

2008/64



Nonnegative factorization
and the maximum edge biclique problem

Nicolas Gillis and François Glineur

CORE

Voie du Roman Pays 34

B-1348 Louvain-la-Neuve, Belgium.

Tel (32 10) 47 43 04

Fax (32 10) 47 43 01

E-mail: corestat-library@uclouvain.be

<http://www.uclouvain.be/en-44508.html>

CORE DISCUSSION PAPER
2008/64

**Nonnegative factorization
and the maximum edge biclique problem**

Nicolas GILLIS¹ and François GLINEUR²

November 2008

Abstract

Nonnegative Matrix Factorization (NMF) is a data analysis technique which allows compression and interpretation of nonnegative data. NMF became widely studied after the publication of the seminal paper by Lee and Seung (Learning the Parts of Objects by Nonnegative Matrix Factorization, *Nature*, 1999, vol. 401, pp. 788--791), which introduced an algorithm based on Multiplicative Updates (MU). More recently, another class of methods called Hierarchical Alternating Least Squares (HALS) was introduced that seems to be much more efficient in practice.

In this paper, we consider the problem of approximating a not necessarily nonnegative matrix with the product of two nonnegative matrices, which we refer to as Nonnegative Factorization (NF); this is the subproblem that HALS methods implicitly try to solve at each iteration. We prove that NF is NP-hard for any fixed factorization rank, using a reduction to the maximum edge biclique problem.

We also generalize the multiplicative updates to NF, which allows us to shed some light on the differences between the MU and HALS algorithms for NMF and give an explanation for the better performance of HALS. Finally, we link stationary points of NF with feasible solutions of the biclique problem to obtain a new type of biclique finding algorithm (based on MU) whose iterations have an algorithmic complexity proportional to the number of edges in the graph, and show that it performs better than comparable existing methods.

Keywords: nonnegative matrix factorization, nonnegative factorization, complexity, multiplicative updates, hierarchical alternating least squares, maximum edge biclique.

¹ CORE, Université catholique de Louvain, Belgium. E-mail: Nicolas.gillis@uclouvain.be

² CORE, Université catholique de Louvain, Belgium. E-mail: francois.glineur@uclouvain.be.

We thank Pr. Paul Van Dooren and Pr. Laurence Wolsey for helpful discussions and advice. The first author is a research fellow of the Fonds de la Recherche Scientifique (F.R.S.-FNRS). The second one is also member of ECORE, the newly created association between CORE and ECARES.

This paper presents research results of the Belgian Program on Interuniversity Poles of Attraction initiated by the Belgian State, Prime Minister's Office, Science Policy Programming. The scientific responsibility is assumed by the authors.

1 Introduction

(Approximate) *Nonnegative Matrix Factorization* (NMF) is the problem of approximating a given nonnegative matrix by the product of two low-rank nonnegative matrices: given a matrix $M \geq 0$, one has to compute two low-rank matrices $V, W \geq 0$ such that

$$M \approx VW. \quad (1.1)$$

This problem was first introduced in 1994 by Paatero and Tapper [25], and more recently received a considerable interest after the publication of two papers by Lee and Seung [21, 22]. It is now well established that NMF is useful in the framework of compression and interpretation of nonnegative data ; it has for example been applied in analysis of image databases, text mining, interpretation of spectra, computational biology and many other applications (see e.g. [2, 9, 11] and references therein).

How can one interpret the outcome of a NMF? Assume each column $M_{:j}$ of matrix M represents an element of a data set: expression (1.1) can be equivalently written as

$$M_{:j} \approx \sum_k V_{:k} W_{kj}, \quad \forall j \quad (1.2)$$

where each element $M_{:j}$ is decomposed into a nonnegative linear combination (with weights W_{kj}) of nonnegative basis elements ($\{V_{:k}\}$, the columns of V). Nonnegativity of V allows interpretation of the basis elements in the same way as the original nonnegative elements in M , which is crucial in applications where the nonnegativity property is a requirement (e.g. where elements are images described by pixel intensities or texts represented by vectors of word counts). Moreover, nonnegativity of the weight matrix W corresponds to an essentially additive reconstruction which leads to a *part-based representation*: basis elements will represent similar parts of the columns of M . Sparsity is another important consideration: finding sparse factors improves compression and leads to a better part-based representation of the data [19].

We start this paper with a brief introduction to the NMF problem: Section 2 recalls existing complexity results, introduces two well-known classes of methods: multiplicative updates [22] and hierarchical alternating least squares [6] and proposes a simple modification to guarantee their convergence. The central problem studied in this paper, Nonnegative Factorization (NF), is a generalization of NMF where the matrix to be approximated with the product of two low-rank nonnegative matrices is not necessarily nonnegative. NF is introduced in Section 3, where it is shown to be NP-hard for any given factorization rank, using a reduction to the problem of finding a maximum edge biclique. Stationary points of the NF problem used in that reduction are also studied. This section ends with a generalization of the NMF multiplicative updates rules to the NF problem and a proof of their convergence. This allows us to shed new light on the standard multiplicative updates for NMF: a new interpretation is given in Section 4, which explains the relatively poor performance of these methods and hints at possible improvements. Finally, Section 5 introduces a new type of biclique finding algorithm that relies on the application of multiplicative updates to the equivalent NF problem considered earlier. This algorithm only requires a number of operations proportional to the number of edges of the graph per iteration, and is shown to perform well when compared to existing methods.

2 Nonnegative Matrix Factorization (NMF)

Given a matrix $M \in \mathbb{R}_+^{m \times n}$ and an integer $r \in \mathbb{N}_0$, the *NMF optimization problem* using the Frobenius norm is defined as

$$\min_{V \in \mathbb{R}^{m \times r}, W \in \mathbb{R}^{r \times n}} \|M - VW\|_F^2 = \sum_{i,j} (M - VW)_{ij}^2 \quad \text{such that } V, W \geq 0 \quad (\text{NMF})$$

$\mathbb{R}^{m \times n}$ denotes the set of real matrices of dimension $m \times n$; $\mathbb{R}_+^{m \times n}$ the set of nonnegative matrices i.e. $\mathbb{R}^{m \times n}$ with every entry nonnegative, and $\mathbf{0}$ the zero matrix of appropriate dimensions.

A wide range of algorithms have been proposed to find approximate solutions for this problem (see e.g. [2, 4, 6, 7, 10, 11, 13, 24]). Most of them use the fact that although problem (NMF) is not convex, its objective function is convex separately in each of the two factors V and W (which implies that finding the optimal factor V corresponding to a fixed factor W reduces to a convex optimization problem, and vice-versa), and try to find good approximate solutions by using alternating minimization schemes. For instance, Nonnegative Least Squares (NNLS) algorithms can be used to minimize (exactly) the cost function alternatively over factors V and W (see e.g. [5, 20]).

Actually, there exist other partitions of the variables that preserve convexity of the alternating minimization subproblems: since the cost function can be rewritten as $\|M - \sum_{i=1}^r V_{:i} W_{i:}\|_F$, it is clearly convex as long as variables do not include simultaneously an element of a column of V and an element of the corresponding row of W (i.e. V_{ki} and W_{il} for the same index i). Therefore, given a subset of indexes $K \subseteq R = \{1, 2, \dots, r\}$, (NMF) is clearly convex for both the following subsets of variables

$$P_K = \left\{ V_{:i} \mid i \in K \right\} \cup \left\{ W_{j:} \mid j \in R \setminus K \right\}$$

and its complement

$$Q_K = \left\{ V_{:i} \mid i \in R \setminus K \right\} \cup \left\{ W_{j:} \mid j \in K \right\}.$$

However, the convexity is lost as soon as one column of V ($V_{:i}$) and the corresponding row of W ($W_{i:}$) are optimized simultaneously, so that the corresponding minimization subproblem can no longer be efficiently solved up to global optimality.

2.1 Complexity

Vavasis studies in [30] the algorithmic complexity of the NMF optimization problem; more specifically, he proves that the following problem, called *Exact Nonnegative Matrix Factorization*¹, is NP-hard:

(Exact NMF) Given a nonnegative matrix $M \geq 0$ of rank k , find, if possible, two nonnegative factors $V \geq 0$ and $W \geq 0$ of rank k such that $M = VW$.

The NMF optimization problem is therefore also NP-hard, since when the rank r is equal to the rank k of the matrix M , any optimal solution to the NMF optimization problem can be used to answer the Exact NMF problem (the answer being positive if and only if the optimal objective value of the NMF optimization problem is equal to zero).

The NP-hardness proof for exact NMF relies on its equivalence with a NP-hard problem in polyhedral combinatorics, and requires both the dimensions of matrix M and its rank k to increase to obtain NP-hardness. In contrast, in the special cases when rank k is equal to 1 or 2, the exact NMF problem can always be answered in the affirmative:

1. When $k = 1$, it is obvious that for any nonnegative rank-one matrix $M \geq 0$ there is nonnegative factors $v \geq 0$ and $w \geq 0$ such that $M = vw^T$.

Moreover, the NMF optimization problem with $r = 1$ can be solved in polynomial time: the Perron-Frobenius theorem implies that the dominant left and right singular vectors of a nonnegative matrix M are nonnegative, while the Eckart-Young theorem states that the outer product of these dominant singular vectors is the best rank-one approximation of M ; these vectors can be computed in polynomial-time using for example the singular value decomposition [15].

¹This is closely related to the *nonnegative rank* of matrix M , which is the minimum value of r for which there exists $V \in \mathbb{R}_+^{m \times r}$ and $W \in \mathbb{R}_+^{r \times n}$ such that $M = VW$ (see [1]).

2. When nonnegative matrix M has rank 2, Thomas has shown [29] that exact NMF is also always possible (see also [8]). The fact that any rank-two nonnegative matrix can be exactly factorized as the product of two rank-two nonnegative matrices can be explained geometrically as follows: viewing columns of M as points in $\mathbb{R}^m$, the fact that M has rank 2 implies that the set of its columns belongs to a two-dimensional subspace. Furthermore, because these columns are nonnegative, they belong to a two-dimensional pointed cone, see Figure 1. Since such a cone is always spanned by two extremes vectors, this implies that all columns of M can be represented exactly as nonnegative linear combinations of two nonnegative vectors, and therefore the exact NMF is always possible².




Figure 1: Rank-two exact NMF ($k = 2$): $m = 3$ and $n = 10$.

Moreover, these two extreme columns can easily be computed in polynomial time (using for example the fact that they define an angle of maximum amplitude among all pairs of columns). Hence, when the optimal rank-two approximation of matrix M is nonnegative, the NMF optimization problem with $r = 2$ can be solved in polynomial time. However, this optimal rank-two approximation is not always nonnegative, so that the complexity of the NMF optimization in the case $r = 2$ is not known. Furthermore, to the best of our knowledge, the complexity of the exact NMF problem and the NMF optimization problem are still unknown for any fixed rank r or k greater than 3.

2.2 Multiplicative Updates (MU)

In their seminal paper [22], Lee and Seung propose multiplicative update rules that aim at minimizing the Frobenius norm between M and VW . To understand the origin of these rules, consider the Karush-Kuhn-Tucker first-order optimality conditions for (NMF)

$$V \geq \mathbf{0}, \quad W \geq \mathbf{0} \quad (2.1)$$

$$\nabla_V \|M - VW\|_F^2 \geq \mathbf{0}, \quad \nabla_W \|M - VW\|_F^2 \geq \mathbf{0} \quad (2.2)$$

$$V \circ \nabla_V \|M - VW\|_F^2 = \mathbf{0}, \quad W \circ \nabla_W \|M - VW\|_F^2 = \mathbf{0} \quad (2.3)$$

where $\circ$ is the Hadamard (component-wise) product between two matrices, and

$$\nabla_V \|M - VW\|_F^2 = -2(M - VW)W^T, \quad \nabla_W \|M - VW\|_F^2 = -2V^T(M - VW). \quad (2.4)$$

²The reason why this property no longer holds for higher values of the rank k is that a k -dimensional cone is not necessarily spanned by a set of k vectors when $k > 2$.

Injecting (2.4) in (2.3), we obtain

$$V \circ (VWW^T) = V \circ (MW^T) \quad (2.5)$$

$$W \circ (V^T VW) = W \circ (V^T M). \quad (2.6)$$

From these equalities, Lee and Seung derive the following simple multiplicative update rules (where $\frac{[\cdot]}{[\cdot]}$ is Hadamard (component-wise) division)

$$V \leftarrow V \circ \frac{[MW^T]}{[VWW^T]}, \quad W \leftarrow W \circ \frac{[V^T M]}{[V^T VW]} \quad (2.7)$$

for which they are able to prove a monotonicity property:

Theorem 1 ([22]). *The Frobenius norm $\|M - VW\|_F$ is nonincreasing under the multiplicative update rules (2.7).*

The algorithm based on the alternated application of these rules is not guaranteed to converge to a first-order stationary point (see e.g. [2], and references therein), but a slight modification proposed in [23] achieves this property (roughly speaking, MU is recast as a variable metric steepest descent method and the step length is modified accordingly). We propose another possibility to overcome this problem by replacing the above updates by the following:

Theorem 2. *For every constant $\epsilon > 0$, $\|M - VW\|_F$ is nonincreasing under*

$$V \leftarrow \max\left(\epsilon, V \circ \frac{[MW^T]}{[VWW^T]}\right), \quad W \leftarrow \max\left(\epsilon, W \circ \frac{[V^T M]}{[V^T VW]}\right) \quad (2.8)$$

for any $(V, W) \geq \epsilon$. Moreover, every limit point of this algorithm is a stationary point of the following optimization problem

$$\min_{V \geq \epsilon, W \geq \epsilon} \|M - VW\|_F^2. \quad (2.9)$$

Proof. See Section 3.3 of this paper, where the more general Theorem 9 is proved. $\square$

2.3 Hierarchical Alternating Least Squares (HALS)

Cichoki et al. [6] and independently several other authors [14, 18] have proposed to solve the problem of Nonnegative Matrix Factorization by considering successively each rank-one factor $V_{:k}W_{k:}$ while keeping the rest of the variables fixed, which can be expressed as

$$M \approx V_{:k}W_{k:} + \sum_{i \neq k} V_{:i}W_{i:} \Leftrightarrow V_{:k}W_{k:} \approx M - \sum_{i \neq k} V_{:i}W_{i:} \quad \text{or} \quad V_{:k}W_{k:} \approx R_k \quad (2.10)$$

where matrix R_k is called the k^{th} residual matrix.

Ideally, one would like to find an optimal rank-one factor $V_{:k}W_{k:}$ according to the Frobenius norm, i.e. solve the following problem

$$\min_{V_{:k} \in \mathbb{R}^m, W_{k:} \in \mathbb{R}^n} \|M - VW\|_F^2 = \|R_k - V_{:k}W_{k:}\|_F^2 \quad \text{such that } V_{:k}, W_{k:} \geq \mathbf{0} \quad (2.11)$$

but, instead of solving this problem directly, these authors propose to optimize column $V_{:k}$ and row $W_{k:}$ separately in an alternating scheme, because the optimal solution to these two (convex) subproblems can be easily computed in closed form, see e.g. [16]:

$$V_{:k}^* = \operatorname{argmin}_{V_{:k} \geq 0} \|R_k - V_{:k}W_{k:}\|_F^2 = \max\left(\mathbf{0}, \frac{R_k W_{k:}^T}{\|W_{k:}\|_2^2}\right) \quad (2.12)$$

$$W_{k:}^* = \operatorname{argmin}_{W_{k:} \geq 0} \|R_k - V_{:k}W_{k:}\|_F^2 = \max\left(\mathbf{0}, \frac{V_{:k}^T R_k}{\|V_{:k}\|_2^2}\right). \quad (2.13)$$

This scheme, which amounts to a block coordinate descent method (for which any cyclic order on the columns of V and the rows of W is admissible), is called *Hierarchical Alternating Least Squares* (HALS)³ and it has been observed to work remarkably well in practice: it outperforms, in most cases, the other algorithms for NMF [7, 14, 16]. Indeed, it combines a low computational cost per iteration (the same as the multiplicative updates) with a relatively fast convergence (significantly faster than the multiplicative updates), see Figure 3 for an example. We will explain later in Section 4 why this algorithm performs much better than the one of Lee and Seung.

A potential issue with this method is that, in the course of the optimization process, one of the vectors $V_{:k}$ (or $W_{k:}$) and the corresponding rank-one factor $V_{:k}W_{k:}$ may become equal to zero (this happens for example if one of the residuals R_k is nonpositive). This then leads to numerical instabilities (the next update is not well-defined) and a rank-deficient approximation (with a rank lower than r). A possible way to overcome this problem is to replace the zero lower bounds on $V_{:k}$ and $W_{k:}$ in (2.12) and (2.13) by a small positive constant, say $\epsilon \ll 1$ (as for the MU), and consider the following subproblems

$$V_{:k}^* = \operatorname{argmin}_{V_{:k} \geq \epsilon} \|R_k - V_{:k}W_{k:}\|_F^2 \text{ and } W_{k:}^* = \operatorname{argmin}_{W_{k:} \geq \epsilon} \|R_k - V_{:k}W_{k:}\|_F^2, \quad (2.14)$$

which lead to the modified closed-form update rules:

$$V_{:k}^* = \max \left(\epsilon, \frac{R_k W_{k:}^T}{\|W_{k:}\|_2^2} \right) \quad \text{and} \quad W_{k:}^* = \max \left(\epsilon, \frac{V_{:k}^T R_k}{\|V_{:k}\|_2^2} \right). \quad (2.15)$$

This idea was already suggested in [6] in order to avoid numerical instabilities. In fact, this variant of the algorithm is now well-defined in all situations because (2.14) guarantees $V_{:k} > 0$ and $W_{k:} > 0$ at each iteration. Furthermore, one can now easily prove that it converges to a stationary point.

Theorem 3. *For every constant $\epsilon > 0$, the limit points of the block coordinate descent algorithm initialized with positive matrices and applied to the optimization problem (2.9) are stationary points.*

Proof. We use the following result of Powell [27] (see also [3, p.268]): the limit points of the iterates of a block coordinate descent algorithm are stationary points provided that the following two conditions hold:

- each block of variables is required to belong to a closed convex set,
- the minimum computed at each iteration for a given block of variables is uniquely attained.

The first condition is clearly satisfied here, since $V_{:k}$ and $W_{k:}$ belong respectively to $([\epsilon, +\infty])^m$ and $([\epsilon, +\infty])^n$, which are closed convex sets. The second condition holds because subproblems (2.14) can be shown to be strictly convex, so that their optimal value is uniquely attained by the solutions provided by rules (2.15). Strict convexity is due to the fact that the objective function of these problems are sums of quadratic terms, each involving a single variable and having a strictly positive coefficient. $\square$

3 Nonnegative Factorization (NF)

Looking back at subproblem (2.11), i.e. approximating the residual R_k with a rank-one term $V_{:k}W_{k:}$, we have seen that the optimal solution separately for both $V_{:k}$ and $W_{k:}$ can be written in a closed form. In the previous section, subproblem (2.11) was then solved by a block coordinate descent algorithm.

A question arises: Is it possible to do better? i.e. *Is it possible to efficiently solve the problem for both vectors simultaneously?* In order to answer this question, we introduce the problem of Nonnegative

³In [16, 18], it is called Rank-one Residue Iteration (RRI) method and in [14] Alternating NMF (ANMF).

Factorization⁴ which is exactly the same as Nonnegative Matrix Factorization except that the matrix to factorize can be *any real matrix*, i.e. is not necessarily nonnegative. Given $M \in \mathbb{R}^{m \times n}$ and $r \in \mathbb{N}_0$, the Nonnegative Factorization optimization problem using the Frobenius norm is:

$$\begin{aligned} \min_{V \in \mathbb{R}^{m \times r}, W \in \mathbb{R}^{r \times n}} & \|M - VW\|_F^2 \\ & V \geq \mathbf{0}, \quad W \geq \mathbf{0} \end{aligned} \quad (\text{NF})$$

Of course, this problem is a generalization of (NMF) and is NP-hard as well. However Nonnegative Factorization will be shown below to be NP-hard for any *fixed* factorization rank (even $r = 1$), which is not the case of (NMF) (cf. Section 2.1). The proof is based on the reduction to the maximum edge biclique problem.

3.1 Complexity

The main result of this section is the NP-hardness result for Nonnegative Factorization for any fixed factorization rank. We first show how the optimization version of the maximum edge biclique problem (MBP) can be formulated as a rank-one Nonnegative Factorization problem (NF-1d). Since the decision version of (MBP) is NP-complete [26], this implies that (NF-1d) is NP-hard. We then prove that (NF) is NP-hard as well using a simple construction.

The Maximum Edge Biclique Problem in Bipartite Graphs

A *bipartite graph* G_b is a graph whose vertices can be divided into two disjoint sets V_1 and V_2 such that there is no edge between two vertices in the same set

$$G_b = (V, E) = (V_1 \cup V_2, E \subseteq (V_1 \times V_2)).$$

A *biclique* K_b is a complete bipartite graph i.e. a bipartite graph where all the vertices are connected

$$K_b = (V', E') = (V'_1 \cup V'_2, E' = (V'_1 \times V'_2)).$$

Finally, the so-called maximum edge biclique problem in a bipartite graph $G_b = (V, E)$ is the problem of finding a biclique $K_b = (V', E')$ in G_b (i.e. $V' \subseteq V$ and $E' \subseteq E$) maximizing the number of edges. The decision problem: *Given B , does G_b contain a biclique with at least B edges?* has been shown to be NP-complete [26]. Therefore the corresponding optimization problem is at least NP-hard.

Let $M_b \in \{0, 1\}^{m \times n}$ be the adjacency matrix of the unweighted bipartite graph $G_b = (V_1 \cup V_2, E)$ i.e. $M_b(i, j) = 1$ if and only if $(V_1(i), V_2(j)) \in E$. In order to avoid trivialities, we will suppose that each vertex of the graph is connected to at least one other vertex i.e. $M_b(i, :) \neq \mathbf{0}, M_b(:, j) \neq \mathbf{0}, \forall i, j$. We denote by $|E|$ the cardinality of E i.e. the number of edges in G_b ; note that $|E| = \|M_b\|_F^2$. The set of zero values will be called $Z = \{(i, j) \mid M_b(i, j) = 0\}$, and its cardinality $|Z|$ satisfies $|E| + |Z| = mn$.

With this notation, the maximum biclique problem in G_b can be formulated as

$$\begin{aligned} \min_{v, w} & \|M_b - vw\|_F^2 \\ & v \in \{0, 1\}^m, w \in \{0, 1\}^n \\ & v_i w_j \leq M_b(i, j), \quad \forall i, j \end{aligned} \quad (\text{MBP})$$

⁴This terminology has already been used for the problem of finding a symmetric nonnegative factorization, i.e. one where $V=W$, but we assign it a different meaning in this paper.

In fact, one can check easily that this objective is equivalent to $\max_{v,w} \sum_{ij} v_i w_j$ since M_b , v and w are binary: instead of maximizing the number of edges inside the biclique, one minimizes the number of edges outside.

Feasible solutions of (MBP) correspond to bicliques of G_b . We will be particularly interested in *maximal* bicliques. A maximal biclique is a biclique which is not contained in a larger biclique: it is a locally optimal solution of (MBP).

The corresponding rank-one Nonnegative Factorization problem is defined as

$$\begin{aligned} \min_{v \in \mathbb{R}^m, w \in \mathbb{R}^n} & \|M_d - vw\|_F^2 \\ & v \geq \mathbf{0}, w \geq \mathbf{0} \end{aligned} \quad (\text{NF-1d})$$

with the matrix M_d defined as

$$M_d = (1 + d)M_b - d \mathbf{1}_{m \times n}, \quad d > 0 \quad (3.1)$$

where $\mathbf{1}_{m \times n}$ is the matrix of all ones with dimension $m \times n$. M_d is the matrix M_b where the zero values have been replaced by $-d$. Clearly M_d is not necessarily a nonnegative matrix.

To prove NP-hardness of (NF-1d), we are going to show that, if d is sufficiently large, optimal solutions of (NF-1d) *coincide* with optimal solutions of the corresponding biclique problem (MBP). From now on, we say that a solution (v, w) coincides with another solution (v', w') if and only if $vw = v'w'$ (i.e. if and only if $v' = \lambda v$ and $w' = \lambda^{-1}w$ for some $\lambda > 0$). We also let $M_+ = \max(\mathbf{0}, M)$ and $M_- = \max(\mathbf{0}, -M)$.

Lemma 1. *Any optimal rank-one approximation with respect to the Frobenius norm of a matrix M for which $\min(M) \leq -\|M_+\|_F$ contains at least one nonpositive entry.*

Proof. If $M = \mathbf{0}$, the result is trivial. If not, $\min(M) < 0$ since $\min(M) \leq -\|M_+\|_F$. Suppose now $(v, w) > 0$ is a best rank-one approximation of M . Therefore, since the negative values of M are approximated by positive ones and since M has at least one negative entry, we have

$$\|M - vw\|_F^2 > \|M_-\|_F^2. \quad (3.2)$$

By the Eckart-Young theorem,

$$\|M - vw\|_F^2 = \|M\|_F^2 - \sigma_{\max}(M)^2 = \|M\|_F^2 - \|M\|_2^2,$$

where $\sigma_{\max}(M)$ is the maximum singular value of M . Clearly,

$$\|M\|_F^2 = \|M_+\|_F^2 + \|M_-\|_F^2 \quad \text{and} \quad \|M\|_2^2 \geq \min(M)^2.$$

So

$$\|M - vw\|_F^2 \leq \|M_+\|_F^2 + \|M_-\|_F^2 - \min(M)^2 \leq \|M_-\|_F^2$$

which is in contradiction with (3.2). $\square$

We restate here a well-known result concerning low-rank approximations (see e.g. [16, p. 29]).

Lemma 2. *The local minima of the best rank-one approximation problem with respect to the Frobenius norm are global minima.*

We can now state the main result about the equivalence of (NF-1d) and (MBP).

Theorem 4. For $d \geq \sqrt{|E|}$, any optimal solution (v, w) of (NF-1d) coincides with an optimal solution of (MBP), i.e. vw is binary and $vw \leq M_b$.

Proof. We focus on the entries of vw which are positive and define

$$K = \left\{ i \in \{1, 2, \dots, m\} \mid v_i > 0 \right\} \quad \text{and} \quad L = \left\{ j \in \{1, 2, \dots, n\} \mid w_j > 0 \right\}. \quad (3.3)$$

$v' = v(K)$, $w' = w(L)$ and $M'_d = M_d(K, L)$ are the submatrices with indexes in (K, L) . Since (v, w) is optimal for M_d , (v', w') must be optimal for M'_d . Suppose there is a $-d$ entry in M'_d , then

$$\min(M'_d) = -d \leq -\sqrt{|E|} = -\|(M_d)_+\|_F \leq -\|(M'_d)_+\|_F,$$

so that Lemma 1 holds for M'_d . Since (v', w') is positive and is an optimal solution of (NF-1d) for M'_d , (v', w') is a local minimum of the unconstrained problem i.e. the problem of best rank-one approximation. By Lemma 2, this must be a global minimum. This is a contradiction with Lemma 1: (v', w') should contain at least one nonpositive entry. Therefore, $M'_d = \mathbf{1}_{|K| \times |L|}$ which implies $v'w' = M'_d$ by optimality and then vw is binary and $vw \leq M_b$. $\square$

Corollary 1. Rank-one Nonnegative Factorization is NP-hard.

Intuitively, the reason (NF-1d) is NP-hard is that if one of the $-d$ entries of M_d is approximated by a positive value, say p , the corresponding error is $d^2 + 2\mathbf{p}d + p^2$. Therefore, the larger d , the more expensive it is to approximate $-d$ by a positive number. Because of that, when d increases, negatives values of M_d will be approximated by smaller values and eventually by zeros.

Therefore, for each negative entry of M , one has to decide whether to approximate it with zero or with a positive value. Moreover, when a value is approximated by zero, one has to choose which entries of V and W will be equal to zero, as in the biclique problem. We suspect that the hardness of Nonnegative Factorization lies in these combinatorial choices.

We can now answer our initial question: Would it be possible to solve efficiently the problem

$$V_{:k}W_{k:} \approx R_k = M - \sum_{i \neq k} V_{:i}W_{i:} \not\approx \mathbf{0}$$

simultaneously for both vectors $(V_{:k}, W_{k:})$? Obviously, unless $P=NP$, we won't be able to find a polynomial-time algorithm to solve this problem. Therefore, it seems hopeless to improve the HALS algorithm using this approach.

Remark 1. Corollary 1 suggests that NMF is a difficult problem for any fixed $r \geq 2$. Indeed, even if one was given the optimal solution of an NMF problem except for one rank-one factor, it is not guaranteed that one would be able to find this last factor in polynomial-time since the corresponding residue is not necessarily nonnegative.

We now generalize Theorem 1 to factorizations of arbitrary rank.

Theorem 5. Nonnegative Factorization (NF) is NP-hard.

Proof. Let $M_b \in \{0, 1\}^{m \times n}$ be the adjacency matrix of a bipartite graph G_b and $r \geq 1$ the factorization rank of (NF). We define the matrix A_b as

$$A_b = \text{diag}(M_b, r) = \begin{pmatrix} M_b & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & M_b & & \mathbf{0} \\ \vdots & & \ddots & \vdots \\ \mathbf{0} & \dots & & M_b \end{pmatrix}$$

which is the adjacency matrix of another bipartite graph G'_b which is nothing but the graph G_b repeated r times. A_d is defined in the same way as M_d , i.e.

$$A_d = (1 + d)A_b - d \mathbf{1}_{m \times n}$$

with $d \geq \sqrt{r|E|}$. Let (V, W) be the optimal rank- r nonnegative factorization of A_d and consider each rank-one factor $V_{:k}W_{k:} \approx R_k = A_d - \sum_{i \neq k} V_{:i}W_{i:}$: each of them must clearly be an optimal rank-one nonnegative factorization of R_k . Since $R_k \leq A_d$,

$$\min(R_k) \leq \min(A_d) = -d \leq -\|(A_d)_+\|_F \leq -\|(R_k)_+\|_F,$$

and Lemma 1 holds. Using exactly the same arguments as in Theorem 4, one can show that, $\forall k$,

$$(V_{:k}W_{k:})_{ij} = 0, \quad \forall (i, j) \text{ s.t. } A_d(i, j) = -d.$$

Therefore, the positive entries of each rank-one factor will correspond to a biclique of G'_b . By optimality of (V, W) , each rank-one factor must correspond to a maximum biclique of G_b since G'_b is the graph G_b repeated r times. Thus (NF-1d) is NP-hard implies (NF) is NP-hard. $\square$

3.2 Stationary points of (NF-1d)

We have shown that optimal solutions of (NF-1d) coincide with optimal solutions of (MBP) for $d \geq \sqrt{E}$, which are NP-hard to find. In this section, we focus on stationary points of (NF-1d) instead: we show how they are related to the feasible solutions of (MBP). This result will be used in Section 5 to design a new type of biclique finding algorithm.

3.2.1 Definitions and Notation

The KKT conditions of (NF-1d), which define the stationary points, are exactly the same as for (NMF): (v, w) is a stationary point of (NF-1d) if and only if

$$v \geq \mathbf{0}, \quad \mu = (vw - M_d)w^T \geq \mathbf{0} \quad \text{and} \quad v \circ \mu = \mathbf{0} \quad (3.4)$$

$$w \geq \mathbf{0}, \quad \lambda = v^T(vw - M_d) \geq \mathbf{0} \quad \text{and} \quad w \circ \lambda = \mathbf{0}. \quad (3.5)$$

Of course, we are especially interested in nontrivial solutions and we then assume $v, w \neq \mathbf{0}$ so that one can check that (3.4)-(3.5) are equivalent to

$$v = \max\left(\mathbf{0}, \frac{M_d w^T}{\|w\|_2^2}\right) \quad \text{and} \quad w = \max\left(\mathbf{0}, \frac{v^T M_d}{\|v\|_2^2}\right). \quad (3.6)$$

Given d , we define three sets of rank-one matrices: S_d , corresponding to the nontrivial stationary points of (NF-1d), with

$$S_d = \{vw \in \mathbb{R}_0^{m \times n} \mid (v, w) \text{ satisfy (3.6)}\},$$

F , corresponding to the feasible solutions of (MBP), with

$$F = \{vw \in \mathbb{R}^{m \times n} \mid (v, w) \text{ is a feasible of (MBP)}\},$$

and B , corresponding to the maximal bicliques of (MBP), i.e. $vw \in B$ if and only if $vw \in F$ and vw coincides with a maximal biclique.

3.2.2 Stationarity of Maximal Bicliques

The next theorem states that, for d sufficiently large, the only nontrivial feasible solutions of (MBP) that are stationary points of (NF-1d) are the maximal bicliques.

Theorem 6. For $d > \max(m, n) - 1$, $F \cap S_d = B$.

Proof. $vw \in B$ if and only if $vw \in F$ and is maximal i.e.

$$(1) \nexists i \text{ such that } v_i = 0 \text{ and } M_d(i, j) = 1, \forall j \text{ s.t. } w_j \neq 0,$$

$$(2) \nexists j \text{ such that } w_j = 0 \text{ and } M_d(i, j) = 1, \forall i \text{ s.t. } v_i \neq 0.$$

Since vw is binary and $v \neq \mathbf{0}$, the nonzero entries of w must be equal to each other. Moreover, $d > \max(m, n) - 1$ so that (1) is equivalent to

$$\begin{aligned} \nexists i \text{ such that } v_i = 0 \text{ and } M_d(i, :)w^T > 0 \\ \iff \\ v_i = 0 \Rightarrow M_d(i, :)w^T < 0 \text{ and } v_i \neq 0 \Rightarrow v_i = \frac{\|M_d(i, :)\|_1}{\|w\|_1} = \frac{M_d(i, :)w^T}{\|w\|_2^2}. \end{aligned}$$

This is equivalent to the stationarity conditions for $v \neq \mathbf{0}$. By symmetry, (2) is equivalent to the stationarity conditions for w . $\square$

Theorem 6 implies that, for d sufficiently large, $B \subset S_d$. It would be interesting to have the opposite affirmation: for d sufficiently large, any stationary point of (NF-1d) corresponds to a maximal biclique of (MBP). Unfortunately, we will see later that this property does not hold.

3.2.3 Limit points of S_d

However, as d goes to infinity, we are going to show that the points in S_d get closer to feasible solutions of (MBP).

Lemma 3. The set S_d is bounded i.e. $\forall d > 0, \forall vw \in S_d$:

$$\|vw\|_2 = \|v\|_2 \|w\|_2 \leq \sqrt{|E|}.$$

Proof. For $vw \in S_d$, by (3.6),

$$\|v\|_2 = \left\| \max\left(\mathbf{0}, \frac{M_d w^T}{\|w\|_2^2}\right) \right\|_2 \leq \frac{\|\max(\mathbf{0}, M_d)w^T\|_2}{\|w\|_2^2} \leq \frac{\|\max(\mathbf{0}, M_d)\|_F}{\|w\|_2} = \frac{\sqrt{|E|}}{\|w\|_2}.$$

$\square$

Lemma 4. For $vw \in S_d$, if $M_d(i, j) = -d$ and if $(vw)_{ij} > 0$, then

$$0 < v_i < \frac{\|v\|_1}{d+1} \quad \text{and} \quad 0 < w_j < \frac{\|w\|_1}{d+1}.$$

Proof. By (3.6),

$$0 < w_j \|v\|_2^2 = v^T M_d(:, j) \leq \|v\|_1 - (d+1)v_i \Rightarrow 0 < v_i < \frac{\|v\|_1}{d+1}.$$

The same can be shown for w by symmetry. $\square$

Theorem 7. As d goes to infinity, stationary points of (NF-1d) get closer to feasible solutions of (MBP) i.e. $\forall \epsilon > 0, \exists D$ s.t. $\forall d > D$:

$$\max_{vw \in S_d} \min_{v_b w_b \in F} \|vw - v_b w_b\|_F < \epsilon. \quad (3.7)$$

Proof. Let $vw \in S_d$ and suppose $vw > 0$. W.l.o.g. $\|w\|_2 = 1$; in fact, if $vw \in S_d$, $(\lambda v \frac{1}{\lambda} w) \in S_d, \forall \lambda > 0$. Note that Lemma 3 implies $\|v\|_2 \leq \sqrt{|E|}$. By (3.6),

$$v = M_d w^T \quad \text{and} \quad w = \frac{v^T M_d}{\|v\|_2^2}.$$

Therefore, $(v/\|v\|_2, w) > 0$ is a pair of singular vectors of M_d associated with the singular value $\|v\|_2 > 0$. If $M_d = \mathbf{1}_{m \times n}$, the only pair of positive singular vectors of M_d is $(\frac{1}{\sqrt{m}} \mathbf{1}_m, \frac{1}{\sqrt{n}} \mathbf{1}_n)$ so that $vw = M_b$ coincides with a feasible solution of (MBP).

Otherwise, we define

$$A = \{i \mid M_d(i, j) = 1, \forall j\} \quad \text{and} \quad B = \{j \mid M_d(i, j) = 1, \forall i\}, \quad (3.8)$$

and their complements $\bar{A} = \{1, 2, \dots, m\} \setminus A$, $\bar{B} = \{1, 2, \dots, n\} \setminus B$; hence,

$$M_d(A, :) = \mathbf{1}_{|A| \times n} \quad \text{and} \quad M_d(:, B) = \mathbf{1}_{m \times |B|}.$$

Using Lemma 4 and the fact that $\|x\|_1 \leq \sqrt{n} \|x\|_2, \forall x \in \mathbb{R}^n$, we get

$$\mathbf{0} < v(\bar{A}) < \frac{\sqrt{m|E|}}{d+1} \mathbf{1}_{|\bar{A}|} \quad \text{and} \quad \mathbf{0} < w(\bar{B}) < \frac{\sqrt{n}}{d+1} \mathbf{1}_{|\bar{B}|}. \quad (3.9)$$

Therefore, since $w \leq \mathbf{1}_n$ and $v \leq \sqrt{|E|} \mathbf{1}_m$, we obtain

$$\|v(\bar{A})w - \mathbf{0}\|_F < \frac{1}{d+1} (m\sqrt{n|E|}) \quad \text{and} \quad \|vw(\bar{B}) - \mathbf{0}\|_F < \frac{1}{d+1} (n\sqrt{m|E|}).$$

It remains to show that $v(A)w(B)$ coincide with a biclique of the (complete) graph generated by $M_b(A, B) = \mathbf{1}_{|A| \times |B|}$. We distinguish three cases:

(1) $A = \emptyset$. (3.9) implies $\|v\|_2 < \frac{1}{d+1} (m\sqrt{m|E|})$ so that

$$w(B) = \frac{v^T \mathbf{1}_{m \times |B|}}{\|v\|_2^2} = \frac{\|v\|_1}{\|v\|_2^2} \mathbf{1}_{|B|} \geq \frac{\sqrt{m}}{\|v\|_2} \mathbf{1}_{|B|} > \frac{d+1}{m\sqrt{|E|}} \mathbf{1}_{|B|} \quad (3.10)$$

which is absurd if $d > m\sqrt{|E|}$ since $\|w\|_2 = 1$.

(2) $B = \emptyset$. Using (3.9), we have $v(A) = M_d(A, :)w^T = \|w\|_1 \mathbf{1}_{|A|} < n \frac{\sqrt{n}}{d+1}$ and then

$$\|v(A)w(B) - \mathbf{0}\|_F < \frac{1}{d+1} (n^2 \sqrt{m}).$$

(3) $A, B \neq \emptyset$. Noting $k_w = \frac{\|v\|_1}{\|v\|_2^2}$, Equation (3.10) gives $w(B) = k_w \mathbf{1}_{|B|}$. Therefore,

$$1 - |\bar{B}| \frac{\sqrt{n}}{d+1} < \|w\|_2^2 - \|w(\bar{B})\|_2^2 = \|w(B)\|_2^2 = |B| k_w^2 \leq \|w\|_2^2 = 1, \quad (3.11)$$

Moreover, $v(A) = \mathbf{1}_{|A| \times m} w^T = \|w\|_1 \mathbf{1}_{|A|}$ so that

$$|B| k_w \leq v(A) = (\|w(B)\|_1 + \|w(\bar{B})\|_1) \mathbf{1}_{|A|} < |B| k_w + |\bar{B}| \frac{\sqrt{n}}{d+1}. \quad (3.12)$$

Finally, combining (3.11) and (3.12) and noting that $k_w \leq 1$ since $\|w\|_2 = 1$,

$$\left(1 - \frac{|\bar{B}|\sqrt{n}}{d+1}\right) \mathbf{1}_{|A| \times |B|} < v(A)w(B) < \left(1 + \frac{|\bar{B}|\sqrt{n}}{d+1}\right) \mathbf{1}_{|A| \times |B|}.$$

We can conclude that, for d sufficiently large, vw is arbitrarily close to a feasible solution of (MBP) which corresponds to the biclique (A, B) .

Recall we supposed $vw > 0$. If $vw \not> 0$, let (K, L) be the indexes defined in (3.3). The above result holds for $v(K)w(L) > 0$ with the matrix $M_d(K, L)$. For d sufficiently large, $v(K)w(L)$ is then close to a feasible solution of (MBP) for $M_d(K, L)$. Adding zero to this feasible solution gives a feasible solution for M_b . $\square$

Example 1. *Let*

$$M_d = \begin{pmatrix} -d & 1 \\ 1 & 1 \end{pmatrix}.$$

Clearly, $\begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix}$ belongs to the set B , i.e. it corresponds to maximal biclique of the graph generated by M_b . By Theorem 6, for $d > 1$, it belongs to S_d i.e. $[(0 \ 1)^T, (0 \ 1)]$ is stationary points of (NF-1d). For $d > 1$, one can also check that the singular values of M_d are disjoint and that the second pair of singular vectors is positive. Since it is a positive stationary point of the unconstrained problem, it is also a stationary point of (NF-1d). As d goes to infinity, it must get closer to a biclique of (MBP) (Theorem 7). Moreover M_d is symmetric so that the right and left singular vectors are equal to each other. Figure 2 shows the evolution⁵ of this positive singular vector of M_d with respect to d . It converges to $(0 \ 1)$ and then the product of the left and right singular vector converges to $\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \in F$.




Figure 2: Evolution of $(v_1 v_2)^T (v_1 v_2) \in S_d$.

⁵By Wedin's theorem (cf. matrix perturbation theory, see e.g. [28]), singular subspaces of M_d associated with a positive singular value are continuously deformed with respect to d .

3.3 Multiplicative Updates for Nonnegative Factorization

In this section, the MU of Lee and Seung presented in Section 2.2 to find approximate solutions of (NMF) are generalized to (NF). Other than providing a way of computing approximate solutions of (NF), this result will also help us to understand why the updates of Lee and Seung are not very efficient in practice.

The Karush-Kuhn-Tucker optimality conditions of the (NF) problem are the same as for (NMF) (see Section 2.2). Of course, any real matrix M can be written as the difference of two nonnegative matrices: $M = P - N$ with $P, N \geq \mathbf{0}$. This can be used to generalize the algorithm of Lee and Seung. In fact, (2.5) and (2.6) become

$$V \circ (VWW^T + NW^T) = V \circ (PW^T) \quad (3.13)$$

$$W \circ (V^T VW + V^T N) = W \circ (V^T P) \quad (3.14)$$

and using the same idea as in Section 2.2, we get the following multiplicative update rules :

Theorem 8. *For $V, W \geq \mathbf{0}$ and $M = P - N$ with $P, N \geq \mathbf{0}$, the cost function $\|M - VW\|_F$ is nonincreasing under the following update rules:*

$$V \leftarrow V \circ \frac{[PW^T]}{[VWW^T + NW^T]}, \quad W \leftarrow W \circ \frac{[V^T P]}{[V^T VW + V^T N]}. \quad (3.15)$$

Proof. We only treat the proof for V since the problem is perfectly symmetric. The cost function can be split into m independent components related to each row of the error matrix, each depending on a specific row of P , N and V , which we call respectively p , n and v . Hence, we can treat each row of V separately, and we only have to show that the function

$$F(v) = \frac{1}{2} \|p - n - vW\|_F^2.$$

is nonincreasing under the following update

$$v_0 \leftarrow v_0 \circ \frac{[pW^T]}{[v_0WW^T + nW^T]}, \quad \forall v_0 > 0. \quad (3.16)$$

F is a quadratic function so that

$$F(v) = F(v_0) + (v - v_0)\nabla F(v_0) + \frac{1}{2}(v - v_0)\nabla^2 F(v_0)(v - v_0)^T, \quad \forall v_0,$$

with $\nabla F(v_0) = (-p + n + v_0W)W^T$ and $\nabla^2 F(v_0) = WW^T$. Let G be a quadratic model of F around v_0 :

$$G(v) = F(v_0) + (v - v_0)\nabla F(v_0) + \frac{1}{2}(v - v_0)K(v_0)(v - v_0)^T$$

with $K(v_0) = \text{diag}\left(\frac{[v_0WW^T + nW^T]}{[v_0]}\right)$. G has the following nice properties (see below):

- (1) G is an upper approximation of F i.e. $G(v) \geq F(v), \forall v$;
- (2) The global minimum of $G(v)$ is nonnegative and given by (3.16).

Therefore, the global minimum of G , given by (3.16), provides a new iterate which guarantee the monotonicity of F . In fact,

$$F(v_0) = G(v_0) \geq \min_v G(v) = G(v^*) \geq F(v^*).$$

It remains to show that (1) and (2) hold.

(1) $G(v) \geq F(v) \forall v$. This is equivalent to $K(v_0) - WW^T$ positive semidefinite (PSD). Lee and Seung have proved [22] that $A = \text{diag}\left(\frac{[v_0 WW^T]}{[v_0]}\right) - WW^T$ is PSD (see also [17]). Since $B = \text{diag}\left(\frac{[nW^T]}{[v_0]}\right)$ is a diagonal nonnegative matrix for $v_0 > 0$ and $nW^T \geq 0$, $A + B = K(v_0) - WW^T$ is also PSD.

(2) The global minimum of G is given by (3.16):

$$\begin{aligned} v^* = \text{argmin}_v G(v) &= v_0 - K^{-1}(v_0) \nabla F(v_0) \\ &= v_0 - v_0 \circ \frac{[-pW^T + (v_0 WW^T + nW^T)]}{[v_0 WW^T + nW^T]} \\ &= v_0 \circ \frac{[pW^T]}{[v_0 WW^T + nW^T]}. \end{aligned}$$

□

As with standard multiplicative updates, convergence can be guaranteed with a simple modification:

Theorem 9. *For every constant $\epsilon > 0$ and for $M = P - N$ with $P, N \geq 0$, $\|M - VW\|_F$ is nonincreasing under*

$$V \leftarrow \max\left(\epsilon, V \circ \frac{[PW^T]}{[VWW^T + NW^T]}\right), \quad W \leftarrow \max\left(\epsilon, W \circ \frac{[V^T P]}{[V^T V W + V^T N]}\right) \quad (3.17)$$

for any $(V, W) \geq \epsilon$. Moreover, every limit point of this algorithm is a stationary point of the optimization problem (2.9).

Proof. We use exactly the same notation as in the proof of Theorem 3.15, so that

$$F(v_0) = G(v_0) \geq \min_{v \geq \epsilon} G(v) = G(v^*) \geq F(v^*), \quad v_0 \geq \epsilon$$

remains valid. By definition, $K(v_0)$ is a diagonal matrix implying that $G(v)$ is the sum of r independent quadratic terms, each depending on a single entry of v . Therefore,

$$\text{argmin}_{v \geq \epsilon} G(v) = \max\left(\epsilon, v_0 \circ \frac{[pW^T]}{[v_0 WW^T + nW^T]}\right),$$

and the monotonicity is proved.

Let $(\bar{V}, \bar{W})$ be a limit point of a sequence $\{(V^k, W^k)\}$ generated by (3.17). The monotonicity implies that $\{\|M - V^k W^k\|_F\}$ converges to $\|M - \bar{V} \bar{W}\|_F$ since the cost function is bounded from below. Moreover,

$$\bar{V}_{ik} = \max\left(\epsilon, \alpha_{ik} \bar{V}_{ik}\right), \quad \forall i, k \quad (3.18)$$

where

$$\alpha_{ik} = \frac{P_{i:} \bar{W}_{k:}^T}{\bar{V}_{i:} \bar{W} \bar{W}_{k:}^T + N_{i:} \bar{W}_{k:}^T},$$

which is well-defined since $\bar{V}_{i:} \bar{W} \bar{W}_{k:}^T > 0$. One can easily check that the stationarity conditions of (2.9) for $\bar{V}$ are

$$\bar{V}_{ik} \geq \epsilon, \quad \alpha_{ik} \leq 1 \quad \text{and} \quad (\bar{V}_{ik} - \epsilon)(\alpha_{ik} - 1) = 0, \quad \forall i, k.$$

Finally, by (3.18), we have either $\bar{V}_{ik} = \epsilon$ and $\alpha_{ik} \leq 1$, or $\bar{V}_{ik} > \epsilon$ and $\alpha_{ik} = 1, \forall i, k$. The same can be done for $\bar{W}$ by symmetry. □

In order to implement the updates (3.15), one has to choose the matrices P and N . It is clear that $\forall P, N \geq \mathbf{0}$ such that $M = P - N$, there exists a matrix $C \geq \mathbf{0}$ such that the two components P and N can be written $P = M_+ + C$ and $N = M_- + C$. When C goes to infinity, the above updates do not change the matrices V and W , which seems to indicate that smaller values of C are preferable. Indeed, in the case $r = 1$, one can prove that $C = \mathbf{0}$ is an optimal choice:

Theorem 10. $\forall P, N \geq \mathbf{0}$ s.t. $M = P - N$, and $\forall v \in \mathbb{R}_+^n, w \in \mathbb{R}_+^m$:

$$\|M - v_1 w\|_F \leq \|M - v_2 w\|_F \leq \|M - v w\|_F, \quad (3.19)$$

for

$$v_1 = v \circ \frac{[M_+ w^T]}{[v w w^T + M_- w^T]} \quad \text{and} \quad v_2 = v \circ \frac{[P w^T]}{[v w w^T + N w^T]}.$$

Proof. The second inequality of (3.19) is a consequence of Theorem 3.15. For the first one, we treat the inequality separately for each entry of v i.e. we prove that

$$\|M_{i:} - v_{1i} w\|_F \leq \|M_{i:} - v_{2i} w\|_F, \quad \forall i.$$

Let define v_i^* as the optimal solution of the unconstrained problem i.e.

$$v_i^* = \operatorname{argmin}_{v_i} \|M_{i:} - v_i w\|_F = \frac{M_{i:} w^T}{w w^T},$$

and $a, b, d \geq 0, e > 0$, as

$$a = (M_+)_{i:} w^T, \quad b = (M_-)_{i:} w^T, \quad d = (P - M_+)_{i:} w^T \quad \text{and} \quad e = v_i w w^T.$$

Noting that $P - M_+ = N - M_-$, we have

$$v_i^* = v_i \left(\frac{a - b}{e} \right), \quad v_{1i} = v_i \left(\frac{a}{e + b} \right) \quad \text{and} \quad v_{2i} = v_i \left(\frac{a + d}{e + b + d} \right).$$

Suppose $v_i^* \geq v_i$. Therefore,

$$\frac{a - b}{e} \geq 1 \Rightarrow a - b - e \geq 0 \Rightarrow \frac{a - b}{e} \geq \frac{a}{e + b}.$$

Moreover, $1 \leq \frac{a + d}{e + b + d}$ since v_{2i} is a better solution than v_i (Theorem 3.15). Finally,

$$1 \leq \frac{a + d}{e + b + d} \leq \frac{a}{e + b} \leq \frac{a - b}{e} \Rightarrow v_i \leq v_{2i} \leq v_{1i} \leq v_i^*.$$

The case $v_i^* \leq v_i$ is similar. □

Unfortunately, this result does not hold for $r > 1$. This is even true for nonnegative matrices, i.e. one can improve the effect of a standard Lee and Seung multiplicative update by using a well-chosen matrix C .

Example 2. With the following matrices

$$M = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 0 \end{pmatrix}, \quad V = \begin{pmatrix} 1 & 1 \\ 1 & 0 \\ 1 & 1 \end{pmatrix}, \quad W = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 1 \end{pmatrix} \quad \text{and} \quad C = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix},$$

we have $\|M - V'W\|_F < \|M - V''W\|_F$ where V' (resp. V'') is updated following (3.15) using $P = M + C$ and $N = C$ (resp. $P = M$ and $N = \mathbf{0}$).

However, in practice, it seems that the choice of a proper matrix C is nontrivial and cannot accelerate significantly the speed of convergence.

4 How good are the Multiplicative Updates (MU) of Lee and Seung?

In this section, we use Theorem 8 to interpret the multiplicative rules for (NMF) and show why the HALS algorithm performs much better in practice.

4.1 An Improved Version of the MU

The aim of the MU is to improve a current solution $(V, W) \geq \mathbf{0}$ by optimizing alternatively V (W fixed), and vice-versa. In order to prove the monotonicity of the MU, $\|M - VW\|_F$ was shown to be nonincreasing under an update of a single row of V (resp. column of W) since the objective function can be split into m (resp. n) independent quadratic terms, each depending on the entries of a row of V (resp. column of W); cf. proof of Theorem 8.

However, there is no guarantee, a priori, that the algorithm is also nonincreasing with respect to an individual update of a column of V (resp. row of W). In fact, each entry of a column of V (resp. row of W) depends on the other entries of the same row (resp. column) in the cost function. The next theorem states that this property actually holds.

Corollary 2. *For $V, W \geq \mathbf{0}$, $\|M - VW\|_F$ is nonincreasing under*

$$V_{:k} \leftarrow V_{:k} \circ \frac{[MW_{k:}^T]}{[VW_{k:}W_{k:}^T]}, \quad W_{k:} \leftarrow W_{k:} \circ \frac{[V_{:k}^T M]}{[V_{:k}^T VW]}, \quad \forall k, \quad (4.1)$$

i.e. under the update of any column of V or any row of W using the MU (2.7).

Proof. This is a consequence of Theorem 8 using $P = M$ and $N = \sum_{i \neq k} V_{:i} W_{i:}$.

In fact, $\|M - VW\|_F = \|(M - \sum_{i \neq k} V_{:i} W_{i:}) - V_{:k} W_{k:}\|_F$. □

Corollary 2 sheds light on a very interesting fact: the multiplicative updates are also trying to optimize alternatively the columns of V and the rows of W using a specific cyclic order: first, the columns of V and then the rows of W . We can now point out two ways of improving the MU:

1. When updating a column of V (resp. a row of W), the columns (resp. rows) already updated are not taken into account: the algorithm uses their old values;
2. The multiplicative updates are not optimal: $P \neq M_+$ and $N \neq M_-$ (cf. Theorem 10). Moreover, there is actually a closed-form solution for these subproblems (cf. HALS algorithm, Section 2.3).

Therefore, using Theorem 10, we have the following new improved updates

Corollary 3. *For $V, W \geq \mathbf{0}$, $\|M - VW\|_F$ is nonincreasing under*

$$V_{:k} \leftarrow V_{:k} \circ \frac{[(R_k)_+ W_{k:}^T]}{[V_{:k} W_{k:} W_{k:}^T + (R_k)_- W_{k:}^T]}, \quad W_{k:} \leftarrow W_{k:} \circ \frac{[V_{:k}^T (R_k)_+]}{[V_{:k}^T V_{:k} W_{k:} + V_{:k}^T (R_k)_-]}, \quad \forall k, \quad (4.2)$$

with $R_k = M - \sum_{i \neq k} V_{:i} W_{i:}$. Moreover, the updates (4.2) perform better than the updates (4.1), but worse than the updates (2.12-2.13) which are optimal.

Proof. This is a consequence of Theorem 8 and 10 using $P = (R_k)_+$ and $N = (R_k)_-$.

In fact, $\|M - VW\|_F = \|R_k - V_{:k} W_{k:}\|_F$. □

4.2 An example

Figure 3 shows an example of the behavior of the different algorithms: the original MU (Section 2.2), the improved version (Corollary 3) and the *optimal* HALS method (Section 2.3). The test was carried out on a commonly used data set for NMF: the cbcl face database⁶; 2429 faces (columns) consisting each of 19×19 pixels (rows) for which we set $r = 40$ and we used the same *scaled* (see Remark 3 below) random initialization and the same cyclic order (same as the MU i.e. first the columns of V then the rows of W) for the three algorithms. We observe that the MU converges significantly less




Figure 3: Comparison of the MU of Lee and Seung (2.7), the optimal rank-one NF multiplicative updates (4.2) and the HALS (2.12-2.13) applied to the cbcl face database.

rapidly than the two other algorithms. There do not seem to be good reasons to use either the MU or the method of Corollary 3 since there is a closed-form solution (2.12-2.13) for the corresponding subproblems.

Finally, the HALS algorithm has the same computational complexity [16] and performs provably much better than the popular multiplicative updates of Lee and Seung. Of course, because of the NP-hardness of (NMF) and the existence of numerous locally optimal solutions, it is not possible to give a theoretical guarantee that HALS will converge to a better solution than the MU: although its iterations are *locally* more efficient, they could still end up at a worse local optimum.

Remark 2. For $r = 1$, one can check that the three algorithms above are equivalent for (NMF). Moreover, they correspond to the power method [13] which converges to the optimal rank-one solution, given that it is initialized with a vector which is not perpendicular to the singular vector corresponding to the maximum singular value.

Remark 3. We say that (V, W) is scaled if the optimal solution to the problem

$$\min_{\alpha \in \mathbb{R}} \|M - \alpha VW\|_F \quad (4.3)$$

is equal to 1. Obviously, any stationary point is scaled ; the next Theorem is an extension of a result of Ho et al. [18].

⁶CBCL Face Database #1, MIT Center For Biological and Computation Learning. Available at <http://cbcl.mit.edu/cbcl/software-datasets/FaceData2.html>.

Theorem 11. *The following statements are equivalent*

- (1) (V, W) is scaled;
- (2) VW is on the boundary of $\mathcal{B}\left(\frac{M}{2}, \frac{1}{2}\|M\|_F\right)$, the ball centered at $\frac{M}{2}$ of radius $\frac{1}{2}\|M\|_F$;
- (3) $\|M - VW\|_F^2 = \|M\|_F^2 - \|VW\|_F^2$ (and then $\|M\|_F^2 \geq \|VW\|_F^2$).

Proof. The solution of (4.3) can be written in the following closed form

$$\alpha = \frac{\langle M, VW \rangle}{\langle VW, VW \rangle}, \quad (4.4)$$

where $\langle A, B \rangle = \sum_{ij} A_{ij} B_{ij} = \text{trace}(AB^T)$ is the scalar product associated with the Frobenius norm. Since $\alpha = 1$,

$$\begin{aligned} \langle VW - M, VW \rangle &= 0 \\ \langle VW - M, VW \rangle + \left\langle \frac{M}{2}, \frac{M}{2} \right\rangle &= \left\langle \frac{M}{2}, \frac{M}{2} \right\rangle \\ \left\langle \frac{M}{2} - VW, \frac{M}{2} - VW \right\rangle &= \left\langle \frac{M}{2}, \frac{M}{2} \right\rangle, \end{aligned}$$

so that (1) and (2) are equivalent. For the equivalence of (1) and (3), we have

$$\begin{aligned} \langle M - VW, M - VW \rangle &= \|M\|_F^2 - 2\langle M, VW \rangle + \|VW\|_F^2 \\ &= \|M\|_F^2 - \|VW\|_F^2 - 2\left(\langle M, VW \rangle - \langle VW, VW \rangle\right) \\ &= \|M\|_F^2 - \|VW\|_F^2 \end{aligned}$$

if and only if $\langle M, VW \rangle = \langle VW, VW \rangle$. □

Theorem 11 can be used as follows: when you compute the error of the current solution, you can scale it without further computational cost. In fact,

$$\begin{aligned} \|M - VW\|_F^2 &= \langle M - VW, M - VW \rangle \\ &= \|M\|_F^2 - 2\langle M, VW \rangle + \|VW\|_F^2. \end{aligned} \quad (4.5)$$

Note that the third term of (4.5) can be computed in $O(\max(m, n)r^2)$ operations since

$$\begin{aligned} \|VW\|_F^2 &= \sum_{ij} \left(\sum_k V_{ik}^2 W_{kj}^2 \right) + 2 \sum_{ij} \left(\sum_{k \neq l} V_{ik} W_{kj} V_{il} W_{lj} \right) \\ &= \sum_k \left(\sum_i V_{ik}^2 \right) \left(\sum_j W_{kj}^2 \right) + 2 \sum_{k \neq l} \left(\sum_i V_{ik} V_{il} \right) \left(\sum_j W_{kj} W_{lj} \right) \\ &= \|(V^T V) \circ (W W^T)\|_1 \end{aligned}$$

where $\|A\|_1 = \sum_{ij} |A_{ij}|$. This is especially interesting for sparse matrices since only a small number of the entries of VW (which could be dense) need to be computed to evaluate the second term of (4.5).

5 Biclique Finding Algorithm

In this section, an algorithm for the maximum edge biclique problem whose main iteration requires $O(|E|)$ operations is presented. It is based on the multiplicative updates for Nonnegative Factorization and the strong relation between these two problems (Theorems 4, 6 and 7). We compare the results with other algorithms with iterates requiring $O(|E|)$ operations using the DIMACS database and random graphs.

5.1 Description

For d sufficiently large, stationary points of (NF-1d) are close to bicliques of (MBP) (Theorem 7). Moreover, the two problems have the same cost function. One could then think of applying an algorithm that finds stationary points of (NF-1d) in order to localize a large biclique of the graph generated by M_b . This is the idea of Algorithm 1 using the multiplicative updates (3.15) with

$$P = (M_d)_+ = M_b \quad \text{and} \quad N = (M_d)_- = d(\mathbf{1}_{m \times n} - M_b).$$

A priori, it is not clear what value d should take. Following the spirit of homotopy methods, we chose to start the algorithm with a small value of d and then to increase it until the algorithm converges to a biclique of (MBP).

Algorithm 1 Biclique Finding Algorithm in $O(|E|)$ operations

Require: $M_b \in \{0, 1\}^{m \times n}$, $v \in \mathbb{R}_{++}^m$, $w \in \mathbb{R}_{++}^n$, $d = d_0 > 0$, $\alpha > 1$.

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

2:

$$v \leftarrow v \circ \frac{[M_b w^T]}{[v \|w\|_2^2 + d(\mathbf{1}_m \|w\|_1 - M_b w^T)]} \quad (5.1)$$

$$w \leftarrow w \circ \frac{[v^T M_b]}{[\|v\|_2^2 w + d(\mathbf{1}_n \|v\|_1 - v^T M_b)]} \quad (5.2)$$

$$d = \alpha d$$

3: **end for**

We observed that initial value of d should not be chosen too large: otherwise, the algorithm often converges to the trivial solution: the empty biclique. In fact, in that case, the denominators in (5.1) and (5.2) will be large, even during the initial steps of the algorithm, and the solution is then forced to converge to zero. Moreover, since the denominators in (5.1) and (5.2) depend on the graph density, the denser the graph is, the greater d_0 can be chosen and vice versa. On the other hand, since our algorithm is equivalent to the power method for $d = 0$ (cf. Remark 2), if d_0 is chosen too small, it will converge to the same solution: the one initialized with the best rank-one approximation of M_b .

For the stopping criterion, one could, for example, wait until the rounding of vw coincides with a feasible solution of (MBP).

5.2 Other Algorithms in $O(|E|)$ operations

We briefly present here two other algorithms to find maximal bicliques using $O(|E|)$ operations per iteration.

5.2.1 Motzkin-Strauss Formalism

In [12], the generalized Motzkin-Strauss formalism for cliques is extended to bicliques by defining the optimization problem

$$\max_{\mathbf{x} \in F_x^\alpha, \mathbf{y} \in F_y^\beta} \mathbf{x}^T M_b \mathbf{y}$$

where $F_x^\alpha = \{x \in \mathbb{R}_+^n \mid \sum_{i=1}^n x_i^\alpha = 1\}$, $F_y^\beta = \{y \in \mathbb{R}_+^m \mid \sum_{i=1}^m y_i^\beta = 1\}$ and $\alpha, \beta \ll 1$.

Nonincreasing multiplicative updates for this problem are then provided:

$$\mathbf{x} \leftarrow \left(\mathbf{x} \circ \frac{M_b \mathbf{y}}{\mathbf{x}^T M_b \mathbf{y}} \right)^{\frac{1}{\alpha}}, \quad \mathbf{y} \leftarrow \left(\mathbf{y} \circ \frac{M_b^T \mathbf{x}}{\mathbf{x}^T M_b \mathbf{y}} \right)^{\frac{1}{\beta}}$$

This algorithm does not necessarily converge to a biclique: if α and β are not sufficiently small, it may converge to a dense bipartite subgraph (a bicluster). In fact, for $\alpha = \beta = 2$, it converges to an optimal rank-one solution of the unconstrained problem as our algorithm does for $d = 0$. In [12], it is suggested to use α and β around 1.05. Finally, $\alpha \neq \beta$ will favor one side of the biclique. We will use $\alpha = \beta$.

5.2.2 Greedy Heuristic

The simplest heuristic one can imagine is to add, at each step, a vertex which is connected to most vertices in the other side of the bipartite graph. Each time a vertex is selected, the next choices are restricted in order to get a biclique eventually: the vertices which are not connected to the one you have just chosen are deleted. The procedure is repeated on the remaining graph until you get a biclique. One can check that this produces a maximal biclique.

5.3 Results

We first present some results for graphs from the DIMACS graph dataset⁷. We extracted bicliques in those (not bipartite) graphs using the preceding algorithms. We performed 100 runs, 200 iterations each, for the two algorithms with the same initializations. We tried to choose appropriate parameters⁸ for both algorithms. Table 1 displays the cardinality of the biclique extracted by the different algorithms.

Table 2 shows the results for random graphs: we have generated randomly 100 graphs with 100 vertices for different densities (the probability of an edge to belong to the graph is equal to the density). The average numbers of edges in the solutions for the different algorithms are displayed for each density. We kept the same configuration as for the DIMACS graphs (same initializations, 100 runs for each graph, 200 iterations). It seems that the multiplicative updates generates, in general, better solutions, especially when dealing with dense graphs. The algorithm based on the Motzkin-Strauss formalism seems less efficient and is more sensitive to the choice of its parameters.

	size (m)	Greedy	M.-S. ($\alpha = 1.01$)		Mult. ($d = 1, \alpha = 1.1$)	
			mean	best	mean	best
ham62	64	304	157	225	269	320
ham64	64	42	21	36	37	42
ham82	256	4672	2839	3920	4569	4770
ham84	256	440	226	506	830	1015
john824	28	36	26	36	28	36
john844	70	182	132	225	220	225
john1624	120	784	457	700	514	675
john3224	496	14400	6294	9129	8722	9108
MANN a9	45	289	272	342	342	342
MANN a27	378	28728	15946	28875	30800	30800

Table 1: Solutions for DIMACS data: number of edges in the bicliques.

⁷[ftp://dimacs.rutgers.edu/pub/challenge/graph/benchmarks/clique](http://dimacs.rutgers.edu/pub/challenge/graph/benchmarks/clique).

⁸For the M.-S. algorithm: the alternative choices of parameters $\alpha = 1.005$ and $\alpha = 1.05$ for the DIMACS graphs and $\alpha = 1.005$ and $\alpha = 1.1$ for the random graphs were tested and all gave worse results. Small changes to the parameters of the Mult. algorithm led to similar results, so that it seems less sensitive to the choice of its parameters than the M.-S. algorithm.

density	Greedy	M.-S. ($\alpha = 1.01$)		M.-S. ($\alpha = 1.05$)		Mult. ($d = 1, \alpha = 1.1$)	
		mean	best	mean	best	mean	best
0.1	8.9	10.4	18.9	14.4	19.2	14.4	19.2
0.2	15.8	14.5	29.8	23.7	31.5	23.9	31.5
0.3	24.2	18.9	38.7	33.8	43.4	34.1	43.3
0.4	37.7	24.0	51.2	46.9	61.3	47.0	61.0
0.5	58.8	31.7	69.0	65.8	86.7	67.6	87.0
0.6	92.5	43.4	93.7	90.5	127.9	101.7	127.8
0.7	158.4	63.6	133.4	117.9	190.0	172.2	202.4
0.8	311.0	106.0	207.1	154.4	261.5	328.0	342.3
0.9	806.8	241.9	431.3	88.1	235.4	828.1	828.1

Table 2: Solutions for random graphs: average number of edges in the bicliques.

Remark 4. *Algorithm 1 enjoys some flexibility:*

- *It is applicable to non-binary matrices i.e. weighted graphs.*
- *It is possible to favor one side of the biclique. In fact, the multiplicative updates for NF can be adapted using the same developments as in Section 3.3 to cost functions with regularization terms, e.g.*

$$\min_{v,w \geq 0} \|M - vw\|_F^2 + \alpha \|v\|_2^2 + \beta \|w\|_2^2.$$

- *If d is kept sufficiently small, for example replacing $d = \alpha d$ by $d = \min(\alpha d, d_m)$ for some $d_m > 0$, there is no guarantee that the algorithm will converge to a biclique. However, the negative entries in M_d will enforce the corresponding entries of the solutions of (NF-1d) to be small (recall that Theorem 7 states that, for d sufficiently large, they will be equal to zero). Therefore, by rounding these solutions, instead of a biclique, one gets a dense submatrix of M_b i.e. a bicluster. Algorithm 1 can then be used as a biclustering algorithm. The density of the corresponding submatrix will depend on the choice of d_m . Table 3 gives an example of such behavior.*

d_m	0.01	0.05	0.1	0.5	1	1.5
size	5412	4428	2952	1073	595	539
density	29%	31%	35%	42%	51%	52%

Table 3: Biclusters for the 'classic' text mining dataset (7094 texts and 41681 words with more than 99.9% of entries equal to zero) with parameters $d_0 = 10^{-5}$, $\alpha = 1.025$, $\text{maxiter} = 500$.

6 Conclusion

We have introduced Nonnegative Factorization (NF), a new variant of Nonnegative Matrix Factorization (NMF), and proved its NP-hardness for any fixed rank by reduction to the maximum edge biclique problem. The multiplicative updates for NMF can be generalized to NF and provide a new interpretation of the algorithm of Lee and Seung, which explains why it does not perform well in practice. We also developed an heuristic algorithm for the biclique problem whose iterations require $O(|E|)$ operations, based on theoretical results about stationary points of a specific rank-one nonnegative factorization problem (NF-1d) and the use of multiplicative updates.

To conclude, we point out that none of the algorithms presented in this paper is guaranteed to converge to a globally optimal solution (and, to the best of our knowledge, such an algorithm has not been proposed yet) ; this is in all likelihood due to the NP-hardness of the NMF and NF problems. Indeed, only convergence to a stationary point has been proved for the algorithms of Sections 2 and 3, a property which, while desirable, provides no guarantee about the quality of the solution obtained (for example, nothing prevents these methods from converging to a stationary but rank-deficient solution, which in most cases could be further improved). Finally, no convergence proof for the biclique finding algorithm introduced in Section 5 is provided (convergence results from the preceding sections no longer hold because of the dynamic updates of parameter d) ; however, this heuristic seems to give very satisfactory results in practice.

References

- [1] A. BERMAN AND R. PLEMMONS, *Rank factorization of nonnegative matrices*, SIAM Review, 15 (3) (1973), p. 655.
- [2] M. BERRY, M. BROWNE, A. LANGVILLE, P. PAUCA, AND R. PLEMMONS, *Algorithms and Applications for Approximate Nonnegative Matrix Factorization*, Computational Statistics and Data Analysis, 52 (2007), pp. 155–173.
- [3] D. P. BERTSEKAS, *Nonlinear Programming: Second Edition*, Athena Scientific, Massachusetts, 1999.
- [4] M. BIGGS, A. GHODSI, AND S. VAVASIS, *Nonnegative Matrix Factorization via Rank-One Down-date*. Preprint, June 2007.
- [5] D. CHEN AND R. PLEMMONS, *Nonnegativity Constraints in Numerical Analysis*. Paper presented at the Symposium on the Birth of Numerical Analysis, Leuven Belgium. To appear in the Conference Proceedings, to be published by World Scientific Press, A. Bultheel and R. Cools, Eds., 2007.
- [6] C. CICHOCKI, R. ZDUNEK, AND S. AMARI, *Hierarchical ALS Algorithms for Nonnegative Matrix and 3D Tensor Factorization*, in ICA07, London, UK, September 9-12, Lecture Notes in Computer Science, Vol. 4666, Springer, pp. 169-176, 2007.
- [7] ———, *Nonnegative Matrix and Tensor Factorization*, IEEE Signal Processing Magazine, (2008), pp. 142–145.
- [8] J. COHEN AND U. ROTHBLUM, *Nonnegative ranks, Decompositions and Factorization of Non-negative Matrices*, Linear Algebra and its Applications, 190 (1993), pp. 149–168.
- [9] K. DEVARAJAN, *Nonnegative Matrix Factorization: An Analytical and Interpretive Tool in Computational Biology*, PLoS Computational Biology, 4(7), e1000029 (2008).
- [10] I. DHILLON, D. KIM, AND S. SRA, *Fast Newton-type Methods for the Least Squares Nonnegative Matrix Approximation Problem*, in Proceedings of SIAM Conf. on Data Mining, 2007.
- [11] I. DHILLON AND S. SRA, *Nonnegative Matrix Approximations: Algorithms and Applications*, tech. report, University of Texas (Austin), 2006. Dept. of Computer Sciences.
- [12] C. DING, Y. ZHANG, T. LI, AND S. HOLBROOK, *Biclustering Protein Complex Interactions with a Biclique Finding Algorithm*, in Sixth IEEE International Conference on Data Mining, 2006, pp. 178–187.

- [13] N. GILLIS, *Approximation et sous-approximation de matrices par factorisation positive: algorithmes, complexité et applications*, master's thesis, Université catholique de Louvain, 2007. In French.
- [14] N. GILLIS AND F. GLINEUR, *Nonnegative Matrix Factorization and Underapproximation*. Communication at 9th International Symposium on Iterative Methods in Scientific Computing, Lille, France, 2008.
- [15] G. GOLUB AND C. VAN LOAN, *Matrix Computation, 3rd Edition*, The Johns Hopkins University Press Baltimore, 1996.
- [16] N.-D. HO, *Nonnegative Matrix Factorization - Algorithms and Applications*, PhD thesis, Université catholique de Louvain, 2008.
- [17] N.-D. HO, P. VAN DOOREN, AND V. BLONDEL, *Weighted Nonnegative Matrix Factorization and Face Feature Extraction*. Submitted to Image and Vision Computing, 2007.
- [18] —, *Descent Type Algorithms for NMF*. arXiv:0801.3199v2, 2008.
- [19] P. HOYER, *Nonnegative Matrix Factorization with Sparseness Constraints*, J. Machine Learning Research, 5 (2004), pp. 1457–1469.
- [20] H. KIM AND H. PARK, *Non-negative Matrix Factorization Based on Alternating Non-negativity Constrained Least Squares and Active Set Method*, SIAM J. Matrix Anal. Appl., 30(2) (2008), pp. 713–730.
- [21] D. LEE AND H. SEUNG, *Learning the Parts of Objects by Nonnegative Matrix Factorization*, Nature, 401 (1999), pp. 788–791.
- [22] —, *Algorithms for Non-negative Matrix Factorization*, In Advances in Neural Information Processing, 13 (2001).
- [23] C.-J. LIN, *On the Convergence of Multiplicative Update Algorithms for Nonnegative Matrix Factorization*, in IEEE Transactions on Neural Networks, 2007.
- [24] —, *Projected Gradient Methods for Nonnegative Matrix Factorization*, Neural Computation, 19 (2007), pp. 2756–2779. MIT press.
- [25] P. PAATERO AND U. TAPPER, *Positive matrix factorization: a non-negative factor model with optimal utilization of error estimates of data values*, Environmetrics, 5 (1994), pp. 111–126.
- [26] R. PEETERS, *The maximum edge biclique problem is NP-complete*, Discrete Applied Mathematics, 131(3) (2003), pp. 651–654.
- [27] M. POWELL, *On Search Directions for Minimization Algorithms*, Mathematical Programming, 4 (1973), pp. 193–201.
- [28] G. STEWART AND J.-G. SUN, *Matrix Perturbation Theory*, Academic Press, San Diego, 1990.
- [29] L. THOMAS, *Rank factorization of nonnegative matrices*, SIAM Review, 16(3) (1974), pp. 393–394.
- [30] S. VAVASIS, *On the Complexity of Nonnegative Matrix Factorization*. arXiv:0708.4149v2, 2007.

Recent titles

CORE Discussion Papers

- 2008/26. Leonidas C. KOUTSOUGERAS and Nicholas ZIROS. Decentralization of the core through Nash equilibrium.
- 2008/27. Jean J. GABSZEWICZ, Didier LAUSSEL and Ornella TAROLA. To acquire, or to compete? An entry dilemma.
- 2008/28. Jean-Sébastien TRANCREZ, Philippe CHEVALIER and Pierre SEMAL. Probability masses fitting in the analysis of manufacturing flow lines.
- 2008/29. Marie-Louise LEROUX. Endogenous differential mortality, non monitored effort and optimal non linear taxation.
- 2008/30. Santanu S. DEY and Laurence A. WOLSEY. Two row mixed integer cuts via lifting.
- 2008/31. Helmuth CREMER, Philippe DE DONDER, Dario MALDONADO and Pierre PESTIEAU. Taxing sin goods and subsidizing health care.
- 2008/32. Jean J. GABSZEWICZ, Didier LAUSSEL and Nathalie SONNAC. The TV news scheduling game when the newscaster's face matters.
- 2008/33. Didier LAUSSEL and Joana RESENDE. Does the absence of competition *in* the market foster competition *for* the market? A dynamic approach to aftermarkets.
- 2008/34. Vincent D. BLONDEL and Yurii NESTEROV. Polynomial-time computation of the joint spectral radius for some sets of nonnegative matrices.
- 2008/35. David DE LA CROIX and Clara DELAVALLADE. Democracy, rule of law, corruption incentives and growth.
- 2008/36. Jean J. GABSZEWICZ and Joana RESENDE. Uncertain quality, product variety and price competition.
- 2008/37. Gregor ZOETTL. On investment decisions in liberalized electricity markets: the impact of price caps at the spot market.
- 2008/38. Helmuth CREMER, Philippe DE DONDER, Dario MALDONADO and Pierre PESTIEAU. Habit formation and labor supply.
- 2008/39. Marie-Louise LEROUX and Grégory PONTIERE. Optimal tax policy and expected longevity: a mean and variance approach.
- 2008/40. Kristian BEHRENS and Pierre M. PICARD. Transportation, freight rates, and economic geography.
- 2008/41. Gregor ZOETTL. Investment decisions in liberalized electricity markets: A framework of peak load pricing with strategic firms.
- 2008/42. Raouf BOUCEKKINE, Rodolphe DESBORDES and Hélène LATZER. How do epidemics induce behavioral changes?
- 2008/43. David DE LA CROIX and Marie VANDER DONCKT. Would empowering women initiate the demographic transition in least-developed countries?
- 2008/44. Geoffrey CARUSO, Dominique PEETERS, Jean CAVAILHES and Mark ROUNSEVELL. Space-time patterns of urban sprawl, a 1D cellular automata and microeconomic approach.
- 2008/45. Taoufik BOUEZMARNI, Jeroen V.K. ROMBOUTS and Abderrahim TAAMOUTI. Asymptotic properties of the Bernstein density copula for dependent data.
- 2008/46. Joe THARAKAN and Jean-Philippe TROPEANO. On the impact of labor market matching on regional disparities.
- 2008/47. Shin-Huei WANG and Cheng HSIAO. An easy test for two stationary long processes being uncorrelated via AR approximations.
- 2008/48. David DE LA CROIX. Adult longevity and economic take-off: from Malthus to Ben-Porath.
- 2008/49. David DE LA CROIX and Gregory PONTIERE. On the Golden Rule of capital accumulation under endogenous longevity.
- 2008/50. Jean J. GABSZEWICZ and Skerdilajda ZANAJ. Successive oligopolies and decreasing returns.
- 2008/51. Marie-Louise LEROUX, Pierre PESTIEAU and Grégory PONTIERE. Optimal linear taxation under endogenous longevity.
- 2008/52. Yuri YATSENKO, Raouf BOUCEKKINE and Natali HRITONENKO. Estimating the dynamics of R&D-based growth models.

Recent titles

CORE Discussion Papers - continued

- 2008/53. Roland Iwan LUTTENS and Marie-Anne VALFORT. Voting for redistribution under desert-sensitive altruism.
- 2008/54. Sergei PEKARSKI. Budget deficits and inflation feedback.
- 2008/55. Raouf BOUCEKKINE, Jacek B. KRAWCZYK and Thomas VALLEE. Towards an understanding of tradeoffs between regional wealth, tightness of a common environmental constraint and the sharing rules.
- 2008/56. Santanu S. DEY. A note on the split rank of intersection cuts.
- 2008/57. Yu. NESTEROV. Primal-dual interior-point methods with asymmetric barriers.
- 2008/58. Marie-Louise LEROUX, Pierre PESTIEAU and Gregory PONTIERE. Should we subsidize longevity?
- 2008/59. J. Roderick McCORIE. The role of Skorokhod space in the development of the econometric analysis of time series.
- 2008/60. Yu. NESTEROV. Barrier subgradient method.
- 2008/61. Thierry BRECHET, Johan EYCKMANS, François GERARD, Philippe MARBAIX, Henry TULKENS and Jean-Pascal VAN YPERSELE. The impact of the unilateral EU commitment on the stability of international climate agreements.
- 2008/62. Giorgia OGGIONI and Yves SMEERS. Average power contracts can mitigate carbon leakage.
- 2008/63. Jean-Sébastien TANCREZ, Philippe CHEVALIER and Pierre SEMAL. A tight bound on the throughput of queueing networks with blocking.
- 2008/64. Nicolas GILLIS and François GLINEUR. Nonnegative factorization and the maximum edge biclique problem.

Books

- Y. POCHET and L. WOLSEY (eds.) (2006), *Production planning by mixed integer programming*. New York, Springer-Verlag.
- P. PESTIEAU (ed.) (2006), *The welfare state in the European Union: economic and social perspectives*. Oxford, Oxford University Press.
- H. TULKENS (ed.) (2006), *Public goods, environmental externalities and fiscal competition*. New York, Springer-Verlag.
- V. GINSBURGH and D. THROSBY (eds.) (2006), *Handbook of the economics of art and culture*. Amsterdam, Elsevier.
- J. GABSZEWICZ (ed.) (2006), *La différenciation des produits*. Paris, La découverte.
- L. BAUWENS, W. POHLMEIER and D. VEREDAS (eds.) (2008), *High frequency financial econometrics: recent developments*. Heidelberg, Physica-Verlag.
- P. VAN HENTENRYCKE and L. WOLSEY (eds.) (2007), *Integration of AI and OR techniques in constraint programming for combinatorial optimization problems*. Berlin, Springer.

CORE Lecture Series

- C. GOURIÉROUX and A. MONFORT (1995), *Simulation Based Econometric Methods*.
- A. RUBINSTEIN (1996), *Lectures on Modeling Bounded Rationality*.
- J. RENEGAR (1999), *A Mathematical View of Interior-Point Methods in Convex Optimization*.
- B.D. BERNHEIM and M.D. WHINSTON (1999), *Anticompetitive Exclusion and Foreclosure Through Vertical Agreements*.
- D. BIENSTOCK (2001), *Potential function methods for approximately solving linear programming problems: theory and practice*.
- R. AMIR (2002), *Supermodularity and complementarity in economics*.
- R. WEISMANTEL (2006), *Lectures on mixed nonlinear programming*.