

Algorithms for Nonnegative Matrix Factorization with the β -Divergence

Cédric Févotte

fevotte@telecom-paristech.fr

*Laboratoire Traitement et Communication de l'Information
(CNRS and Télécom Paris Tech), 75014 Paris, France*

Jérôme Idier

Jerome.Idier@ircsyn.ec-nantes.fr

*Institute de Recherche en Communications et Cybernetique de Nantes
(CNRS, Ecole Centrale de Nantes, Ecole des Mines de Nantes, and Université de Nantes), 44000 Nantes, France*

This letter describes algorithms for nonnegative matrix factorization (NMF) with the β -divergence (β -NMF). The β -divergence is a family of cost functions parameterized by a single shape parameter β that takes the Euclidean distance, the Kullback-Leibler divergence, and the Itakura-Saito divergence as special cases ($\beta = 2, 1, 0$ respectively). The proposed algorithms are based on a surrogate auxiliary function (a local majorization of the criterion function). We first describe a majorization-minimization algorithm that leads to multiplicative updates, which differ from standard heuristic multiplicative updates by a β -dependent power exponent. The monotonicity of the heuristic algorithm can, however, be proven for $\beta \in (0, 1)$ using the proposed auxiliary function. Then we introduce the concept of the majorization-equalization (ME) algorithm, which produces updates that move along constant level sets of the auxiliary function and lead to larger steps than MM. Simulations on synthetic and real data illustrate the faster convergence of the ME approach. The letter also describes how the proposed algorithms can be adapted to two common variants of NMF: penalized NMF (when a penalty function of the factors is added to the criterion function) and convex NMF (when the dictionary is assumed to belong to a known subspace).

1 Introduction ---

Given a data matrix \mathbf{V} of dimensions $F \times N$ with nonnegative entries, non-negative matrix factorization (NMF) is the problem of finding a factorization

$$\mathbf{V} \approx \mathbf{W}\mathbf{H}, \tag{1.1}$$

where \mathbf{W} and \mathbf{H} are nonnegative matrices of dimensions $F \times K$ and $K \times N$, respectively. K is usually chosen such that $F K + K N \ll F N$, hence reducing the data dimension. The factorization is in general only approximate, so that the terms *approximate nonnegative matrix factorization* and *nonnegative matrix approximation* also appear in the literature. NMF has been used for various problems in diverse fields. To cite a few, we mention the problems of learning parts of faces and semantic features of text (Lee & Seung, 1999), polyphonic music transcription (Smaragdis & Brown, 2003), object characterization by reflectance spectra analysis (Berry, Browne, Langville, Pauca, & Plemmons, 2007), portfolio diversification (Drakakis, Rickard, de Frein, & Cichocki, 2007), DNA gene expression analysis (Brunet, Tamayo, Golub, & Mesirov, 2004; Gao & Church, 2005), clustering of protein interactions (Greene, Cagney, Krogan, & Cunningham, 2008), and image denoising and inpainting (Mairal, Bach, Ponce, & Sapiro, 2010). The factorization, equation 1.1, is usually sought after through the minimization problem

$$\min_{\mathbf{W}, \mathbf{H}} D(\mathbf{V}|\mathbf{WH}) \text{ subject to } \mathbf{W} \geq 0, \mathbf{H} \geq 0, \quad (1.2)$$

where the notation $\mathbf{A} \geq 0$ expresses nonnegativity of the entries of matrix \mathbf{A} (and not semidefinite positiveness), and where $D(\mathbf{V}|\mathbf{WH})$ is a separable measure of fit such that

$$D(\mathbf{V}|\mathbf{WH}) = \sum_{f=1}^F \sum_{n=1}^N d([\mathbf{V}]_{fn} | [\mathbf{WH}]_{fn}), \quad (1.3)$$

where $d(x|y)$ is a scalar cost function. What we intend by “cost function” is a positive function of $y \in \mathbb{R}_+$ given $x \in \mathbb{R}_+$, with a single minimum for $x = y$.

A popular cost function in NMF is the β -divergence $d_\beta(x|y)$ of Basu, Harris, Hjort, and Jones (1998), Eguchi and Kano (2001), and Cichocki and Amari (2010), defined rigorously in section 2.1. In essence, it is a parameterized cost function with a single parameter β , which takes the Euclidean distance, the generalized Kullback-Leibler (KL) divergence, and the Itakura-Saito (IS) divergence as special cases ($\beta = 2, 1$, and 0 , respectively). NMF with the β -divergence has been widely used in music signal processing, in particular, for transcription and source separation (O’Grady, 2007; O’Grady & Pearlmutter, 2008; FitzGerald, Cranitch, & Coyle, 2009; Bertin, Févotte, & Badeau, 2009; Févotte, Bertin, & Durrieu, 2009; Vincent, Bertin, & Badeau, 2010; Dessein, Cont, & Lemaître, 2010; Hennequin, Badeau, & David, 2010). In these works the nonnegative data matrix \mathbf{V} is a spectrogram that is decomposed into elementary spectra with NMF. The parameter β can be tuned so as to optimize transcription or separation accuracy on training data. While popular in music signal processing, NMF with the β -divergence

(shortened as β -NMF in the rest of the letter) can be of interest to any field: the parameter β essentially controls the assumed statistics of the observation noise and can be either fixed or learned from training data or by cross-validation. As Févotte and Cemgil (2009) noted, the values $\beta = 2, 1, 0$, respectively underlie gaussian additive, Poisson, and multiplicative gamma observation noise. The β -divergence offers a continuum of noise statistics that interpolate among these three specific cases (Basu et al., 1998; Eguchi & Kano, 2001; Minami & Eguchi, 2002; Cichocki & Amari, 2010).

The standard β -NMF algorithm used in the above-mentioned papers is presented as a gradient-descent algorithm where the step size is set adaptively and chosen such that the updates are multiplicative, as originally described by Cichocki, Zdunek, and Amari (2006). The same algorithm can be derived from the following heuristic, proposed by Févotte et al. (2009). Let θ be a coefficient of \mathbf{W} or \mathbf{H} . As will be seen later, when using the β -divergence, the derivative $\nabla_{\theta} D(\theta)$ of the criterion $D(\mathbf{V}|\mathbf{WH})$ with respect to (w.r.t) θ can be expressed as the difference of two nonnegative functions such that $\nabla_{\theta} D(\theta) = \nabla_{\theta}^{+} D(\theta) - \nabla_{\theta}^{-} D(\theta)$. Then a heuristic multiplicative algorithm simply writes

$$\theta \leftarrow \theta \cdot \frac{\nabla_{\theta}^{-} D(\theta)}{\nabla_{\theta}^{+} D(\theta)}, \quad (1.4)$$

which ensures the nonnegativity of the parameter updates, provided initialization with a nonnegative value. It produces a descent algorithm in the sense that θ is updated toward left (resp. right) when the gradient is positive (resp. negative). A fixed point θ^* of the algorithm implies either $\nabla_{\theta} D(\theta^*) = 0$ or $\theta^* = 0$. Monotonicity of this algorithm has been proven by Kompass (2007) for the specific range of values of β for which the divergence $d_{\beta}(x|y)$ is convex w.r.t y (i.e., $\beta \in [1, 2]$; see section 2.1). The proof is based on a majorization-minimization (MM) procedure: an auxiliary function is built and iteratively minimized for each column of \mathbf{H} (given \mathbf{W}) and each row of \mathbf{W} (given \mathbf{H}). The auxiliary function is built using Jensen's inequality, thanks to the convexity of the cost for $\beta \in [1, 2]$. However, it was observed in practice that the multiplicative algorithm, 1.4, is still monotone (i.e., decreases the criterion function at each iteration) for values of β out of the "convexity" interval $[1, 2]$, though no proof is available.

This letter studies three descent algorithms for β -NMF, based on an auxiliary function that unifies existing auxiliary functions for the Euclidean distance and KL divergence (De Pierro, 1993; Lee & Seung, 2001), the "generalized divergence" of Kompass (2007), and the IS divergence (Cao, Eggermont, & Terebey, 1999). This auxiliary function was also recently proposed independently by Nakano et al. (2010). The construction of the auxiliary function relies on the decomposition of the criterion function into its convex and concave parts, following the approach of Cao et al.

(1999) for the IS divergence. An auxiliary function to the convex part is constructed using Jensen’s inequality, while the concave part is locally majorized by its tangent. It is shown that MM algorithms based on the latter auxiliary function yield multiplicative updates that coincide with the heuristic described by equation 1.4 for $\beta \in [1, 2]$, but differ from a β -dependent power exponent when $\beta \notin [1, 2]$, a result also obtained by Nakano et al. (2010). Additionally, we show that the monotonicity of the heuristic algorithm can be proven for $\beta \in (0, 1)$, using the proposed auxiliary function (it is shown to produce a descent algorithm, though it does not fully minimize the auxiliary function). Then we introduce the concept of the maximization-equalization (ME) algorithm, which produces updates that move along constant level sets of the auxiliary function and leads to larger steps than MM. This is akin to overrelaxation and is shown experimentally to produce faster convergence. Finally, we show how the described MM, ME, and heuristic algorithms can be adapted to two common variants of NMF: penalized NMF (i.e., when a penalty function of \mathbf{W} or \mathbf{H} is added to the criterion function) and convex NMF (when the dictionary is assumed to belong to a known subspace, as proposed by Ding, Li, and Jordan, 2010).

The letter is organized as follows. Section 2 defines and discusses the β -divergence and then exposes in detail the optimization task addressed in this letter. Section 3 recalls the concept of auxiliary function and then introduces a general auxiliary function for the β -NMF problem. Section 4 describes algorithms based on the proposed auxiliary function, namely, MM and ME algorithms, and describes how they relate to the heuristic update, equation 1.4. Section 5 reports simulations and convergence behaviors on synthetic and real data (with audio transcription and face interpolation examples). Section 6 describes extensions of the proposed algorithms to penalized and convex NMF. Section 7 concludes and discusses open questions.

2 Preliminaries

In this section we present the β -divergence and more precisely specify the task that is addressed in this letter. A detailed exposition of the β -divergence can be found in Cichocki & Amari (2010).

2.1 Definition of the β -Divergence. The β -divergence was introduced by Basu et al. (1998) and Eguchi and Kano (2001) and can be defined as

$$d_\beta(x|y) \stackrel{\text{def}}{=} \begin{cases} \frac{1}{\beta(\beta-1)} (x^\beta + (\beta-1)y^\beta - \beta x y^{\beta-1}) & \beta \in \mathbb{R} \setminus \{0, 1\} \\ x \log \frac{x}{y} - x + y & \beta = 1 \\ \frac{x}{y} - \log \frac{x}{y} - 1 & \beta = 0 \end{cases} \quad (2.1)$$

Basu et al. (1998) and Eguchi and Kano (2001) assume $\beta \geq 1$, but the definition domain can be extended to $\beta \in \mathbb{R}$, as suggested by Cichocki et al. (2006), which is the definition domain that is considered in this letter. The β -divergence can be shown continuous in β by using the identity $\lim_{\beta \rightarrow 0} (x^\beta - y^\beta)/\beta = \log(x/y)$. The limit cases $\beta = 0$ and $\beta = 1$ correspond to the IS and KL divergences, respectively. The β -divergence coincides up to a factor $1/\beta$ with the “generalized divergence” of Kompass (2007), which, in the context of NMF as well, was separately constructed so as to interpolate between the KL divergence ($\beta = 1$) and the Euclidean distance ($\beta = 2$). The β -divergence is plotted for various values of β on Figure 1. Note that in this letter, we will abusively refer to $d_{\beta=2} = (x - y)^2/2$ as the Euclidean distance, though the latter is formally defined with a square root, and for vectors.

The first and second derivative of $d_\beta(x|y)$ w.r.t y are continuous in β :

$$\begin{aligned} d'_\beta(x|y) &= y^{\beta-2} (y - x), \\ d''_\beta(x|y) &= y^{\beta-3} [(\beta - 1)y - (\beta - 2)x]. \end{aligned} \tag{2.2}$$

The derivative shows that $d_\beta(x|y)$, as a function of y , has a single minimum in $y = x$ and that it increases with $|y - x|$, justifying its relevance as a measure of fit. The second derivative shows that the β -divergence is convex w.r.t y for $\beta \in [1, 2]$. Outside this interval, the divergence can always be expressed as the sum of a convex, concave, and constant part, such that

$$d_\beta(x|y) = \check{d}(x|y) + \hat{d}(x|y) + \bar{d}(x), \tag{2.3}$$

where $\check{d}(x|y)$ is a convex function of y , $\hat{d}(x|y)$ is a concave function of y , and $\bar{d}(x)$ is a constant of y . The decomposition is not unique, since constant or linear terms (w.r.t y) are both convex and concave or, less trivially, since any convex term can be added to $\check{d}(x|y)$ while subtracted from $\hat{d}(x|y)$. In the following, we will use the natural conventions given in Table 1.

As noted by Févotte et al. (2009), a noteworthy property of the β -divergence is its behavior w.r.t to scale, as the following equation holds for any value of β :

$$d_\beta(\lambda x|\lambda y) = \lambda^\beta d_\beta(x|y). \tag{2.4}$$

It implies that factorizations obtained with $\beta > 0$ (such as with 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 low-power components, and, conversely, factorizations obtained with $\beta < 0$ will rely more heavily on smallest data values. The IS divergence ($\beta = 0$) is scale invariant, that is, $d_{IS}(\lambda x|\lambda y) = d_{IS}(x|y)$, and is the only one in the family of β -divergences to possess this property.

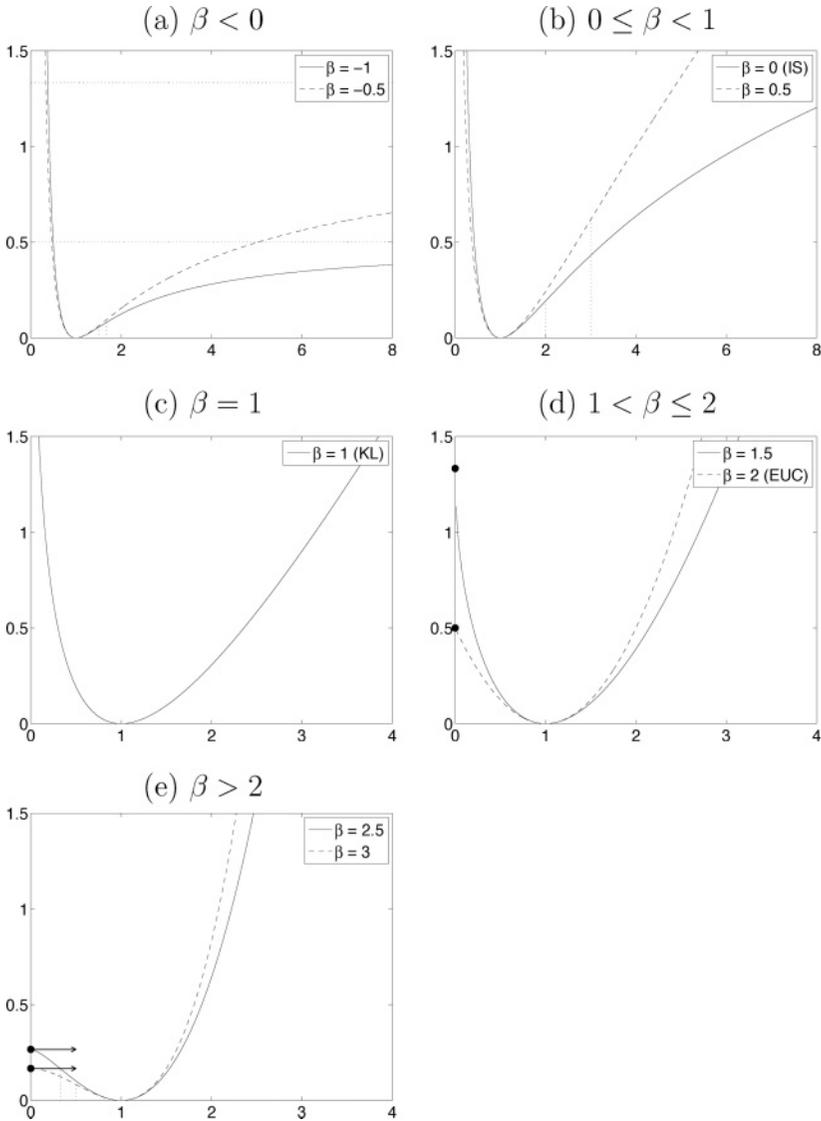


Figure 1: β -divergence $d_\beta(x|y)$ as a function of y (with $x = 1$). Panels a–e illustrate the regimes of the β -divergence for its five characteristic ranges of values of β . The divergence is convex for $1 \leq \beta \leq 2$, as seen in panels c and d. In the other panels, the inflection points are indicated with vertical dotted lines. For $\beta < 0$, the divergence possesses horizontal asymptotes of coordinate $x^\beta / (\beta(\beta - 1))$ as $y \rightarrow \infty$. For $\beta > 1$, the divergence takes finite value $x^\beta / (\beta(\beta - 1))$ at $y = 0$, where the derivative is zero for $\beta > 2$.

Table 1: Example of Differentiable Convex-Concave-Constant Decomposition of the β -Divergence under the Form 2.3.

	$\tilde{d}(x y)$	$\tilde{d}'(x y)$	$\hat{d}(x y)$	$\hat{d}'(x y)$	$\bar{d}(x)$
$\beta < 1$ and $\beta \neq 0$	$-\frac{1}{\beta-1}x y^{\beta-1}$	$-x y^{\beta-2}$	$\frac{1}{\beta}y^\beta$	$y^{\beta-1}$	$\frac{1}{\beta(\beta-1)}x^\beta$
$\beta = 0$	$x y^{-1}$	$-x y^{-2}$	$\log y$	y^{-1}	$x(\log x - 1)$
$1 \leq \beta \leq 2$	$d'_\beta(x y)$	$d'_\beta(x y)$	0	0	0
$\beta > 2$	$\frac{1}{\beta}y^\beta$	$y^{\beta-1}$	$-\frac{1}{\beta-1}x y^{\beta-1}$	$-x y^{\beta-2}$	$\frac{1}{\beta(\beta-1)}x^\beta$

Factorizations with small, positive values of β are relevant to decomposition of audio spectra, which typically exhibit exponential power decrease along frequency f and also usually comprise low-power transient components such as note attacks, together with higher-power components such as tonal parts of sustained notes. For example, Févotte et al. (2009) present the results of the decomposition of a piano power spectrogram with IS-NMF and show that components corresponding to very low residual noise and hammer hits on the strings are extracted with great accuracy, while these components are either ignored or severely degraded when using Euclidean or KL divergences. Similarly, the value $\beta = 0.5$ is advocated by FitzGerald et al. (2009) and Hennequin et al. (2010) and has been shown to give optimal results in music transcription based on NMF of the magnitude spectrogram by Vincent et al. (2010).

The β -divergence belongs to the family of Bregman divergences. For $\beta \notin \{0, 1\}$, a suitable Bregman generating function is $\phi(y) = y^\beta / (\beta(\beta - 1))$, as noted by Févotte and Cemgil (2009). This function, however, cannot generate the IS and KL divergences by continuity when β tends to 0 or 1. The latter divergences may nonetheless be generated separately, using the functions $\phi(y) = -\log y$ and $\phi(y) = y \log y$, respectively. Cichocki and Amari (2010) give a general Bregman generating function of the β -divergence, continuously defined for all $\beta \in \mathbb{R}$, in the form of $\phi_{\beta \neq 0,1}(y) = (y^\beta - \beta y + \beta - 1) / (\beta(\beta - 1))$, $\phi_{\beta=0}(y) = y - \log y - 1$ and $\phi_{\beta=1}(y) = y \log y - y + 1$. NMF with Bregman divergences has been considered by Dhillon and Sra (2005), where the lack of results about the monotonicity of multiplicative algorithms in general has been noted.¹ This letter fills this gap for the specific case of β -divergence.

¹More precisely, Dhillon and Sra (2005) give proofs of monotonicity for the “reverse” problem of minimizing $D(\mathbf{WH}|\mathbf{V})$ instead of $D(\mathbf{V}|\mathbf{WH})$, while pointing out that monotonicity of multiplicative algorithms based on the heuristic 1.4 for the latter problem is, however, observed in practice.

2.2 Task.

2.2.1 *Core Optimization Problem.* Like most algorithms in the literature to date, the NMF algorithms we describe in this letter sequentially update \mathbf{H} given \mathbf{W} and then \mathbf{W} given \mathbf{H} . These two steps are essentially the same, by symmetry of the factorization ($\mathbf{V} \approx \mathbf{W}\mathbf{H}$ is equivalent to $\mathbf{V}^T \approx \mathbf{H}^T\mathbf{W}^T$ and the roles of \mathbf{W} and \mathbf{H} are simply exchanged) and because we are not making any assumption on the relative values of F and N . Hence, we may concentrate on solving the following subproblem,

$$\min_{\mathbf{H}} C(\mathbf{H}) \stackrel{\text{def}}{=} D(\mathbf{V}|\mathbf{W}\mathbf{H}) \text{ subject to } \mathbf{H} \geq \mathbf{0}, \tag{2.5}$$

with fixed \mathbf{W} and where in the rest of the letter $D(\mathbf{V}|\mathbf{W}\mathbf{H})$ is, as of equation 1.3, with $d(x|y) = d_\beta(x|y)$. The criterion function $C(\mathbf{H})$ separates into $\sum_n D(\mathbf{v}_n|\mathbf{W}\mathbf{h}_n)$, where \mathbf{v}_n and \mathbf{h}_n are the n th column of \mathbf{V} and \mathbf{H} , respectively, so that we are essentially left with solving the problem,

$$\min_{\mathbf{h}} C(\mathbf{h}) = D(\mathbf{v}|\mathbf{W}\mathbf{h}) \text{ subject to } \mathbf{h} \geq \mathbf{0}, \tag{2.6}$$

where $\mathbf{v} \in \mathbb{R}_+^F$, $\mathbf{W} \in \mathbb{R}_+^{F \times K}$ and $\mathbf{h} \in \mathbb{R}_+^K$.

2.2.2 *KKT Necessary Conditions.* An admissible solution \mathbf{h}^* to problem 2.6 must satisfy the Karush-Kuhn-Tucker (KKT) first-order optimality conditions,

$$\nabla_{\mathbf{h}} C(\mathbf{h}^*) \cdot \mathbf{h}^* = \mathbf{0}, \tag{2.7}$$

$$\nabla_{\mathbf{h}} C(\mathbf{h}^*) \geq \mathbf{0}, \tag{2.8}$$

$$\mathbf{h}^* \geq \mathbf{0}, \tag{2.9}$$

where the dot notation \cdot denotes entrywise operations (here term-to-term multiplication) and $\nabla_{\mathbf{h}} C(\mathbf{h})$ denotes the gradient of $C(\mathbf{h})$, given by

$$\nabla_{\mathbf{h}} C(\mathbf{h}) = \mathbf{W}^T [d'(v_f|[\mathbf{W}\mathbf{h}]_f)]_f \tag{2.10}$$

$$= \mathbf{W}^T [(\mathbf{W}\mathbf{h})^{(\beta-2)}(\mathbf{W}\mathbf{h} - \mathbf{v})], \tag{2.11}$$

where the notation $[x_f]_f$ refers to the column vector $[x_1, \dots, x_F]^T$. The KKT conditions 2.7 to 2.9 can be summarized as

$$\min\{\mathbf{h}^*, \nabla_{\mathbf{h}} C(\mathbf{h}^*)\} = \mathbf{0}_K, \tag{2.12}$$

where the min operator is entrywise and $\mathbf{0}_K$ is a null vector of dimension K .

2.2.3 Algorithms. In the following, we will say that an algorithm is monotone if it produces a sequence of iterates $\{\mathbf{h}^{(i)}\}_{i \geq 0}$, such that $C(\mathbf{h}^{(i+1)}) \leq C(\mathbf{h}^{(i)})$ for all $i \geq 0$. An algorithm is said to be convergent if it produces a sequence of iterates $\{\mathbf{h}^{(i)}\}_{i \geq 0}$ that converges to a limit point \mathbf{h}^* satisfying the KKT conditions 2.7 to 2.9. Monotonicity does not imply convergence in general, nor is it necessary to convergence.

3 An Auxiliary Function for β -NMF

In this section we define the concept of auxiliary function and then exhibit a separable auxiliary function for the β -NMF problem.

3.1 Definition of Auxiliary Function.

Definition 1. (*auxiliary function*). The $\mathbb{R}_+^K \times \mathbb{R}_+^K \rightarrow \mathbb{R}_+$ mapping $G(\mathbf{h}|\tilde{\mathbf{h}})$ is said to be an auxiliary function to $C(\mathbf{h})$ if and only if

- $\forall \mathbf{h} \in \mathbb{R}_+^K, C(\mathbf{h}) = G(\mathbf{h}|\mathbf{h})$
- $\forall (\mathbf{h}, \tilde{\mathbf{h}}) \in \mathbb{R}_+^K \times \mathbb{R}_+^K, C(\mathbf{h}) \leq G(\mathbf{h}|\tilde{\mathbf{h}})$.

In other words an auxiliary function $G(\mathbf{h}|\tilde{\mathbf{h}})$ is a majorizing function (or upper bound) of $C(\mathbf{h})$, which is tight for $\mathbf{h} = \tilde{\mathbf{h}}$. The optimization of $C(\mathbf{h})$ can be replaced by iterative optimization of $G(\mathbf{h}|\tilde{\mathbf{h}})$. Indeed, any iterate $\mathbf{h}^{(i+1)}$ satisfying

$$G(\mathbf{h}^{(i+1)}|\mathbf{h}^{(i)}) \leq G(\mathbf{h}^{(i)}|\mathbf{h}^{(i)}) \tag{3.1}$$

satisfies $C(\mathbf{h}^{(i+1)}) \leq C(\mathbf{h}^{(i)})$, because we have

$$C(\mathbf{h}^{(i+1)}) \leq G(\mathbf{h}^{(i+1)}|\mathbf{h}^{(i)}) \leq G(\mathbf{h}^{(i)}|\mathbf{h}^{(i)}) = C(\mathbf{h}^{(i)}). \tag{3.2}$$

The iterate $\mathbf{h}^{(i+1)}$ is typically chosen as

$$\mathbf{h}^{(i+1)} = \underset{\mathbf{h} \geq 0}{\operatorname{arg\,min}} G(\mathbf{h}|\mathbf{h}^{(i)}), \tag{3.3}$$

which forms the basis of MM algorithms (see, e.g., Hunter & Lange, 2004, for a tutorial). However, any other iterate $\mathbf{h}^{(i+1)}$ satisfying equation 3.1 produces a monotone algorithm. As such, Figure 2 illustrates the three updates strategies that will be developed in this letter.

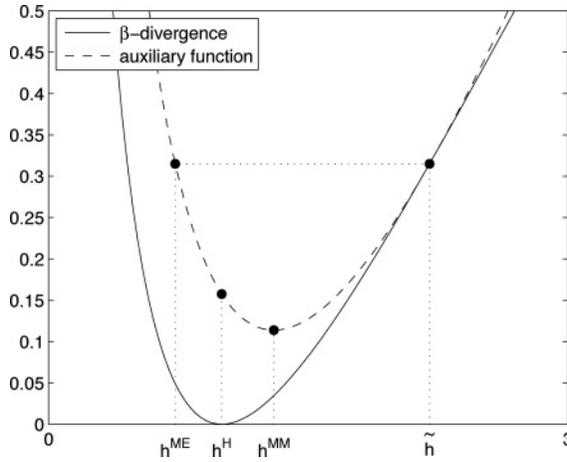


Figure 2: The β -divergence $d_\beta(x|y)$ for $\beta = 0.5$ (with $x = 1$) and its auxiliary function in dimension 1 (with $\tilde{h} = 2.2$). The MM update h^{MM} corresponds to the minimum of the auxiliary function (see section 4.1). The heuristic update h^H given by equation 1.4 is discussed in section 4.2 (the heuristic update minimizes the criterion function in the simple one-dimensional case, but this is not true in larger dimensions). The ME update h^{ME} consists of selecting the next update on the other side of the “valley” defined by the auxiliary function, opposite the point defined by the current solution \tilde{h} (see section 4.3).

3.2 Separable Auxiliary Function for β -NMF. In this section, we construct an auxiliary function to $C(\mathbf{h})$ for the specific case of the β -divergence. Our approach follows the one of Cao et al. (1999) for IS divergence and consists of majorizing the convex part of the criterion using Jensen’s inequality and majorizing the concave part by its tangent, as detailed in the proof of the following theorem. Here and henceforth, we denote $\mathbf{W}\tilde{\mathbf{h}}$ by $\tilde{\mathbf{v}}$, with entries $[\mathbf{W}\tilde{\mathbf{h}}]_f = \tilde{v}_f$.

Theorem 1 (auxiliary function for β -NMF). Let $\tilde{\mathbf{h}}$ be such that

- (i) $\forall f, \tilde{v}_f > 0,$
- (ii) $\forall k, \tilde{h}_k > 0.$

Then the function $G(\mathbf{h}|\tilde{\mathbf{h}})$ defined by

$$\begin{aligned}
 G(\mathbf{h}|\tilde{\mathbf{h}}) = & \sum_f \left[\sum_k \frac{w_{fk}\tilde{h}_k}{\tilde{v}_f} \tilde{d} \left(v_f | \tilde{v}_f \frac{h_k}{\tilde{h}_k} \right) \right] \\
 & + \left[\widehat{d}'(v_f|\tilde{v}_f) \sum_k w_{fk}(h_k - \tilde{h}_k) + \widehat{d}(v_f|\tilde{v}_f) \right] + \bar{d}(v_f) \quad (3.4)
 \end{aligned}$$

is an auxiliary function to $C(\mathbf{h}) = \sum_f d(v_f | [\mathbf{W}\mathbf{h}]_f)$, where $\check{d}(x|y) + \widehat{d}(x|y) + \bar{d}(x)$ is any differentiable convex-concave-constant decomposition of the β -divergence, such as the one defined in Table 1.

Proof. The condition $G(\mathbf{h}|\mathbf{h}) = C(\mathbf{h})$ is trivially met. The criterion $C(\mathbf{h})$ may be written as

$$C(\mathbf{h}) = \sum_f C_f(\mathbf{h}), \quad (3.5)$$

where $C_f(\mathbf{h}) \stackrel{\text{def}}{=} d(v_f | [\mathbf{W}\mathbf{h}]_f)$. We prove $C(\mathbf{h}) \leq G(\mathbf{h}|\tilde{\mathbf{h}})$ by constructing an auxiliary function to each part $C_f(\mathbf{h})$ of the criterion and, more precisely, by treating the convex and concave part separately. Let us define $\check{C}_f(\mathbf{h}) \stackrel{\text{def}}{=} \check{d}(v_f | [\mathbf{W}\mathbf{h}]_f)$ and $\widehat{C}_f(\mathbf{h}) \stackrel{\text{def}}{=} \widehat{d}(v_f | [\mathbf{W}\mathbf{h}]_f)$, so that we can write

$$C_f(\mathbf{h}) = \check{C}_f(\mathbf{h}) + \widehat{C}_f(\mathbf{h}) + \bar{d}(v_f). \quad (3.6)$$

Convex part: We first prove that

$$\check{G}_f(\mathbf{h}|\tilde{\mathbf{h}}) = \sum_k \frac{w_{fk} \tilde{h}_k}{\check{v}_f} \check{d}\left(v_f | \check{v}_f \frac{h_k}{\tilde{h}_k}\right) \quad (3.7)$$

is an auxiliary function to $\check{C}_f(\mathbf{h})$. The condition $\check{G}_f(\mathbf{h}|\mathbf{h}) = \check{C}_f(\mathbf{h})$ is trivially met. The condition $\check{G}_f(\mathbf{h}|\tilde{\mathbf{h}}) \geq \check{C}_f(\tilde{\mathbf{h}})$ is proven as follows. Let \mathcal{K} be the set of indices k such that $w_{fk} \neq 0$. Define $\forall k \in \mathcal{K}$,

$$\tilde{\lambda}_{fk} = \frac{w_{fk} \tilde{h}_k}{\check{v}_f} = \frac{w_{fk} \tilde{h}_k}{\sum_{\ell \in \mathcal{K}} w_{f\ell} \tilde{h}_\ell}. \quad (3.8)$$

We have $\sum_{k \in \mathcal{K}} \tilde{\lambda}_{fk} = 1$ and

$$\check{G}_f(\mathbf{h}|\tilde{\mathbf{h}}) = \sum_{k \in \mathcal{K}} \tilde{\lambda}_{fk} \check{d}\left(v_f | \frac{w_{fk} h_k}{\tilde{\lambda}_{fk}}\right) \quad (3.9)$$

$$\geq \check{d}\left(v_f | \sum_{k \in \mathcal{K}} \tilde{\lambda}_{fk} \frac{w_{fk} h_k}{\tilde{\lambda}_{fk}}\right) \quad (3.10)$$

$$= \check{d}\left(v_f | \sum_{k=1}^K w_{fk} h_k\right) \quad (3.11)$$

$$= \widetilde{C}_f(\mathbf{h}), \tag{3.12}$$

where we used Jensen’s inequality, by convexity of $\widetilde{d}(x|y)$.

Concave part: An auxiliary function $\widehat{G}_f(\mathbf{h}|\tilde{\mathbf{h}})$ to the concave part $\widehat{C}_f(\mathbf{h})$ can be taken as the first-order Taylor approximation to $\widehat{C}_f(\mathbf{h})$ in the vicinity of $\tilde{\mathbf{h}}$:

$$\widehat{G}_f(\mathbf{h}|\tilde{\mathbf{h}}) = \widehat{C}_f(\tilde{\mathbf{h}}) + \nabla^T \widehat{C}_f(\tilde{\mathbf{h}}) (\mathbf{h} - \tilde{\mathbf{h}}). \tag{3.13}$$

The function satisfies $\widehat{G}_f(\mathbf{h}|\mathbf{h}) = \widehat{C}_f(\mathbf{h})$ by construction and $\widehat{G}_f(\mathbf{h}|\tilde{\mathbf{h}}) \geq \widehat{C}_f(\mathbf{h})$ by concavity of $\widehat{C}_f(\mathbf{h})$, using the property that the tangent to any point is an upper bound of a concave function.² Using

$$\nabla_{h_k} \widehat{C}_f(\mathbf{h}) = w_{fk} \widehat{d}'(v_f | [\mathbf{W}\mathbf{h}]_f), \tag{3.14}$$

we give the explicit form for $\widehat{G}_f(\mathbf{h}|\tilde{\mathbf{h}})$ by

$$\widehat{G}_f(\mathbf{h}|\tilde{\mathbf{h}}) = \widehat{d}(v_f | \tilde{v}_f) + \widehat{d}'(v_f | \tilde{v}_f) \sum_k w_{fk} (h_k - \tilde{h}_k). \tag{3.15}$$

In the end, a suitable auxiliary function $G(\mathbf{h}|\tilde{\mathbf{h}})$ to $C(\mathbf{h})$ is obtained by summing up the auxiliary functions constructed for each individual part of the criterion,

$$G(\mathbf{h}|\tilde{\mathbf{h}}) = \sum_f \left(\widetilde{G}_f(\mathbf{h}|\tilde{\mathbf{h}}) + \widehat{G}_f(\mathbf{h}|\tilde{\mathbf{h}}) + \widehat{d}(v_f) \right), \tag{3.16}$$

which leads to equation 3.4.

3.2.1 Properties of the Auxiliary Function. $G(\mathbf{h}|\tilde{\mathbf{h}})$ is by construction separable in functions of the individual coefficients h_k of \mathbf{h} , which allows decoupling the optimization. It is convenient to rewrite the auxiliary function as such in order to derive some of the algorithms of section 4. We may write

$$G(\mathbf{h}|\tilde{\mathbf{h}}) = \sum_k G_k(h_k|\tilde{h}_k) + cst, \tag{3.17}$$

² $\widehat{C}_f(\mathbf{h}) = \widehat{d}(v_f | [\mathbf{W}\mathbf{h}]_f)$ is concave as the composition of a concave function and a linear function.

where cst is a constant w.r.t \mathbf{h} and

$$G_k(h_k|\tilde{\mathbf{h}}) \stackrel{\text{def}}{=} \tilde{h}_k \left[\sum_f \frac{w_{fk}}{\tilde{v}_f} \tilde{d} \left(v_f | \tilde{v}_f \frac{h_k}{\tilde{h}_k} \right) \right] + h_k \left[\sum_f w_{fk} \hat{d}'(v_f | \tilde{v}_f) \right]. \tag{3.18}$$

The gradient of the auxiliary function is given by

$$\nabla_{h_k} G(\mathbf{h}|\tilde{\mathbf{h}}) = \sum_f w_{fk} \left[\tilde{d}' \left(v_f | \tilde{v}_f \frac{h_k}{\tilde{h}_k} \right) + \hat{d}'(v_f | \tilde{v}_f) \right]. \tag{3.19}$$

Thanks to the separability of the auxiliary function into its variables, the Hessian matrix is diagonal with

$$\nabla_{h_k}^2 G(\mathbf{h}|\tilde{\mathbf{h}}) = \sum_f \tilde{v}_f \frac{w_{fk}}{\tilde{h}_k} \tilde{d}'' \left(v_f | \tilde{v}_f \frac{h_k}{\tilde{h}_k} \right). \tag{3.20}$$

By convexity of $\tilde{d}(x|y)$, we have $\tilde{d}''(x|y) \geq 0$, which implies positive definiteness of the Hessian matrix and, hence, convexity of the auxiliary function $G(\mathbf{h}|\tilde{\mathbf{h}})$ (convexity more simply derives from the fact that the auxiliary function is built as a sum of convex functions).

3.2.2 Connections with Other Works. The construction of $G(\mathbf{h}|\tilde{\mathbf{h}})$ employs standard mathematical tools (Jensen’s inequality, Taylor approximation) that are well known from the MM literature (see, e.g., Hunter & Lange, 2004). For $\beta \in [1, 2]$, $G(\mathbf{h}|\tilde{\mathbf{h}})$ coincides with the auxiliary function built by Kompass (2007), who proposed a generalization of the auxiliary functions proposed by Lee and Seung (2001) for the Euclidean distance ($\beta = 2$) and the generalized KL divergence ($\beta = 1$). For $\beta = 0$ (IS divergence), $G(\mathbf{h}|\tilde{\mathbf{h}})$ coincides with the auxiliary function proposed by Cao et al. (1999). It is worth recalling that in the algorithms proposed by Lee and Seung (2001) the update of \mathbf{W} given \mathbf{H} or \mathbf{H} given \mathbf{W} are instances of well-known algorithms for image restoration (for which \mathbf{W} acts as a fixed, known blurring matrix and \mathbf{H} is a vectorized image to be reconstructed). These algorithms are the iterative space reconstructing algorithm (ISRA) of Daube-Witherspoon and Muehlehner (1986) and the Richardson-Lucy (RL) algorithm of Richardson (1972) and Lucy (1974), which perform nonnegative linear regression with the Euclidean distance and KL divergence, respectively. The ISRA and RL algorithms are shown to be MM algorithms by De Pierro (1993). Similarly, the algorithms proposed by Cao et al. (1999) for nonnegative linear

Table 2: Exponent in the Multiplicative Updates Given by the MM Algorithm.

	$\beta < 1$	$1 \leq \beta \leq 2$	$\beta > 2$
$\gamma(\beta)$	$\frac{1}{2-\beta}$	1	$\frac{1}{\beta-1}$

regression with the IS divergence were designed in the image restoration setting. Finally, let us mention that an auxiliary function based on Jensen's inequality for NMF with the α -divergence (which is always convex w.r.t. to its second argument) is given by Cichocki, Lee, Kim, and Choi (2008).

4 Algorithms for β -NMF

In this section, we describe algorithms for β -NMF based on the auxiliary function constructed in section 3. In the following, $\tilde{\mathbf{h}}$ should be understood as the current iterate $\mathbf{h}^{(i)}$, and we are seeking to obtain $\mathbf{h}^{(i+1)}$ such that equation 3.1 is satisfied.

4.1 Maximization-Minimization Algorithm. An MM algorithm can be derived by minimizing the auxiliary function $G(\mathbf{h}|\tilde{\mathbf{h}})$ w.r.t. \mathbf{h} . Given the convexity and the separability of the auxiliary function, the optimum is obtained by canceling the gradient given by equation 3.19. This is trivially done and leads to the following update,

$$h_k^{\text{MM}} = \tilde{h}_k \left(\frac{\sum_f w_{fk} v_f \tilde{v}_f^{\beta-2}}{\sum_f w_{fk} \tilde{v}_f^{\beta-1}} \right)^{\gamma(\beta)}, \quad (4.1)$$

where $\gamma(\beta)$ is given in Table 2. Note that $\gamma(\beta) \leq 1, \forall \beta$. As suggested in section 1, the gradient of the criterion may be written as the difference of two nonnegative functions such that

$$\nabla_{h_k} C(\tilde{\mathbf{h}}) = \nabla_{h_k}^+ C(\tilde{\mathbf{h}}) - \nabla_{h_k}^- C(\tilde{\mathbf{h}}), \quad (4.2)$$

$$\nabla_{h_k}^+ C(\tilde{\mathbf{h}}) = \sum_f w_{fk} \tilde{v}_f^{\beta-1}, \quad (4.3)$$

$$\nabla_{h_k}^- C(\tilde{\mathbf{h}}) = \sum_f w_{fk} v_f \tilde{v}_f^{\beta-2}, \quad (4.4)$$

so that the update, equation 4.1, can be rewritten in the more interpretable form

$$h_k^{\text{MM}} = \tilde{h}_k \left(\frac{\nabla_{h_k}^- C(\tilde{\mathbf{h}})}{\nabla_{h_k}^+ C(\tilde{\mathbf{h}})} \right)^{\gamma(\beta)}. \tag{4.5}$$

The conclusion is thus that the MM algorithm leads to multiplicative updates, but the latter differ from the “usual ones,” obtained by setting $\gamma(\beta) = 1$ for all β and derived heuristically by Cichocki et al. (2006) through gradient descent with adaptative step or by Févotte et al. (2009) by splitting the gradient into two nonnegative functions as discussed above and in section 1. The MM update differs from the heuristic update by the exponent $\gamma(\beta)$, which is not equal to 1 for $\beta \notin [1, 2]$.

4.2 Heuristic Algorithm. This section discusses the properties of the heuristic update proposed by Cichocki et al. (2006) and Févotte et al. (2009) and defined for all β by

$$h_k^{\text{H}} = \tilde{h}_k \left(\frac{\sum_f w_{fk} v_f \tilde{v}_f^{\beta-2}}{\sum_f w_{fk} \tilde{v}_f^{\beta-1}} \right). \tag{4.6}$$

Very few mathematical results exist for the heuristic update when β falls outside $[1, 2]$, that is, when the β -divergence $d_\beta(x|y)$ is not convex. In such a case, the heuristic update can be erroneously interpreted as an MM algorithm by wrongly applying Jensen’s inequality to $C(\mathbf{h})$. Yet in the particular case $\beta = 0$, it holds that each heuristic update produces a decrease of $C(\mathbf{h})$ (Cao et al., 1999). One objective of this section is to extend this result to all values of β between 0 and 1.

Let us first introduce a scalar auxiliary function $g(y|\tilde{y}; x)$:

$$\forall y, \tilde{y}, x > 0, \quad g(y|\tilde{y}; x) = \check{d}(x|y) + \hat{d}(x|\tilde{y}) + (y - \tilde{y})\hat{d}'(x|\tilde{y}) + \bar{d}(x), \tag{4.7}$$

where $\check{d}(x|y)$, $\hat{d}(x|y)$, and $\bar{d}(x|y)$ are defined in Table 1. By immediate application of theorem 1 to the scalar case, $g(y|\tilde{y}; x)$ is an auxiliary function to $d(x|y)$. In particular, $g(\tilde{y}|\tilde{y}; x) = d(x|\tilde{y})$. Then we have the following preliminary result:

Lemma 1. For all $\beta \in \mathbb{R}$,

$$G_k(h_k|\tilde{\mathbf{h}}) = \frac{1}{\tilde{h}_k^{\beta-1}} \left(\sum_f w_{fk} \tilde{v}_f^{\beta-1} \right) g(h_k|\tilde{h}_k; h_k^{\text{H}}) + \text{cst}. \tag{4.8}$$

Proof. For each of the four possible expressions of (\widehat{d}, \check{d}) given in Table 1, the validity of equation 4.8 can be checked straightforwardly by direct verification.

As already mentioned in section 3.1, the MM update, equation 3.3, is not the only way of taking advantage of the auxiliary function $G(\mathbf{h}|\tilde{\mathbf{h}})$ to obtain a decrease of $C(\mathbf{h})$: any update satisfying equation 3.1 also ensures that $C(\mathbf{h})$ does not increase. This is a key remark to understand the behavior of the heuristic algorithm for $\beta \in (0, 1)$, given the following property:

Theorem 2. For all $\beta \in (0, 1)$, and all $\tilde{\mathbf{h}}$ such that conditions (i) and (ii) of theorem 1 hold, the heuristic algorithm produces nonincreasing values of $C(\mathbf{h})$, according to the following inequality:

$$G(\mathbf{h}^H|\tilde{\mathbf{h}}) \leq G(\tilde{\mathbf{h}}|\tilde{\mathbf{h}}). \tag{4.9}$$

Proof. For all $\beta \in (0, 1)$, straightforward calculations yield

$$g(\tilde{y}|\tilde{y}; x) - g(x|\tilde{y}; x) = \check{d}(x|\tilde{y}) - \check{d}(x|x) - (x - \tilde{y})\widehat{d}'(x|\tilde{y}) \tag{4.10}$$

$$= \frac{1}{1 - \beta} \tilde{y}^\beta (1 - \beta + \beta\theta - \theta^\beta), \tag{4.11}$$

where $\theta = x/\tilde{y}$. Since $f(\theta) = \theta^\beta$ is a concave function of θ , we have $f(\theta) \leq f(1) + (\theta - 1)f'(1)$, which also reads $\theta^\beta \leq 1 + (\theta - 1)\beta$. Hence, $g(\tilde{y}|\tilde{y}; x) - g(x|\tilde{y}; x) \geq 0$ for all x, \tilde{y} . The latter inequality implies $\forall k, g(h_k^H|\tilde{h}_k, h_k^H) \leq g(\tilde{h}_k|\tilde{h}_k, h_k^H)$, so that we have $G_k(h_k^H|\tilde{\mathbf{h}}) \leq G_k(\tilde{h}_k|\tilde{\mathbf{h}})$ according to equation 4.8, which leads to the result by summation over k .

Cao et al. (1999) show that inequality 4.9 becomes an equality in the case $\beta = 0$, so that each heuristic update yields $G(\mathbf{h}^H|\tilde{\mathbf{h}}) = G(\tilde{\mathbf{h}}|\tilde{\mathbf{h}})$. In this particular case, the heuristic algorithm can be called a majorization-equalization (ME) algorithm, a class of algorithms described in next section. For values of β outside the range $[0, 2]$, inequality 4.9 no longer holds.³ Of course, this does not mean that the heuristic updates produce increasing values of $C(\mathbf{h})$. On the contrary, numerical simulations tend to indicate that they always produce nonincreasing values of $C(\mathbf{h})$, but proving this is still an open issue for $\beta \notin [0, 2]$. Compared to MM updates, heuristic updates produce larger or equal steps for all β , since it can trivially be shown that

$$\forall k, |h_k^H - \tilde{h}_k| \geq |h_k^{MM} - \tilde{h}_k|. \tag{4.12}$$

³Indeed, we can prove that the reversed inequality holds for all $\beta < 0$, while no systematic result is known for $\beta > 2$.

For $\beta \notin [1, 2]$, numerical simulations indicate that the heuristic algorithm is faster than the MM algorithm (and we recall that the two algorithms coincide for $\beta \in [1, 2]$). Given equation 4.12, skipping from the latter to the former has an effect comparable to that of overrelaxation: on average, stretching the steps allows reducing their number to reach convergence. This will be discussed in more detail in section 4.4.

In order to produce even larger steps for $\beta \in [0, 2]$ and yet nonincreasing values of $C(\mathbf{h})$, the following section explores the concept of majorization-equalization.

4.3 Majorization-Equalization Algorithm. Let us introduce the general notion of ME update by the fact that the new iterate \mathbf{h}^{ME} fulfills

$$G(\mathbf{h}^{\text{ME}}|\tilde{\mathbf{h}}) = G(\tilde{\mathbf{h}}|\tilde{\mathbf{h}}). \tag{4.13}$$

Equation 4.13 actually defines a level set rather than a single point. Let us concentrate on the following more constrained and manageable condition, given the separability of $G(\mathbf{h}|\tilde{\mathbf{h}})$:

$$\forall k, \quad G_k(h_k^{\text{ME}}|\tilde{\mathbf{h}}) = G_k(\tilde{h}_k|\tilde{\mathbf{h}}).$$

Given equation 4.8, this amounts to solving the following equation for y , for any $\tilde{y}, x > 0$:

$$g(y|\tilde{y}; x) = g(\tilde{y}|\tilde{y}; x). \tag{4.14}$$

Since $g(y|\tilde{y}; x)$ is strictly convex w.r.t. y , equation 4.14 has no more than two solutions, one of them being \tilde{y} . By construction, the selection of the other solution (provided that it exists) will provide ME steps that are larger than MM updates—

$$\forall k, \quad |h_k^{\text{ME}} - \tilde{h}_k| \geq |h_k^{\text{MM}} - \tilde{h}_k|, \tag{4.15}$$

as illustrated by Figure 2. To go further on the determination of this solution, a case-by-case analysis must be performed, depending on the range of β .

Case 1. $\beta \in [0, 1)$ In that case, we have

$$g(y|\tilde{y}; x) = \frac{1}{1-\beta} x y^{\beta-1} + y \tilde{y}^{\beta-1} + cst. \tag{4.16}$$

Let us remark that

$$\forall \tilde{y}, x > 0, \quad \lim_{y \rightarrow 0} g(y|\tilde{y}; x) = \lim_{y \rightarrow \infty} g(y|\tilde{y}; x) = \infty, \tag{4.17}$$

Table 3: Values of β for Which ME Updates are Closed Form, by Root Extraction of Polynomials of Degree d .

$\beta \leq 0$	$0 \leq \beta \leq 1$	$1 \leq \beta \leq 2$	$\beta \geq 2$	d
0	0	2	2	1
-1	1/2	3/2	3	2
-2	2/3	4/3	4	3
-3	3/4	5/4	5	4

so that equation 4.14 always admits two positive solutions (or one double-positive solution if $\tilde{y} = x$), one of the two being $y = \tilde{y}$. The other one is the solution of interest. However, it is not closed form, except for specific values of β (see Table 3). More precisely, when $\beta = 1 - 1/d$ and d is an integer, the solution can be found by solving the following polynomial equation of degree d , for $z = y^{1/d}$:

$$(1 - \beta) \sum_{\ell=1}^d \tilde{z}^{d-\ell} z^\ell - x = 0, \tag{4.18}$$

where $\tilde{z} = \tilde{y}^{1/d}$. Not surprisingly, the simplest case $\beta = 0$ ($d = 1$) leads us to $y = x$, and thus to $h_k^{ME} = h_k^H$. The case $\beta = 0.5$ ($d = 2$) is more interesting. The extraction of the positive root of equation 4.18 then provides the following update formula:

$$h_k^{ME} = \frac{\tilde{h}_k}{4} \left(\sqrt{1 + 8 \frac{h_k^H}{\tilde{h}_k}} - 1 \right)^2. \tag{4.19}$$

This expression does not correspond to a multiplicative update, although it ensures that positivity is maintained.

Case 2. $\beta \in (1, 2]$ In this case, we have

$$g(y|\tilde{y}; x) = \frac{1}{\beta} y^\beta - \frac{1}{\beta - 1} x y^{\beta-1} + cst. \tag{4.20}$$

$g(y|\tilde{y}; x)$ tends toward ∞ for $y \rightarrow \infty$, but it remains finite for $y \rightarrow 0$. As a consequence, equation 4.14 admits the trivial solution $y = \tilde{y}$ only if $g(\tilde{y}|\tilde{y}; x) > g(0|\tilde{y}; x)$, and also the unwanted solution 0 if $g(\tilde{y}|\tilde{y}; x) = g(0|\tilde{y}; x)$. It is only when $g(\tilde{y}|\tilde{y}; x) < g(0|\tilde{y}; x)$ that a positive, nontrivial solution exists. This solution is closed form for specific values of β given in Table 3. They correspond to $\beta = 1 + 1/d$, where d is an integer. Equation 4.14

then amounts to solving the following polynomial equation of degree d , for $z = y^{1/d}$:

$$\sum_{\ell=0}^d \bar{z}^{d-\ell} z^\ell - (d + 1)x = 0, \tag{4.21}$$

with $\bar{z} = \tilde{y}^{1/d}$. The simplest case is $\beta = 2$ ($d = 1$), and the solution is then given by $y = 2x - \tilde{y}$ if $\tilde{y} < 2x$, which yields the overrelaxed update

$$h_k^{\text{ME}} = 2h_k^{\text{MM}} - \tilde{h}_k \tag{4.22}$$

provided that $\tilde{h}_k < 2h_k^{\text{MM}}$. Note that this result more simply stems from the fact that when $\beta = 2$, the auxiliary function is parabolic and thus symmetric with respect to h_k^{MM} . In the case $\beta = 1.5$ ($d = 2$), a positive ME update exists if $\tilde{h}_k < 3h_k^{\text{MM}}$, and it takes the following form:

$$h_k^{\text{ME}} = \frac{\tilde{h}_k}{4} \left(\sqrt{12 \frac{h_k^{\text{MM}}}{\tilde{h}_k} - 3 - 1} \right)^2. \tag{4.23}$$

Because we need an update strategy that is defined everywhere, we propose to rely on a linear mixture between the MM update and a prolonged version of ME, defined as

$$h_k^\theta = \theta h_k^{\text{pME}} + (1 - \theta)h_k^{\text{MM}}, \tag{4.24}$$

where $\theta \in (0, 1)$ and h_k^{pME} prolongs the ME update by 0 when the latter does not exist:

$$h_k^{\text{pME}} = \begin{cases} h_k^{\text{ME}} & \text{if } h_k^{\text{ME}} \text{ is defined} \\ 0 & \text{otherwise.} \end{cases} \tag{4.25}$$

It is mathematically easy to check that h_k^θ fulfills equation 3.1 for all $\theta \in [0, 1]$ and that positivity is maintained for all $\theta \in [0, 1)$. In practice, values of θ near 1 may be favored to produce larger steps.

When $\beta < 0$ or $\beta > 2$, similar analyses can be conducted. In particular, there are specific values of β for which a closed-form expression of ME updates is available according to Table 3.

When $\beta < 0$, ME updates always exist since equations 4.16 and 4.17 still hold. Moreover, they provide nonincreasing values of $C(\mathbf{h})$, while the latter monotonicity property is not yet proved for the heuristic algorithm. However, simulations tend to indicate that the heuristic algorithm is faster

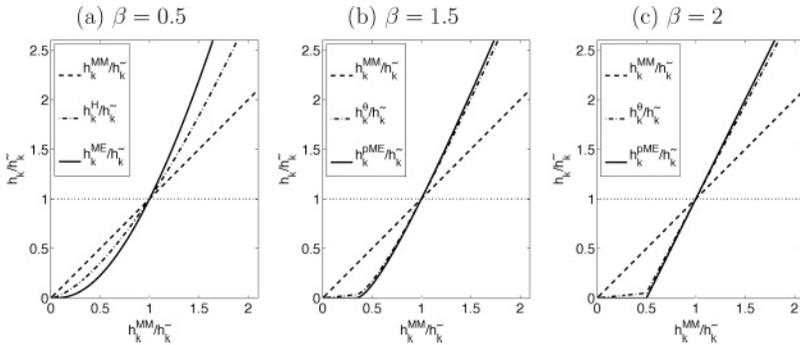


Figure 3: Normalized updates h_k/\tilde{h}_k as functions of h_k^{MM}/\tilde{h}_k ($\theta = 0.9$). The region between the dotted horizontal line and the solid line corresponds to the steps that fulfill equation 3.1. The larger the departure from the horizontal line, the larger the step is.

than the ME algorithm (which itself is faster than the MM algorithm) in the case $\beta < 0$. This is in conformity with the fact that ME steps can then be proved to be smaller than heuristic steps (on the basis of the reversed inequality mentioned in note 3).

When $\beta > 2$, ME updates do not necessarily exist, akin to the case $\beta \in (1, 2]$. When they exist, they provide nonincreasing values of $C(\mathbf{h})$, while the latter is not yet proved for the heuristic formula. However, since this range of β values is not of practical interest, we will not go further into a detailed analysis here.

4.4 Overrelaxation Properties of the Heuristic and ME Updates. The heuristic and ME updates produce larger steps than the MM update, that is, $|h_k^H - \tilde{h}_k|$ and $|h_k^{ME} - \tilde{h}_k|$ are larger than $|h_k^{MM} - \tilde{h}_k|$, for all values of $\beta \in \mathbb{R}$. This is a form of overrelaxation, which will be shown in section 5 to produce faster convergence in practice. The normalized ME (or pME) and heuristic updates studied in previous sections can be written as a function of h_k^{MM}/\tilde{h}_k , such that

$$\frac{h_k}{\tilde{h}_k} = f\left(\frac{h_k^{MM}}{\tilde{h}_k}\right), \tag{4.26}$$

where h_k is either $h_k^H, h_k^{ME}, h_k^{pME}$, or h_k^θ . For the heuristic update, the function is simply given by $f(x) = x^{1/\gamma(\beta)}$. For $\beta \in \{0.5, 1.5, 2\}$, the function f corresponding to the ME/pME update is easily derived from equations 4.19, 4.23, and 4.22. Figure 3 displays the latter functions for the updates studied in this work when $\beta \in \{0.5, 1.5, 2\}$. Overrelaxation appears from the fact that

$h_k < h_k^{\text{MM}}$ whenever $h_k^{\text{MM}} < \tilde{h}_k$ (steps toward left) and $h_k > h_k^{\text{MM}}$ whenever $h_k^{\text{MM}} > \tilde{h}_k$ (steps toward right).

General results about overrelaxation of MM algorithms are given by Salakhutdinov and Roweis (2003), and in particular in the case of NMF. The authors consider the specific case of KL divergence but their study holds for any divergence. They show that in a neighborhood of a stationary point, for any $\eta \in (0, 2)$, relaxed updates h_k^{R} of the form

$$\frac{h_k^{\text{R}}}{\tilde{h}_k} = \left(\frac{h_k^{\text{MM}}}{\tilde{h}_k} \right)^\eta \quad (4.27)$$

will converge to the same point as h_k^{MM} , with a different, possibly better, rate of convergence. In particular, the optimal learning rate η , providing the largest rate of convergence, can be computed from the eigenvalues of the Jacobian, at convergence, of the mapping that relates h_k^{MM} at iteration (i) to h_k^{MM} at iteration ($i + 1$). The optimal learning rate is shown to be always greater than or equal to 1. A similar result was recently obtained by Badeau, Bertin, and Vincent (2010). However, these results do not translate into a practical algorithm, because the latter relaxation property holds only locally, and the computation of the optimal learning rate requires the stationary point to be known. As such, Salakhutdinov and Roweis (2003) propose an adaptive scheme that incrementally proposes values of η greater than 1 at each iteration, and backtrack to $\eta = 1$ when the criterion ceases decreasing.

Our results show that for $\beta \in (0, 1)$, the learning rate $\eta = 1/\gamma(\beta) = 2 - \beta$, corresponding to the heuristic update, ensures descent of the criterion everywhere. The results of Salakhutdinov and Roweis (2003) indicate that the learning rate can be increased to $\eta = 2$ when the algorithm approaches the solution. Note that in the neighborhood of the solution, the Taylor approximation $f(x) \approx f(1) + f'(1)(x - 1)$ applied to $f(x) = x^\eta$ implies that

$$(h_k^{\text{R}} - \tilde{h}_k) \approx \eta(h_k^{\text{MM}} - \tilde{h}_k). \quad (4.28)$$

A similar approximation carried out with the ME/pME updates defined by equations 4.19, and 4.23 for $\beta \in \{0.5, 1.5\}$ reveals that in these two cases, $f'(1) = 2$ (and by construction $f(1) = 1$), so that in a neighborhood of the solution, we have

$$(h_k^{\text{ME}} - \tilde{h}_k) \approx 2(h_k^{\text{MM}} - \tilde{h}_k). \quad (4.29)$$

This means that the ME algorithms produce the largest admissible learning rate $\eta = 2$ in the neighborhood of the solution, while avoiding adaptation of the learning rate so as to ensure monotonicity of the criterion. This result holds everywhere for $\beta = 2$, see equation 4.22, by symmetry of the auxiliary function w.r.t to h_k^{MM} . The interested reader may also refer to Lantéri, Roche,

Cuevas, and Aime (2001); and Lanteri, Theys, Richard, and Févotte (2010), for relaxation of multiplicative algorithms using adaptative learning rates computed through line search.

4.5 Implementation and Complexity of the Algorithms. As seen in section 4.4, the update rules of all the studied algorithms can be expressed as functions of the ratio $\nabla_{h_k}^- C(\hat{\mathbf{h}})/\nabla_{h_k}^+ C(\hat{\mathbf{h}})$, which dominates the algorithmic complexities. Fortunately, the latter ratio takes a simple matrix form that leads to efficient implementations. As such, getting back to the original factorization problem, the heuristic update 4.6 for factors \mathbf{H} and \mathbf{W} can conveniently be expressed in the following matrix form,

$$\mathbf{H} \leftarrow \mathbf{H} \cdot \frac{\mathbf{W}^T [(\mathbf{WH})^{(\beta-2)} \cdot \mathbf{V}]}{\mathbf{W}^T [\mathbf{WH}]^{(\beta-1)}}, \quad (4.30)$$

$$\mathbf{W} \leftarrow \mathbf{W} \cdot \frac{[(\mathbf{WH})^{(\beta-2)} \cdot \mathbf{V}] \mathbf{H}^T}{[\mathbf{WH}]^{(\beta-1)} \mathbf{H}^T}, \quad (4.31)$$

where the division \cdot/ \cdot is here taken entrywise. The MM update simply involves bringing the corrective ratio to the power $\gamma(\beta)$, and the ME update involves applying a function specific to the value of β . Hence, the algorithms have similar complexity $\mathcal{O}(FKN)$, and their implementation takes simple forms. (Matlab implementations of the algorithms discussed in this letter are available online at <http://perso.telecom-paristech.fr/~fevotte/Code/code.beta.nmf.zip>.)

5 Simulations

In this section, we report performance results of β -NMF algorithms for the specific values $\beta \in \{0.5, 1.5, 2\}$. These values are chosen for their practical interest and because simple ME algorithms exist in their case. This section will evidence the performance improvement brought by the ME approach over the MM or heuristic approaches, with similar computational burden. More precisely, the ME algorithm considered in this section is a mixture of prolonged ME and MM, defined by equation 4.24 and with $\theta = 0.95$, but we will still refer to it as ME for simplicity. The algorithms for all three considered values of β are compared on small-sized synthetic data in section 5.1. The algorithms for $\beta = 0.5$ are analyzed in section 5.2 on the basis of a small music transcription example as this specific value of β has proven efficient for this task (FitzGerald et al., 2009; Vincent et al., 2010; Hennequin et al., 2010).

In the following results we will display the cost values through iterations as well as, following Gonzalez and Zhang (2005), KKT residuals. The residuals allow us to monitor convergence to a stationary point and are here

defined as

$$\text{KKT}(\mathbf{W}) = \|\min \{\mathbf{W}, [(\mathbf{W}\mathbf{H})^{(\beta-2)} \cdot (\mathbf{W}\mathbf{H} - \mathbf{V})]\mathbf{H}^T\}\|_1 / FK, \quad (5.1)$$

$$\text{KKT}(\mathbf{H}) = \|\min \{\mathbf{H}, \mathbf{W}^T [(\mathbf{W}\mathbf{H})^{(\beta-2)} \cdot (\mathbf{W}\mathbf{H} - \mathbf{V})]\}\|_1 / KN. \quad (5.2)$$

They are meant to converge to 0 by equation 2.12. Again, the monotonicity of the heuristic, MM, and ME algorithms does not imply convergence of the iterates to a stationary point. Hence, displaying the KKT residuals allows experimentally checking whether convergence is achieved in practice.

One iteration of each algorithm consists of updating \mathbf{W} given $\mathbf{H}^{(i-1)}$ and \mathbf{H} given $\mathbf{W}^{(i)}$ and then normalizing $\mathbf{W}^{(i)}$ and $\mathbf{H}^{(i)}$ to eliminate trivial scale indeterminacies that leave the cost function unchanged. The normalization step consists of rescaling each column of \mathbf{W} so that $\|\mathbf{w}_k\|_1 = 1$ and rescaling the k th row of \mathbf{H} accordingly. The normalization step is not required but is useful to display and compare the KKT residuals, which are scale sensitive.

5.1 Factorization of Synthetic Data. We consider a synthetic data matrix \mathbf{V} constructed as $\mathbf{V} = \mathbf{W}^* \mathbf{H}^*$ where the ground truth factors are generated as the absolute values of gaussian noise.⁴ The matrix can be exactly factorized so that all algorithms should converge to a solution such that $D(\mathbf{V}|\mathbf{W}\mathbf{H}) = 0$. The dimensions are $F = 10$, $N = 25$, and $K = 5$. The algorithms (heuristic, MM, and ME for $\beta = 0.5$, MM and ME for $\beta \in \{1.5, 2\}$) are run for 10^5 iterations and initialized with positive random values. Figures 4, 5, and 6 display for each of the three values of β the normalized cost values $D(\mathbf{V}|\mathbf{W}\mathbf{H})/FN$, the KKT residuals, and the fit residuals computed as $\|\mathbf{W}^{(i)} - \hat{\mathbf{W}}\|_F / FK$ and $\|\mathbf{H}^{(i)} - \hat{\mathbf{H}}\|_F / KN$, where $\hat{\mathbf{W}}$ and $\hat{\mathbf{H}}$ are the factor estimates at the end of the 10^5 iterations and $\|\cdot\|_F$ is the Frobenius norm. The fit residuals allow measuring the closeness of the current iterates to their end value.

The cost values in all three cases converge to 0 as an exact factorization is reached (oscillations appear in the end iterations as machine precision is reached). Convergence is achieved in all three cases, as shown by both the cost values and KKT residuals. We visually inspected the factorizations returned by the algorithms. For each value of $\beta \in \{0.5, 1.5\}$, the different algorithms appeared to converge to the same solution (and solutions obtained for the two values of β appeared comparable). This was less clear for $\beta = 2$, where ME appeared to reach out a different solution from MM. Still, in this run, ME provides the fastest convergence for every considered value of β . Other runs, obtained from other starting points (obtained randomly), tend to show that when the compared algorithms converge to the same solution, ME converges more quickly. Convergence to a common solution can

⁴For example, in Matlab notation $\mathbf{V} = \text{abs}(\text{randn}(F, K)) * \text{abs}(\text{randn}(K, N))$.

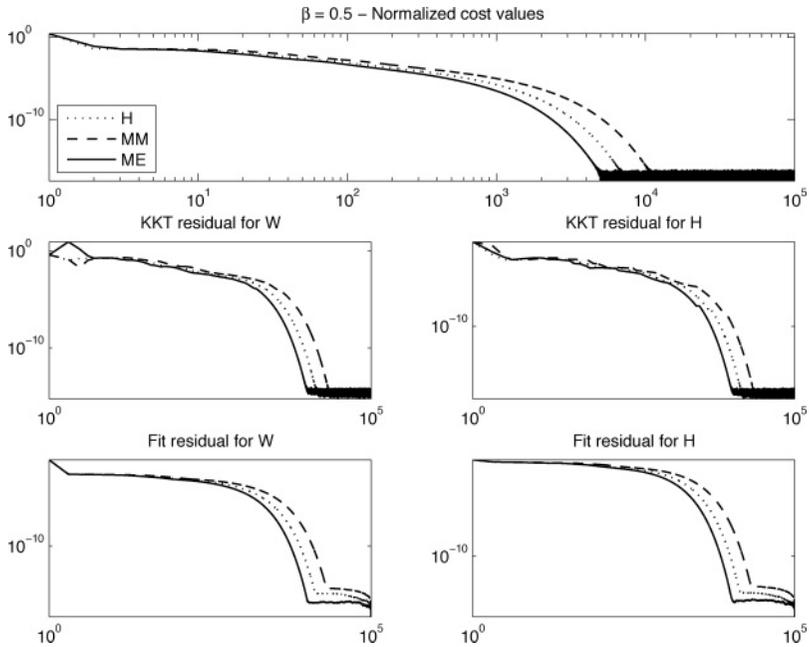


Figure 4: One run of the heuristic (H), ME, and MM algorithms on synthetic data with $\beta = 0.5$. Logarithmic scales for both x - and y -axes.

be controlled in the specific case where \mathbf{W} is fixed and $\beta \in \{1.5, 2\}$, because the objective function is then convex w.r.t. \mathbf{H} . In this scenario, ME was found to always converge faster than MM. (These simulations are reported in the companion report available online at http://perso.telecom-paristech.fr/~fevotte/Samples/neco11/beta_nmf_supp.pdf.)

The fit residuals in Figures 4, 5, and 6 show that full convergence will not need to be attained to obtain satisfying solutions for most applications as the fit residual will be considered sufficiently small after a few hundred iterations. Note that the factor iterates do not necessarily converge to the ground-truth values \mathbf{W}^* and \mathbf{H}^* (and this is what we observed) because of the identifiability ambiguities inherent in NMF (Donoho & Stodden, 2004; Laurberg, Christensen, Plumbley, Hansen, and Jensen, 2008).

Using a Matlab implementation run on a Mac 2.6 GHz computer with 2 GB RAM, the CPU time required by each algorithm for the 10^5 iterations is about 60 s for $\beta \in \{0.5, 1.5\}$ and 20 s for $\beta = 2$, including the computation of the cost values and KKT residuals. The ME algorithm is marginally more expensive than MM, itself only slightly more expensive than the heuristic algorithm, for $\beta = 0.5$. The CPU times incurred by the algorithms when $\beta = 2$ is considerably lower thanks to simplifications in equations 4.30

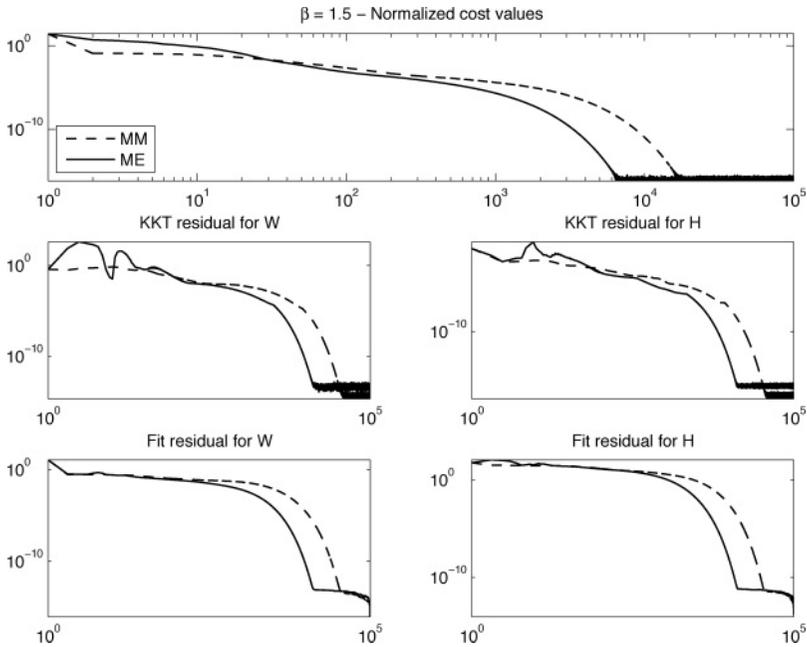


Figure 5: One run of the ME and MM algorithms on synthetic data with $\beta = 1.5$. Logarithmic scales for both x - and y -axes.

and 4.31. Indeed, in the latter case, the term $(\mathbf{W}\mathbf{H})\mathbf{H}^T$ appearing at the denominator can more efficiently be computed as $\mathbf{W}(\mathbf{H}\mathbf{H}^T)$, which involves a multiplication of matrices with smaller sizes.

5.2 Audio Spectrogram Decomposition. This section addresses the comparison of the heuristic, MM, and ME algorithms for $\beta = 0.5$ applied to an audio spectrogram. We consider the short piano sequence of Févotte et al. (2009), recorded in live conditions, composed of four musical notes, played all at once in the first measure and then played by pairs in all possible combinations in the subsequent measures. A magnitude spectrogram of the audio signal is computed, leading to nonnegative matrix data \mathbf{V} of size $F = 513$ frequency bins by $N = 674$ time frames. The data are represented in the top left of Figure 7.

As discussed in Févotte et al. (2009), K was set to 6 so as to retrieve in \mathbf{W} the individual spectra of each of the four notes and supplementary spectra corresponding to transients and residual noise. The three algorithms were initialized with common positive random values and run for 10^5 iterations. Figure 7 displays the cost values and KKT residuals along the 10^5 iterations. It was manually checked that the algorithms converged to the

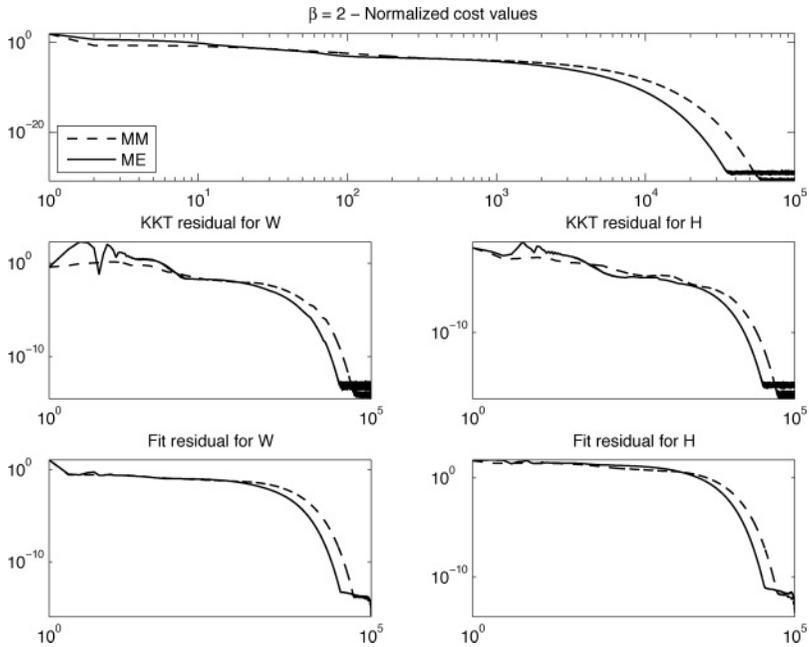


Figure 6: One run of the ME and MM algorithms on synthetic data with $\beta = 2$. Logarithmic scales for both x - and y -axes.

desired ground-truth solution, that is, the notes, transients, and residual noise spectra are correctly unmingled. The three plots show that the ME provides fastest convergence overall, though, judging from the KKT residuals, it appears that convergence is not achieved within the 10^5 iterations. However, the musical pitch values (computed from \mathbf{W} at every iteration) converge to their ground-truth values after only 30, 50, and 580 iterations for ME, heuristic, and MM, respectively. Other initializations yielded two types of results. In a minority of cases, either the heuristic and MM update, on one side, or the ME update, on the other side, converged to a local solution. In the large majority of cases, the three algorithms converge to the same solution, and the results are similar to those of Figure 7: the heuristic algorithm produces the largest decreases of the objective function in the early iterations and is then supplanted by ME. In some runs, the pitch values converged faster with the heuristic algorithm than with ME, and it was found that MM is generally slower than the other two algorithms. These results suggest a mixed update of the form of equation 4.24 where the mixture parameter θ could be made iteration dependent so as to give more weight to the heuristic update in the early iterations and then to ME.

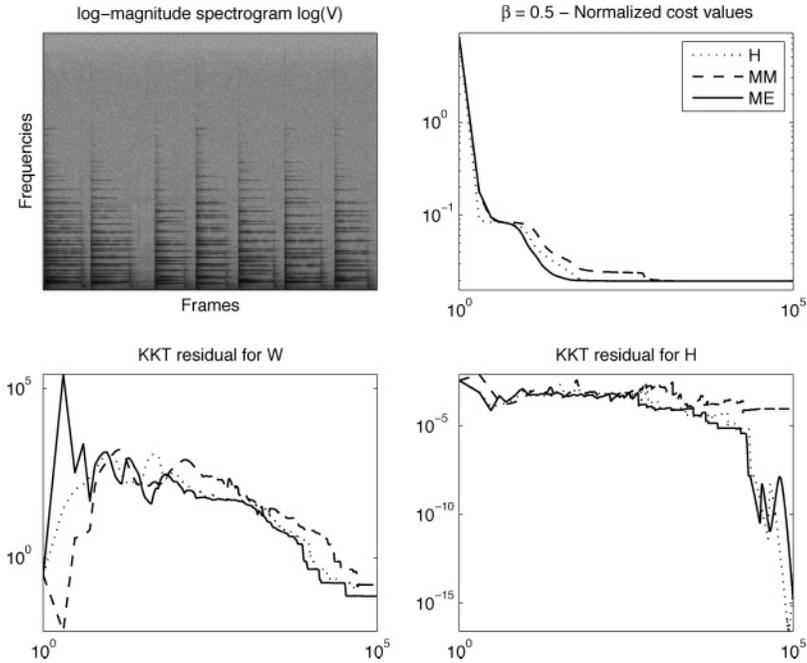


Figure 7: One run of the heuristic (H), MM, and ME algorithms on the piano magnitude spectrogram with $\beta = 0.5$. Logarithmic scales for both x - and y -axes.

5.3 Face Data Decomposition. Finally, in this section we consider decomposition of face data using β -NMF. We use the Olivetti data set, composed of 10 grayscale 8 bits 64×64 face images of 40 people. We retrieved the data in Matlab format from <http://cs.nyu.edu/~roweis/data.html>. The images are vectorized and form the columns of \mathbf{V} , with dimensions $F = 4096$ and $N = 400$. Figure 8 displays the objective functions of one run of the ME and MM algorithms for $\beta \in \{1.5, 2\}$ and illustrates the faster convergence of ME. Other runs led to sensibly similar plots.

As stated in section 1, β -NMF is popular in audio signal processing where the value of β can be controlled so as to improve transcription or separation accuracy. The idea of tuning the value of β so as to optimize performance applies to any NMF-based method for any type of data. As such, to motivate the use of β -NMF in a nonaudio setting, we propose an image interpolation example, inspired by Cichocki et al. (2008) and Cemgil (2009), where we show the influence of β on the reconstruction of missing data. We discard 25% of the Olivetti data randomly and produce NMF decompositions using the available data for $\beta \in \{-1, 0, 1, 2, 3\}$ and $K \in \{50, 100, 200\}$. Accounting for the missing data requires minor modifications in the algorithms,

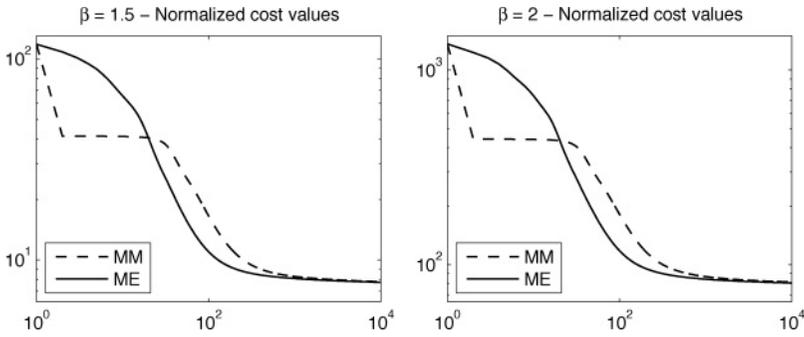


Figure 8: One run of the MM and ME algorithms on the Olivetti data set with $\beta = 1.5$ (left) and $\beta = 2$ (right). Logarithmic scales for both x - and y -axes.

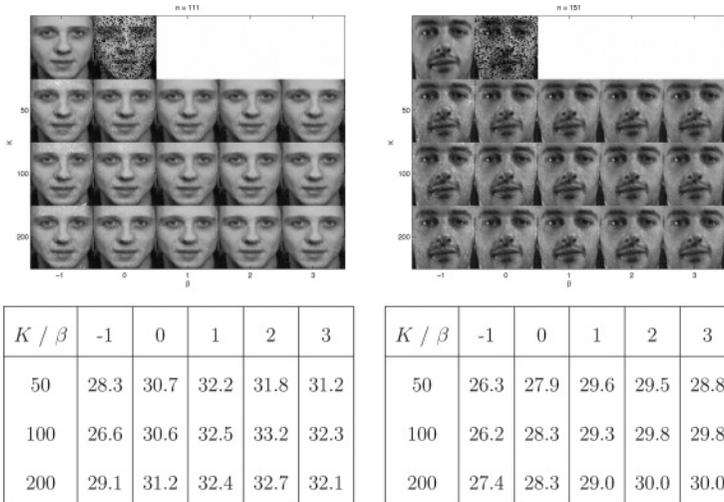


Figure 9: Interpolation results with the Olivetti data set. Original and corrupted data are shown at the top left of each plot. Below are the reconstructions obtained for $K \in \{50, 100, 200\}$ and $\beta \in \{-1, 0, 1, 2, 3\}$. Tables report PSNRs (in dB) of the reconstructions.

basically multiplying \mathbf{V} and its approximate \mathbf{WH} with a binary mask in which zeroes indicate missing pixels (see Ho, 2008; Cichocki et al., 2008; Le Roux, Kameoka, Ono, de Cheveigné, & Sagayama, 2008; Cemgil, 2009; Smaragdis, Raj, & Shashanka, 2009) for similar setups. For simplicity we considered only the MM algorithm, as it is consistently defined for all values of β . It was run from five different initializations for every combination (K, β) , and the factorization yielding the lowest end cost value was selected. Figure 9 displays the original image, missing pixels, and reconstructions for

two of the images in the data set. We have also computed the peak signal-to-noise ratio (PSNR) between original and reconstructed images.⁵ The maximum mean PSNR value (averaged over all 400 images) is obtained for $\beta = 2$ and $K = 200$. However, since the PSNR value is equivalent to the Euclidean distance between the original and reconstructed image, it is expected that the optimal value of β is biased toward the metric used to assess the quality of reconstruction. Perceptually, we often found the reconstruction obtained with $\beta = 3$ to be more satisfying than with $\beta = 2$.

6 Variants of β -NMF

In this section we briefly discuss how some common variants of NMF, penalized NMF and convex NMF, can be handled under the β -divergence.

6.1 Penalized β -NMF. Supplementary functions of \mathbf{W} or \mathbf{H} (or both) are often added to the cost function, equation 1.3, so as to induce some sort of regularization of the factor estimates or so as to reflect prior belief, for example, in Bayesian maximum a posteriori (MAP) estimation. When such penalty terms are separable in the columns of \mathbf{H} or in the rows of \mathbf{W} , penalized NMF essentially amounts to solving the following optimization problem:

$$\min_{\mathbf{h}} C_P(\mathbf{h}) \stackrel{\text{def}}{=} D(\mathbf{v}|\mathbf{W}\mathbf{h}) + L(\mathbf{h}) \text{ subject to } \mathbf{h} \geq 0, \quad (6.1)$$

where $L(\mathbf{h})$ is the penalty term. An auxiliary function to $C_P(\mathbf{h})$ is readily given by

$$G_P(\mathbf{h}|\tilde{\mathbf{h}}) \stackrel{\text{def}}{=} G(\mathbf{h}|\tilde{\mathbf{h}}) + L(\mathbf{h}), \quad (6.2)$$

where $G(\mathbf{h}|\tilde{\mathbf{h}})$ is any auxiliary function to $C(\mathbf{h}) = D(\mathbf{v}|\mathbf{W}\mathbf{h})$. MM or ME algorithms can then be designed on a case-by-case basis. Let us consider a short example for illustration: ℓ_1 -norm regularization. In that case, we have

$$L(\mathbf{h}) = \lambda \sum_k h_k, \quad (6.3)$$

where λ is a positive weight parameter. Using the auxiliary function designed in section 3.2 and equation 3.19, the gradient of the penalized

⁵PSNR is a standard evaluation criterion in image reconstruction, defined as $20 \log_{10}(FP/\|\mathbf{v} - \hat{\mathbf{v}}\|_2)$, where \mathbf{v} and $\hat{\mathbf{v}}$ denote the vectorized original and reconstructed images and P is the maximum pixel possible value ($P = 255$ in our case).

auxiliary function is

$$\nabla_{h_k} G_L(\mathbf{h}|\tilde{\mathbf{h}}) = \sum_f w_{fk} \left[\tilde{d}' \left(v_f | \tilde{v}_f \frac{h_k}{\tilde{h}_k} \right) + \widehat{d}'(v_f | \tilde{v}_f) \right] + \lambda.$$

The MM algorithm for ℓ_1 -regularized β -NMF takes a very simple form for $\beta \leq 1$, such that

$$h_k = \tilde{h}_k \left(\frac{\sum_f w_{fk} v_f \tilde{v}_f^{\beta-2}}{\sum_f w_{fk} \tilde{v}_f^{\beta-1} + \lambda} \right)^{\gamma(\beta)}. \tag{6.4}$$

This in particular leads to ℓ_1 -regularized NMF algorithms for KL-NMF and IS-NMF with proven monotonicity. An update similar to equation 6.4 is obtained for $\beta \geq 2$, but the λ term appears through its sign opposite at the numerator instead of appearing at the denominator. Hence the nonnegativity constraint may become active and must be treated carefully; in that case, our result coincides with similar findings of Pauca, Piper, and Plemmons (2006) and Mørup and Clemmensen (2007) for the specific case of ℓ_1 -regularized NMF with the Euclidean distance ($\beta = 2$). In the case $\beta \in (1, 2)$, the MM algorithm does not come up with a simple closed-form update, which supports the fact that in the penalized case, handy algorithms may come only on a case-by-case basis. This is similar to expectation-maximization (EM) procedures for MAP estimation, in which the E-step is essentially unchanged but where the M-step might become intractable because of the penalty term. ME algorithms can also be designed for the ℓ_1 -regularized problem, and it can be shown that the results of Table 3 (i.e., the values of β for which a closed-form update exists) still hold in that case.

6.2 Convex β -NMF. In some recent NMF-related works, the dictionary \mathbf{W} is constrained to belong to a known subspace $\mathbf{S} \in \mathbb{R}_+^{F \times M}$ such that

$$\mathbf{W} = \mathbf{S}\mathbf{L}, \tag{6.5}$$

where $\mathbf{L} \in \mathbb{R}_+^{M \times K}$. For example, Ding et al. (2010) assume the columns of \mathbf{W} to be linear combinations (with unknown expansion coefficients) of data points (columns of \mathbf{V}), so as to enforce the dictionary to be composed of data centroids, while Vincent et al. (2010) assume the dictionary element to be linear combinations of narrow-band spectra so as to enforce harmonicity and smoothness of the dictionary. The term *convex NMF* was introduced by Ding et al. (2010) to express the idea that \mathbf{W} belongs to the convex set of all nonnegative linear combinations of elements of \mathbf{S} , but this does not make the optimization problem convex in itself, in the general case.

In this setting, the dictionary update is tantamount to solving

$$\begin{aligned} \min_{\mathbf{L}} C_{cv}(\mathbf{L}) &\stackrel{\text{def}}{=} D(\mathbf{V}|\mathbf{SLH}) \\ &= \sum_{fn} d \left(v_{fn} \mid \sum_{mk} s_{fm} l_{mk} h_{kn} \right) \quad \text{subject to } \mathbf{L} \geq 0. \end{aligned} \tag{6.6}$$

In fact, this matricial optimization problem can be turned into vectorial nonnegative linear regression so that the results of section 4 hold. Given some mappings $(f, n) \in \{1, F\} \times \{1, N\} \rightarrow p \in \{1, FN\}$ and $(m, k) \in \{1, M\} \times \{1, K\} \rightarrow q \in \{1, MK\}$, let us introduce the following variables: \mathbf{T} is the matrix of dimension $FN \times MK$ with coefficients $t_{pq} = s_{fm} h_{kn}$, $\underline{\mathbf{v}}$ is the column vector of size FN with coefficients $v_p = v_{fn}$, and $\underline{\mathbf{l}}$ is the column vector of size MK with coefficients $l_q = l_{mk}$. Then we have

$$D(\mathbf{V}|\mathbf{SLH}) = \sum_p d \left(v_p \mid \sum_q t_{pq} l_q \right), \tag{6.7}$$

and thus the estimation of \mathbf{L} amounts to the approximation $\underline{\mathbf{v}} \approx \mathbf{T}\underline{\mathbf{l}}$. As such, any of the algorithms described in section 4 can be employed for this task. As before, the resulting vectorial updates can be turned into matricial updates, leading to simple and efficient implementations. For example, the MM update reads

$$\mathbf{L} \leftarrow \mathbf{L} \cdot \left(\frac{\mathbf{S}^T [(\mathbf{SLH})^{(\beta-2)} \cdot \mathbf{V}] \mathbf{H}^T}{\mathbf{S}^T [(\mathbf{SLH})^{(\beta-1)}] \mathbf{H}^T} \right)^{\cdot \gamma(\beta)}. \tag{6.8}$$

This result proves the monotonicity of some of the algorithms derived heuristically in Vincent et al. (2010) and also extends the results of Ding et al. (2010) for convex NMF with the Euclidean distance to the more general β -divergence.⁶

7 Conclusion

This letter has addressed NMF with the β -divergence. The problem may be reduced to a mere nonnegative linear regression problem, and our approach is based on the construction of an auxiliary function $G(\mathbf{h}|\hat{\mathbf{h}})$ that majorizes

⁶More precisely, Ding et al. (2010) consider a “semi”-NMF version where $\mathbf{S} = \mathbf{V}$ and the data are allowed to be real valued, while the nonnegativity constraint is solely imposed on \mathbf{L} and \mathbf{H} . Our results do not apply to this more general framework, only to the special case where \mathbf{V} is nonnegative.

the objective function $C(\mathbf{h})$ everywhere and is tight for $\mathbf{h} = \tilde{\mathbf{h}}$. The auxiliary function unifies existing auxiliary functions for the Euclidean distance and the KL divergence (Lee & Seung, 2001), for the generalized divergence of Kompass (2007) (in essence, the β -divergence on its convex part, that is, $\beta \in [1, 2]$) and for the IS divergence (Cao et al., 1999). Various descent algorithms, free of tuning parameters, may then be derived from this auxiliary function. As such, the findings of this letter may be summarized as follows:

- The MM algorithm based on the described auxiliary function is shown to yield multiplicative algorithms for $\beta \in \mathbb{R}$, as described by equation 4.1. For $\beta \in [1, 2]$ (interval of values for which the β -divergence is convex), the MM algorithm coincides with the heuristic algorithm given by equation 4.6, as already known from Kompass (2007).
- In section 4.2, we prove the monotonicity of the heuristic algorithm for $\beta \in (0, 1)$ by proving the inequality $G(\mathbf{h}^H | \tilde{\mathbf{h}}) \leq G(\tilde{\mathbf{h}} | \tilde{\mathbf{h}})$. Hence, with the existing monotonicity results for $\beta = 0$ and $\beta \in [1, 2]$, aggregated, it can now be claimed that the heuristic algorithm is monotone for $\beta \in [0, 2]$, which is the range of values of practical interest that has been considered in the literature.
- In section 4.3, we introduced the concept of maximization-equalization (ME) algorithms. Such algorithms are exhibited for specific values of β , in particular for $\beta \in \{0, 0.5, 1.5, 2\}$, which are values of practical interest. For $\beta = 0$ (IS divergence), the ME algorithm coincides with the heuristic algorithm, whose monotonicity already holds from Cao et al. (1999). For other values of β , the ME algorithms are nonmultiplicative. For $\beta \in \{0.5, 1.5, 2\}$ they amount to solving polynomial equations of order 1 or 2. Simulations have illustrated the faster convergence of the ME approach w.r.t to MM or heuristic, with equivalent complexity.
- In section 6 we considered variants of NMF with the β -divergence. We explained how penalty terms may be handled in the auxiliary function setting; in particular we presented simple multiplicative algorithms for ℓ_1 regularized KL or IS NMF. Then we showed how the algorithms constructed for plain NMF hold for convex NMF, generalizing and proving the monotonicity of existing algorithms.

As for perspectives, this work leaves two important questions unanswered. The first one is the monotonicity of the heuristic algorithm for $\beta \notin [0, 2]$. The monotonicity is observed in practice, but we have not been able to come up with proofs in the presented setting. Either other approaches need to be followed, or a different type of auxiliary function from the one presented here needs to be envisaged. As we suggested in section 2.1, the convex-concave decomposition of the β -divergence is not unique, and decompositions other than the “natural” one employed in this letter may lead

to auxiliary functions that more closely fit the criterion. The second, probably more ambitious, question is the convergence of the sequence of iterates produced by the proposed algorithms to a stationary point. Partial results exist for Euclidean NMF (Lin, 2007a), convergence of multiplicative rules for nonnegative linear regression (i.e., when only one of the two matrices is updated) has been studied in a few cases; see (Titterton, 1987; De Pierro, 1993; Eggermont & LaRiccia, 1998), but general results for NMF with the β -divergence are still lacking. A noteworthy attempt has recently been made by Badeau et al. (2010), which points to difficulties in the convergence study due to the inherent scale ambiguity of factorization models.

Finally, another relevant perspective is the design of new types of β -NMF algorithms. In the Euclidean case, projected gradient methods (Lin, 2007b), second-order active sets methods (Kim, Sra, & Dhillon, 2008), and block-coordinate descent methods (Mairal et al., 2010) have recently been shown to outperform standard multiplicative updates; see also Mørup and Hansen (2009) for a comparison of a selection of algorithms. As such it would be interesting to study how these approaches may extend to the more general β -NMF framework.

Acknowledgments

We thank Francis Bach, Henri Lantéri, Augustin Lefèvre, Cédric Richard, and Céline Theys for inspiring discussions related to optimization in NMF. Many thanks to the reviewers for helpful comments. This work is supported by project ANR-09-JCJC-0073-01 TANGERINE (Theory and Applications of Nonnegative Matrix Factorization).

References

- Badeau, R., Bertin, N., & Vincent, E. (2010). Stability analysis of multiplicative update algorithms and application to non-negative matrix factorization. *IEEE Transactions on Neural Networks*, 21(12), 1869–1881.
- Basu, A., Harris, I. R., Hjort, N. L., & Jones, M. C. (1998). Robust and efficient estimation by minimising a density power divergence. *Biometrika*, 85(3), 549–559.
- Berry, M. W., Browne, M., Langville, A. N., Pauca, V. P., & Plemmons, R. J. (2007). Algorithms and applications for approximate nonnegative matrix factorization. *Computational Statistics and Data Analysis*, 52(1), 155–173.
- Bertin, N., Févotte, C., & Badeau, R. (2009). A tempering approach for Itakura-Saito non-negative matrix factorization. With application to music transcription. In *Proc. International Conference on Acoustics, Speech and Signal Processing* (pp. 1545–1548). Piscataway, NJ: IEEE.
- Brunet, J.-P., Tamayo, P., Golub, T. R., & Mesirov, J. P. (2004). Metagenes and molecular pattern discovery using matrix factorization. In *Proceedings of the National Academy of Sciences*, 101, 4164–4169.

- Cao, Y., Eggermont, P. P. B., & Terebey, S. (1999). Cross Burg entropy maximization and its application to ringing suppression in image reconstruction. *IEEE Transactions on Image Processing*, 8(2), 286–292.
- Cemgil, A. T. (2009). Bayesian inference for nonnegative matrix factorisation models. *Computational Intelligence and Neuroscience*, 2009, 785152. doi:10.1155/2009/785152.
- Cichocki, A., & Amari, S. (2010). Families of alpha- beta- and gamma-divergences: Flexible and robust measures of similarities. *Entropy*, 12(6), 1532–1568.
- Cichocki, A., Lee, H., Kim, Y.-D., & Choi, S. (2008). Non-negative matrix factorization with α -divergence. *Pattern Recognition Letters*, 29(9), 1433–1440.
- Cichocki, A., Zdunek, R., & Amari, S. (2006). Csiszar's divergences for non-negative matrix factorization: Family of new algorithms. In *Proc. 6th International Conference on Independent Component Analysis and Blind Signal Separation (ICA'06)* (pp. 32–39). Berlin: Springer.
- Daube-Witherspoon, M., & Muehllehner, G. (1986). An iterative image space reconstruction algorithm suitable for volume ECT. *IEEE Transactions on Medical Imaging*, 5(5), 61–66.
- De Pierro, A. R. (1993). On the relation between the ISRA and the EM algorithm for positron emission tomography. *IEEE Trans. Medical Imaging*, 12(2), 328–333.
- Dessein, A., Cont, A., & Lemaitre, G. (2010). Real-time polyphonic music transcription with non-negative matrix factorization and beta-divergence. In *Proc. 11th International Society for Music Information Retrieval Conference*. Society for Music Information Retrieval.
- Dhillon, I. S., & Sra, S. (2005). Generalized nonnegative matrix approximations with Bregman divergences. In B. Schölkopf, J. Platt, & T. Hoffman (Eds.), *Advances in neural information processing systems*, 19. Cambridge, MA: MIT Press.
- Ding, C. H. Q., Li, T., & Jordan, M. I. (2010). Convex and semi-nonnegative matrix factorizations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(1), 45–55.
- Donoho, D., & Stodden, V. (2004). When does non-negative matrix factorization give a correct decomposition into parts? In S. Thrun, L. Saul, & B. Schölkopf (Eds.), *Advances in neural information processing systems*, 16. Cambridge, MA: MIT Press.
- Drakakis, K., Rickard, S., de Frein, R., & Cichocki, A. (2007). Analysis of financial data using non-negative matrix factorization. *International Mathematical Forum*, 3, 1853–1870.
- Eggermont, P. P. B., & LaRiccia, V. N. (1998). *On EM-like algorithms for minimum distance estimation*. <http://www.udel.edu/FREC/eggermont/Preprints/emlike.pdf>.
- Eguchi, S., & Kano, Y. (2001). *Robustifying maximum likelihood estimation*. (Tech. Rep. Research Memo 802). Tokyo: Institute of Statistical Mathematics.
- Févotte, C., Bertin, N., & Durrieu, J.-L. (2009). Nonnegative matrix factorization with the Itakura-Saito divergence. With application to music analysis. *Neural Computation*, 21(3), 793–830.
- Févotte, C., & Cemgil, A. T. (2009). Nonnegative matrix factorisations as probabilistic inference in composite models. In *Proc. 17th European Signal Processing Conference (EUSIPCO)* (pp. 1913–1917). EURASIP.

- FitzGerald, D., Cranitch, M., & Coyle, E. (2009). On the use of the beta divergence for musical source separation. In *Proc. Irish Signals and Systems Conference*.
- Gao, Y., & Church, G. (2005). Improving molecular cancer class discovery through sparse non-negative matrix factorization. *Bioinformatics*, 21, 3970–3975.
- Gonzalez, E. F., & Zhang, Y. (2005). *Accelerating the Lee-Seung algorithm for non-negative matrix factorizations* (Tech. Rep.). Houston, TX: Rice University.
- Greene, D., Cagney, G., Krogan, N., & Cunningham, P. (2008). Ensemble non-negative matrix factorization methods for clustering protein-protein interactions. *Bioinformatics*, 24(15), 1722–1728.
- Hennequin, R., Badeau, R., & David, B. (2010). NMF with time-frequency activations to model non stationary audio events. In *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'10)* (pp. 445–448). Piscataway, NJ: IEEE.
- Ho, N.-D. (2008). *Nonnegative matrix factorization algorithms and applications*. Unpublished doctoral dissertation, Université Catholique de Louvain.
- Hunter, D. R., & Lange, K. (2004). A tutorial on MM algorithms. *American Statistician*, 58, 30–37.
- Kim, D., Sra, S., & Dhillon, I. S. (2008). Fast projection-based methods for the least squares nonnegative matrix approximation problem. *Statistical Analysis and Data Mining*, 1, 38–51.
- Kompass, R. (2007). A generalized divergence measure for nonnegative matrix factorization. *Neural Computation*, 19(3), 780–791.
- Lantéri, H., Roche, M., Cuevas, O., & Aïme, C. (2001). A general method to devise maximum-likelihood signal restoration multiplicative algorithms with non-negativity constraints. *Signal Processing*, 81(5), 945–974.
- Lantéri, H., Theys, C., Richard, C., & Févotte, C. (2010). Split gradient method for nonnegative matrix factorization. In *Proc. 18th European Signal Processing Conference (EUSIPCO)*. EURASIP.
- Laurberg, H., Christensen, M., Plumbley, M. D., Hansen, L. K., & Jensen, S. H. (2008). Theorems on positive data: On the uniqueness of NMF. *Computational Intelligence and Neuroscience*, article 764206.
- Le Roux, J., Kameoka, H., Ono, N., de Cheveigné, A., & Sagayama, S. (2008). Computational auditory induction by missing-data non-negative matrix factorization. In *Proc. ISCA Workshop on Statistical and Perceptual Audition (SAPA)* (pp. 1–6).
- Lee, D. D., & Seung, H. S. (1999). Learning the parts of objects with nonnegative matrix factorization. *Nature*, 401, 788–791.
- Lee, D. D., & Seung, H. S. (2001). Algorithms for non-negative matrix factorization. In T. K. Leen, T. G. Dietterich, & K.-R. Müller (Eds.), *Advances in neural and information processing systems*, 13 (pp. 556–562). Cambridge, MA: MIT Press.
- Lin, C.-J. (2007a). On the convergence of multiplicative update algorithms for non-negative matrix factorization. *IEEE Transactions on Neural Networks*, 18, 1589–1596.
- Lin, C.-J. (2007b). Projected gradient methods for nonnegative matrix factorization. *Neural Computation*, 19, 2756–2779.
- Lucy, L. B. (1974). An iterative technique for the rectification of observed distributions. *Astronomical Journal*, 79, 745–754.
- Mairal, J., Bach, F., Ponce, J., & Sapiro, G. (2010). Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*, 11, 10–60.

- Minami, M., & Eguchi, S. (2002). Robust blind source separation by beta-divergence. *Neural Computation*, 14, 1859–1886.
- Mørup, M., & Clemmensen, L. H. (2007). Multiplicative updates for the LASSO. In *Proc. IEEE Workshop on Machine Learning for Signal Processing*. Piscataway, NJ: IEEE.
- Mørup, M., & Hansen, L. K. (2009). Tuning pruning in sparse non-negative matrix factorization. In *Proc. 17th European Signal Processing Conference (EUSIPCO'09)*. EURASIP.
- Nakano, M., Kameoka, H., Le Roux, J., Kitano, Y., Ono, N., & Sagayama, S. (2010). Convergence-guaranteed multiplicative algorithms for non-negative matrix factorization with beta-divergence. In *Proc. IEEE International Workshop on Machine Learning for Signal Processing*. Piscataway, NJ: IEEE.
- O'Grady, P. D. (2007). *Sparse separation of under-determined speech mixtures*. Unpublished doctoral dissertation, National University of Ireland Maynooth.
- O'Grady, P. D., & Pearlmutter, B. A. (2008). Discovering speech phones using convolutive non-negative matrix factorisation with a sparseness constraint. *Neuro-computing*, 72(1–3), 88–101.
- Pauca, V. P., Piper, J., & Plemmons, R. J. (2006). Nonnegative matrix factorization for spectral data analysis. *Linear Algebra and Its Applications*, 416, 29–47.
- Richardson, W. H. (1972). Bayesian-based iterative method of image restoration. *Journal of the Optical Society of America*, 62, 55–59.
- Salakhutdinov, R., & Roweis, S. (2003). Adaptive overrelaxed bound optimization methods. In *Proc. International Conference on Machine Learning* (pp. 664–671). CSREA Press.
- Smaragdis, P., & Brown, J. C. (2003). Non-negative matrix factorization for polyphonic music transcription. In *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*. Piscataway, NJ: IEEE Press.
- Smaragdis, P., Raj, B., & Shashanka, M. (2009). Missing data imputation for spectral audio signals. In *IEEE International Workshop on Machine Learning for Signal Processing*. Piscataway, NJ: IEEE Press.
- Titterton, D. M. (1987). On the iterative image space reconstruction algorithm for ECT. *IEEE Trans. Medical Imaging*, 6(1), 52–56.
- Vincent, E., Bertin, N., & Badeau, R. (2010). Adaptive harmonic spectral decomposition for multiple pitch estimation. *IEEE Trans. Audio, Speech and Language Processing*, 18, 528–537.