



HAL
open science

Stochastic Subsampling for Factorizing Huge Matrices

Arthur Mensch, Julien Mairal, Bertrand Thirion, Gaël Varoquaux

► **To cite this version:**

Arthur Mensch, Julien Mairal, Bertrand Thirion, Gaël Varoquaux. Stochastic Subsampling for Factorizing Huge Matrices. 2017. hal-01431618v1

HAL Id: hal-01431618

<https://hal.science/hal-01431618v1>

Preprint submitted on 11 Jan 2017 (v1), last revised 30 Oct 2017 (v3)

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.

Stochastic Subsampling for Factorizing Huge Matrices

Arthur Mensch
Inria Parietal
Saclay, France
arthur.mensch@m4x.org

Julien Mairal
Inria Thoth
Grenoble, France
julien.mairal@inria.fr

Bertrand Thirion
Inria Parietal
Saclay, France
bertrand.thirion@inria.fr

Gaël Varoquaux
Inria Parietal
Saclay, France
gael.varoquaux@inria.fr

January 11, 2017

Abstract

We present a matrix-factorization algorithm that scales to input matrices with both huge number of rows and columns. Learned factors may be sparse or dense and/or non-negative, which makes our algorithm suitable for dictionary learning, sparse component analysis, and non-negative matrix factorization. Our algorithm streams matrix columns while subsampling them to iteratively learn the matrix factors. At each iteration, the row dimension of a new sample is reduced by subsampling, resulting in lower time complexity compared to a simple streaming algorithm. Our method comes with convergence guarantees to reach a stationary point of the matrix-factorization problem. We demonstrate its efficiency on massive functional Magnetic Resonance Imaging data (2 TB), and on patches extracted from hyperspectral images (103 GB). For both problems, which involve different penalties on rows and columns, we obtain significant speed-ups compared to state-of-the-art algorithms.

Index words. Matrix factorization, dictionary learning, non-negative, stochastic optimization, majorization minimization, randomized methods, functional MRI, hyperspectral imaging

1 Introduction

Matrix factorization is a flexible approach to uncover latent factors in low-rank or sparse models. With sparse factors, it is used in dictionary learning, and has proven very effective for denoising and visual feature encoding in signal and computer vision [see e.g., 1]. When the data admit a low-rank structure, matrix factorization has proven very powerful for various tasks such as matrix completion [2, 3], word embedding [4, 5], or network models [6]. It is flexible enough to accommodate a large set of constraints and regularizations, and has gained significant attention in scientific domains where interpretability is a key aspect, such as genetics [7] and neuroscience [8]. In this paper, our goal is to adapt matrix-factorization techniques to huge-dimensional datasets, i.e., with large number of columns n and large number of rows p . Specifically, our work is motivated by the rapid increase in sensor resolution, as in hyperspectral imaging or fMRI, and the challenge that the resulting high-dimensional signals pose to current algorithms.

As a widely-used model, the literature on matrix factorization is very rich and two main classes of formulations have emerged. The first one addresses a convex-optimization problem with a penalty promoting low-rank structures, such as the trace or max norms [2]. This formulation has strong theoretical

The research leading to these results was supported by the ANR (MACARON project, ANR-14-CE23-0003-01 NiConnect project, ANR-11-BINF-0004NiConnect).

guarantees [3], but lacks scalability for huge datasets or sparse factors. For these reasons, our paper is focused on a second type of approach, which relies on nonconvex optimization. Stochastic (or online) optimization methods have been developed in this setting. Unlike classical alternate minimization procedures, they learn matrix decompositions by observing a single matrix column (or row) at each iteration. In other words, they stream data along one matrix dimension. Their cost per iteration is significantly reduced, leading to faster convergence in various practical contexts. More precisely, two approaches have been particularly successful: stochastic gradient descent [9] and stochastic majorization-minimization methods [10, 11]. The former has been widely used for matrix completion [see 12, 13, 14, and references therein], while the latter has been used for dictionary learning with sparse and/or structured regularization [15]. Despite those efforts, stochastic algorithms for dictionary learning are currently unable to deal efficiently with matrices that are large in both dimensions.

We propose a new matrix-factorization algorithm that can handle such matrices. It builds upon the stochastic majorization-minimization framework of [10], which we generalize for our problem. In this framework, the objective function is minimized by iteratively improving an upper-bound surrogate of the function (*majorization* step) and minimizing it to obtain new estimates (*minimization* step). The core idea of our algorithm is to approximate these steps to perform them faster. We carefully introduce and control approximations, so to extend [10] convergence results when neither the majorization nor the minimization step is performed exactly.

For this purpose, we borrow ideas from *randomized* methods in machine learning and signal processing. Indeed, quite orthogonally to stochastic optimization, efficient approaches to tackle the growth of dataset dimension have exploited random projections [16, 17] or sampling, reducing data dimension while preserving signal content. Large-scale datasets often have an intrinsic dimension which is significantly smaller than their ambient dimension. Good examples are biological datasets [18] and physical acquisitions with an underlying sparse structure enabling compressed sensing [19]. In this context, models can be learned using only random data summaries, also called sketches. For instance, randomized methods [see 20, for a review] are efficient to compute PCA [21], a classic matrix-factorization approach, and to solve constrained or penalized least-square problems [22, 23]. On a theoretical level, recent works on *sketching* [24, 25] have provided bounds on the risk of using random summaries in learning.

Using random projections as a pre-processing step is not appealing in our applicative context since factors learned on reduced data are not interpretable. On the other hand, it is possible to exploit *random sampling* to approximate the steps of online matrix factorization. Factors are learned in the original space whereas the dimension of each iteration is reduced together with the computational cost per iteration.

Contribution

The contribution of this paper is both practical and theoretical. We introduce a new matrix factorization algorithm, called *subsampling online matrix factorization* (SOMF), which is faster than state-of-the-art algorithms by an order of magnitude on large real-world datasets (hyperspectral images, large fMRI data). It leverages random sampling with stochastic optimization to learn sparse and dense factors more efficiently. To prove the convergence of SOMF, we extend the stochastic majorization-minimization framework [10] and make it robust to some time-saving approximations. We then show convergence guarantees for SOMF under reasonable assumptions. Finally, we propose an extensive empirical validation of the subsampling approach.

In a first version of this work [26] presented at the International Conference in Machine Learning (ICML), we proposed an algorithm similar to SOMF, without any theoretical guarantees. The algorithm that we present here has such guarantees, which we express in a more general framework, stochastic majorization-minimization. It is validated for new sparsity settings and a new domain of application. An open-source efficient Python package is provided.

Notations

Matrices are written using bold capital letters and vectors using bold small letters (e.g., \mathbf{X} , $\boldsymbol{\alpha}$). We use superscript to specify the column (sample or component) number, and write $\mathbf{X} = [\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}]$. We use subscripts to specify the *iteration* number, as in \mathbf{x}_t . The floating bar, as in \bar{g}_t , is used to stress that a given value is an average over iterations, or an expectation. The superscript $*$ is used to denote an exact value, when it has to be compared to an inexact value, e.g., to compare $\boldsymbol{\alpha}_t^*$ (exact) to $\boldsymbol{\alpha}_t$ (approximation).

2 Prior art: matrix factorization with stochastic majorization-minimization

Below, we introduce the matrix-factorization problem and recall a specific stochastic algorithm to solve it observing one column (or a minibatch) at every iteration. We cast this algorithm in the stochastic majorization-minimization framework [10], which we will use in the convergence analysis.

2.1 Problem statement

In our setting, the goal of matrix factorization is to decompose a matrix $\mathbf{X} \in \mathbb{R}^{p \times n}$ — typically n signals of dimension p — as a product of two smaller matrices:

$$\mathbf{X} \approx \mathbf{D}\mathbf{A} \quad \text{with} \quad \mathbf{D} \in \mathbb{R}^{p \times k} \text{ and } \mathbf{A} \in \mathbb{R}^{k \times n},$$

with potential sparsity or structure requirements on \mathbf{D} and \mathbf{A} . In signal processing, sparsity is often enforced on the code \mathbf{A} , in a problem called *dictionary learning* [27]. In such a case, \mathbf{D} is called the “dictionary” and \mathbf{A} the sparse code. We use this terminology throughout the paper.

Learning the factorization is typically performed by minimizing a quadratic data-fitting term, with constraints and/or penalties over the code and the dictionary:

$$\min_{\substack{\mathbf{D} \in \mathcal{C} \\ \mathbf{A} \in \mathbb{R}^{k \times n}}} \sum_{i=1}^n \frac{1}{2} \|\mathbf{x}^{(i)} - \mathbf{D}\boldsymbol{\alpha}^{(i)}\|_2^2 + \lambda \Omega(\boldsymbol{\alpha}^{(i)}), \quad (1)$$

where $\mathbf{A} \triangleq [\boldsymbol{\alpha}^{(1)}, \dots, \boldsymbol{\alpha}^{(n)}]$, \mathcal{C} is a column-wise separable convex set of $\mathbb{R}^{p \times k}$ and $\Omega : \mathbb{R}^p \rightarrow \mathbb{R}$ is a penalty over the code. Both constraint set and penalty may enforce structure or sparsity, though \mathcal{C} has traditionally been used as a technical requirement to ensure that the penalty on \mathbf{A} does not vanish with \mathbf{D} growing arbitrarily large. Two choices of \mathcal{C} and Ω are of particular interest. The problem of dictionary learning sets \mathcal{C} as the ℓ_2 ball for each atom and Ω to be the ℓ_1 norm. Due to the sparsifying effect of ℓ_1 penalty [28], the dataset admits a *sparse* representation in the dictionary. On the opposite, finding a *sparse set* in which to represent a given dataset, with a goal akin to sparse PCA [29], requires to set as the ℓ_1 ball for each atom and Ω to be the ℓ_2 norm. Our work considers the *elastic-net* constraints and penalties [30], which encompass both special cases. Fixing ν and μ in $[0, 1]$, we denote by $\Omega(\cdot)$ and $\|\cdot\|$ the elastic-net penalty in \mathbb{R}^p and \mathbb{R}^k :

$$\begin{aligned} \Omega(\boldsymbol{\alpha}) &\triangleq (1 - \nu)\|\boldsymbol{\alpha}\|_1 + \frac{\nu}{2}\|\boldsymbol{\alpha}\|_2^2, \\ \mathcal{C} &\triangleq \left\{ \mathbf{D} \in \mathbb{R}^{p \times k} / \|\mathbf{d}^{(j)}\| \triangleq (1 - \mu)\|\mathbf{d}^{(j)}\|_1 + \frac{\mu}{2}\|\mathbf{d}^{(j)}\|_2^2 \leq 1 \right\}. \end{aligned} \quad (2)$$

Following [15], we can also enforce the positivity of \mathbf{D} and/or \mathbf{A} by replacing \mathbb{R} by \mathbb{R}^+ in \mathcal{C} , and adding positivity constraints on \mathbf{A} in (1), as in non-negative sparse coding [31]. We rewrite (1) as an empirical risk minimization problem depending on the dictionary only. The matrix \mathbf{D} solution of (1) is indeed obtained by minimizing the empirical risk \bar{f}

$$\mathbf{D} \in \operatorname{argmin}_{\mathbf{D} \in \mathcal{C}} \left(\bar{f}(\mathbf{D}) \triangleq \frac{1}{n} \sum_{i=1}^n f(\mathbf{D}, \mathbf{x}^{(i)}) \right), \quad (3)$$

$$\text{where } f(\mathbf{D}, \mathbf{x}) \triangleq \min_{\boldsymbol{\alpha} \in \mathbb{R}^k} \frac{1}{2} \|\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}\|_2^2 + \lambda \Omega(\boldsymbol{\alpha}),$$

and the matrix \mathbf{A} is obtained by solving the linear regression

$$\min_{\mathbf{A} \in \mathbb{R}^{k \times n}} \sum_{i=1}^n \frac{1}{2} \|\mathbf{x}^{(i)} - \mathbf{D}\boldsymbol{\alpha}^{(i)}\|_2^2 + \lambda \Omega(\boldsymbol{\alpha}^{(i)}). \quad (4)$$

The problem (1) is non-convex in the parameters (\mathbf{D}, \mathbf{A}) , and hence (3) is not convex. However, the problem (1) is convex in both \mathbf{D} and \mathbf{A} when fixing one variable and optimizing with respect to the other. As such, it is naturally solved by alternate minimization over \mathbf{D} and \mathbf{A} , which asymptotically provides a stationary point of (3). Yet \mathbf{X} has typically to be observed hundred of times before obtaining a good dictionary. Alternate minimization is therefore not adapted to datasets with many samples. These are often rich enough to learn a good decomposition in one or a few passes.

Algorithm 1 Online matrix factorization (OMF) [15]

Input: Initial iterate \mathbf{D}_0 , sample stream $(\mathbf{x}_t)_{t>0}$, number of iterations T .

for t from 1 to T **do**

 Draw $\mathbf{x}_t \sim \mathcal{P}$.

 Compute $\boldsymbol{\alpha}_t = \operatorname{argmin}_{\boldsymbol{\alpha} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{x}_t - \mathbf{D}_{t-1} \boldsymbol{\alpha}\|_2^2 + \lambda \Omega(\boldsymbol{\alpha})$.

 Update the *parameters* of aggregated surrogate \bar{g}_t :

$$\begin{aligned}\bar{\mathbf{C}}_t &= \left(1 - \frac{1}{t}\right) \bar{\mathbf{C}}_{t-1} + \frac{1}{t} \boldsymbol{\alpha}_t \boldsymbol{\alpha}_t^\top. \\ \bar{\mathbf{B}}_t &= \left(1 - \frac{1}{t}\right) \bar{\mathbf{B}}_{t-1} + \frac{1}{t} \mathbf{x}_t \boldsymbol{\alpha}_t^\top.\end{aligned}\tag{8}$$

 Compute (using block coordinate descent):

$$\mathbf{D}_t = \operatorname{argmin}_{\mathbf{D} \in \mathcal{C}} \frac{1}{2} \operatorname{Tr}(\mathbf{D}^\top \mathbf{D} \bar{\mathbf{C}}_t) - \operatorname{Tr}(\mathbf{D}^\top \bar{\mathbf{B}}_t).$$

Output: Final iterate \mathbf{D}_T .

2.2 Online matrix factorization

When \mathbf{X} has a large number of columns but a limited number of rows, the stochastic optimization method of [15] outputs a good dictionary much more rapidly than alternate-minimization. In this setting [see 32], learning the dictionary is naturally formalized as an expected risk minimization

$$\min_{\mathbf{D} \in \mathcal{C}} \bar{f}(\mathbf{D}) \triangleq \mathbb{E}_{\mathbf{x}}[f(\mathbf{D}, \mathbf{x})],\tag{5}$$

where \mathbf{x} is drawn from the data distribution and forms an i.i.d. stream $(\mathbf{x}_t)_t$. In the finite-sample setting, (5) reduces to (3) when \mathbf{x}_t is drawn uniformly at random from $\{\mathbf{x}^{(i)}, i \in [1, n]\}$. We then write i_t the sample number selected at time t .

The online matrix factorization algorithm proposed in [15] draws a sample \mathbf{x}_t at each iteration, and uses it to improve the current iterate \mathbf{D}_{t-1} . For this, it first computes the code $\boldsymbol{\alpha}_t$ associated to \mathbf{x}_t on the current dictionary:

$$\boldsymbol{\alpha}_t \triangleq \operatorname{argmin}_{\boldsymbol{\alpha} \in \mathbb{R}^k} \frac{1}{2} \|\mathbf{x}_t - \mathbf{D}_{t-1} \boldsymbol{\alpha}\|_2^2 + \lambda \Omega(\boldsymbol{\alpha}).\tag{6}$$

Then, it updates \mathbf{D}_t to make it optimal in reconstructing past samples $(\mathbf{x}_s)_{s \leq t}$ from previously computed codes $(\boldsymbol{\alpha}_s)_{s \leq t}$:

$$\mathbf{D}_t \in \operatorname{argmin}_{\mathbf{D} \in \mathcal{C}} \left(\bar{g}_t(\mathbf{D}) \triangleq \frac{1}{t} \sum_{s=1}^t \frac{1}{2} \|\mathbf{x}_s - \mathbf{D} \boldsymbol{\alpha}_s\|_2^2 + \lambda \Omega(\boldsymbol{\alpha}_s) \right).\tag{7}$$

Importantly, minimizing \bar{g}_t is equivalent to minimizing the quadratic function

$$\mathbf{D} \rightarrow \frac{1}{2} \operatorname{Tr}(\mathbf{D}^\top \mathbf{D} \bar{\mathbf{B}}_t^\top) - \operatorname{Tr}(\mathbf{D}^\top \bar{\mathbf{C}}_t),\tag{9}$$

where $\bar{\mathbf{B}}_t$ and $\bar{\mathbf{C}}_t$ are small matrices that summarize previously seen samples and codes:

$$\bar{\mathbf{B}}_t = \frac{1}{t} \sum_{s=1}^t \mathbf{x}_s \boldsymbol{\alpha}_s^\top \quad \bar{\mathbf{C}}_t = \frac{1}{t} \sum_{s=1}^t \boldsymbol{\alpha}_s \boldsymbol{\alpha}_s^\top.\tag{10}$$

As the constraints \mathcal{C} have a separable structure per atom, [15] uses projected block coordinate descent to minimize \bar{g}_t . The function gradient writes $\nabla \bar{g}_t(\mathbf{D}) = \mathbf{D} \bar{\mathbf{C}}_t - \bar{\mathbf{B}}_t$, and it is therefore enough to maintain $\bar{\mathbf{B}}_t$ and $\bar{\mathbf{C}}_t$ in memory to solve (7). $\bar{\mathbf{B}}_t$ and $\bar{\mathbf{C}}_t$ are updated online, using the rules (8) (Alg. 1).

The function \bar{g}_t is an upper-bound surrogate of the true current empirical risk, whose definition involves the regression minima computed on current dictionary \mathbf{D} :

$$\bar{f}_t(\mathbf{D}) \triangleq \frac{1}{t} \sum_{s=1}^t \min_{\boldsymbol{\alpha} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{x}_s - \mathbf{D} \boldsymbol{\alpha}\|_2^2 + \lambda \Omega(\boldsymbol{\alpha}) \leq \bar{g}_t(\mathbf{D}).\tag{11}$$

Algorithm 2 Stochastic majorization-minimization [SMM 10]

Input: Initial iterate θ_0 , weight sequence $(w_t)_{t>0}$, sample stream $(\mathbf{x}_t)_{t>0}$, number of iteration T .

for t from 1 to T **do**

 Draw $x_t \sim \mathcal{P}$, get $f_t : \theta \in \Theta \rightarrow f(\mathbf{x}_t, \theta)$.

 Construct a surrogate of f_t near θ_{t-1} , that meets

$$g_t \geq f_t, \quad g_t(\theta_{t-1}) = f_t(\theta_{t-1}). \quad (12)$$

 Update the aggregated surrogate:

$$\bar{g}_t = (1 - w_t)\bar{g}_{t-1} + w_t g_t.$$

 Compute

$$\theta_t = \underset{\theta \in \Theta}{\operatorname{argmin}} \bar{g}_t(\theta). \quad (13)$$

Output: Final iterate θ_T .

Using empirical processes theory [33], it is possible to show that minimizing \bar{f}_t at each iteration asymptotically yields a stationary point of the expected risk (5). Unfortunately, minimizing (11) is expensive as it involves the computation of optimal current codes for every previously seen sample at each iteration, which boils down to naive alternate-minimization.

In contrast, \bar{g}_t is much cheaper to minimize than \bar{f}_t , using block coordinate descent. It is possible to show that \bar{g}_t converges towards a locally tight upper-bound of the objective \bar{f}_t and that minimizing \bar{g}_t at each iteration also asymptotically yields a stationary point of the expected risk (5). This establishes the correctness of the *online matrix factorization* algorithm (OMF), which we summarize in Alg. 1. In practice, the OMF algorithm performs a single pass of block coordinate descent: the minimization step is inexact. This heuristic will be justified by our theoretical contribution in Section 4.

Extensions

For efficiency, it is essential to use mini-batches $\{\mathbf{x}_s, s \in \mathcal{T}_t\}$ of size η instead of single samples in the iterations [15]. The surrogate parameters $\bar{\mathbf{B}}_t, \bar{\mathbf{C}}_t$ are then updated by the mean value of $\{(\mathbf{x}_s \boldsymbol{\alpha}_s^\top, \boldsymbol{\alpha}_s \boldsymbol{\alpha}_s^\top)\}_{s \in \mathcal{T}_t}$ over the batch. The optimal size of the mini-batches is usually close to k . (8) uses the sequence of weights $(\frac{1}{t})_t$ to update parameters $\bar{\mathbf{B}}_t$ and $\bar{\mathbf{C}}_t$. [15] replaces these weights with a sequence $(w_t)_t$, which can decay more slowly to give more importance to recent samples in \bar{g}_t . These weights will prove important in our analysis.

2.3 Stochastic majorization-minimization

Online matrix factorization belongs to a wider category of algorithms introduced in [10] that minimize locally tight upper-bounding surrogates instead of a more complex objective, in order to solve an expected risk minimization problem. Generalizing online matrix factorization, we introduce the *stochastic majorization-minimization* (SMM) algorithm, which is at the core of our theoretical contribution.

In online matrix factorization, the true empirical risk functions \bar{f}_t and their *surrogates* \bar{g}_t follow the update rules, with generalized weight $(w_t)_t$ set to $(\frac{1}{t})_t$ in (7) – (11):

$$\bar{f}_t \triangleq (1 - w_t)\bar{f}_{t-1} + w_t f_t, \quad \bar{g}_t \triangleq (1 - w_t)\bar{g}_{t-1} + w_t g_t, \quad (14)$$

where the *pointwise* loss function and its surrogate are

$$\begin{aligned} f_t(\mathbf{D}) &\triangleq \min_{\boldsymbol{\alpha} \in \mathbb{R}^k} \frac{1}{2} \|\mathbf{x}_t - \mathbf{D}\boldsymbol{\alpha}\|_2^2 + \lambda\Omega(\boldsymbol{\alpha}), \\ g_t(\mathbf{D}) &\triangleq \frac{1}{2} \|\mathbf{x}_t - \mathbf{D}\boldsymbol{\alpha}_t\|_2^2 + \lambda\Omega(\boldsymbol{\alpha}_t). \end{aligned} \quad (15)$$

The function g_t is a majorizing surrogate of f_t : $g_t \geq f_t$, and g_t is tangent to f_t in \mathbf{D}_{t-1} , i.e., $g_t(\mathbf{D}_{t-1}) = f_t(\mathbf{D}_{t-1})$ and $\nabla(g_t - f_t)(\mathbf{D}_{t-1}) = 0$. At each step of online matrix factorization:

- The surrogate g_t is computed along with α_t , using (6).
- The parameters $\bar{\mathbf{B}}_t, \bar{\mathbf{C}}_t$ are updated following (8). They define the *aggregated* surrogate \bar{g}_t up to a constant.
- The quadratic function \bar{g}_t is minimized efficiently by block coordinate descent, using parameters $\bar{\mathbf{B}}_t$ and $\bar{\mathbf{C}}_t$ to compute its gradient.

The stochastic majorization-minimization framework simply formalizes the three steps above, for a larger variety of loss functions $f_t(\theta) \triangleq f(\theta, \mathbf{x}_t)$, where θ is the parameter we want to learn (\mathbf{D} in the online matrix factorization setting). At iteration t , a surrogate g_t of the loss f_t is computed to update the aggregated surrogate \bar{g}_t following (14). The surrogate functions $(g_t)_t$ should be upper-bounds of loss functions $(f_t)_t$, tight in the current iterate θ_{t-1} (e.g., the dictionary \mathbf{D}_{t-1}). This simply means that $f_t(\theta_{t-1}) = g_t(\theta_{t-1})$ and $\nabla(f_t - g_t)(\theta_{t-1}) = 0$. Computing \bar{g}_t can be done if g_t is defined simply, as in OMF where it is linearly parametrized by $(\alpha_t \alpha_t^\top, \mathbf{x}_t \alpha_t^\top)$. \bar{g}_t is then minimized to obtain a new iterate θ_t .

It can be shown following [10] that stochastic majorization-minimization algorithms find asymptotical stationary point of the expected risk $\mathbb{E}_{\mathbf{x}}[f(\theta, \mathbf{x})]$ under mild assumptions recalled in Section 4. The generic stochastic majorization-minimization algorithm is presented in Alg. 2, where it comes across as a generalization of online matrix factorization. SMM admits the same minibatch and decaying weight extensions (used in Alg. 2) as online matrix factorization.

In this work, we extend the SMM framework and allow both majorization and minimization steps to be approximated. As a side contribution, our extension proves that performing a single pass of block coordinate descent to update the dictionary, an important heuristic in [15], is indeed correct. We first introduce the new matrix factorization algorithm at the core of this paper and then present the extended SMM framework.

3 Stochastic subsampling for high dimensional data decomposition

The online algorithm presented in Section 2 is very efficient to factorize matrices that have a large number of columns (i.e., with a large number of samples n), but a reasonable number of rows — the dataset is not very high dimensional. However, it is not designed to deal with very high number of rows: the cost of a single iteration depends linearly on p . On terabyte-scale datasets from fMRI with $p = 2 \cdot 10^5$ features, the original online algorithm requires several weeks to reach convergence. This is a major motivation for designing new matrix factorization algorithms that scale in *both directions*.

In the large-sample regime $p \gg k$, the underlying dimensionality of columns may be much lower than the actual p : the rows of a single column drawn at random are therefore correlated and redundant. This guides us on how to scale online matrix factorization with regard to the number of rows:

- The online algorithm OMF uses a *single* column of (or mini-batch) of \mathbf{X} at each iteration to enrich the average surrogate and update the *whole* dictionary.
- We go a step beyond and use a *fraction* of a single column of \mathbf{X} to refine a fraction of the dictionary.

More precisely, we draw a column and observe only *some* of its rows at each iteration, to refine the same rows of the dictionary, as illustrated in Figure 1. To take into account all features from the dataset, rows are selected at random at each iteration: we call this technique *stochastic subsampling*. Stochastic subsampling reduces the efficiency of the dictionary update *per iteration*, as less information is incorporated in the current iterate \mathbf{D}_t . On the other hand, with a correct design, the cost of a single iteration can be considerably reduced, as it grows with the number of observed features. Section 5 shows that the proposed algorithm is an order of magnitude faster than the original OMF on large and redundant datasets.

First, we formalize the idea of working with a fraction of the p rows at a single iteration. We adapt the online matrix factorization algorithm, to reduce the iteration cost by a factor close to the ratio of selected rows. This defines a new online algorithm, called *subsampling online matrix factorization* (SOMF). At each iteration, it uses q rows of the column \mathbf{x}_t to update the sequence of iterates $(\mathbf{D}_t)_t$. As in Section 2, we introduce a more general algorithm, *stochastic approximate majorization-minimization* (SAMM), of

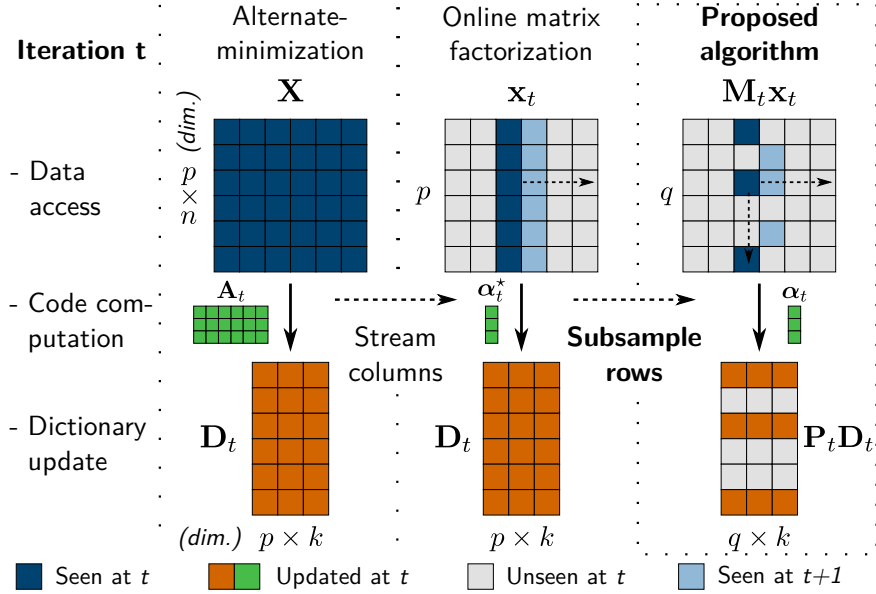


Figure 1: *Stochastic subsampling* further improves online matrix factorization to handle datasets with large number of columns and rows. \mathbf{X} is the input $p \times n$ matrix, \mathbf{D}_t and \mathbf{A}_t are respectively the dictionary and code at time t .

which SOMF is an instance. It extends the stochastic majorization-minimization framework, with similar theoretical guarantees but potentially faster convergence.

3.1 Subsampled online matrix factorization

Formally, as in online matrix factorization, we consider a sample stream $(\mathbf{x}_t)_t$ in \mathbb{R}^p that cycles into a finite sample set $\{\mathbf{x}^{(i)}, i \in [1, n]\}$. We minimize the empirical risk (3).

3.1.1 Stochastic subsampling and algorithm outline

We want to reduce the time complexity of a single iteration. In the original algorithm, the complexity depends linearly on the sample dimension p in three aspects:

- $\mathbf{x}_t \in \mathbb{R}^p$ is used to compute the code α_t ,
- it is used to update the surrogate parameters $\bar{\mathbf{B}}_t \in \mathbb{R}^{p \times k}$,
- $\mathbf{D}_t \in \mathbb{R}^{p \times k}$ is fully updated at each iteration.

Our algorithm reduces the dimensionality of these steps at each iteration, such that p becomes $q = \frac{p}{r}$ in the time complexity analysis, where $r > 1$ is a *reduction factor*. Formally, we randomly draw, at iteration t , a mask \mathbf{M}_t that “selects” a random subset of \mathbf{x}_t . We use it to drop a part of the features of \mathbf{x}_t and to “freeze” these features in dictionary \mathbf{D} at iteration t .

It is convenient to consider \mathbf{M}_t as a $\mathbb{R}^{p \times p}$ random diagonal matrix, such that each coefficient is a Bernoulli variable with parameter $\frac{1}{r}$, normalized to be 1 in expectation. $\forall j \in [0, p - 1]$,

$$\mathbb{P}[\mathbf{M}_t[j, j] = r] = \frac{1}{r}, \quad \mathbb{P}[\mathbf{M}_t[j, j] = 0] = 1 - \frac{1}{r}. \quad (16)$$

Thus, r describes the average proportion of observed features and $\mathbf{M}_t \mathbf{x}_t$ is a non-biased, low-dimensional estimator of \mathbf{x}_t :

$$\mathbb{E}[\|\mathbf{M}_t \mathbf{x}_t\|_0] = \frac{p}{r} = q \quad \mathbb{E}[\mathbf{M}_t \mathbf{x}_t] = \mathbf{x}_t. \quad (17)$$

with $\|\cdot\|_0$ counting the number of non-zero coefficients. We define the pair of orthogonal projectors $\mathbf{P}_t \in \mathbb{R}^{q \times p}$ and $\mathbf{P}_t^\perp \in \mathbb{R}^{(p-q) \times p}$ that project \mathbb{R}^p onto $\text{Im}(\mathbf{M}_t)$ and $\text{Ker}(\mathbf{M}_t)$. In other words, $\mathbf{P}_t \mathbf{Y}$ and

Algorithm 3 Subsampled online matrix factorization (SOMF)

Input: Initial iterate \mathbf{D}_0 , weight sequences $(w_t)_{t>0}$, $(\gamma_c)_{c>0}$, sample set $\{\mathbf{x}^{(i)}\}_{i>0}$, number of iterations T .

for t from 1 to T **do**

 Draw $\mathbf{x}_t = \mathbf{x}^{(i)}$ at random and \mathbf{M}_t following (16).

 Update the regression parameters for sample i :

$$\begin{aligned} c^{(i)} &\leftarrow c^{(i)} + 1, & \gamma &\leftarrow \gamma_{c^{(i)}}. \\ \boldsymbol{\beta}_t^{(i)} &\leftarrow (1 - \gamma)\mathbf{G}_{t-1}^{(i)} + \gamma\mathbf{D}_{t-1}^\top \mathbf{M}_t \mathbf{x}^{(i)}, & \boldsymbol{\beta}_t &\leftarrow \boldsymbol{\beta}_t^{(i)}. \\ \mathbf{G}_t^{(i)} &\leftarrow (1 - \gamma)\mathbf{G}_{t-1}^{(i)} + \gamma\mathbf{D}_{t-1}^\top \mathbf{M}_t \mathbf{D}_{t-1}, & \mathbf{G}_t &\leftarrow \bar{\mathbf{G}}_t^{(i)}. \end{aligned}$$

 Compute the approximate code for \mathbf{x}_t :

$$\boldsymbol{\alpha}_t \leftarrow \underset{\boldsymbol{\alpha} \in \mathbb{R}^k}{\operatorname{argmin}} \frac{1}{2} \boldsymbol{\alpha}^\top \mathbf{G}_t \boldsymbol{\alpha} - \boldsymbol{\alpha}^\top \boldsymbol{\beta}_t + \lambda \Omega(\boldsymbol{\alpha}). \quad (18)$$

 Update the parameters of the aggregated surrogate \bar{g}_t :

$$\begin{aligned} \bar{\mathbf{C}}_t &\leftarrow (1 - w_t)\bar{\mathbf{C}}_t + w_t \boldsymbol{\alpha}_t \boldsymbol{\alpha}_t^\top, \\ \mathbf{P}_t \bar{\mathbf{B}}_t &\leftarrow (1 - w_t)\mathbf{P}_t \bar{\mathbf{B}}_t + w_t \mathbf{P}_t \mathbf{x}_t \boldsymbol{\alpha}_t^\top. \end{aligned} \quad (19)$$

 Compute simultaneously (using Alg. 4 for 1st line):

$$\begin{aligned} \mathbf{P}_t \mathbf{D}_t &\leftarrow \underset{\mathbf{D}^r \in \mathcal{C}^r}{\operatorname{argmin}} \frac{1}{2} \operatorname{Tr}(\mathbf{D}^{r\top} \mathbf{D}^r \bar{\mathbf{C}}_t) - \operatorname{Tr}(\mathbf{D}^{r\top} \mathbf{P}_t \bar{\mathbf{B}}_t). \\ \mathbf{P}_t^\perp \bar{\mathbf{B}}_t &\leftarrow (1 - w_t)\mathbf{P}_t^\perp \bar{\mathbf{B}}_{t-1} + w_t \mathbf{P}_t^\perp \mathbf{x}_t \boldsymbol{\alpha}_t^\top. \end{aligned} \quad (20)$$

Output: Final iterate \mathbf{D}_T .

$\mathbf{P}_t^\perp \mathbf{Y}$ are the submatrices of $\mathbf{Y} \in \mathbb{R}^{p \times y}$ with rows respectively selected and not selected by \mathbf{M}_t . In algorithms, $\mathbf{P}_t \mathbf{Y} \leftarrow \mathbf{Z} \in \mathbb{R}^{q \times n}$ assigns the rows of \mathbf{Z} to the rows of \mathbf{Y} selected by \mathbf{P}_t , by an abuse of notation.

In brief, subsampled online matrix factorization, defined in Alg. 3, follows the outer loop of online matrix factorization, with the following major modifications at iteration t :

- it uses $\mathbf{M}_t \mathbf{x}_t$ and low-size statistics instead of \mathbf{x}_t to estimate the code $\boldsymbol{\alpha}_t$ and the surrogate g_t ,
- it updates a subset of the dictionary $\mathbf{P}_t \mathbf{D}_{t-1}$ to reduce the surrogate value $\bar{g}_t(\mathbf{D})$. Relevant parameters of \bar{g}_t are computed using $\mathbf{P}_t \mathbf{x}_t$ and $\boldsymbol{\alpha}_t$ only.

We now present SOMF in details. For comparison purpose, we write all variables that would be computed following the OMF rules at iteration t with a $*$ superscript. For simplicity, in Alg. 3 and in the following paragraphs, we assume that we use one sample per iteration —in practice, we use mini-batches of size η . The next derivations are transposable when a batch I_t is drawn at iteration t instead of a single sample i_t .

3.1.2 Code computation

In the OMF algorithm presented in Section 2, $\boldsymbol{\alpha}_t^*$ is obtained solving (6), which amounts to set

$$\boldsymbol{\alpha}_t^* = \underset{\boldsymbol{\alpha}}{\operatorname{argmin}} \frac{1}{2} \boldsymbol{\alpha}^\top \mathbf{G}_t^* \boldsymbol{\alpha} - \boldsymbol{\alpha}^\top \boldsymbol{\beta}_t^* + \lambda \Omega(\boldsymbol{\alpha}), \quad (21)$$

where $\mathbf{G}_t^* = \mathbf{D}_{t-1}^\top \mathbf{D}_{t-1}$ and $\boldsymbol{\beta}_t^* = \mathbf{D}_{t-1}^\top \mathbf{x}_t$. For large p , the computation of \mathbf{G}_t^* and $\boldsymbol{\beta}_t^*$ dominates the complexity of the regression step, which depends almost linearly on p . To reduce this complexity, we

use *estimators* for \mathbf{G}_t^* and β_t^* , computed at a cost proportional to the reduced dimension q . We propose three kind of estimators with different properties.

The most simple q -costly *unbiased* estimation of \mathbf{G}_t^* and β_t^* is obtained by subsampling matrix products with \mathbf{M}_t :

$$\begin{aligned}\mathbf{G}_t &= \mathbf{D}_{t-1}^\top \mathbf{M}_t \mathbf{D}_{t-1} \\ \beta_t &= \mathbf{D}_{t-1}^\top \mathbf{M}_t \mathbf{x}_t.\end{aligned}\tag{a}$$

This is the strategy proposed in [26]. We use \mathbf{G}_t and β_t in (18) (Alg. 3), which amounts to minimize the *masked loss*

$$\min_{\alpha \in \mathbb{R}^k} \frac{1}{2} \|\mathbf{M}_t(\mathbf{x}^t - \mathbf{D}_{t-1}^\top \alpha)\|_2^2 + \lambda \Omega(\alpha).\tag{22}$$

\mathbf{G}_t and β_t are computed in a number of operations proportional to q , which brings a speed-up factor of almost r in the code computation for large p . On large data, using estimators (a) instead of exact \mathbf{G}_t^* and β_t^* proves very efficient during the first epochs (cycles over the columns).¹ However, due to the masking, \mathbf{G}_t and β_t are not *consistent* estimators: they do not converge to \mathbf{G}_t^* and β_t^* for large t , which breaks theoretical guarantees on the algorithm output. Empirical results in Section 5.5 show that the sequence of iterates approaches a critical point of the risk (3), but may then drift from it.

At iteration t , the sample \mathbf{x}_t is drawn from a finite set of samples $\{\mathbf{x}^{(i)}\}_i$. This allows to average estimators over previously seen samples and address the non-consistency issue of (a). Namely, we keep in memory $2n$ estimators, written $(\mathbf{G}_t^{(i)}, \beta_t^{(i)})_{1 \leq i \leq n}$. We observe the sample $i = i_t$ at iteration t and use it to update the i -th estimators $\bar{\mathbf{G}}_t^{(i)}, \bar{\beta}_t^{(i)}$ following

$$\begin{aligned}\beta_t^{(i)} &= (1 - \gamma) \mathbf{G}_{t-1}^{(i)} + \gamma \mathbf{D}_{t-1}^\top \mathbf{M}_t \mathbf{x}^{(i)} \\ \mathbf{G}_t^{(i)} &= (1 - \gamma) \mathbf{G}_{t-1}^{(i)} + \gamma \mathbf{D}_{t-1}^\top \mathbf{M}_t \mathbf{D}_t^{(i)},\end{aligned}\tag{23}$$

where γ is a weight factor determined by the number of time sample i has been previously observed at time t . Precisely, given $(\gamma_c)_c$ a decreasing sequence of weights,

$$\gamma = \gamma_{c_t^{(i)}} \quad \text{where} \quad c_t^{(i)} = \left| \left\{ s \leq t, \mathbf{x}_s = \mathbf{x}^{(i)} \right\} \right|.$$

All others estimators $\{\mathbf{G}_t^{(j)}, \beta_t^{(j)}\}_{j \neq i}$ are left unchanged from iteration $t - 1$. The set $\{\mathbf{G}_t^{(i)}, \beta_t^{(i)}\}_{1 \leq i \leq n}$ is used to define the *averaged* estimators

$$\begin{aligned}\mathbf{G}_t &\triangleq \mathbf{G}_t^{(i)} = \sum_{s \leq t, \mathbf{x}_s = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \mathbf{D}_{s-1}^\top \mathbf{M}_s \mathbf{D}_{s-1} \\ \beta_t &\triangleq \beta_t^{(i)} = \sum_{s \leq t, \mathbf{x}_s = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \mathbf{D}_{s-1}^\top \mathbf{M}_s \mathbf{x}^{(i)},\end{aligned}\tag{b}$$

where $\gamma_{s,t}^{(i)} = \gamma_{c_t^{(i)}} \prod_{s < t, \mathbf{x}_s = \mathbf{x}^{(i)}} (1 - \gamma_{c_s^{(i)}})$. Using β_t and \mathbf{G}_t in (18), α_t minimizes the masked loss averaged over the previous iterations where sample i appeared:

$$\min_{\alpha \in \mathbb{R}^k} \sum_{\substack{s \leq t \\ \mathbf{x}_s = \mathbf{x}^{(i)}}} \frac{\gamma_{s,t}^{(i)}}{2} \|\mathbf{M}_s(\mathbf{x}^{(i)} - \mathbf{D}_{s-1}^\top \alpha)\|_2^2 + \lambda \Omega(\alpha).\tag{24}$$

The sequences $(\mathbf{G}_t)_t$ and $(\beta_t)_t$ are *consistent* estimations of $(\mathbf{G}_t^*)_t$ and $(\beta_t^*)_t$ — consistency arises from the fact that a single sample $\mathbf{x}^{(i)}$ is observed with different masks along iterations. Solving (24) is made closer and closer to solving (21), to ensure the correctness of the algorithm (see Section 4). Yet, computing the estimators (b) is no more costly than computing (a) and still permits to speed up a single iteration close to r times. In the mini-batch setting, for every $i \in I_t$, we use the estimators $\mathbf{G}_t^{(i)}$ and $\beta_t^{(i)}$ to compute α_i ,

¹Estimators (a) are also available in the infinite sample setting, when minimizing expected risk (5) from a i.i.d sample stream $(\mathbf{x}_t)_t$.

Table 1: Comparison of estimators used for code computation

Est.	β_t	\mathbf{G}_t	Convergence	Extra mem. cost	1 st epoch perform.
(a)	Masked	Masked			✓
(b)	Averaged	Averaged	✓	nk^2	✓
(c)	Averaged	Exact	✓	nk	

This method has a memory cost of $\mathcal{O}(nk^2)$. This is reasonable² compared to the dataset size if $p \gg k^2$. To reduce the memory usage, another strategy is to use the *true* Gram matrix \mathbf{G}_t and the estimator β_t from (b):

$$\begin{aligned} \mathbf{G}_t &\triangleq \mathbf{G}_t^* = \mathbf{D}_{t-1}^\top \mathbf{D}_{t-1} \\ \beta_t &\triangleq \sum_{s \leq t, \mathbf{x}_s = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \mathbf{D}_{s-1}^\top \mathbf{M}_s \mathbf{x}^{(i)} \end{aligned} \quad (\text{c})$$

As previously, the consistency of $(\beta_t)_t$ ensures that (5) is correctly solved despite the approximation in $(\alpha_t)_t$ computation. With the partial dictionary update step we propose, it is possible to maintain \mathbf{G}_t at a cost proportional to q . The time complexity of the coding step is thus similarly reduced when replacing (b) or (c) estimators in (21), but the latter option has a memory usage in $\mathcal{O}(nk)$. Although estimators (c) are slightly less performant in the first epochs, they are a good compromise between resource usage and convergence. We summarize the characteristics of the three estimators (a)–(c) in Table 1, anticipating their empirical comparison in Section 5.

Surrogate computation The computation of α_t using one of the estimators above defines a surrogate $g_t(\mathbf{D}) \triangleq \frac{1}{2} \|\mathbf{x}_t - \mathbf{D}\alpha_t\|_2^2 + \lambda\Omega(\alpha)$, which we use to update the aggregated surrogate $\bar{g}_t \triangleq (1-w_t)\bar{g}_{t-1} + w_t g_t$, as in online matrix factorization. We follow (8) (with weights $(w_t)_t$) to update the matrices \mathbf{B}_t and \mathbf{C}_t , which define \bar{g}_t up to constant factors. The update of \mathbf{B}_t requires a number of operations proportional to p . Fortunately, we will see in the next paragraph that it is possible to update $\mathbf{P}_t^\perp \mathbf{B}_t$ in the main thread with a number of operation proportional to q and to complete the update of $\mathbf{P}_t^\perp \mathbf{B}_t$ in parallel with the dictionary update step.

Weight sequences Specific $(w_t)_t$ and $(\gamma_c)_c$ in Alg. 3 are required. We provide then in Assumption **(I)** of the analysis: $w_t = \frac{1}{t^u}$ and $\gamma_c = \frac{1}{c^v}$, where $u \in (\frac{11}{12}, 1)$ and $v \in (\frac{3}{4}, 3u - 2)$ to ensure convergence. Weights have little impact on convergence speed in practice.

3.1.3 Dictionary update

In the original online algorithm, the whole dictionary \mathbf{D}_{t-1} is updated at iteration t . To reduce the time complexity of this step, we add a “freezing” constraint to the minimization (7) of \bar{g}_t . Every row r of \mathbf{D} that corresponds to an *unseen* row r at iteration r (such that $\mathbf{M}_t[r, r] = 0$) remains unchanged. This casts the problem (7) into a lower dimensional space. Formally, the freezing operation comes out as a additional constraint in (7):

$$\mathbf{D}_t = \underset{\substack{\mathbf{D} \in \mathcal{C} \\ \mathbf{P}_t^\perp \mathbf{D} = \mathbf{P}_t^\perp \mathbf{D}_{t-1}}}{\operatorname{argmin}} \frac{1}{2} \operatorname{Tr}(\mathbf{D}^\top \mathbf{D} \bar{\mathbf{C}}_t) - \operatorname{Tr}(\mathbf{D}^\top \bar{\mathbf{B}}_t). \quad (25)$$

The constraints are separable into two blocks of rows. For each atom $\mathbf{d}^{(j)}$, the rules $\|\mathbf{d}^{(j)}\| \leq 1$ and $\mathbf{P}_t^\perp \mathbf{d}^{(j)} = \mathbf{P}_t^\perp \mathbf{d}_{t-1}^{(j)}$ can indeed be rewritten

$$\begin{cases} \|\mathbf{P}_t \mathbf{d}^{(j)}\| &\leq 1 - \|\mathbf{d}_{t-1}^{(j)}\| + \|\mathbf{P}_t \mathbf{d}_{t-1}^{(j)}\| \triangleq r_t^{(j)} \\ \mathbf{P}_t^\perp \mathbf{d}^{(j)} &= \mathbf{P}_t^\perp \mathbf{d}_{t-1}^{(j)}. \end{cases} \quad (26)$$

²It is also possible to efficiently swap the estimators $(\mathbf{G}_t^{(i)})_i$ on disk, as they are only accessed for $i = i_t$ at iteration t .

Algorithm 4 Partial dictionary update

Input: Dictionary \mathbf{D}_{t-1} , projector \mathbf{P}_t , statistics $\bar{\mathbf{C}}_t, \bar{\mathbf{B}}_t$, norms $(n_{t-1}^{(j)})_{0 \leq j < k}$, Gram matrix \mathbf{G}_t (optional).

$\mathbf{D}_t \leftarrow \mathbf{D}_{t-1}, \quad \mathbf{G}_t \leftarrow \mathbf{G}_t - \mathbf{D}_{t-1}^\top \mathbf{P}_t \mathbf{D}_{t-1}.$

for $j \in \text{permutation}([1, k])$ **do**

$r_t^{(j)} \leftarrow n_{t-1}^{(j)} + \|\mathbf{P}_t \mathbf{d}_{t-1}^{(j)}\|.$

$\mathbf{u} \leftarrow \mathbf{P}_t \mathbf{d}_{t-1}^{(j)} + \frac{1}{\bar{\mathbf{C}}_t[j, j]} (\mathbf{P}_t \bar{\mathbf{b}}_t^{(j)} - \mathbf{P}_t \mathbf{D}_t \bar{\mathbf{c}}_t^{(j)}). \quad \triangleright \text{in } \mathbb{R}^q$

$\mathbf{P}_t \mathbf{d}_t^{(j)} \leftarrow \text{enet_projection}(\mathbf{u}, r_t^{(j)}). \quad \triangleright \text{in } \mathbb{R}^q$

$n_t^{(j)} \leftarrow r_t^{(j)} - \|\mathbf{P}_t \mathbf{d}_t^{(j)}\|.$

$\mathbf{G}_{t+1} \leftarrow \mathbf{G}_t + \mathbf{D}_t^\top \mathbf{P}_t \mathbf{D}_t.$

Output: Dictionary \mathbf{D}_t , norms $(n_t^{(j)})_j$, Gram matrix $\mathbf{G}_{t+1}.$

Solving (25) is therefore equivalent to solving the following problem in $\mathbb{R}^{q \times k}$, with $\mathbf{B}_t^r \triangleq \mathbf{P}_t \mathbf{B}_t$,

$$\mathbf{D}^r \in \underset{\mathbf{D}^r \in \mathcal{C}^r}{\text{argmin}} \frac{1}{2} \text{Tr}(\mathbf{D}^{r \top} \mathbf{D}^r \bar{\mathbf{C}}_t) - \text{Tr}(\mathbf{D}^{r \top} \bar{\mathbf{B}}_t^r) \quad (27)$$

$$\text{where } \mathcal{C}^r = \{\mathbf{D}^r \in \mathbb{R}^{q \times k} / \forall j \in [0, k-1], \|\mathbf{d}^{r(j)}\| \leq r_t^{(j)}\}.$$

The rows of \mathbf{D}_t selected by \mathbf{P}_t are then replaced with \mathbf{D}^r , while the other rows of \mathbf{D}_t are unchanged from iteration $t-1$. Formally, $\mathbf{P}_t \mathbf{D}_t = \mathbf{D}^r$ and $\mathbf{P}_t^\perp \mathbf{D}_t = \mathbf{P}_t^\perp \mathbf{D}_{t-1}$. We solve (27) by a projected block coordinate descent (BCD) similar to the one used in the original algorithm, but performed in a subspace of size q . We compute each column j of the gradient that we use in the block coordinate descent loop with $q \times k$ operations, as it writes $\mathbf{D}^r \bar{\mathbf{c}}_t^{(j)} - \bar{\mathbf{b}}_t^{r(j)} \in \mathbb{R}^q$, where $\bar{\mathbf{c}}_t^{(j)}$ and $\bar{\mathbf{b}}_t^{r(j)}$ are the j -th columns of $\bar{\mathbf{C}}_t$ and $\bar{\mathbf{B}}_t^r$. Each reduced atom $\mathbf{d}^{r(j)}$ is projected onto the elastic-net ball of radius $r_t^{(j)}$, at an average cost in $\mathcal{O}(q)$ following [15]. This makes the complexity of a single-column update proportional to q . Performing the projection requires to keep in memory the values $\{n_t^{(j)} \triangleq 1 - \|\mathbf{d}_t^{(j)}\|\}_j$, which can be updated online at a negligible cost.

We provide the reduced dictionary update step in Alg. 4, where we use the function `enet_projection`(\mathbf{u}, r) that performs the orthogonal projection of $\mathbf{u} \in \mathbb{R}^q$ onto the elastic-net ball of radius r . As in the original algorithm, we perform a single pass over columns to solve (27). Dictionary update is now performed with a number of operations proportional to q , instead of p in the original algorithm. Thanks to the random nature of $(\mathbf{M}_t)_t$, updating \mathbf{D}_{t-1} into \mathbf{D}_t reduces \bar{g}_t enough to ensure convergence.

Gram matrix computation Performing partial updates of \mathbf{D}_t makes it possible to maintain the full Gram matrix $\mathbf{G}_t = \mathbf{G}_t^*$ with a cost in $\mathcal{O}(qk^2)$ per iteration. \mathbf{G}_t can indeed be updated by computing the reduced Gram matrix $\mathbf{D}^\top \mathbf{P}_t \mathbf{D}$ before and after the dictionary update:

$$\mathbf{G}_{t+1} = \mathbf{D}_t^\top \mathbf{D}_t = \mathbf{G}_t - \mathbf{D}_{t-1}^\top \mathbf{P}_t \mathbf{D}_{t-1} + \mathbf{D}_t^\top \mathbf{P}_t \mathbf{D}_t. \quad (28)$$

As mentioned previously, this update rule allows to use the true Gram matrix \mathbf{G}_t in (21) with a q -dependent complexity.

Parallel surrogate computation Performing block coordinate descent on \bar{g}_t^r requires to access $\bar{\mathbf{B}}_t^r = \mathbf{P}_t \bar{\mathbf{B}}_t$ only. Assuming we may use more than two threads, this allows to parallelize the dictionary update step with the update of $\mathbf{P}_t^\perp \bar{\mathbf{B}}_t$. In the main thread, we compute $\mathbf{P}_t \bar{\mathbf{B}}_t$ following

$$\mathbf{P}_t \bar{\mathbf{B}}_t \leftarrow (1 - w_t) \bar{\mathbf{P}}_t \bar{\mathbf{B}}_{t-1} + w_t \mathbf{P}_t \mathbf{x}_t \boldsymbol{\alpha}_t^\top. \quad (19 - \text{Alg. 3})$$

which has a cost proportional to q . Then, we update in parallel the dictionary and the rows of $\bar{\mathbf{B}}_t$ that are not selected by \mathbf{M}_t :

$$\mathbf{P}_t^\perp \bar{\mathbf{B}}_t \leftarrow (1 - w_t) \mathbf{P}_t^\perp \bar{\mathbf{B}}_{t-1} + w_t \mathbf{P}_t^\perp \mathbf{x}_t \boldsymbol{\alpha}_t^\top. \quad (20 - \text{Alg. 3})$$

This update requires $k(p - q)\eta$ operations (one matrix-matrix product) for a mini-batch of size η . In contrast, with appropriate implementation, the dictionary update step requires $4kq^2$ to $6kq^2$ operations, among which $2kq^2$ come from slower matrix-vector products. Assuming $k \sim \eta$, updating $\bar{\mathbf{B}}_t$ is faster than updating the dictionary up to $r \sim 10$, and performing (20) on a second thread is seamless in term of wall-clock time. More threads may be used for larger reduction or batch size.

3.1.4 Subsampling and time complexity

Subsampling may be used in only *some* of the steps of Alg. 3, with the other steps following Alg. 1. Whether to use subsampling or not in each step depends on the trade-off between the computational speed-up it brings and the approximations it makes. It is useful to understand how complexity of OMF evolves with p . We write s the average number of non-zero coefficients in $(\boldsymbol{\alpha}_t)_t$ ($s = k$ when $\Omega = \|\cdot\|_2^2$). OMF complexity has three terms:

- (i) $\mathcal{O}(pk^2)$: computation of the Gram matrix \mathbf{G}_t , update of the dictionary \mathbf{D}_t with block coordinate descent,
- (ii) $\mathcal{O}(pk\eta)$: computation of $\boldsymbol{\beta}_t = \mathbf{D}_{t-1}^\top \mathbf{x}_t$ and of $\bar{\mathbf{B}}_t$ using $\mathbf{x}_t \boldsymbol{\alpha}_t^\top$,
- (iii) $\mathcal{O}(ks^2\eta)$: computation of $\boldsymbol{\alpha}_t$ using \mathbf{G}_t and $\boldsymbol{\beta}_t$, using matrix inversion or elastic-net regression.

Using subsampling turns p into $q = \frac{p}{r}$ in the expressions above. It improves single iteration time when the cost of regression $\mathcal{O}(ks^2\eta)$ is dominated by another term. This happens whenever $\frac{p}{r} > s^2$, where r is the reduction factor used in the algorithm. Subsampling can bring performance improvement up to $r \sim \frac{p}{s^2}$. It can be introduced in either computations from (i) or (ii), or both. When using small batch size, i.e., when $\eta < k$, computations from (i) dominates complexity, and subsampling should be first introduced in dictionary update (i), and for code computation (ii) beyond a certain reduction ratio. On the other hand, with large batch size $\eta > k$, subsampling should be first introduced in code computation, then in the dictionary update step. The reasoning above ignore potentially large constants. The best trade-offs in using subsampling must be empirically determined, which we do in Section 5.

3.2 Stochastic approximate majorization-minimization

The SOMF algorithm can be understood within the stochastic majorization-minimization framework. The modifications that we propose are indeed perturbations to the first and third steps of the SMM presented in Algorithm 2:

- The code is computed approximately: the surrogate is only an *approximate* majorizing surrogate of f_t near \mathbf{D}_{t-1} .
- The surrogate objective is only *reduced* and not minimized, due to the added constraint and the fact that we perform only one pass of block coordinate descent.

We propose a new *stochastic approximate majorization-minimization* (SAMM) framework handling these perturbations:

- A majorization step (12 – Alg. 2), computes an *approximate surrogate* of f_t near θ_{t-1} : $g_t \approx g_t^*$, where g_t is a true upper-bounding surrogate of \bar{f}_t .
- A minimization step (13 – Alg. 2), finds θ_t by reducing *enough* the objective \bar{g}_t : $\theta_t \approx \theta_t^* \triangleq \operatorname{argmin}_{\theta \in \Theta} \bar{g}_t(\theta)$, which implies $\bar{g}_t(\theta_t) \gtrsim \bar{g}_t(\theta_t^*)$.

The SAMM framework is general, in the sense that approximations are not specified. The next section provides a theoretical analysis of the approximation of SAMM and establishes how SOMF is an instance of SAMM. It concludes by establishing Proposition 3, which provides convergence guarantees for SOMF, under the same assumptions made for OMF in [15].

4 Convergence analysis

We prove the correctness of the SMM and SOMF algorithms under reasonable assumptions. First, we recall the theoretical guarantees of the stochastic majorization-minimization algorithm [10]. Then, we show how to make it robust to various approximations, leading to the SMM algorithm. Finally, we show how SOMF is a special case of SMM.

4.1 Basic assumptions and results on SMM convergence

We recall the main results on stochastic majorization-minimization algorithms, established in [10], under assumptions that we slightly tighten for our purpose. In our setting, we consider the empirical risk minimization problem

$$\min_{\theta \in \Theta} \left(\bar{f}(\theta) \triangleq \frac{1}{n} \sum_{i=1}^n f(\theta, \mathbf{x}^{(i)}) \right), \quad (29)$$

where $f : \mathbb{R}^K \times \mathcal{X} \rightarrow \mathbb{R}$ is a loss function and

(A) $\Theta \subset \mathbb{R}^K$ and the support \mathcal{X} of the data are compact.

This is a special case of (5) where the samples $(\mathbf{x}_t)_t$ are drawn uniformly from the set $\{\mathbf{x}^{(i)}\}_i$. The loss functions $f_t \triangleq f(\cdot, \mathbf{x}_t)$ defined on \mathbb{R}^K can be non-convex. We instead assume that they meet reasonable regularity conditions:

(B) $(f_t)_t$ is uniformly R -Lipschitz continuous on \mathbb{R}^K and uniformly bounded on Θ .

(C) The directional derivatives [34] $\nabla f_t(\theta, \theta' - \theta)$ and $\nabla \bar{f}(\theta, \theta' - \theta)$ exist for all θ and θ' in \mathbb{R}^K .

Assumption **(C)** allows to characterize the *stationary points* of problem (29), namely $\theta \in \Theta$ such that $\nabla \bar{f}(\theta, \theta' - \theta) \geq 0$ for all $\theta' \in \Theta$ — intuitively a point is stationary when there is no direction in which the objective can be improved.

Let us now recall the definition of first-order surrogate functions used in the SMM algorithm. $(g_t)_t$ are selected in the set $\mathcal{S}_{\rho,L}(f_t, \theta_{t-1})$, hereby introduced.

Definition 1 (First-order surrogate function). Given a function $f : \mathbb{R}^K \rightarrow \mathbb{R}$, $\theta \in \Theta$ and $\rho, L > 0$, we define $\mathcal{S}_{\rho,L}(f, \theta)$ as the set of functions $g : \mathbb{R}^K \rightarrow \mathbb{R}$ such that

- g is majorizing f on Θ and g is ρ -strongly convex,
- g and f are tight at θ — i.e., $g(\theta) = f(\theta)$, $g - f$ is differentiable, $\nabla(g - f)$ is L -Lipschitz, $\nabla(g - f)(\theta) = 0$.

We stipulate that \bar{g}_t should be *parametrized* and thus representable in memory. The following assumption is met in OMF, as \bar{g}_t is parametrized by the matrices $\bar{\mathbf{C}}_t$ and $\bar{\mathbf{B}}_t$.

(D) Parametrized surrogates. The surrogates $(\bar{g}_t)_t$ are parametrized by vectors in a compact set $\mathcal{K} \subset \mathbb{R}^P$. Namely, for all $t > 0$, there exists $\kappa_t \in \mathcal{K}$ such that \bar{g}_t is unequivocally defined as $g_t \triangleq \bar{g}_{\kappa_t}$.

Finally, we ensure that the weights $(w_t)_t$ used in Alg. 2 decrease at a certain rate.

(E) There exists $u \in (\frac{3}{4}, 1)$ such that $w_t = t^{-u}$.

When $(\theta_t)_t$ is the sequence yielded by Alg. 2, the following result (Proposition 3.4 in [10]) establishes the convergence of $(\bar{f}(\theta_t))_t$ and states that θ_t is asymptotically a stationary point of the finite sum problem (29), as a special case of the *expected* risk minimization problem (5).

Proposition 1 (Convergence of SMM, from [10]). Under assumptions **(A)** – **(E)**, $(\bar{f}(\theta_t))_{t \geq 1}$ converges with probability one. Every limit point θ_∞ of $(\theta_t)_t$ is a stationary point of the risk \bar{f} defined in (29). That is,

$$\forall \theta \in \Theta, \quad \nabla \bar{f}(\theta_\infty, \theta - \theta_\infty) \geq 0. \quad (30)$$

The correctness of online matrix factorization can be deduced from this proposition.

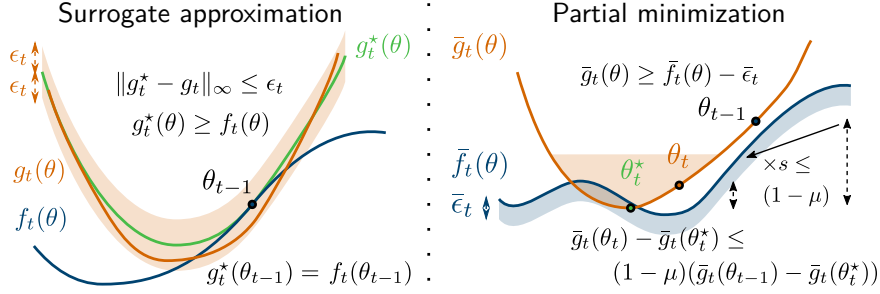


Figure 2: Both steps of SAMM make well-behaved approximations. The operations that are performed in exact SMM are in green and superscripted by $*$, while the actual computed values are in orange. Light bands recall the bounds on approximations assumed in **(F)** and **(G)**.

4.2 Convergence of SAMM

We now introduce assumptions on the approximations made in SAMM, before extending the result of Proposition 1. We make hypotheses on both the surrogate computation (*majorization*) step and the iterate update (*minimization*) step.

4.2.1 Approximate surrogate computation

The SMM algorithm selects a surrogate for f_t at point θ_{t-1} within the set $\mathcal{S}_{\rho,L}(f_t, \theta_{t-1})$. Surrogates within this set are *tight* at θ_{t-1} and greater than f_t everywhere. In SAMM, we allow the use of surrogates that are only *approximately majorizing* f_t and *approximately tight* at θ_{t-1} . This is indeed what SOMF does when using estimators in the code computation step.

For that purpose, we introduce the set $\mathcal{T}_{\rho,L}(f, \theta, \epsilon)$, that contains all functions ϵ -close of a surrogate in $\mathcal{S}_{\rho,L}(f, \theta)$ for the ℓ_∞ -norm:

Definition 2 (Approximate first-order surrogate function). Given a function $f : \mathbb{R}^K \rightarrow \mathbb{R}$ and $\theta \in \Theta$, $\mathcal{T}_{\rho,L}(f, \theta, \epsilon)$ is the set of ρ -strongly convex functions $g : \mathbb{R}^K \rightarrow \mathbb{R}$ such that

- g is ϵ -majorizing f on Θ : $\forall \kappa \in \Theta, g(\kappa) - f(\kappa) \geq -\epsilon$,
- g and f are ϵ -tight at θ — i.e., $g(\theta) - f(\theta) \leq \epsilon$, $g - f$ is differentiable, $\nabla(g - f)$ is L -lipschitz.

We assume that SAMM selects an approximative surrogate in $\mathcal{T}_{\rho,L}(f_t, \theta_{t-1}, \epsilon_t)$ at each iteration, where $(\epsilon_t)_t$ is a deterministic or random sequence that vanishes at a sufficient rate.

(F) For all $t > 0$, $g_t \in \mathcal{T}_{\rho,L}(f_t, \theta_{t-1}, \epsilon_t)$. There exists a constant $\gamma > 0$ such that $\mathbb{E}[\epsilon_t] \in \mathcal{O}(t^{2(u-1)-\gamma})$ and $\epsilon_t \rightarrow_\infty 0$ almost surely.

4.2.2 Approximate surrogate minimization

We do not require θ_t to be the minimizer of \bar{g}_t any longer, but assume that the surrogate objective function \bar{g}_t decreases “fast enough”. Namely, θ_t obtained from partial minimization should be closer to a minimizer of \bar{g}_t than θ_{t-1} . Precisely, we assume, writing \mathcal{F}_{t-1} the filtration of all known variables up to the end of iteration $t - 1$,

(G) For all $t > 0$, $\bar{g}_t(\theta_t) < \bar{g}_t(\theta_{t-1})$. Moreover, there exists $\mu > 0$ such that, for all $t > 0$

$$\begin{aligned} \mathbb{E}[\bar{g}_t(\theta_t) - \bar{g}_t(\theta_t^*) | \mathcal{F}_{t-1}] &\leq (1 - \mu)(\bar{g}_t(\theta_{t-1}) - \bar{g}_t(\theta_t^*)) \\ \text{where } \theta_t^* &= \underset{\theta \in \Theta}{\operatorname{argmin}} \bar{g}_t(\theta), \end{aligned} \quad (31)$$

Assumption **(G)** is met by choosing an appropriate method for the inner \bar{g}_t minimization step — a large variety of gradient-descent algorithms indeed have convergence rates of the form (31).

The principles of SAMM are illustrated in Figure 2, which provides a geometric interpretation of the approximations introduced in **(F)** and **(G)**. With both assumptions at hand, we can state the theoretical guarantees of SAMM.

4.2.3 Convergence result

The following proposition guarantees that the stationary point condition of Proposition 1 holds for the SAMM algorithm, despite the use of approximate surrogates and approximate minimization.

Proposition 2 (Convergence of SAMM). Under assumptions **(A)** – **(G)**, $(\bar{f}(\theta_t))_{t \geq 1}$, the conclusion of Proposition 1 holds for SAMM.

The proof of this proposition relies on a central lemma, that shows that the iterate sequence $(\theta_t)_t$ has bounded asymptotic variations. This lemma also theoretically justifies to perform of a single pass of block coordinate descent in OMF [15].

Lemma 3 (Estimate stability under SAMM approximations). Consider any sequence of iterates $(\theta_t)_t$ provided by SAMM. Assume there exists $\epsilon > 0$ such that $g_t \in \mathcal{T}_{L,\rho}(f_t, \theta_{t-1}, \epsilon)$ for all $t \geq 1$. Assume **(A)** – **(E)** and **(G)**. Then the sequence $(\|\theta_t - \theta_{t-1}\|_2)_t$ is bounded, and θ_t is asymptotically an exact minimizer. Namely, almost surely,

$$\|\theta_t - \theta_{t-1}\|_2 \in \mathcal{O}(w_t), \quad \|\bar{g}_t(\theta_t) - \bar{g}_t(\theta_t^*)\|_2 \in \mathcal{O}(w_t^2), \quad (32)$$

where we recall $\theta_t^* \triangleq \operatorname{argmin}_{\theta \in \Theta} \bar{g}_t(\theta)$.

Assumption **(F)** is essential to bound the errors introduced by the sequence $(\epsilon_t)_t$ in the proof of Proposition 2, while **(G)** is the key element of Lemma 3. Both proofs follow — algebraic derivations are detailed in Appendix A.

Proof of Lemma 3. We first show the result when the decrease rate (31) holds deterministically. We focus on bounding four terms: $A_t \triangleq \|\theta_t - \theta_{t-1}\|_2$, $B_t \triangleq \|\theta_t - \theta_t^*\|_2$, $C_t \triangleq \|\theta_t^* - \theta_{t-1}^*\|_2$, $D_t \triangleq \bar{g}_t(\theta_t) - \bar{g}_t(\theta_t^*)$. Using the assumptions, the convexity of the norms and the strong-convexity of \bar{g}_t ,

$$\begin{aligned} A_t^2 &\leq \frac{6}{\rho}(D_t + D_{t-1}) + \mathcal{O}(w_t^2) \\ \tilde{D}_t &\leq (1 - \mu)\tilde{D}_{t-1} \frac{w_{t-1}^2}{w_t^2} + o(\tilde{D}_t + \tilde{D}_{t-1}), \end{aligned} \quad (33)$$

where $\tilde{D}_t \triangleq \frac{D_t}{w_t^2}$ and $\frac{o(\tilde{D}_t + \tilde{D}_{t-1})}{\tilde{D}_t + \tilde{D}_{t-1}} \rightarrow 0$ when $\tilde{D}_t \rightarrow \infty$. From basic considerations on sequences, this is enough to show that \tilde{D}_t is bounded and the result follows. In the general case **(G)** where the decrease rate only holds in expectation, we show that (33) is met replacing \tilde{D}_t by $\mathbb{E}[\tilde{D}_t]$, and therefore $\mathbb{E}[D_t] = \mathcal{O}(w_t^2)$. Moreover, we now have $\mathbb{E}[\tilde{D}_t | \mathcal{F}_{t-1}] \leq (1 - \mu)\tilde{D}_{t-1} + o(\tilde{D}_t, \tilde{D}_{t-1})$, using previous notations. *Doob's forward convergence theorem* [35] ensures that \tilde{D}_t is bounded almost surely and the lemma follows. \square

Proof of Proposition 2. We adapt the proof of Proposition 3.3 from [10] (reproduced as Proposition 1 in our work). On one hand, relaxing tightness and majorizing hypotheses introduces some extra terms in the derivations. Assumption **(F)** allows to control these extra terms. On the other hand, Lemma 3 allows to use the inequality $\|\theta_t - \theta_{t-1}\|_2 = \mathcal{O}(w_t)$, without assuming that $(\theta_t)_t$ are minimizers of $(\bar{g}_t)_t$. This inequality ensures that \bar{g}_t and \bar{f}_t are asymptotically tight at θ_{t-1} , which permit to conclude.

In the first part of the proof, we control the positive expected variation of $(\bar{g}_t(\theta_t))_t$ to show that it is a converging quasi-martingale [see 36]. By construction of \bar{g}_t and properties of the surrogates $(g_t)_t$,

$$\begin{aligned} \bar{g}_t(\theta_t) - \bar{g}_{t-1}(\theta_{t-1}) &\leq w_t(g_t(\theta_{t-1}) - \bar{g}_{t-1}(\theta_{t-1})) \\ &\leq w_t(f_t(\theta_{t-1}) - \bar{f}_{t-1}(\theta_{t-1})) + w_t(\bar{\epsilon}_{t-1} + \epsilon_t), \end{aligned} \quad (34)$$

where $\bar{\epsilon}_t \triangleq (1 - w_t)\bar{\epsilon}_{t-1} + w_t\epsilon_t$ and $\bar{f}_t = (1 - w_t)\bar{f}_{t-1} + w_t f_t$, introduced in (14). The inequality (34) involves two terms: the first one, controlled by empirical process theory, is the same as in [10]. The second one arises from the ϵ_t -approximation in the construction of g_t . Assumption **(F)** provides tight-enough bounds to show that the expected total positive variation of $(\bar{g}_t(\theta_t))_t$ is bounded. $(\bar{g}_t(\theta_t))_t$ is therefore a quasi-martingale as in [10], and as such converges almost surely towards an integrable random variable g^* .

Lemma 3, which is available as assumption **(G)** is met, ensures that $\|\theta_t - \theta_{t-1}\|_2 \in \mathcal{O}(w_t)$. Summing over iterations t in (34) and taking the expectation of the sum, we obtain $\sum w_t(f_t(\theta_{t-1}) - \bar{g}_t(\theta_{t-1})) < \infty$.

We may therefore use the numerical Lemma A.6 from [10] and obtain $\bar{f}_t(\theta_t) - \bar{g}_t(\theta_t) \rightarrow 0$. Lemma B.7 from [10], that we do not need to adapt, ensures that $\bar{f}_t \rightarrow \bar{f}$ uniformly. Combined, both convergence results show that $\bar{f}(\theta_t) \rightarrow g^*$ almost surely.

The smoothness of $\bar{h}_t \triangleq \bar{g}_t - \bar{f}_t$ combined with $|\bar{h}_t(\theta_t)| < \bar{\epsilon}_t$ shows $\|\nabla \bar{h}_t(\theta_t)\|_2 \rightarrow 0$. As a consequence of Lemma 3, $\|\theta_t - \theta_t^*\|_2 \rightarrow 0$. We easily show that $(\bar{h}_t)_t$ is uniformly Lipschitz, which yields $\|\nabla \bar{h}_t(\theta_t^*)\|_2 \rightarrow 0$. We assume without loss of generality that θ_t converges toward one of its accumulation points θ_∞ , which implies $\theta_t^* \rightarrow \theta_\infty$. Writing \bar{g}_∞ the limit of $(\bar{g}_t)_t$, we immediately obtain $\nabla \bar{g}_\infty(\theta_\infty, \theta - \theta_\infty) \geq 0$ for all $\theta \in \Theta$, as θ_t^* are minimizers of $(\bar{g}_t)_t$. We then perform a first-order development of \bar{h}_t near θ_t^* instead of θ_t in [10], which shows $\nabla \bar{h}_\infty(\theta_\infty, \theta - \theta_\infty) = 0$, where \bar{h}_∞ is the uniform limit of $(\bar{h}_t)_t$. This is enough to conclude. \square

It now remains to establish that SOMF is indeed an instance of stochastic approximate majorization-minimization.

4.3 SOMF is an instance of SAMM

We introduce the same hypothesis on $\bar{\mathbf{C}}_t$ as in [15] and a similar one on $\mathbf{D}_t^\top \mathbf{D}_t$ — they ensure strong convexity of the surrogate and boundedness of $(\boldsymbol{\alpha}_t)_t$ but do not cause any loss of generality as they are met in practice after a few iterations, and can be guaranteed by adding small ℓ_2 regularizations to f .

(H) There exists ρ such that for all $t > 0$, $\bar{\mathbf{C}}_t, \mathbf{D}_t^\top \mathbf{D}_t \succ \rho \mathbf{I}$.

We further assume that the weights $(w_t)_t$ and $(\gamma_c)_c$ decay at specific rates, which will prove important to meet **(F)**. The proofs can be adapted for more complex weights.

(I) There exists $u \in (\frac{11}{12}, 1)$ and $v \in (\frac{3}{4}, 3u - 2)$ such that, for all $t > 0$, $c > 0$, $w_t = t^{-u}$, $\gamma_c \triangleq c^{-v}$.

We can now state our main practical convergence result, that applies on any sequence $(\mathbf{D}_t)_t$ produced by SOMF, using estimators (b) or (c). \bar{f} is the empirical risk defined in (3).

Proposition 3 (SOMF convergence). Under assumptions **(H)** and **(I)**, $\bar{f}(\mathbf{D}_t)$ converges with probability one and every limit point \mathbf{D}_∞ of $(\mathbf{D}_t)_t$ is a stationary point of \bar{f} : for all $\mathbf{D} \in \mathcal{C}$

$$\nabla \bar{f}(\mathbf{D}_\infty, \mathbf{D} - \mathbf{D}_\infty) \geq 0 \quad (35)$$

Proof of Proposition 3. From assumption **(H)** and the definition of f , $(f_t)_t$ verify **(A)**–**(D)** while **(I)** implies **(E)**. We separately show that the majorization step of SOMF meets assumptions **(F)**, and that its minimization step meets **(G)**.

4.3.1 The update of \mathbf{D}_t verifies **(G)**

Immediately, the function value decrease deterministically at each minimization step: $\bar{g}_t(\mathbf{D}_t) \leq \bar{g}_t(\mathbf{D}_{t-1})$. Moreover, with non-zero probability d , $\mathbf{M}_t = \mathbf{I}_p$ and Alg. 4 performs a single pass of simple block coordinate descent on \bar{g}_t . In this case, as \bar{g}_t is strongly convex from **(H)**, [37, 38] ensures that the suboptimality decreases at least of factor $1 - \mu$, where μ is a constant independent of t . This is enough to show that $\mathbb{E}[\bar{g}_t(\mathbf{D}_t) - \bar{g}_t(\mathbf{D}_t^*) | \mathcal{F}_{t-1}] \leq (1 - d\mu)(\bar{g}_t(\mathbf{D}_{t-1}) - \bar{g}_t(\mathbf{D}_t^*))$.

4.3.2 The approximate surrogates $(g_t)_t$ verify **(F)**

We write $g_t^* \in \mathcal{S}_{\rho,L}(f_t, \mathbf{D}_{t-1})$ the surrogate of OMF at iteration t , that uses the exact loadings $\boldsymbol{\alpha}_t^*$ in (15). g_t is the surrogate of SOMF, that uses $\boldsymbol{\alpha}_t$ defined in (b) or (c). Defining $\epsilon_t \triangleq \sup_{\mathbf{D} \in \mathcal{C}} |g_t^*(\mathbf{D}) - g_t(\mathbf{D})|$, we have $g_t \in \mathcal{T}_{\rho,L}(f_{t-1}, \theta_{t-1}, \epsilon_t)$. We show that $(\epsilon_t)_t$ meets the conditions of **(F)**. To do so, we first use Lemma B.6 from [39] to show that

$$\mathbf{G}, \boldsymbol{\beta} \rightarrow \operatorname{argmin}_{\boldsymbol{\alpha} \in \mathbb{R}^k} \frac{1}{2} \boldsymbol{\alpha}^\top \mathbf{G} \boldsymbol{\alpha} - \boldsymbol{\alpha}^\top \boldsymbol{\beta} + \lambda \Omega(\boldsymbol{\alpha}) \quad (36)$$

is Lipschitz continuous. Moreover, as $\mathbf{D} \in \mathcal{C}$, the functions $\boldsymbol{\alpha} \rightarrow \frac{1}{2} \boldsymbol{\alpha}^\top \mathbf{D}^\top \mathbf{D} \boldsymbol{\alpha} - \boldsymbol{\alpha}^\top \mathbf{D}^\top \mathbf{x}_t + \lambda \Omega(\boldsymbol{\alpha})$ are Lipschitz continuous, with a Lipschitz constant independent of \mathbf{D} . Composing both functions, we have for all $\mathbf{D} \in \mathcal{C}$,

$$g_t^*(\mathbf{D}) - g_t(\mathbf{D}) \leq C(\|\mathbf{G}_t^* - \mathbf{G}_t\|_F + \|\boldsymbol{\beta}_t^* - \boldsymbol{\beta}_t\|_F), \quad (37)$$

where C is a constant. As a first consequence, $\|g_t - g_t^*\|_\infty$ is bounded, hence Lemma 3 holds and the estimates verify $\mathbf{D}_t - \mathbf{D}_{t-1} = \mathcal{O}(w_t)$. Establishing that **(F)** is met reduces to show that both $\|\mathbf{G}_t^* - \mathbf{G}_t\|_F$ and $\|\beta_t^* - \beta_t\|_F$ are properly bounded. Using the triangular inequality in (b),

$$\begin{aligned} \|\beta_t - \beta_t^*\|_F &\leq B \sum_{s \leq t, \mathbf{x}_s = \mathbf{x}_t = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \|\mathbf{D}_{s-1} - \mathbf{D}_{t-1}\|_F \\ &\quad + C \left\| \left(\sum_{s \leq t, \mathbf{x}_s = \mathbf{x}_t = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \mathbf{M}_s \right) - \mathbf{I} \right\|_F, \end{aligned} \quad (38)$$

where an identical result holds for $\|\mathbf{G}_t - \mathbf{G}_t^*\|_F$. The left term accounts for the aggregation of past information in building the estimates β_t . The right term is the norm of a vector whose coefficients are the running average of i.i.d zero-mean Bernoulli variables. We define $\gamma \triangleq \frac{1}{2} \min(v - \frac{3}{4}, (3u-2) - v)$, which is positive from assumption **(I)**. We show that the expectation of each terms of (38) is in $\mathcal{O}(t^{2(u-1)-\gamma})$, and that each term converges to 0 almost surely. For the left term, this is a consequence of Lemma B.7 from [10], applied to constant value functions. Bounding the right term simply uses $\mathbf{D}_t - \mathbf{D}_{t-1} = \mathcal{O}(w_t)$ with rearrangement in the sum — see Appendix A for details. The proposition follows. \square

5 Experiments

The SOMF algorithm is designed for datasets with large number of samples n and large dimensionality p . Indeed, as detailed in Section 3.1, subsampling removes the computational bottlenecks that arise from high dimensionality. Proposition 3 establishes that the subsampling used in SOMF is safe: that the final result is indeed a stationary point of the empirical risk defined in (3). In this section, we provide a strong empirical validation of the efficiency of subsampling.

We tackle two different problems, in functional Magnetic Resonance Imaging (fMRI) and hyperspectral imaging. Both involve the factorization of very large matrices \mathbf{X} with sparse factors. As the data we consider is huge, subsampling reduces the time of a single iteration by a factor close to $\frac{p}{q}$. Yet it is also much redundant: SOMF makes little approximations and accessing only a fraction of the features per iteration should not hinder much the refinement of the dictionary. Hence high speed-ups are expected — and indeed obtained. All experiments can be reproduced using open-source code.

5.1 Problems and datasets

5.1.1 Functional MRI

Matrix factorization has long been used on functional Magnetic Resonance Imaging [18]. Data are temporal series of 3D images of brain activity and are decomposed into spatial modes capturing regions that activate synchronously. They form a matrix \mathbf{X} where columns are the 3D images, and rows corresponds to voxels. Interesting dictionaries for neuroimaging capture spatially-localized components, with a few brain regions. This can be obtained by enforcing sparsity on the dictionary: we use a ℓ_2 penalty and the elastic-net constraint. SOMF streams subsampled 3D brain records to learn the sparse dictionary \mathbf{D} . Data can be huge: we use the whole HCP dataset [40], with $n = 2.4 \cdot 10^6$ (2000 records, 1 200 time points) and $p = 2 \cdot 10^5$, totaling 2 TB of dense data. For comparison, we also use a smaller public dataset (ADHD200 [41]) with 40 records, $n = 7000$ samples and $p = 6 \cdot 10^4$ voxels. Historically, brain decomposition have been obtained by minimizing the classical dictionary learning objective on *transposed* data [42]: the code \mathbf{A} holds sparse spatial maps and voxel time-series are streamed. This is not a natural streaming order for fMRI data as \mathbf{X} is stored columnwise on disk, which makes the sparse dictionary formulation more appealing. Importantly, we seek a *low-rank* factorization, to keep the decomposition interpretable — $k \sim 100 \ll p$.

5.1.2 Hyperspectral imaging

Hyperspectral cameras acquire images with many channels that correspond to different spectral bands. They are used heavily in remote sensing (satellite imaging), and material study (microscopic imaging). They yield digital images with around 1 million pixels, each associated with hundreds of spectral channels. Sparse matrix factorization has been widely used on these data for image classification [43, 44] and

Table 2: Summary of experimental settings

Field	Functional MRI		Hyperspectral imaging
Dataset	ADHD	HCP	Patches from AVIRIS
Factors	D sparse, A dense		D dense, A sparse
# samples n	$7 \cdot 10^3$	$2 \cdot 10^6$	$2 \cdot 10^6$
# features p	$6 \cdot 10^4$	$2 \cdot 10^5$	$6 \cdot 10^4$
X size	2 GB	2 TB	103 GB
Use case ex.	Extracting predictive feature		Recognition / denoising

denoising [45, 46]. All methods rely on the extraction of full-band patches representing a local image neighborhood with all channels included. These patches are very high dimensional, due to the number of spectral bands. From one image of the AVIRIS project [47], we extract $n = 2 \cdot 10^6$ patches of size 16×16 with 224 channels, hence $p = 6 \cdot 10^4$. A dense dictionary is learned from these patches. It should allow a sparse representation of samples: we either use the classical dictionary learning setting (ℓ_1 /elastic-net penalty), or further add positive constraints to the dictionary and codes: both methods may be used and deserved to be benchmarked. We seek a dictionary of reasonable size: we use $k \sim 256 \ll p$.

5.2 Experimental design

To validate the introduction of subsampling and the usefulness of SOMF, we perform two major experiments.

- We measure the performance of SOMF when increasing the reduction factor, and show benefits of stochastic dimension reduction on all datasets.
- We assess the importance of subsampling in each of the steps of SOMF. We compare the different approaches proposed for code computation.

Validation

We compute the objective function (3) over a test set to rule out any overfitting effect — a dictionary should be a good representation of unseen samples. This criterion is always plotted against wall-clock time, as we are interested in the performance of SOMF for practitioners.




Figure 3: Subsampling provides significant speed-ups on all fMRI and hyperspectral datasets. A reduction factor of 12 is a good overall choice. With larger data, larger reduction factors can be used for better performance — convergence is reached 13× faster than state-of-the-art methods on the 2TB HCP dataset.

Tools

To perform a valid benchmark, we implement OMF and SOMF using Cython [48]. We use coordinate descent [49] to solve Lasso problems with optional positivity constraints. Code computation is parallelized to handle minibatches. Experiments use *scikit-learn* [50] for numerics, and *nilearn* [51] for handling fMRI




Figure 4: Given a 3 minute time budget, the atoms learned by SOMF are more focal and less noisy than those learned by OMF. They are closer to the dictionary of first line, for which convergence has been reached.

data. We have released the code in an open-source Python package³. Experiments were run on 3 cores of an Intel Xeon 2.6GHz, in which case computing $\mathbf{P}_t^\perp \bar{\mathbf{B}}_t$ is faster than updating $\mathbf{P}_t \mathbf{D}_t$.

Parameter setting

Setting the number of components k and the amount of regularization λ is a hard problem in the absence of ground truth. Those are typically set by cross-validation when matrix factorization is part of a supervised pipeline. For fMRI, we set $k = 70$ to obtain interpretable networks, and set λ so that the decomposition approximately covers the whole brain (i.e., every map is $\frac{k}{70}$ sparse). For hyperspectral images, we set $k = 256$ and select λ to obtain a dictionary on which codes are around 3% sparse. We cycle randomly through the data (fMRI records, image patches) until convergence, using minibatches of size $\eta = 200$ for HCP and AVIRIS, and $\eta = 50$ for ADHD (small number of samples). Hyperspectral patches are normalized in the dictionary learning setting, but not in the non-negative setting — the classical pre-conditioning for each case. We use $u = 0.917$ and $v = 0.751$ for weight sequences.

5.3 Reduction brings speed-up at all data scales

We benchmark SOMF for various reduction factors against the original online matrix factorization algorithm OMF [15], on the three presented datasets. We stream data in the same order for all reduction factors. Using variant (c) (true Gram matrix, averaged β_t) performs slightly better on fMRI datasets, whereas (b) (averaged Gram matrix and β_t) is slightly faster for hyperspectral decomposition. For comparison purpose, we display results using estimators (b) only.

Figure 3 plots the test objective against CPU time. First, we observe that all algorithms find dictionaries with very close objective function values for all reduction factors, on each dataset. This is not a trivial observation as the matrix factorization problem (3) is not convex and different runs of OMF and SOMF may converge towards minima with different values. Second, and most importantly, SOMF provides significant improvements in convergence speed for three different sizes of data and three different factorization settings. Both observations confirm the relevance of the subsampling approach.

Quantitatively, we summarize the speed-ups obtained in Table 3. On fMRI data, on both large and medium datasets, SOMF provides more than an order of magnitude speed-up. Practitioners working on datasets akin to HCP can decompose their data in 20 minutes instead of 4 h previously, while working on a single machine. We obtain the highest speed-ups for the largest dataset — accounting for the extra redundancy that usually appears when dataset size increase. Up to $r \sim 8$, speed-up is of the order of r — subsampling induces little noise in the iterate sequence, compared to OMF. Hyperspectral decomposition is performed near $7\times$ faster than with OMF in the classical dictionary learning setting, and $3\times$ in the non-negative setting, which further demonstrates the versatility of SOMF. Qualitatively, given a certain time budget, Figure 4 compares the results of OMF and the results of SOMF with a subsampling ratio $r = 24$, in the non-negative setting. Our algorithm yields a valid smooth bank of filters much faster. The same comparison has been made for fMRI in [26].

³<https://github.com/arthurmensch/modl>

Table 3: Time to reach convergence ($< 1\%$ test objective)

Dataset	ADHD		AVIRIS (NMF)		AVIRIS (DL)		HCP	
	OMF	SOMF	OMF	SOMF	OMF	SOMF	OMF	SOMF
Conv. time	6 min	28 s	2 h 30	43 min	1 h 16	11 min	3 h 50	17 min
Speed-up		11.8		3.36		6.80		13.31

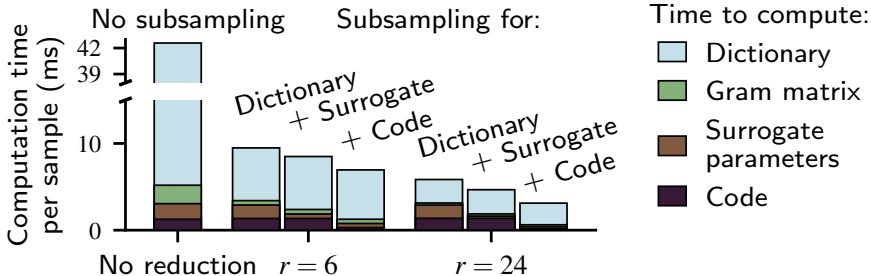


Figure 5: Profiling OMF and SOMF on HCP. Partial dictionary update removes the major bottleneck of online matrix factorization for small reductions. For higher reduction, parameter update and code computation must be subsampled to further reduce the iteration time.

Limitations

Table 3 reports convergence time within 1%, which is enough for application in practice. SOMF is less beneficial when setting very high precision: for convergence within 0.01%, speed-up for HCP is 3.4. This is expected as SOMF trades speed for approximation. For high precision convergence, the reduction ratio can be reduced after a few epochs. As expected, there exists an optimal reduction ratio, depending on the problem and precision, beyond which performance reduces: $r = 12$ yields better results than $r = 24$ when performing dictionary learning on AVIRIS for 1% precision.

Our first experiment establishes the power of stochastic subsampling as a whole. In the following two experiments, we refine our analysis to show that subsampling is indeed useful in the three steps of online matrix factorization.

5.4 For each step of SOMF, subsampling removes a bottleneck

In Section 3, we have provided theoretical guidelines on when to introduce subsampling in each of the three steps of an iteration of SOMF. This analysis predicts that, for $\eta \sim k$, we should first use partial dictionary update, before using approximate code computation and asynchronous parameter aggregation. We verify this by measuring the time spent by SOMF on each of the updates for various reduction factors, on the HCP dataset. Results are presented in Figure 5. We observe that block coordinate descent is indeed the bottleneck in OMF. Introducing partial dictionary update removes this bottleneck, and as the reduction factor increases, code computation and surrogate aggregation becomes the major bottlenecks. Introducing subsampling as described in SOMF overcomes these bottlenecks, which rationalizes all steps of SOMF from a computational point of view.

5.5 Code subsampling becomes useful for high reduction

It remains to assess the performance of approximate code computation and averaging techniques used in SOMF. Indeed, subsampling for code computation introduces noise that may undermine the computational speed-up. To understand the impact of approximate code computation, we compare three strategies to compute $(\alpha_t)_t$ on the HCP dataset. First we follow (21) and use the full stream $(\mathbf{x}_t)_t$. Subsampling is thus disabled for code computation and used only in dictionary update. Second, we rely on masked, non-consistent estimators (a), as in [26] — this breaks convergence guarantees. Third, we use averaged estimators (β_t, \mathbf{G}_t) from (c) to reduce the variance of $(\alpha_t)_t$ computation.

Fig. 6 compares the three strategies for $r \in \{12, 24\}$. Partial minimization at each step is the




Figure 6: Approximating code computation with the proposed subsampling method further accelerates the convergence of SOMF. Refining code computation using past iterations (averaged estimates) performs better than simply performing a subsampled linear regression as in [26]

most important part to accelerate convergence: the SOMF algorithm subsampling only in the dictionary update step already largely outperforms OMF. This is expected, as dictionary update constitutes the main bottleneck of OMF in large-scale settings. Yet, for large reduction factors, using subsampling in code computation is important to further accelerate convergence. This clearly appears when comparing the plain and dashed black curves. Using past estimates to better approximate $(\alpha_t)_t$ yields faster convergence than the non-converging, masked loss strategy (a) proposed in [26]. This latter strategy remains a good option, as it empirically converges and is simpler to implement.

6 Conclusion

In this paper, we introduce SOMF, a matrix-factorization algorithm that can handle input data with very large number of rows and columns. It leverages subsampling within the inner loop of a streaming algorithm to make iterations faster and accelerate convergence. We show that SOMF provides a stationary point of the non-convex matrix factorization problem. To prove this result, we extend the stochastic majorization-minimization framework to two major approximations. We assess the performance of SOMF on real-world large-scale problems, with different sparsity/positivity requirements on learned factors. In particular, on fMRI and hyperspectral data decomposition, we show that the use of subsampling can speed-up decomposition up to 13 times. The larger the dataset, the more SOMF outperforms state-of-the-art techniques, which is very promising for future applications. This calls for adaptation of our approach to learn more complex models.

References

- [1] Julien Mairal, Francis Bach, and Julien Ponce. Sparse Modeling for Image and Vision Processing. *Foundations and Trends in Computer Graphics and Vision*, 8(2-3):85–283, 2014.
- [2] Nathan Srebro, Jason Rennie, and Tommi S. Jaakkola. Maximum-margin matrix factorization. In *Advances in Neural Information Processing Systems*, pages 1329–1336, 2004.
- [3] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.
- [4] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global Vectors for Word Representation. In *EMNLP*, volume 14, pages 1532–43, 2014.
- [5] Omer Levy and Yoav Goldberg. Neural word embedding as implicit matrix factorization. In *Advances in Neural Information Processing Systems*, pages 2177–2185, 2014.
- [6] Yin Zhang, Matthew Roughan, Walter Willinger, and Lili Qiu. Spatio-temporal compressive sensing and internet traffic matrices. 39(4):267–278, 2009.
- [7] Hyunsoo Kim and Haesun Park. Sparse non-negative matrix factorizations via alternating non-negativity-constrained least squares for microarray data analysis. *Bioinformatics*, 23(12):1495–1502, 2007.
- [8] Gaël Varoquaux, Alexandre Gramfort, Fabian Pedregosa, Vincent Michel, and Bertrand Thirion. Multi-subject dictionary learning to segment an atlas of brain spontaneous activity. In *Information Processing in Medical Imaging Conference*, 2011.

- [9] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT*, pages 177–186. Springer, 2010.
- [10] Julien Mairal. Stochastic majorization-minimization algorithms for large-scale optimization. In *Advances in Neural Information Processing Systems*, 2013.
- [11] Meisam Razaviyayn, Mingyi Hong, and Zhi-Quan Luo. A unified convergence analysis of block successive minimization methods for nonsmooth optimization. *SIAM Journal on Optimization*, 23(2):1126–1153, 2013.
- [12] Samuel Burer and Renato D. C. Monteiro. Local Minima and Convergence in Low-Rank Semidefinite Programming. *Mathematical Programming*, 103:427, 2004.
- [13] Benjamin Recht and Christopher Ré. Parallel stochastic gradient algorithms for large-scale matrix completion. *Mathematical Programming Computation*, 5:201–226, 2013.
- [14] Robert M. Bell and Yehuda Koren. Lessons from the Netflix prize challenge. *ACM SIGKDD Explorations Newsletter*, 9(2):75–79, 2007.
- [15] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*, 11:19–, 2010.
- [16] William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- [17] Ella Bingham and Heikki Mannila. Random projection in dimensionality reduction: Applications to image and text data. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2001.
- [18] M. J. McKeown, S. Makeig, G. G. Brown, T. P. Jung, S. S. Kindermann, A. J. Bell, and T. J. Sejnowski. Analysis of fMRI Data by Blind Separation into Independent Spatial Components. *Human Brain Mapping*, 6(3):160–188, 1998.
- [19] Emmanuel J. Candès and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Inf. Theory*, 52:5406, 2006.
- [20] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [21] Vladimir Rokhlin, Arthur Szlam, and Mark Tygert. A randomized algorithm for principal component analysis. *SIAM J. Matrix Anal. A.*, 31:1100–1124, 2009.
- [22] Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In *IEEE Symposium on Foundations of Computer Science*, 2006.
- [23] Yichao Lu, Paramveer Dhillon, Dean P. Foster, and Lyle Ungar. Faster ridge regression via the subsampled randomized hadamard transform. In *NIPS*, 2013.
- [24] Mert Pilanci and Martin J Wainwright. Iterative hessian sketch: Fast and accurate solution approximation for constrained least-squares. *JMLR*, pages 1–33, 2015.
- [25] Garvesh Raskutti and Michael Mahoney. Statistical and algorithmic perspectives on randomized sketching for ordinary least-squares. In *ICML*, 2015.
- [26] Arthur Mensch, Julien Mairal, Bertrand Thirion, and Gaël Varoquaux. Dictionary learning for massive matrix factorization. In *ICML*, 2016.
- [27] Bruno A. Olshausen and David J. Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, 37:3311–3325, 1997.
- [28] Robert Tibshirani. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Series B Stat. Methodol.*, pages 267–288, 1996.
- [29] Hui Zou, Trevor Hastie, and Robert Tibshirani. Sparse principal component analysis. *Journal of computational and graphical statistics*, 15:265–286, 2006.
- [30] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Series B Stat. Methodol.*, 67(2):301–320, 2005.

- [31] Patrik O Hoyer. Non-negative matrix factorization with sparseness constraints. *Journal of Machine Learning Research*, 5(Nov):1457–1469, 2004.
- [32] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization Methods for Large-Scale Machine Learning. *arXiv preprint arXiv:1606.04838*, 2016.
- [33] Aad W. Van der Vaart. *Asymptotic Statistics*, volume 3. CUP, 2000.
- [34] Jonathan M. Borwein and Adrian S. Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Springer Science & Business Media, 2010.
- [35] Joseph Leo Doob. *Stochastic processes*. John Wiley & Sons, 1990.
- [36] Michel Métivier. *Semimartingales: A Course on Stochastic Processes*, volume 2. Walter de Gruyter, 1982.
- [37] Amir Beck and Luba Tetruashvili. On the convergence of block coordinate descent type methods. *SIAM Journal on Optimization*, 23(4):2037–2060, 2013.
- [38] Stephen J. Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, 2015.
- [39] Julien Mairal. Optimization with first-order surrogate functions. *ICML*, 2013.
- [40] David C. Van Essen, Stephen M. Smith, Deanna M. Barch, Timothy E. J. Behrens, Essa Yacoub, and Kamil Ugurbil. The WU-Minn Human Connectome Project: An overview. *NeuroImage*, 80:62–79, 2013.
- [41] Michael P. Ph D. Milham, Damien PA-C. Fair, Maarten Ph D. Mennes, and Stewart H. M. D. Mostofsky. The adhd-200 consortium: A model to advance the translational potential of neuroimaging in clinical neuroscience. *Front. Syst. Neurosci.*, 6, 2012.
- [42] Gaël Varoquaux, Yannick Schwartz, Philippe Pinel, and Bertrand Thirion. Cohort-level brain mapping: Learning cognitive atoms to single out specialized regions. In *IPMI*, 2013.
- [43] Yi Chen, Nasser M. Nasrabadi, and Trac D. Tran. Hyperspectral image classification using dictionary-based sparse representation. *IEEE GRS*, 49:3973, 2011.
- [44] A. Soltani-Farani, H. R. Rabiee, and S. A. Hosseini. Spatial-Aware Dictionary Learning for Hyperspectral Image Classification. *IEEE GRS*, 53:527, 2015.
- [45] Matteo Maggioni, Vladimir Katkovnik, Karen Egiazarian, and Alessandro Foi. Nonlocal transform-domain filter for volumetric data denoising and reconstruction. *IEEE Transactions on Image Processing*, 22(1):119–133, 2013.
- [46] Yi Peng, Deyu Meng, Zongben Xu, Chenqiang Gao, Yi Yang, and Biao Zhang. Decomposable nonlocal tensor dictionary learning for multispectral image denoising. In *IEEE CVPR*, 2014.
- [47] Gregg Vane. First results from the airborne visible/infrared imaging spectrometer (AVIRIS). In *SPIE Annual Technical Symposium*, 1987.
- [48] Stefan Behnel, Robert Bradshaw, Craig Citro, Lisandro Dalcin, Dag Sverre Seljebotn, and Kurt Smith. Cython: The best of both worlds. *Comput. Sci. Eng.*, 13:31, 2011.
- [49] Jerome Friedman, Trevor Hastie, Holger Höfling, Robert Tibshirani, and others. Pathwise coordinate optimization. *The Annals of Applied Statistics*, 1(2):302–332, 2007.
- [50] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825, 2011.
- [51] Alexandre Abraham, Fabian Pedregosa, Michael Eickenberg, Philippe Gervais, Andreas Mueller, Jean Kossaifi, Alexandre Gramfort, Bertrand Thirion, and Gael Varoquaux. Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics*, 8:14, 2014.

A Proofs of convergence — detailed derivations

This appendix contains the detailed algebraic derivations for the proofs of Lemma 3, Proposition 2 and Proposition 3. For readability, we reproduce the structure of the proofs presented in the main text.

A.1 Basic properties of the surrogates, estimate stability and optimality — Proof of Lemma 3

We derive an important result on the stability and optimality of the sequence $(\theta_t)_t$, formalized in Lemma 3 — introduced in the main text. Let us first introduce a numerical lemma on the boundedness of well-behaved deterministic and random sequence.

Lemma 1 (Boundedness of quasi-geometric decreasing functions). Consider $(x_t)_{t \geq 0} \in (\mathbb{R}^+)^{\mathbb{N}}$, $u : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, $t_0 \in \mathbb{N}$ and $\alpha \in [0, 1)$ such that

$$\forall t \geq t_0, x_t \leq \alpha x_{t-1} + u(x_t, x_{t-1}), \quad \text{where } u(x_t, x_{t-1}) = o(x_t + x_{t-1}) \quad \text{if } x_t + x_{t-1} \rightarrow \infty \quad (39)$$

Then $(x_t)_t$ is bounded. Assume now that $(X_t)_t$ is a random sequence in \mathbb{R}^+ , such that

$$\mathbb{E}[X_t] < \infty, \quad \text{and } \forall t > t_0, \mathbb{E}[X_t | \mathcal{F}_{t-1}] \leq \alpha X_{t-1} + u(X_t, X_{t-1}), \quad (40)$$

where \mathcal{F}_{t-1} is the filtration of all information up to time $t - 1$. Then $(X_t)_t$ is bounded almost surely.

Proof. We first focus on the deterministic case. Assume that $(x_t)_t$ is not bounded. Then there exists a subsequence of $(x_t)_t$ that diverges towards $+\infty$. We assume without loss of generality that $(x_t)_t \rightarrow \infty$. Then, $x_t + x_{t-1} \rightarrow \infty$ and for all $\epsilon > 0$, using the asymptotic bounds on u , there exists $t_1 \geq t_0$ such that

$$\forall t \geq t_1, x_t \leq \frac{\alpha + \epsilon}{1 - \epsilon} x_{t-1} \quad (41)$$

Setting ϵ small enough, we have that x_t is bounded by a geometrically decreasing sequence after t_1 , and converges to 0, which contradicts our hypothesis. This is enough to conclude.

In the random case, we consider a realisation of $(X_t)_t$ that is not bounded, and assumes without loss of generality that it diverges to $+\infty$. Following the reasoning above, there exists $\beta < 1$, $t_1 > 0$, such that $\mathbb{E}[X_t | \mathcal{F}_{t-1}] \leq \beta X_{t-1}$ for all $t > t_1$. Therefore X_t is a supermartingale beyond a certain time. As $\mathbb{E}[X_t] < \infty$, Doob's forward convergence lemma on discrete martingales [35] ensures that $(X_t)_t$ converges almost surely. Therefore the event $\{(X_t)_t \text{ is not bounded}\}$ cannot happen on a set with non-zero probability, less it would lead to a contradiction. The lemma follows. \square

We now derive some basic properties of the approximate surrogate functions used in SAMM. The proof is adapted from [10].

Lemma 2 (Basic properties of approximate surrogate functions). Consider any sequence of iterates $(\theta_t)_t$ and assume there exists $\epsilon > 0$ such that $g_t \in \mathcal{T}_{L, \rho}(f_t, \theta_{t-1}, \epsilon)$ for all $t \geq 1$. Define $h_t = g_t - f_t$ for all $t \geq 1$, $\bar{h}_0 \triangleq h_0$ and $\bar{h}_t \triangleq (1 - w_t)\bar{h}_{t-1} + w_t h_t$. Under assumptions **(B)** – **(E)**,

- (i) $(\nabla h_t(\theta_{t-1}))_{t > 0}$ is uniformly bounded and there exists R' such that $\{\nabla h_t\}_t$ is uniformly bounded by R' .
- (ii) $(h_t)_t$ and $(\bar{h}_t)_t$ are uniformly R' -Lipschitz, $(g_t)_t$ and $(\bar{g}_t)_t$ are uniformly $(R + R')$ -Lipschitz.

Proof. We first prove (i). We set $\alpha > 0$ and define $\theta' = \theta_t - \alpha \frac{\nabla h_t(\theta_t)}{\|\nabla h_t(\theta_t)\|_2}$. As h_t has a L -Lipschitz gradient on \mathbb{R}^K , we have

$$\begin{aligned} h_t(\theta') &\leq h_t(\theta_t) - \alpha \|\nabla h_t(\theta_t)\|_2 + \frac{L\alpha^2}{2} \\ \|\nabla h_t(\theta_t)\|_2 &\leq \frac{1}{\alpha} (h_t(\theta_t) - h_t(\theta')) + \frac{L\alpha^2}{2} \leq \frac{2}{\alpha} \epsilon + \frac{L\alpha^2}{2}, \end{aligned} \quad (42)$$

where we use $h_t(\theta_t) < \epsilon$ and $-h_t(\theta'_t) \leq \epsilon$ from the assumption $g_t \in \mathcal{T}_{L,\rho}(f_t, \theta_{t-1}, \epsilon)$. Moreover, by definition, ∇h_t exists and is L -lipschitz for all t . Therefore, $\forall t \geq 1$,

$$\|\nabla h_t(\theta_{t-1})\|_2 - L\|\theta_{t-1} - \theta\|_2 \leq \|\nabla h_t(\theta)\|_2 \leq \|\nabla h_t(\theta_{t-1})\|_2 + L\|\theta_{t-1} - \theta\|_2. \quad (43)$$

Since Θ is compact and $(\|\nabla h_t(\theta_{t-1})\|_2)_{t \geq 1}$ is bounded in (42), ∇h_t is bounded by R' independent of t . (ii) follows by basic considerations on Lipschitz functions. \square

Using previous lemmas, we prove a result on the stability of the estimates, that derives from combining the properties of $(g_t)_t$ and the geometric decrease assumption **(G)**.

Lemma 3 (Estimate stability under SAMM approximationsversion). In the same setting as Lemma 2, with the additional assumption **(G)** (expected linear decrease of \bar{g}_t suboptimality), the sequence $\|\theta_t - \theta_{t-1}\|_2$ converges to 0 as fast as $(w_t)_t$, and θ_t is asymptotically an exact minimizer. Namely, almost surely,

$$\|\theta_t - \theta_{t-1}\|_2 \in \mathcal{O}(w_t) \quad \text{and} \quad \bar{g}_t(\theta_t) - \bar{g}_t(\theta_t^*) \in \mathcal{O}(w_t^2). \quad (44)$$

Proof. We first establish the result when a deterministic version of **(G)** holds, as it makes derivations simpler to follow.

A.1.1 Determistic decrease rate

We temporarily make a stronger assumption than **(G)** and assume that decays are deterministic.

(G_{det}) For all $t > 0$, $\bar{g}_t(\theta_t) < \bar{g}_t(\theta_{t-1})$. Moreover, there exists $\mu > 0$ such that, for all $t > 0$

$$\begin{aligned} \bar{g}_t(\theta_t) - \bar{g}_t(\theta_t^*) &\leq (1 - \mu)(\bar{g}_t(\theta_{t-1}) - \bar{g}_t(\theta_t^*)) \\ \text{where } \theta_t^* &= \underset{\theta \in \Theta}{\operatorname{argmin}} \bar{g}_t(\theta), \end{aligned} \quad (45)$$

We introduce the following auxiliary positive values, that we will seek to bound in the proof:

$$A_t \triangleq \|\theta_t - \theta_{t-1}\|_2, \quad B_t \triangleq \|\theta_t - \theta_t^*\|_2, \quad C_t \triangleq \|\theta_t^* - \theta_{t-1}^*\|_2, \quad D_t \triangleq \bar{g}_t(\theta_t) - \bar{g}_t(\theta_t^*). \quad (46)$$

Our goal is to bound A_t . We first relate it to C_t and B_t using convexity of ℓ_2 norm:

$$A_t^2 \leq 3B_t^2 + 3B_{t-1}^2 + 3C_t^2. \quad (47)$$

As θ_t^* is the minimizer of \bar{g}_t , by strong convexity of $(\bar{g}_t)_t$,

$$\frac{\rho}{2} B_t^2 = \frac{\rho}{2} \|\theta_t - \theta_t^*\|_2^2 \leq D_t \quad (48)$$

$$\begin{aligned} \frac{\rho}{2} \|\theta_t^* - \theta_{t-1}^*\|_2^2 &\leq \bar{g}_t(\theta_{t-1}^*) - \bar{g}_t(\theta_t^*) \leq (1 - w_t)(\bar{g}_{t-1}(\theta_{t-1}^*) - \bar{g}_{t-1}(\theta_t^*)) + w_t(g_t(\theta_{t-1}^*) - g_t(\theta_t^*)) \\ &\leq w_t(R + R')\|\theta_t^* - \theta_{t-1}^*\|_2, \quad \text{and therefore } C_t \leq w_t \frac{2Q}{\rho}, \end{aligned} \quad (49)$$

where the second inequalities holds because θ_{t-1}^* is a minimizer of \bar{g}_t and $(g_t)_t$ are Q -Lipschitz, where $Q \triangleq R + R'$, using Lemma 2. Replacing (48) and (49) in (47) yields

$$A_t^2 \leq \frac{6}{\rho}(D_t + D_{t-1}) + \frac{12Q^2}{\rho} w_t^2, \quad (50)$$

and we are left to show that $D_t \in \mathcal{O}(w_t^2)$ to conclude. For this, we decompose the inequality from **(G_{det})**:

$$\begin{aligned} D_t &\leq (1 - \mu)(\bar{g}_t(\theta_{t-1}) - \bar{g}_t(\theta_t^*)) \\ &= (1 - \mu) \left(w_t(g_t(\theta_{t-1}) - g_t(\theta_t)) + w_t(g_t(\theta_t) - g_t(\theta_t^*)) \right) \\ &\quad + (1 - \mu) \left((1 - w_t)(\bar{g}_{t-1}(\theta_{t-1}) - \bar{g}_{t-1}(\theta_{t-1}^*)) + (1 - w_t)(\bar{g}_{t-1}(\theta_{t-1}^*) - \bar{g}_{t-1}(\theta_t^*)) \right) \\ &\leq (1 - \mu)(w_t Q(A_t + B_t) + D_{t-1}), \end{aligned} \quad (51)$$

where the second inequality holds for the same reasons as in (49). Injecting (48) and (50) in (51), we obtain

$$\begin{aligned} \tilde{D}_t &\leq (1 - \mu)\tilde{D}_{t-1}\frac{w_{t-1}^2}{w_t^2} + u(\tilde{D}_t, \tilde{D}_{t-1}), \quad \text{where} \\ u(\tilde{D}_t, \tilde{D}_{t-1}) &\triangleq (1 - \mu)\tilde{Q}\left(\sqrt{3(\tilde{D}_t + \tilde{D}_{t-1}\frac{w_{t-1}^2}{w_t^2})} + \tilde{Q} + \sqrt{\tilde{D}_t}\right), \end{aligned} \quad (52)$$

defining $\tilde{D}_t \triangleq \frac{D_t}{w_t^2}$ and $\tilde{Q} \triangleq \sqrt{\frac{2}{\rho}}Q$. Since $\frac{w_{t-1}}{w_t} \rightarrow 1$ and the perturbation term $u(\tilde{D}_t, \tilde{D}_{t-1}) = o(\tilde{D}_t + \tilde{D}_{t-1})$ for $\tilde{D}_t \rightarrow \infty$, the deterministic result of Lemma 1 holds. This ensures that \tilde{D}_t is bounded, which combined with (48) allows to conclude.

A.1.2 Stochastic decrease rates

In the general case **(G)**, the inequalities (50), (48), (49) holds, and (51) is replaced by

$$\mathbb{E}[\tilde{D}_t | \mathcal{F}_{t-1}] \leq (1 - \mu)\tilde{D}_{t-1}\frac{w_{t-1}^2}{w_t^2} + u(\tilde{D}_t, \tilde{D}_{t-1}), \quad (53)$$

where \mathcal{F}_{t-1} is the filtration on the algorithm past until $t - 1$. Taking the expectation of this inequality and using Jensen inequality, we show that (51) holds when replacing \tilde{D}_t by $\mathbb{E}[\tilde{D}_t]$. This shows that $\mathbb{E}[D_t] = \mathcal{O}(w_t^2)$. We may therefore use the stochastic section of Lemma 1, and the result follows. \square

A.2 Convergence of SAMM — Proof of Proposition 2

We now proceed to prove the Proposition 2, that extends the stochastic majorization-minimization framework to allow approximations in both majorization and minimizations steps.

Proof of Proposition 2. We adapt the proof of Proposition 3.3 from [10] (reproduced as Proposition 1 in our work). Relaxing tightness and majorizing hypotheses introduces some extra error terms in the derivations. Assumption **(F)** allows to control these extra terms without breaking convergence. The stability lemma is crucial there.

A.2.1 Almost sure convergence of $(\bar{g}_t(\theta_t))$

We control the positive expected variation of $(g_t(\theta_t))_t$ to show that it is a converging quasi-martingale. By construction of \bar{g}_t and properties of the surrogates $(g_t)_t$:

$$\begin{aligned} \bar{g}_t(\theta_t) - \bar{g}_{t-1}(\theta_{t-1}) &= (\bar{g}_t(\theta_t) - \bar{g}_t(\theta_{t-1})) + w_t(g_t(\theta_{t-1}) - \bar{g}_{t-1}(\theta_{t-1})) \\ &\leq w_t(g_t(\theta_{t-1}) - \bar{g}_{t-1}(\theta_{t-1})) \\ &\leq w_t(g_t(\theta_{t-1}) - f_t(\theta_{t-1})) + w_t(f_t(\theta_{t-1}) - \bar{f}_{t-1}(\theta_{t-1})) \\ &\quad + w_t(\bar{f}_{t-1}(\theta_{t-1}) - \bar{g}_{t-1}(\theta_{t-1})) \\ &\leq w_t(f_t(\theta_{t-1}) - \bar{f}_{t-1}(\theta_{t-1})) + w_t(\bar{\epsilon}_{t-1} + \epsilon_t), \end{aligned} \quad (54)$$

where the average error sequence $(\bar{\epsilon}_t)_t$ is defined recursively: $\bar{\epsilon}_0 \triangleq \epsilon_0$ and $\bar{\epsilon}_t \triangleq (1 - w_t)\bar{\epsilon}_{t-1} + w_t\epsilon_t$. The first inequality uses $\bar{g}_t(\theta_t) \leq \bar{g}_t(\theta_{t-1})$. To obtain the forth inequality we observe $g_t(\theta_{t-1}) - f_t(\theta_{t-1}) < \epsilon_t$ by definition of ϵ_t and $\bar{f}_{t-1}(\theta_{t-1}) - \bar{g}_{t-1}(\theta_{t-1}) \leq \bar{\epsilon}_t$, which can easily be shown by induction on t . It follows, taking the conditional expectation with respect to \mathcal{F}_{t-1} the past information before time t :

$$\mathbb{E}[\bar{g}_t(\theta_t) - \bar{g}_{t-1}(\theta_{t-1}) | \mathcal{F}_{t-1}] \leq w_t \sup_{\theta \in \Theta} |f(\theta) - \bar{f}_{t-1}(\theta)| + w_t(\bar{\epsilon}_{t-1} + \mathbb{E}[\epsilon_t | \mathcal{F}_{t-1}]), \quad (55)$$

Here, we make use of the fact that ϵ_{t-1} is deterministic with respect to \mathcal{F}_{t-1} . To ensure convergence, we must bound both terms in (55): the first term is the same as in the original proof with exact surrogate,

while the second is the perturbative term introduced by the approximation sequence $(\epsilon_t)_t$. We use Lemma B.7 from [10], issued from the theory of empirical processes:

$$\mathbb{E}[\sup_{\theta \in \Theta} |f(\theta) - \bar{f}_{t-1}(\theta)|] = \mathcal{O}(w_t t^{1/2}), \quad \text{and thus} \quad \sum_{t=1}^{\infty} w_t \mathbb{E}[\sup_{\theta \in \Theta} |f(\theta) - \bar{f}_{t-1}(\theta)|] < \sum_{t=1}^{\infty} t^{1/2-2u} < \infty, \quad (56)$$

as **(E)** assumes $u \in (\frac{3}{4}, 1)$. Let us now focus on the second term of (55). We define, for all $1 \leq i \leq t$, $w_i^t = w_i \prod_{j=i+1}^t (1 - w_j)$. Then

$$\mathbb{E}[\bar{\epsilon}_t] = \sum_{i=1}^t w_i^t \mathbb{E}[\epsilon_t] \leq w_t \sum_{i=1}^t \mathbb{E}[\epsilon_t] \quad (57)$$

We can safely assume $2(u-1) - \gamma > -1$ by setting a sufficiently small γ . As $\mathbb{E}[\epsilon_t] \in \mathcal{O}(t^{2(u-1)-\gamma})$, this allows to bound the partial sum $\sum_{i=1}^t \mathbb{E}[\epsilon_i] = \mathcal{O}(t^{2u-1-\gamma})$. Therefore

$$w_t \mathbb{E}[\bar{\epsilon}_{t-1} + \mathbb{E}[\epsilon_t | \mathcal{F}_{t-1}]] \leq w_t^2 \left(\sum_{i=1}^t \mathbb{E}[\epsilon_i] \right) + w_t \mathbb{E}[\epsilon_t] \leq A t^{-1-\gamma} + B t^{u-2-\gamma} \leq C t^{-1-\gamma}, \quad (58)$$

where we use $u < 1$. Therefore $\sum_{t=1}^{\infty} w_t \mathbb{E}[\bar{\epsilon}_{t-1} + \mathbb{E}[\epsilon_t | \mathcal{F}_{t-1}]] < \infty$. We use quasi-martingale theory to conclude, as in [10]. We define the variable δ_t as 1 if $\mathbb{E}[\bar{g}_t(\theta_t) - \bar{g}_{t-1}(\theta_{t-1}) | \mathcal{F}_{t-1}] \geq 0$, and 0 otherwise. As all terms of (55) are positive:

$$\begin{aligned} \sum_{t=1}^{\infty} \mathbb{E}[\delta_t (\bar{g}_t(\theta_t) - \bar{g}_{t-1}(\theta_{t-1}))] &= \sum_{t=1}^{\infty} \mathbb{E}[\delta_t \mathbb{E}[\bar{g}_t(\theta_t) - \bar{g}_{t-1}(\theta_{t-1}) | \mathcal{F}_{t-1}]] \\ &\leq \sum_{t=1}^{\infty} w_t \mathbb{E}[\sup_{\theta \in \Theta} |f(\theta) - \bar{f}_{t-1}(\theta)|] + \sum_{t=1}^{\infty} w_t \mathbb{E}[\bar{\epsilon}_{t-1} + \mathbb{E}[\epsilon_t | \mathcal{F}_{t-1}]] < \infty. \end{aligned} \quad (59)$$

As \bar{g}_t are bounded from below (\bar{f}_t is bounded from **(B)** and we easily show that $\bar{\epsilon}_t$ is bounded), we can apply Theorem A.1 from [10], that is a quasi-martingale convergence theorem originally found in [36]. It ensures that $(g_t(\theta_t))_{t \geq 1}$ converges almost surely to an integrable random variable g^* , and that $\sum_{t=1}^{\infty} \mathbb{E}[|\mathbb{E}[\bar{g}_t(\theta_t) - \bar{g}_{t-1}(\theta_{t-1}) | \mathcal{F}_{t-1}]|] < \infty$ almost surely.

A.2.2 Almost sure convergence of $\bar{f}(\theta_t)$

We rewrite the second inequality of (54), adding $\bar{\epsilon}_t$ on both sides:

$$\begin{aligned} 0 \leq w_t (\bar{g}_{t-1}(\theta_{t-1}) - \bar{f}_{t-1}(\theta_{t-1}) + \bar{\epsilon}_{t-1}) &\leq w_t (g_t(\theta_{t-1}) - f_t(\theta_{t-1})) + w_t (f_t(\theta_{t-1}) - \bar{f}_{t-1}(\theta_{t-1})) \\ &\quad + (\bar{g}_{t-1}(\theta_{t-1}) - \bar{g}_t(\theta_t)) + w_t \bar{\epsilon}_{t-1} \\ &\leq w_t (f_t(\theta_{t-1}) - \bar{f}_{t-1}(\theta_{t-1})) + (\bar{g}_{t-1}(\theta_{t-1}) - \bar{g}_t(\theta_t)) \\ &\quad + w_t (\epsilon_t + \bar{\epsilon}_{t-1}), \end{aligned} \quad (60)$$

where the left side bound has been obtained in the last paragraph by induction and the right side bound arises from the definition of ϵ_t . Taking the expectation of (60) conditioned on \mathcal{F}_{t-1} , almost surely,

$$\begin{aligned} 0 \leq w_t (\bar{g}_{t-1}(\theta_{t-1}) - \bar{f}_{t-1}(\theta_{t-1}) + \bar{\epsilon}_{t-1}) &\leq w_t (f(\theta_{t-1}) - \bar{f}_{t-1}(\theta_{t-1})) - \mathbb{E}[\bar{g}_t(\theta_t) - \bar{g}_{t-1}(\theta_{t-1}) | \mathcal{F}_{t-1}] \\ &\quad + w_t (\bar{\epsilon}_{t-1} + \mathbb{E}[\epsilon_t | \mathcal{F}_{t-1}]), \end{aligned} \quad (61)$$

We separately study the three terms of the previous upper bound. The first two terms can undergo the same analysis as in [10]. First, almost sure convergence of $\sum_{t=1}^{\infty} \mathbb{E}[|\mathbb{E}[\bar{g}_t(\theta_t) - \bar{g}_{t-1}(\theta_{t-1}) | \mathcal{F}_{t-1}]|]$ implies that $\mathbb{E}[\bar{g}_t(\theta_t) - \bar{g}_{t-1}(\theta_{t-1}) | \mathcal{F}_{t-1}]$ is the summand of an almost surely converging sum. Second, $w_t (f(\theta_{t-1}) - \bar{f}_{t-1}(\theta_{t-1}))$ is the summand of an absolutely converging sum with probability one, less it would contradict (56). To bound the third term, we have once more to control the perturbation introduced by $(\epsilon_t)_t$. We have $\sum_{t=1}^{\infty} w_t \bar{\epsilon}_{t-1} + w_t \mathbb{E}[\epsilon_t | \mathcal{F}_{t-1}] < \infty$ almost surely, otherwise Fubini theorem would invalidate (58).

As the three terms are the summand of absolutely converging sums, the positive term $w_t(\bar{g}_{t-1}(\theta_{t-1}) - \bar{f}_{t-1}(\theta_{t-1}) + \bar{\epsilon}_{t-1})$ is the summand of an almost surely convergent sum. This is not enough to prove that $\bar{h}_t(\theta_t) \triangleq \bar{g}_t(\theta_t) - \bar{f}_t(\theta_t) \rightarrow_{\infty} 0$, hence we follow [10] and make use of its Lemma A.6. We define $X_t \triangleq \bar{h}_{t-1}(\theta_{t-1}) + \bar{\epsilon}_{t-1}$. As **(F)** holds, we can use Lemma 3, which ensures that $(\bar{h}_t)_{t \geq 1}$ are uniformly R' -Lipschitz and $\|\theta_t - \theta_{t-1}\|_2 = \mathcal{O}(w_t)$. Hence,

$$\begin{aligned} |X_{t+1} - X_t| &\leq |\bar{h}_t(\theta_t) - \bar{h}_{t-1}(\theta_{t-1})| + |\bar{\epsilon}_t - \bar{\epsilon}_{t-1}| \\ &\leq R' \|\theta_t - \theta_{t-1}\|_2 + |\bar{\epsilon}_t - \bar{\epsilon}_{t-1}|, \quad \text{as } \bar{h}_t \text{ is } R' \text{-Lipschitz} \\ &\leq \mathcal{O}(w_t) + |\bar{\epsilon}_t - \bar{\epsilon}_{t-1}|, \quad \text{as } \|\theta_t - \theta_{t-1}\|_2 = \mathcal{O}(w_t) \end{aligned} \quad (62)$$

From assumption **(F)**, $(\epsilon_t)_t$ and $(\bar{\epsilon}_t)_t$ are bounded and therefore

$$|\bar{\epsilon}_t - \bar{\epsilon}_{t-1}| \leq w_t(|\epsilon_t| + |\bar{\epsilon}_{t-1}|) \in \mathcal{O}(w_t), \quad \text{hence } |X_{t+1} - X_t| \leq \mathcal{O}(w_t). \quad (63)$$

Lemma A.6 from [10] then ensures that X_t converges to zero with probability one. Assumption **(F)** ensures that $\epsilon_t \rightarrow_{\infty} 0$ almost surely, from which we can easily deduce $\bar{\epsilon}_t \rightarrow_{\infty} 0$ almost surely. Therefore $\bar{h}_t(\theta_t) \rightarrow 0$ with probability one and $(\bar{f}_t(\theta_t))_{t \geq 1}$ converges almost surely to g^* .

A.2.3 Almost sure convergence of $f(\theta_t)$

Lemma B.7 from [10], based on empirical process theory [33] ensures that \bar{f}_t uniformly converges to \bar{f} . Therefore, $(\bar{f}(\theta_t))_{t \geq 1}$ converges almost surely to g^* .

A.2.4 Asymptotic stationary point condition

Preliminary to the final result, we establish the asymptotic stationary point condition (65) as in [10]. This requires to adapt the original proof to take into account the errors in surrogate computation and minimization. We set $\alpha > 0$. By definition, $\nabla \bar{h}_t$ is L -Lipschitz over \mathbb{R}^K . Therefore, setting $\theta' = \theta_t - \alpha \frac{\nabla \bar{h}_t(\theta_t)}{\|\nabla \bar{h}_t(\theta_t)\|_2}$,

$$\begin{aligned} \bar{h}_t(\theta') &\leq \bar{h}_t(\theta_t) - \alpha \|\nabla \bar{h}_t(\theta_t)\|_2 + \frac{L\alpha^2}{2} \\ \|\nabla \bar{h}_t(\theta_t)\|_2 &\leq \frac{1}{\alpha} (\bar{h}_t(\theta_t) - \bar{h}_t(\theta')) + \frac{L\alpha}{2}, \quad \text{and thus } \leq \frac{2}{\alpha} \bar{\epsilon}_t + \frac{L\alpha}{2}, \end{aligned} \quad (64)$$

where we use $|\bar{h}_t(\theta)| \leq \bar{\epsilon}_t$ for all $\theta \in \mathbb{R}^K$. As $\bar{\epsilon}_t \rightarrow 0$ and the inequality (64) is true for all α , we draw $\|\nabla \bar{h}_t(\theta_t)\|_2 \rightarrow_{\infty} 0$ almost surely. From the strong convexity of \bar{g}_t and Lemma 3, $\|\theta_t - \theta_t^*\|_2$ converges to zero, which ensures

$$\|\nabla \bar{h}_t(\theta_t^*)\|_2 \leq \|\nabla \bar{h}_t(\theta_t)\|_2 + L \|\theta_t - \theta_t^*\|_2 \rightarrow_{\infty} 0. \quad (65)$$

A.2.5 Parametrized surrogates

We use assumption **(D)** to finally prove the property, adapting the proof of Proposition 3.4 in [10]. We first recall the derivations of [10] for obtaining (66). We define $(\kappa_t)_t$ such that $\bar{g}_t = g_{\kappa_t}$ for all $t > 0$. We assume that θ_{∞} is a limit point of $(\theta_t)_t$. As Θ is compact, there exists an increasing sequence $(t_k)_k$ such that $(\theta_{t_k})_k$ converges toward θ_{∞} . As \mathcal{K} is compact, a converging subsequence of $(\kappa_{t_k})_k$ can be extracted, that converges towards $\kappa_{\infty} \in \mathcal{K}$. From the sake of simplicity, we drop subindices and assume without loss of generality that $\theta_t \rightarrow \theta_{\infty}$ and $\kappa_t \rightarrow \kappa_{\infty}$. From the compact parametrization assumption, we easily show that $(\bar{g}_{\kappa_t})_t$ uniformly converges towards $\bar{g}_{\infty} \triangleq \bar{g}_{\kappa_{\infty}}$. Then defining $\bar{h}_{\infty} = \bar{g}_{\infty} - \bar{f}$,

$$\forall \theta \in \Theta, \quad \nabla \bar{f}(\theta_{\infty}, \theta - \theta_{\infty}) = \nabla \bar{g}_{\infty}(\theta_{\infty}, \theta - \theta_{\infty}) - \nabla \bar{h}_{\infty}(\theta_{\infty}, \theta - \theta_{\infty}). \quad (66)$$

Let us turn to show $\nabla \bar{f}(\theta_{\infty}, \theta - \theta_{\infty}) \geq 0$ for all $\theta \in \Theta$. We consider the sequence $(\theta_t^*)_t$. From Lemma 3, $\|\theta_t - \theta_t^*\|_2 \rightarrow 0$, which implies $\theta_t^* \rightarrow \theta_{\infty}$. \bar{g}_t converges uniformly towards \bar{g}_{∞} , which implies $(g_t(\theta_t^*))_t \rightarrow g_{\infty}(\theta_{\infty})$. Yet for all $t > 0$, $\theta \in \Theta$,

$$g_t(\theta_t^*) \leq g_t(\theta), \quad \text{which implies } g_{\infty}(\theta_{\infty}) \leq \inf_{\theta \in \Theta} g_{\infty}(\theta), \quad (67)$$

by taking the limit for $t \rightarrow \infty$. Therefore θ_∞ is the minimizer of \bar{g}_∞ and thus $\nabla \bar{g}_\infty(\theta_\infty, \theta - \theta_\infty) \geq 0$.

Adapting [10], we perform the first-order expansion of \bar{h}_t around θ_t^* (instead of θ_t in the original proof) and show that $\nabla \bar{h}_\infty(\theta_\infty, \theta - \theta_\infty) = 0$, as \bar{h}_t differentiable, $\|\nabla \bar{h}_t(\theta_t^*)\|_2 \rightarrow 0$ and $\theta_t^* \rightarrow \theta_\infty$. This is sufficient to conclude. \square

A.3 Convergence of SOMF — Proof of Proposition 3

Proof of Proposition 3. From assumption **(B)**, $(x_t)_t$ is ℓ_2 -bounded by a constant X . With assumption **(H)**, it implies that $(\alpha_t)_t$ is ℓ_2 -bounded by a constant A . This is enough to show that $(g_t)_t$ and $(\theta_t)_t$ meet basic assumptions **(A)**–**(D)**. Assumption **(E)** immediately implies **(I)**. It remains to show that $(g_t)_t$ and $(\theta_t)_t$ meets the assumptions **(F)** and **(G)**. This will allow to cast SOMF as an instance of SAMM and conclude.

A.3.1 The computation of \mathbf{D}_t verifies **(G)**

We write $\mathbf{D}_t^* = \operatorname{argmin}_{\mathbf{D} \in \mathcal{C}} \bar{g}_t(\mathbf{D})$. We want to show that performing subsampled block coordinate descent on \bar{g}_t is sufficient to meet assumption **(G)**.

Let us first focus on the case where $\mathbf{M}_t = \mathbf{P}_t = \mathbf{I}$, i.e., the setting with no subsampling. In this setting, at iteration t , \mathbf{D}_t is obtained by performing k updates of projected block coordinate descent on function \bar{g}_t , starting from \mathbf{D}_{t-1} . For all $j \in [0, k-1]$, for all $\mathbf{D} \in \mathbb{R}^{p \times k}$,

$$\nabla \bar{g}_t(\mathbf{D}) = \mathbf{D} \bar{\mathbf{C}}_t - \bar{\mathbf{B}}_t \quad (68)$$

We make use of convergence rates in projected block coordinate descent, established in [37, 38]. For $\mathbf{D}^{(j)} \in \mathbb{R}^{p \times k}$ any matrix with non-zero j -th column $\mathbf{d}^{(j)}$ and zero elsewhere

$$\nabla \bar{g}_t(\mathbf{D} + \mathbf{D}^{(j)}) - \nabla \bar{g}_t(\mathbf{D}) = \bar{\mathbf{C}}_t[j, j] \mathbf{d}^{(j)} \quad (69)$$

and hence \bar{g}_t gradient has *component Lipschitz constant* $L_j = \bar{\mathbf{C}}_t[j, j]$ for component j , as already noted in [15]. Using [38] terminology, $\nabla \bar{g}_t$ has *coordinate Lipschitz constant* $L_{\max} \triangleq \max_{0 \leq j < k} \bar{\mathbf{C}}_t[j, j] \leq \max_{t > 0, 0 \leq j < k} \alpha_t[j]^2 \leq A^2$, as $(\alpha_t)_t$ is bounded from **(H)**. As a consequence, \bar{g}_t gradient is also L -Lipschitz continuous, where [38] note that $L \leq \sqrt{k} L_{\max}$. Moreover, \bar{g}_t is strongly convex with strong convexity modulus $\rho > 0$ by hypothesis **(H)**. Then, [37] ensures that after one cycle over the k blocks

$$\begin{aligned} \mathbb{E}[\bar{g}_t(\mathbf{D}_t) - \bar{g}_t(\mathbf{D}_t^*) | \mathcal{F}_{t-1}, \mathbf{M}_t = \mathbf{I}_p] &\leq \left(1 - \frac{\rho}{2L_{\max}(1 + kL^2/L_{\max}^2)}\right) (\bar{g}_t(\mathbf{D}_{t-1}) - \bar{g}_t(\mathbf{D}_t^*)) \\ &\leq (1 - \mu) (\bar{g}_t(\mathbf{D}_{t-1}) - \bar{g}_t(\mathbf{D}_t^*)) \quad \text{where } \mu \triangleq \frac{\rho}{2A^2(1 + k^2)} \end{aligned} \quad (70)$$

In the general case where $\mathbf{M}_t \neq \mathbf{I}_p$, we only write the deterministic decrease of the surrogates, i.e., $\bar{g}_t(\mathbf{D}_t) \leq \bar{g}_t(\mathbf{D}_{t-1})$. Then, writing $d = \mathbb{P}[\mathbf{M}_t = \mathbf{I}_p]$,

$$\begin{aligned} \mathbb{E}[\bar{g}_t(\mathbf{D}_t) - \bar{g}_t(\mathbf{D}_t^*) | \mathcal{F}_{t-1}] &= d \mathbb{E}[\bar{g}_t(\mathbf{D}_t) - \bar{g}_t(\mathbf{D}_t^*) | \mathcal{F}_{t-1}, \mathbf{M}_t = \mathbf{I}_p] \\ &\quad + (1 - d) \mathbb{E}[\bar{g}_t(\mathbf{D}_t) - \bar{g}_t(\mathbf{D}_t^*) | \mathcal{F}_{t-1}, \mathbf{M}_t \neq \mathbf{I}_p] \\ &\leq (1 - d\mu) (\bar{g}_t(\mathbf{D}_{t-1}) - \bar{g}_t(\mathbf{D}_t^*)). \end{aligned} \quad (71)$$

Therefore, the surrogates $(g_t)_t$ used in SOMF meets assumption **(G)**.

A.3.2 The surrogates $(g_t)_t$ verify **(F)**

We define $g_t^* \in \mathcal{S}_{\rho, L}(f_t, \mathbf{D}_{t-1})$ as the surrogate of OMF at iteration t , that uses the exact computation of α_t (21), and g_t the surrogate used in SOMF.⁴ Namely, using the auxiliary loss function $l(\alpha, \mathbf{G}, \beta) \triangleq \frac{1}{2} \alpha^\top \mathbf{G} \alpha - \alpha^\top \beta + \lambda \Omega(\alpha)$, we define g_t^* and g_t as

$$\begin{aligned} \alpha_t^* &= \operatorname{argmin}_{\alpha \in \mathbb{R}^k} l(\alpha, \mathbf{G}_t^*, \beta_t^*), & \alpha_t &= \operatorname{argmin}_{\alpha \in \mathbb{R}^k} l(\alpha, \mathbf{G}_t, \beta_t), \\ g_t^*(\mathbf{D}) &\triangleq l(\alpha_t^*, \mathbf{D}^\top \mathbf{D}, \mathbf{D}^\top \mathbf{x}_t), & g_t(\mathbf{D}) &\triangleq l(\alpha_t, \mathbf{D}^\top \mathbf{D}, \mathbf{D}^\top \mathbf{x}_t), \end{aligned} \quad (72)$$

⁴In the following, we abusely write that the pointwise surrogates g_t^* and g_t are ρ -strongly-convex to match the Definition 1 and 2, whereas **(H)** ($\bar{\mathbf{C}}_t \succ \rho \mathbf{I}_k$) only ensures that \bar{g}_t^* and \bar{g}_t are ρ -strongly convex. It is not a problem as we only use the strong convexity of the average surrogates $(\bar{g}_t)_t$ in all proofs.

where the matrices \mathbf{G}_t^* , β_t^* are defined in (21) and \mathbf{G}_t , β_t in (b) or (c). We define $\epsilon_t \triangleq \|g_t^* - g_t\|_\infty = \sup_{\mathbf{D} \in \mathcal{C}} |g_t^*(\mathbf{D}) - g_t(\mathbf{D})|$. We define $\gamma \triangleq \frac{1}{2} \min(v - \frac{3}{4}, (3u - 2) - v)$. We verify $\gamma > 0$ from assumption **(I)**. By definition, $g_t \in \mathcal{T}_{\rho, L}(f, \theta_{t-1}, \epsilon_t)$. To ensure that assumption **(F)** is met, we show $\mathbb{E}[\epsilon_t] \in \mathcal{O}(t^{2(u-1)-\gamma})$ and $\epsilon_t \rightarrow 0$ almost surely.

We define $\mathcal{A} = \{\alpha \in \mathbb{R}^k, \|\alpha\|_2 \leq A\}$. We first establish that the function $\phi : \mathbf{G}, \beta \rightarrow \operatorname{argmin}_{\alpha \in \mathcal{A}} l(\alpha, \mathbf{G}, \beta)$ is $\frac{A}{\rho}$ -Lipschitz continuous. We introduce the set Φ , image of $\mathcal{C} \times \mathcal{X}$ by the function $(\mathbf{D}, \mathbf{x}) \rightarrow \mathbf{D}^\top \mathbf{D}, \mathbf{D}^\top \mathbf{x}$. The function $(\mathbf{G}, \beta) \rightarrow l(\alpha, \mathbf{G}, \beta)$ is defined on Φ and differentiable for all $\alpha \in \mathbb{R}^k$. Its gradient writes $\nabla_{(\mathbf{G}, \beta)} l((\mathbf{G}, \beta), \alpha) = (\frac{1}{2} \alpha^\top \alpha, -\beta)$. As a function of α , it is therefore A -Lipschitz continuous on \mathcal{A} , for all $(\mathbf{G}, \beta) \in \Phi$. Finally, from **(H)**, $\alpha \rightarrow l((\mathbf{G}, \beta), \alpha)$ is ρ -strongly convex for all $(\mathbf{G}, \beta) \in \Phi$. Therefore, using Lemma B.6 from [39] that guarantees the Lipschitz continuity of minimizers for parametrized functions, ϕ is $\frac{A}{\rho}$ -Lipschitz. As the function $\alpha \rightarrow l(\alpha, \mathbf{D}^\top \mathbf{D}, \mathbf{D}^\top \mathbf{x}_t)$ is $(A + X)$ -Lipschitz, for all $\mathbf{D} \in \mathcal{C}$,

$$|g_t^*(\mathbf{D}) - g_t(\mathbf{D})| \leq (A + X) \|\phi(\mathbf{G}_t^*, \beta_t^*) - \phi(\mathbf{G}_t, \beta_t)\|_2 \leq \frac{A}{\mu} (A + X) (\|\mathbf{G}_t^* - \mathbf{G}_t\|_2 + \|\beta_t^* - \beta_t\|_2). \quad (73)$$

Showing **(F)** thus reduces to bound $\|\mathbf{G}_t^* - \mathbf{G}_t\|_2$ and $\|\beta_t^* - \beta_t\|_2$ in expectation and almost surely, for both update rules (b) or (c). We focus on bounding $\|\beta_t - \beta_t^*\|_2$ — bounding $\|\mathbf{G}_t - \mathbf{G}_t^*\|_2$ in the case (b) is done in the exact same way. Without loss of generality, we assume that the streaming order $(i_t)_t$ is known, and condition all expectations on this knowledge — extending to the full random case simply requires to integrate over $(i_t)_t$. For $t > 0$, we write $i \triangleq i_t$ and recall

$$\beta_t \triangleq \beta_t^{(i)} = \sum_{s \leq t, \mathbf{x}_s = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \mathbf{D}_{s-1}^\top \mathbf{M}_s \mathbf{x}^{(i)},$$

where $\gamma_{s,t}^{(i)} = \gamma_{c_t^{(i)}} \prod_{s < t, \mathbf{x}_s = \mathbf{x}^{(i)}} (1 - \gamma_{c_s^{(i)}})$ and $c_t^{(i)} = |\{s \leq t, \mathbf{x}_s = \mathbf{x}^{(i)}\}|$. We can decompose $\beta_t - \beta_t^*$ as

$$\begin{aligned} \beta_t - \beta_t^* &= \sum_{s \leq t, \mathbf{x}_s = \mathbf{x}_t = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} (\mathbf{D}_{s-1} - \mathbf{D}_{t-1})^\top \mathbf{M}_s \mathbf{x}^{(i)} + \mathbf{D}_{t-1}^\top \left(\sum_{s \leq t, \mathbf{x}_s = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \mathbf{M}_s - \mathbf{I} \right) \mathbf{x}^{(i)} \\ \|\beta_t - \beta_t^*\|_2 &\leq B \sum_{s \leq t, \mathbf{x}_s = \mathbf{x}_t = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \|\mathbf{D}_{s-1} - \mathbf{D}_{t-1}\|_2 + C \left\| \left(\sum_{s \leq t, \mathbf{x}_s = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \mathbf{M}_s \right) - \mathbf{I} \right\|_2, \end{aligned} \quad (74)$$

where B and C are constants. The inequality involves a left term due to the averaging in the estimator construction, and a right term that is a norm of a vector whose coefficients are the running average of zero-mean Bernoulli variables. The latter term indeed rewrites as

$$B \|\bar{\mathbf{M}}_t^{(i)} - \mathbf{I}\|_2 \quad \text{where} \quad \bar{\mathbf{M}}_t^{(i)} \triangleq \sum_{s \leq t, \mathbf{x}_s = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \mathbf{M}_s = (1 - \gamma_{c_t^{(i)}}) \bar{\mathbf{M}}_{t-1}^{(i)} + \gamma_{c_t^{(i)}} \mathbf{M}_t. \quad (75)$$

We may use Lemma B.7 from [10], applied to constant functions, to obtain the following rates in expectation:

$$\mathbb{E}[\|\bar{\mathbf{M}}_t^{(i_t)} - \mathbf{I}\|_2 | (i_t)_t] \leq C \gamma_{c_t^{(i_t)}}^{(i_t)} \sqrt{c_t^{(i_t)}} \leq C' \gamma_t \sqrt{t} \quad \text{almost surely} \quad (76)$$

with C, C' constants, as $(i_t)_t$ is drawn uniformly at random. Therefore $\mathbb{E}[\|\bar{\mathbf{M}}_t^{(i_t)} - \mathbf{I}\|_2] \leq C' t^{1/2-v} \leq C' t^{2(u-1)-\gamma}$, as $\gamma \leq v - \frac{3}{4} \leq v - (\frac{5}{2} - 2u)$ by construction and assumption **(I)**. Following the same lemma from [10], we also have $\epsilon_t \rightarrow 0$ almost surely.

We now turn to bound the left term. For simplicity, we drop the dependency in i when defining the next values. We write $(t_k)_{k \geq 0}$ the times at which $\mathbf{x}^{(i)}$ is drawn, $j = c_t^{(i)}$ the current counter of $\mathbf{x}^{(i)}$ apparitions, i.e, such that $t_j = t$. We then define $G_{k,j} \triangleq \gamma_{t_k, t_j}^{(i)}$ for all $k \leq j$. We set $w \triangleq 3u - 2 - \gamma \in (v, 1)$ by definition of γ and split the sum at $k_j \triangleq j + 1 - \lfloor j^w \rfloor$:

$$\begin{aligned} \sum_{s \leq t, \mathbf{x}_s = \mathbf{x}_t = \mathbf{x}^{(i)}} \gamma_{s,t}^{(i)} \|\mathbf{D}_{s-1} - \mathbf{D}_{t-1}\|_2 &= \sum_{k=1}^j G_{k,j} \|\mathbf{D}_{t_k-1} - \mathbf{D}_{t_{j-1}-1}\|_2 \\ &\leq 2\mathcal{D} \sum_{k=1}^{k_j} G_{k,j} + \sum_{k=k_j}^j G_{k,j} \sum_{s=t_k-1}^{t_j-1} w_s. \end{aligned} \quad (77)$$

We have bounded $\|\mathbf{D}_t\|_2^2$ by the diameter \mathcal{D} of \mathcal{C} in the left part of the split, and used Lemma 3 in the right part. It applies as **(G)** is met and (73) ensures that $(\|g_t - g_t^*\|_\infty)_t$ is bounded by $\epsilon \geq 0$. Bounding (77) requires to further study the two part of (77). For the left part, we use the exponential decay in the reuse of past iterates that arises from $v < w < 1$. For all $j > 0$,

$$\begin{aligned} \sum_{k=1}^{k_j} G_{k,j} &= \sum_{k=0}^{k_j} \gamma_k \prod_{p=k+1}^{k_j} (1 - \gamma_p) \leq \sum_{k=1}^{k_j} \gamma_k (1 - \gamma_j)^{j-k} \leq \frac{(1 - \gamma_j)^{\lfloor j^w \rfloor}}{\gamma_j} = j^v \exp\left(\ln\left(1 - \frac{1}{j^v}\right) \lfloor j^w \rfloor\right) \\ &\in \mathcal{O}(t^{2(u-1)-\gamma}), \end{aligned} \quad (78)$$

as $w > v$, from basic asymptotic comparisons. For the right part, we use the fact that $(w_t)_t$ decays sufficiently rapidly. As $\sum_{k=1}^j G_{k,j} = 1$, we have, with C constant

$$\begin{aligned} \sum_{k=k_j}^j G_{k,j} \sum_{s=t_k-1}^{t_j-1} w_s &\leq \max_{k_j \leq k \leq j} \left(\sum_{s=t_k-1}^{t_j-1} w_s \right) = \sum_{s=t_{k_j}-1}^{t_j-1} w_s \leq w_{j-\lfloor j^w \rfloor} \lfloor j^w \rfloor = \frac{\lfloor j^w \rfloor}{j^u \left(1 - \frac{\lfloor j^w \rfloor}{j}\right)^u} \leq C j^{w-u} \\ &\leq C j^{2(u-1)-\gamma}, \end{aligned} \quad (79)$$

from the definition of w . This is enough to show $(\epsilon_t)_t$ verifies **(F)**. The SOMF algorithm can therefore be casted as stochastic approximate majorization-minimization algorithm, where all assumptions are met for convergence. We conclude using Proposition 2. \square