Blind Audio Source Separation with Minimum-Volume Beta-Divergence NMF

Valentin Leplat, Nicolas Gillis, Andersen M.S. Ang*

Abstract—Considering a mixed signal composed of various audio sources and recorded with a single microphone, we consider on this paper the blind audio source separation problem which consists in isolating and extracting each of the sources. To perform this task, nonnegative matrix factorization (NMF) based on the Kullback-Leibler and Itakura-Saito β -divergences is a standard and state-of-theart technique that uses the time-frequency representation of the signal. We present a new NMF model better suited for this task. It is based on the minimization of β -divergences along with a penalty term that promotes the columns of the dictionary matrix to have a small volume. Under some mild assumptions and in noiseless conditions, we prove that this model is provably able to identify the sources. In order to solve this problem, we propose multiplicative updates whose derivations are based on the standard majorizationminimization framework. We show on several numerical experiments that our new model is able to obtain more interpretable results than standard NMF models. Moreover, we show that it is able to recover the sources even when the number of sources present into the mixed signal is overestimated. In fact, our model automatically sets sources to zero in this situation, hence performs model order selection automatically.

Index Terms—nonnegative matrix factorization, β -divergences, minimum-volume regularization, identifiability, blind audio source separation, model order selection

I. INTRODUCTION

Blind audio source separation concerns the techniques used to extract unknown signals called sources from a mixed audio signal x. In this paper, we assume that the audio signal is recorded with a single microphone. Considering a mixed signal composed of various audio sources, the blind audio source separation consists in isolating and extracting each of the sources on the

well posed. The most common technique used for this kind of problem is to get some form of redundancy in the mixed signal in order to make it overdetermined. This is typically done by computing the spectrogram which represents the signal in the time and frequency domains simultaneously (splitting the signals into overlapping time frames). The computation of spectrograms can be summarized as follows: short time segments are extracted from the signal and multiplied element wise by a window function or "smoothing" window of size F. Successive windows overlap by a fraction of their length, which is usually taken as 50%. On each of these segments, a discrete Fourier transform is computed and stacked column-by-column in a matrix X. Thus, from a one-dimensional signal $x \in \mathbb{R}^T$, we obtain a complex matrix $X \in \mathbb{C}^{F \times N}$ called spectrogram where $F \times N \simeq$ 2T (due to the 50% overlap between windows). Note that the length of the window determines the shape of the spectrogram. These preliminary operations correspond to computing the short time Fourier transform (STFT), which is given by the following formula: for $1 \le f \le F$ and $1 \le n \le N$, $X_{f,n} = \sum_{j=0}^{F-1} w_j x_{nL+j} e^{(-i\frac{2\pi f j}{F})}$, where $w \in \mathbb{R}^F$ is the smoothing window of size F, L is a shift parameter (also called hop size), and H = F - L is the overlap parameter. The number of rows corresponds to the frequency resolution. Letting f_s be the sampling rate of the audio signal, consecutive rows correspond to frequency bands that are f_s/F Hz apart.

basis of the single recording. Usually, the only known

information is the number of estimated sources present

in the mixed signal. The blind source separation problem

is said to be underdetermined as there are fewer sensors

(only one in our case) than sources. It then appears nec-

essary to find additional information to make the problem

The time-frequency representation of a signal highlights two of its fundamental properties: sparsity and redundancy. Sparsity comes from the fact that most real signals are not active at all frequencies at all time points. Redundancy comes from the fact that frequency patterns of the sources repeat over time. Mathematically, this means that the spectrogram is a low-rank matrix. These

^{*} Department of Mathematics and Operational Research, Faculté Polytechnique, Université de Mons, Rue de Houdain 9, 7000 Mons, Belgium. Authors acknowledge the support by the Fonds de la Recherche Scientifique - FNRS and the Fonds Wetenschappelijk Onderzoek - Vlanderen (FWO) under EOS Project no 0005318F-RG47, and by the European Research Council (ERC starting grant no 679515). E-mails: {valentin.leplat, nicolas.gillis, manshun.ang}@umons.ac.be. Manuscript received in July 2019.

two fundamental properties led sound source separation techniques to integrate algorithms such as nonnegative matrix factorization (NMF). Such techniques retrieve sensible solutions even for single-channel signals.

A. Mixing assumptions

Given K source signals $s^{(k)} \in \mathbb{R}^T$ for $1 \le k \le K$, we assume the acquisition process is well modelled by a linear instantaneous mixing model:

$$x(t) = \sum_{k=1}^{K} s^{(k)}(t) \quad \text{with } t = 0, ..., T - 1.$$
 (1)

Therefore, for each time index t, the mixed signal x(t)from a single microphone is the sum of the K source signals. It is standard to assume that microphones are linear as long as the recorded signals are not too loud. If signals are too loud, they are usually clipped. The mixing process is modelled as instantaneous as opposed to convolutive used to take into account sound effects such as reverberation. The source separation problem consist in finding source estimates $\hat{s}^{(k)}$ of $s^{(k)}$ sources for all $k \in \{1, \ldots, K\}$. Let us denote S the linear STFT operator, and let S^{\dagger} be its conjugate transpose. We have $S^{\dagger}S = FI$, where I is the identity matrix of appropriate dimension. For the remainder of this paper, S^{\dagger} stands for the inverse short time Fourier transform. Note that the term inverse is not meant in a mathematical sense. Indeed the STFT is not a surjective transformation from \mathbb{R}^T to $\mathbb{C}^{F \times N}$. In other words, each spectrogram or each matrix with complex entries is not necessarily the STFT of a real signal; see [1] and [2] for more details. By applying the STFT operator S to (1), we obtain the mixing model in the time-frequency domain :

$$X = S(x(t)) = S\left(\sum_{k=1}^{K} s^{(k)}(t)\right) = \sum_{k=1}^{K} S^{(k)},$$

where $S^{(k)}$ is the STFT of the source k, that is, the spectrogram of source k. To identify the sources, we use in this paper the amplitude spectrogram $V = |X| \in \mathbb{R}_+^{F \times N}$ defined as $V_{fn} = |X_{fn}|$ for all f, n. We assume that $V = \sum_{k=1}^{K} |S^{(k)}|$, which means that there is no sound cancellation between the sources, which is usually the case in most signals. Finally, we assume that the source spectrograms $|S^{(k)}|$ are well approximated by nonnegative rank-one matrices. This leads to the NMF model described in the next section. Note that a source can be made of several rank-one factors in which case a post-processing step will have to recombine them a posteriori (e.g., looking at the correspondence in the activation of the sources over time). Note also that we focus on the NMF stage of the source separation which factorizes V into the source spectrograms. For the phases reconstruction, which is a highly non-trivial problem, we consider a naive reconstruction procedure consisting in keeping the same phase as the input mixture for each source [1].

B. NMF for audio source separation

Given a non-negative matrix $V \in \mathbb{R}^{F \times N}_+$ (the spectrogram) and a positive integer $K \ll \min(F, N)$ (the number of sources, called the factorization rank), NMF aims to compute two non-negative matrices W with K columns and H with K rows such that $V \approx WH$. NMF approximate each column of V by a linear combination of the columns of W weighted by the components of the corresponding column of H [3]. When the matrix V corresponds to the amplitude spectrogram or the power spectrogram of an audio signal, we have that

• W is referred as the dictionary matrix and each column corresponds to the spectral content of a source, and

• H is the activation matrix specifying if a source is active at a certain time frame and in which intensity.

In other words, each rank-one factor W(:, k)H(k, :) will correspond to a source: the *k*th column W(:, k) of *W* is the spectral content of source *k*, and the *k*th row H(k, :)of *H* is its activation over time. To compute *W* and *H*, NMF requires to solve the following optimization problem

$$\min_{W \ge 0, H \ge 0} D(V|WH) = \sum_{f,n} d(V_{fn}|[WH]_{fn}),$$

where $A \ge 0$ means that A is component-wise nonnegative, and d(x|y) is an appropriate measure of fit. In audio source separation, a common measure of fit is the discrete β -divergence denoted $d_{\beta}(x|y)$ and equal to

$$\frac{1}{\beta(\beta-1)} \left(x^{\beta} + (\beta-1) y^{\beta} - \beta x y^{\beta-1} \right) \text{ for } \beta \neq 0, 1,$$

$$x \log \frac{x}{y} - x + y \text{ for } \beta = 1,$$

$$\frac{x}{y} - \log \frac{x}{y} - 1 \text{ for } \beta = 0.$$

For $\beta = 2$, this the standard squared Euclidean distance, that is, the squared Frobenius norm $||V - WH||_F^2$. For $\beta = 1$ and $\beta = 0$, the β -divergence corresponds to the Kullback-Leibler (KL) divergence and the Itakura-Saito (IS) divergence, respectively. The error measure which should be chosen accordingly with the noise statistic assumed on the data. The Frobenius norm assumes i.i.d. Gaussian noise, KL divergence assumes additive Poisson noise, and the IS divergence assumes multiplicative Gamma noise [4]. The β -divergence $d_{\beta}(x|y)$ is homogeneous of degree β : $d_{\beta}(\lambda x | \lambda y) = \lambda^{\beta} d_{\beta}(x | y)$. It implies that factorizations obtained with $\beta > 0$ (such as the Euclidean distance or the KL divergence) will rely more heavily on the largest data values and less precision is to be expected in the estimation of the lowpower components. The IS divergence ($\beta = 0$) is scaleinvariant that is $d_{IS}(\lambda x | \lambda y) = d_{IS}(x | y)$ [5]. The IS divergence is the only one in the β -divergences family to possess this property. It implies that time-frequency areas of low power are as important in the divergence computation as the areas of high power. This property is interesting in audio source separation as low-power frequency bands can perceptually contribute as much as high-power frequency bands. Note that both KL and IS divergences are more adapted to audio source separation than Euclidean distance as it is built on logarithmic scale as human perception; see [1] and [5]. Moreover, the β -divergence is only convex with respect to W (or H) if $\beta \geq 1$. Otherwise, the objective function is nonconvex. This implies that, for $\beta < 1$, even the problem of inferring H with W fixed is non-convex. For more details on β -divergences; see [5].

C. Contribution and outline of the paper

In Section II, we propose a new NMF model, referred to as minimum-volume β -NMF (min-vol β -NMF), to tackle the audio source separation problem. This model penalizes the columns of the dictionary matrix W so that their convex hull has a small volume. To the best of our knowledge, this model is novel in two aspects: (1) it is the first time a minimum-volume penalty is associated with a β -divergence for $\beta \neq 2$ and it is the first time such models are used in the context of audio source separation, and (2) as opposed to most previously proposed minimum-volume NMF models, our model imposes a normalization constraints on the factor W instead of H. As far as we know, the only other paper that used a normalization of W is [6] but the authors did not justify this choice compared to the normalization of H (the choice seems arbitrary, motivated by the 'elimination of the norm indeterminacy'), nor provided theoretical guarantees. In this paper, we explain why normalization of W is a better choice in practice, and we prove that, under some mild assumptions and in the noiseless case, this model provably identify the sources; see Theorem 1. To the best of our knowledge, this is the first result of this type in the audio source separation literature. In Section III, we propose an algorithm to tackle min-vol β -NMF, focusing on the KL and IS divergences. The algorithm is based on multiplicative updates (MU) that are derived using the standard majorization-minimization framework, and that monotonically decrease the objective function. In Section IV, we present several numerical experiments, comparing min-vol β -NMF with standard NMF and sparse NMF. The two mains conclusions are that (1) minimum-volume β -NMF performs consistently better to identify the sources, and (2) as opposed to NMF and sparse NMF, min-vol β -NMF is able to detect when the factorization rank is overestimated by automatically setting sources to zero.

II. MINIMUM-VOLUME NMF WITH β -divergences

In this section, we present a new model of separation based on the minimization of β -divergences including a penalty term promoting solutions with minimum volume spanned by the columns of the dictionary matrix W. Section II-A recalls the geometric interpretation of NMF which motivated the use of a minimum volume penalty on the dictionary W. Section II-B discusses the new proposed normalization compared to previous minimum volume NMF models, and proves that min-vol β -NMF provably recovers the true factors (W, H) under mild conditions and in the noiseless case; see Theorem 1.

A. Geometry and the min-vol β -NMF model

As mentioned earlier, V = WH means that each column of V is a linear combination of the columns of W weighted by the components of the corresponding column of H; in fact, $v_n = Wh_n$ for n = 1, ..., N, where v_n denotes the *n*th column of data matrix V. This gives to NMF a nice geometric interpretation: for all n

$$v_n \in \operatorname{cone}(W) = \left\{ v \in \mathbb{R}^F | v = W\theta, \theta \ge 0 \right\},\$$

meaning that the columns of V are contained in the convex cone generated by the columns of W; see Figure 1 for an illustration. From this interpretation, it follows that, in general, NMF decompositions are not unique because there exists several (often, infinitely many) sets of columns of W that span the convex cone generated by the data points; see for example [8] for more details. Hence, NMF is in most cases ill-posed because the optimal solution is not unique. In order to make the solution unique (up to permutation and scaling of the columns of W and the rows of H) hence making the problem well-posed and the parameters (W, H) of the problem identifiable, a key idea is to look for a solution W with minimum volume. Intuitively, we will look for the cone cone(W) containing the data points and as close



Fig. 1: Geometric interpretation of NMF for K = 3 [7].

as possible to these data points. The use of minimumvolume NMF has lead to a new class of NMF methods that outperforms existing ones in many applications such as document analysis and blind hyperspectral unmixing; see the recent survey [9]. Note that minimum-volume NMF implicitly enhances the factor H to be sparse: the fact that W has a small volume implies that many data points will be located on the facets of the cone(W) hence H will be sparse.

Hence, in this paper, we consider the following model, referred to as min-vol β -NMF:

$$\min_{W(:,j)\in\Delta^F\forall j,H\geq 0} \quad D_{\beta}(V|WH) + \lambda \mathrm{vol}(W), \qquad (2)$$

where $\Delta^F = \left\{ x \in \mathbb{R}^F_+ \big| \sum_{i=1}^F x_i = 1 \right\}$ is the unit simplex, λ is a penalty parameter and vol(W) is a function that measures the volume spanned by the columns of W. In this paper, we use $vol(W) = logdet(W^TW + \delta I)$, where δ is a small positive constant that prevents $logdet(W^TW)$ to go to $-\infty$ when W tends to a rank-deficient matrix (that is, when $r = \operatorname{rank}(W) <$ K). The reason for using such a measure is that $\sqrt{\det(\overline{W^TW})}/K!$ is the volume of the convex hull of the columns of W and the origin. This measure is one of the most widely used ones, and has been shown to perform very well in practice [10], [11]. Moreover, the criterion logdet $(W^TW + \delta I)$ is able to distinguish two rank-deficient solutions and favour solutions for W with smaller volume [12]. Finally, as we will illustrate in Section IV, this criterion is able to identify the right number of source even when K is overestimated, by putting some rank-one factors to zero.

B. Normalization and identifiability of min-vol β -NMF

As mentioned above, under some appropriate conditions on V = WH, minimum-volume NMF models will provably recover the ground-truth (W, H) that generated V, up to permutation and scaling of the rank-one factors. The first identifiability results for minimum-volume NMF models assumed that the entries in each column of H sum to one, that is, that $H^T e = e$ where e is the all-one column vector whose dimension is clear from the context, meaning that H is column stochastic [8], [13]. Under this condition, each column of V lies in the *convex hull* of the columns of W; see Figure 2 for an illustration.



Fig. 2: Geometric interpretation of NMF when K = 3 and H is column stochastic [7].

Under the three assumptions that (1) H is column stochastic, (2) W is full column rank, and (3) H satisfies the sufficiently scattered condition, minimizing the volume of conv(W) such that V = WH recover the true underlying factors, up to permutation and scaling. Intuitively, the sufficiently scattered condition requires H to be sparse enough so that data points are located on the facets of conv(W); see Appendix A for a formal definition. The sufficiently scattered condition makes sense for most audio source data sets as it is reasonable to assume that, for most time points, only a few sources are active hence H is sparse; see [9] for more details on the sufficiently scattered condition. Note that the sufficiently scattered condition is a generalization of the separability condition which requires $W = V(:, \mathcal{J})$ for some index set \mathcal{J} of size K [14]. However, separability is a much stronger assumption as it requires that, for each sources, there exists a time point where only that source is active. Note that although min-vol NMF guarantees identifiability, the corresponding optimization problem (2) is still hard to solve in general, as for the original NMF problem [15].

Despite this nice result, the constraint $H^T e = e$ makes the NMF model less general and does not apply to all data sets. In the case where the data does not naturally belong to a convex hull, one needs to normalize the data points so that their entries sum to one so that $H^T e = e$ can be assumed without loss of generality (in the noiseless case). This normalization can sometimes increase the noise and might greatly influence the solution, hence are usually not recommended in practice; see the discussion in [9].

In [7], authors show that identifiability still holds when the condition that H is column stochastic is relaxed to H being row stochastic. As opposed to column stochasticity, row stochasticity of H can be assumed without loss of generality since any factorization WHcan be properly normalized so that this assumption holds. In fact, $WH = \sum_{k=1}^{K} (a_k W(:, k))(H(k, :)/a_k)$ for any $a_k > 0$ for k = 1, ..., K. In other terms, letting A be the diagonal matrix with $A(k, k) = a_k = \sum_{j=1}^{n} H(k, j)$ for k = 1, ..., K, we have $WH = (WA)(A^{-1}H) = W'H'$ where $H' = A^{-1}H$ is row stochastic.

Similarly as done in [7], we prove in this paper that requiring W to be column stochastic (which can also be made without loss of generality) also leads to identifiability. Geometrically, the columns of W are constrained to be on the unit simplex. Minimizing the volume still makes a lot of sense: we want the columns of W to be as close as possible to one another within the unit simplex. In Appendix A, we prove the following theorem.

Theorem 1. Assume $V = W^{\#}H^{\#}$ with rank(V) = K, $W^{\#} \ge 0$ and $H^{\#}$ satisfies the sufficiently scattered condition (Definition 2 in Appendix A). Then the optimal solution of

$$\min_{W \in \mathbb{R}^{F \times K}, H \in \mathbb{R}^{K \times N}} \operatorname{logdet} (W^{T}W)$$
(3)
such that $V = WH, W^{T}e = e, H \ge 0,$

recovers $(W^{\#}, H^{\#})$ up to permutation and scaling.

Proof. See Appendix A.

In noiseless conditions, replacing $W^T e = e$ with He = e in (3) leads to the same identifiability result; see [7, Theorem 1]. Therefore, in noiseless conditions and under the conditions of Theorem 1, both models return the same solution up to permutation and scaling. However, in the presence of noise, we have observed that the two models may behave very differently. In fact, we advocate

that the constraint $W^T e = e$ is better suited for noisy real-world problems, which we have observed on many numerical examples. In fact, we have observed that the normalization $W^T e = e$ is much less sensitive to noise and returns much better solutions. The reason is mostly twofold:

(i) As described above, using the normalization He = eamounts to multiply W by a diagonal matrix whose entries are the ℓ_1 norms of the rows of H. Therefore, the columns of W that correspond to dominating (resp. dominated) sources, that is, sources with much more (resp. less) power and/or active at many (resp. few) time points, will have much higher (resp. lower) norm. Therefore, the term $logdet(W^TW + \delta I)$ is much more influenced by the dominating sources and will have difficulties to penalize the dominated sources. In other terms, the use of the term $logdet(W^TW + \delta I)$ with the normalization He = e implicitly requires that the rankone factors W(:,k)H(k,:) for $k = 1, \ldots, K$ are well balanced, that is, have similar norms. This is not the case for many real (audio) signals.

(ii) As it will be explained in Section III, the update of W needs the computation of the matrix Y which is the inverse of $W^TW + \delta I$ -this terms appears in the gradient with respect to W of the objective function. The numerical stability for such operations is related to the condition number of $W^TW + \delta I$. For a ℓ_1 normalization on the columns of W, the condition number is bounded above as follows: $\operatorname{cond}(W^T W + \delta I) = \frac{\sigma_{\max}(W^T W + \delta I)}{\sigma_{\min}(W)^2 + \delta} = \frac{\sigma_{\max}(W)^2 + \delta}{\delta} \leq \frac{(\sqrt{K} \max_k ||W(:,k)||_2)^2 + \delta}{\delta} \leq 1 + \frac{K}{\delta}$, where $\sigma_{\min}(W)$ and $\sigma_{\max}(W)$ are the smallest and largest singular values of W, respectively. In the numerical experiments, we use $\delta = 1$. On the other hand, the normalization He = e may lead to arbitrarily large values for the condition number of $W^T W + \delta I$, which we have observed numerically on several examples. This issue can be mitigated with the use of the normalization $He = \rho e$ for some $\rho > 0$ sufficiently large for which identifiability still holds [7]. However, it still performs worse because of the first reason explained above.

For these reasons, we believe that the model (3) would also be better suited (compared to the normalization on H) in other contexts; for example for document classification [16].

III. Algorithm for min-vol β -NMF

Most NMF algorithms alternatively update H for W fixed and vice versa, and we adopt this strategy in this paper. For W fixed, (2) is equivalent to standard NMF

and we will use the MU that have already been derived in the literature [3], [5].

To tackle (2) for H fixed, let us consider

$$\min_{W \ge 0} \quad F(W) = D_{\beta}(V|WH) + \lambda \operatorname{logdet}(W^{T}W + \delta I).$$
(4)

Note that, for now, we have discarded the normalization on the columns of W. In our algorithm, we will use the update for W obtained by solving (4) as a descent direction along with a line search procedure to integrate the constraint on W. This will ensure that the objective function F is non-increasing at each iteration. In the following sections we derive MU for W that decrease the objective in (4). We follow the standard majorizationminimization framework [17]. First, an auxiliary function, which we denote \overline{F} , is constructed so that it majorizes the objective. An auxiliary function for F at point \widetilde{W} is defined as follows.

Definition 1. The function $\overline{F}(W|\tilde{W}) : \Omega \times \Omega \to \mathbb{R}$ is an auxiliary function for $F(W) : \Omega \to \mathbb{R}$ at $\tilde{W} \in \Omega$ if the conditions $\overline{F}(W|\tilde{W}) \ge F(W)$ for all $W \in \Omega$ and $\overline{F}(\tilde{W}|\tilde{W}) = F(\tilde{W})$ are satisfied.

Then, the optimization of F can be replaced by an iterative process that minimizes \overline{F} . More precisely, the new iterate $W^{(i+1)}$ is computed by minimizing exactly the auxiliary function at the previous iterate $W^{(i)}$. This guarantees F to decrease at each iteration.

Lemma 1. Let $W, W^{(i)} \ge 0$, and let \overline{F} be an auxiliary function for F at $W^{(i)}$. Then F is non-increasing under the update $W^{(i+1)} = \underset{W \ge 0}{\operatorname{argmin}} \overline{F}(W|W^{(i)})$.

Proof. In fact, we have by definition that
$$F(W^{(i)}) = \overline{F}(W^{(i)}|W^{(i)}) \ge \min_{W} \overline{F}(W|W^{(i)}) = \overline{F}(W^{(i+1)}|W^{(i)}) \ge F(W^{(i+1)}).$$

The most difficult part in using the majorizationminimization framework is to design an auxiliary function that is easy to optimize. Usually such auxiliary functions are separable (that is, there is no interaction between the variables so that each entry of W can be updated independently) and convex.

A. Separable auxiliary functions for β -divergences

For the sake of completeness, we briefly recall the auxiliary function proposed in [5] for the data fitting term. It consists in majorizing the convex part of the β -divergence using Jensen's inequality and majorizing the

concave part by its tangent (first-order Taylor approximation). We have

$$d_{\beta}(x|y) = \check{d}_{\beta}(x|y) + \hat{d}_{\beta}(x|y) + \bar{d}_{\beta}(x|y), \qquad (5)$$

where d is convex function of y, d is a concave function of y and \overline{d} is a constant of y; see Table I.

TABLE I: Differentiable convex-concave-constant decomposition of the β -divergence under the form (5) [5].

	$\check{d}(x y)$	$\hat{d}(x y)$	$ar{d}(x)$
$\beta = 0$	xy^{-1}	$\log(y)$	$x(\log(x) - 1)$
$\beta \in [1,2]$	$d_{\beta}(x y)$	0	0

The function $D_{\beta}(V|WH)$ can be written as $\sum_{f} D_{\beta}(v_{f}|w_{f}H)$ where v_{f} and w_{f} are respectively the *f*th row of *V* and *W*. Therefore we only consider the optimization over one specific row of *W*. To simplify notation, we denote iterates $w^{(i+1)}$ (next iterate) and $w^{(i)}$ (current iterate) as *w* and \tilde{w} , respectively.

Lemma 2 ([5]). Let $\tilde{v} = \tilde{w}H$ and \tilde{w} be such that $\tilde{v_n} > 0$ for all n and $\tilde{w_k} > 0$ for all k. Then the function

$$G(w|\tilde{w}) = \sum_{n} \left[\sum_{k} \frac{\tilde{w}_{k}h_{kn}}{\tilde{v}_{n}} \check{d}(v_{n}|\tilde{v}_{n}\frac{w_{k}}{\tilde{w}_{k}}) \right] + \bar{d}(v_{n}) \\ + \left[\hat{d}'(v_{n}|\tilde{v}_{n}) \sum_{k} (w_{k} - \tilde{w}_{k})h_{kn} + \hat{d}(v_{n}|\tilde{v}_{n}) \right]$$
(6)

is an auxiliary function for $\sum_{n} d(v_n | [wH]_n)$ at \tilde{w} .

B. A separable auxiliary function for the minimumvolume regularizer

The minimum-volume regularizer $logdet(W^TW + \delta I)$ is a non-convex function. However, it can be upperbounded using the fact that logdet(.) is a concave function so that its first-order Taylor approximation provides an upper bound; see for example [10]. For any positivedefinite matrices A and $B \in \mathbb{R}^{K \times K}$, we have:

$$\log\det(B) \le \log\det(A) + \operatorname{trace}\left(A^{-1}(B-A)\right)$$
$$= \operatorname{trace}\left(A^{-1}B\right) + \log\det(A) - K.$$

This implies that for any $W, Z \in \mathbb{R}^{F \times K}$, we have

1

$$\operatorname{ogdet}(W^T W + \delta I) \le l(W, Z),$$
(7)

where $l(W, Z) = \text{trace}(YW^TW) + \text{logdet}(Y^{-1}) - K$, $Y = (Z^TZ + \delta I)^{-1}$ with $\delta > 0$. Note that $Z^TZ + \delta I$ is positive definite hence is invertible and its inverse Y is also positive definite. Finally l(W, Z) is an auxiliary function for $\text{logdet}(W^TW + \delta I)$ at Z. However, it is quadratic and not separable hence non-trivial to optimize over the nonnegative orthant. The non-constant part of l(W,Z) can be written as $\sum_f w_f Y w_f^T$ where w_f is the *f*th row of *W*. Henceforth we will focus on one particular row vector *w* with $l(w) = w^T Y w$ which will be further considered as a column vector of size $K \times 1$.

Lemma 3. Let $w, \tilde{w} \in \mathbb{R}_+^K$ be such that $\tilde{w}_k > 0$ for all $k, Y = Y^+ - Y^-$ with $Y^+ = \max(Y, 0)$ and $Y^- = \max(-Y, 0)$, and $\Phi(\tilde{w})$ be the diagonal matrix $\Phi(\tilde{w}) = \text{Diag}\left(2\frac{[Y^+\tilde{w}+Y^-\tilde{w}]}{[\tilde{w}]}\right)$ where $\frac{[A]}{[B]}$ is the component-wise division between A and B, and $\Delta w = w - \tilde{w}$. Then

$$\bar{l}(w|\tilde{w}) = l(\tilde{w}) + \Delta w^T \nabla l\left(\tilde{w}\right) + \frac{1}{2} \Delta w^T \Phi(\tilde{w}) \Delta w,$$
(8)

is a separable auxiliary function for $l(w)=w^TYw$ at \tilde{w} .

Proof. See Appendix B.
$$\Box$$

Remark 1 (Choice of the auxiliary function). A simpler choice for the auxiliary function would be to replace $\Phi(\tilde{w})$ with $2\lambda_{\max}(Y)I$ where $\lambda_{\max}(Y)$ is the largest eigenvalue of Y (the constant 2 appears because $l(w) = w^T Y w$ while there is a factor 1/2 in front of $\Phi(\tilde{w})$). However, it would lead to a worse approximation. In particular if Y is a diagonal matrix (since $Y \succ 0$, these diagonal elements are positive), our choice gives $\Phi(\tilde{w}) = 2Y$ for any $\tilde{w} > 0$, meaning that the auxiliary function matches perfectly the function l(w). This would not be the case for the choice $2\lambda_{\max}(Y)I$ (unless Y is a scaling of the identity matrix).

C. Auxiliary function for min-vol β -NMF

Based on the auxiliary functions presented in Sections III-A and III-B, we can directly derive a separable auxiliary function $\overline{F}(W|\tilde{W})$ for min-vol β -NMF (2).

Corollary 1. For $W, H \ge 0$, $\lambda > 0$, $Y = (\tilde{W}^T \tilde{W} + \delta I)^{-1}$ with $\delta > 0$ and the constant $c = \text{logdet}(Y^{-1}) + K$, the function

$$\bar{F}(W|\tilde{W}) = \sum_{f} G\left(w_{f}|\tilde{w}_{f}\right) + \lambda\left(\sum_{f} \bar{l}\left(w_{f}|\tilde{w}_{f}\right) + c\right)$$

where G is given by (6) and \bar{l} by (8), is a convex and separable auxiliary function for $F(W) = D_{\beta}(V|WH) + \lambda \log\det(W^TW + \delta I)$ at \tilde{W} .

Proof. This follows directly from Lemma 2, Equation (7) and Lemma 3. \Box

In the following section, we provide explicitly MU for the KL divergence ($\beta = 1$) by finding a closed-form solution for the minimization of \overline{F} . In Appendix C, we provide the MU for the IS divergence ($\beta = 0$). Due to the lack of space, the other cases are not treated explicitly but can be in a similar way. For the same reason, we will only compare KL NMF models in the numerical experiments (Section IV).

D. Algorithm for min-vol KL-NMF

As before, let us focus on a single row of W, denoted w, as the objective function F(W) is separable by row. For $\beta = 1$, the derivative of the auxiliary function $\overline{F}(w|\tilde{w})$ with respect to a specific coefficient w_k is given by $\nabla_{w_k} \overline{F}(w|\tilde{w}) = \sum_n h_{kn} - \sum_n h_{kn} \frac{\tilde{w}_k v_n}{w_k \tilde{v}_n} + 2\lambda \left[\text{Diag} \left(\frac{Y + \tilde{w} + Y - \tilde{w}}{\tilde{w}} \right) \right]_k w_k - 2\lambda \left[\text{Diag} \left(\frac{Y + \tilde{w} + Y - \tilde{w}}{\tilde{w}} \right) \right]_k \tilde{w}_k$. Due to the separability, we set the derivative to zero to obtain the closed-form solution, which is given in Table II in matrix form.

Note that although the closed-form solution has a negative term in the numerator of the multiplicative factor (see Table II), they always remain nonnegative given that V, H and \tilde{W} are nonnegative. In fact, the term before the minus sign is always larger than the term after the minus sign: $J_{F,N}H^T - 4\lambda(\tilde{W}Y^-)$ is squared (component wise) and added a positive term, hence the component-wise square root of that result is larger than $J_{F,N}H^T - 4\lambda(\tilde{W}Y^-)$.

Algorithm 1 summarizes our algorithm to tackle (2) for $\beta = 1$ which we refer to as min-vol KL-NMF. Note that the update of H (step 4) is the one from [3]. More importantly, note that we have incorporated a line-search for the update of W. In fact, although the MU for Ware guaranteed to decrease the objective function, they do not guarantee that W remains normalized, that is, that $||W(:,k)||_1 = 1$ for all k. Hence, we normalize W after it is updated (step 10), and we normalize Haccordingly so that WH remains unchanged. When this normalization is performed, the β -divergence part of F is unchanged but the minimum-volume penalty will change so that the objective function F might increase. In order to guarantee non-increasingness, we integrate a simple backtracking line-search procedure; see steps 11-16 of Algorithm 1. In summary, our MU provide a descent direction that preserved nonnegativity of the iterates, and we use a projection and a simple backtracking linesearch to guarantee the monotonicity of the objective TABLE II: Multiplicative update for min-vol KL-NMF.

$$W = \tilde{W} \odot \frac{\left[\left[\left[J_{F,N} H^T - 4\lambda (\tilde{W}Y^-) \right]^{\cdot 2} + 8\lambda \tilde{W} (Y^+ + Y^-) \odot \left(\frac{[V]}{[\tilde{W}H]} H^T \right) \right]^{\cdot \frac{1}{2}} - \left(J_{F,N} H^T - 4\lambda (\tilde{W}Y^-) \right) \right]}{[4\lambda \tilde{W} (Y^+ + Y^-)]}$$

where $A \odot B$ (resp. $\frac{[A]}{[B]}$) is the Hadamard product (resp. division) between A and B, $A^{(.\alpha)}$ is the element-wise α exponent of A, $J_{F,N}$ is the F-by-N all-one matrix, and $Y = Y^+ - Y^- = (\tilde{W}^T \tilde{W} + \delta I)^{-1}$ with $\delta > 0$, $Y^+ \ge , Y^- \ge 0$, $\lambda > 0$.

function, as in standard projected gradient descent methods.

Algorithm 1 min-vol KL-NMF

- **Input:** A matrix $V \in \mathbb{R}^{M \times T}$, an initialization $H \in \mathbb{R}^{K \times T}_+$, an initialization $W \in \mathbb{R}^{M \times K}$, a factorization rank K, a maximum number of iterations maxiter, min-vol weight $\lambda > 0$ and $\delta > 0$
- **Output:** A rank-K NMF (W, H) of $V \approx WH$ with $W \ge 0$ and $H \ge 0$.

1: $\gamma = 1, Y = (W^T W + \delta I)^{-1}$ 2: for i = 1 : maxiter do % Update of matrix H 3: $\begin{array}{l} H \leftarrow H \odot \frac{\left[W^T\left(\frac{|V|}{|WH|}\right)\right]}{\left[W^T J_{F,N}\right]} \\ \% \ Update \ of \ matrix \ W \end{array}$ 4: 5: $Y \leftarrow (W^T W + \delta I)^{-1}$ 6: $Y^+ \leftarrow \max(Y, 0)$ 7: $Y^- \leftarrow \max(-Y, 0)$ 8: W^+ is updated according to Table II 9: $(W^+_{\gamma}, H_{\gamma}) =$ normalize (W^+, H) 10: % Line-search procedure 11: while $F(W_{\gamma}^+, H_{\gamma}) > F(W, H)$ do 12: $\begin{array}{l} \gamma \leftarrow \gamma \times 0.8 \\ W_{\gamma}^{+} \leftarrow (1 - \gamma) W + \gamma W^{+} \\ (W_{\gamma}^{+}, H_{\gamma}) \leftarrow \text{normalize} \left(W_{\gamma}^{+}, H \right) \end{array}$ 13: $14 \cdot$ 15: end while 16: $(W, H) \leftarrow (W_{\gamma}^+, H_{\gamma})$ 17: % Update of γ to avoid a vanishing stepsize 18: 19: $\gamma \leftarrow \min(1, \gamma \times 1.2)$ 20: end for

It can be verified that the computational complexity of the min-vol KL-NMF is asymptotically equivalent to the standard MU for β -NMF, that is, it requires $\mathcal{O}(FNK)$ operations per iteration. Indeed, all the main operations include matrix products with a complexity of $\mathcal{O}(FNK)$ and element-wise operations on matrices of size $F \times K$ or $K \times N$. Note that the inversion of the K-by-K matrix $(W^TW + \delta I)$ requires $\mathcal{O}(K^3)$ operations which is dominated by $\mathcal{O}(FNK)$ since $K \leq \min(F, N)$ (in fact, typically $K \ll \min(F, N)$ hence this term is negligible). Therefore, although Algorithm 1 will be slower than the baseline KL-NMF (that is, the standard MU) because of the additional terms to be computed and the linesearch, the asymptotical computational cost is the same; see Table IV for runtime comparison.

IV. NUMERICAL EXPERIMENTS

In this section we report an experimental comparative study of baseline KL-NMF, min-vol KL-NMF (Algorithm 1) and sparse KL-NMF [18] applied to the spectrogram of two monophonic piano sequences and a synthetic mix of a bass and drums. For the two monophonic piano sequences, the audio signals are true life signals with standard quality. Note that the sequences are made of pure piano notes, the number K should therefore correspond to the number of notes present into the mixed signals. The comparative study is performed for several values of K with a focus on the case where the factorization rank K is overestimated. For all simulations, random initializations are used for Wand H, the best results among 5 runs are kept for the comparative study. In all cases, we use a Hamming window of size F=1024, and 50% overlap between two frames. Sparse KL-NMF has a similar structure as minvol KL-NMF, with a penalty parameter for the sparsity enforcing regularization. To tune these two parameters, we have used the same strategy for both methods: we manually tried a wide range of values and report the best results. The code is available from bit.ly/minvolKLNMF (code written in MATLAB R2017a), and can be used to rerun directly all experiments below. They were run on a laptop computer with Intel Core i7-7500U CPU @ 2.70GHz 4 and 32GB memory.

a) Mary had a little lamb: The first audio sample is the first measure of "Mary had a little lamb". The sequence is composed of three notes; E_4 , D_4 and C_4 , played all at once. The recorded signal is 4.7 seconds long and downsampled to $f_s = 16000$ Hz yielding T=75200 samples. STFT of the input signal x yields a temporal resolution of 16ms and a frequency resolution of 31.25Hz, so that the amplitude spectrogram V has N=294 frames and F=257 frequency bins. The musical score is shown on Figure 3. All NMF algorithms were run for 200 iterations which allowed them to converge. Figure 4 presents the columns of W (dictionary matrix)



Fig. 3: Musical score of "Mary had a little lamb".



Fig. 4: Comparative study of baseline KL-NMF (top), min-vol KL-NMF (middle) and sparse KL-NMF (bot-tom) applied to "Mary had a little lamb" amplitude spectrogram with K=3.

and the rows of H for baseline KL-NMF and min-vol KL-NMF with K = 3. Figure 5 presents the time-frequency masking coefficients. These coefficients are computed as follows

$$\mathrm{mask}_{f,n}^{(k)} = \frac{\hat{X}_{f,n}^{(k)}}{\sum_k \hat{X}_{f,n}^{(k)}} \qquad \text{with } k = 1, ..., K \,,$$

where $\hat{X}^{(k)} = W(:,k)H(k,:)$ is the estimated source k. The masks are nonnegative and sum to one for each pair (f,n). This representation allows to identify visually whether the NMF algorithm was able to separate the sources properly. All the simulations give a nice separation with similar results for W and H. The activations are coherent with the sequences of the notes. However, Figure 5 shows that min-vol KL-NMF and sparse KL-NMF provide a better separation in terms of time-frequency localization compared to the baseline KL-NMF.



(c) sparse KL-NMF

Fig. 5: Masking coefficients obtained with baseline KL-NMF (top), min-vol KL-NMF (middle) and sparse KL-NMF (bottom) applied to "Mary had a little lamb" amplitude spectrogram with K=3.

We now perform the same experiment but using K=7. Figure 6 presents the results. This situation corresponds to the situation where the factorization rank is overestimated. Figure 7 presents the time-frequency masking coefficients. We observe that min-vol KL-NMF is able to extract the three notes correctly and set automatically to zero three source estimates (more precisely, three rows of H are set to zero, while the corresponding columns of W have entries equal to one another as $||W(:,k)||_1 = 1$ for all k) while baseline KL-NMF and sparse KL-NMF split the notes in all the sources. One can observe that a fourth note is identified in all simulations (see isolated



Fig. 6: Comparative study of baseline KL-NMF (top), min-vol KL-NMF (middle) and sparse KL-NMF (bot-tom) applied to "Mary had a little lamb" amplitude spectrogram with K=7.

peaks on Figure 7-(b), second row of H from the top) and corresponds to the noise within the piano just before triggering a specific note (in particular, the hammer noise). This observation is confirmed by the fact that the amplitude is proportional to the natural strength of the fingers playing the notes. In this scenario, with K is overestimated, min-vol KL-NMF outperforms baseline KL-NMF and sparse KL-NMF.

b) Prelude of Bach: The second audio sample corresponds to the first 30 seconds of "Prelude and Fugue No.1 in C major" from J. S. Bach played by Glenn Gould¹. The audio sample is a sequence of 13 notes: $B_3, C_4, D_4, E_4, F_4^{\#}, G_4, A_4, C_5, D_5, E_5, F_5, G_5, A_5$. The recorded signal is downsampled to $f_s = 11025$ Hz yielding T=330750 samples. STFT of the input signal x yields a temporal resolution of 46ms and a frequency resolution of 10.76Hz, so that the amplitude spectrogram V has N=647 frames and F=513 frequency bins. The musical score is presented on Figure 8. All NMF algorithms were run for 300 iterations which allowed



(c) sparse KL-NMF

Fig. 7: Masking coefficients obtained with baseline KL-NMF (top), min-vol KL-NMF (middle) and sparse KL-NMF (bottom) applied to "Mary had a little lamb" amplitude spectrogram with K=7.

them to converge. Figure 9 presents the results obtained for W and H with a factorization rank K = 16, hence overestimated. We observe that min-vol KL-NMF automatically sets three components to zero (with * symbol on Figure 9) while 13 source estimates are determined. The analysis of the fundamentals (maximum peak frequency) of the 13 source estimates correspond to the theoretical fundamentals of the 13 notes mentioned earlier. Note that using baseline KL-NMF or sparse KL-NMF led to same conclusions as for the first audio

¹https://www.youtube.com/watch?v=ZlbK5r5mBH4



Fig. 8: Musical score of the sample "Prelude and Fugue No.1 in C major".

sample; these two algorithms generate as many source estimates as imposed by the rank of factorization while min-vol KL-NMF algorithm preserves the integrity of the 13 sources. Additionally, the activations are coherent with the sequences of the notes. Figure 10 shows (on a limited time interval) that the estimate sequence follows the sequence defined in the score. Note that a threshold and permutations on rows of H was used to improve visibility.

c) Bass and drums: The third audio signal is a synthetic mix of a bass and drums². The audio signal is downsampled to f_s =16000Hz yielding T=104821 samples. STFT of the input signal x yields a temporal resolution of 32ms and a frequency resolution of 15.62Hz, so that the amplitude spectrogram V has N=206 frames and F=513frequency bins. For this synthetic mix, we have access to the true sources under the form of two audio files. Therefore, we can estimate the quality of the separation with standard metrics, namely the signal to distortion ratios (SDR), the source to interference ratios (SIR) and the sources to artifacts ratios (SAR) [19]. They have been computed with the toolbox BSS Eval³. The metrics are expressed in dB and the higher they are the better is the separation. Algorithms min-vol KL-NMF, baseline KL-NMF and sparse KL-NMF have been considered for this comparative study. A factorization rank equal to two is used. It is clear that the rank-one approximation is too simplistic for these sources but the goal is to compare the algorithms and show that min-vol KL-NMF is able to find a better solution even in this simplified context. All NMF algorithms were run for 400 iterations which allowed them to converge. Table III shows the results.

Except for SAR metric for the second source (drums), min-vol KL-NMF outperforms baseline KL-NMF and sparse KL-NMF.

d) Runtime performance: Let us compare the runtime of baseline KL-NMF, min-vol KL-NMF (Algorithm 1) and sparse KL-NMF [18]. The algorithms are compares



Fig. 9: Factors matrices W and H obtained with min-vol KL-NMF with factorization rank K=16 on the sample "Prelude and Fugue No.1 in C major".

on the three examples presented in paragraphs IV-0a and IV-0b:

- Setup \sharp 1: sample "Mary had a little lamb" with K = 3,200 iterations.
- Setup $\sharp 2$: sample "Mary had a little lamb" with K = 7,200 iterations.
- Setup \sharp 3: "Prelude and Fugue No.1 in C major" with K = 16, 300 iterations.

For each test setup, the algorithms are run for the same 20 random initializations of W and H. Table IV reports the average and standard deviation of the runtime (in seconds) over these 20 runs. We observe that the runtime of min-vol KL-NMF (Algorithm 1) is slower but not significantly so, as expected. In particular, on the larger setup $\sharp 3$, it is less than three times slower than the standard MU.

²http://isse.sourceforge.net/demos.html

³http://bass-db.gforge.inria.fr/bss_eval/

TABLE III: SDR, SIR and SAR metrics comparison for results obtained with baseline KL-NMF and min-vol KL-NMF on a synthetic mix of bass and drums

Algorithms	Source 1: bass		Source 2: drums			
-	SDR(dB)	SIR(dB)	SAR(dB)	SDR(dB)	SIR(dB)	SAR(dB)
min-vol KL-NMF	-1.14	0.12	7.78	9.60	19.8	10.09
baseline KL-NMF	-4.26	-1.39	2.64	7.97	9.00	15.25
sparse KL-NMF	-4.69	-1.73	2.33	7.89	8.96	14.98



Fig. 10: Validation of the estimate sequence obtained with min-vol KL-NMF with factorization rank K=16 on the sample "Prelude and Fugue No.1 in C major".

TABLE IV: Runtime performance in seconds of baseline KL-NMF, min-vol KL-NMF (Algorithm 1) and sparse KL-NMF [18]. The table reports the average and standard deviation over 20 random initializations for three experimental setups described in the text.

Algorithms	runtime in seconds			
	setup #1	setup #2	setup #3	
baseline KL-NMF	$0.44{\pm}0.03$	0.43 ± 0.01	3.81±0.19	
min-vol KL-NMF	3.79 ± 0.13	2.39 ± 0.30	10.19 ± 1.28	
sparse KL-NMF	0.20 ± 0.02	0.20 ± 0.01	2.21 ± 0.26	

V. CONCLUSION AND PERSPECTIVES

In this paper, we have presented a new NMF model of audio source separation based on the minimization of a cost function that includes a β -divergence (data fitting term) and a penalty term that promotes solutions W with minimum volume. We have proved the identifiability of the model in the exact case, under the sufficiently scattered condition for the activation matrix H. We have provided multiplicative updates to tackle this problem and have illustrated the behaviour of the method on realworld audio signals. We highlighted the capacity of the model to deal with the case where K is overestimated by setting automatically to zero some components and give good results for the source estimates.

Further work includes tackling the following questions:

- Under which conditions can we prove the identifiability of min-vol β -NMF in the presence of noise, and the rank-deficient case?
- Can we prove that min-vol β-NMF performs model order selection automatically? Under which conditions? We have observed this behaviour on many examples, but the proof remains elusive.
- Can we design more efficient algorithms?

Further work also includes the use of our new model min-vol β -NMF for other applications and the design of more efficient algorithms (for example, that avoid using a line-search procedure) with stronger convergence guarantees (beyond the monotonicity of the objective function).

Acknowledgments: We thank Kejun Huang and Xiao Fu for helpful discussion on Theorem 1, and giving us the insight to adapt their proof from [7] to our model (2). We also thank the reviewers for their insightful comments that helped us improve the paper.

APPENDIX

A. Sufficiently scattered condition and identifiability

Before giving the definition of the sufficiently scattered condition from [8], let us first recall an important property of the duals of nested cones.

Lemma 4. Let C_1 and C_2 be convex cones such that $C_1 \subseteq C_2$. Then $C_2^* \subseteq C_1^*$ where C_2^* and C_1^* are respectively the dual cones of C_1 and C_2 . The dual of a cone C is defined as $C^* = \{y | x^T y \ge 0 \text{ for all } x \in C\}$.

Definition 2. (Sufficiently Scattered) A matrix $H \in \mathbb{R}^{K \times N}_+$ is sufficiently scattered if

- 1) $\mathcal{C} \subseteq \operatorname{cone}(H)$, and
- 2) cone $(H)^* \cap bd\mathcal{C}^* = \{\lambda e_k | \lambda \ge 0, k = 1, ..., K\},\$

where $C = \{x | x^T e \ge \sqrt{K-1} ||x||_2\}$ is a second order cone, $C^* = \{x | x^T e \ge ||x||_2\}$, cone $(H) = \{x | x = H\theta, \theta \ge 0\}$ is the conic hull of the columns of H, and bd denotes the boundary of a set.

We can now prove Theorem 1.

Proof of Theorem 1. Recall that $W^{\#}$ and $H^{\#}$ are the true latent factors that generated V, with $\operatorname{rank}(V) = K$ and $H^{\#}$ is sufficiently scattered. Let us consider \hat{W} and \hat{H} a feasible solution of (3). Since $\operatorname{rank}(V) = K$ and $V = \hat{W}\hat{H}$, we must have $\operatorname{rank}(\hat{W}) = \operatorname{rank}(\hat{H}) = K$. Hence there exists an invertible matrix $A \in \mathbb{R}^{K \times K}$ such that $\hat{W} = W^{\#}A^{-1}$ and $\hat{H} = AH^{\#}$. Since \hat{W} is a feasible solution of problem (3), we have

$$e^T \hat{W} = e^T W^{\#} A^{-1} = e^T A^{-1} = e^T$$

where we assumed $e^T W^{\#} = e^T$ without loss of generality since $W^{\#} \ge 0$ and $\operatorname{rank}(W^{\#}) = K$. Note that $e^T A^{-1} = e^T$ is equivalent to $e^T A = e^T$. This means that matrix A is column stochastic. Therefore we have that $e^T A e = K$. Since \hat{H} is a feasible solution, we also have $\hat{H} = AH^{\#} \ge 0$. Let us denote by a_j the *j*th row of A, and by a_k^T the *k*th column of A^T . By the definition of the a dual cone, $AH^{\#} \geq 0$ means that the rows $a_i \in \operatorname{cone}(H^{\#})^*$ for j = 1, ..., K. Since $H^{\#}$ is sufficiently scattered, cone $(H)^* \subseteq \mathcal{C}^*$ (by Lemma 4) hence $a_j \in \mathcal{C}^*$. Therefore we have $||a_j||_2 \leq$ $a_i e$ by definition of C. This leads to the following: $\begin{aligned} \|\det(A)\| &= |\det(A^T)| \le \prod_{k=1}^K \|a_k^T\|_2 = \prod_{j=1}^K \|a_j\|_2 \le \\ \prod_{j=1}^K a_j e &\le \left(\frac{\sum_j^K a_j e}{K}\right)^K = \left(\frac{e^T A e}{K}\right)^K = 1. \text{ The} \end{aligned}$ first inequality is the Hadamard inequality, the second inequality is due to $a_j \in \mathcal{C}^*$, the third inequality is the arithmetic-geometric mean inequality. Now we can conclude exactly as is done in [7, Theorem 1] by showing that matrix A can only be a permutation matrix for an optimal solution (W, H) of (3), and therefore identifiability for model (3) holds. \square

B. Proof of Lemma 3

Separability of $\bar{l}(w|\tilde{w})$ holds since $\Phi(\tilde{w})$ is diagonal. The condition $\bar{l}(\tilde{w}|\tilde{w}) = l(\tilde{w})$ from Definition 1 can be checked easily. It remains to prove that $\bar{l}(w|\tilde{w}) \ge l(w)$ for all w. Let us first rewrite the quadratic function l(w) using its Taylor expansion at $w = \tilde{w}$: $l(w) = l(\tilde{w}) + (w - \tilde{w})^T \nabla l(\tilde{w}) + \frac{1}{2}(w - \tilde{w})^T \nabla^2 l(\tilde{w})(w - \tilde{w}) = l(\tilde{w}) + (w - \tilde{w})^T 2Y\tilde{w} + \frac{1}{2}(w - \tilde{w})^T 2Y(w - \tilde{w})$. Proving that $\overline{l}(w|\tilde{w}) \ge l(w)$ is equivalent to proving that $\frac{1}{2}(w - \tilde{w})^T [\Phi(\tilde{w}) - 2Y](w - \tilde{w}) \ge 0$, which boils down to proving that the matrix $[\Phi(\tilde{w}) - 2Y]$ is positive semi-definite. We have $\Phi_{ij}(\tilde{w}) = 2\delta_{ij}\frac{(Y+\tilde{w})_i + (Y-\tilde{w})_i}{\tilde{w}_i}$, where δ_{ij} is the Kronecker symbol. Let us consider the following matrix: $M_{ij}(\tilde{w}) = \tilde{w}_i [\Phi(\tilde{w}) - 2Y]_{ij}\tilde{w}_j$, which is a rescaling of $[\Phi(\tilde{w}) - 2Y]$. It remains to show that M is positive semi-definite⁴. Since M is symmetric and its diagonal entries are non-negative, it is sufficient to show that M is diagonally dominant [20, Proposition 7.2.3], that is,

$$|M_{ii}| \ge \sum_{j \ne i} |M_{ij}|$$
 for all *i*.

We have for all *i* that

$$\begin{split} M_{ii} &= 2w_i \sum_j \left(Y_{ij}^+ + Y_{ij}^-\right) w_j - 2w_i Y_{ii} w_i, \text{ and} \\ M_{ij} &= -2w_i Y_{ij} w_j \quad \text{for } j \neq i. \end{split}$$

Since $Y_{ij}^+ + Y_{ij}^- = |Y_{ij}|$, we have

$$M_{ii} - \sum_{j \neq i} |M_{ij}| = 2w_i \sum_j |Y_{ij}| w_j - 2w_i Y_{ii} w_i - 2w_i \sum_{j \neq i} |Y_{ij}| w_j = 2w_i |Y_{ii}| w_i - 2w_i Y_{ii} w_i \ge 0,$$

implying that M is diagonally dominant.

C. Algorithm for min-vol IS-NMF

For $\beta = 0$ (IS divergence), the derivative of the auxiliary function $\overline{F}(w|\tilde{w})$ with respect to a specific coefficient w_k is given by:

$$\begin{split} \nabla_{w_k} \bar{F}(w|\tilde{w}) &= \sum_n \frac{h_{kn}}{\tilde{v}_n} - \sum_n h_{kn} \frac{\tilde{w}_k^2 v_n}{w_k^2 \tilde{v}_n^2} + 2\lambda \left[Y \tilde{w} \right]_k \\ &+ 2\lambda \left[\text{Diag} \left(\frac{Y^+ \tilde{w} + Y^- \tilde{w}}{\tilde{w}} \right) \right]_k w_k \\ &- 2\lambda \left[\text{Diag} \left(\frac{Y^+ \tilde{w} + Y^- \tilde{w}}{\tilde{w}} \right) \right]_k \tilde{w}_k. \end{split}$$

⁴The remainder of the proof was suggested to us by one of the reviewers, it is more elegant and simpler than our original proof.

Let

$$\tilde{a} = 2\lambda \left[\text{Diag} \left(\frac{Y^+ \tilde{w} + Y^- \tilde{w}}{\tilde{w}} \right) \right]_k,$$

$$\tilde{b} = \sum_n \frac{h_{kn}}{\tilde{v}_n} - 4\lambda \left[Y^- \tilde{w} \right]_k,$$

$$\tilde{d} = -\sum_n h_{kn} \frac{\tilde{w}_k^2 v_n}{\tilde{v}_n^2}.$$
(9)

Setting the derivative to zero requires to compute the roots of the following degree-three polynomial $\tilde{a}w_k^3$ + $bw_k^2 + d$. We used the procedure developed in [21] which is based on the explicit calculation of the intermediary root of a canonical form of cubic. This procedure is able to provide highly accurate numerical results even for badly conditioned polynomials. The algorithm for minvol IS-NMF follows the same steps as for min-vol KL-NMF: only the two steps corresponding to the updates of W and H have to be modified. For the update of H (step 4), use the standard MU. For the update of W (step 9). use

for $f \leftarrow 1$ to F

for $k \leftarrow 1$ to K

Compute \tilde{a} , \tilde{b} and \tilde{d} according to equations (9)

Compute the roots of $\tilde{a}w_k^3 + \tilde{b}w_k^2 + \tilde{d}$

Pick y among these roots and zero that minimizes the objective

 $W_{f,k}^+ \leftarrow \max\left(10^{-16}, y\right)$ end for

end for

REFERENCES

- [1] A. Lefèvre, "Méthode d'apprentissage de dictionnaire pour la séparation de sources audio avec un seul capteur," Ph.D. dissertation, Ecole Normale Supérieure de Cachan, 2012.
- [2] P. Magron, "Reconstruction de phase par modèles de signaux : application à la séparation de sources audio," Ph.D. dissertation, TELECOM ParisTech, 2016.
- [3] D. Lee and H. Seung, "Algorithms for non-negative matrix factorization," in NIPS'00 Proceedings of the 13th International Conference on Neural Information Processing Systems, NIPS. MIT Press Cambridge, 2000, pp. 535-541.
- [4] 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.
- [5] C. Févotte and J. Idier, "Algorithms for nonnegative matrix factorization with the β -divergence," Neural computation, vol. 23, no. 9, pp. 2421-2456, 2011.
- [6] G. Zhou, S. Xie, Z. Yang, J.-M. Yang, and Z. He, "Minimumvolume-constrained nonnegative matrix factorization: Enhanced ability of learning parts," IEEE Transactions on Neural Networks, vol. 22, no. 10, pp. 1626-1637, 2011.
- [7] X. Fu, K. Huang, and N. D. Sidiropoulos, "On identifiability of nonnegative matrix factorization," IEEE Signal Processing Letters, vol. 25, no. 3, pp. 328-332, 2018.

- [8] K. Huang, N. Sidiropoulos, and A. Swami, "Non-negative matrix factorization revisited: Uniqueness and algorithm for symmetric decomposition," IEEE Transactions on Signal Processing, vol. 62, no. 1, pp. 211-224, 2014.
- [9] X. Fu, K. Huang, N. Sidiropoulos, and W.-K. Ma, "Nonnegative matrix factorization for signal and data analytics: Identifiability, algorithms, and applications," IEEE Signal Processing Magazine, vol. 36, pp. 59-80, 2019.
- [10] X. Fu, K. Huang, B. Yang, W.-K. Ma, and N. Sidiropoulos, "Robust volume minimization-based matrix factorization for remote sensing and document clustering," IEEE Transactions on Signal Processing, vol. 64, no. 23, p. 6254-6268.
- A. Ang and N. Gillis, "Algorithms and comparisons of non-[11] negative matrix factorizations with volume regularization for hyperspectral unmixing," IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, vol. 12, no. 12, pp. 4843-4853, 2019.
- V. Leplat, A. Ang, and N. Gillis, "Minimum-volume rank-[12] deficient nonnegative matrix factorizations," in IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2019, pp. 3402-3406.
- [13] C.-H. Lin, W.-K. Ma, W.-C. Li, C.-Y. Chi, and A. Ambikapathi, "Identifiability of the simplex volume minimization criterion for blind hyperspectral unmixing: The no-pure-pixel case," IEEE Transactions on Geoscience and Remote Sensing, vol. 53, no. 10, pp. 5530-5546, 2015.
- [14] S. Arora, R. Ge, R. Kannan, and A. Moitra, "Computing a nonnegative matrix factorization-provably," SIAM Journal on Computing, vol. 45, no. 4, pp. 1582-1611, 2016.
- [15] S. Vavasis, "On the complexity of nonnegative matrix factorization," SIAM Journal on Optimization, vol. 20, no. 3, pp. 1364-1377, 2010.
- [16] X. Fu, K. Huang, N. D. Sidiropoulos, Q. Shi, and M. Hong, "Anchor-free correlated topic modeling," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 41, no. 5, pp. 1056-1071, 2019.
- [17] Y. Sun, P. Babu, and D. Palomar, "Majorization-minimization algorithms in signal processing, communications, and machine learning," IEEE Transactions on Signal Processing, vol. 65, no. 3, pp. 794-816, 2017.
- [18] J. L. Roux, F. J. Weninger, and J. R. Hershey, "Sparse NMF half-baked or well done?" Mitsubishi Electric Research Laboratories (MERL), Tech. Rep., 2015.
- [19] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," IEEE Transations on Audio, Speech, and Language Processing, vol. 14, no. 4, pp. 1462 - 1469, June 2006.
- [20] R. A. Horn and C. R. Johnson, Matrix analysis. Cambridge university press, 1985.
- [21] E. Rechtschaffen, "Real roots of cubics: explicit formula for quasi-solutions," The Mathematical Gazette, no. 524, p. 268-276, 2008.