Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

ADAM NIELSEN AND MARCUS WEBER

Computing the nearest reversible Markov chain

ZIB Report 14-48 (December 2014)

Herausgegeben vom
Konrad-Zuse-Zentrum für Informationstechnik Berlin
Takustraße 7
D-14195 Berlin-Dahlem

Telefon: 030-84185-0
Telefax: 030-84185-125

e-mail: bibliothek@zib.de
URL: http://www.zib.de

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

Computing the nearest reversible Markov chain

Adam Nielsen*and Marcus Weber†

Abstract

Reversible Markov chains are the basis of many applications. However, computing transition probabilities by a finite sampling of a Markov chain can lead to truncation errors. Even if the original Markov chain is reversible, the approximated Markov chain might be non-reversible and will lose important properties, like the real valued spectrum. In this paper, we show how to find the closest reversible Markov chain to a given transition matrix. It turns out that this matrix can be computed by solving a convex minimization problem.

1 Introduction

A Markov chain with $n \in \mathbb{N}$ states is described through a transition matrix $P \in \mathbb{R}^{n\times n}$, i.e.

$$\sum_{j=1}^{n} P_{ij} = 1 \quad \text{and} \quad P_{ij} \geq 0$$

for $i, j = 1, \ldots, n$. A vector $\pi \in \mathbb{R}^n$ is called a probability distribution if

$$\sum_{i=1}^{n} \pi_i = 1 \quad \text{and} \quad \pi_j \geq 0$$

for $j = 1, \ldots, n$ holds. A Markov chain or its corresponding transition matrix P is called reversible according to a probability distribution π if the following equation

$$DP = P^T D$$

is valid for the diagonal matrix $D = \text{diag}(\pi_1, \ldots, \pi_n)$. In this case, π is a stationary distribution of the Markov chain, i.e.

$$\pi^T P = \pi^T.$$

The main result of this article is that for any transition Matrix P, any probability distribution π, and any norm $||\cdot||$ on $\mathbb{R}^{n\times n}$ which is induced by a scalar product, there exists a unique transition matrix $\widetilde{P}$ which is reversible according to π and has minimal distance to P with respect to the norm $||\cdot||$.

This article is structured as follows. First, we will prove the above conjecture. Then, we show how to numerically obtain the reversible matrix $\widetilde{P}$, and discuss the computational cost and include a perturbation analysis. Finally, we will give an application and a numerical example.

*nielsen@zib.de
†weber@zib.de

2 Existence Proof

Given a stochastic matrix P and a stationary distribution π, it will be proven in the following that there exists a unique transition matrix $\widetilde{P}$ which minimizes $||\widetilde{P} - P||$ and is reversible according to π.

First, let us consider

$$U = \{A \in \mathbb{R}^{n\times n} \mid DA = A^T D \text{ and } \exists k \in \mathbb{R} \text{ with } \sum_{j=1}^{n} A_{ij} = k \text{ for all } i = 1, \dots, n\}$$

where $D = \text{diag}(\pi_1, \dots, \pi_n)$ denotes the diagonal matrix with values π_i on the diagonal. Notice that U is a subspace of $\mathbb{R}^{n\times n}$ because for $A, B \in U$ with $\sum_{j=1}^{n} a_{1j} = k_1$ and $\sum_{i=1}^{n} b_{1j} = k_2$ we obtain for any $\alpha, \beta \in \mathbb{R}$ that

$$D(\alpha A + \beta B) = \alpha DA + \beta DB = \alpha A^T D + \beta B^T D = (\alpha A + \beta B)^T D$$

and

$$\sum_{j=1}^{n} \alpha a_{ij} + \beta b_{ij} = \alpha k_1 + \beta k_2 \quad \text{for all } i = 1, \dots, n.$$

We can show the following property:

Lemma 2.1. *For $A \in U$ with $\sum_{j=1}^{n} A_{ij} = k$ it holds $\pi^T A = k\pi^T$.*

Proof. Since $A \in U$, we have $\pi_i A_{ij} = \pi_j A_{ji}$. Therefore, we obtain

$$(\pi^T A)_i = \sum_{l=1}^{n} \pi_l A_{li} = \pi_i \sum_{l=1}^{n} A_{il} = \pi_i k.$$

□

For an understanding of the space U, the following matrices are of interest:

$$A_{ij}^{[r,s]} = \begin{cases} \pi_s & \text{if } i = r \text{ and } j = s, \\ \pi_r & \text{if } i = s \text{ and } j = r, \\ 1 - \pi_s & \text{if } i = j = r, \\ 1 - \pi_r & \text{if } i = j = s, \\ 1 & \text{if } i = j \text{ and } s \neq i \neq r, \\ 0 & \text{else,} \end{cases}$$

and

$$\delta_{ij}^{[r,s]} = \begin{cases} 1 & \text{if } i = j \text{ and } i \neq r, \\ 1 & \text{if } i = r \text{ and } j = s, \\ 0 & \text{else.} \end{cases}$$

Proposition 2.1. *The set*

$$\{A^{[r,s]} \mid (r,s) \in I_A\} \cup \{\delta^{[r,s]}, \delta^{[s,r]} \mid (r,s) \in I_B\} \cup \{\text{Id}\}$$

is a basis of U *where* Id *is the identity matrix and*

$$\begin{aligned} I &= \{(r,s) \mid 1 \leq r < s \leq n\}, \\ A &= \{i \colon \pi_i \neq 0\}, \\ B &= \{i \colon \pi_i = 0\}, \\ I_A &= \{(r,s) \in I \mid r \in A \text{ or } s \in A\}, \\ I_B &= \{(r,s) \in I \mid r,s \in B\}. \end{aligned}$$

The dimension of U *is given by*

$$dim(U) = \binom{n}{2} + 1 + \binom{|B|}{2}.$$

Proof. One may notice that the row sum of Id, $A^{[r,s]}$ and $\delta^{[r,s]}$ is always one and that

$$DA^{[r,s]} = \left(A^{[r,s]}\right)^T D$$

holds. The latter can be seen by

$$\begin{aligned} \left(DA^{[r,s]}\right)_{ij} &= \sum_{k=1}^{n} D_{ik} A^{[r,s]}_{kj} \\ &= \pi_i A^{[r,s]}_{ij} \\ &\overset{(*)}{=} \pi_j A^{[r,s]}_{ji} \\ &= \sum_{k=1}^{n} A^{[r,s]}_{ki} D_{kj} \\ &= \left((A^{[r,s]})^T D\right)_{ij}, \end{aligned}$$

where $(*)$ holds because for $i = r$ and $j = s$ or vice versa we have

$$\pi_i A^{[r,s]}_{ij} = \pi_r \pi_s = \pi_s \pi_r = \pi_j A^{[r,s]}_{ji},$$

for $i = j$ the equation in $(*)$ is trivial and in all other cases we have $A^{[r,s]}_{ij} = A^{[r,s]}_{ji} = 0$, and therefore $A^{[r,s]} \in U$. Also, we have

$$D\delta^{[r,s]} = {\delta^{[r,s]}}^T D$$

for all $(r,s) \in I_B$. This is because $\pi_r = 0$ and, therefore,

$$D\delta^{[r,s]}_{i,j} = \begin{cases} \pi_i & \text{if } i = j, \\ 0 & \text{else.} \end{cases}$$

Since $D\delta^{[r,s]}$ is just a diagonal matrix, it is in particular symmetric and we obtain

$$D\delta^{[r,s]} = (D\delta^{[r,s]})^T = {\delta^{[r,s]}}^T D,$$

the argument is analogous for $\delta^{[s,r]}$. To prove that these matrices are indeed a basis of U, it remains to show that they are linearly independent and that they span the subspace U.

To see the linearly independence, consider now an arbitrary linear combination of zero:

$$\sum_{(r,s)\in I_A} \alpha_{r,s} A^{[r,s]} + \sum_{(r,s)\in I_B} \alpha_{r,s}\delta^{[r,s]} + \beta_{r,s}\delta^{[s,r]} + \alpha I = 0.$$

For $(r,s) \in I_A$ the matrix $A^{[r,s]}$ is the only matrix in the above linear combination that could have a non-zero entry in row r and column s and in row s and column r. Therefore, we obtain

$$0 = \alpha_{r,s}\pi_s \quad \text{and} \quad 0 = \alpha_{r,s}\pi_r.$$

Thus, $\pi_s \neq 0$ or $\pi_r \neq 0$ provides $\alpha_{r,s} = 0$. For $(r,s) \in I_B$, we obtain analogously

$$\alpha_{r,s} = 0 \quad \text{and} \quad \beta_{r,s} = 0.$$

The linear combination reduces to $\alpha \operatorname{Id} = 0$ which, finally, leads us to $\alpha = 0$.

It remains to show that the given matrices span the subspace U. Thus, let us consider a matrix $C \in U$ with $\sum_{j=1}^{n} C_{ij} = k$ for some $k \in \mathbb{R}$. For $(r,s) \in I_A$ define $\alpha_{r,s} := \frac{C_{sr}}{\pi_r}$ if $\pi_r \neq 0$ and otherwise $\alpha_{r,s} := \frac{C_{rs}}{\pi_s}$. From $\pi_r C_{rs} = \pi_s C_{sr}$ we get that

$$\alpha_{r,s} A^{[r,s]}_{rs} = C_{rs} \quad \text{and} \quad \alpha_{r,s} A^{[r,s]}_{sr} = C_{sr}.$$

For $(r,s) \in I_B$ choose

$$\alpha_{r,s} = C_{r,s} \quad \text{and} \quad \beta_{r,s} = C_{s,r}.$$

Since each off-diagonal element appears in exactly one matrix, C differs from

$$\widetilde{C} := \sum_{(r,s)\in I_A} \alpha_{r,s} A^{[r,s]} + \sum_{(r,s)\in I_B} \alpha_{r,s}\delta^{[r,s]} + \beta_{r,s}\delta^{[s,r]}$$

only in the diagonal. We also know that there exists $l \in \mathbb{R}$ with $\sum_{j=1}^{n} \widetilde{C}_{ij} = l$ for $i = 1, \ldots, n$ since $\widetilde{C} \in U$. Therefore, the matrix $\hat{C} := \widetilde{C} + (k-l)\operatorname{Id}$ has row-sum k and

$$\hat{C}_{ii} = k - \sum_{j=1, j\neq i}^{n} \hat{C}_{ij} = k - \sum_{j=1, j\neq i}^{n} C_{ij} = C_{ii},$$

holds. Therefore,

$$C = \hat{C} \in U,$$

which shows that the given matrices span U. Furthermore,

$$\dim U = |I_A| + |I_B| + |I_B| + 1 = |I| + |I_B| + 1.$$

The statement follows from

$$|I| = \binom{n}{2}$$

and

$$|I_B| = \binom{|B|}{2}.$$

□

Our goal is to find a matrix

$$\widetilde{P} \in X = \{A \in U \mid A_{ij} \geq 0 \text{ for } i,j = 1,\dots,n \text{ and } \sum_{j=1}^{n} a_{1j} = 1\}$$

such that:

$$||\widetilde{P} - P|| \leq ||A - P|| \quad \text{for all } A \in X. \tag{1}$$

To do so, let us characterize the set X. To simplify notation, let us denote the basis of Proposition 2.1 by $(v_i)_{i=1,\dots,m}$ with $m = \dim U$ and $v_m = \mathrm{Id}$.

For any $A \in X$ we have a unique coefficent vector $\mathbf{x} \in R^m$ with $A = \sum_{i=1}^{m} x_i v_i$. From $\sum_{j=1}^{n} a_{ij} = 1$ we get that

$$\begin{aligned} 1 = \sum_{j=1}^{n} a_{ij} &= \sum_{j=1}^{n} \left(\sum_{l=1}^{m} x_l v_l(i,j) \right) \\ &= \sum_{l=1}^{m} x_l \left(\sum_{j=1}^{n} v_l(i,j) \right) \\ &= \sum_{l=1}^{m} x_l. \end{aligned}$$

This can be rewritten in $\mathbb{1}^T \mathbf{x} = 1$ where $\mathbb{1} \in \mathbb{R}^m$ is the constant vector $\mathbb{1}_i = 1$ for $i = 1,\dots,m$. Further, we have $a_{ij} \geq 0$ for all $i \neq j$ if and only if $x_l \geq 0$ for all $l = 1,\dots,m-1$. This can be rewritten as $-\mathbf{x}e_i \leq 0$ for $i = 1,\dots,m-1$. The diagonal entries of A can be positive even if x_m is negative. Let us quickly go back to the old notation to see how we have to define the restriction here. So let A be given by

$$A = \sum_{(r,s)\in I_A} \alpha_{r,s} A^{[r,s]} + \sum_{(r,s)\in I_B} \alpha_{r,s}\delta^{[r,s]} + \beta_{r,s}\delta^{[s,r]} + \alpha Id,$$

thus,

$$a_{ii} = \sum_{(r,s)\in I_A} \alpha_{r,s} A_{ii}^{[r,s]} + \sum_{(r,s)\in I_B} \alpha_{r,s}\delta_{ii}^{[r,s]} + \beta_{r,s}\delta_{ii}^{[s,r]} + \alpha I_{ii}.$$

Which leads to

$$a_{ii} = \sum_{\substack{(r,s)\in I_A \\ r=i}} \alpha_{r,s}(1-\pi_s) + \sum_{\substack{(r,s)\in I_A \\ s=i}} \alpha_{r,s}(1-\pi_r) + \sum_{\substack{(r,s)\in I_A \\ r\neq i\neq s}} \alpha_{r,s} + \sum_{\substack{(r,s)\in I_B \\ r\neq i}} \alpha_{r,s} + \sum_{\substack{(r,s)\in I_B \\ s\neq i}} \beta_{r,s} + \alpha.$$

The condition $a_{ii} \geq 0$ is equivalent to $-\mathbf{x}g_i \leq 0$ where

$$g_i(j) = \begin{cases} 1-\pi_s & \text{if } v_j = A^{[i,s]} \text{ for some } s > i, \\ 1-\pi_r & \text{if } v_j = A^{[r,i]} \text{ for some } r < i, \\ 0 & \text{if } v_j = \delta^{[i,s]} \text{ for some } s, \\ 1 & \text{else,} \end{cases}$$

for $i = 1, \ldots, n$ and $j = 1, \ldots, m$. Given the matrix

$$C = -1 \cdot \begin{bmatrix} — & e_1^T & — \\ & \vdots & \\ — & e_{m-1}^T & — \\ — & g_1^T & — \\ & \vdots & \\ — & g_n^T & — \end{bmatrix} \in \mathbb{R}^{(n+m-1)\times m},$$

the condition that $A = \sum_{i=1}^m x_i v_i$ is in set X is equivalent to

$$Cx \leq 0 \quad \text{and} \quad \mathbb{1}^T x = 1.$$

In order to find a matrix $\widetilde{P} \in X$ which satisfies inequality (1), let us recall that the norm $||\cdot||$ is induced by a scalar product $\langle \cdot, \cdot \rangle$, i.e.

$$||A|| = \sqrt{\langle A, A \rangle}$$

for any matrix $A \in \mathbb{R}^{n\times n}$. We want to minimize the term

$$\begin{aligned} \left\| \sum_{i=1}^m x_i v_i - P \right\|^2 &= \sum_{i,j=1}^m x_i x_j \langle v_i, v_j \rangle - 2 \sum_{i=1}^m x_i \langle v_i, P \rangle + \langle P, P \rangle \\ &= \frac{1}{2} x^T Q x + x^T f + c \end{aligned} \tag{2}$$

with

$$Cx \leq 0 \quad \text{and} \quad \mathbb{1}^T x = 1,$$

where

$$Q(i,j) := 2 \langle v_i, v_j \rangle, \quad f(i) = -2 \langle v_i, P \rangle$$

and

$$c = \langle P, P \rangle .$$

Since Q is a Gram matrix of linear independent vectors, it is positive definite. This follows because for any $x \in \mathbb{R}^n$ it holds

$$x^T Q x = \sum_{i,j} x_i Q_{ij} x_j = \left\langle \sum_{i=1}^n x_i v_i, \sum_{i=1}^n x_i v_i \right\rangle .$$

Since $v_1, \ldots, v_n$ is a basis, we have that $\sum_{i=1}^n x_i v_i \neq 0$ for $x \neq 0$ and, in consequence, $x^T Q x > 0$ for $x \neq 0$. Since Q is positive definite, the quadratic function (2) is strongly convex. Therefore, we have formulated the problem into a strongly convex quadratic programming problem that attains its global minimum, since the quadratic function is coercive, continuous and the set X is non-empty because of $\mathrm{Id} \in X$, see [9] Theorem 1.15. Also, the global minimum is unique because the quadratic function is strongly convex.

Thus, we have proven the existence of a unique matrix $\widetilde{P} \in X$ which fulfills inequality (1). We will now discuss how to compute this matrix.

3 Complexity and Perturbation

To avoid technical difficulties, we will assume in this chapter that $\pi_i > 0$ for $i = 1, \ldots, n$. For $A \in \mathbb{R}^{n\times n}$ the trace $tr(A)$ is given by the sum of the diagonal elements

$$tr(A) = \sum_{i=1}^{n} A_{ii}.$$

It is known that

$$\langle A, A\rangle_F := tr(A^T B)$$

is a scalar product on $\mathbb{R}^{n\times n}$ for $A, B \in \mathbb{R}^{n\times n}$. This scalar product induces the Frobenius norm

$$||A||_F = \sqrt{\langle A, A\rangle_F} = \sqrt{\sum_{i=1}^{n}\sum_{j=1}^{n} |a_{ij}|^2}.$$

The following complexity analysis will be given according to the Frobenius norm.

3.1 Complexity

Unfortunately, the matrix Q is quite large, in particular we have $Q \in \mathbb{R}^{m\times m}$ where

$$m = \dim(U) = 1 + \frac{n^2}{2}.$$

Each entry of Q is given by a trace of two sparse matrices which can be computed by the following formula

$$\left\langle A^{[r,s]}, A^{[r',s']}\right\rangle_F = \begin{cases} n - \pi_r - \pi_{r'} - \pi_s - \pi_{s'} & \text{if } r, r', s, s' \text{ are distinct,} \\ n - 1 - 2\pi_s + (1-\pi_r)(1-\pi_{r'}) & \text{if } r \neq r', s = s', \\ n - 1 - 2\pi_r + (1-\pi_s)(1-\pi_{s'}) & \text{if } r = r', s \neq s, \\ n - 1 + (1-\pi_s)(1-\pi_{s'}) + (1-\pi_r)(1-\pi_{r'}) & \text{if } r = s' \text{ or } r' = s, \\ n - 2 + (1-\pi_r)^2 + \pi_r^2 + (1-\pi_s)^2 + \pi_s^2 & \text{if } r = r' \text{ and } s = s'. \end{cases}.$$

Note that because of $r < s$ and $r' < s'$ other cases are not possible. We only give a proof for the case $r \neq r'$ and $s = s'$, all other cases can be obtained analogously. First, if $A_{ij}^{[r,s]} \neq 0$ for $i \neq j$ it follows that either $i = r, j = s$ or $i = s, j = r$ holds. Since $r \neq r', s = s'$ we have $A_{ij}^{[r',s']} = 0$, therefore,

$$\begin{aligned} \left\langle A^{[r,s]}, A^{[r',s']}\right\rangle_F &= \sum_{i=1}^{n} A_{ii}^{[r,s]} A_{ii}^{[r',s']} \\ &= n - 3 + A_{rr}^{[r,s]} A_{rr}^{[r',s']} + A_{r'r'}^{[r,s]} A_{r'r'}^{[r',s']} + A_{ss}^{[r,s]} A_{ss}^{[r',s']} \\ &= n - 3 + 1 - \pi_s + 1 - \pi_{s'} + (1-\pi_r)(1-\pi_{r'}) \\ &= n - 1 - 2\pi_s + (1-\pi_r)(1-\pi_{r'}). \end{aligned}$$

Therefore, the computational cost to evaluate a single entry of Q does not increase with n and the effort to compute and store Q is $O(n^4)$.

The convex minimization problem can be solved with a barrier method, e.g. the interior point method. The method consist of N Newton iterations.

The number N of Newton iterations to find a strictly feasible point is bounded by

$$N \leq \sqrt{n + n^2/2 - 1} \log \left(\frac{n + n^2/2 - 1}{\varepsilon} \right) \gamma$$

where $\varepsilon > 0$ is the demanded accuracy and γ is a constant depending on the choice of backtracking parameter, see [1], section 11.5.5. Therefore, the amount of Newton iterations is bounded by $O(n \log n)$. Unfortunately, each Newton iteration has to solve a linear equation system and cost $O((n^2)^3)$ since $Q \in \mathbb{R}^{(\frac{n^2}{2}+1)\times(\frac{n^2}{2}+1)}$. Therefore, the total cost for the optimization problem is bounded by $O(n^7 \log n)$.

In the field of convex optimization, the upper bound for the amount of Newton iterations is known to be a large overestimation [1]. If we assume this amount to be independent from n, then we can expect that the time to find a solution with the convex optimization problem should be given by

$$g(n) = \alpha n^6$$

where n is the size of the matrix $P \in \mathbb{R}^{n \times n}$. If we include the bad estimation for the Newton iteration, then the time should be represented by

$$f(n) = \beta n^7 \log(n).$$

In order to explore the computation time of the convex optimization problem, we generated for each $n = 10, \ldots, 100$ a stochastic matrix $P \in \mathbb{R}^{n \times n}$ and a stochastic vector $\pi \in \mathbb{R}^n$ where each entry was drawn from the standard uniform distribution on the open interval (0,1) and then we normalized P and π. We used Matlab R2012b on a 3 gigahertz computer with 8 gigabyte memory. We solved the convex optimization problem with the interior-point-convex algorithm from the provided Matlab method quadprog with default options, i.e. relative dual feasibility $= 2.31e-15$ with TolFun$=1e-0.8$, complementarity measure $= 1.68e-10$ with TolFun$=1e-0.8$ and relative max constraint violoation $= 0$ with TolCon$=1e-0.8$. For $n = 45$ the execution time was about 12.46 seconds. The scalars α, β where chosen such that $f(45) = g(45) = \frac{12.46}{60}$ holds. In Figure 1 one can see that g seems to be a reasonable approximation of the execution time. Also one can see that the execution time takes only seconds for matrices in $\mathbb{R}^{50\times50}$, but already 32 minutes for matrices in $\mathbb{R}^{100\times100}$.

3.2 Perturbation analysis

For

$$k := \min_{r,r',s,s'} \left\langle A^{[r,s]}, A^{[r',s']} \right\rangle_F$$

we obtain

$$k \operatorname{Id} \leq Q \leq n \operatorname{Id}$$

where the inequality has to be read componentwise. The upper bound follows from

$$\langle v_i, v_j \rangle_F \leq \sqrt{\langle v_i, v_i \rangle_F} \sqrt{\langle v_j, v_j \rangle_F} \leq \max_{l=1,\ldots,m} \{\langle v_l, v_l \rangle_F\}$$

and the fact that

$$\left\langle A^{[r,s]}, A^{[r,s]} \right\rangle_F = n - 2 + (1 - \pi_r)^2 + \pi_r^2 + (1 - \pi_s)^2 + \pi_s^2 \leq n$$

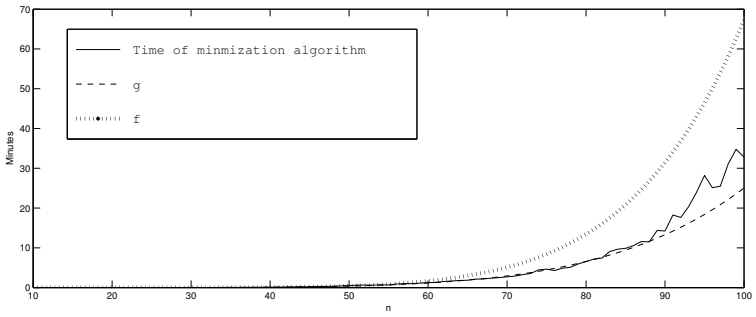

Figure 1: Duration of convex minimization problem to find closest reversible Markov chain.

and

$$\langle \mathrm{Id}, \mathrm{Id} \rangle_F = n$$

holds. Since Q is symmetric, its condition number according to the spectral norm is given by

$$\kappa(Q) = \frac{\lambda_{\max}}{\lambda_{\min}} \leq \frac{n}{k}.$$

Since $k = n - c$ for some $0 < c < 4$ we have that

$$\kappa(Q) \leq \frac{1}{1 - c/n} \to 1$$

for $n \to \infty$. Therefore, Q is well conditioned.

We assume now that we are interested in finding a reversible matrix $\hat{P} \in X$, but that we only have a perturbed version $P = \hat{P} + E$ of $\hat{P}$ which is not necessarily reversible. With the above scheme, we can find a reversible matrix $\widetilde{P} \in X$ which is closest to P according to the Frobenius norm. The question arises, how eigenvalues change between $\widetilde{P}$ and $\hat{P}$.

In order to answer this question, the weighted Frobenius norm

$$||A||_{\widetilde{F}} := ||D^{\frac{1}{2}} A D^{-\frac{1}{2}}||_F$$

is introduced with

$$D = diag(\pi_1, \ldots, \pi_n), \quad D^{\frac{1}{2}} D^{\frac{1}{2}} = D$$

and

$$D^{-\frac{1}{2}} := (D^{\frac{1}{2}})^{-1}.$$

Note that the weighted Frobenius norm is given by

$$||A||^2_{\widetilde{F}} = \sum_{i,j=1}^{n} a_{i,j}^2 \frac{\pi_i}{\pi_j}.$$

We can then give the following estimation about the eigenvalues:

Theorem 3.1. *Let $A, B \in X$, let $\lambda_1, \ldots, \lambda_n$ be the eigenvalues of A and $\hat{\lambda}_1, \ldots, \hat{\lambda}_n$ be the eigenvalues of B. Then there exists a permutation σ of the integers $1, 2, \ldots, n$ such that*

$$\sum_{i=1}^{n} |\hat{\lambda}_{\sigma(i)} - \lambda_i|^2 \leq ||A - B||_{\widetilde{F}}^2.$$

Proof. The matrices $A, B \in X$ are self-adjoint according to the scalar product $\langle x, y \rangle_\pi := x^T D y$. Lets denote with $\{w_1, \ldots, w_n\}$ a $\langle , \rangle_\pi$-orthonormal basis, and denote with W the matrix whose columns contain the vectors w_i, i.e.

$$W = \begin{pmatrix} | & | & & | \\ w_1 & w_2 & \ldots & w_n \\ | & | & & | \end{pmatrix}.$$

Then

$$A' := W^{-1} A W \quad \text{and } B' = W^{-1} B W$$

are symmetric, see [3] Chapter 5.6.1. By the Hoffman and Wielandt Theorem (see Theorem 6.3.5 in [5]) we obtain

$$\sum_{i=1}^{n} |\hat{\lambda}_{\sigma(i)} - \lambda_i|^2 \leq ||A' - B'||_F^2$$

for a permutation σ, because similar matrices have the same eigenvalues. It remains to show

$$||W^{-1} C W||_F^2 = ||C||_{\widetilde{F}}^2$$

or equivalently

$$||C||_F^2 = ||W C W^{-1}||_{\widetilde{F}}^2$$

for any matrix $C \in \mathbb{R}^n$. By construction of W we have

$$W^T D W = I \quad \text{and} \quad W^{-1} D^{-1} (W^T)^{-1} = I. \tag{3}$$

Therefore,

$$\begin{aligned} ||W C W^{-1}||_{\widetilde{F}}^2 &= ||D^{\frac{1}{2}} W C W^{-1} D^{-\frac{1}{2}}||_F^2 \\ &= \operatorname{tr}(D^{-\frac{1}{2}} (W^{-1})^T C^T W^T D^{\frac{1}{2}} D^{\frac{1}{2}} W C W^{-1} D^{-\frac{1}{2}}) \\ &\overset{(*)}{=} \operatorname{tr}(W^{-1} D^{-1} (W^{-1})^T C^T W^T D W C) \\ &\overset{(3)}{=} \operatorname{tr}(C^T C) \\ &= ||C||_F^2, \end{aligned}$$

where in $(*)$ it is used that the trace is invariant under cyclic permutations . □

This shows that in order to guarantee good approximations for the eigenvalues, one has to assure a good approximation of a_{ij} for those i, j where $\pi_j << \pi_i$. These transitions are also known as rare events and often difficult to compute.

We will end this chapter with a short experiment motivated by Theorem 3.1. We create 100 reversible Markov chains $(\hat{P}_i)_{i=1,\ldots,100} \subset \mathbb{R}^{5\times 5}$ as follows. First,

we generate a symmetric matrix A where each entry A_{ij} with $i \leq j$ is chosen equally distributed from $[0,1]$ and $A_{ij} := A_{ji}$ for $i > j$. After normalizing A the resulting matrix is reversible [7] and will be denoted as $\hat{P}_i$. We perturb and normalize $\hat{P}_i$ to obtain a perturbed version P_i of $\hat{P}_i$. For each P_i we compute the unique reversible matrix $\widetilde{P}_i$ which is closest to P_i according to weighted Frobenius norm and according to the stationary distribution of $\hat{P}_i$. The matrix $\widetilde{P}_i$ exists and can be computed by the introduced convex optimization problem, because the weighted Frobenius norm is induced by the scalar product

$$\langle A, B \rangle_{\widetilde{F}} = tr(DAD^{-1}B^T).$$

Also we compute the unique reversible matrix $\check{P}_i$ which is closest to P_i according to standard Frobenius norm and according to the stationary distribution of $\hat{P}_i$. For each tuple $(\hat{P}, \widetilde{P}, P, \check{P})$ we compute the corresponding eigenvalues $(\hat{\lambda}_j, \widetilde{\lambda}_j, \lambda_j, \check{\lambda}_j)_{j=1,\ldots,5}$. Then, we compute the numbers

$$c_1 = \min_{\sigma \in \Pi_5} \sqrt{\sum_{j=1}^{5} |\hat{\lambda}_j - \widetilde{\lambda}_{\sigma(j)}|^2}, \quad c_2 = \min_{\sigma \in \Pi_5} \sqrt{\sum_{j=1}^{5} |\hat{\lambda}_j - \lambda_{\sigma(j)}|^2}$$

and

$$c_3 = \min_{\sigma \in \Pi_5} \sqrt{\sum_{j=1}^{5} |\hat{\lambda}_j - \check{\lambda}_{\sigma(j)}|^2},$$

where Π_5 denotes the set of all permutations of the numbers $\{1, \ldots, 5\}$. The experiment is visualized in Figure 2. The result is that the closest reversible matrix $\check{P}$ according the Frobenius norm does not maintain the spectrum of $\hat{P}$. However, the closest reversible matrix according the weighted Frobenius norm $\widetilde{P}$ gives a good approximation for the spectrum of $\hat{P}$ and improves always the spectrum of the perturbed version P.

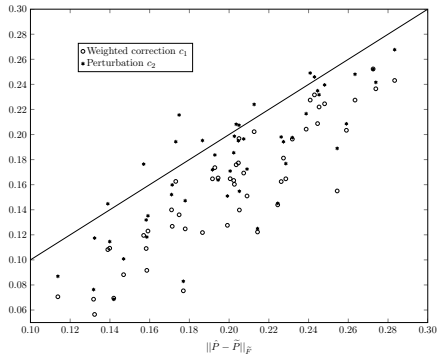

Figure 2: Comparison of c_1 and c_2.

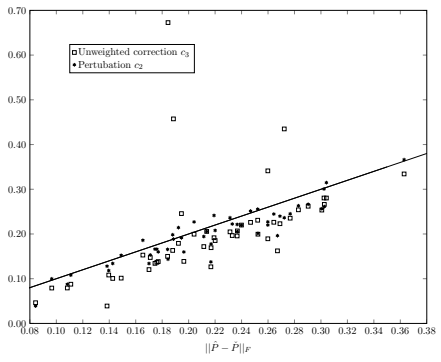

Figure 3: Comparison of c_2 and c_3.

Three remarks may be in order:

- The resulting matrix

$$\widetilde{P} = \sum_{(r,s) \in I} \alpha_{r,s} A^{[r,s]} + \alpha \, \mathrm{Id}$$

is reversible up to machine precision. In fact, it is reversible if $\pi_r \alpha_{r,s} \pi_s = \pi_s \alpha_{r,s} \pi_r$ holds[1].

- A numerical error in the solution $(x)_{i=1,\dots,m-1} = (\alpha_{r,s})_{(r,s)\in I}$ of the convex optimization problem results in an error of the entry $\widetilde{P}(r,s) = \alpha_{r,s}\pi_s$ and $\widetilde{P}(s,r) = \alpha_{r,s}\pi_r$ for $r \neq s$.
- In order to guarantee that $\widetilde{P}$ has row sum one, it might be advantages to set

$$\widetilde{p}_{ii} = 1 - \sum_{j \neq i} \widetilde{p}_{ij}.$$

4 Application and numerical example

In the past decades, analysis of a certain class of stochastic processes has been formulated in terms of approximating a transfer operator $\mathcal{T}\colon L^2(\mu) \to L^2(\mu)$ [12, 8, 11] which is self-adjoint according to the scalar product

$$\langle f, g \rangle_\mu = \int_E f(x)\, g(x)\, \mu(dx),$$

i.e.

$$\langle \mathcal{T} f, g \rangle_\mu = \langle f, \mathcal{T} g \rangle_\mu$$

for all $f, g \in L^2(\mu)$, where (E, Σ, μ) is a measure space for some set $E \subset \mathbb{R}^n$, and

$$L^2(\mu) = \{f\colon E \to \mathbb{R} \mid \langle f, f \rangle_\mu < \infty,\ f\mu\text{-measurable}\}/N$$

with

$$N = \{f\colon E \to \mathbb{R} \mid \exists A \in \Sigma \text{ with } \mu(A) = 0 \text{ and } f(x) = 0 \text{ for all } x \in E \backslash A\}.$$

Furthermore, the operator has the property that for $f \in L^2(\mu)$ with $f \geq 0$ it follows $\mathcal{T} f \geq 0$ almost surely.

Let us assume that we have a set of non-negative functions $\{\phi_1, \dots, \phi_n\}$ given, such that $\sum_{i=1}^n \phi_i(x) = 1$ for all $x \in E$ and $\langle \phi_i, \mathbb{1} \rangle_\mu > 0$ for $i = 1, \dots, n$ where $\mathbb{1}(x) := 1$ for all $x \in E$. If we define

$$\pi_i := \langle \phi_i, \mathbb{1} \rangle_\mu$$

and

$$T_{ij} := \frac{\langle \mathcal{T}\phi_i, \phi_j \rangle_\mu}{\langle \phi_i, \mathbb{1} \rangle_\mu},$$

then it is straightforward to verify the properties

(i) $\pi T = \pi$,

(ii) $DT = T^T D$,

(iii) $\sum_{j=1}^n T_{ij} = 1$ for all $i = 1, \dots, n$,

(iv) $T_{ij} \geq 0$,

[1] One may note that $I_A = I$ because of the assumption at the beginning of this chapter.

where $D = diag(\pi_1, \ldots, \pi_n)$ denotes the diagonal matrix of π. In other words, T is a reversible Markov chain. This matrix plays an essential role in the Galerkin discretization of $\mathcal{T}$.

Theorem 4.1. *Let $\{\phi_1, \ldots, \phi_n\} \subset L^2(\mu)$ be a basis with $\langle \phi_i, \mathbb{1} \rangle_\mu > 0$ of a subspace D, and $Q\colon L^2(\mu) \to D$ the orthogonal projection onto D. For any self-adjoint continuous operator $\mathcal{T}\colon L^2(\mu) \to L^2(\mu)$ we have*

$$\mathcal{M} = S^{-1}T, \quad T_{ij} = \frac{\langle \mathcal{T}\phi_i, \phi_j \rangle_\mu}{\langle \phi_i, \mathbb{1} \rangle_\mu}, \quad S_{ij} = \frac{\langle \phi_i, \phi_j \rangle_\mu}{\langle \phi_i, \mathbb{1} \rangle_\mu}$$

is a right matrix representation of $Q\mathcal{T}Q$ according to the basis $A = \{\phi_1, \ldots, \phi_n\}$, i.e. for any

$$f = \sum_{i=1}^n \alpha_i \phi_i, \quad Q\mathcal{T}Qf = \sum_{i=1}^n \beta_i \phi_i$$

it holds

$$\mathcal{M}(\alpha_1, \ldots, \alpha_n)^T = (\beta_1, \ldots, \beta_n)^T.$$

Proof. Consider the Gram matrix of $\{\phi_1, \ldots, \phi_n\}$

$$\hat{S}_{ij} = \langle \phi_i, \phi_j \rangle_\mu .$$

This matrix is invertible since $\{\phi_1, \ldots, \phi_n\}$ is a basis and the orthogonal projection Q can be represented as

$$Qv = \sum_{i,j=1}^n \hat{S}_{ij}^{-1} \langle v, \phi_i \rangle_\mu \phi_j.$$

This can be verified by checking $\langle Qv - v, g \rangle_\mu = 0$ for all $g \in D$, $v \in L^2(\mu)$. From

$$S = D^{-1}\hat{S} \quad \text{with} \quad D = diag\left(\langle \phi_1, \mathbb{1} \rangle_\mu, \ldots, \langle \phi_n, \mathbb{1} \rangle_\mu \right)$$

we obtain

$$S^{-1} = \hat{S}^{-1}D \quad \text{and, therefore,} \quad \hat{S}_{ij}^{-1} = S_{ij}^{-1} \frac{1}{\langle \phi_j, \mathbb{1} \rangle_\mu} = S_{ji}^{-1} \frac{1}{\langle \phi_i, \mathbb{1} \rangle_\mu},$$

in the last step it was used that $\hat{S}^{-1}$ is symmetric since $\hat{S}$ is symmetric. This implies

$$Qv = \sum_{i,j=1}^n S_{ji}^{-1} \langle v, \phi_i \rangle_\mu \frac{\phi_j}{\langle \phi_i, 1 \rangle_\mu}.$$

Therefore,

$$\begin{aligned} Q\mathcal{T}Q\phi_k = Q\mathcal{T}\phi_k &= \sum_{i,j=1}^n S_{ij}^{-1} \langle \mathcal{T}\phi_k, \phi_i \rangle_\mu \frac{\phi_j}{\langle \phi_i, 1 \rangle_\mu} \\ &= \sum_{i,j=1}^n S_{ij}^{-1} \frac{\langle T\phi_i, \phi_k \rangle_\mu}{\langle \phi_i, \mathbb{1} \rangle_\mu} \phi_j = \sum_{j=1}^n \left(\sum_{i=1}^n S_{ji}^{-1} T_{ik} \right) \phi_j \\ &= \sum_{j=1}^n \left(S^{-1}T \right)_{kj} \phi_j. \end{aligned}$$

□

For the more general case that the operator is not self-adjoint, one can find a similar Galerkin approximation, see [10], Theorem 3.

Note that S is a special case of the matrix T with $\mathcal{T} = I$, where I denotes the identity operator. Therefore, we can use our machinery to correct T and S if π is given.

One possible choice for the functions ϕ_i are set based functions, i.e. we have disjoint sets $A_1, \ldots, A_n$ with $\cup_{i=1}^n A_i = E$ and

$$\phi_i(x) = \mathbb{1}_{A_i}(x) = \begin{cases} 1 & \text{if } x \in A_i \\ 0 & \text{else.} \end{cases}$$

In this case, the matrix S from Theorem 4.1 turns out to be the identity matrix and we are only left with the task of computing T. Other possible choices are radial basis functions or commitor functions which can be found in [13, 10]

So far, we know how to maintain reversibility for S and T separately. The question arises, if $S^{-1}T$ also inherits some properties that might get lost due to numerical errors. First, we can rewrite

$$S = D^{-1}\hat{S}, \quad T = D^{-1}\hat{T}$$

with $\hat{S}_{ij} = \langle \phi_i, \phi_j \rangle_\mu$ and $\hat{T}_{ij} = \langle \mathcal{T}\phi_i, \phi_j \rangle_\mu$, hence

$$S^{-1}T = \hat{S}^{-1}DD^{-1}\hat{T} = \hat{S}^{-1}\hat{T}.$$

Analogously as shown for Q also $\hat{S}$ and $\hat{S}^{-1}$ are symmetric positive definite matrices. Now, since $\hat{S}^{-1}$ is positive definite , we know the existence of a symmetric square matrix A such that $A^2 = \hat{S}^{-1}$. Thus,

$$A^{-1}\hat{S}^{-1}\hat{T}A = A\hat{T}A.$$

Consequently, $\hat{S}^{-1}\hat{T}$ is similar to a symmetric matrix and hence diagonalizable. This shows that the spectrum of $\hat{S}^{-1}\hat{T}$ is real and that we know the existence of a basis of eigenvectors of $S^{-1}T$. For numerical estimation, the following procedure may be advantageous: First, approximate T^* and S^* by the convex minimization problem stated above. Then, calculate DS^* which is symmetric up to machine precision. For the symmetric matrix $(DS^*)^{-1}$ one can solve a convex minimization problem that finds the closest symmetric positive semi definite matrix $((DS^*)^{-1})^*$ according to the Frobenius norm, see [4]. Finally, $((DS^*)^{-1})^* T^*$ is the corrected matrix.

Also it would be possible to correct S^{-1} with our explained optimization scheme, because this matrix has also row sum one and is reversible, this follows from

$$S\mathbb{1} = \mathbb{1} \quad \Rightarrow \quad \mathbb{1} = S^{-1}\mathbb{1}$$

and

$$DS = S^T D \quad \Rightarrow \quad (S^T)^{-1}D = DS^{-1} \quad \Rightarrow \quad (S^{-1})^T D = DS^{-1}.$$

Therefore, by dropping the constraints given by the matrix C, one can also approximate S^{-1} as a convex quadratic programing problem with only linear constrains.

A Galerkin approximation of the transfer operator can be used to apply spectral clustering algorithms such as PCCA+ [2] which relies on the real spectrum of TS^{-1}. Especially for a set-based reduction, we can guarantee with our machinery that T will have a real spectrum, and, therefore, spectral clustering such as PCCA+ is applicable.

4.1 Numerical example

Consider the one dimensional, 2π-periodic function $V_B \colon \mathbb{R} \to \mathbb{R}$

$$V_B(x) = a + b\,\cos(x) + c\,\cos^2(x) + d\,\cos^3(x)$$

with a=2.0567, b=-4.0567, c=0.3133 and d=6.4267. This can be seen as an approximation of a butane potential energy function, where x is the central dihedral angle. We are going to realize a trajectory of the dihedral angles

$$X_t = \widetilde{X}_t \mod 2\pi$$

of butane from the stochastic differential equation

$$d\widetilde{X}_t = -\nabla V_B(\widetilde{X}_t)dt + \sigma dB_t.$$

If we divide $[0, 2\pi]$ in the sets $A_i = [\frac{i-1}{30}, \frac{i}{30}]$ for $i = 1, \ldots, 30$, then the Galerkin discretization of $\mathcal{T}$ reduces to

$$T_{ij} = \mathbb{P}(X_\tau \in A_j \mid X_0 \in A_i),$$

see [8].

To gain the associated Markov State Model $T \in \mathbb{R}^{30\times 30}$, we compute a long-term trajectory $(X_i)_{i=0,\ldots,n-1}$ by performing n timesteps of size dt using the Euler-Maruyama discretization

$$\widetilde{X}_{i+1} = \widetilde{X}_i - \nabla V_B(\widetilde{X}_i)dt + \sigma\sqrt{dt}\,\eta_i,$$

where η_i are i.i.d random variables distributed according to the standard normal distribution. This trajectory is chopped into pieces of length 400 yielding M subtrajectories $(X_i^k)_{i=1,\ldots,l} := (X_{lk}, \ldots, X_{l(k+1)-1})$ for $k = 0, \ldots, M-1$. We can estimate T by counting transitions

$$C_{ij} = \sum_{k=0}^{M-1} \mathbb{1}_{A_i}(X_1^k)\,\mathbb{1}_{A_j}(X_l^k)$$

and

$$T_{ij} = \frac{C_{ij}}{\sum_{i=1}^{30} C_{ij}}.$$

The Markov State Model created in this way becomes only reasonable when considered for a trajectory longer then 10^6 timesteps, due to rare transition events.

We will now compare the eigenvalues of the Markov State Model T to the eigenvalues of the above introduced corrected estimation $\widetilde{T}$ with respect to the weighted Frobenius norm $||\cdot||_{\widetilde{F}}$. We will also compare them for different length of a given trajectory, starting from length $n = 10^6$ until $n = 5 \cdot 10^8$. We used the timestep $dt = 0.001$. Since the eigenvalues of T turned out to be complex sometimes, we simply set the imaginary part to zero, in order to compare the eigenvalues with $\widetilde{T}$.

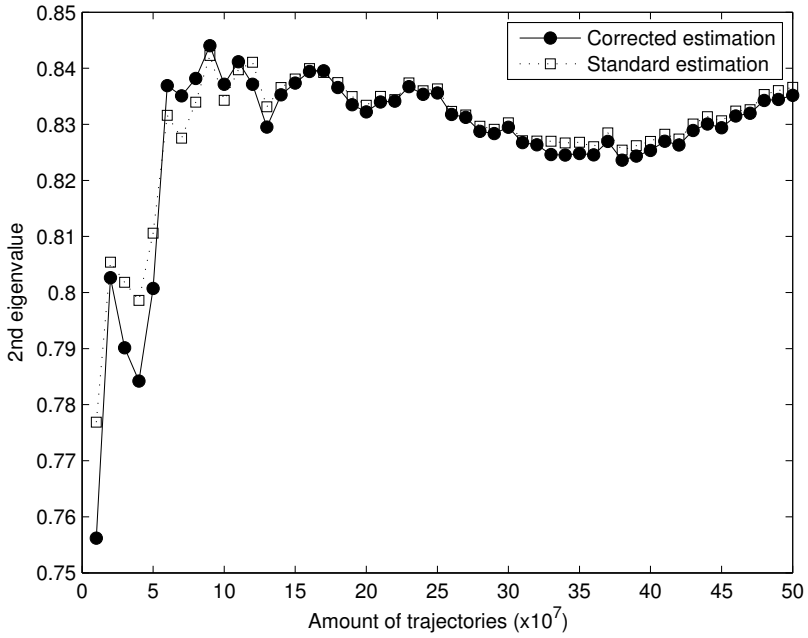

Discussion

In the figures the standard estimation and the corrected estimation do not essentially differ for the eigenvalues close to 1. Our algorithm of finding the nearest reversible Markov chain especially preserves the dominant eigenvalues of the underlying operator and, thus, is suitable for cluster analysis[2] of Molecular simulation data.

In contrast to T, the resulting transition matrix $\widetilde{T}$ from our algorithm is guaranteed to be applicable to spectral clustering methods, such as PCCA+, because of its real spectrum.

Acknowledgment We gratefully thank the Berlin Mathematical School for financial support of Adam Nielsen.

[2]For the connection between cluster analysis and eigenvalues we refer to [6].

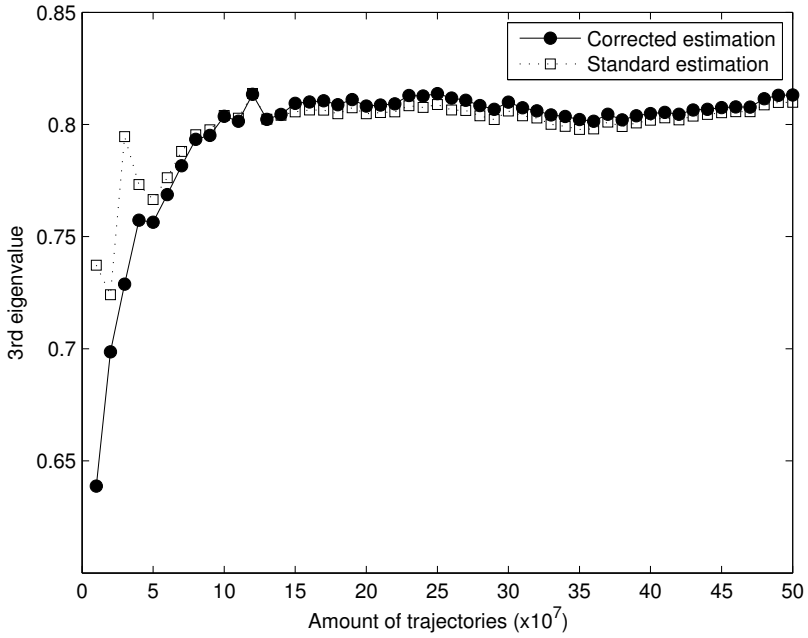
Corrected estimation
Standard estimation
3rd eigenvalue
Amount of trajectories (x10^7)

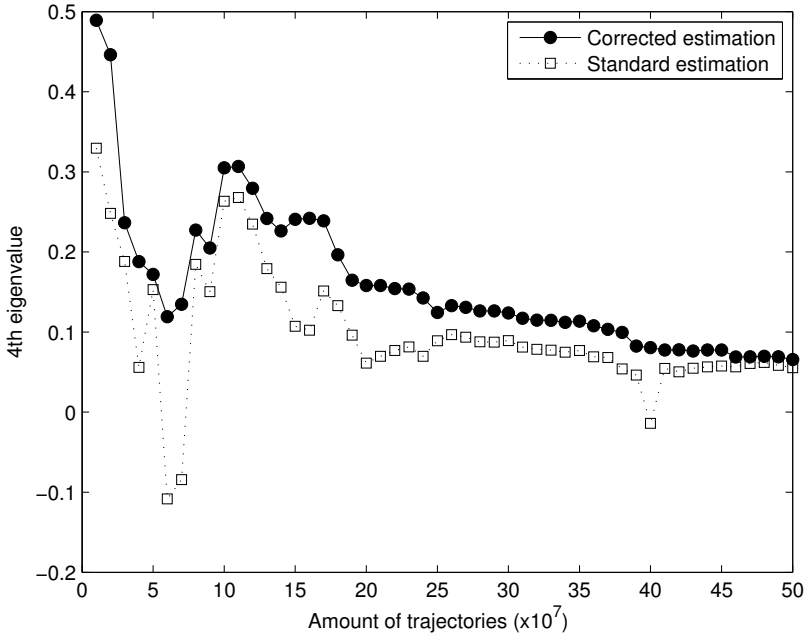
Corrected estimation
Standard estimation
4th eigenvalue
Amount of trajectories (x10^7)

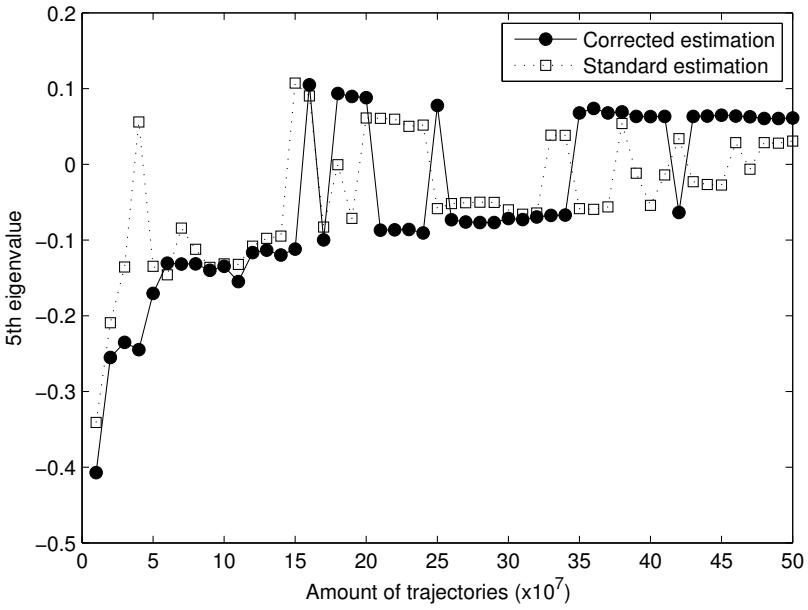
Corrected estimation
Standard estimation
5th eigenvalue
Amount of trajectories (x10^7)
0.2
0.1
0
−0.1
−0.2
−0.3
−0.4
−0.5
0
5
10
15
20
25
30
35
40
45
50

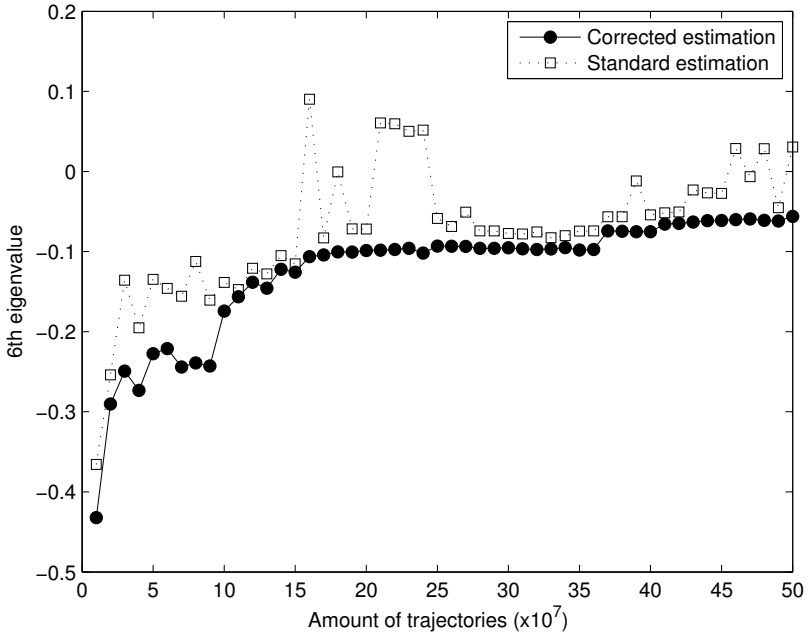
Corrected estimation
Standard estimation
6th eigenvalue
Amount of trajectories (x10^7)
0.2
0.1
0
−0.1
−0.2
−0.3
−0.4
−0.5
0
5
10
15
20
25
30
35
40
45
50

References

[1] S. Boyd and L. Vandenberghe, Convex optimization, Cambridge University Press, 2004.

[2] Peter Deuflhard and M. Weber, Robust Perron cluster analysis in conformation dynamics, In M. Dellnitz, S. Kirkland, M. Neumann, and C. Schütte, editors, Linear Algebra and Its Applications - Special Issue on Matrices and Mathematical Biology,volume 398C (2005), 161–184.

[3] G. Fischer, Lineare algebra, 17 ed., Vieweg + Teubner, 2010.

[4] Nicholas J. Higham, Computing a nearest symmetric positive semidefinite matrix, Linear Algebra and its Applications **103** (1988), ii–iii.

[5] R. Horn and C. Johnson, Matrix analysis, Cambridge University Press, 1985.

[6] Wilhelm Huisinga and Bernd Schmidt, Metastability and dominan eigenvaues of transfer operators, Lecture Notes in Computational Science and Engineering **49** (2006), 167–182.

[7] Ph. Metzner, F. Noé, and Ch. Schütte, Estimating the sampling error: Distribution of transition matrices and functions of transition matrices for given trajectory data, Physical Review E, 80 (2). 021106.

[8] A. Nielsen, K.Fackeldey, and M.Weber, Transfer operator for arbitary measures, in submission.

[9] N. V. Thoai R. Horst, P. M. Pardalos, Introduction to global optimization, Kluwer Academic Publishers, 1995.

[10] M. Sarich, Projected transfer operators - discretization of Markov processes in high-dimensional state space., Ph.D. thesis, Freie Universität Berlin, 2011.

[11] C. Schütte, Conformational dynamics: Modelling, theory, algorithm, and application to biomolecules., habilitation, Freie Universität Berlin, 1999.

[12] Christof Schütte, Wilhelm Huisinga, and Peter Deuflhard, Transfer operator approach to conformational dynamics in biomolecular systems, Applied Mathematics Entering the 21th Century (2003).

[13] Marcus Weber, A subspace approach to molecular Markov state models via a new infinitesimal generator., habilitation, Freie Universität Berlin, 2011.