



HAL
open science

A new Sinkhorn algorithm with Deletion and Insertion operations

Luc Brun, Benoit Gaüzère, Sébastien Bogleux, Florian Yger

► **To cite this version:**

Luc Brun, Benoit Gaüzère, Sébastien Bogleux, Florian Yger. A new Sinkhorn algorithm with Deletion and Insertion operations. [Research Report] GREYC, UMR 6072. 2021. hal-03454896v2

HAL Id: hal-03454896

<https://hal.science/hal-03454896v2>

Submitted on 18 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

A new Sinkhorn algorithm with Deletion and Insertion operations

Luc Brun¹, Benoit Gaüzère², Sébastien Bougleux¹, Florian Yger³

¹Normandie Univ, CNRS, ENSICAEN, UNICAEN, GREYC, France

²Normandie Univ, UNIROUEN, UNIHAVRE, INSA Rouen, LITIS, Rouen, France

³PSL Université Paris-Dauphine, LAMSADE, Paris, France

October 2021

1 Introduction

This report is devoted to the continuous estimation of an ϵ -assignment (Definition 1). Roughly speaking, an ϵ -assignment between two sets V_1 and V_2 may be understood as a bijective mapping between a sub part of V_1 and a sub part of V_2 . The remaining elements of V_1 (not included in this mapping) are mapped onto an ϵ pseudo element of V_2 . We say that such elements are deleted. Conversely, the remaining elements of V_2 correspond to the image of the ϵ pseudo element of V_1 (Figure 1). We say that these elements are inserted.

Let us note that if V_1 and V_2 have the same size, the bijective mapping induced by an ϵ -assignment may involve all elements of V_1 , each element being mapped onto a single element of V_2 . In this sense, an ϵ -assignment is more general than a bijective mapping. Moreover, the main advantage of an ϵ -assignment is that it provides us the freedom to not map any element which is then assigned to the ϵ element of V_2 or belong to the image of the ϵ element of V_1 . This last property allows us to reject some mappings if for example, these mappings are associated to a large cost.

An ϵ -assignment function may be associated to an ϵ -assignment matrix (Figure 1(b)) just like any bijective mapping is associated to a permutation matrix. Given two sets, V_1 and V_2 of respective sizes n and m , an ϵ -assignment matrix is encoded by a $(n+1) \times (m+1)$ matrix, where $n+1$ and $m+1$ play respectively the roles of the ϵ element of V_1 and the one of V_2 . The last column of index $m+1$ of such a matrix encodes the deletions while the last line encodes the insertions. By construction, there is a single 1 in each of the first n rows and m columns, the remaining elements being set to 0.




Figure 1: (a) An example of ϵ -assignment function. 1 is mapped onto b, 2 onto a, 3 is deleted. (b) its associated ϵ assignment matrix

Given V_1 and V_2 , one can define a $(n+1) \times (m+1)$ cost matrix encoding the cost of the mapping of any element of V_1 onto an element of V_2 as well as the cost of deleting each element of V_1 and inserting each element of V_2 . Finding an ϵ -assignment minimizing the sum of mappings, deletions and insertions costs is a direct extension of the Linear Sum Assignment Problem (LSAP) called the Linear Sum Assignment Problem with Edition [1] (LSAPE). Given an ϵ -assignment matrix X and a cost matrix C , this cost may be formulated as:

$$\min_X \sum_{i=1}^{n+1} \sum_{j=1}^{m+1} c_{i,j} x_{i,j}$$

where X is taken over all ϵ -assignment matrices.

We define in previous works [2, 1], an adaptation of the Hungarian algorithm which allows to find an optimal solution to the above problem in $\mathcal{O}(\min(n, m)^2 \max(n, m))$. However, while providing an optimal solution, this algorithm does not readily allow the computation of the gradient of the associated operation. This last drawback, does not allow to easily insert such an algorithm into a deep learning pipeline. On the other hand, the Sinkhorn algorithm [6], is based on a continuous relaxation of the problem where permutation matrices are replaced by bi-stochastic matrices with an entropic regularization. This algorithm is the workhorse of computational optimal transport [4] and is based on iterative matrix multiplications hereby allowing the backpropagation of the gradient [3]. The aim of this technical report is to transpose the results of the Sinkhorn algorithm to ϵ assignment matrices. Just like the Sinkhorn algorithm which does not provide a permutation matrix but rather a bi-stochastic matrix, our algorithm will provide an ϵ bi-stochastic matrix (Definition 3). This last point may be of advantage within the Neural Network framework where the hard decisions corresponding to ϵ -assignment matrices may not allow a proper propagation of the gradient.

More formally, given a similarity matrix S (which may be easily deduced from a cost matrix), we aim at finding two diagonal matrices D_1 and D_2 such that $D_1 S D_2$ is a ϵ bi-stochastic matrix. Section 2 provides the main definitions and notations used in the remaining part of this report. The existence and uniqueness of a solution is demonstrated in Section 3 while Section 4 provides a constructive algorithm which convergence is demonstrated. Let us note that while Section 3 is a simple adaption of the original proof [6], Section 4 is significantly different from [6] since the arguments used for bi-stochastic matrices in the original proof do not hold for ϵ bi-stochastic matrices.

2 Definitions and notations

Definition 1 (ϵ -assignment).

Let n and m be two strictly positive integers. An ϵ -assignment is a mapping $\varphi : \{1, \dots, n+1\} \rightarrow \mathcal{P}(\{1, \dots, m+1\})$ satisfying the following constraints:

$$\begin{cases} \forall i \in \{1, \dots, n\}, & |\varphi(i)| = 1 \\ \forall j \in \{1, \dots, m\}, & |\varphi^{-1}(j)| = 1 \\ m+1 \in \varphi(n+1) \end{cases}$$

where $\mathcal{P}(\{1, \dots, m+1\})$ is the power set of $\{1, \dots, m+1\}$.

Each element of $i \in \{1, \dots, n\}$ is thus mapped onto a set composed of a single element of $\{1, \dots, m+1\}$ ($|\varphi(i)| = 1$) and in the same way the set of antecedents of each $j \in \{1, \dots, m\}$ is reduced to one element ($|\varphi^{-1}(j)| = 1$). Hence the only element of $\{1, \dots, n+1\}$ which can be mapped onto a set composed of several elements is $n+1$. In the same way, $m+1$ is the only element which may have several antecedents. The constraint $m+1 \in \varphi(n+1)$ ensures that $n+1$ is mapped to at least one element and that $m+1$ has at least an antecedent.

In the example of Figure 1 we have $n = 3$ and $m = 2$. Elements 1, 2, 3 are respectively mapped onto $\{b\}, \{a\}, \{3\}$. Where the last mapping corresponds to a deletion of 3 (which is mapped onto $m+1 = 3$). Consequently $m+1 = 3$ has two antecedents 3 and $4 = n+1$.

Definition 2 (ϵ -row/column stochastic matrix).

A non negative $(n+1) \times (m+1)$ matrix X is called an ϵ -row stochastic matrix iff:

$$\forall i \in \{1, \dots, n\} \quad \sum_{j=1}^{m+1} X_{i,j} = 1$$

X is called an ϵ -column stochastic matrix iff:

$$\forall j \in \{1, \dots, m\} \quad \sum_{i=1}^{n+1} X_{i,j} = 1$$

Definition 3 (ϵ -bi-stochastic matrix).

A non negative $(n+1) \times (m+1)$ matrix X is called an ϵ -bi-stochastic matrix iff:

$$\begin{cases} \sum_{j=1}^{m+1} X_{i,j} = 1 & \forall i \in \{1, \dots, n\} \\ \sum_{i=1}^{n+1} X_{i,j} = 1 & \forall j \in \{1, \dots, m\} \\ x_{n+1, m+1} = 1 \end{cases}$$

If $X \in \{0, 1\}^{(n+1) \times (m+1)}$, X is called an ϵ -assignment matrix and there is a one-to-one mapping between ϵ -assignments and ϵ -assignment matrices.

Let us note that any ϵ -bi-stochastic matrix is a bi-stochastic matrix on which the bi-stochastic constraints are relaxed on the last line and last column. So any squared bi-stochastic matrix is also an ϵ -bi-stochastic matrix (the reverse being obviously false).

Definition 4. ϵ -diagonal

If A is a $(n+1) \times (m+1)$ matrix and φ an ϵ -assignment then the set $A_{1, \varphi(1)}, \dots, A_{n+1, \varphi(n+1)}$ is called an ϵ -**diagonal** of A corresponding to φ . if A is squared and φ is the identity, the diagonal is called the main diagonal.

Note that $a_{1, \varphi(1)}, \dots, a_{n, \varphi(n)}$ is a sequence (as $\varphi(i)$ is unique for $i \in \{1, \dots, n\}$ while $a_{n+1, \varphi(n+1)}$ is a set). The above definition is a straightforward of the usual notion of diagonal where φ is required to be a permutation. In the following we will only consider ϵ -diagonals of matrices which will be simply called diagonal.

Definition 5. total support

If A is a nonnegative matrix, A is said to have total support if $A \neq 0$ and if every positive element of A lies on a positive ϵ -diagonal. A nonnegative matrix that contains a positive diagonal is said to have a support.

If μ and ν define set of indices respectively contained in $\{1, \dots, n+1\}$ and $\{1, \dots, m+1\}$ then :

- $A[\mu, \nu]$ denotes the sub matrix of A restricted to indices μ and ν ,
- $A(\mu, \nu)$ denotes the sub matrix of A restricted to indices not contained in μ , i.e. $\{1, \dots, n+1\} - \mu$, and to the indices contained in ν ,
- $A[\mu, \nu)$ denotes the sub matrix of A restricted to indices contained in μ and not contained in ν , i.e. $\{1, \dots, m+1\} - \nu$.
- $A(\mu, \nu)$ denotes the sub matrix of A restricted to the indices not contained in μ and ν .

Definition 6. *Secable rectangular matrix*

A rectangular $n \times m$ non negative matrix A is said to be secable if one can find :

- a partition of $\{1, \dots, n\}$ into two sets X and Y and
- a partition of $\{1, \dots, m\}$ into two sets Z and T

such that:

$$A[X, T] = 0 \text{ and } A[Y, Z] = 0$$

$$\begin{array}{c|cc} & Z & T \\ \hline X & A[X, Z] & 0 \\ Y & 0 & A[Y, T] \end{array}$$

Let us note that this notion of secable matrix is quite close from the one of block diagonal matrix. However, $A[X, T]$ is not required to be squared.

3 Existence and uniqueness

Theorem 3.1. *Let A be a nonnegative $(n+1) \times (m+1)$ matrix such that $A[\{1, \dots, n\}, \{1, \dots, m\}]$ does not contain any line or column filled with 0. A necessary and sufficient condition that there exists an ϵ bi-stochastic matrix B of the form $D_1 A D_2$ where D_1 and D_2 are diagonal matrices with positive main diagonals and a last entry equal to 1 is that A has total support. If B exists then it is unique. Also D_1 and D_2 are unique if and only if A is non secable.*

Proof. Let us suppose that $B = D_1 A D_2$ and $B' = D'_1 A D'_2$ are ϵ -bi-stochastic matrices where $D_1 = \text{diag}(x_1, \dots, x_n, 1)$, $D_2 = \text{diag}(y_1, \dots, y_m, 1)$, $D'_1 = \text{diag}(x'_1, \dots, x'_n, 1)$ and $D'_2 = \text{diag}(y'_1, \dots, y'_m, 1)$. If $p_i = \frac{x'_i}{x_i}$ and $q_j = \frac{y'_j}{y_j}$:

$$\sum_{i=1}^{n+1} x_i a_{ij} y_j = 1, \forall j = 1, \dots, m; \quad (1)$$

$$\sum_{j=1}^{m+1} x_i a_{ij} y_j = 1, \forall i = 1, \dots, n \quad (2)$$

$$\sum_{i=1}^{n+1} p_i x_i a_{ij} q_j y_j = 1, \forall j = 1, \dots, m; \quad (3)$$

$$\sum_{j=1}^{m+1} p_i x_i a_{ij} q_j y_j = 1, \forall i = 1, \dots, n \quad (4)$$

Let $E_j = \{i | a_{ij} > 0\}$, $F_i = \{j | a_{ij} > 0\}$ and put

$$\alpha = \{i \in \{1, \dots, n\} | p_i = \min_i p_i = \underline{p}\}, \quad \beta = \{j \in \{1, \dots, m\} | q_j = \max_j q_j = \bar{q}\}$$

Let us note that since $x_{n+1} = y_{m+1} = x'_{n+1} = y'_{m+1} = 1$ we have $p_{n+1} = q_{m+1} = 1$, $\underline{p} \leq 1$ and $\bar{q} \geq 1$. Moreover, we have by hypothesis $E_j \cap \{1, \dots, n\} \neq \emptyset$ and $F_i \cap \{1, \dots, m\} \neq \emptyset$ for all $(i, j) \in \{1, \dots, n\} \times \{1, \dots, m\}$.

Let us fist show that $\alpha = \emptyset \iff \beta = \emptyset$.

If $\alpha = \emptyset$, then For all $i \in \{1, \dots, n\}$ $p_i > 1$. Then using equation (3), we have for any $j \in \{1, \dots, m\}$:

$$1 = \sum_{i=1}^{n+1} p_i x_i a_{ij} q_j y_j > q_j \sum_{i=1}^{n+1} x_i a_{ij} y_j = q_j$$

Hence $q_j < 1$ for all $j \in \{1, \dots, m\}$ and thus $\beta = \emptyset$.

Conversely, if $\beta = \emptyset$, we have $q_j < 1$ for all $j \in \{1, \dots, m\}$. Using equation (4) for $i \in \{1, \dots, n\}$:

$$1 = \sum_{j=1}^{m+1} p_i x_i a_{ij} q_j y_j < p_i \sum_{j=1}^{m+1} x_i a_{ij} y_j = p_i$$

Hence, $p_i > 1$ for all $i \in \{1, \dots, n\}$ and $\alpha = \emptyset$.

In this case we consider the alternative definitions for α and β :

$$\begin{cases} \alpha &= \{i \in \{1, \dots, n\} | p_i = \max_i p_i = \bar{p}\}, \\ \beta &= \{j \in \{1, \dots, m\} | q_j = \min_j q_j = \underline{q}\} \end{cases}$$

Since $p_i > 1$ for all $i \in \{1, \dots, n\}$ and $q_j < 1$ for all $j \in \{1, \dots, m\}$, α and β are non empty.

Using initial definitions for α and β , let us thus consider $i_0 \in \alpha$ and $j_0 \in \beta$. Then using (3):

$$q_{j_0} = \frac{1}{\sum_{i=1}^{n+1} p_i x_i a_{ij_0} y_{j_0}} \leq \frac{1}{p_{i_0} \sum_{i=1}^{n+1} x_i a_{ij_0} y_{j_0}} = p_{i_0}^{-1}$$

where the last equality comes from (1). Similarly, using (4):

$$p_{i_0} = \frac{1}{\sum_{j=1}^{m+1} x_i a_{ij} q_j y_j} \geq \frac{1}{q_{j_0} \sum_{j=1}^{m+1} x_i a_{ij} y_j} = q_{j_0}^{-1}$$

where the last equality comes from (2). Whence $q_{j_0} = p_{i_0}^{-1} = \underline{p}^{-1}$. But in this case, we have using (3):

$$\sum_{i=1}^{n+1} p_i x_i a_{ij_0} q_{j_0} y_{j_0} = \sum_{i \in E_{j_0}} p_i x_i a_{ij_0} q_{j_0} y_{j_0} = \sum_{i \in E_{j_0}} \frac{p_i}{\underline{p}} x_i a_{ij_0} y_{j_0} = 1$$

This last equality is compatible with (1) only if $p_i = \underline{p}$ for all $i \in E_{j_0}$. Dropping sub indices, we have for all $j \in \beta$ and all $i \in E_j$ $p_i = \underline{p}$. Thus

$$\bigcup_{j \in \beta} E_j \subseteq \alpha \cup \{n+1\}.$$

Hence if $j \in \beta$ and $i \notin \alpha \cup \{n+1\}$, $i \notin \bigcup_{j \in \beta} E_j$. So $a_{ij} = 0$. More concisely, we have: $A[\alpha \cup \{n+1\}, \beta] = 0$.

In the same way, $p_{i_0} = q_{j_0}^{-1} = \bar{q}^{-1}$ implies using (4):

$$\sum_{j=1}^{m+1} p_{i_0} x_{i_0} a_{i_0 j} q_j y_j = \sum_{j \in F_{i_0}} \frac{q_j}{\bar{q}} x_{i_0} a_{i_0 j} y_j = 1$$

which is compatible with (2) only if $q_j = \bar{q}$ for all $j \in F_{i_0}$. Thus for all $i \in \alpha$ and for all $j \in F_i$ we have $q_j = \bar{q}$. Thus

$$\bigcup_{i \in \alpha} F_i \subseteq \beta \cup \{m+1\} \text{ and } A[\alpha, \beta \cup \{m+1\}] = 0$$

On $\alpha \times \beta$ we have $p_i q_j = \underline{p} \bar{q} = 1$. Thus:

$$a_{ij} = \frac{b_{ij}}{x_i y_j} = \frac{b'_{ij}}{x'_i y'_j} \Rightarrow p_i q_j b_{ij} = b'_{ij} \Rightarrow b_{ij} = b'_{ij}$$

Hence $B[\alpha, \beta] = B'[\alpha, \beta]$

Moreover, for any $j \in \beta$ we have:

$$\sum_{i=1}^{n+1} b_{i,j} = \sum_{i=1}^{n+1} x_i a_{i,j} y_j = \sum_{i \in E_j \subset \alpha \cup \{n+1\}} x_i a_{i,j} y_j = \sum_{i \in \alpha \cup \{n+1\}} x_i a_{i,j} y_j = \sum_{i \in \alpha \cup \{n+1\}} b_{i,j} = 1$$

In the same way, we have for any $j \in \beta$, $\sum_{i \in \alpha \cup \{n+1\}} b'_{i,j} = 1$. But in this case using $B[\alpha, \beta] = B'[\alpha, \beta]$, we have for $j \in \beta$:

$$1 = \sum_{i \in \alpha \cup \{n+1\}} b'_{i,j} = \sum_{i \in \alpha} b'_{i,j} + b'_{n+1,j} = \sum_{i \in \alpha} b_{i,j} + b'_{n+1,j}$$

Thus $1 - b_{n+1,j} + b'_{n+1,j} = 1$ which induces $b_{n+1,j} = b'_{n+1,j}$ which imposes $\bar{q} = 1$. Indeed, since $j \in \beta$, we have $q_j = \bar{q}$ and $q_j b_{n+1,j} = \bar{q} b_{n+1,j} = b'_{n+1,j}$.

In the same way for $i \in \alpha$:

$$\sum_{j=1}^{m+1} b_{i,j} = \sum_{j=1}^{m+1} x_i a_{i,j} y_j = \sum_{j \in F_i \subset \beta \cup \{m+1\}} x_i a_{i,j} y_j = \sum_{j \in \beta \cup \{m+1\}} x_i a_{i,j} y_j = \sum_{j \in \beta \cup \{m+1\}} b_{i,j} = 1$$

and we have in the same way for $i \in \alpha$, $\sum_{j \in \beta \cup \{m+1\}} b'_{i,j} = 1$. We then obtain for $i \in \alpha$:

$$\sum_{j \in \beta \cup \{m+1\}} b'_{i,j} = \sum_{j \in \beta} b'_{i,j} + b'_{i,m+1} = \sum_{j \in \beta} b_{i,j} + b'_{i,m+1}$$

Using the previous equality $1 - b_{i,m+1} + b'_{i,m+1} = 1$, hence $b_{i,m+1} = b'_{i,m+1}$ and $\underline{p} = 1$ (since $\underline{p} b_{i,n+1} = b'_{i,n+1}$).

Thus $B[\alpha \cup \{n+1\}, \beta \cup \{m+1\}] = B'[\alpha \cup \{n+1\}, \beta \cup \{m+1\}]$ is an ϵ bi-stochastic matrix (where $n+1$ and $m+1$ plays the role of the last row and column respectively).

Let us briefly show that $\{1, \dots, n\} - \alpha$ and $\{1, \dots, m\} - \beta$ are simultaneously empty or non empty. Let us first suppose that $\alpha = \{1, \dots, n\}$ and let us consider $j \notin \beta \cup \{m+1\}$. Since $E_j \cap \{1, \dots, n\} \neq \emptyset$ by hypothesis, it exists $i \in \{1, \dots, n\} = \alpha$ such that $j \in F_i$. But since $i \in \alpha$, we have $F_i \subset \beta \cup \{m+1\}$ and thus a contradiction. In the same way, if $\beta = \{1, \dots, m\}$, let us consider $i \notin \alpha \cup \{n+1\}$. Since $F_i \cap \{1, \dots, m\} \neq \emptyset$, it exists $j \in \{1, \dots, m\} = \beta$ such that $i \in E_j \subset \alpha \cup \{n+1\}$. Again a contradiction.

If A is non secable the configuration where both $\{1, \dots, n\} - \alpha$ and $\{1, \dots, m\} - \beta$ are non empty correspond to a partition of $\{1, \dots, n+1\}$ into $\alpha \cup \{n+1\}$ and its complementary and a partition of $\{1, \dots, m+1\}$ into $\beta \cup \{m+1\}$ and its complementary with no connections between $\alpha \cup \{n+1\}$ and the complementary of $\beta \cup \{m+1\}$ nor any connection between $\beta \cup \{m+1\}$ and the complementary of $\alpha \cup \{n+1\}$ (Figure 2). Such a decomposition being refused, we have $\alpha = \{1, \dots, n\}$ and $\beta = \{1, \dots, m\}$ Hence $A[\alpha \cup \{n+1\}, \beta \cup \{m+1\}] = A$ and $D_1 A D_2 = D'_1 A D'_2$ and D_1 and D_2 are unique ($\underline{p} = \bar{q} = p_{n+1} = q_{m+1} = 1$).

If the non secable property of A does not hold and $A(\alpha \cup \{n+1\}, \beta)$ and $A(\alpha, \beta \cup \{m+1\})$ exist, $B(\alpha, \beta)$ and $B'(\alpha, \beta)$ exist, include the row $n+1$ and the column $m+1$, are ϵ bi-stochastic matrices and have a size lower than the one of A . Furthermore, $B(\alpha, \beta) = D''_1 A(\alpha, \beta) D''_2$ and $B'(\alpha, \beta) = D'''_1 A(\alpha, \beta) D'''_2$ where D''_i and D'''_j have like D_i and D'_j (from which they are derived) a positive main diagonal with a 1 at last position. The argument may be repeated on these submatrices until $D_1 A D_2 = D'_1 A D'_2$ is established. Given that $B(\alpha, \beta) = B'(\alpha, \beta)$, we already know that $B[\alpha \cup \{n+1\}, \beta \cup \{m+1\}] = B'[\alpha \cup \{n+1\}, \beta \cup \{m+1\}]$ and that A and hence B and B' are zeros elsewhere (Figure 2). Hence B is equal to B' . Note however, that since $\alpha \cup \{n+1\} \neq \{1, \dots, n+1\}$ (and the same for β) D_1 and D_2 are no longer unique. \square




Figure 2: Decomposition of matrix B .

4 A constructive algorithm

For any $i \in \{1, \dots, n\}$ and any $j \in \{1, \dots, m\}$ let us consider the series $(x_{i,p})_{p \in \mathbb{N}}$ and $(y_{j,p})_{p \in \mathbb{N}}$ defined as follows:

$$\begin{cases} \forall i \in \{1, \dots, n\} & x_{i,0} = \left(\sum_{j=1}^{m+1} a_{ij} \right)^{-1} & x_{i,p+1} = \chi_{i,p}^{-1} x_{i,p} \\ \forall j \in \{1, \dots, m\} & y_{j,0} = 1 & y_{j,p+1} = \gamma_{j,p}^{-1} y_{j,p} \end{cases}$$

Moreover we also define:

$$\begin{cases} \forall i \in \{1, \dots, n\} & \chi_{i,p} = \sum_{j=1}^{m+1} x_{i,p} a_{i,j} y_{j,p} = \sum_{j=1}^{m+1} \gamma_{j,p-1}^{-1} x_{i,p} a_{i,j} y_{j,p-1} \\ \forall j \in \{1, \dots, m\} & \gamma_{j,p} = \sum_{i=1}^{n+1} \chi_{i,p}^{-1} x_{i,p} a_{i,j} y_{j,p} \end{cases}$$

with for all p :

$$\begin{aligned} x_{n+1,p} &= y_{m+1,p} = 1 \\ \chi_{n+1,p} &= \gamma_{m+1,p} = 1 \end{aligned}$$

Let us denote by L_p the $(n+1) \times (m+1)$ matrix whose entries in $n \times (m+1)$ are equal to $(x_{i,p} a_{i,j} y_{j,p-1})$ and whose last row is filled with zeros but a 1 at position $(n+1, m+1)$.

In the same way let us denote by C_p the $(n+1) \times (m+1)$ matrix whose entries in $(n+1) \times m$ are equal to $x_{i,p} a_{i,j} y_{j,p}$ and whose last column is filled with zeros except a 1 at position $(n+1) \times (m+1)$.

If $\chi_p \in \mathbb{R}^{n+1}$ and $\gamma_p \in \mathbb{R}^{m+1}$ denote the vectors encoding respectively $(\chi_{i,p})_{i \in \{1, \dots, n+1\}}$ and $(\gamma_{j,p})_{j \in \{1, \dots, m+1\}}$ we have for $p \geq 1$:

$$\begin{cases} \chi_p &= L_p \gamma_{p-1}^{-1} \\ \gamma_p &= C_p^T \chi_p^{-1} \end{cases} \quad (5)$$

L_p is row stochastic. Indeed, for any $i \in \{1, \dots, n\}$ and $p \geq 1$:

$$\begin{aligned} \sum_{j=1}^{m+1} (L_p)_{i,j} &= \sum_{j=1}^{m+1} x_{i,p} a_{i,j} y_{j,p-1} \\ &= \sum_{j=1}^{m+1} \chi_{i,p-1}^{-1} x_{i,p-1} a_{i,j} y_{j,p-1} \\ &= \chi_{i,p-1}^{-1} \sum_{j=1}^{m+1} x_{i,p-1} a_{i,j} y_{j,p-1} \\ &= 1 \end{aligned}$$

and the last line of L_p contains a single entry equal to 1.

Moreover, C_p is column stochastic for $p \geq 1$. Indeed for each $j \in \{1, \dots, m\}$:

$$\begin{aligned} \sum_{i=1}^{n+1} (C_p)_{i,j} &= \sum_{i=1}^{n+1} x_{i,p} a_{i,j} y_{j,p} \\ &= \sum_{i=1}^{n+1} \chi_{i,p-1}^{-1} \gamma_{j,p-1}^{-1} x_{i,p-1} a_{i,j} y_{j,p-1} \\ &= \gamma_{j,p-1}^{-1} \sum_{i=1}^{n+1} \chi_{i,p-1}^{-1} x_{i,p-1} a_{i,j} y_{j,p-1} \\ &= 1 \end{aligned}$$

$$\begin{aligned}
L_p &= \begin{pmatrix} x_{1,p}a_{1,1}y_{1,p-1} & x_{1,p}a_{1,2}y_{2,p-1} & \dots & x_{1,p}a_{1,m}y_{m,p-1} & x_{1,p}a_{1,m+1} \\ x_{2,p}a_{2,1}y_{1,p-1} & x_{2,p}a_{2,2}y_{2,p-1} & \dots & x_{2,p}a_{2,m}y_{m,p-1} & x_{2,p}a_{2,m+1} \\ & & \vdots & & \\ x_{n,p}a_{n,1}y_{1,p-1} & x_{n,p}a_{n,2}y_{2,p-1} & \dots & x_{n,p}a_{n,m}y_{m,p-1} & x_{n,p}a_{n,m+1} \\ & 0 & & 0 & 1 \end{pmatrix} \\
C_p &= \begin{pmatrix} x_{1,p}a_{1,1}y_{1,p} & x_{1,p}a_{1,2}y_{2,p} & \dots & x_{1,p}a_{1,m}y_{m,p} & 0 \\ x_{2,p}a_{2,1}y_{1,p} & x_{2,p}a_{2,2}y_{2,p} & \dots & x_{2,p}a_{2,m}y_{m,p} & 0 \\ & & \vdots & & \\ x_{n,p}a_{n,1}y_{1,p} & x_{n,p}a_{n,2}y_{2,p} & \dots & x_{n,p}a_{n,m}y_{m,p} & 0 \\ a_{n+1,1}y_{1,p} & a_{n+1,2}y_{2,p} & \dots & a_{n+1,m}y_{m,p} & 1 \end{pmatrix} \\
C_p^T &= \begin{pmatrix} x_{1,p}a_{1,1}y_{1,p} & x_{2,p}a_{2,1}y_{1,p} & \dots & x_{n,p}a_{n,1}y_{1,p} & a_{n+1,1}y_{1,p} \\ x_{1,p}a_{1,2}y_{2,p} & x_{2,p}a_{2,2}y_{2,p} & \dots & x_{n,p}a_{n,2}y_{2,p} & a_{n+1,2}y_{2,p} \\ & & \vdots & & \\ x_{1,p}a_{1,m}y_{m,p} & x_{2,p}a_{2,m}y_{m,p} & \dots & x_{n,p}a_{n,m}y_{m,p} & a_{n+1,m}y_{m,p} \\ & 0 & & 0 & 1 \end{pmatrix}
\end{aligned}$$

Figure 3: L_p and C_p matrices.

Note that C_0 is not column stochastic. One noticeable effect of this negative property is that $\gamma_0 \neq 1$ while $\chi_0 = 1$.

Moreover the last column contains a single positive entry equal to 1. Hence C_p^T is row stochastic and equation 5 involves two row stochastic matrices.

Combining both equations of 5 we have:

$$\chi_p = L_p (C_{p-1}^T \chi_{p-1}^{-1})^{-1}$$

where the inverse notation applied to a vector denotes the element-wise inverse operation.

Since L_p is row stochastic, we have:

$$\chi_p - 1 = L_p \left[(C_{p-1}^T \chi_{p-1}^{-1})^{-1} - 1 \right]$$

Using $\frac{1}{x} - 1 = \frac{1-x}{x}$ we obtain:

$$\chi_p - 1 = L_p \left[(1 - C_{p-1}^T \chi_{p-1}^{-1}) \odot (C_{p-1}^T \chi_{p-1}^{-1})^{-1} \right]$$

where \odot is the element-wise product also known as Hadamard product. Since, for $p \geq 2$, C_{p-1}^T is row stochastic we have $C_{p-1}^T \mathbf{1} = 1$ and thus:

$$\begin{aligned}
\chi_p - 1 &= L_p \left[(C_{p-1}^T (1 - \chi_{p-1}^{-1})) \odot (C_{p-1}^T \chi_{p-1}^{-1})^{-1} \right] \\
&= L_p \text{diag} (C_{p-1}^T \chi_{p-1}^{-1})^{-1} C_{p-1}^T (1 - \chi_{p-1}^{-1})
\end{aligned}$$

Using $1 - \frac{1}{x} = \frac{x-1}{x}$ we obtain:

$$\begin{aligned}
\chi_p - 1 &= L_p \text{diag} (C_{p-1}^T \chi_{p-1}^{-1})^{-1} C_{p-1}^T (\chi_{p-1} - 1) \odot \chi_{p-1}^{-1} \\
&= L_p \text{diag} (C_{p-1}^T \chi_{p-1}^{-1})^{-1} C_{p-1}^T \text{diag}(\chi_{p-1})^{-1} (\chi_{p-1} - 1) \\
&= L_p \text{diag} (\gamma_{p-1})^{-1} C_{p-1}^T \text{diag}(\chi_{p-1})^{-1} (\chi_{p-1} - 1)
\end{aligned}$$

Moreover the left and right multiplications of C_{p-1}^T by diagonal matrices is equivalent to a multiplications of its lines by γ_{p-1}^{-1} and its columns by χ_{p-1}^{-1} . More precisely we have:

$$\begin{aligned} & \text{diag}(\gamma_{p-1})^{-1} C_{p-1}^T \text{diag}(\chi_{p-1})^{-1} = \\ & \begin{pmatrix} \chi_{1,p-1}^{-1} \gamma_{1,p-1}^{-1} x_{1,p-1} a_{1,1} y_{1,p-1} & \cdots & \chi_{n,p-1}^{-1} \gamma_{1,p-1}^{-1} x_{n,p-1} a_{n,1} y_{1,p-1} & \gamma_{1,p-1}^{-1} a_{n+1,1} y_{1,p} \\ \chi_{1,p-1}^{-1} \gamma_{2,p-1}^{-1} x_{1,p-1} a_{1,2} y_{2,p-1} & \cdots & \chi_{n,p-1}^{-1} \gamma_{2,p-1}^{-1} x_{n,p-1} a_{n,2} y_{2,p-1} & \gamma_{2,p-1}^{-1} a_{n+1,2} y_{2,p} \\ & & \vdots & \\ \chi_{1,p-1}^{-1} \gamma_{m,p-1}^{-1} x_{1,p-1} a_{1,m} y_{m,p-1} & \cdots & \chi_{n,p-1}^{-1} \gamma_{m,p-1}^{-1} x_{n,p-1} a_{n,m} y_{m,p-1} & \gamma_{m,p-1}^{-1} a_{n+1,m} y_{m,p} \\ & & 0 & 1 \end{pmatrix} \\ & = \begin{pmatrix} x_{1,p} a_{1,1} y_{1,p} & x_{2,p} a_{2,1} y_{1,p} & \cdots & x_{n,p} a_{n,1} y_{1,p} & a_{n+1,1} y_{1,p} \\ x_{1,p} a_{1,2} y_{2,p} & x_{2,p} a_{2,2} y_{2,p} & \cdots & x_{n,p} a_{n,2} y_{2,p} & a_{n+1,2} y_{2,p} \\ & & \vdots & & \\ x_{1,p} a_{1,m} y_{m,p} & x_{2,p} a_{2,m} y_{m,p} & \cdots & x_{n,p} a_{n,m} y_{m,p} & a_{n+1,m} y_{m,p} \\ & & 0 & 0 & 1 \end{pmatrix} = C_p^T \end{aligned}$$

Hence we have:

$$\forall p \geq 2 \quad \chi_p - 1 = L_p C_p^T (\chi_{p-1} - 1) \quad (6)$$

As L_p and C_p^T are row stochastic matrices, so is $L_p C_p^T$. Moreover, $L_p C_p^T$ is a square $(n+1) \times (n+1)$ matrix. Let us note that due to this row stochastic property the last equality is equivalent to:

$$\forall p \geq 2 \quad \chi_p = L_p C_p^T \chi_{p-1}$$

Examining more precisely the matrix $L_p C_p^T$ we have:

$$\begin{cases} (L_p C_p^T)_{i,j} & = x_{i,p} x_{j,p} \sum_{k=1}^m a_{i,k} a_{j,k} y_{k,p} y_{k,p-1} & \forall (i,j) \in \{1, \dots, n\}^2 \\ (L_p C_p^T)_{n+1,j} & = 0 & \forall j \in \{1, \dots, n\} \\ (L_p C_p^T)_{i,n+1} & = x_{i,p} \sum_{k=1}^m a_{i,k} a_{n+1,k} y_{k,p} y_{k,p-1} + x_{i,p} a_{i,m+1} & \forall i \in \{1, \dots, n\} \\ (L_p C_p^T)_{n+1,n+1} & = 1 \end{cases}$$

Let us note that i and j may be interchanged in the first equation above. Hence:

$$\forall (i,j) \in \{1, \dots, n\}^2 \quad (L_p C_p^T)_{i,j} = (L_p C_p^T)_{j,i}$$

Moreover we have for all $(i,j) \in \{1, \dots, n\} \times \{1, \dots, m\}$:

$$\begin{cases} \chi_{i,0} & = x_{i,0} \sum_{j=1}^{m+1} a_{i,j} y_{j,0} & = \frac{\sum_{j=1}^{m+1} a_{i,j}}{\sum_{j=1}^{m+1} a_{i,j}} = 1 \\ \gamma_{j,0} & = \sum_{i=1}^{n+1} \chi_{i,0}^{-1} x_{i,0} a_{i,j} y_{j,0} & = \sum_{i=1}^{n+1} \frac{a_{i,j}}{\sum_{j=1}^{m+1} a_{i,j}} \end{cases}$$

Since the sum of a line (or a column) of A cannot be equal to zeros, it exists two positive numbers $\underline{\gamma}, \bar{\gamma}$ such that:

$$0 < \underline{\gamma} 1 \leq \gamma_0 \leq \bar{\gamma} 1$$

We have thus:

$$\begin{aligned} 0 < \underline{\gamma} 1 & \leq \gamma_0 & \leq \bar{\gamma} 1 \\ \bar{\gamma}^{-1} 1 & \leq \gamma_0^{-1} & \leq \underline{\gamma}^{-1} 1 \\ \bar{\gamma}^{-1} 1 & \leq L_1 \gamma_0^{-1} & \leq \underline{\gamma}^{-1} 1 \\ 0 < \bar{\gamma}^{-1} 1 & \leq \chi_1 & \leq \underline{\gamma}^{-1} 1 \end{aligned}$$

where the last inequality is deduced from the fact that $\chi_1 = L_1 \gamma_0^{-1}$, and the fact that all the entries of L_1 are non negative.

Hence using (6) and a basic recursion we have :

$$\forall p \geq 1, \quad 0 < \bar{\gamma}^{-1} 1 \leq \chi_p \leq \underline{\gamma}^{-1} 1$$

Using $\gamma_p = C_p^T \chi_p^{-1}$ which is row stochastic we obtain:

$$\forall p \geq 1, \quad 0 < \underline{\gamma} \leq \gamma_p \leq \bar{\gamma} 1$$

From now on, let us suppose that all entries of the last column and the last line of A are positive.

Let us suppose that for some $i \in \{1, \dots, n\}$ we have $\lim_{p \rightarrow +\infty} x_{i,p} = +\infty$. Since we have:

$$\chi_{i,p} = x_{i,p} \left[\sum_{j=1}^m a_{i,j} y_{j,p} + a_{i,m+1} \right]$$

we also have $\lim_{p \rightarrow +\infty} \chi_{i,p} = +\infty$ which is impossible since χ is bounded. It exists thus an upper bound M_x such that:

$$\forall i \in \{1, \dots, n+1\} \forall p \geq 2 \quad x_{i,p} \leq M_x$$

In this case we have:

$$1 = y_{j,p+1} \sum_{i=1}^{n+1} a_{i,j} x_{i,p+1} \leq y_{j,p+1} M_x \sum_{i=1}^{n+1} a_{i,j} \leq y_{j,p+1} M_x \|A\|_1$$

Hence we have:

$$y_{j,p+1} \geq \frac{1}{M_x \|A\|_1}$$

In the same way, let us suppose that for some $j \in \{1, \dots, m\}$ we have $\lim_{p \rightarrow +\infty} y_{j,p} = +\infty$. Since we have:

$$\gamma_{j,p} = y_{j,p} \left[\sum_{i=1}^n \chi_{i,p}^{-1} a_{i,j} x_{i,p} + a_{n+1,j} \right]$$

we also have $\lim_{p \rightarrow +\infty} \beta_{j,p} = +\infty$ which is forbidden since γ is also upper bounded.

So, it exists an upper bound M_y such that:

$$\forall j \in \{1, \dots, m+1\}, \forall p \geq 2 \quad y_{j,p} \leq M_y$$

In this case we have:

$$1 = x_{i,p+1} \sum_{j=1}^{m+1} a_{i,j} y_{j,p} \leq x_{i,p+1} M_y \sum_{j=1}^{m+1} a_{i,j} \leq x_{i,p+1} M_y \|A\|_\infty$$

Hence:

$$x_{i,p+1} \geq \frac{1}{M_y \|A\|_\infty}$$

1. The series $(x_{i,p})_p$ and $(y_{j,p})_p$ being both lower bounded, all non zeros entries of $L_p C_p^T$ are lower bounded by a positive value. Moreover since $L_p C_p^T$ is row stochastic, all its entries are bounded by 1. We say that the entries of $L_p C_p^T$ are *uniformly positive*.

2. Moreover:

$$\forall i \in \{1, \dots, n\} \quad (L_p C_p)_{i,i} = x_{i,p}^2 \sum_{k=1}^m (a_{i,k})^2 y_{k,p} y_{k,p-1}$$

Since both series $(x_{i,p})_p$ and $(y_{j,p})_p$ are lower bounded and that a line of A cannot be equal to zeros, we have $(L_p C_p^T)_{i,i} > 0$. Let us additionally note that $(L_p C_p^T)_{n+1,n+1} = 1 > 0$.

3. Let us consider the Graph $\mathcal{G}(L_p C_p^T) = (V, E)$ where $V = \{1, \dots, n+1\}$ and an edge (i, j) connects node i to node j iff $(L_p C_p^T)_{j,i} > 0$. In this case since for any $i \in \{1, \dots, n\}$, $(L_p C_p^T)_{i,n+1} \geq x_{i,p} a_{i,m+1} > 0$. The node $n+1$ is adjacent to all nodes in $\{1, \dots, n\}$. Conversely, since $(L_p C_p^T)_{n+1,j} = 0$ for all $j \in \{1, \dots, n\}$, no node j is incident to $n+1$. The node $n+1$ defines a *source component*. Moreover, for any pair $(i, j) \in \{1, \dots, n\}^2$, since $(L_p C_p^T)_{i,j} = (L_p C_p^T)_{j,i}$, node i is adjacent to node j and vice versa. Hence $n+1$ is the only source component of $\mathcal{G}(L_p C_p^T)$ which is thus *quasi-strongly connected*.

Using [5], Lemma 3 with $T=1$ we can conclude that χ converges towards a consensus of the form $c1$. Since the last entry of χ is a constant equal to 1 we have $c = 1$ and:

$$\lim_{p \rightarrow +\infty} \chi_p = 1$$

Since C_p^T is row stochastic we have:

$$\lim_{p \rightarrow +\infty} \gamma_p = \lim_{p \rightarrow +\infty} C_p^T \chi_p^{-1} = 1$$

Since $\chi_{i,p} = \frac{x_{i,p}}{x_{i,p+1}}$ converges towards 1, $(x_{i,p})_p$ is a Cauchy serie in a complete space (\mathbb{R}^+) . Hence $(x_{i,p})_p$ converges. The same argument holds for $(y_{j,p})_p$.

Let us consider the two diagonal matrices $D_{1,p} = \text{diag}(x_{1,p}, \dots, x_{n+1,p})$ and $D_{2,p} = \text{diag}(y_{1,p}, \dots, y_{m+1,p})$ together with the matrix $S_p = D_{1,p} A D_{2,p}$. By construction we have for any $i \in \{1, \dots, n\}$:

$$(S_p 1)_i = \sum_{j=1}^{m+1} x_{i,p} a_{i,j} y_{j,p} = \chi_{i,p}$$

Thus: $\lim_{p \rightarrow \infty} (S_p 1)_i = \lim_{p \rightarrow \infty} \chi_{i,p} = 1$. The matrix S_p converges thus towards an ϵ row stochastic matrix.

Moreover, for any $j \in \{1, \dots, m\}$:

$$(S_p^T 1)_j = \sum_{i=1}^{n+1} x_{i,p} a_{i,j} y_{j,p} = \sum_{i=1}^{n+1} x_{i,p+1} a_{i,j} y_{j,p} + (x_{i,p} - x_{i,p+1}) a_{i,j} y_{j,p}$$

Since $x_{i,p}$ converges for any $i \in \{1, \dots, n+1\}$, $y_{j,p}$ is bounded, and the above sums are finite, it exists for any $\eta > 0$ a value p_0 such that for any $p \geq p_0$, we have :

$$\left| \sum_{i=1}^{n+1} x_{i,p} a_{i,j} y_{j,p} - \sum_{i=1}^{n+1} x_{i,p+1} a_{i,j} y_{j,p} \right| < \eta$$

Both sums converge (or diverge) thus toward a same value. Moreover:

$$\sum_{i=1}^{n+1} x_{i,p+1} a_{i,j} y_{j,p} = \gamma_{j,p}$$

We have thus: $\lim_{p \rightarrow \infty} (S_p^T 1)_j = \lim_{p \rightarrow \infty} \gamma_{j,p} = 1$. The matrix S_p converges thus toward an ϵ column stochastic matrix and hence an ϵ bi-stochastic matrix.

5 Two iterative algorithms

The code (in python) corresponding to the construction of matrices $D_{1,p}$ and $D_{2,p}$ is provided in Figure 4. You may note the fact that we set x_{n+1} and y_{m+1} to 1 respectively on line 9 and 15. This point together with the use of rectangular matrices is the main

```

1     def sinkhorn_D1D2(S, nb_iter, eps):
2         ones_n = torch.ones(S.shape[0], device=S.device)
3         ones_m = torch.ones(S.shape[1], device=S.device)
4         c=ones_m
5         converged=False
6         i=0
7         while i <=nb_iter and not converged:
8             xp=1.0/(S@c)
9             xp[-1]=1.0
10            if i>=1:
11                # computation of ||x_{p+1}/x_{p-1}||
12                norm_x=t1.norm(x/xp-torch.ones_like(x/xp), ord=float('inf'))
13            x=xp
14            yp=1.0/(S.T@r)
15            yp[-1]=1.0
16            # computation of ||y_{p+1}/y_{p-1}||
17            norm_y=t1.norm(y/yp-torch.ones_like(y/yp), ord=float('inf'))
18            y=yp
19            if i>=1:
20                converged= (norm_x <= eps) and (norm_y<=eps)
21            i+=1
22
23     return torch.diag(x)@S@torch.diag(y)

```

Figure 4: The code corresponding to the construction scheme described in Section 4

```

1     def sinkhorn_Sp(S, nb_iter, eps):
2         ones_n = torch.ones(S.shape[0], device=S.device)
3         ones_m = torch.ones(S.shape[1], device=S.device)
4
5         i=1
6         Sp=S
7         while i<=nb_iter and not converged:
8             D=torch.diag(1.0/(Sp@ones_m))
9             D[D.shape[0]-1,D.shape[1]-1]=1.0
10            Sp1 = D@Sp
11            norm_col=t1.norm((ones_n@Sk1-ones_m)[0:-1], ord=float('inf'))
12            D=torch.diag(1.0/(ones_n@Sp1))
13            D[D.shape[0]-1,D.shape[1]-1]=1.0
14            Sp = Sp1@D
15            norm_line=t1.norm((Sp@ones_m-ones_n)[0:-1], ord=float('inf'))
16            converged=(norm_col <= eps) and (norm_line<=eps)
17            i+=1
18            D=torch.diag(1.0/(Sp@ones_m))
19            D[D.shape[0]-1,D.shape[1]-1]=1.0
20            Sp1[0:n,:]=(D@Sp)[0:n,:]
21            Sp1[-1,0:m]=ones_m[0:m]-torch.sum((D@Sp)[0:n,0:m], dim=0)
22
23
24     return Sp1

```

Figure 5: A version of our computation of an ϵ assignment matrix which updates directly the matrix S_p

difference between this algorithm and the "classical" Sinkhorn algorithm. The convergence criterion which allows to avoid to loop up to the maximum number of iterations is based on the fact that both χ_p and γ_p converge toward a vector of 1.

An equivalent code computing directly the matrix S_p is provided in Figure 5. In this case the setting of x_{n+1} and y_{m+1} to 1 is performed on line 9 and 19. The stopping criterion is based on the computation of the distance of the current matrix to the set of ϵ assignment matrices. To do so, we compute the distance between the vector of 1 and $S_p^T \mathbf{1}$ after each row normalization (line 11). In the same way, we compute the distance between a vector of 1 and $S_p \mathbf{1}$ after each column normalization (line 15). After convergence, we apply a last row normalization before setting the last line of our ϵ assignment matrix to the complement to 1 of each column.

6 From similarity to cost matrices and vice versa

The Sinkhorn algorithm is well known for providing an approximation of the Linear Sum Assignment Problem (Section 7) which can be formulated as:

$$\max_X \sum_{i=1}^n \sum_{j=1}^n s_{i,j} x_{i,j}$$

where $S = (s_{i,j})$ is our similarity matrix and $X = (x_{i,j})$ is taken over all bi stochastic matrices. The optimal solution being a permutation matrix, hence a binary matrix.

This maximization problem may be translated into a minimization problem by considering the matrix $c \mathbf{1}_{n \times n} - S$, where $\mathbf{1}_{n \times n}$ is a $n \times n$ matrix filled of 1 and c is a positive constant greater than all values of S . We have indeed:

$$\begin{aligned} \sum_{i=1}^n \sum_{j=1}^n (c - s_{i,j}) x_{i,j} &= \sum_{i=1}^n \sum_{j=1}^n c x_{i,j} - \sum_{i=1}^n \sum_{j=1}^n s_{i,j} x_{i,j} \\ &= c \sum_{i=1}^n \sum_{j=1}^n x_{i,j} - \sum_{i=1}^n \sum_{j=1}^n s_{i,j} x_{i,j} \\ &= c \sum_{i=1}^n \mathbf{1} - \sum_{i=1}^n \sum_{j=1}^n s_{i,j} x_{i,j} \\ &= cn - \sum_{i=1}^n \sum_{j=1}^n s_{i,j} x_{i,j} \end{aligned}$$

Hence c and n being constant, minimize $\sum_{i=1}^n \sum_{j=1}^n (c - s_{i,j}) x_{i,j}$ is equivalent to maximize $\sum_{i=1}^n \sum_{j=1}^n s_{i,j} x_{i,j}$. The matrix $c \mathbf{1}_{n \times n} - S$ is usually interpreted as a cost matrix. This last point is important if one wants to compare the Sinkhorn algorithm to an optimal Hungarian algorithm which performs a minimization of costs instead of a maximization of similarities.

As stated in Section 1, our algorithms being an extension of the Sinkhorn algorithm we expect them to converge to :

$$\max_x \sum_{i=1}^{n+1} \sum_{j=1}^{m+1} s_{i,j} x_{i,j}$$

where $X = (x_{i,j})$ is taken over all ϵ -bi stochastic matrices. However, the transformation of this maximization of similarities into a minimization of costs, is slightly more complex in the case of ϵ assignment matrices. To do so, let us consider a $(n+1) \times (m+1)$ matrix $C = (c_{i,j})$ with:

$$c_{i,j} = \begin{cases} 2c & \text{if } i \leq n \wedge j \leq m \\ c_{lr} & \text{if } i = n+1 \wedge j \leq m \\ c_{lc} & \text{if } i \leq n \wedge j = m+1 \\ 0 & \text{if } i = n+1 \wedge j = m+1 \end{cases}$$

where c, c_{lr}, c_{lc} are three positive constants. Considering the cost matrix $C - S$ we have:

$$\sum_{i=1}^{n+1} \sum_{j=1}^{m+1} (c_{i,j} - s_{i,j}) x_{i,j} = \sum_{i=1}^{n+1} \sum_{j=1}^{m+1} c_{i,j} x_{i,j} - \sum_{i=1}^{n+1} \sum_{j=1}^{m+1} s_{i,j} x_{i,j}$$

If $c_{lc} = 2c$ and $c_{lr} = 0$: We have:

$$\begin{aligned}\sum_{i=1}^{n+1} \sum_{j=1}^{m+1} c_{i,j} x_{i,j} &= \sum_{i=1}^n \sum_{j=1}^{m+1} 2c x_{i,j} \\ &= 2c \sum_{i=1}^n 1 \\ &= 2cn\end{aligned}$$

If $c_{lr} = 2c$ and $c_{lc} = 0$: We have:

$$\begin{aligned}\sum_{i=1}^{n+1} \sum_{j=1}^{m+1} c_{i,j} x_{i,j} &= \sum_{j=1}^m \sum_{i=1}^{n+1} 2c x_{i,j} \\ &= 2c \sum_{j=1}^m 1 \\ &= 2cm\end{aligned}$$

If $c_{lr} = c$ and $c_{lc} = c$: We have:

$$\begin{aligned}\sum_{i=1}^{n+1} \sum_{j=1}^{m+1} c_{i,j} x_{i,j} &= 2c \sum_{i=1}^n \sum_{j=1}^m x_{i,j} + c \sum_{i=1}^n x_{i,m+1} + c \sum_{j=1}^m x_{n+1,j} \\ &= c \sum_{i=1}^n \sum_{j=1}^{m+1} x_{i,j} + c \sum_{j=1}^m \sum_{i=1}^{n+1} x_{i,j} \\ &= cn + cm\end{aligned}$$

In all cases we have thus:

$$\sum_{i=1}^{n+1} \sum_{j=1}^{m+1} (c_{i,j} - s_{i,j}) x_{i,j} = Q - \sum_{i=1}^{n+1} \sum_{j=1}^{m+1} s_{i,j} x_{i,j} \quad (7)$$

where $Q = cn$ if $c_{lc} = 2c$ and $c_{lr} = 0$, $Q = cm$ if $c_{lr} = 2c$ and $c_{lc} = 0$ and finally $Q = cn + cm$ if $c_{lr} = c_{lc} = c$. The minimization of the left part of equation 7 (minimization of costs) is thus equivalent to a maximization of the similarities.

Let us note that the trivial solution consisting to take $C = c\mathbb{1}_{(n+1) \times (m+1)}$ does not provide an equivalence between both problems since additional terms related either to the last column or the last row forbid to state that one problem is equal to a constant minus the other problem. Moreover, the last solution ($c_{lr} = c$ and $c_{lc} = c$) is the only one allowing to ensure that all coefficients of the similarity matrix are positive when transforming a cost matrix into a similarity matrix.

7 Experiments

We proposed in Section 4 two algorithms converging toward an unique solution if the conditions defined in Section 3 are satisfied. The aim of this section is to measure experimentally the convergence of our algorithms toward a solution maximizing :

$$\sum_{i=1}^{n+1} \sum_{j=1}^{m+1} s_{i,j} x_{i,j}$$

over all ϵ bi stochastic matrices X . Where S is the input matrix. Such a problem is called a Linear Sum Assignment Problem with Edition (LSAPE).

7.1 Deviation of the Sinkhorn algorithm from the optimal solution

Sinkhorn algorithm provides an approximate solution to the well known Linear Sum Assignment Problem (LSAP):

$$\max_X \sum_{i=1}^n \sum_{j=1}^m s_{i,j} x_{i,j}$$

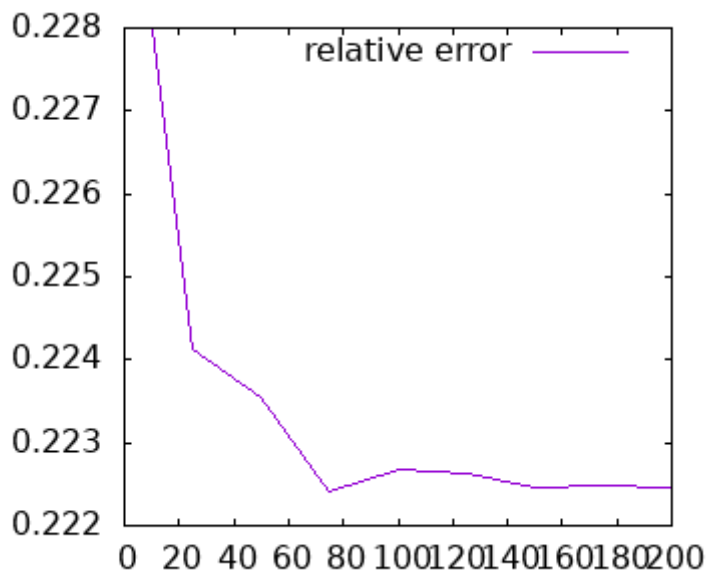


Figure 6: Relative Error between the results provided by the Sinkhorn Algorithm and the optimal solution provided by an Hungarian method.

where X is taken over the set of bi stochastic matrices. From a certain point of view, LSAP may be considered as a restriction of the LSAP with squared matrices and no deletions/insertions. Let us first evaluate the error induced by the use of the Sinkhorn algorithm. To do this, we define matrices filled by random number in the interval $[1, 2]$. For each matrix size we compute 100 matrices and compute for each matrix both the solution produced by the Sinkhorn algorithm and the optimal one produced by an Hungarian algorithm. The results of this experiment are displayed in Figure 6 for $n \in \{10, \dots, 200\}$.

Considering Figure 6, the error is approximately constant for all sizes of matrices lies between 20% and 23%.

7.2 Deviation of our algorithms from the optimal solution

In order to test our algorithm we use the same kind of random matrices but with a specific procedure for the last row and column which encode respectively the affinity of each element toward insertions and deletions:

$$S_{i,j} := \begin{cases} \text{rand}()+1 & \text{if } i < n \wedge j < m \\ 0 & \text{if } i = n \wedge j = m \\ h \text{rand}() & \text{else} \end{cases} \quad (8)$$

For $h \leq 0.5$ we can insure that for any $(i, j) \in \{1, \dots, n\} \times \{1, \dots, m\}$, $s_{i,j} \geq s_{n+1,j} + s_{j,m+1}$. In other terms we always get a greater sum by substituting i onto j than by deleting i and then inserting j . Conversely, if $s_{i,j} < s_{n+1,j} + s_{j,m+1}$ the substitution of i onto j will never be part of an optimal ϵ -assignment since this operation can be replaced, with a greater value of the sum, by the removal of i and the insertion of j .

Let us first focus on values of h lower than .5. Figure 7 shows on the first two columns the relative error of both sinkhorn_Sp and sinkhorn_D1D2 according to an optimal LSAP algorithm [2] for increasing sizes of the matrix. Let us note that [2] minimizes a sum of costs. We compare both algorithms using the results of Section 6. For each matrix size 100 random matrices are generated and the results are averaged for all three algorithm (sinkhorn_Sp, sinkhorn_D1D2 and the optimal one). Our algorithms provide an approximation of the LSAP which is slightly above 20%, hence comparable with

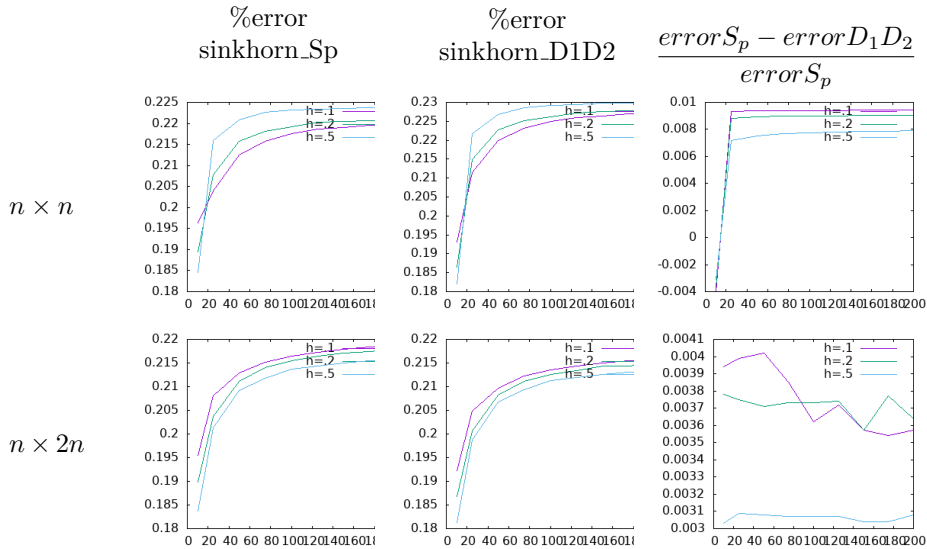


Figure 7: Relative errors of sinkhorn_Sp and sinkhorn_D1D2 according to the optimal solution.

the one provided by the Sinkhorn algorithm for the LSAP problem. Interestingly, our algorithm provides better approximations for small matrix sizes while Figure 6 suggest an opposite behavior for the Sinkhorn algorithm.

The last column of Figure 7 allows to compare more precisely sinkhorn_Sp and sinkhorn_D1D2. Both algorithms seems to be equivalent since the relative error between both methods do not exceed 1%.

For h greater than .5 we observe in Figure 8 that we keep an error of about 20% for $h = 1.0$ and $h = 2.0$. In these cases the values of the last line and the last column remain comparable with the inner values of the random matrix. However, for larger values of h , namely $h \in \{4.0, 6.0, 8.0\}$ we observe a large increase of the relative error especially in the case of squared matrices. We can conclude from these experiments that our algorithms do not converge to the expected value when the values of the last column/line are very large compared to the inner values. More precisely, when we have:

$$s_{i,j} \ll s_{n+1,i} + s_{i,m+1} \text{ for } (i,j) \in \{1, \dots, n\} \times \{1, \dots, m\}$$

A simple solution to fix this problem, consists in simplifying the similarity matrix by removing (setting to a low value) any entry (i,j) , $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$ such that (i,j) cannot belong to any optimal solution. Given the similarity matrix S , such entries are characterized by $s_{i,j} < s_{n+1,j} + s_{i,m+1}$. In such cases, the substitution of i onto j may be advantageously replaced by the removal of i and the insertion of j . This last point forbids the assignment of i onto j in any optimal ϵ -assignment.

Figure 9 represents the relative errors according to the optimal solution performed by Sinkhorn_Sp and sinkhorn_D1D2 using this simplification of the similarity matrix. In this experiment the "low value" replacing any entry of the matrix S which can not be included in any optimal solution has been fixed to 10^{-4} . One can first observe that we get the same behavior for the squared ($n \times n$) case and the rectangular one ($n \times 2n$). We can further observe that all errors remain below 20% for all sizes. Moreover, the relative error appear to be decreasing as a function of h for both algorithms. This may be explained by the fact that as h get higher, the simplified matrix becomes more and more trivial. Indeed for largest values of h , simplified similarity matrices correspond to trivial matrices with a constant value (equal to 10^{-4} in this experiment) for all entries (i,j) in $\{1, \dots, n\} \times \{1, \dots, m\}$ and a last row and column which remains unchanged and

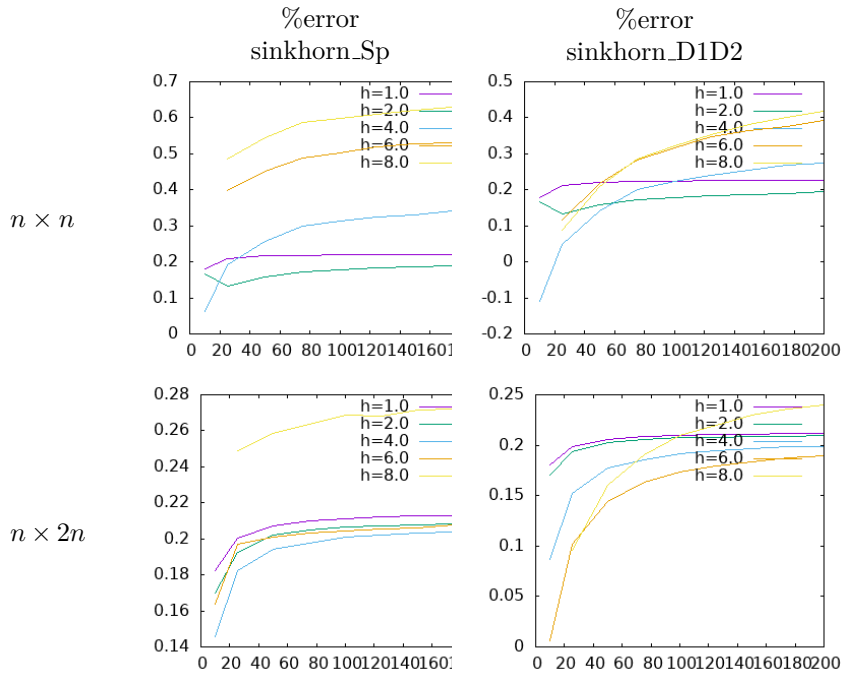


Figure 8: Relative errors of sinkhorn_Sp and sinkhorn_D1D2 according to the optimal solution.




Figure 9: Relative errors of sinkhorn_Sp and sinkhorn_D1D2 according to the optimal solution after a simplification of the similarity matrix.

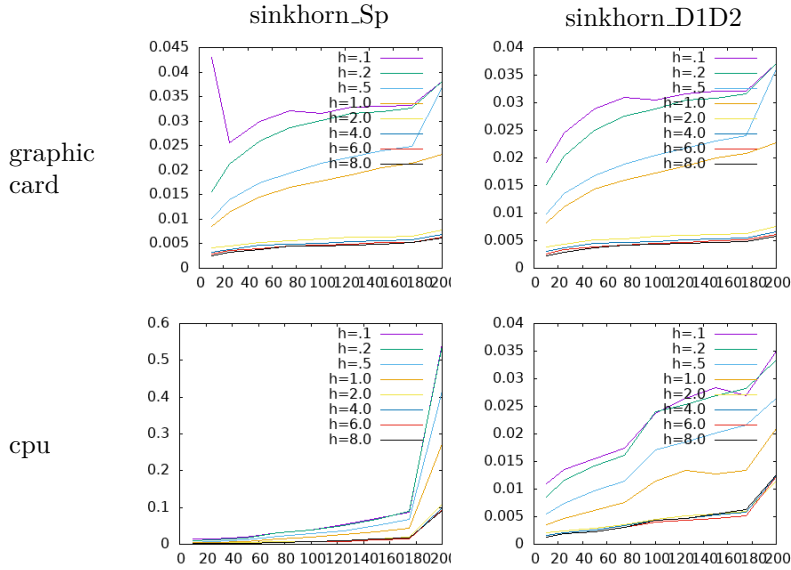


Figure 10: Execution times of `sinkhorn_Sp` and `sinkhorn_D1D2` on squared $n \times n$ matrices using both graphic card and cpu computation.

greater than 10^{-4} by several orders of magnitude. In such cases our algorithms converge immediately to the optimal solution which correspond to the removal of all elements in $\{1, \dots, n\}$ and the insertion of all elements in $\{1, \dots, m\}$.

7.3 Execution times of our algorithms

The execution times of both `sinkhorn_Sp` and `sinkhorn_D1D2` computed either on a graphic card (Nvidia Quadro P2000) or on a CPU (Intel Core i5 650@3.2GHz) are displayed in Figures 10 for $n \times n$ squared matrices and in Figure 11 for $n \times 2n$ rectangular matrices. In both figures and for each size, both algorithms have been run 100 times and the execution times have been averaged.

Considering the squared case (Figure 10) we can see that on graphics cards both `sinkhorn_Sp` and `sinkhorn_D1D2` take roughly the same amount of time. This is clearly not the case on CPU where the updates of the matrix S_p instead of the two diagonal matrices D_1 and D_2 induce a large difference, by a factor greater than 10, between the execution times of `sinkhorn_Sp` and `sinkhorn_D1D2`. We can further observe that if we do not observe major differences between the cuda and cpu execution times for `sinkhorn_D1D2` this is clearly not the case for `sinkhorn_Sp` which is more than 10 times accelerated by the use of the graphic card. Finally, we can note that the execution times are decreasing as a function of h . As previously, this last point is due to the fact that as h get higher the simplified similarity matrices become more and more trivial and our iterative algorithms need less and less iterations to converge.

Concerning the rectangular case (Figure 11), we observe the same trends than in the squared case. However the ratio between the execution times of `sinkhorn_Sp` and `sinkhorn_D1D2` on CPU is, in this case, about 6. Moreover, we observe for such matrices a factor approximately equal to 2 between the execution times of `sinkhorn_D1D2` on GPU and CPU. As in the squared case the CUDA implementation provides a large speedup for `sinkhorn_Sp`. Finally, the execution times of both algorithms decrease as h increase. This phenomenon already encountered in the squared case is due to the same reasons.

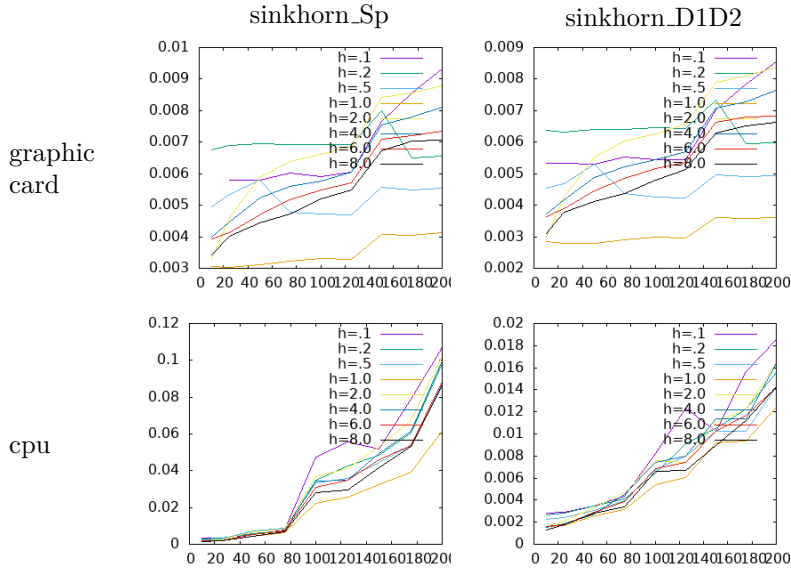


Figure 11: Execution times of `sinkhorn_Sp` and `sinkhorn_D1D2` on $n \times 2n$ rectangular matrices using both graphic card and cpu computation.

8 Conclusion

We have presented in this report (Section 5) two algorithms. The proof of their convergence is provided in Section 4 while conditions of the existence and uniqueness of the limits are provided in Section 3. Section 6 provides simple methods to transform the sum maximization problem addressed by our algorithms into sum minimization (minimization of a sum of costs).

As shown in Section 7, these algorithms provide an approximate solution to the Linear Sum Assignment Problem with Edition (LSAPE). The relative error of these algorithms compared to the optimal solutions is similar, and even much lower in some cases, to the relative error between the classical Sinkhorn and the optimal solutions to the Linear Sum Assignment Problem (LSAP). The main difference between these algorithms and the Hungarian based algorithms providing the optimal solution is that our algorithms are iterative and differentiable and may thus be easily inserted within a backpropagation based learning framework such as artificial neural networks.

Compared to the LSAP, the LSAPE problem allows to manage assignments between sets of different sizes by allowing the possibility to reject some elements from the matching through the use of insertion/deletion operations. Let us note that these insertion/deletion operations are integrated to the matching algorithm. The LSAPE thus avoids any artificial preprocessing step consisting in the selection of the more promising objects of both sets for matching. Our algorithms output an ϵ bi-stochastic matrix which mainly differs from the results provided by [4, 3] through the explicit notions of insertion and deletion.

Let us finally note that we have required in Section 5 that all the coefficients of the last line and column of the similarity matrix should be positive. This requirement is certainly too strong and weaker conditions may certainly be established by future works. An alternative future research direction, consists in addressing the LSAPE problem using the recent advances in Graph Neural Networks (GNN).

References

- [1] Sébastien Bougleux and Luc Brun. Linear Sum Assignment with Edition. Research report, Normandie Université ; GREYC CNRS UMR 6072, March 2016.
- [2] Sébastien Bougleux, Benoit Gaüzère, and Luc Brun. A Hungarian Algorithm for Error-Correcting Graph Matching. In Pasquale Foggia, Cheng-Lin Liu, and Mario Vento, editors, *11th IAPR-TC-15 International Workshop on Graph-Based Representation in Pattern Recognition (GbrPR 2017)*, volume 10310 of *Lecture notes in Computer Sciences (LNCS)*, pages 118–127, AnaCapri, Italy, May 2017. Pasquale Foggia, Springer.
- [3] Aude Genevay, Gabriel Peyre, and Marco Cuturi. Learning generative models with sinkhorn divergences. In Amos Storkey and Fernando Perez-Cruz, editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 1608–1617. PMLR, 09–11 Apr 2018.
- [4] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [5] Anton V. Proskurnikov, Giuseppe Carlo Calafiore, and Ming Cao. Recurrent averaging inequalities in multi-agent control and social dynamics modeling. *Annu. Rev. Control.*, 49:95–112, 2020.
- [6] Richard Dennis Sinkhorn and Paul Joseph Knopp. concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348, 1967.