

**Max-Planck-Institut
für Mathematik
in den Naturwissenschaften
Leipzig**

**A note on overrelaxation in the
Sinkhorn algorithm**

(revised version: December 2021)

by

*Tobias Lehmman, Max von Renesse,
Alexander Sambale, and André Uschmajew*

Preprint no.: 110

2020



A note on overrelaxation in the Sinkhorn algorithm

Tobias Lehmann* Max-K. von Renesse*

Alexander Sambale* André Uschmajew†

Abstract

We derive an a priori parameter range for overrelaxation of the Sinkhorn algorithm, which guarantees global convergence and a strictly faster asymptotic local convergence. Guided by the spectral analysis of the linearized problem we pursue a zero cost procedure to choose a near optimal relaxation parameter.

1 Introduction and statement of result

The Sinkhorn algorithm is the benchmark approach to fast computation of the entropic regularization of optimal transportation [4]. Ultimately, one is faced with the following numerical problem: Given two probability vectors $a \in \mathbb{R}_+^m$, $b \in \mathbb{R}_+^n$ and a matrix $K \in \mathbb{R}_+^{m \times n}$, the goal is to find a pair of vectors $(u, v) \in \mathbb{R}_+^m \times \mathbb{R}_+^n$ such that

$$u \circ Kv = a \quad \text{and} \quad v \circ K^\top u = b, \quad (1)$$

where $x \circ y$ denotes the componentwise multiplication (Hadamard product) of vectors of equal dimension. Here \mathbb{R}_+ refers to the positive reals. We assume $\min(m, n) \geq 2$.

In the standard Sinkhorn algorithm an approximating sequence (u_ℓ, v_ℓ) starting from an initial vector $v_0 \in \mathbb{R}_+^n$ is constructed via the update rule

$$u_{\ell+1} = \frac{a}{Kv_\ell}, \quad v_{\ell+1} = \frac{b}{K^\top u_{\ell+1}},$$

where $\frac{x}{y}$ denotes the componentwise division of vectors of equal dimension. It is a classic result by Sinkhorn [16] that for any initial point $v_0 \in \mathbb{R}_+^n$ the algorithm converges to a solution (u^*, v^*) of (1), which is unique modulo rescaling $(tu^*, t^{-1}v^*)$, $t > 0$. Moreover, the convergence, e.g. of suitably normalized iterates $u_\ell/\|u_\ell\|$ and $v_\ell/\|v_\ell\|$, or using other equivalent distance measures like the Hilbert metric, is R-linear with an asymptotic rate at least $\Lambda(K)^2$, where $\Lambda(K) < 1$ is the Birkhoff contraction ratio defined in (8) further below [8]. See also [15] for an overview.

In this note we discuss a modified version of the Sinkhorn algorithm employing relaxation, which was recently proposed in [18] and [14]. It uses the update rule

$$u_{\ell+1} = u_\ell^{1-\omega} \circ \left(\frac{a}{Kv_\ell} \right)^\omega, \quad v_{\ell+1} = v_\ell^{1-\omega} \circ \left(\frac{b}{K^\top u_{\ell+1}} \right)^\omega, \quad (2)$$

*Universität Leipzig, Fakultät für Mathematik und Informatik, Augustusplatz 10, 04109 Leipzig, Germany

†Max Planck Institute for Mathematics in the Sciences, 04103 Leipzig, Germany

where $\omega > 0$ is a suitably chosen relaxation parameter, and exponentiation is understood componentwise. In a log-domain formulation such as (7) further below, the relation to the classic concept of relaxation in (nonlinear) fixed point iterations will become immediately apparent. Note that the iteration (2) still has the solution of (1) as its unique (modulo scaling) fixed point. As illustrated in [18] and [14], choosing the parameter ω larger than one can significantly accelerate the convergence speed compared to the standard Sinkhorn method, which sometimes can be slow. For optimal transport, such an improvement could be in particular relevant in the regime of small regularization, or when a high target precision is needed, such as in applications in density functional theory [3].

While global convergence for $\omega \neq 1$ is not obvious anymore, local convergence of the modified method is ensured for all $0 < \omega < 2$, and the asymptotically optimal relaxation parameter can be determined from its linearization at a fixed point (u_*, v_*) . In logarithmic coordinates, the linearization of the standard Sinkhorn method has the iteration matrix

$$M = \text{diag}\left(\frac{1}{a}\right) P_* \text{diag}\left(\frac{1}{b}\right) P_*^\top, \quad \text{where } P_* = \text{diag}(u^*) K \text{diag}(v^*). \quad (3)$$

The local convergence rate equals the second largest eigenvalue

$$0 \leq \vartheta^2 < 1$$

of that matrix; see [11]. Note that M has real and nonnegative eigenvalues since it is similar to a positive semidefinite matrix, and its largest eigenvalue equals one (the eigenvector having constant entries), which accounts for the scaling indeterminacy in the problem formulation. For the modified method with relaxation, the local rate is also related to ϑ^2 , which has been worked out in [18] and is summarized in the following theorem. For convenience, we provide a brief outline how this result can be obtained at the end of section 2.

Theorem 1 (cf. [18]). *Assume $\vartheta^2 > 0$. For all choices of $0 < \omega < 2$ the modified Sinkhorn algorithm (2) is locally convergent in some neighborhood of (u^*, v^*) . Its asymptotic (R-linear) convergence rate is*

$$\rho_\vartheta(\omega) := \begin{cases} \frac{1}{4} \left(\omega \vartheta + \sqrt{\omega^2 \vartheta^2 - 4(\omega - 1)} \right)^2, & \text{if } 0 < \omega \leq \omega^{\text{opt}}, \\ \omega - 1, & \text{if } \omega^{\text{opt}} \leq \omega < 2, \end{cases} \quad (4)$$

where

$$\omega^{\text{opt}} = \frac{2}{1 + \sqrt{1 - \vartheta^2}} > 1. \quad (5)$$

It holds $\rho_\vartheta(\omega) < 1$ for all $0 < \omega < 2$, and ω^{opt} provides the minimal possible rate (independent of the starting point) on that interval, namely

$$\rho^{\text{opt}} = \omega^{\text{opt}} - 1 = \frac{1 - \sqrt{1 - \vartheta^2}}{1 + \sqrt{1 - \vartheta^2}} < \vartheta^2.$$

By the above theorem, the optimal relaxation parameter ω^{opt} is always larger than one (if $\vartheta^2 > 0$). In fact, by the exact formula (4) for the convergence rate, the range of ω for which the modified method is asymptotically strictly faster than the standard Sinkhorn method, that is, $\rho_\vartheta(\omega) < \vartheta^2 = \rho_\vartheta(1)$, is precisely the interval

$$1 < \omega < 1 + \vartheta^2. \quad (6)$$

However, the value of ϑ^2 depends on the solution and is therefore not known in advance. To deal with this problem, an adaptive procedure for choosing ω is proposed in [18].

As our contribution, the main goal in this note is to provide an a priori interval for the relaxation parameter ω for which the modified iteration is both globally convergent and locally faster than the standard Sinkhorn method. In Theorem 3 we first prove global convergence of the modified method for parameters in the interval $0 < \omega < \frac{2}{1+\Lambda(K)}$. In Theorem 4 we then provide an a priori lower bound $\vartheta^2 \geq \delta_{K,a,b} > 0$, which depends only on the data of the problem, but requires a full rank assumption on K . By (6), any $\omega \in (1, 1 + \delta_{K,a,b})$ then satisfies $\rho_\vartheta(\omega) < \vartheta^2$. Taken together this yields the following result.

Theorem 2. *Assume $\text{rank}(K) = \min(m, n) \geq 2$. For any $1 < \omega < 1 + \vartheta^2$ the asymptotic local convergence rate of the modified Sinkhorn method (2) is faster than for the standard Sinkhorn method. For $1 < \omega < \min\left(1 + \delta_{K,a,b}, \frac{2}{1+\Lambda(K)}\right)$ the modified method is both globally convergent and asymptotically faster than the standard method.*

We remark that our derived a priori interval for ω is usually very small, and hence our result is of rather theoretical interest. In the relevant cases, when ϑ^2 is close to one, significant acceleration is achieved only when ω is close to ω^{opt} (which tends to two for $\vartheta^2 \rightarrow 1$). A possible heuristic to select a nearly optimal relaxation is to approximate the second largest eigenvalue of M based on the current iterate. After a similarity transform, this requires to compute the spectral norm of a symmetric matrix. An even simpler approach, as suggested in [18], is to directly estimate ϑ^2 , and hence ω^{opt} , by monitoring the convergence rate of the standard Sinkhorn method in terms of a suitable residual. In the final section 4 we include numerical illustrations, which indicate that in certain cases such heuristics can be quite precise already in the initial phase of the algorithm, resulting in the almost optimal convergence rate at almost no additional cost. This confirms that overrelaxation is a simple way to significantly accelerate the Sinkhorn method in cases where it is slow. For completeness, we should mention that alternative approaches for solving problem (1) and aiming at fast convergence have been proposed based on Newton's method, see, e.g., [12, 2] and references therein.

The convergence analysis of the Sinkhorn method is usually carried out in a log-domain formulation [15]. We choose the closely related framework of compositional data space used, e.g., in statistics [13], which we think could be of independent interest in this context. In this space, which is introduced in the next section, the Sinkhorn algorithm with a positive matrix K reads as a nonlinear fixed point iteration for an essentially contractive iteration function, as is known from the Birkhoff–Hopf theorem. The main results are then presented in Section 3. Let us note that the assumption that K has strictly positive entries is not essential for all of the results. While global convergence of the standard Sinkhorn method to a unique (up to scaling) positive solution (u, v) of (1) can be shown under several weaker assumptions, most notably when $a = b = \mathbf{1}$ and K is square, nonnegative and has total support [17], we require the global contractivity of the process in Hilbert metric (which holds for positive K) in our proof that global convergence can still be ensured for some $\omega > 1$ (Theorem 3). The idea of accelerating convergence by overrelaxation, on the other hand, is very general and the local spectral analysis provided by Theorem 1 applies whenever the iteration (2) is locally well defined around a (positive) fixed point (u^*, v^*) and $\vartheta^2 < 1$. Correspondingly, Theorem 4 on a lower bound for ϑ^2 does not require K to be positive. Hence one has guaranteed acceleration of local convergence for $1 < \omega < 1 + \delta_{K,a,b}$ in several scenarios where K is only nonnegative.

2 Formulation in compositional data space

The problem (1) as well as the Sinkhorn algorithm and its modified variant inherit a natural scaling indeterminacy of the variables u and v . It can be therefore formulated in a suitable equivalence space. Here we recast the algorithm in the framework of what is called *compositional data space*; see, e.g., [13, 1]. To this aim, let

$$\mathcal{C}^m := \mathbb{R}_+^m / \sim,$$

where

$$x \sim x' \iff \exists t > 0 : x = tx'.$$

The resulting equivalence class of x will be denoted by \underline{x} . One specifies a vector addition and a scalar multiplication on \mathcal{C}^m via

$$\underline{x} + \underline{y} := \underline{x \circ y}, \quad \gamma \cdot \underline{x} := \underline{x^\gamma}, \quad \gamma \in \mathbb{R},$$

where x^γ has the components $(x_1^\gamma, \dots, x_n^\gamma)$. As a result $(\mathcal{C}^m, +, \cdot)$ becomes a real vector space of dimension $m - 1$. In this space we consider the so called *Hilbert norm*

$$\|\underline{x}\|_H := \log \max_{i,j} \frac{x_i}{x_j},$$

turning $(\mathcal{C}^m, +, \cdot, \|\cdot\|_H)$ into a finite dimensional Banach space. Note that this norm on the equivalence classes coincides with the well-known *Hilbert distance* on the representatives:

$$d_H(x, y) = \|\underline{x} - \underline{y}\|_H.$$

Similarly we construct a Banach space $\mathcal{C}^n = \mathbb{R}_+^n / \sim$.

The modified Sinkhorn algorithm (2) can be interpreted as an iteration in the space $\mathcal{C}^m \times \mathcal{C}^n$ and reads

$$\begin{aligned} \underline{u}_{\ell+1} &= (1 - \omega) \cdot \underline{u}_\ell + \omega \cdot \underline{a} - \omega \cdot \mathcal{K}(\underline{v}_\ell), \\ \underline{v}_{\ell+1} &= (1 - \omega) \cdot \underline{v}_\ell + \omega \cdot \underline{b} - \omega \cdot \mathcal{K}_\top(\underline{u}_{\ell+1}), \end{aligned} \tag{7}$$

where $\mathcal{K}: \mathcal{C}^n \rightarrow \mathcal{C}^m$ and $\mathcal{K}_\top: \mathcal{C}^m \rightarrow \mathcal{C}^n$ are now the *nonlinear* maps given by

$$\mathcal{K}(\underline{v}) = \underline{Kv}, \quad \mathcal{K}_\top(\underline{u}) = \underline{K^\top u}.$$

The convergence of the standard Sinkhorn algorithm ($\omega = 1$) is based on a famous result of Birkhoff and Hopf on the contractivity of \mathcal{K} and \mathcal{K}_\top . To state it, define the quantities

$$\eta(K) := \max_{i,j,k,\ell} \frac{K_{ik}K_{j\ell}}{K_{jk}K_{i\ell}} \quad \text{and} \quad \Lambda(K) := \frac{\sqrt{\eta(K)} - 1}{\sqrt{\eta(K)} + 1}. \tag{8}$$

Then the following holds; for a proof, see, e.g., [5, Theorems 3.5 & 6.2].

Theorem (Birkhoff–Hopf). *For any $K \in \mathbb{R}_+^{m \times n}$ and $v, v' \in \mathbb{R}_+^n$ let $\Lambda(K)$ be defined as above. Then*

$$\sup_{v, v' \in \mathbb{R}_+^n} \frac{d_H(Kv, Kv')}{d_H(v, v')} = \Lambda(K).$$

Note that $\Lambda(K) = \Lambda(K^\top) < 1$. As a result, both \mathcal{K} and \mathcal{K}_\top are contractive maps in the Hilbert norm with Lipschitz constant $\Lambda(K)$, which is also called the *Birkhoff contraction ratio*

of K . Based on this, it is not difficult to establish the global convergence of the standard Sinkhorn algorithm in the space $\mathcal{C}^m \times \mathcal{C}^n$ at a rate $O(\Lambda(K)^2)$.

It is important to emphasize that studying the convergence in $\mathcal{C}^m \times \mathcal{C}^n$, that is, convergence of equivalence classes, is sufficient for understanding the method in $\mathbb{R}_+^m \times \mathbb{R}_+^n$. Indeed, a pair $(\underline{u}^*, \underline{v}^*)$ is a fixed point of (7) if and only if for any choice of representatives (u^*, v^*) there exist λ, μ such that $\lambda u \circ K v = a$ and $\mu v \circ K^\top u = b$. From $\mathbf{1}_m^\top a = \mathbf{1}_n^\top b$ (here $\mathbf{1}$ denotes a vector of all ones) it follows that $\lambda = \mu$, and hence, e.g. $u^+ := \lambda^{-1/2} u^*$ and $v^+ := \lambda^{-1/2} v^*$ solve the initial problem (1), where $\lambda = u^\top K v$. Moreover, choosing representatives (u_ℓ, v_ℓ) of the iterates $(\underline{u}_\ell, \underline{v}_\ell)$ such that $\mathbf{1}_m^\top u_\ell = \mathbf{1}_n^\top v_\ell = 1$, and setting $u_\ell^+ := \lambda_\ell^{-\frac{1}{2}} u_\ell$, $v_\ell^+ := \lambda_\ell^{-\frac{1}{2}} v_\ell$ with $\lambda_\ell = u_\ell^\top K v_\ell$, yields a sequence which converges exponentially fast to (u^*, v^*) .

We now briefly outline how the local convergence analysis for (7) can be conducted [18], leading to Theorem 1. By combining both steps of the iteration (7) into a nonlinear fixed point iteration $(\underline{u}_{\ell+1}, \underline{v}_{\ell+1}) = \mathcal{F}(\underline{u}_\ell, \underline{v}_\ell)$ in the space $\mathcal{C}^m \times \mathcal{C}^n$, one finds that its derivative at the fixed point $(\underline{u}^*, \underline{v}^*)$ takes the form

$$M_\omega := (I_{m+n} - \omega \cdot L)^{-1}[(1 - \omega) \cdot I_{m+n} + \omega \cdot U], \quad (9)$$

where

$$L = \begin{pmatrix} 0 & 0 \\ -\mathcal{K}'_\top(\underline{u}^*) & 0 \end{pmatrix}, \quad U = \begin{pmatrix} 0 & -\mathcal{K}'(\underline{v}^*) \\ 0 & 0 \end{pmatrix}.$$

Matrices of the form M_ω are well known as error iteration matrices of block SOR methods for linear systems. The spectral radius of M_ω can be computed exactly from formula (4), if the spectral radius ϑ of $L + U$ is known; see [19, Sec. 6.2] or [9, Thm. 4.27]. The eigenvalues of $L + U$, however, are square roots of the eigenvalues of the composition of derivatives $\mathcal{K}'(\underline{v}^*)\mathcal{K}'_\top(\underline{u}^*)$, which is a linear map on \mathcal{C}^m . It remains to show that the largest eigenvalue of that operator is precisely the second largest eigenvalue of the matrix M in (3). Indeed, by elementary calculations, M is the matrix representation of $\mathcal{K}'(\underline{v}^*)\mathcal{K}'_\top(\underline{u}^*)$ under the isomorphism $u \mapsto \underline{\exp}(u)$ between the subspace $\{u \in \mathbb{R}^m : \mathbf{1}_m^\top u = 0\} \subseteq \mathbb{R}^m$ and \mathcal{C}^m .

3 Main results

We prove the global convergence of the modified method for a range of values ω larger than one.

Theorem 3. *Let $\Lambda = \Lambda(K)$ be the Birkhoff contraction ratio of K . For $0 < \omega < \frac{2}{1+\Lambda}$, the modified Sinkhorn algorithm (7) converges, for any starting point, to $(\underline{u}^*, \underline{v}^*)$ exponentially fast.*

Proof. Starting from (7), using the triangle inequality and the contractivity of \mathcal{K} and \mathcal{K}_\top provided by the Birkhoff-Hopf theorem, we obtain

$$\begin{aligned} \|\underline{u}_{\ell+1} - \underline{u}^*\|_H &\leq |1 - \omega| \|\underline{u}_\ell - \underline{u}^*\|_H + \omega \Lambda \|\underline{v}_\ell - \underline{v}^*\|_H, \\ \|\underline{v}_{\ell+1} - \underline{v}^*\|_H &\leq |1 - \omega| \|\underline{v}_\ell - \underline{v}^*\|_H + \omega \Lambda \|\underline{u}_{\ell+1} - \underline{u}^*\|_H \\ &\leq (|1 - \omega| + (\omega \Lambda)^2) \|\underline{v}_\ell - \underline{v}^*\|_H + \omega \Lambda |1 - \omega| \|\underline{u}_\ell - \underline{u}^*\|_H. \end{aligned}$$

As a consequence, for $\Delta u_\ell := \|\underline{u}_{\ell+1} - \underline{u}^*\|_H$ and $\Delta v_\ell := \|\underline{v}_{\ell+1} - \underline{v}^*\|_H$ we obtain

$$\begin{pmatrix} \Delta u_{\ell+1} \\ \Delta v_{\ell+1} \end{pmatrix} \leq T_\omega \begin{pmatrix} \Delta u_\ell \\ \Delta v_\ell \end{pmatrix}, \quad \text{where} \quad T_\omega = \begin{pmatrix} |1 - \omega| & \omega \Lambda \\ \omega \Lambda |1 - \omega| & |1 - \omega| + (\omega \Lambda)^2 \end{pmatrix},$$

and the vector inequality is understood entry-wise. Since all involved quantities are non-negative the inequality can be iterated, which gives

$$\begin{pmatrix} \Delta u_{\ell+1} \\ \Delta v_{\ell+1} \end{pmatrix} \leq (T_\omega)^{\ell+1} \begin{pmatrix} \Delta u_0 \\ \Delta v_0 \end{pmatrix}.$$

Hence, to prove exponential convergence it suffices to show that the spectral radius of T_ω is strictly less than one. Since the spectral radius equals $|1 - \omega| + \frac{(\omega\Lambda)^2}{2} + \sqrt{\frac{(\omega\Lambda)^4}{4} + (\omega\Lambda)^2|1 - \omega|}$ this is the case if and only if $0 < \omega < 2/(1 + \Lambda)$. \square

Next we provide a lower bound for the second largest eigenvalue ϑ^2 of the matrix M in (3), which by (6) then yields an interval for ω such that the modified method has a strictly faster asymptotic convergence rate than the standard Sinkhorn method.

Theorem 4. *Let $\text{rank}(K) = \min(m, n) \geq 2$ and*

$$\delta_1 = \frac{a_{\min}}{b_{\max}} \cdot \frac{1 - b_{\max}}{\left(\frac{\|K\|_\infty}{\sigma_{\min}(K)}\right)^2 - a_{\min}} > 0, \quad \delta_2 = \frac{b_{\min}}{a_{\max}} \cdot \frac{1 - a_{\max}}{\left(\frac{\|K^\top\|_\infty}{\sigma_{\min}(K)}\right)^2 - b_{\min}} > 0,$$

where $\sigma_{\min}(K)$ is the smallest positive singular value of K , $\|K\|_\infty = \max_{\|v\|_\infty=1} \|Kv\|_\infty$, and the subscripts min, max denote the smallest and largest entry of the corresponding vector. Then it holds

$$\vartheta^2 \geq \delta_{K,a,b} := \begin{cases} \delta_1 & \text{if } m > n, \\ \delta_2 & \text{if } m < n, \\ \max(\delta_1, \delta_2) & \text{if } m = n. \end{cases}$$

Note that for a positive matrix $\|K\|_\infty > \sigma_{\min}(K)$. Moreover, $a_{\min} \leq \frac{1}{m} \leq \frac{1}{n} \leq b_{\max} < 1$ if $m \geq n$, and vice versa if $m \leq n$. Hence $\delta_{K,a,b}$ is indeed smaller than one, which is in line with the bound $\vartheta^2 \leq \Lambda(K)^2$.

Proof. We consider the case $m \geq n$. Instead of matrix M we consider the positive semidefinite matrix

$$H = \text{diag}\left(\frac{u^*}{a^{1/2}}\right) K \text{diag}\left(\frac{v^* \circ v^*}{b}\right) K^\top \text{diag}\left(\frac{u^*}{a^{1/2}}\right) \in \mathbb{R}^{m \times m},$$

which is obtained from M by a similarity transformation (and using (1)). Since the dominant eigenvector of H (with eigenvalue one) is $a^{1/2}$, we have

$$\vartheta^2 = \max \left\{ \frac{\langle w, Hw \rangle}{\langle w, w \rangle} : \langle w, a^{1/2} \rangle = 0 \right\}.$$

By projecting on the orthogonal complement of $a^{1/2}$, and noting that $\|a^{1/2}\|^2 = 1$, we first rewrite this as

$$\vartheta^2 = \max \frac{\langle w, Hw \rangle - \langle w, a^{1/2} \rangle^2}{\langle w, w \rangle - \langle w, a^{1/2} \rangle^2},$$

where the maximum is taken over all w that are not collinear to $a^{1/2}$. For such w the numerator is always nonnegative and the denominator is positive. Next we substitute

$$w = a^{-1/2} \circ K v^* \circ z$$

with a new variable z . This yields

$$\vartheta^2 = \max \frac{\langle K^\top z, \frac{v^* \circ v^*}{b} \circ K^\top z \rangle - \langle K^\top z, v^* \rangle^2}{\langle z, \frac{K v^* \circ K v^*}{a} \circ z \rangle - \langle K^\top z, v^* \rangle^2},$$

where the maximum is taken over all z not collinear with u^* (the numerator is then nonnegative and the denominator is positive). To obtain a lower bound, we now evaluate the expression at z satisfying

$$K^\top z = e_j$$

where e_j denotes the j -th unit vector. Note that such z exists (K^\top has full row rank) and is indeed not collinear to u^* , since otherwise $K^\top u^*$ would be collinear with e_j , which contradicts $K^\top u^* \circ v^* = b$. Therefore, using this z , we get

$$\vartheta^2 \geq \frac{\left(\frac{1}{b_{\max}} - 1\right) (v^*)_j^2}{\langle z, \frac{K v^* \circ K v^*}{a} \circ z \rangle - (v^*)_j^2}.$$

We can choose j as the position of a largest entry of the vector v^* . Then in the denominator

$$\left\langle z, \frac{K v^* \circ K v^*}{a} \circ z \right\rangle \leq \max_i \frac{(K v^*)_i^2}{a_i} \|z\|^2 \leq \frac{\|K\|_\infty^2 (v^*)_j^2}{a_{\min}} \frac{1}{\sigma_{\min}(K)^2}.$$

This leads to the asserted lower bound $\vartheta^2 \geq \delta_1$.

When $m \leq n$, we can simply interchange the roles of K and K^\top , a and b , as well as u^* and v^* in this proof to obtain $\vartheta^2 \geq \delta_2$. \square

Taken together, Theorem 3 and Theorem 4 result in Theorem 2.

4 Numerical illustration

We illustrate the effect of overrelaxation by two numerical experiments related to optimal transport. The first is motivated by an application to color transfer between images [6]. The matrix $K = K_\varepsilon$ is generated as

$$K_{ij} = \exp\left(-\frac{\|x_i - y_j\|^2}{\varepsilon}\right)$$

where $x_i, y_j \in \mathbb{R}^3$ are RGB values (scaled to $[0, 1]$) of $m = n = 1000$ randomly sampled pixels in two different color images, respectively.¹ The vectors a and b are chosen as uniform distributions, i.e. $a = \mathbf{1}_m/m$ and $b = \mathbf{1}_n/n$. We choose $\varepsilon = 0.01$. In this scenario the standard Sinkhorn method is reasonably fast, but still can be accelerated using overrelaxation. A typical outcome for different relaxation strategies is shown in Fig. 1 left, where we plot for 500 iterations the ℓ_1 -distance $\|P_\ell - P_*\|_1$ between the matrices $P_\ell = \text{diag}(u_\ell) K \text{diag}(v_\ell)$ and a numerical reference solution $P_* = \text{diag}(u^*) K \text{diag}(v^*)$. This error corresponds to the total variation distance of the corresponding transport plan. Even if this quantity (specifically P_*) is not available in a practical computation it is a natural measure for the convergence of the method. Besides the standard Sinkhorn method ($\omega = 1$), we run the method with a fixed relaxation $\omega = 1.5$, and with the ‘optimal’ relaxation ω^{opt} , which is computed via formula (5) from the second largest singular value ϑ of matrix $\text{diag}(1/a^{1/2}) P_* \text{diag}(1/b^{1/2})$ (then ϑ^2 is the second largest eigenvalue of (3)). We do not consider relaxation based on the lower bound on ϑ^2 in Theorem 4, since the resulting ω is too close to one. In all variants of the algorithm the same (uniformly) random starting vectors u_0 and v_0 are used.

¹The setup follows the *OT for image color adaptation* example from the Python Optimal Transport toolbox [7]. The used images `ocean_day.jpg` and `ocean_sunset.jpg` are contained in the toolbox.

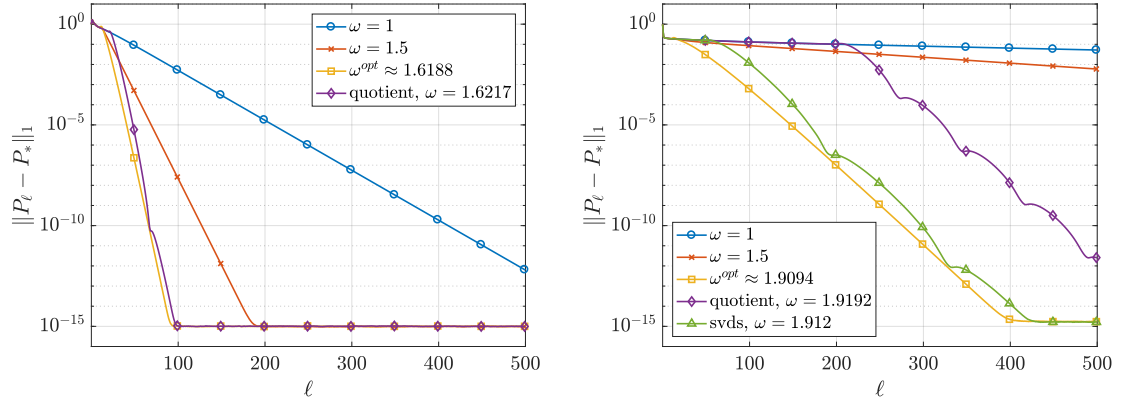


Figure 1: Effect of different relaxation strategies in two examples.

As can be seen, using ω^{opt} significantly accelerates the convergence speed. Moreover, although ω^{opt} only provides the optimal local rate, the positive effect shows quite immediately. However, the value of ω^{opt} is a priori unknown in practice. Therefore we also tested a simple heuristic, similar to one suggested in [18]. It is known that the convergence of the Sinkhorn method can be monitored, e.g., through the error $\|P_\ell \mathbf{1}_n - a\|_1$; cf. [15, Remark 4.14]. Therefore, since ϑ^2 equals the asymptotic convergence rate of the standard Sinkhorn method, we may take

$$\hat{\vartheta}^2 = \sqrt{\frac{\|P_\ell \mathbf{1}_n - a\|_1}{\|P_{\ell-2} \mathbf{1}_n - a\|_1}}$$

as a current approximation for ϑ^2 . In the purple curve (diamond markers) in Fig. 1 left, we updated ω a single time after 20 steps of the standard method based on this quantity, and using formula (5). This comes at almost no additional cost, but yields the near optimal rate in this example. Of course such a heuristic could be applied in a more systematic way, e.g., by monitoring the changes of $\left(\frac{\|P_\ell \mathbf{1}_n - a\|_1}{\|P_{\ell-p} \mathbf{1}_n - a\|_1}\right)^{1/p}$ for a suitable value of p over several iterations. We note that adapting ω in (linear and nonlinear) SOR methods based on currently observed convergence rates is a classical idea and has been proposed, e.g., in [19] or [10].

As a second example we consider a 1D transport problem between two random measures a and b (generated from a uniform distribution) on an equidistant grid in $[0, 1]$, and with ℓ_1 -norm as a cost. The matrix K in this case is given as

$$K_{ij} = \exp\left(-\frac{\left|\frac{i}{m-1} - \frac{j}{n-1}\right|}{\varepsilon}\right).$$

Again we choose $m = n = 1000$ and $\varepsilon = 0.01$, and then compare different relaxation strategies, but starting from the same random initialization (u_0, v_0) . As can be seen in Fig. 1 right, which shows 500 iterations with different relaxation strategies, this problem seems to be more difficult and the standard Sinkhorn method is extremely slow. A suitable relaxation compensates this and restores fast convergence, however, as illustrated by the slow convergence of the curve for $\omega = 1.5$, the estimation of ω^{opt} , and hence of ϑ^2 , needs to be rather precise. Since here the convergence rate of the standard method stabilizes later, we apply the above heuristic of estimating ϑ^2 only after 200 iterations of the standard iteration, resulting in the purple curve (diamond markers). The oscillatory behavior occurs because ω is estimated larger than ω^{opt} , in

which case the spectral radius $\omega - 1$ of the linearized iteration matrix M_ω in (9) is achieved at complex eigenvalues. It is possible in this example to update ω earlier using computationally more expensive heuristics. For instance, the green curve (triangle markers) is obtained by computing after 50 iterations of the standard method an approximation of ϑ as the second largest singular value of the matrix $\text{diag}(1/a_\ell^{1/2})P_\ell \text{diag}(1/b_\ell^{1/2})$, where $a_\ell = u_\ell \circ K v_\ell$ and $b_\ell = v_\ell \circ K^\top u_\ell$. This could be done iteratively, we used the Matlab function `svds`. This results in an almost optimal convergence rate in this example. Of course, several similar strategies could be devised.

References

- [1] C. Barcelo-Vidal and J.-A. Martin-Fernandez. The mathematics of compositional analysis. *Aust. J. Stat.*, 45(4):57–71, 2016.
- [2] C. Brauer, C. Clason, D. Lorenz, and B. Wirth. A Sinkhorn-Newton method for entropic optimal transport. *arXiv 1710.06635*, 2017.
- [3] C. Cotar, G. Friesecke, and C. Klüppelberg. Density functional theory and optimal transportation with Coulomb cost. *Comm. Pure Appl. Math.*, 66(4):548–599, 2013.
- [4] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In C. J. C. Burges et al., editors, *Advances in Neural Information Processing Systems 26*, pages 2292–2300. Curran Associates, Inc., 2013.
- [5] S. P. Eveson and R. D. Nussbaum. An elementary proof of the Birkhoff-Hopf theorem. *Math. Proc. Cambridge Philos. Soc.*, 117(1):31–55, 1995.
- [6] S. Ferradans, N. Papadakis, G. Peyré, and J.-F. Aujol. Regularized discrete optimal transport. *SIAM J. Imaging Sci.*, 7(3):1853–1882, 2014.
- [7] R. Flamary et al. POT: Python Optimal Transport. *J. Mach. Learn. Res.*, 22(78):1–8, 2021. Website: <https://pythonot.github.io/>
- [8] J. Franklin and J. Lorenz. On the scaling of multidimensional matrices. *Linear Algebra Appl.*, 114/115:717–735, 1989.
- [9] W. Hackbusch. *Iterative solution of large sparse systems of equations*. Springer, [Cham], second edition, 2016.
- [10] L. A. Hageman and T. A. Porsching. Aspects of nonlinear block successive overrelaxation. *SIAM J. Numer. Anal.*, 12:316–335, 1975.
- [11] P. A. Knight. The Sinkhorn-Knopp algorithm: convergence and applications. *SIAM J. Matrix Anal. Appl.*, 30(1):261–275, 2008.
- [12] P. A. Knight and D. Ruiz. A fast algorithm for matrix balancing. *IMA J. Numer. Anal.*, 33(3):1029–1047, 2013.
- [13] V. Pawlowsky-Glahn, J. J. Egozcue, and R. Tolosana-Delgado. *Modeling and analysis of compositional data*. John Wiley & Sons, Ltd., Chichester, 2015.
- [14] G. Peyré, L. Chizat, F.-X. Vialard, and J. Solomon. Quantum entropic regularization of matrix-valued optimal transport. *European J. Appl. Math.*, 30(6):1079–1102, 2019.
- [15] G. Peyré and M. Cuturi. Computational optimal transport. *Found. Trends Mach. Learn.*, 11(5-6):355–607, 2019.
- [16] R. Sinkhorn. Diagonal equivalence to matrices with prescribed row and column sums. *Amer. Math. Monthly*, 74:402–405, 1967.
- [17] R. Sinkhorn and P. Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific J. Math.*, 21:343–348, 1967.
- [18] A. Thibault, L. Chizat, C. Dossal, and N. Papadakis. Overrelaxed Sinkhorn-Knopp algorithm for regularized optimal transport. *Algorithms (Basel)*, 14(5), 143, 2021.
- [19] D. M. Young. *Iterative solution of large linear systems*. Academic Press, New York-London, 1971.