

Distributionally Robust and Multi-Objective Nonnegative Matrix Factorization

Nicolas Gillis, *Member, IEEE*, Le Thi Khanh Hien, Valentin Leplat,
and Vincent Y. F. Tan, *Senior Member, IEEE*

Abstract—Nonnegative matrix factorization (NMF) is a linear dimensionality reduction technique for analyzing nonnegative data. A key aspect of NMF is the choice of the objective function that depends on the noise model (or statistics of the noise) assumed on the data. In many applications, the noise model is unknown and difficult to estimate. In this paper, we define a multi-objective NMF (MO-NMF) problem, where several objectives are combined within the same NMF model. We propose to use Lagrange duality to judiciously optimize for a set of weights to be used within the framework of the weighted-sum approach, that is, we minimize a single objective function which is a weighted sum of the all objective functions. We design a simple algorithm based on multiplicative updates to minimize this weighted sum. We show how this can be used to find distributionally robust NMF (DR-NMF) solutions, that is, solutions that minimize the largest error among all objectives, using a dual approach solved via a heuristic inspired from the Frank-Wolfe algorithm. We illustrate the effectiveness of this approach on synthetic, document and audio data sets. The results show that DR-NMF is robust to our incognizance of the noise model of the NMF problem.

Index Terms—Nonnegative matrix factorization, Multiple objectives, Distributional robustness, Multiplicative updates

1 INTRODUCTION

Nonnegative matrix factorization (NMF) consists in the following problem: Given a nonnegative matrix $X \in \mathbb{R}_+^{m \times n}$ and a factorization positive rank $r \ll \min(m, n)$, find two nonnegative matrices $W \in \mathbb{R}_+^{m \times r}$ and $H \in \mathbb{R}_+^{r \times n}$ such that $WH \approx X$. NMF is a linear dimensionality reduction technique for nonnegative data. In fact, assuming each column of X is a data point, it is reconstructed via a linear combination of r basis elements given by the columns of W while the columns of H provide the weights (or coefficients) to reconstruct each column of X within that basis, that is, for all j ,

$$X(:, j) \approx \sum_{k=1}^r W(:, k)H(k, j).$$

NMF has attracted a lot of attention since the seminal paper of Lee and Seung [1], with applications in image analysis, document classification and music analysis. See for example [2], [3] and the references therein. Many NMF models have been proposed over the years. They mostly differ in two aspects:

- 1) *Additional constraints* are added to the factor matrices W and H such as sparsity [4], spatial coherence [5] or

smoothness [6]. These constraints are motivated by a priori information on the sought solution and depend on the application at hand. Note that these additional constraints are in most cases imposed via a penalty term in the objective function.

- 2) The choice of the *objective function* that assesses the quality of an approximation by evaluating some distance between WH and X differs. This choice is usually motivated by the noise model/statistics assumed on the data matrix X . The most widely used class of objective functions are component-wise and based on the β -divergences defined as follows: for $x, y \in \mathbb{R}_+$,

$$D_\beta(x, y) = \begin{cases} \frac{x}{y} - \log \frac{x}{y} - 1 & \text{for } \beta = 0, \\ x \log \frac{x}{y} - x + y & \text{for } \beta = 1, \\ \frac{1}{\beta(\beta-1)} (x^\beta + (\beta-1)y^\beta - \beta xy^{\beta-1}) & \text{for } \beta \neq 0, 1. \end{cases}$$

We will use the following matrix-wise notation,

$$D_\beta(X, WH) = \sum_{i,j} D_\beta(X_{ij}, (WH)_{ij}).$$

Minimizing the β -divergence in NMF is equivalent to maximizing the log-likelihood of the NMF model under different noise distributions [7], [8]. The following special cases are of particular interest (see for example [7] for a discussion):

- $D_2(X, WH) = \frac{1}{2} \|X - WH\|_F^2$ is the Frobenius norm (additive Gaussian noise).
- $D_1(X, WH) = \text{KL}(X, WH)$ is the Kullback-Leibler (KL) divergence (Poisson noise).
- $D_0(X, WH) = \text{IS}(X, WH)$ is the Itakura-Saito (IS) divergence (multiplicative Gamma noise).

• N. Gillis, L. T. K. Hien and V. Leplat are with Department of Mathematics and Operational Research, Faculté Polytechnique, Université de Mons, Rue de Houdain 9, 7000 Mons, Belgium. This work was supported by the Fonds de la Recherche Scientifique - FNRS and the Fonds Wetenschappelijk Onderzoek - Vlaanderen (FWO) under EOS Project no O005318F-RG47, and by the European Research Council (ERC starting grant no 679515).

E-mails: {nicolas.gillis, thikhanhien.le, valentin.leplat}@umons.ac.be

• V. Y. F. Tan is with the Department of Electrical and Computer Engineering, Department of Mathematics, National University of Singapore, Singapore 119077. This work is also supported by a Singapore National Research Foundation (NRF) Fellowship (R-263-000-D02-281).

E-mail: vtan@nus.edu.sg

Manuscript accepted, February 2021.

In this paper, we focus on the second aspect, namely, the choice of the objective function. We will consider a multi-objective NMF (MO-NMF) formulation. More precisely, we will consider a weighted sum of the different objective functions, which is arguably one of the most widely used approach in multi-objective optimization [9]. Our main motivation to consider this class of models is that in many applications it is not clear which objective function to use because the statistics of the noise is unknown. To the best of our knowledge, there are currently three main classes of methods to handle this situation:

- The user chooses the objective function she/he believes is the most suitable for the application at hand. This is, as far as we know, the simplest and most widely-used approach. However, this approach is an *ad hoc* one.
- The objective function is automatically selected using cross-validation, where the training is done on a subset of the entries of the input data matrix and the testing on the remaining entries [10], [11].
- The most suitable objective function is chosen using some statistically motivated criteria such as score matching [12] or maximum likelihood [13].

However, in all the above approaches, if the choice of the objective function is wrong, the NMF solution provided could be far from the desired solution (as we will show in our numerical experiments in Section 5). Another possibility which we propose in this paper is to compute an NMF solution that is robust to different types of noise distributions; this is referred to as *distributionally robust*, and is closely related to robust optimization [14]. In mathematical terms, we will consider the problem

$$\min_{(W,H) \geq 0} \max_{\beta \in \Omega} D_{\beta}(X, WH),$$

As we will see, this problem can be tackled by minimizing a weighted sum of the different objective functions [9], exactly as for MO-NMF, but where the weights assigned to the different objective functions are *automatically tuned* within the iterative process.

Outline of the paper: In Section 2, we first define MO-NMF and explain how to scale the objective functions to make the comparison between the constituent NMF objective functions. Then we give our main motivation to consider MO-NMF, namely to be able to compute distributionally robust NMF (DR-NMF) solutions, that is, solutions that minimize the largest objective function value. In Section 3, we propose simple multiplicative updates (MU) to solve a weighted-sum approach for MO-NMF. In Section 4, we propose a heuristic scheme to solve DR-NMF which updates the primal variables using the MU and the dual variable using the Frank-Wolfe descent direction. Finally, we illustrate in Section 5 the effectiveness of our approach on synthetic, document and audio data sets.

2 MULTI-OBJECTIVE NMF (MO-NMF)

Let Ω be a finite subset of \mathbb{R}_+ . We consider in this paper the following MO-NMF problem:

$$\min_{(W,H) \geq 0} \{D_{\beta}(X, WH)\}_{\beta \in \Omega}.$$

Note that we focus on β -divergences to simplify our presentation and because these are the most widely-used divergences to measure the “distance” between the given matrix X and its approximation WH in the NMF literature. However, our approach can be adapted to be used for other objectives functions (for example, α -divergences [15]). To tackle this problem, we consider the standard weighted-sum approach [16] which consists in solving the following minimization problem which involves a *single* objective function:

$$\min_{(W,H) \geq 0} D_{\Omega}^{\lambda}(X, WH),$$

where $D_{\Omega}^{\lambda}(X, WH) = \sum_{\beta \in \Omega} \lambda_{\beta} D_{\beta}(X, WH)$, $\lambda \in \mathbb{R}_+^{|\Omega|}$, and $\|\lambda\|_1 = \sum_{\beta \in \Omega} \lambda_{\beta} = 1$. Using different values for λ allows to generate different Pareto-optimal solutions. See Section 5.1 for some examples. Note, however, that it does not allow to generate all Pareto-optimal solutions [16]. A Pareto-optimal solution is a solution that is not dominated by any other solution. That is, (W, H) is a Pareto-optimal solution if there does not exist a feasible solution (W', H') such that

- $D_{\beta}(X, W'H') \leq D_{\beta}(X, WH)$ for all $\beta \in \Omega$, and
- there exists $\beta \in \Omega$ such that $D_{\beta}(X, W'H') < D_{\beta}(X, WH)$.

Multi-objective optimization has already been considered for NMF problems. However, most of the existing literature considers combining a single data fitting term with penalty terms on the factor matrices; for example, an ℓ_1 penalty to obtain sparse solutions [17]. As far as we know, the only paper where several objectives are used to balance different data fitting terms is [18]. The authors combined two objectives, one being a standard data fitting term (more precisely, they used the Frobenius norm $\|X - WH\|_F^2$) and the other being a data fitting term in a feature space obtained using a nonlinear kernel (that is, a term of the form $\|\Phi(X) - \Phi(W)H\|_{\mathcal{H}}^2$ where $\|\cdot\|_{\mathcal{H}}$ corresponds to the norm in the feature space). Hence this approach is rather different than ours where we allow more than two objectives and where we only focus on the input space. Moreover, we will optimize the weights in a principled optimization-theoretic fashion, whereas [18] uses an *ad hoc* manner to combine the two terms. Another related work [19] considers a data fusion problem where several data sets, denoted X_1, X_2, \dots, X_p , share the same factor H . Their goal is to compute $H \geq 0$ and $W_i \geq 0$ such that $X_i \approx W_i H$ for $i = 1, 2, \dots, p$. To achieve this goal, the authors use a weighted objective function $\sum_{i=1}^p \lambda_i D_{\beta_i}(X_i, W_i H)$ for some well-chosen weights λ_i 's, and some parameters β_i 's that depend on the noise statistic of the corresponding data set. Again, this is a rather different setup that ours as there is no distributionally robust aspect.

2.1 Scaling of the objectives

It can be easily checked that for any constant $\alpha > 0$, we have

$$D_{\beta}(\alpha X, \alpha WH) = \alpha^{\beta} D_{\beta}(X, WH).$$

Hence the values of the divergences for different values of β depend highly on the scaling of the input matrix. This is usually not a desirable property in practice, since most data sets are not particularly properly scaled and

since scaling simply multiplies the noise by a constant which in most cases does not change its distribution (only its parameters). Therefore, we will scale the objectives to have a meaningful linear combination, in the sense that each term in the sum has the same importance. It will be particularly crucial for our DR-NMF model described in the next section. In fact, as we will see in Section 5, DR-NMF will generate solutions that have small error for all objectives instead of just one; and as such, the solutions inherit superior qualities of the ones generated by different divergences. We will use the following approach to scale the different objective functions. First, we compute a solution (W_β, H_β) for $\min_{(W,H) \geq 0} D_\beta(X, WH)$ to obtain the error $e_\beta = D_\beta(X, W_\beta H_\beta)$. Note that we can only compute this minimization in an approximate fashion because the NMF problem is NP-hard [20]. Then, we define

$$\bar{D}_\beta(X, WH) = \frac{D_\beta(X, WH)}{e_\beta},$$

so that $\bar{D}_\beta(X, W_\beta H_\beta) = 1$. Finally, we will only consider the MO-NMF problem where the objectives $D_\beta(X, WH)$ are replaced by their normalized versions $\bar{D}_\beta(X, WH)$, that is,

$$\min_{(W,H) \geq 0} \bar{D}_\Omega^\lambda(X, WH), \quad (1)$$

where $\bar{D}_\Omega^\lambda(X, WH) = \sum_{\beta \in \Omega} \lambda_\beta \bar{D}_\beta(X, WH)$. In Section 3, we propose a MU algorithm to tackle this problem.

2.2 Main motivation: Distributionally robust NMF

If the noise model on the data is unknown, but it is known that it corresponds to a distribution associated with a β -divergence with $\beta \in \Omega$ (for example, the Tweedie distribution as discussed in [8]), it makes sense to consider the following distributionally robust NMF (DR-NMF) problem

$$\min_{(W,H) \geq 0} \max_{\beta \in \Omega} \bar{D}_\beta(X, WH). \quad (2)$$

We use $\bar{D}_\beta(\cdot, \cdot)$, not $D_\beta(\cdot, \cdot)$, because otherwise, in most cases, the above problem amounts to minimizing a single objective corresponding to the β -divergence with the largest value; see the discussion in Section 2.1 where Ω is a subset of β 's of interest. In Section 4, we will design an algorithm to tackle this problem based on MO-NMF. We remark that, as mentioned in the introduction, Problem (2) is intrinsically a deterministic robust optimization problem. However, since each β -divergence is associated to a distribution of noise (see some examples in the introduction), we prefer using the name DR-NMF for Problem (2) to emphasize its essence, which is finding a solution that is robust to different types of noise distributions.

3 MULTIPLICATIVE UPDATES FOR (1)

In this section, we propose MU for (1) which we will be able to use as a subroutine to tackle MO-NMF and DR-NMF. As with most NMF algorithms, we use an alternating strategy; that is, we will first optimize over the variable W for fixed H and then reverse their roles. By the symmetry of the problem ($X^T = H^T W^T$), we will focus on the update of H ; the update of W can be obtained similarly.

3.1 Deriving MU

Let us recall the standard way MU are derived (see for example [7], [21], [22]) on the following general optimization problem with nonnegativity constraints

$$\min_{x \geq 0} f(x). \quad (3)$$

Let us apply a rescaled gradient descent method to (3), that is, use the following update

$$x^+ = x - B \nabla f(x),$$

where x is the current iterate, x^+ is the next iterate, and B is a diagonal matrix with positive diagonal elements. Let $\nabla_+ f(x) > 0$ and $\nabla_- f(x) > 0$ be such that $\nabla f(x) = \nabla_+ f(x) - \nabla_- f(x)$. Taking $B_{ii} = \frac{x_i}{\nabla_+ f(x)_i}$ for all i , we obtain the following MU rule:

$$\begin{aligned} x^+ &= x - \frac{[x]}{[\nabla_+ f(x)]} \circ (\nabla_+ f(x) - \nabla_- f(x)) \\ &= x \circ \frac{[\nabla_- f(x)]}{[\nabla_+ f(x)]}, \end{aligned} \quad (4)$$

where \circ (resp. $[\cdot]/[\cdot]$) refers to component-wise multiplication (resp. division) between two vectors or matrices. Note that we need strict positivity of $\nabla_+ f(x)$ and $\nabla_- f(x)$, otherwise we would encounter problems involving division by zero or a variable directly set to zero, which is not desirable. Using the above simple rule with proper choices for $\nabla_+ f(x)$ and $\nabla_- f(x)$ leads to algorithms that are, in many cases, guaranteed to not increase the objective function, that is, $f(x^+) \leq f(x)$; see below for some examples, and [22] for a discussion and an unified rule to design such updates. This is a desirable property since it avoids any line-search procedure and also preserves non-negativity naturally. If we cannot guarantee that the updates are non-increasing, the step length can be reduced, that is, use

$$x_\gamma^+ = x - \gamma B \nabla f(x),$$

for some $0 < \gamma \leq 1$ which leads to

$$x_\gamma^+ = (1 - \gamma)x + \gamma x^+.$$

For example, one can set the step size $\gamma = 1/2^k$ for the smallest k such that the error decreases; such a k is guaranteed to exist since the rescaled gradient direction is a descent direction. We implemented such a line search; see Algorithm 1 below. This idea is similar to that in [23]. Moreover, it would be worth investigating the use of regularizers to guarantee convergence to stationary points without the use of a line search [24].

For $x_i = 0$, we have that $B_{ii} = 0$ and the MU are not able to modify x_i : this is the so-called zero-locking phenomenon [25]. A possible way to fix this issue in practice is to use a lower bound ϵ on the entries of x , say $\epsilon = 10^{-16}$, replacing x^+ with $\max(\epsilon, x^+)$. This allows such algorithms to be guaranteed to converge to a stationary point of $\min_{x \geq \epsilon} f(x)$ [26], [27]. More precisely, any sequence of solutions generated by the modified MU has at least one convergent subsequence and the limit of any convergent subsequence is a stationary point [27]. Moreover, it can also be shown [26, Chap. 4.1] that such stationary points are close to stationary points of the original problem in (3). We will use this simple strategy in this paper.

3.2 Multiplicative Updates for (1)

We now provide more details on how to choose $\nabla_- f(x)$ and $\nabla_+ f(x)$ for the family of β -divergences in order to tackle (1). For all β , we have

$$\nabla^H D_\beta(X, WH) = \nabla_+^H D_\beta(X, WH) - \nabla_-^H D_\beta(X, WH),$$

where ∇^H denotes the gradient with respect to variable H , and

$$\begin{aligned} \nabla_+^H D_\beta(X, WH) &= W^T (WH)^{\circ(\beta-1)} \quad \text{and} \\ \nabla_-^H D_\beta(X, WH) &= W^T \left((WH)^{\circ(\beta-2)} \circ X \right), \end{aligned}$$

where $A^{\circ k}$ is the component-wise exponentiation by k of the matrix A . To solve (1) using MU, we simply use the linear combination of the above standard choice [8]; see Algorithm 1 for the update of H (the update for W is obtained in the same way by symmetry). Note that the line-search procedure (steps 3 to 6) is very rarely entered (we have only observed it in all our numerical experiments described in Section 5 when $\Omega = \{0\}$, that is, only for IS-NMF alone). Note also that the only difference between \bar{D}_β and D_β is a constant term; see Section 2.1. In the case of a single objective (i.e., that $|\Omega| = 1$), Algorithm 1 particularizes to the standard MU algorithm for NMF; see for example [7] and the references therein.

Algorithm 1 MU for H to solve (1)

Input: The matrices $X \in \mathbb{R}_+^{m \times n}$ and $W \in \mathbb{R}_+^{m \times r}$, an initialization $H \in \mathbb{R}_+^{m \times r}$, a finite set Ω of nonnegative real numbers, and $\lambda \in \mathbb{R}_+^{|\Omega|}$.

Output: $H_\gamma^+ \in \mathbb{R}_+^{m \times r}$ such that $\bar{D}_\Omega^\lambda(X, WH_\gamma^+) \leq \bar{D}_\Omega^\lambda(X, WH)$.

- 1: $H^+ = H \circ \left[\frac{\sum_{\beta \in \Omega} \lambda_\beta (\nabla_-^H \bar{D}_\beta(X, WH))}{\sum_{\beta \in \Omega} \lambda_\beta (\nabla_+^H \bar{D}_\beta(X, WH))} \right]$.
 - 2: $\gamma = 1, H_1^+ = H^+$.
 - 3: **while** $\bar{D}_\Omega^\lambda(X, WH_\gamma^+) > \bar{D}_\Omega^\lambda(X, WH)$ **do**
 - 4: $\gamma = \frac{\gamma}{2}$.
 - 5: $H_\gamma^+ = (1 - \gamma)H + \gamma H^+$.
 - 6: **end while**
-

Because of the step length procedure that guarantees the objective function to not increase (steps 3-6), the use of Algorithm 1 in an alternating scheme to solve (1) by updating W and H alternatively is guaranteed to not increase the objective function. Since the objective function is bounded below, this guarantees that the objective function values converge as k goes to infinity.

4 ALGORITHM FOR DR-NMF

As $(W, H) \mapsto \max_{\beta \in \Omega} \bar{D}_\beta(X, WH)$ is a non-convex function, obtaining a global solution (W^*, H^*) for (2) efficiently is not possible in general. In particular, deciding whether the minimum in (2) is equal to zero (that is, deciding whether there exists W and H such that $X = WH$) is NP-hard [20]. In the following, we propose to find an approximate solution for the DR-NMF problem via a weighted sum of the different objective functions. We first observe that

$$\max_{\beta \in \Omega} \bar{D}_\beta(X, WH) = \max_{\lambda \geq 0, \sum_{\beta \in \Omega} \lambda_\beta = 1} \sum_{\beta \in \Omega} \lambda_\beta \bar{D}_\beta(X, WH).$$

Hence (2) can be reformulated as

$$\min_{(W, H) \geq 0} \max_{\lambda \geq 0, \|\lambda\|_1 = 1} \sum_{\beta \in \Omega} \lambda_\beta \bar{D}_\beta(X, WH). \quad (5)$$

4.1 Related works on min-max problems

The problem in (5) is a min-max problem. Let us present a brief review of well-known methods for solving a general min-max problem, also known as a *saddle point problem* (SSP), of the form

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \Phi(x, y), \quad (6)$$

where \mathcal{X} and \mathcal{Y} are closed convex sets. SPPs are abundant in game theory, machine learning and statistics. A special class of SPPs is the class of bilinear SPPs which assume that the objective can be expressed as $\Phi(x, y) = f(x) + \langle Ax, y \rangle + g(y)$ where A is a linear operator, f and g are differentiable functions, and the coupling between x and y is linear in x and linear in y . Bilinear SPPs have been extensively studied and can be solved efficiently by several methods such as Nesterov's smoothing method [28] and the primal-dual hybrid gradient method [29], [30]. For non-bilinear SPPs, the proximal mirror descent method (which subsumes the proximal gradient descent method as a special case [31], [32], [33]), is often the method of choice since it is a direct method applied to the underlying SPP (while [28] requires advanced smoothing techniques) and it can be adapted to the case in which regularizers or constraints are present, assuming the involving proximal maps can be computed. In another line of works, the approach of sequentially solving auxiliary sub-problems (which can have closed-form solutions or can be approximately solved by suitable solvers) to alternatively update x (while fixing y) and y (while fixing x) has also been developed in the literature [34], [35], [36].

To establish convergence guarantees of algorithms for SPPs, the following three typical assumptions are made in the literature: (A) Φ is convex in x and concave in y , (B) the gradient map $(x, y) \mapsto [\nabla_x \Phi(x, y), -\nabla_y \Phi(x, y)]$ is Lipschitz continuous, and (C) the SPP has at least one saddle-point, that is, there exists (x^*, y^*) such that

$$\Phi(x^*, y) \leq \Phi(x^*, y^*) \leq \Phi(x, y^*) \quad \text{for all } x \in \mathcal{X} \text{ and } y \in \mathcal{Y}. \quad (7)$$

Assumption (C) is satisfied when \mathcal{X} and \mathcal{Y} are convex compact sets and Φ is a continuous convex-concave function [37, Proposition 5.5.3], hence Assumption (C) can be omitted for SPPs under these settings [31], [32]. For SPPs under other settings (for example when \mathcal{X} or \mathcal{Y} is unbounded, or when $\Phi(x, y)$ is not convex-concave), Assumption (C) concerning the existence of saddle points is a standard assumption for the development of numerical algorithms for solving SPPs; see [30], [33], [34] and references therein. Our SPP (5) neither satisfies Assumption (A) nor Assumption (B) since $(W, H) \mapsto \bar{D}_\beta(X, WH)$ is not convex and $(W, H) \mapsto \nabla \bar{D}_\beta(X, WH)$ is not Lipschitz continuous. These facts prevents us from applying standard SPP algorithms with convergence guarantees to solve (5).

4.2 Preliminaries

The following proposition provides some properties of saddle points of Φ . Its proof can be derived from [37, Proposition 3.4.1] and the definition of subgradients.

Proposition 1. Consider the SPP in (6). Suppose that for each $(x, y) \in \mathcal{X} \times \mathcal{Y}$, $\Phi(\cdot, y)$ and $-\Phi(x, \cdot)$ are subdifferentiable on \mathcal{X} and \mathcal{Y} respectively. Then,

(I) (x^*, y^*) is a saddle point of (6) if and only if there exist a subgradient $\Phi'_x(x^*, y^*)$ of $\Phi(\cdot, y^*)$ at x^* and a subgradient $-\Phi'_y(x^*, y^*)$ of $-\Phi(x^*, \cdot)$ at y^* such that

$$x^* = \Pi_{\mathcal{X}}(x^* - \Phi'_x(x^*, y^*)) \text{ and } y^* = \Pi_{\mathcal{Y}}(y^* + \Phi'_y(x^*, y^*)).$$

(II) (x^*, y^*) is a saddle point of (6) if and only if strong duality holds, that is,

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \Phi(x, y) = \max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} \Phi(x, y),$$

and

$$x^* \in \arg \min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \Phi(x, y), \quad y^* \in \arg \max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} \Phi(x, y).$$

Suppose Φ has a saddle point. Proposition 1 shows that if we can find a solution of the dual problem $y^* \in \arg \max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} \Phi(x, y)$ and a solution x^* of $\min_{x \in \mathcal{X}} \Phi(x, y^*)$, that is, $x^* = \Pi_{\mathcal{X}}(x^* - \Phi'_x(x^*, y^*))$, such that the equation $y^* = \Pi_{\mathcal{Y}}(y^* + \Phi'_y(x^*, y^*))$ also holds, then (x^*, y^*) is a saddle point of Φ , which then leads to the fact that x^* is a solution of the primal problem $\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \Phi(x, y)$. This motivates us to use a dual subgradient method that solves the dual problem of (5), which is the maximization problem of a concave function. This is described in the next subsection.

4.3 A dual subgradient method

Define the functions $L(W, H; \lambda) = \bar{D}_{\Omega}^{\lambda}(X, WH)$, and $g(\lambda) = \min_{(W, H) \geq 0} L(W, H; \lambda)$, and the set

$$\Lambda = \{\lambda : \lambda \geq 0, \|\lambda\|_1 = 1\}.$$

It then follows from Danskin's theorem [37, Proposition B.25] that the vector $-g'(\lambda)$ with

$$g'(\lambda) = [\bar{D}_{\beta}(X, W_{\lambda} H_{\lambda})]_{\beta \in \Omega}, \quad (8)$$

where $(W_{\lambda}, H_{\lambda}) \in \arg \min_{W \geq 0, H \geq 0} \bar{D}_{\Omega}^{\lambda}(X, WH)$, is a subgradient of $-g$ at λ . We now solve the dual problem

$$\max_{\lambda \geq 0, \|\lambda\|_1 = 1} g(\lambda), \quad (9)$$

to obtain an optimal solution λ^* . We observe that $g(\lambda)$ is concave and, as such, a subgradient method with a suitable choice of step sizes guarantees the convergence to a global optimal solution of the concave maximization problem in (9); see, for example, [38].

Algorithm 2 describes a dual subgradient method. It is worth noting that Problem (9) can also be solved by a mirror descent method; see for example [39, Chapter 4]. Note that we take Π_{Λ} in Step 4 of Algorithm 2 to be the Euclidean projection operator. This allows us to establish a convergence guarantee for the sequence of dual parameters $\{\lambda^{(k)}\}_{k \in \mathbb{N}}$ generated in Step 4 of Algorithm 2 by applying [38, Theorem 3]. Specifically, suppose the step sizes satisfy $\rho_k \rightarrow 0$, $\sum_{k=1}^{\infty} \rho_k = +\infty$ and $\sum_{k=1}^{\infty} \rho_k^2 < \infty$. Then the sequence $\{\lambda^{(k)}\}_{k \in \mathbb{N}}$ generated by Algorithm 2 converges to a solution λ^* of (9). Furthermore, suppose (W^*, H^*) is a limit point of $(W^{(k)}, H^{(k)})$ and assume that $(W, H, \lambda) \mapsto L(W, H; \lambda)$ is lower semicontinuous at (W^*, H^*, λ^*) . Then

Algorithm 2 A dual subgradient method

Input: The matrices $X \in \mathbb{R}_+^{m \times n}$, a finite set Ω of nonnegative real numbers.

Output: (W, H) , an approximate solution to DR-NMF (2).

1: Initialize $\lambda_{\beta}^{(1)} = \frac{1}{|\Omega|}$ for all $\beta \in \Omega$.

2: **for** $k = 1, 2, \dots$ **do**

3: Compute $(W^{(k)}, H^{(k)})$ as an optimal solution to

$$\min_{(W, H) \geq 0} \bar{D}_{\Omega}^{\lambda^{(k)}}(X, WH). \quad (10)$$

4: Update

$$\lambda^{(k+1)} = \Pi_{\Lambda}(\lambda^{(k)} + \rho_k g'(\lambda^{(k)})),$$

where $g'(\lambda^{(k)})$ is computed as in (8).

5: **end for**

we have $(W^*, H^*) \in \arg \min_{(W, H) \geq 0} \bar{D}_{\Omega}^{\lambda^*}(X, WH)$. Indeed, let $\{k_n\}_{n \in \mathbb{N}}$ be such that $(W^{(k_n)}, H^{(k_n)}) \rightarrow (W^*, H^*)$ as $n \rightarrow \infty$. Step 3 of Algorithm 2 implies that

$$L(W^{(k_n)}, H^{(k_n)}; \lambda^{(k_n)}) \leq L(W, H; \lambda^{(k_n)}), \quad \forall (W, H) \geq 0.$$

Taking $n \rightarrow \infty$, we obtain

$$\begin{aligned} L(W^*, H^*; \lambda^*) &\leq \liminf_{n \rightarrow \infty} L(W^{(k_n)}, H^{(k_n)}; \lambda^{(k_n)}) \\ &\leq L(W, H; \lambda^*) \text{ for all } (W, H) \geq 0. \end{aligned}$$

Hence $(W^*, H^*) \in \arg \min_{(W, H) \geq 0} \bar{D}_{\Omega}^{\lambda^*}(X, WH)$.

We have shown that the iterates generated by Algorithm 2 converge to a solution λ^* of the dual problem and to a solution (W^*, H^*) of $\min_{(W, H) \geq 0} L(W, H; \lambda^*)$. After obtaining λ^* , the discussion after Proposition 1 indicates that if we assume that L has a saddle point (which is a common assumption as mentioned in Section 4.1) and that we can find a solution (W^*, H^*) of $\min_{(W, H) \geq 0} L(W, H; \lambda^*)$ that satisfies the condition $\lambda^* = \Pi_{\Lambda}(\lambda^* + L'_{\lambda}(W^*, H^*; \lambda^*))$, then we can recover a solution (W^*, H^*) of the primal problem (5). Hence we can regard the output $(W^{(k)}, H^{(k)})$ of Algorithm 2 as an approximate solution of (5). In practice, we can run Algorithm 2 until we observe that the change between two consecutive iterates is negligible; for example stop the algorithm when $\|\lambda^{(k+1)} - \lambda^{(k)}\|_1 \leq \varepsilon$, for a predetermined tolerance $\varepsilon > 0$. Since $\|\lambda^{(k)}\|_1 = 1$ for all k , choosing for example $\varepsilon = 0.001$ means that we stop the algorithm when $\lambda^{(k)}$ is modified by less than 0.1% (compared to the previous iterate).

The performance of Algorithm 2 critically depends on the solver for solving the weighted-sum minimization in Step 3, which itself is a difficult non-convex optimization problem. We can use Algorithm 1 to find an approximate solution $(W^{(k)}, H^{(k)})$ in Step 3. However, subgradient methods are often slow in practice. Indeed, we observe that Algorithm 2 combined with Algorithm 1 is very slow for the data sets we use in our experiments (see Section 4.5 and Figure 1). Therefore, although Algorithm 2 provides some convergence guarantees for the primal and dual problems, we are motivated to propose another practical approach for finding an approximate solution to (2). In the following, we present a heuristic scheme that performs very well,

significantly better than Algorithm 2 where Step 3 is approximately solved via Algorithm 1.

4.4 A Frank-Wolfe heuristic scheme for DR-NMF

Unfortunately, Algorithm 2 is not practical because Step 3 requires one to solve (10) which is NP-hard in general (since it is a generalization of NMF). Hence we instead propose a heuristic scheme described in Algorithm 3.

Algorithm 3 Heuristic Frank-Wolfe scheme for DR-NMF (2)

Input: The matrices $X \in \mathbb{R}_+^{m \times n}$ and an initialization $W^{(0)} \in \mathbb{R}_+^{m \times r}$ and $H^{(0)} \in \mathbb{R}_+^{n \times r}$, a finite set Ω of nonnegative real numbers.

Output: (W, H) , an approximate solution to DR-NMF (2).

- 1: Initialize $\lambda_\beta = \frac{1}{|\Omega|}$ for all $\beta \in \Omega$.
- 2: **for** $k = 1, 2, \dots$ **do**
- 3: Update $H^{(k+1)}$ using Algorithm 1 with initialization $H = H^{(k)}$, and with $W = W^{(k)}$ and $\lambda = \lambda^{(k)}$.
- 4: Update $W^{(k+1)}$ using Algorithm 1 and the symmetry of the NMF problem, with initialization $W = W^{(k)}$, and with $H = H^{(k+1)}$ and $\lambda = \lambda^{(k)}$.
- 5: Let $\beta^* \in \arg \max_{\beta \in \Omega} \bar{D}_\beta(X, W^{(k+1)} H^{(k+1)})$ and

$$(\lambda_*^{(k)})_\beta = \begin{cases} 1 & \text{if } \beta = \beta^*, \\ 0 & \text{if } \beta \neq \beta^*. \end{cases} \quad (11)$$

- 6: Choose a step size γ_k , and update

$$\lambda^{(k+1)} = \lambda^{(k)} + \gamma_k (\lambda_*^{(k)} - \lambda^{(k)}).$$

- 7: **end for**

- 8: $(W, H) = (W^{(k+1)}, H^{(k+1)})$.
-

Let us explain the main ideas behind Algorithm 3. Steps 3 and 4 are designed to decrease $(W, H) \mapsto L(W, H; \lambda^{(k)})$. Note that if we perform Steps 3 and 4 repeatedly, its output will approximate the output of Step 3 of Algorithm 2. However, this would be computationally rather expensive as it would require many updates of (W, H) for each $\lambda^{(k)}$, and the MU constitutes the most expensive steps of Algorithm 3.

The descent direction $\lambda_*^{(k)}$ used to update λ at iteration k is the one from the Frank-Wolfe (FW) algorithm [40], which is also known as the conditional gradient method; see for example [41] and the references therein. In the context of solving DR-NMF (2), let us explain the intuition behind the descent direction $\lambda_*^{(k)}$, defined in Step 5 of Algorithm 3. Letting $\beta^* \in \arg \max_{\beta \in \Omega} \bar{D}_\beta(X, W^{(k+1)} H^{(k+1)})$, we have for all $\beta \in \Omega$ that

$$\bar{D}_{\beta^*}(X, W^{(k+1)} H^{(k+1)}) \geq \bar{D}_\beta(X, W^{(k+1)} H^{(k+1)}).$$

Defining $\lambda_*^{(k)}$ as the vector with a single non-zero entry equal to one at position β^* , see (11), we have $\lambda_*^{(k)} = \arg \max_{\lambda \in \Lambda} \bar{D}_\beta^\lambda(X, W^{(k+1)} H^{(k+1)})$. Therefore, since we are trying to solve the optimization problem $\min_{(W, H) \geq 0} \max_{\beta \in \Omega} \bar{D}_\beta(X, WH)$, the β^* -divergence should be given more importance at the next iteration in order for $\bar{D}_{\beta^*}(X, WH)$ to decrease, and hence the maximum among the β -divergences to decrease as well at the

next iteration. Finally, Algorithm 3 updates $\lambda^{(k)}$ in Step 6 as follows

$$\begin{aligned} \lambda^{(k+1)} &= \lambda^{(k)} + \gamma_k (\lambda_*^{(k)} - \lambda^{(k)}), \\ &= (1 - \gamma_k) \lambda^{(k)} + \gamma_k \lambda_*^{(k)}. \end{aligned} \quad (12)$$

In our experiments, we choose the step sizes $\gamma_k = \frac{1}{k+1}$. We leave the fine-tuning of the step sizes as a future direction of research, although we have tried different step sizes, and we were not able to find step sizes that perform significantly better than $\gamma_k = \frac{1}{k+1}$. In fact, the performance for γ_k of similar order is very similar. For example, the standard FW parameter choice of $\gamma_k = \frac{2}{k+2}$ which yields essentially the same but slightly worse performance.

We emphasize that Algorithm 3 is a heuristic algorithm, and we leave convergence guarantees as a research topic for future work. We use Algorithm 3 in our experiments and observe that it performs very well in terms of simultaneously decreasing all β -divergences for $\beta \in \Omega$. Thus, we do not need prior knowledge on the noise distribution or, equivalently, on the value of β .

4.5 Comparison between the dual subgradient method and the heuristic scheme for DR-NMF

The update in Step 5 of Algorithm 3 results in faster convergence compared to the subgradient direction (see Step 4 of Algorithm 2) because it gives much more importance to the β -divergence that is maximal at the current iteration. In fact, the entries of the subgradient are always all positive (unless $X = WH$ in which case the problem is solved). Thus, the direction that places weight *only* at the current maximum β -divergence (as is done in Step 5 of Algorithm 3) outperforms the subgradient direction empirically, leading to much faster convergence in practical problems.

Let us compare the dual subgradient method and the heuristic scheme for DR-NMF on a simple synthetic data set. However, we have made the same observations on all other data sets we have experimented with such as the ones presented in Section 5. Figure 1 illustrates the distinction between the two algorithms with a synthetic experiment that compares Algorithm 3 with the variant in which Step 5 is replaced with a standard subgradient step, that is, Step 4 of Algorithm 2. In this illustrative experiment, the entries of a 100-by-100 matrix X are generated uniformly at random in the interval $[0, 1]$, and we use $r = 10$ and $\Omega = \{0, 1, 2\}$, that is, DR-NMF with the IS-divergence, the KL-divergence and the Frobenius norm. Both variants are initialized with the same matrices $(W^{(0)}, H^{(0)})$ whose entries are also generated uniformly at random in $[0, 1]$.

We observe that the variant using the subgradient converges very slowly. Indeed, the maximum of the three objectives (the IS-divergence) is far from convergence, even after 1000 iterations.¹ In contrast, Algorithm 3 converges much faster. In particular, the values of the IS-divergence and the Frobenius norm quickly converge to one another. All in all, Algorithm 3 finds a solution with scaled β -divergence within 2% of the smallest possible values

1. It required 3000 iterations to make the IS-divergence and the Frobenius norm intersect, but then the Frobenius norm becomes larger and, within the next 7000 iterations (for a total of 10000 iterations), the IS divergence remains larger hence convergence is not attained.




Fig. 1: Comparison of Algorithm 3 with its variant where Step 5 is replaced by a standard subgradient step (i.e., Step 4 of Algorithm 2). The curves display the evolution of the scaled β -divergences (referred to as “error” in the legend).

for the three β -divergences within 240 iterations, that is, $\max_{\beta \in \Omega} \bar{D}_{\beta}(X, W^{(k)}H^{(k)}) \leq 1.02$ for $k \geq 240$. This is made possible because of our more aggressive heuristic strategy to update λ . In contrast, if one uses the subgradient direction, about 800 iterations are required to obtain the same approximation guarantee.

Remark 1 (Property of DR-NMF solutions). *We observe on Figure 1 that the two scaled β -divergences with the largest values are equal to each another. The same behaviour will often be observed in the extensive sets of experiments in Section 5. Let us explain why this is expected to happen. First, recall that the lowest possible value of a single scaled β -divergence is one; see Section 2.1. Second, the maximum scaled β -divergence is attained at a single scaled β -divergence when it is strictly larger than all the other scaled β -divergences. In that case, the maximum scaled β -divergence will be larger than one, and hence it can typically² be reduced locally while ensuring that the other β -divergences remain smaller, hence reducing the maximum scaled β -divergence.*

5 NUMERICAL EXPERIMENTS

In this section, we apply DR-NMF on several data sets. In all cases, we perform 1000 iterations. All tests are preformed using Matlab R2015a on a laptop Intel CORE i7-7500U CPU @2.9GHz 24GB RAM. The code is available on Code Ocean via <https://doi.org/10.24433/CO.7769595.v1>.

5.1 MO-NMF: Examples of the Pareto frontier on synthetic data

In this section, we illustrate the use of Algorithm 1 to compute Pareto-optimal solutions. We will focus on the case $\beta = 0, 1, 2$, that is, IS- and KL-divergences and the Frobenius

norm. Note, however, that our algorithm and code can deal with any $\beta \geq 0$ and any finite set Ω .

We generate the input matrix X as follows: $X = \max(0, \tilde{W}\tilde{H} + N)$ where the component matrices \tilde{W} , \tilde{H} and the noise matrix N are generated as follows:

- The entries of $\tilde{W} \in \mathbb{R}^{200 \times 10}$ and $\tilde{H} \in \mathbb{R}^{10 \times 200}$ are generated using the uniform distribution in the interval $[0,1]$. We define $\tilde{X} = \tilde{W}\tilde{H}$ which is the noiseless low-rank matrix.
- Let us define $x_{\beta} = 1$ if $\beta \in \Omega$, $x_{\beta} = 0$ otherwise. Let also

$$\tilde{N} = x_0 \frac{N_{\text{IS}}}{\|N_{\text{IS}}\|_F} + x_1 \frac{N_{\text{KL}}}{\|N_{\text{KL}}\|_F} + x_2 \frac{N_{\text{F}}}{\|N_{\text{F}}\|_F},$$

where

- $N_{\text{IS}} = \tilde{X} \circ G$ is multiplicative Gamma noise where each entry of G is generated using the normal distribution of mean 0 and variance 1,
- each entry of N_{KL} is generated according to the Poisson distribution of parameter 1 (for simplicity, since the expected value of $(\tilde{W}\tilde{H})_{i,j}$ is the same for all (i, j)),
- each entry of N_{F} is generated using the normal distribution of mean 0 and variance 1.

We set $N = \epsilon \frac{\|\tilde{X}\|_F}{\|\tilde{N}\|_F} \tilde{N}$ with $\epsilon = 0.2$.

Finally, $X = \max(0, \tilde{X} + N)$ is a low-rank matrix to which had been contaminated with 20% of noise (that is, $\|N\|_F = 0.2\|\tilde{X}\|_F$) and then was projected onto the nonnegative orthant. The noise is constructed using the distributions corresponding to $\beta \in \Omega$.

Figure 2 shows the Pareto-optimal solutions for MO-NMF. More precisely, it provides the solution for the problems

$$\min_{(W,H) \geq 0} \bar{D}_{\Omega}^{\lambda}(X, WH),$$

where $\lambda = (\ell, 1 - \ell)$ for $\ell = 0, 0.1, \dots, 1$, and for $\Omega = \{0, 1\}, \{0, 2\}, \{1, 2\}$. To simplify computation, we have used the true underlying solution (W_t, H_t) as the initialization (using random or other initializations sometimes generate solution which are more often not on the Pareto frontier because NMF may have many local minima). The Pareto frontier is as expected: the smallest possible value for each objective is 1 (because of the scaling), for which the other objective function is the largest. As λ changes, one objective increases while the other decreases. The DR-NMF solution computed with Algorithm 3 finds the point on the Pareto frontier such that $\bar{D}_{\beta_2}(X, WH) = \bar{D}_{\beta_1}(X, WH)$ for $\beta_1 \neq \beta_2 \in \Omega$.

For DR-NMF, we observe that

- The solution of DR-NMF does not necessarily coincide with a value of λ close to $(0.5, 0.5)$. For example, for the case of the IS divergence with the Frobenius norm, it is close to $\lambda = (0.9, 0.1)$.
- Using DR-NMF allows to obtain a solution with low error for both objectives, always at most 2% worse than the lowest error. Minimizing a single objective sometimes leads to solution with error up to 35% higher than the lowest (in the case IS divergence with Frobenius norm). We will observe a similar behaviour on real data sets.

². Because of the non-convexity of the objectives of DR-NMF, such a descent direction is not guaranteed to exist.




Fig. 2: Pareto-optimal solutions for $\lambda = (\ell, 1 - \ell)$ for $\ell = 0, 0.1, \dots, 1$ and $\Omega = \{0, 1\}, \{0, 2\}, \{1, 2\}$ (top to bottom), and solution computed by DR-NMF (Algorithm 3).

5.2 Sparse document data sets: $\Omega = \{1, 2\}$

For sparse data sets, it is known that only the β -divergence for $\beta = 1, 2$ can exploit the sparsity structure. In fact, in all other cases, all entries of the product WH have to be computed explicitly which is impractical for large sparse matrices since WH can be dense. In other words, let K denote the number of non-zero entries of X . Then the MU for NMF with the β -divergence for $\beta = 1, 2$ can be run in $O(Kr)$ operations, while for the other values of β , it requires $O(mnr)$ operations.

As explained in [42], for sparse word-count matrices, Poisson noise is the most appropriate model; in fact, Gaussian noise (and any dense noise) does not make much sense on sparse data sets. Hence we expect KL-NMF to provide better results than Fro-NMF. However, we believe it is rather interesting to run DR-NMF with $\Omega = \{1, 2\}$ on such data sets to see how it performs. One should expect DR-NMF to perform on average worse than KL-NMF (since

it has to take into account the Frobenius norm which is not appropriate) but better than the Frobenius norm (since it takes into account the more appropriate KL-NMF).

In this section, we use the 15 sparse document data sets from [43]. These are large and highly sparse matrices whose entries $X(i, j)$ is the number of times word j appears in document i . We apply KL-NMF, Fro-NMF and DR-NMF with $\Omega = \{1, 2\}$. To simplify the comparison, reduce the computational load and to have a good initial solution, we use the same initial matrices $(W^{(0)}, H^{(0)})$ in all cases, namely the solution obtained by the successive projection algorithm [44] that has provable guarantee under the separability condition [45], [46]. We perform rank- r factorization where r is the number of classes reported for these data sets.

Table 1 reports the results. The first and second columns report the name of the data set and the number of classes, respectively. The next four columns report the accuracies of the clustering obtained with the factorizations (W, H) produced by KL-NMF, Fro-NMF, and DR-NMF with $\Omega = \{1, 2\}$ solved via Algorithm 3. Given the true disjoint clusters $C_i \subset \{1, 2, \dots, m\}$ for $1 \leq i \leq r$ and given a computed disjoint clustering $\{\tilde{C}_i\}_{i=1}^r$, its accuracy is defined as

$$\text{accuracy}(\{\tilde{C}_i\}_{i=1}^r) = \min_{\pi \in [1, 2, \dots, r]} \frac{1}{m} \sum_{i=1}^r |C_i \cap \tilde{C}_{\pi(i)}|,$$

where $[1, 2, \dots, r]$ is the set of permutations of $\{1, 2, \dots, r\}$. For simplicity, given an NMF (W, H) where each row of H corresponds to a topic, we cluster the documents by selecting its closest topic, that is, document j is assigned to the topic k that maximizes $\frac{X(i, j) H(k, j)}{\|H(k, :)\|_2}$. The next three columns report how much higher the KL error (in percent) of the solutions of Fro-NMF and DR-NMF are compared to KL-NMF, that is, it reports

$$\bar{D}_1(X, WH) - 1 = \frac{\text{KL}(X, WH)}{\text{KL}(X, W_1 H_1)} - 1$$

where (W_1, H_1) is the solution computed by KL-NMF. The last three columns report how much higher the Frobenius error (in percent) of the solutions of KL-NMF and DR-NMF are compared to Fro-NMF, that is,

$$\bar{D}_2(X, WH) - 1 = \frac{\|X - WH\|_F^2}{\|X - W_2 H_2\|_F^2} - 1$$

where (W_2, H_2) is the solution computed by Fro-NMF.

We observe the following:

- In terms of clustering, DR-NMF in fact allows us to be robust in the sense that it is able to provide *in all cases* at least the second highest clustering accuracy. On four data sets, it is even able to provide the highest accuracy, sometimes by a large margin. Globally, DR-NMF does not perform as well as KL-NMF although on average their accuracy only differs by 3.32%. However, DR-NMF performs better than Fro-NMF, with 6.44% higher accuracy on average.
- In terms of error, as already noted in the previous section, DR-NMF is able to simultaneously provide solutions with small KL and Frobenius error, on average 2.65% higher than the solution computed with

TABLE 1: Comparison of NMF with KL-divergence and Frobenius norm, and DR-NMF with $\Omega = \{1, 2\}$ on text mining data sets from [43]. Bold numbers indicate the best accuracy, underlined numbers indicate the second best accuracy.

Data set	r	Clustering accuracy (%)			$D_1(X, WH) - 1$ (%)		$D_2(X, WH) - 1$ (%)	
		KL-NMF	Fro-NMF	DR-NMF	Fro-NMF	DR-NMF	KL-NMF	DR-NMF
NG20	20	42.15	23.08	<u>28.74</u>	21.47	3.85	149.48	3.83
ng3sim	3	63.48	38.06	<u>49.87</u>	16.82	2.70	17.32	2.70
classic	4	83.66	55.64	<u>78.46</u>	13.19	0.74	2.44	0.74
ohscal	10	37.45	30.50	<u>32.13</u>	10.03	1.76	9.60	1.75
k1b	6	64.27	59.19	<u>60.30</u>	9.00	1.32	5.02	1.32
hitech	6	43.29	<u>46.94</u>	48.02	8.27	1.12	3.98	1.13
reviews	5	75.65	<u>51.19</u>	<u>74.88</u>	7.89	1.03	7.70	1.03
sports	7	<u>43.93</u>	40.37	50.26	9.60	1.24	7.10	1.24
la1	6	<u>65.95</u>	65.04	67.98	9.26	1.03	3.61	1.03
la12	6	56.25	54.80	<u>54.29</u>	7.32	0.70	2.76	0.70
la2	6	54.96	49.17	<u>52.07</u>	9.21	0.82	3.04	0.82
tr11	9	62.32	50.48	<u>51.45</u>	22.88	4.48	97.27	4.47
tr23	6	34.80	<u>35.29</u>	38.73	56.04	3.83	47.36	3.78
tr41	10	54.33	44.99	<u>53.08</u>	24.38	4.93	46.17	4.90
tr45	10	46.81	38.26	<u>39.13</u>	42.52	10.15	50.14	10.15
Average		55.29	45.53	51.96	17.86	2.65	30.20	2.64

a single objective. On the other hand, optimizing a single objective often leads to very large errors for the other one, up to 149% on NG20, with an average of 17.86% for Fro-NMF and 30.20% for KL-NMF.

5.3 Dense time-frequency matrices of audio signals: $\Omega = \{0, 1\}$

NMF has been used successfully to separate sources from a single audio recording. However, there is a debate in the literature as to whether the KL or the IS divergence should be used; see [47], [48] and the references therein. In fact, as we will see, IS-NMF and KL-NMF provide rather different results on different audio data sets. On one hand, due to its insensitivity to scaling (see Section 2.1), IS-NMF gives the same relative importance to all entries of the data matrix. For example, the error for approximating 1 by 10 is the same as for approximating 10 by 100, that is, $D_0(1, 10) = D_0(10, 100)$. On the other hand, KL-NMF gives more importance to larger entries as it is (linearly) sensitive to scaling; for example, the error for approximating 1 by 10 is ten times smaller than approximating 10 by 100, that is, $10D_1(1, 10) = D_1(10, 100)$.

5.3.1 Quantitative results

Our DR-NMF approach overcomes the issue of having to choose between the IS- and KL-divergences by generating solutions which possess small IS and KL errors simultaneously. We use 10 diverse audio data sets:

- voice_cell, syntBassDrum and syntCCcyGC were downloaded from <http://isse.sourceforge.net/demos.html>.
- prelude_JSB is the the well-tempered Clavier performed by Glenn Gould 1/13 between the 19th et 49th seconds, downloaded from https://www.youtube.com/watch?v=IrJjPYi_vhM.
- ShanHur_sunrise was downloaded from <http://bass-db.gforge.inria.fr/fasst/>.
- trio_Brahms and trio_bapitru were derived from the TRIOS data set [49]; see <https://c4dm.eecs.qmul.ac.uk/rdr/handle/123456789/27>.

- sisec_mixdrums and sisec_mixfemale come from the SISEC data set; see <http://sisec.wiki.irisa.fr/tiki-indexbfd7.html?page=Underdetermined+speech+and+music+mixtures>.
- piano_Mary is a recording at the third author's house.

Table 2 reports the results, exactly as done in the last columns of Table 1, except that it also reports the standard deviation among 10 random initializations.

For these data sets, the results are even more striking than for the sparse text data sets in Section 5.2. In particular, DR-NMF has on average an error higher by about 10% compared to both IS-NMF and KL-NMF, while KL-NMF (resp. IS-NMF) has on average an increase in IS error of 149% (resp. 179%). Moreover, DR-NMF is more robust in the sense that its standard deviation is significantly lower. This shows that by taking into account different objectives, DR-NMF is less sensitive to initialization.

As we will see in the next section, using DR-NMF allows to obtain more robust results than using IS-NMF or KL-NMF alone.


5.3.2 Qualitative results

In the previous section, we have shown quantitative results showing that DR-NMF is able to obtain solutions with low KL and IS divergence simultaneously. In this section, we investigate the data set piano_Mary in more detail and show that DR-NMF also leads to better separation for three comparative studies described in detail below: (1) no noise added to the signal, (2) Poisson noise added and (3) Gamma noise added. This data set is the first 4.7 seconds of "Mary had a little lamb". The sequence is composed of three notes, namely, E_4 , D_4 and C_4 . The recorded signal is downsampled to $f_s = 16000\text{Hz}$ yielding $T = 75200$ samples. The short-time Fourier transform (STFT) of the input signal x is computed using a Hamming window of size $F = 512$ leading to a temporal resolution of 32ms and a frequency resolution of 31.25Hz. We use 50% overlap between two frames, leading to $n = 294$ frames and $m = 257$ frequency bins. Figure 3 displays the musical score.

There are three notes plus a fourth source. This last source is the very first offset of each note in the musical

TABLE 2: Comparison of NMF with the IS- and KL-divergences, and DR-NMF with $\Omega = \{0, 1\}$ on audio data sets with $m = 149$ and $r = 10$. The table reports the averages and standard deviations over 10 initializations.

Data set	n	$D_0(X, WH) - 1$ (%)		$D_1(X, WH) - 1$ (%)	
		KL-NMF	DR-NMF	IS-NMF	DR-NMF
syntBassDrum	543	39.54 ± 3.28	7.06 ± 3.01	108.09 ± 16.78	7.06 ± 3.01
piano_Mary	586	387.51 ± 253.67	9.73 ± 2.74	177.71 ± 22.79	9.73 ± 2.74
prelude_JSB	2582	31.81 ± 3.57	13.05 ± 3.55	185.93 ± 54.84	13.04 ± 3.55
syntCCcyGC	1377	9.79 ± 0.64	2.63 ± 0.38	42.21 ± 10.29	2.63 ± 0.38
trio_Brahms	14813	360.49 ± 44.65	14.74 ± 1.99	257.61 ± 130.29	14.74 ± 1.99
trio_bapitru	6200	354.66 ± 25.31	9.16 ± 2.03	249.99 ± 28.72	9.15 ± 2.03
voice_cell	2181	186.46 ± 2.20	13.98 ± 3.75	191.12 ± 23.43	13.98 ± 3.75
ShanHur_sunrise	4102	53.51 ± 8.32	12.31 ± 1.36	184.72 ± 34.74	12.30 ± 1.35
sisec_mixdrums	1249	25.61 ± 0.87	12.74 ± 1.38	292.40 ± 56.62	12.74 ± 1.38
sisec_mixfemale	1249	37.68 ± 2.88	12.12 ± 0.85	100.67 ± 8.88	12.12 ± 0.85
Average		148.71 ± 34.54	10.75 ± 2.10	179.05 ± 38.74	10.75 ± 2.10


 Fig. 3: Musical score of “Mary had a little lamb”. The notes activate as follows: $E_4, D_4, C_4, D_4, E_4, E_4, E_4$.


sequence, that is, some common mechanical vibration acting in the piano just before triggering a specific note, which can be associated to the hammer noise (denoted H_N), hence the correct rank is $r = 4$; see [50] for more details.

5.3.2.1 No added noise: Figure 4 displays the evolution of the scaled IS- and KL-divergences along iterations. DR-NMF is able to compute a solution with low IS and KL error, which is not the case of IS-NMF and DR-NMF (in particular, KL-NMF has IS error almost 9 times larger than IS-NMF). However, the three solutions generated by IS-NMF, KL-NMF and DR-NMF all give a correct separation. The reason is that this recording is of good quality hence the noise is rather low.

5.3.2.2 Poisson noise: The second comparative study is performed on the same data set with Poisson noise added to the input audio spectrogram following the methodology described in Section 5.1. We use $N = \epsilon \frac{\|X\|_F}{\|\tilde{N}\|_F} \tilde{N}$ with $\epsilon = 0.6$ and $\tilde{N} = \frac{N_{KL}}{\|N_{KL}\|_F}$. Figure 5 displays the rows of H (that is, the activations of the notes over time) for NMF with IS- and KL-divergences, and for DR-NMF with $\Omega = \{0, 1\}$ with $r = 4$.

As expected with this noise model and high noise level, IS-NMF is not able to extract the three notes, while KL-NMF and DR-NMF identify them. In fact, the recovered activations, that is, the rows of H , correspond to the activations of the notes from the musical score shown on Figure 3: C_4 is activated once, D_4 twice and E_4 four times. Note that the hammer noise (H_N) is not extracted (a source is set to zero) but is mixed with C_4 and to a smaller extent with D_4 . This illustrates that DR-NMF is robust to different types of noises (in this case, Poisson noise).

5.3.2.3 Gamma noise: The third comparative study is performed on the same data set with multiplicative Gamma noise, accordingly to the methodology described in Section 5.1. We use $N = \epsilon \frac{\|X\|_F}{\|\tilde{N}\|_F} \tilde{N}$ with $\epsilon = 0.4$


 Fig. 4: Comparative study of NMF with IS- and KL-divergences, and DR-NMF with $\Omega = \{0, 1\}$ applied to the amplitude spectrogram of “Mary had a little lamb” with $r = 4$. The figure shows the evolution of the scaled β -divergences for the different NMF models.

and $\tilde{N} = \frac{N_{IS}}{\|N_{IS}\|_F}$. For this experiment, we overestimate the number of sources present into the input spectrogram by choosing $r = 5$; this allows to highlight the differences between the different NMF variants better. Figure 6 displays the rows of H for NMF with IS- and KL-divergences, and for DR-NMF with $\Omega = \{0, 1\}$.

KL-NMF identifies five sources among which the third one has no physical meaning and seems to be a mixture of several notes. IS-NMF correctly identifies the three notes, the fourth estimate (the hammer) is less accurately estimated in terms of amplitude for the activations but IS-NMF is able to set to zero the fifth estimate which is appealing as it automatically remove an unnecessary component. DR-NMF again takes advantage from both divergences as it is able to extract the three notes correctly, the fourth estimate (the hammer) is well extracted and the fifth estimate is close to zero. This again illustrates that DR-NMF is robust to




Fig. 5: Comparative study of NMF with IS- and KL-divergences, and DR-NMF with $\Omega = \{0, 1\}$ applied to “Mary had a little lamb” amplitude spectrogram with $r = 4$ and Poisson noise. The figure shows the activations (that is, the rows of H) of the recovered sources over time, and indicates which source it corresponds to (the notes C_4 , D_4 , E_4 , and the hammer noise H_N).




Fig. 6: Comparative study of NMF with IS- and KL-divergences, and DR-NMF with $\Omega = \{0, 1\}$ applied to “Mary had a little lamb” amplitude spectrogram with $r = 5$ and Gamma noise. The figure shows the activations (that is, the rows of H) of the recovered sources over time. Note that the hammer noise is also extracted with C_4 and D_4 , although with a small intensity.

different types of noises (in this case, multiplicative Gamma noise).

6 CONCLUSION AND FURTHER WORK

In this paper, we have proposed an NMF model that takes into account several data fitting terms. We then proposed

to tackle this problem with a weighted-sum approach with carefully chosen weights, and designed variations of MU algorithm to minimize the corresponding objective function. We used this model to design a DR-NMF algorithm, inspired from the Frank-Wolfe algorithm, that allows to obtain NMF solutions with low reconstruction errors with respect to several objective functions. We illustrated the effectiveness of this approach on synthetic, document and audio data sets. For audio data sets, DR-NMF provided particularly stunning results, being able to obtain solutions with significantly lower IS and KL errors (simultaneously), while generating meaningful solutions under different noise models or statistics. It is our hope that the proposed algorithms for DR-NMF (Algorithm 3) resolve the long-standing debate [47], [48] on whether to use IS- or KL-NMF for audio data sets. Using DR-NMF provides a safe alternative when one is uncertain of the noise statistics of audio data sets. Indeed, the noise statistics is rarely, if at all, known in practice.

Possible further research include the design of more efficient algorithms to solve multi-objective NMF, the extension of our distributionally robust model to low-rank tensor decompositions, and the refinement of our model by adding additional penalty terms or constraints to exploit properties, such as sparsity, smoothness or minimum volume [3], [15], [51], in the decompositions. Another challenging direction of research is to consider the DR-NMF problem with an uncountably infinite uncertainty set Ω such as $\Omega = [0, 2]$. Finally, an important direction of research that we plan to investigate is the design of an efficient algorithm for DR-NMF with convergence guarantees; see Section 4.

ACKNOWLEDGMENTS

We thank the reviewers and the handling editor for their insightful comments that helped us improve the paper.

REFERENCES

- [1] D. D. Lee and H. S. Seung, “Learning the parts of objects by non-negative matrix factorization,” *Nature*, vol. 401, no. 6755, p. 788, 1999.
- [2] A. Cichocki, R. Zdunek, and S.-i. Amari, “New algorithms for non-negative matrix factorization in applications to blind source separation,” in *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, vol. 5. IEEE, 2006, pp. V-V.
- [3] N. Gillis, “The why and how of nonnegative matrix factorization,” *Regularization, Optimization, Kernels, and Support Vector Machines*, vol. 12, no. 257, pp. 257–291, 2014.
- [4] P. O. Hoyer, “Non-negative matrix factorization with sparseness constraints,” *Journal of Machine Learning Research*, vol. 5, no. Nov, pp. 1457–1469, 2004.
- [5] X. Liu, W. Xia, B. Wang, and L. Zhang, “An approach based on constrained nonnegative matrix factorization to unmix hyperspectral data,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 2, pp. 757–772, 2011.
- [6] S. Essid and C. Févotte, “Smooth nonnegative matrix factorization for unsupervised audiovisual document structuring,” *IEEE Transactions on Multimedia*, vol. 15, no. 2, pp. 415–425, 2013.
- [7] C. Févotte and J. Idier, “Algorithms for nonnegative matrix factorization with the β -divergence,” *Neural Computation*, vol. 23, no. 9, pp. 2421–2456, 2011.
- [8] V. Y. F. Tan and C. Févotte, “Automatic relevance determination in nonnegative matrix factorization with the β -divergence,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 7, pp. 1592–1605, 2013.

- [9] R. T. Marler and J. S. Arora, "The weighted sum method for multi-objective optimization: new insight," *Structural and Multidisciplinary Optimization*, vol. 41, no. 6, pp. 853–862, 2010.
- [10] M. N. H. Mollah, S. Eguchi, and M. Minami, "Robust prewhitening for ICA by minimizing β -divergence and its application to fastica," *Neural Processing Letters*, vol. 25, no. 2, pp. 91–110, 2007.
- [11] H. Choi, S. Choi, A. Kataké, and Y. Choe, "Learning α -integration with partially-labeled data," in *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*. IEEE, 2010, pp. 2058–2061.
- [12] Z. Lu, Z. Yang, and E. Oja, "Selecting β -divergence for non-negative matrix factorization by score matching," in *International Conference on Artificial Neural Networks*. Springer, 2012, pp. 419–426.
- [13] O. Dikmen, Z. Yang, and E. Oja, "Learning the information divergence," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 7, pp. 1442–1454, 2015.
- [14] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust optimization*. Princeton University Press, 2009, vol. 28.
- [15] A. Cichocki, R. Zdunek, A. H. Phan, and S.-i. Amari, *Nonnegative matrix and tensor factorizations: applications to exploratory multi-way data analysis and blind source separation*. John Wiley & Sons, 2009.
- [16] K. Deb, "Multi-objective optimization," in *Search methodologies*. Springer, 2014, pp. 403–449.
- [17] M. Gong, X. Jiang, H. Li, and K. C. Tan, "Multiobjective sparse non-negative matrix factorization," *IEEE Transactions on Cybernetics*, vol. 49, no. 8, pp. 2941–2954, 2018.
- [18] F. Zhu and P. Honeine, "Biobjective nonnegative matrix factorization: Linear versus kernel-based models," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 7, pp. 4012–4022, 2016.
- [19] U. Şimşekli, A. T. Cemgil, and B. Ermiş, "Learning mixed divergences in coupled matrix and tensor factorization models," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 2120–2124.
- [20] S. A. Vavasis, "On the complexity of nonnegative matrix factorization," *SIAM J. Optim.*, vol. 20, no. 3, pp. 1364–1377, 2010.
- [21] D. D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Advances in neural information processing systems*, 2001, pp. 556–562.
- [22] Z. Yang and E. Oja, "Unified development of multiplicative algorithms for linear and quadratic nonnegative matrix factorization," *IEEE Transactions on Neural Networks*, vol. 22, no. 12, pp. 1878–1891, 2011.
- [23] C.-J. Lin, "On the convergence of multiplicative update algorithms for nonnegative matrix factorization," *IEEE Transactions on Neural Networks*, vol. 18, no. 6, pp. 1589–1596, 2007.
- [24] R. Zhao and V. Y. F. Tan, "A unified convergence analysis of the multiplicative update algorithm for regularized nonnegative matrix factorization," *IEEE Transactions on Signal Processing*, vol. 66, no. 1, pp. 129–138, 2018.
- [25] M. W. Berry, M. Browne, A. N. Langville, V. P. Pauca, and R. J. Plemmons, "Algorithms and applications for approximate nonnegative matrix factorization," *Computational statistics & data analysis*, vol. 52, no. 1, pp. 155–173, 2007.
- [26] N. Gillis, "Nonnegative matrix factorization: Complexity, algorithms and applications," *PhD thesis, Université catholique de Louvain*, 2011.
- [27] N. Takahashi and R. Hibi, "Global convergence of modified multiplicative updates for nonnegative matrix factorization," *Computational Optimization and Applications*, vol. 57, no. 2, pp. 417–440, 2014.
- [28] Y. Nesterov, "Smooth minimization of non-smooth functions," *Math. Program.*, vol. 103, no. 1, pp. 127–152, 2005.
- [29] A. Chambolle and T. Pock, "On the ergodic convergence rates of a first-order primal-dual algorithm," *Math. Program.*, vol. 159, no. 1, pp. 253–287, 2016.
- [30] Y. Chen, G. Lan, and Y. Ouyang, "Optimal primal-dual methods for a class of saddle point problems," *SIAM J. Optim.*, vol. 24, no. 4, pp. 1779–1814, 2014.
- [31] A. Nemirovski, "Prox-method with rate of convergence $O(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems," *SIAM J. Optim.*, vol. 15, no. 1, pp. 229–251, 2005.
- [32] A. Juditsky and A. Nemirovski, *First-Order Methods for Nonsmooth Convex Large-Scale Optimization, II: Utilizing Problem's Structure*. MIT Press, 2012, pp. 149–184.
- [33] P. Mertikopoulos, B. Lecouat, H. Zenati, C.-S. Foo, V. Chandrasekhar, and G. Piliouras, "Optimistic mirror descent in saddle-point problems: Going the extra(-gradient) mile," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=Bkg8jC9KQ>
- [34] E. Y. Hamedani and N. S. Aybat, "A primal-dual algorithm for general convex-concave saddle point problems," arXiv:1803.01401.
- [35] L. T. K. Hien, R. Zhao, and W. B. Haskell, "An inexact primal-dual smoothing framework for large-scale non-bilinear saddle point problems," arXiv:1711.03669.
- [36] S. Lu, I. Tsaknakis, and M. Hong, "Block alternating optimization for non-convex min-max problems: Algorithms and applications in signal processing and communications," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 4754–4758.
- [37] D. P. Bertsekas, *Convex Optimization Theory*. Athena Scientific, Belmont, Massachusetts, 2009.
- [38] K. M. Anstreicher and L. A. Wolsey, "Two "well-known" properties of subgradient optimization," *Math. Prog.*, vol. 120, no. 1, pp. 213–220, 2009.
- [39] S. Bubeck, "Convex optimization: Algorithms and complexity," *Found. Trends Mach. Learn.*, vol. 8, no. 3–4, p. 231–357, Nov. 2015. [Online]. Available: <https://doi.org/10.1561/22000000050>
- [40] M. Frank and P. Wolfe, "An algorithm for quadratic programming," *Naval Research Logistics Quarterly*, vol. 3, no. 1–2, pp. 95–110, 1956.
- [41] M. Jaggi, "Revisiting Frank-Wolfe: Projection-free sparse convex optimization," in *Proceedings of the 30th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, S. Dasgupta and D. McAllester, Eds., vol. 28, no. 1. Atlanta, Georgia, USA: PMLR, 17–19 Jun 2013, pp. 427–435. [Online]. Available: <http://proceedings.mlr.press/v28/jaggi13.html>
- [42] E. C. Chi and T. G. Kolda, "On tensors, sparsity, and nonnegative factorizations," *SIAM Journal on Matrix Analysis and Applications*, vol. 33, no. 4, pp. 1272–1299, 2012.
- [43] S. Zhong and J. Ghosh, "Generative model-based document clustering: a comparative study," *Knowledge and Information Systems*, vol. 8 (3), pp. 374–384, 2005.
- [44] U. Araújo, B. Saldanha, R. Galvão, T. Yoneyama, H. Chame, and V. Visani, "The successive projections algorithm for variable selection in spectroscopic multicomponent analysis," *Chemometrics and Intelligent Laboratory Systems*, vol. 57, no. 2, pp. 65–73, 2001.
- [45] N. Gillis and S. A. Vavasis, "Fast and robust recursive algorithms for separable nonnegative matrix factorization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 4, pp. 698–714, 2013.
- [46] S. Arora, R. Ge, Y. Halpern, D. Mimno, A. Moitra, D. Sontag, Y. Wu, and M. Zhu, "A practical algorithm for topic modeling with provable guarantees," in *International Conference on Machine Learning*, 2013, pp. 280–288.
- [47] T. Virtanen, "Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 3, pp. 1066–1074, 2007.
- [48] C. Févotte, N. Bertin, and J.-L. Durrieu, "Nonnegative matrix factorization with the Itakura-Saito divergence: With application to music analysis," *Neural computation*, vol. 21, no. 3, pp. 793–830, 2009.
- [49] J. Fritsch, "High quality musical audio source separation," Master's thesis, UPMC / IRCAM / Telecom Paristech, 2012.
- [50] N. Gillis, "Learning with nonnegative matrix factorizations," *SIAM News*, vol. 25, no. 2, pp. 1–3, 2019.
- [51] X. Fu, K. Huang, N. D. Sidiropoulos, and W.-K. Ma, "Nonnegative matrix factorization for signal and data analytics: Identifiability, algorithms, and applications." *IEEE Signal Process. Mag.*, vol. 36, no. 2, pp. 59–80, 2019.