD | $\underbrace{KM}_{\text{Gradient}} + \underbrace{KM \log(M)}_{\text{FFTs}} + \underbrace{KM}_{\text{Shrinkage}}$ | $K \left(\prod_{q=1}^p n_q \right) \log \left(\prod_{q=1}^p n_q \right)$ | $K \left(\prod_{q=1}^p n_q \right)$ |
| T-ConvFISTA | $\underbrace{(KR)^2 \sum_{q=1}^p n_q}_{\text{Gradient}} + \underbrace{KR \sum_{q=1}^p n_q \log(M)}_{\text{FFTs}} + \underbrace{KR \sum_{q=1}^p n_q}_{\text{Shrinkage}}$ | $(KR)^2 \sum_{q=1}^p n_q$ | $KR \max(n_q)$ |

Efficient computation of the gradient via the Gram matrix. In the case where $\prod_{i=1}^p n_i \gg KRn_q$ (low rank assumption), two terms of (5) can be precomputed efficiently: $(\hat{\mathbf{A}}^H \otimes \mathbf{I})\hat{\Gamma}^H \hat{\mathbf{y}}^{(q)}$ and the Gram matrix $\mathbf{G} \triangleq (\hat{\mathbf{A}}^H \otimes \mathbf{I})\hat{\Gamma}^H \hat{\Gamma}(\hat{\mathbf{A}} \otimes \mathbf{I})$. For the first term, we have

$$(\hat{\mathbf{A}} \otimes \mathbf{I})\hat{\mathbf{y}}^{(q)} = (\hat{\mathbf{A}} \otimes \mathbf{I})\text{vec}(\mathbf{Y}^{(q)}) = \text{vec}(\mathbf{Y}^{(q)} \hat{\mathbf{A}}^T),$$

whose computation can be performed in $\mathcal{O}(KR \prod_{i=1}^p n_i)$ operations, instead of $\mathcal{O}(KRn_q \prod_{i=1}^p n_i)$ for a naive strategy. Moreover, the following proposition shows that the Gram matrix \mathbf{G} can be computed independently and efficiently.

Proposition 1. *The matrix $\mathbf{G} \triangleq (\hat{\mathbf{A}}^H \otimes \mathbf{I})\hat{\Gamma}^H \hat{\Gamma}(\hat{\mathbf{A}} \otimes \mathbf{I})$ can be obtained by computing K^2 blocks $\mathbf{G}_{k,\ell}$, $1 \leq k, \ell \leq K$, where*

$$\mathbf{G}_{k,\ell} = \left(\left(\overset{\leftarrow}{\odot}_{i=1, i \neq q} \hat{\mathbf{Z}}_{k,i} \right)^H \otimes \mathbf{I} \right) \overline{\text{diag}(\hat{\mathbf{d}}_k^{(q)})} \text{diag}(\hat{\mathbf{d}}_\ell^{(q)}) \left(\left(\overset{\leftarrow}{\odot}_{i=1, i \neq q} \hat{\mathbf{Z}}_{\ell,i} \right) \otimes \mathbf{I} \right), \quad (6)$$

and each of these blocks can be computed in $\mathcal{O}(R^2 \prod_{i=1, i \neq q}^p n_i)$ operations.

As a consequence of proposition 1, the full matrix \mathbf{G} can be computed in $\mathcal{O}((KR)^2 \prod_{i=1, i \neq q}^p n_i)$ operations. Moreover, \mathbf{G} is a particular $(KRn_q \times KRn_q)$ banded matrix (see appendix), and consequently its product with $\hat{\mathbf{z}}^{(q)}$ can be made in only $\mathcal{O}((KR)^2 n_q)$ operations instead of $\mathcal{O}((KRn_q)^2)$ with a naive strategy. The resulting total complexity of the \mathcal{Z} -step of TC-FISTA and classical CSC approaches are reported in Table 1.

4 Experiments

All numerical experiments were run on a Linux laptop with 12-core 2.5GHz Intel CPUs using standard Python libraries, Tensorly [Kossaifi et al., 2019], Sporc0 [Wohlberg, 2017], and all the speed-ups from Section 3. Extended results and additional experiments can be found in the appendix. Our code is publicly available at https://github.com/pierreHmbt/Tensor_CDL.

4.1 Evaluation on synthetic data

We first evaluated the performance of TC-FISTA on a synthetic dataset, in order to illustrate the advantages of the low rank approach and highlight the benefits of the gradient computation optimization proposed in Section 3. In these experiments, TC-FISTA was compared with the two state-of-the-art methods for multivariate CSC, **FCSC** with Sherman-Morrison iterations (FCSC-ShM) [Bristow et al., 2013; Wohlberg, 2015] and **ConvFISTA** in the Fourier domain (ConvFISTA-FD) [Wohlberg, 2015; Garcia-Cardona and Wohlberg, 2018a]. To allow for fair speed comparison, we used the strategy proposed in [Mairal et al., 2009, 2010], and re-implemented the other two methods in Python.

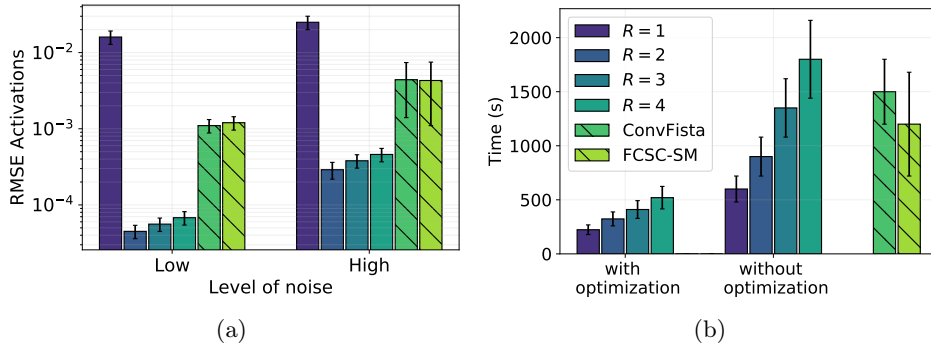


Figure 1: **Left:** Average RMSE (log-scale) between true and reconstructed activation tensors (\mathcal{Z} -step) on the synthetic dataset. Low (resp. High) noise denotes a SNR of 25 (resp. 10). **Right:** Time until convergence of the dictionary learning process ($\mathcal{Z} + \mathcal{D}$ steps), with and without the optimizations discussed in Section 3. The standard deviation are indicated using black lines.

Synthetic simulation setup. Experiments were conducted on a random synthetic dataset, containing 10 independent signals of size $128 \times 128 \times 128$ generated using $K = 3$ atoms of size $5 \times 5 \times 5$ and activation tensors of rank $R^* \leq 2$. Following Wohlberg [2015], we added a multivariate Gaussian noise to each signal to obtain a low level of noise (SNR of 25dB) or a high level of noise (SNR of 10dB). Each hyperparameter was chosen between $[0.0001, 100]$ for all methods. The gradient step of TC-FISTA was set to the inverse of the Lipschitz constant of each subproblem (4).

TC-FISTA with known dictionary. First, we evaluated the performance of the different algorithms on the \mathcal{Z} -step – i.e. the reconstruction of the activation tensors using a known dictionary – for different values of the noise ratio, and rank constraint R (for TC-FISTA). Each method was run with five different random initializations, and the one returning the best result was kept. We measured the performance using the Root Mean Square Error (RMSE) between the true values of the activations \mathcal{Z} and their estimated values. The results are reported in Figure 1 (left). First, note that if the rank is underestimated ($R = 1$), then TC-FISTA does not learn a suitable estimation of the activations. However, interestingly, if the rank is overestimated ($R = 3$ or 4), then TC-FISTA only suffers from a small drop in performance. These two behavior are in line with previous results on low rank tensors estimation [Zhou et al., 2013]. We also see that TC-FISTA (with $R > 1$) performs significantly better than its counterparts without rank constraints, particularly in presence of strong noise, justifying the low rank activation approach. Finally, note that despite using different initializations, TC-FISTA consistently learns a good estimator of \mathcal{Z} , resulting in a low standard deviation.

TC-FISTA with unknown dictionary. We then compared the performance of the aforementioned algorithms on the complete K-CSC problem, which includes the simultaneous learning of the activations (\mathcal{Z} -step) and the atoms (\mathcal{D} -step), on signals with high level of noise. Figure 1(right) shows the average time until convergence (i.e. until the relative convergence tolerance becomes lower than $1e - 4$ [Boyd et al., 2011]). While it is important to remark that the relative speeds of each methods are dependent of their choice of hyperparameters as well as on the sparsity of the signals, we observe that A) TC-FISTA with the optimizations discussed in Section 3 is significantly faster than its regular counterpart and B) that TC-FISTA is faster than FCSC-ShM and ConvFista, even if the advantage decreases as R increases. This is in line with the time complexity of each algorithm, as presented in Table 1.

4.2 Multichannel EEG signals during a General Anesthesia

In this experiment, we study multichannel EEG signals during a General Anesthesia induced by propofol and maintained by sevoflurane. As this setting is well-studied, the expected behavior of EEG is known, and could thus validate our approach in order to use it in other less studied configurations.

Dataset. We now apply TC-FISTA to the real dataset of 32-channel EEG signals recorded at 250 Hz during a General Anesthesia (GA) issued from [Dubost et al., 2019]. Compared to usual MEG datasets

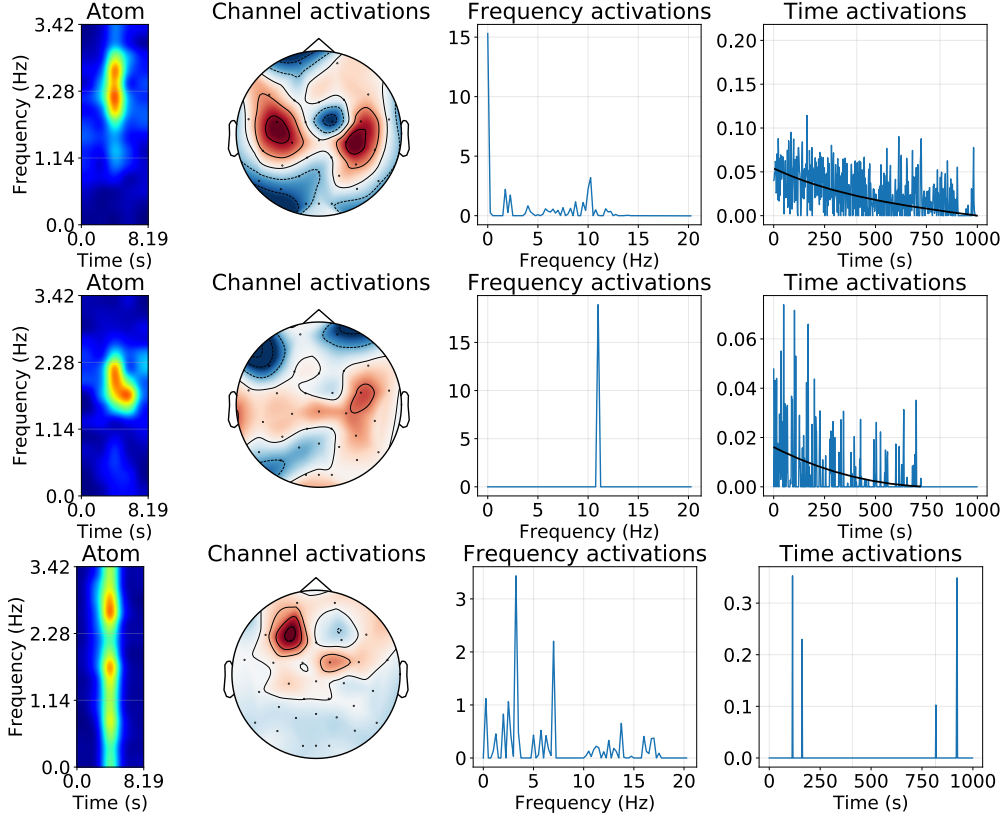


Figure 2: Three atoms of interest with their partial activations. From left to right: the time-frequency atom, the channel activations (mode 1) on the scalp draw with MNE [Gramfort et al., 2014], the frequency activations (mode 2), and the time activations (mode 3).

(see e.g. [Gramfort et al., 2013]), signals exhibit low SNR, impulsive noise and artefacts [Dubost et al., 2019]. We crop the signal to focus on the "Recovery of consciousness" (ROC), a crucial and difficult part of GA Purdon et al. [2013], for a resulting duration of 1000 seconds per signal. Following usual EEG preprocessing [Becker et al., 2010; Zhao et al., 2011], the signals were filtered using a bandpass filter between 1 and 20Hz, and then transformed using a Short time Fourier Transform representation with window size equals to 1024 samples and 50% overlap (bins of size 0.17Hz), resulting in a tensor \mathcal{Y} of size $(32 \times 82 \times 490)$. During a GA, patients are known to be static and EEG signals present a small number of patterns. As a consequence, we set $R = 3$, and learn $K = 4$ atoms of size $(1 \times 15 \times 5)$ corresponding to time-frequency atoms covering 8.19 seconds and a band of frequencies of 3.42Hz. Hyperparameters were chosen similarly to the Section 4.1. Finally, as previously mention, we add a non-negative constraint on the activations to fit the non-negativeness of spectrograms.

Results: α and δ waves. We present in Figure 2 three atoms² learnt by TC-FISTA, together with their partial activations. First, it is interesting to note that the learnt activations of the first two atoms (resp. atom 3) are rank 2 (resp. 1) instead of 3 – the actual constraint. This illustrates the robustness of the TC-FISTA representation with respect to mildly overestimated optimal rank– which may be explained by the combination of rank constraint and sparsity inducing penalty. Second, due to resulting low rank, these activations can be interpreted by looking at their behavior along the different modes (channel \times frequency \times time), as represented in Figure 2. Hence, by studying the frequency activations (mode 2), we see that the first two atoms encodes α and δ -waves, two important frequencies that occur during GA maintained by sevoflurane Purdon et al. [2013]. Similarly, we observe that their time activations (mode 3) are decreasing with a behavior that was suggested by previous work during the ROC phase [Purdon et al., 2013].

²In our experiments on this dataset, these atoms are consistently learnt by TC-FISTA, while the fourth atom varies and does not appear to encode any meaningful information.

Detecting impulsive noise. Interestingly, the third atom is of completely different nature, as it spans a wide range of frequencies. This atom may be seen as an encoding of impulsive noise, as its activations are mostly concentrated around one channel (mode 1). By using the activation of the third atom learnt by TC-FISTA, this impulsive noise can be removed from the final reconstruction by vanishing its contribution, highlighting the robust behavior of TC-FISTA.

Robustness to noise and reconstruction. By studying the channel activations (mode 1) of atoms 1 and 2, we observe the presence of three deficient channels: 10 (CP1), 21 (CP2), and 28 (F4). This is coherent with clinical context, as these channels are located in the frontal and central area – where individual anatomic variations can explain a contact defect – resulting in sensors recording only noise at these positions. Interestingly, due to the low-rank constraint, our model is robust to this problem – as it assumes strong correlation between the channels – and is able to reconstruct the missing or bad signal. This property emphasizes the robust nature of TC-FISTA with respect to unstructured noise and sensor defect.

5 Conclusion

In this paper, we developed a tensor based CSC framework, named Kruskal CSC, where signals, atoms and activations are assumed to be tensors, but activations are constrained to have a low CP-rank. We proposed an efficient optimisation strategy to solve this framework, named TC-FISTA, that entails a reduced complexity and lower computation time. We experimentally evaluated our method, and showed that our TC-FISTA is significantly more robust to noise and artefacts than its unconstrained counterparts, and as a result is able to produce a sparse, accurate and interpretable decomposition of noisy tensor signals such as EEG spectrograms recorded during propofol induced GA. Future work may include the use of TC-FISTA to other types of GA, as well as EEG recordings, or other types of tensor based signals such as multi-source music.

Acknowledgement

The authors would like to thank the Service d’Anesthésie-Réanimation of the Hôpital d’Instruction des Armées Bégin and especially Dr. Clément Dubost for the access to the data, the fruitful discussions, and interpretations of the recovered patterns.

References

- E. Acar, D. M. Dunlavy, and T. G. Kolda. A scalable optimization approach for fitting canonical tensor decompositions. *Journal of Chemometrics*, 25(2):67–86, 2011.
- Q. Barthélemy, A. Larue, A. Mayoue, D. Mercier, and J. I. Mars. Shift & 2d rotation invariant sparse coding for multivariate signals. *IEEE Transactions on Signal Processing*, 60(4):1597–1611, 2012.
- 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.
- H. Becker, P. Comon, L. Albera, M. Haardt, and I. Merlet. Multiway space-time-wave-vector analysis for source localization and extraction. In *2010 18th European Signal Processing Conference*, pages 1349–1353. IEEE, 2010.
- H. Becker, L. Albera, P. Comon, M. Haardt, G. Birot, F. Wendling, M. Gavaret, C. G. Bénar, and I. Merlet. Eeg extended source localization: tensor-based vs. conventional methods. *NeuroImage*, 96: 143–157, 2014.
- H. Becker, L. Albera, P. Comon, R. Gribonval, F. Wendling, and I. Merlet. Brain-source imaging: From sparse to tensor models. *IEEE Signal Processing Magazine*, 32(6):100–112, 2015.
- S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011.

- H. Bristow, A. Eriksson, and S. Lucey. Fast convolutional sparse coding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 391–398, 2013.
- R. Chalasani, J. C. Principe, and N. Ramakrishnan. A fast proximal method for convolutional sparse coding. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pages 1–5. IEEE, 2013.
- F. Cong, Q.-H. Lin, L.-D. Kuang, X.-F. Gong, P. Astikainen, and T. Ristaniemi. Tensor decomposition of eeg signals: a brief review. *Journal of Neuroscience Methods*, 248:59–69, 2015.
- C. Dubost, P. Humbert, A. Benizri, J.-P. Tourtier, N. Vayatis, and P.-P. Vidal. Selection of the best electroencephalogram channel to predict the depth of anesthesia. *Frontiers in Computational Neuroscience*, 13, 2019.
- D. E. Dudgeon. Multidimensional digital signal processing. *Engewood Cliffs*, 1983.
- B. A. Fritz, P. L. Kalarickal, H. R. Maybrier, M. R. Muench, D. Dearth, Y. Chen, K. E. Escallier, A. B. Abdallah, N. Lin, and M. S. Avidan. Intraoperative electroencephalogram suppression predicts postoperative delirium. *Anesthesia and analgesia*, 122(1):234, 2016.
- C. Garcia-Cardona and B. Wohlberg. Convolutional dictionary learning: A comparative review and new algorithms. *IEEE Transactions on Computational Imaging*, 4(3):366–381, 2018a.
- C. Garcia-Cardona and B. Wohlberg. Convolutional dictionary learning for multi-channel signals. In *2018 52nd Asilomar Conference on Signals, Systems, and Computers*, pages 335–342. IEEE, 2018b.
- R. Glowinski and A. Marroco. Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de dirichlet non linéaires. *Revue Française d’Automatique, Informatique, et Recherche Opérationnelle*, 9(R2):41–76, 1975.
- A. Gramfort, M. Luessi, E. Larson, D. A. Engemann, D. Strohmeier, C. Brodbeck, R. Goj, M. Jas, T. Brooks, L. Parkkonen, et al. Meg and eeg data analysis with mne-python. *Frontiers in Neuroscience*, 7:267, 2013.
- A. Gramfort, M. Luessi, E. Larson, D. A. Engemann, D. Strohmeier, C. Brodbeck, L. Parkkonen, and M. S. Hämäläinen. Mne software for processing meg and eeg data. *Neuroimage*, 86:446–460, 2014.
- L. He, K. Chen, W. Xu, J. Zhou, and F. Wang. Boosted sparse and low-rank tensor regression. In *Advances in Neural Information Processing Systems*, pages 1009–1018, 2018.
- C. Hildreth. A quadratic programming procedure. *Naval research logistics quarterly*, 4(1):79–85, 1957.
- P. Humbert, J. Audiffren, L. Oudre, and N. Vayatis. Low rank activations for tensor-based convolutional sparse coding. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3252–3256. IEEE, 2020.
- M. Jas, T. D. La Tour, U. Simsekli, and A. Gramfort. Learning the morphology of brain signals using alpha-stable convolutional sparse coding. In *Advances in Neural Information Processing Systems*, pages 1099–1108, 2017.
- G. Karvat, A. Schneider, M. Alyahyay, F. Steenbergen, M. Tangermann, and I. Diester. Real-time detection of neural oscillation bursts allows behaviourally relevant neurofeedback. *Communications biology*, 3(1):1–10, 2020.
- T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- J. Kossaifi, Y. Panagakis, A. Anandkumar, and M. Pantic. Tensorly: Tensor learning in python. *The Journal of Machine Learning Research*, 20(1):925–930, 2019.
- J. B. Kruskal. Rank, decomposition, and uniqueness for 3-way and n-way arrays. *Multway data analysis*, pages 7–18, 1989.
- T. D. La Tour, T. Moreau, M. Jas, and A. Gramfort. Multivariate convolutional sparse coding for electromagnetic brain signals. In *Advances in Neural Information Processing Systems*, pages 3292–3302, 2018.

- J. Liu, C. Garcia-Cardona, B. Wohlberg, and W. Yin. First-and second-order methods for online convolutional dictionary learning. *SIAM Journal on Imaging Sciences*, 11(2):1589–1628, 2018.
- N. Liu and J. Rinehart. Closed-loop propofol administration: routine care or a research tool? what impact in the future?, 2016.
- J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online dictionary learning for sparse coding. In *Proceedings of the 26th annual international conference on machine learning*, pages 689–696, 2009.
- 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.
- M. Nikolova and P. Tan. Alternating proximal gradient descent for nonconvex regularised problems with multiconvex coupling terms. 2017.
- P. Paatero. A weighted non-negative least squares algorithm for three-way ‘parafac’ factor analysis. *Chemometrics and Intelligent Laboratory Systems*, 38(2):223–242, 1997.
- N. Parikh, S. Boyd, et al. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014.
- A.-H. Phan, P. Tichavský, and A. Cichocki. Low rank tensor deconvolution. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2169–2173. IEEE, 2015.
- Y. Punjasawadwong, W. Chau-in, M. Laopaiboon, S. Punjasawadwong, and P. Pin-on. Processed electroencephalogram and evoked potential techniques for amelioration of postoperative delirium and cognitive dysfunction following non-cardiac and non-neurosurgical procedures in adults. *Cochrane Database of Systematic Reviews*, (5), 2018.
- P. L. Purdon, E. T. Pierce, E. A. Mukamel, M. J. Prerau, J. L. Walsh, K. F. K. Wong, A. F. Salazar-Gomez, P. G. Harrell, A. L. Sampson, A. Cimenser, et al. Electroencephalogram signatures of loss and recovery of consciousness from propofol. *Proceedings of the National Academy of Sciences*, 110(12): E1142–E1151, 2013.
- J. Quesada, P. Rodriguez, and B. Wohlberg. Separable dictionary learning for convolutional sparse coding via split updates. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4094–4098. IEEE, 2018.
- J. Quesada, G. Silva, P. Rodriguez, and B. Wohlberg. Combinatorial separable convolutional dictionaries. In *2019 XXII Symposium on Image, Signal Processing and Artificial Vision (STSIVA)*, pages 1–5. IEEE, 2019.
- G. Rabusseau and H. Kadri. Low-rank regression with tensor responses. In *Advances in Neural Information Processing Systems*, pages 1867–1875, 2016.
- R. Rigamonti, A. Sironi, V. Lepetit, and P. Fua. Learning separable filters. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2754–2761, 2013.
- G. Silva, J. Quesada, and P. Rodríguez. Efficient separable filter estimation using rank-1 convolutional dictionary learning. In *2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6. IEEE, 2018.
- D. Simon and M. Elad. Rethinking the csc model for natural images. In *Advances in Neural Information Processing Systems*, pages 2271–2281, 2019.
- A. Sironi, B. Tekin, R. Rigamonti, V. Lepetit, and P. Fua. Learning separable filters. *IEEE transactions on pattern analysis and machine intelligence*, 37(1):94–106, 2014.
- F. van Ede, A. J. Quinn, M. W. Woolrich, and A. C. Nobre. Neural oscillations: sustained rhythms or transient burst-events? *Trends in Neurosciences*, 41(7):415–417, 2018.
- Y. Wang, J. T. Kwok, and L. M. Ni. Generalized convolutional sparse coding with unknown noise. *IEEE Transactions on Image Processing*, 29:5386–5395, 2020.

- B. Wohlberg. Efficient algorithms for convolutional sparse representations. *IEEE Transactions on Image Processing*, 25(1):301–315, 2015.
- B. Wohlberg. Convolutional sparse representation of color images. In *2016 IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI)*, pages 57–60. IEEE, 2016.
- B. Wohlberg. SPORCO: A Python package for standard and convolutional sparse representations. In *Proceedings of the 15th Python in Science Conference*, pages 1–8, Austin, TX, USA, July 2017. doi: 10.25080/shinma-7f4c6e7-001.
- Q. Zhao, C. F. Caiafa, D. P. Mandic, L. Zhang, T. Ball, A. Schulze-Bonhage, and A. S. Cichocki. Multilinear subspace regression: An orthogonal tensor decomposition approach. In *Advances in Neural Information Processing Systems*, pages 1269–1277, 2011.
- H. Zhou, L. Li, and H. Zhu. Tensor regression with applications in neuroimaging data analysis. *Journal of the American Statistical Association*, 108(502):540–552, 2013.

APPENDIX

A Pseudo code and implementation

The pseudo-code of the \mathbf{Z} -step of TC-FISTA for one mode can be found in Algorithm 1. The implementation of TC-FISTA using Python is available in the file `code/TC_FISTA.py`, together with notebooks that provide illustrations of the experiments.

Algorithm 1 TC-FISTA (sub-problem)

Input: signal \mathcal{Y} , dictionary $\mathcal{D}_1, \dots, \mathcal{D}_K$, regularization and step parameters α, β, η ($\eta = 1/L$, the inverse of Lipschitz constant if calculate), tolerance ε

Initialization: $\mathbf{Z}^{(0)}$

Precompute: $\hat{\mathbf{Y}}, \{\hat{\mathbf{D}}_k\}, \mathbf{G}$ and $(\hat{\mathbf{A}} \otimes \mathbf{I})\hat{\mathbf{y}}^{(q)}$

$t^{(0)} \leftarrow 1$

repeat

▷ Update of \mathbf{W} via a proximal gradient step (ISTA)

Compute $\hat{\mathbf{Z}}^{(s)}$ using the DFT

$\hat{\mathbf{z}}^{(s)} \leftarrow \text{vec}(\hat{\mathbf{Z}}^{(s)})$

$\hat{\mathbf{w}}^{(s+1/2)} \leftarrow \hat{\mathbf{z}}^{(s)} - \eta \left(\mathbf{G}\hat{\mathbf{z}}^{(s)} - (\hat{\mathbf{A}} \otimes \mathbf{I})\hat{\mathbf{y}}^{(q)} \right)$

$\hat{\mathbf{W}}^{(s+1/2)} \leftarrow$ Matricization of $\hat{\mathbf{w}}^{(s+1/2)}$

Compute $\mathbf{W}^{(s+1/2)}$ using the IDFT

▷ Update of \mathbf{W} via a proximal step (ISTA)

$\mathbf{W}^{(s+1)} \leftarrow \text{prox}_{\eta, \alpha, \beta} \left(\mathbf{W}_k^{(s+1/2)} \right)$

▷ Nesterov momentum step (FISTA)

$t^{(s+1)} \leftarrow \frac{1 + \sqrt{1 + 4 \cdot t^{(s)^2}}}{2}$

$\mathbf{Z}^{(s+1)} \leftarrow \mathbf{W}^{(s+1)} + \frac{t^{(s)} - 1}{t^{(s+1)} + 1} (\mathbf{W}^{(s+1)} - \mathbf{W}^{(s)})$

until $\|\mathbf{Z}^{(s+1)} - \mathbf{Z}^{(s)}\|_\infty \leq \varepsilon$

B Additional results on synthetic data

To further illustrate and compare the effectiveness and efficiency of TC-FISTA, we consider in this section one additional (smaller) dataset and other noise levels. We also provide additional information on the data sets considered in the main paper and their generation.

For the convenience of the reader, we list here the CSC/CDL algorithms compared and the acronyms we use throughout this section: FISTA with tensor-based rank constraint (TC-FISTA), ADMM with iterative application of the Sherman-Morrison equation (FCSC-ShM) [Bristow et al., 2013; Wohlberg, 2015], FISTA in the Fourier domain (ConvFISTA-FD) [Chalasanani et al., 2013; Wohlberg, 2015].

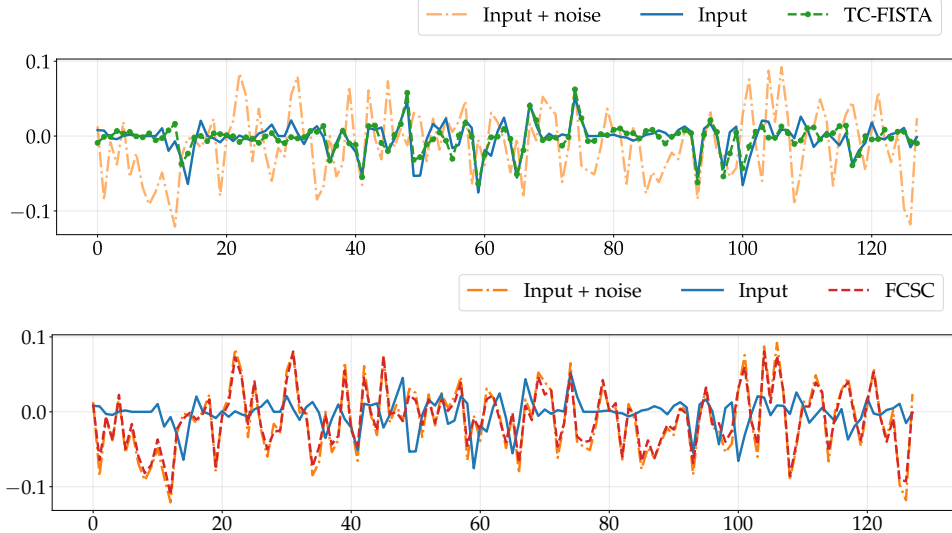


Figure 3: One tube of 3-rd order (Top) input and reconstruction with TC-FISTA, and (Bottom) input + noise and reconstruction with FCSC-ShM. ConvFISTA-FD leads to a similar reconstruction.

B.1 Setup

Dataset. Small-scale and large-scale experiments are performed by considering two main different datasets

- A *small-scale dataset* which contains 10 independent input signals of size $(25 \times 25 \times 25)$. Each signal is generated as follows. We draw $K = 3$ atoms of size $(5 \times 5 \times 5)$ according to an Uniform distribution with values in $[-1, 1]$ and normalize them. Then, we set the maximal CP-rank to $R^* \leq 2$ and draw the sparse activations from a Bernoulli-Uniform distribution with Bernoulli parameter equals to 0.2, and range of values in $[-1, 1]$. Finally, we generate the input tensor according to the convolutional model induced by the K-CDL model. Following Wohlberg [2015], we added a multivariate Gaussian noise to each signal to obtain a Signal to Noise Ratio (SNR) of 25dB, 10dB, and 5dB.
- A *large-scale dataset* – the one from the main paper – which is generated as the small-scale dataset but with input signals of size $(128 \times 128 \times 128)$ and Bernoulli parameter equals to 0.02.

Recall that the definition of the SNR between a signal \mathbf{y}_{ref} and a comparison one $\mathbf{y}_{noisy} = \mathbf{y}_{ref} + \varepsilon$ is

$$\text{SNR}(\mathbf{y}_{ref}, \mathbf{y}_{noisy}) = 10 \log_{10} \left(\frac{\text{Var}(\mathbf{y}_{ref})}{\text{MSE}(\mathbf{y}_{ref}, \mathbf{y}_{noisy})} \right),$$

where MSE denotes the Mean Squared Error. SNR is an asymmetric decibel measurement (dB) used to compare the level of a signal to the level of background noise [Wohlberg, 2015].

Metrics. We use four metrics to evaluate our methods

- The Root Mean Square Error (RMSE) between the true input signal (resp. the true activations) and the reconstruction (resp. the learnt activations). The lowest the better. We denote them $\text{RMSE}(\mathcal{Y})$ and $\text{RMSE}(\mathcal{Z})$.
- The number of times, among all the initializations, our method reaches a “correct” minimizers, e.g. RMSE under $\varepsilon = 1.e-6$. This metric reflects the dependency of an algorithm to its initializations. The higher the better. We denoted them $\#\{\text{RMSE}(\mathcal{Y}) < \varepsilon\}$ and $\#\{\text{RMSE}(\mathcal{Z}) < \varepsilon\}$ where ε need to be set.

B.2 Evaluation of the K-CSC

On this additional experiment, we compare the performances of TC-FISTA on the CSC task and show the importance of the rank structure to get rid of the noise. The true dictionary is therefore given at the beginning. The $\{\mathcal{Z}_{k,q}\}$ are initialized with random Uniform matrices.

Table 2: Results return on the CSC task on the two datasets with noise. For TC-FISTA, R is set to the true value, $R^* = 2$. Mean and standard deviation are reported.

| CSC SNR | R = 2 Metrics | <i>Noisy small-scale dataset</i> | <i>Noisy large-scale dataset</i> |
|------------|-------------------------------------|---|---|
| | | TC-FISTA | TC-FISTA |
| 25dB | RMSE(\mathcal{Y}) ↓ | 3.999 · e-4 (±4.427 · e-5) | 7.606 · e-5 (±1.253 · e-6) |
| | RMSE(\mathcal{Z}) ↓ | 2.403 · e-4 (±2.528 · e-5) | 4.469 · e-5 (±1.1666 · e-6) |
| | #{RMSE(\mathcal{Y}) < 1.e-3} ↑ | 94% | 90% |
| | #{RMSE(\mathcal{Z}) < 1.e-3} ↑ | 96% | 90% |
| 10dB | RMSE(\mathcal{Y}) ↓ | 2.492 · e-3 (±9.024 · e-5) | 4.958 · e-4 (±7.733 · e-6) |
| | RMSE(\mathcal{Z}) ↓ | 1.495 · e-3 (±5.066 · e-5) | 2.913 · e-4 (±7.194 · e-4) |
| | #{RMSE(\mathcal{Y}) < 2.5e-3} ↑ | 84% | 88% |
| | #{RMSE(\mathcal{Z}) < 2.5e-3} ↑ | 98% | 90% |
| 5dB | RMSE(\mathcal{Y}) ↓ | 4.847 · e-3 (±3.166 · e-4) | 9.918 · e-4 (±1.529 · e-5) |
| | RMSE(\mathcal{Z}) ↓ | 2.894 · e-3 (±1.805 · e-4) | 5.828 · e-4 (±1.434 · e-5) |
| | #{RMSE(\mathcal{Y}) < 5.e-3} ↑ | 52% | 88% |
| | #{RMSE(\mathcal{Z}) < 4.e-3} ↑ | 86% | 88% |
| CSC SNR | Metrics | <i>Noisy small-scale dataset</i> | <i>Noisy large-scale dataset</i> |
| | | ConvFISTA-FD [Chalasanani et al., 2013; Wohlberg, 2015] | ConvFISTA-FD [Chalasanani et al., 2013; Wohlberg, 2015] |
| 25dB | RMSE(\mathcal{Y}) ↓ | 2.109 · e-3 (±2.547 · e-4) | 1.732 · e-3 (±8.703 · e-6) |
| | RMSE(\mathcal{Z}) ↓ | 1.311 · e-3 (±2.027 · e-4) | 1.050 · e-3 (±2.791 · e-5) |
| | #{RMSE(\mathcal{Y}) < 1.e-3} ↑ | 0% | 0% |
| | #{RMSE(\mathcal{Z}) < 1.e-3} ↑ | 0% | 10% |
| 10dB | RMSE(\mathcal{Y}) ↓ | 6.673 · e-3 (±6.847 · e-4) | 6.689 · e-3 (±3.344 · e-4) |
| | RMSE(\mathcal{Z}) ↓ | 4.367 · e-3 (±5.919 · e-4) | 4.406 · e-3 (±3.010 · e-3) |
| | #{RMSE(\mathcal{Y}) < 2.5e-3} ↑ | 0% | 0% |
| | #{RMSE(\mathcal{Z}) < 2.5e-3} ↑ | 0% | 0% |
| 5dB | RMSE(\mathcal{Y}) ↓ | 1.209 · e-2 (±1.202 · e-3) | 1.211 · e-2 (±6.242 · e-4) |
| | RMSE(\mathcal{Z}) ↓ | 7.774 · e-3 (±1.009 · e-3) | 7.813 · e-3 (±5.251 · e-4) |
| | #{RMSE(\mathcal{Y}) < 5.e-3} ↑ | 0% | 0% |
| | #{RMSE(\mathcal{Z}) < 4.e-3} ↑ | 0% | 0% |
| CSC SNR | Metrics | <i>Noisy small-scale dataset</i> | <i>Noisy large-scale dataset</i> |
| | | FCSC-ShM [Bristow et al., 2013; Wohlberg, 2015] | FCSC-ShM [Bristow et al., 2013; Wohlberg, 2015] |
| 25dB | RMSE(\mathcal{Y}) ↓ | 2.292 · e-3 (±1.220 · e-5) | 1.732 · e-3 (±8.707 · e-6) |
| | RMSE(\mathcal{Z}) ↓ | 1.454 · e-3 (±7.326 · e-5) | 1.050 · e-3 (±2.794 · e-6) |
| | #{RMSE(\mathcal{Y}) < 1.e-3} ↑ | 0% | 0% |
| | #{RMSE(\mathcal{Z}) < 1.e-3} ↑ | 0% | 10% |
| 10dB | RMSE(\mathcal{Y}) ↓ | 6.734 · e-3 (±7.117 · e-4) | 6.689 · e-3 (±3.344 · e-5) |
| | RMSE(\mathcal{Z}) ↓ | 4.393 · e-3 (±6.060 · e-4) | 4.405 · e-3 (±3.011 · e-3) |
| | #{RMSE(\mathcal{Y}) < 2.5e-3} ↑ | 0% | 0% |
| | #{RMSE(\mathcal{Z}) < 2.5e-3} ↑ | 0% | 0% |
| 5dB | RMSE(\mathcal{Y}) ↓ | 1.215 · e-2 (±1.252 · e-3) | 1.211 · e-2 (±6.243 · e-4) |
| | RMSE(\mathcal{Z}) ↓ | 7.800 · e-3 (±1.033 · e-3) | 7.812 · e-3 (±5.252 · e-4) |
| | #{RMSE(\mathcal{Y}) < 5.e-3} ↑ | 0% | 0% |
| | #{RMSE(\mathcal{Z}) < 4.e-3} ↑ | 0% | 0% |

Results with noise. For each input signal, we run TC-FISTA with five different initializations. The metric $\#\{\text{RMSE}(\cdot) < \varepsilon\}$ is therefore calculated on 50 initializations. Each time, the reconstruction giving the lowest RMSE(\mathcal{Y}) among the five tries is kept. We set $R = 2$ during all the experiments.

Quantitative results are collected in Table 2. The most remarkable result is that even under a very strong noise (SNR of 5dB), TC-FISTA yields good reconstructions and is not too dependant on the initialization as shown by the two last metrics. Figure 3 provides a visual example of this important property. We see that TC-FISTA reconstructs the input signal with high accuracy while FCSC-ShM completely fails and mostly overfits the noise. This was expected for TC-FISTA because the noise does not share the low-rank structure of the signal. Hence, the K-CSC model which includes a low-rank constraint “manages” to not capture it and thus recovers the true signal with accuracy. In other words, taking into account the low-rank structure of the signal eliminates the noise and allows a better recovery of the activations. Furthermore, notice that since for both datasets $R^* = 2$, the larger the signal, the more restrictive the rank constraint is. This leads to lower RMSEs on the large-scale dataset than on the small-scale one.

C Additional results on the EEG dataset

We provide additional results and illustrations for the EEG dataset.

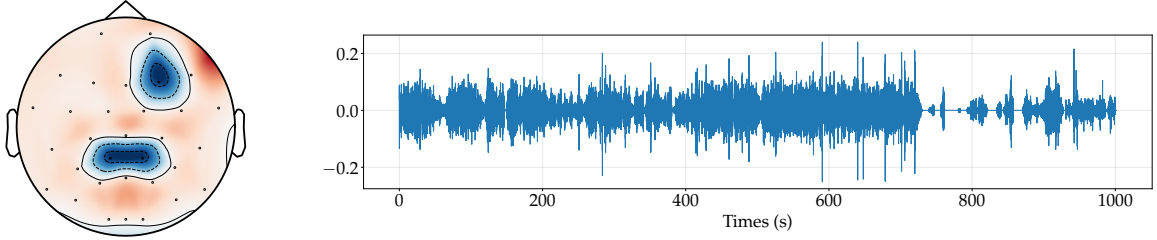


Figure 4: (Left) Spatial activations of the on of the learnt atom before removing the bad channels. (Right) Raw signal at one of the bad channel.

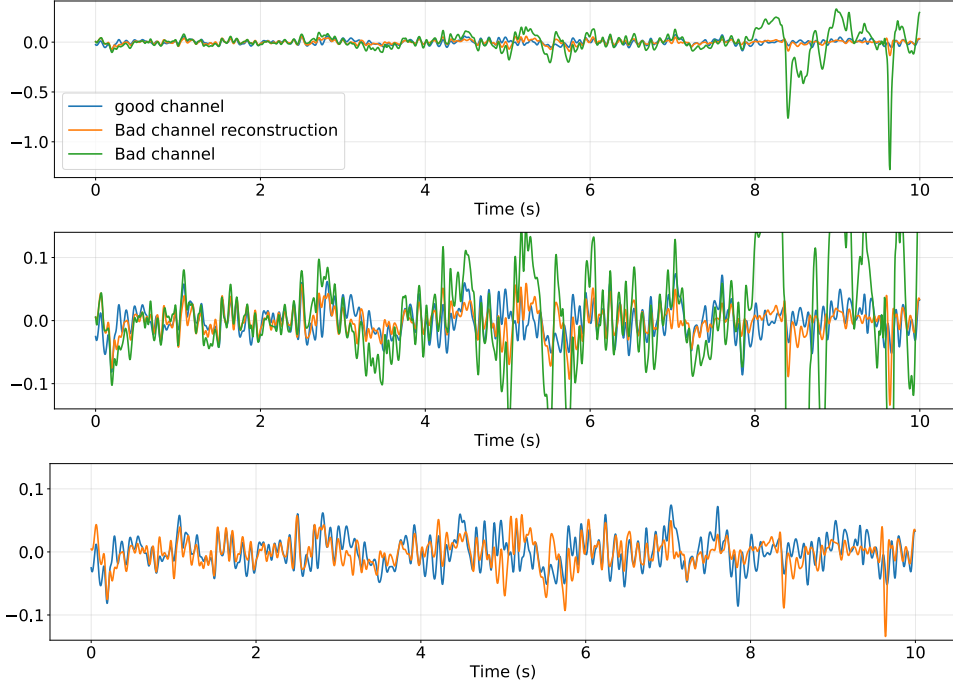


Figure 5: On the top, raw signal of a good channel (blue), a bad channel (green), and a reconstruction of the bad channel (orange). On the middle, a closer look to the good channel. On the bottom, only the good channel (blue) and the reconstruction of the bad channel (orange).

C.1 Identification of bad channels

As stated in the main paper, the K-CDL model provides a highly interpretable decomposition of activations. In our setting, the first mode corresponds to the “channel activations”. Looking at this mode, we see from the activation of one of the learnt atom that three common channels are defective (see Figure 4). Fortunately, due to the rank constraint, we are able to reconstruct these channels. To emphasize this phenomenon, we choose a non-deficient channel which is spatially close to a deficient one. We assume that this good channel approximately reflects the values that the bad one would have taken. Therefore, we expect that the reconstruction of the bad channel is similar to the values of the good one. An illustration of the reconstruction is given in Figure 5 meaning that the two band of frequencies become less important with time.

C.2 δ and α -waves during a general anesthesia

From the second mode (relative to the frequency domain), we have identified atoms link to δ and α -waves, the two most important band of frequencies in General Anesthesia (GA) [Purdon et al., 2013]. We know from [Purdon et al., 2013] that the importance of these bands during the Recovery Of Consciousness (ROC) tends to decrease in time. Interestingly, this phenomenon is captured by TC-FISTA. Indeed, Figure 6 shows that the time activations related to these two bands decreases with time.

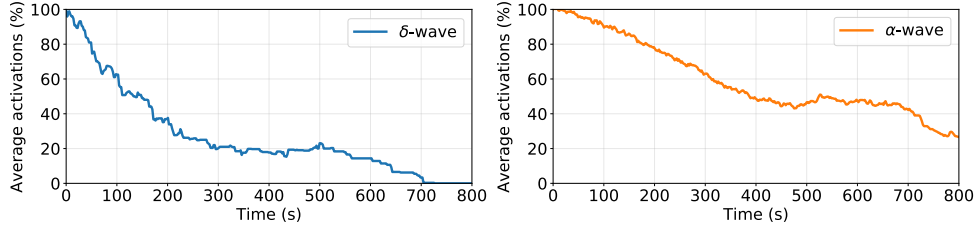


Figure 6: Evolution of the time activations for the first and second atoms of Figure 8 which are relative to the δ and α waves.

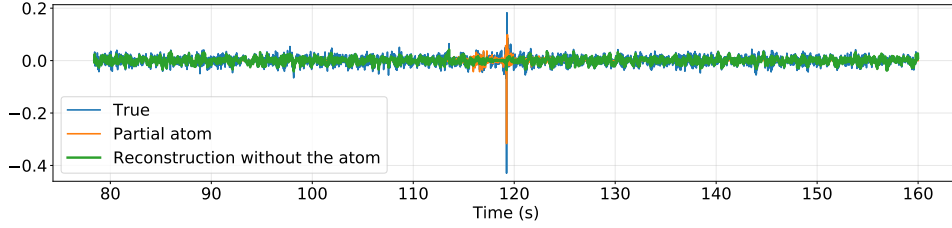


Figure 7: Reconstruction of the signal when one atom relative to the impulsion noise (at 120s) is removed.

C.3 Removing atoms from the final reconstruction

In the main paper, we identified the third atom as an atom relative to some impulsion noise. Thanks to its identification, we can remove it from the final reconstruction. We therefore obtain a signal which do not contain this impulsion noise. An illustration is given in Figure 7. Note that, this particular noise is low-rank. Hence, the K-CDL is not robust to it, contrary to the Gaussian noise from the synthetic section. Atoms and activations are displayed in Figure 8.

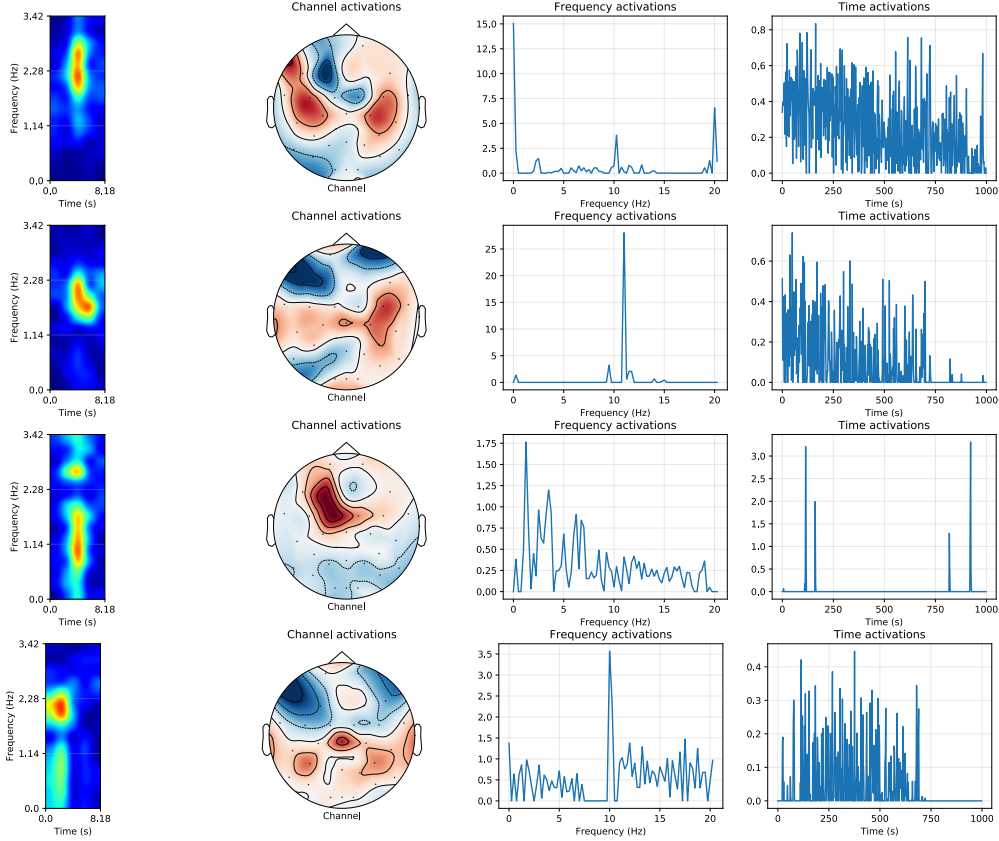


Figure 8: The four atoms with their partial activations. From left to right: the time-frequency atom, the channel activations (mode 1) on the scalp, the frequency activations (mode 2), and the time activations (mode 3). Scale not the same for each figure.

D Technical proofs

This section provides the technical proofs of the different propositions exposed in the main paper. To be as complete as possible, we recall well-known but important definitions and theoretical properties from multidimensional Fourier analysis. We begin with two important definitions.

Definition 2. (Discrete Fourier Transform (DFT)) – Let consider a function \mathcal{F} defined on $\{0 \dots, N_1 - 1\} \times \dots \times \{0 \dots, N_p - 1\}$ with period (N_1, \dots, N_p) . The Discrete Fourier Transform (DFT) of \mathcal{F} is given by

$$\hat{\mathcal{F}}[k_1, \dots, k_p] = \sum_{n_1=0}^{N_1-1} \dots \sum_{n_p=0}^{N_p-1} \mathcal{F}[n_1, \dots, n_p] \exp \left(-i2\pi \left(\frac{k_1 n_1}{N_1}, \dots, \frac{k_p n_p}{N_p} \right) \right),$$

and the Inverse DFT (IDFT) of $\hat{\mathcal{F}}$ is given by

$$\mathcal{F}[n_1, \dots, n_p] = \left(\frac{1}{\prod_{i=1}^p N_i} \right) \cdot \sum_{k_1=0}^{N_1-1} \dots \sum_{k_p=0}^{N_p-1} \hat{\mathcal{F}}[k_1, \dots, k_p] \exp \left(i2\pi \left(\frac{k_1 n_1}{N_1}, \dots, \frac{k_p n_p}{N_p} \right) \right).$$

Let now consider the periodization of two discrete function \mathcal{F} and \mathcal{G} ,

$$\begin{aligned} \tilde{\mathcal{F}}[n_1, \dots, n_p] &= \mathcal{F}[n_1 \bmod N_1, \dots, n_p \bmod N_p] \\ \tilde{\mathcal{G}}[n_1, \dots, n_p] &= \mathcal{G}[n_1 \bmod N_1, \dots, n_p \bmod N_p]. \end{aligned}$$

The two functions $\tilde{\mathcal{F}}$ and $\tilde{\mathcal{G}}$ are now two discrete functions with period (N_1, \dots, N_p) (each of the modes are periodic one-dimensional signals). The *circular convolution* is defined as follow.

Definition 3. (Circular discrete convolution) – Let consider two functions \mathcal{F}, \mathcal{G} defined on $\{0 \dots, N_1 - 1\} \times \dots \times \{0 \dots, N_p - 1\}$ with both a period of (N_1, \dots, N_p) . The circular convolution between $\tilde{\mathcal{F}}$ and $\tilde{\mathcal{G}}$ is given by

$$(\tilde{\mathcal{F}} \otimes \tilde{\mathcal{G}})[n_1, \dots, n_p] = \sum_{k_1=0}^{N_1-1} \dots \sum_{k_p=0}^{N_p-1} \tilde{\mathcal{F}}[k_1, \dots, k_p] \tilde{\mathcal{G}}[n_1 - k_1, \dots, n_p - k_p].$$

$\tilde{\mathcal{F}} \otimes \tilde{\mathcal{G}}$ is a signal of period (N_1, \dots, N_p) and can be decomposed in a Fourier basis like classical periodic signals which give rises to the following important theorem.

Theorem 2. (Discrete convolution theorem) – If \mathcal{F} and \mathcal{G} have period (N_1, \dots, N_p) , then the DFT of $\mathcal{H} = \mathcal{F} \otimes \mathcal{G}$ is

$$\widehat{\mathcal{H}}[n_1, \dots, n_p] = \widehat{\mathcal{F}}[n_1, \dots, n_p] * \widehat{\mathcal{G}}[n_1, \dots, n_p], \quad \text{or in tensor notation} \quad \widehat{\mathcal{H}} = \widehat{\mathcal{F}} * \widehat{\mathcal{G}},$$

where $*$ is the component-wise product or Hadamard product.

We now present a simple lemma showing the important advantage of separable signals (i.e. low-rank signal) over non-separable ones in term of complexity.

Lemma 1. (Mode-wise DFT) – Let \mathcal{X} be a tensor in \mathbb{X} with CP-decomposition $\llbracket \mathbf{X}^{(1)}, \dots, \mathbf{X}^{(p)} \rrbracket$. The DFT of \mathcal{X} can be performed mode-wise i.e.

$$\widehat{\mathcal{X}} = \sum_{r=1}^R \widehat{\mathbf{x}}_r^{(1)} \circ \dots \circ \widehat{\mathbf{x}}_r^{(p)} = \llbracket \widehat{\mathbf{X}}^{(1)}, \dots, \widehat{\mathbf{X}}^{(p)} \rrbracket, \quad (7)$$

where the DFT is taken along each columns of the factors. The complexity of the computation of $(\widehat{\mathbf{X}}^{(1)}, \dots, \widehat{\mathbf{X}}^{(p)})$ using the Fast Fourier Transform algorithm (FFT) is of $\mathcal{O}(R \sum_{i=1}^p n_i \log(n_i))$ instead of $\mathcal{O}(\prod_{i=1}^p n_i \log(\prod_{i=1}^p n_i))$.

Proof. Using the definition of both the DFT and the CP-decomposition, the proof is straightforward. Furthermore, as we only perform 1-D FFT, we obtain the given complexity. \square

To prove the new results, we will extensively used matricization and vectorization techniques. One important formula in tensor algebra is given by the above proposition.

Proposition 2. (Matricization of the Kruskal operator [Kolda and Bader, 2009]) – Let \mathcal{X} be a tensor in \mathbb{X} with CP-decomposition $\llbracket \mathbf{X}^{(1)}, \dots, \mathbf{X}^{(p)} \rrbracket$. Then,

$$\mathbf{X}_{(q)} = \mathbf{X}_q (\mathbf{X}_p \odot \dots \odot \mathbf{X}_{q+1} \odot \mathbf{X}_{q-1} \odot \dots \odot \mathbf{X}_1)^\top = \mathbf{X}_q \left(\overset{\leftarrow p}{\odot}_{i=1, i \neq q} \mathbf{X}_i \right)^\top,$$

where \odot is the Khatri–Rao product and $\overset{\leftarrow p}{\odot}_{i=1}$ denotes the product of p Khatri–Rao products in reverse order. We can also also vectorized this formula which gives

$$\begin{aligned} \text{vec}(\mathbf{X}_{(q)}) &= (\mathbf{X}_p \odot \dots \odot \mathbf{X}_{q+1} \odot \mathbf{X}_{q-1} \odot \dots \odot \mathbf{X}_1 \otimes \mathbf{I}_{n_q}) \text{vec}(\mathbf{X}_q) \\ &= \left(\overset{\leftarrow p}{\odot}_{i=1, i \neq q} \mathbf{X}_i \otimes \mathbf{I}_{n_q} \right) \text{vec}(\mathbf{X}_q), \end{aligned}$$

where \mathbf{I}_{n_q} is an identity matrix of size $(n_q \times n_q)$.

Lemma 2. (Equality in the Fourier domain) – The fidelity term $f(\cdot)$ from Equation (3) of the main paper is equals in the Fourier domain to

$$f(\{\mathbf{Z}_{k,q}\}_{k=1}^K) = \frac{1}{2 \prod_{i=1}^p N_i} \left\| \widehat{\mathcal{Y}} - \sum_{k=1}^K \widehat{\mathcal{D}}_k * \llbracket \widehat{\mathbf{Z}}_k^{(1)}, \dots, \widehat{\mathbf{Z}}_k^{(p)} \rrbracket \right\|_F^2. \quad (8)$$

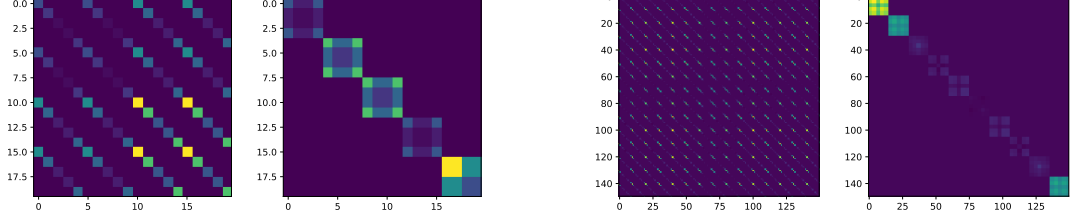


Figure 9: Visualization of $G = (\widehat{\mathbf{A}}^H \otimes \mathbf{I}) \widehat{\mathbf{\Gamma}}^H \widehat{\mathbf{\Gamma}} (\widehat{\mathbf{A}} \otimes \mathbf{I})$ before and after a reordering. The two left matrices correspond to the Gram matrix without and with reordering. The two right matrices also correspond to the Gram matrix without and with reordering but for a higher dimension.

Proof. The proof rests on several equalities and properties.

$$\begin{aligned}
& \frac{1}{2} \left\| \mathbf{y} - \sum_{k=1}^K \mathcal{D}_k \otimes \sum_{r=1}^R \mathbf{z}_{k,r}^{(1)} \circ \dots \circ \mathbf{z}_{k,r}^{(p)} \right\|_F^2 \\
&= \frac{1}{2 \prod_{i=1}^p N_i} \left\| \widehat{\mathbf{y}} - \sum_{k=1}^K DFT(\mathcal{D}_k \otimes \sum_{r=1}^R \mathbf{z}_{k,r}^{(1)} \circ \dots \circ \mathbf{z}_{k,r}^{(p)}) \right\|_F^2 \quad (\text{Parseval's theorem - Plancherel}) \\
&= \frac{1}{2 \prod_{i=1}^p N_i} \left\| \widehat{\mathbf{y}} - \sum_{k=1}^K \widehat{\mathcal{D}}_k * \sum_{r=1}^R DFT(\mathbf{z}_{k,r}^{(1)} \circ \dots \circ \mathbf{z}_{k,r}^{(p)}) \right\|_F^2 \quad (\text{Convolution theorem (2)}) \\
&= \frac{1}{2 \prod_{i=1}^p N_i} \left\| \widehat{\mathbf{y}} - \sum_{k=1}^K \widehat{\mathcal{D}}_k * \sum_{r=1}^R \widehat{\mathbf{z}}_{k,r}^{(1)} \circ \dots \circ \widehat{\mathbf{z}}_{k,r}^{(p)} \right\|_F^2 \quad (\text{Lemma (1)}) \\
&= \frac{1}{2 \prod_{i=1}^p N_i} \left\| \widehat{\mathbf{y}} - \sum_{k=1}^K \widehat{\mathcal{D}}_k * [\widehat{\mathbf{Z}}_k^{(1)}, \dots, \widehat{\mathbf{Z}}_k^{(p)}] \right\|_F^2 \quad (\text{Kruskal operator}).
\end{aligned}$$

□

We now state the main theorem of the paper.

Theorem 1. (Compact vectorized formulation) – Let $\widehat{\mathbf{y}}^{(q)}$ (resp. $\widehat{\mathbf{d}}_k^{(q)}$) be the vectorization of the folding of $\widehat{\mathbf{y}}$ (resp. $\widehat{\mathcal{D}}_k$) along the dimension q , $\widehat{\mathbf{Z}}_k^{(q)} = [\widehat{\mathbf{Z}}_k^{(q)}(:, 1) \mid \dots \mid \widehat{\mathbf{Z}}_k^{(q)}(:, R)]$ be the DFT of $\mathbf{Z}_k^{(q)}$ along its columns and $\widehat{\mathbf{z}}_k^{(q)} = \text{vec}(\widehat{\mathbf{Z}}_k^{(q)})$. Then

$$f(\{\mathbf{Z}_k^{(r)}\}_{k=1, r=1}^{K,p}) \propto \frac{1}{2} \left\| \widehat{\mathbf{y}}^{(q)} - \widehat{\mathbf{\Gamma}}^{(q)} (\widehat{\mathbf{A}}^{(q)} \otimes \mathbf{I}^{(q)}) \widehat{\mathbf{z}}^{(q)} \right\|_F^2, \quad (9)$$

where $\widehat{\mathbf{z}}^{(q)} = [\widehat{\mathbf{z}}_1^{(q)\top}, \dots, \widehat{\mathbf{z}}_K^{(q)\top}]^\top \in \mathbb{C}^{KRn_q}$, $\widehat{\mathbf{\Gamma}}^{(q)} = [\text{diag}(\widehat{\mathbf{d}}_1^{(q)}), \dots, \text{diag}(\widehat{\mathbf{d}}_K^{(q)})] \in \mathbb{C}^{n_1 \dots n_p \times Kn_1 \dots n_p}$, $\mathbf{I}^{(q)} \in \mathbb{R}^{n_q \times n_q}$ is the identity matrix and

$$\widehat{\mathbf{A}}^{(q)} = \begin{pmatrix} \widehat{\mathbf{B}}_1^{(q)} & & 0 \\ & \ddots & \\ 0 & & \widehat{\mathbf{B}}_K^{(q)} \end{pmatrix} \in \mathbb{C}^{K \prod_{1, i \neq q} n_i \times KR} \quad \text{with} \quad \widehat{\mathbf{B}}_k^{(q)} = \left(\overset{\leftarrow p}{\odot}_{i=1, i \neq q} \widehat{\mathbf{Z}}_k^{(i)} \right).$$

Proof. The proof mainly rests on the Proposition 2 and on the formulation of the previous lemma.

$$\begin{aligned}
& \left\| \widehat{\mathbf{y}} - \sum_{k=1}^K \widehat{\mathcal{D}}_k * [\widehat{\mathbf{Z}}_{k,1}, \dots, \widehat{\mathbf{Z}}_{k,p}] \right\|_F^2 = \left\| \widehat{\mathbf{Y}}^{(q)} - \sum_{k=1}^K \widehat{\mathcal{D}}_k^{(q)} * \widehat{\mathbf{Z}}_{k,q} \left(\overset{\leftarrow p}{\odot}_{i=1} \widehat{\mathbf{Z}}_k^{(i)} \right)^\top \right\|_F^2 \quad (\text{matricization}) \\
& = \left\| \widehat{\mathbf{y}}^{(q)} - \sum_{k=1}^K \widehat{\mathbf{d}}_k^{(q)} * \left(\overset{\leftarrow p}{\odot}_{i=1} \widehat{\mathbf{Z}}_k^{(i)} \otimes I \right) \text{vec}(\widehat{\mathbf{Z}}_{k,q}) \right\|_F^2 \quad (\text{vectorization}) \\
& = \left\| \widehat{\mathbf{y}}^{(q)} - \sum_{k=1}^K \text{diag}(\widehat{\mathbf{d}}_k^{(q)}) \left(\overset{\leftarrow p}{\odot}_{i=1} \widehat{\mathbf{Z}}_k^{(i)} \otimes I \right) \text{vec}(\widehat{\mathbf{Z}}_{k,q}) \right\|_F^2 \quad (\mathbf{x} * \mathbf{y} = \text{diag}(\mathbf{x})\mathbf{y}) \\
& = \left\| \widehat{\mathbf{y}}^{(q)} - \sum_{k=1}^K \text{diag}(\widehat{\mathbf{d}}_k^{(q)}) \widehat{\mathbf{C}}_k \widehat{\mathbf{z}}_k \right\|_F^2,
\end{aligned}$$

where the last line is just notations. To obtain the final equality, we stack the matrices $\{\text{diag}(\widehat{\mathbf{d}}_k^{(q)})\}$ and construct a block-diagonal matrix such that the block are the $\{\widehat{\mathbf{C}}_k\}$. Finally we obtain the following equality.

$$\begin{pmatrix} \widehat{\mathbf{C}}_1 & & \\ & \ddots & \\ & & \widehat{\mathbf{C}}_K \end{pmatrix} = \begin{pmatrix} \widehat{\mathbf{B}}_1 \otimes I & & \\ & \ddots & \\ & & \widehat{\mathbf{B}}_K \otimes I \end{pmatrix} = \begin{pmatrix} \widehat{\mathbf{B}}_1 & & \\ & \ddots & \\ & & \widehat{\mathbf{B}}_K \end{pmatrix} \otimes I,$$

with $\widehat{\mathbf{B}}_k = \left(\overset{\leftarrow p}{\odot}_{i=1, i \neq q} \widehat{\mathbf{Z}}_{k,i} \right)$. This end the proof. \square

Corollary 1. (Gradient of f) – With the notation of Theorem 2, the gradient of f with respect to $\mathbf{z}^{(q)} = [\text{vec}(\mathbf{Z}_1^{(q)})^\top, \dots, \text{vec}(\mathbf{Z}_K^{(q)})^\top]^\top$ is given by

$$\nabla_{\mathbf{z}^{(q)}} f \left(\{\mathbf{Z}_k^{(r)}\}_{k=1, r=1}^{K,p} \right) = \text{IDFT} \left[\left((\widehat{\mathbf{A}}^{(q)} \otimes I) \widehat{\mathbf{\Gamma}}^{(q)} \right)^H \left(\widehat{\mathbf{\Gamma}}^{(q)} (\widehat{\mathbf{A}}^{(q)} \otimes I) \widehat{\mathbf{z}}^{(q)} - \widehat{\mathbf{y}}^{(q)} \right) \right], \quad (10)$$

where $\text{IDFT}[\cdot]$ stands for the Inverse Discrete Fourier Transform.

Proof. See [Liu et al., 2018] Section 2.3.2. for a complete and detailed proof. \square

Proposition 1. The matrix $\mathbf{G} \triangleq (\widehat{\mathbf{A}}^H \otimes I) \widehat{\mathbf{\Gamma}}^H \widehat{\mathbf{\Gamma}} (\widehat{\mathbf{A}} \otimes I)$ can be obtained by computing K^2 blocks $\mathbf{G}_{k,\ell}, 1 \leq k, \ell \leq K$, where

$$\mathbf{G}_{k,\ell} = \left(\left(\overset{\leftarrow p}{\odot}_{i=1, i \neq q} \widehat{\mathbf{Z}}_{k,i} \right)^H \otimes I \right) \overline{\text{diag}(\widehat{\mathbf{d}}_k^{(q)})} \text{diag}(\widehat{\mathbf{d}}_\ell^{(q)}) \left(\left(\overset{\leftarrow p}{\odot}_{i=1, i \neq q} \widehat{\mathbf{Z}}_{\ell,i} \right) \otimes I \right), \quad (11)$$

and each of these blocks can be computed in $\mathcal{O}(R^2 \prod_{i=1, i \neq q}^p n_i)$ operations.

Proof. The first step of the proof requires to write $\widehat{\mathbf{\Gamma}}$ as the Kronecker product of two specific matrices in order to use the equality $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC} \otimes \mathbf{BD})$. Recall that $\widehat{\mathbf{\Gamma}}$ is a block-diagonal matrix, i.e. $\widehat{\mathbf{\Gamma}} = [\text{diag}(\widehat{\mathbf{d}}_1^{(q)}), \dots, \text{diag}(\widehat{\mathbf{d}}_K^{(q)})]$. We can decompose each diagonal-block, $\text{diag}(\widehat{\mathbf{d}}_k^{(q)})$, into smaller diagonal matrices using the kronecker product as follow

$$\text{diag}(\widehat{\mathbf{d}}_k^{(q)}) = \sum_{i=1}^{N_{\setminus q}} \text{diag}(e_i) \otimes \Delta_{k,i} \quad \text{with} \quad N_{\setminus q} = \prod_{i=1, i \neq q}^p n_i,$$

where $\text{diag}(e_i) \in \mathbb{R}^{N_{\setminus q} \times N_{\setminus q}}$ and $\Delta_{k,i} \in \mathbb{C}^{n_q \times n_q}$ being the i -th diagonal block of $\text{diag}(\widehat{\mathbf{d}}_k^{(q)})$ (i.e. $\Delta_{k,i} = \text{diag}(\widehat{\mathbf{d}}_k^{(q)})_{(i \cdot n_q : (i+1) \cdot n_q), (i \cdot n_q : (i+1) \cdot n_q)}$). As $(\text{diag}(e_i) \otimes \Delta_{k,i})$ is decomposed into two matrices of the proper

dimension, we can use the equality $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC} \otimes \mathbf{BD})$ leading to

$$\begin{aligned}
& \left((\odot_{i=1, i \neq q}^p \widehat{\mathbf{Z}}_{k,i})^H \otimes I \right) \overline{\text{diag}(\widehat{\mathbf{d}}_k^{(q)})} \text{diag}(\widehat{\mathbf{d}}_\ell^{(q)}) \left((\odot_{i=1, i \neq q}^p \widehat{\mathbf{Z}}_{\ell,i}) \otimes I \right) \\
&= \left(\widehat{\mathbf{B}}_k^H \otimes I \right) \sum_{i=1}^{N_{\setminus q}} (\text{diag}(e_i) \otimes \overline{\Delta_{k,i}}) \sum_{j=1}^{N_{\setminus q}} (\text{diag}(e_j) \otimes \Delta_{\ell,j}) \left(\widehat{\mathbf{B}}_\ell \otimes I \right) \\
&= \sum_{i=1}^{N_{\setminus q}} \sum_{j=1}^{N_{\setminus q}} \left(\widehat{\mathbf{B}}_k^H \otimes I \right) (\text{diag}(e_i) \otimes \overline{\Delta_{k,i}}) (\text{diag}(e_j) \otimes \Delta_{\ell,j}) \left(\widehat{\mathbf{B}}_\ell \otimes I \right) \\
&= \sum_{i=1}^{N_{\setminus q}} \sum_{j=1}^{N_{\setminus q}} \left(\widehat{\mathbf{B}}_k^H \text{diag}(e_i) \text{diag}(e_j) \widehat{\mathbf{B}}_\ell \otimes \overline{\Delta_{k,i}} \Delta_{\ell,j} \right) \\
&= \sum_{i=1}^{N_{\setminus q}} \left(\widehat{\mathbf{B}}_k^H \text{diag}(e_i) \text{diag}(e_i) \widehat{\mathbf{B}}_\ell \otimes \overline{\Delta_{k,i}} \Delta_{\ell,i} \right) \\
&= \sum_{i=1}^{N_{\setminus q}} \left((\text{diag}(e_i) \widehat{\mathbf{B}}_k)^H \text{diag}(e_i) \widehat{\mathbf{B}}_\ell \otimes \overline{\Delta_{k,i}} \Delta_{\ell,i} \right) = \sum_{i=1}^{N_{\setminus q}} \left(\widehat{\mathbf{B}}_k(i, :) \circ \widehat{\mathbf{B}}_\ell(i, :) \otimes \overline{\Delta_{k,i}} \Delta_{\ell,i} \right).
\end{aligned}$$

The outer product of two vectors in $\mathbb{C}^{1 \times R}$ is of complexity $\mathcal{O}(R^2)$. This product is made for each $1 \leq i \leq N_{\setminus q}$ and for each K^2 blocks. Hence, the overall complexity is $\mathcal{O}((KR)^2 \prod_{i=1, i \neq q}^p n_i)$. In addition, as this matrix has a particular block structure, i.e. this is a band-matrix (see Figure 9), its product with a vector of size KRn_q is of small complexity equals to $\mathcal{O}((KR)^2 n_q)$. \square