

Research Repository

Tensor cross interpolation for global discrete optimization with application to Bayesian network inference

Accepted for publication in Computational Mathematics and Mathematical Physics.

Research Repository link: <https://repository.essex.ac.uk/40582/>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the published version if you wish to cite this paper.

<https://doi.org/10.1134/S0965542525700666>

Tensor cross interpolation for global discrete optimization with application to Bayesian network inference [★]

Sergey Dolgov¹, Dmitry Savostyanov²

¹University of Bath, Claverton Down, Bath, BA2 7AY, United Kingdom

²University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, United Kingdom

Abstract

Global discrete optimization is notoriously difficult due to the lack of gradient information and the curse of dimensionality, making exhaustive search infeasible. Tensor cross approximation is an efficient technique to approximate multivariate tensors (and discretized functions) by tensor product decompositions based on a small number of tensor elements, evaluated on adaptively selected fibers of the tensor, that intersect on submatrices of (nearly) maximum volume. The submatrices of maximum volume are empirically known to contain large elements, hence the entries selected for cross interpolation can also be good candidates for the globally maximal element within the tensor. In this paper we consider evolution of epidemics on networks, and infer the contact network from observations of network nodal states over time. By numerical experiments we demonstrate that the contact network can be inferred accurately by finding the global maximum of the likelihood using tensor cross interpolation. The proposed tensor product approach is flexible and can be applied to global discrete optimization for other problems, e.g. discrete hyperparameter tuning.

Keywords: epidemiological modelling, networks, tensor train, cross approximation, Bayesian inference.

1. Introduction

Global optimization of functions that may have multiple local optima is an infamously challenging problem, particularly in high dimensions. Continuous global optimization is more researched, and many well-performing methods exist, such as genetic algorithms [64, 67, 70], consensus-based optimization [6, 20, 31, 37, 44, 63], or Markov Chain Monte Carlo [42, 43]. In contrast, discrete global optimization of multivariate functions is short of state of the art methods, for the following two main reasons:

1. the absence of gradients — and more generally the lack of continuity in the discrete setting, which makes many efficient methods from the list above impossible to apply;
2. the curse of dimensionality — the search space size grows exponentially with the number of variables, hence a brute force search is unfeasible in more than a dozen of dimensions.

Methods based on the Markov process framework (now over a discrete state set) are one of the few feasible approaches. However, without gradient enhancements they may converge slowly and get stuck in local optima frequently.

A novel and promising approach to multivariate optimization, and generally to solving problems in high dimensions, is based on tensor product approximations. To introduce the concept, consider a function $f(\mathbf{x})$ defined on a Cartesian product of N discrete sets $\mathbb{X} \times \cdots \times \mathbb{X} =: \mathbb{X}^N \ni \mathbf{x}$. The function values $f(\mathbf{x})$ can be collected into a N -dimensional array $[f(\mathbf{x})]_{\mathbf{x} \in \mathbb{X}^N}$, also called a *tensor*. The tensor has $(\#\mathbb{X})^N$ entries, making it unfeasible to compute or store for large N , which is known as the curse of dimensionality. However, if the tensor admits a low-rank approximation, it may be possible to compute it efficiently and work with this representation instead.

In this paper we use the Tensor-Train (TT) approximation [52], noting that other tensor product formats and algorithms are available, see e.g. [32]. Even if the tensor itself is large and inaccessible in full, its TT

[★]Supported by the Engineering and Physical Sciences Research Council New Investigator Award EP/T031255/1.

Email addresses: S.Dolgov@bath.ac.uk (Sergey Dolgov), D.Savostyanov@essex.ac.uk (Dmitry Savostyanov)

approximation can be computed efficiently by tensor cross interpolation [17, 53, 56] from a number of adaptively selected samples, which shape one-dimensional lines, or *fibers*, in the given tensor. A fiber is a subset of positions $\mathbf{x}^{(s)} = (x_1, \dots, x_{n-1}, x_n^{(s)}, x_{n+1}, \dots, x_N)$, $s = 1, \dots, (\#\mathbb{X})$ in the tensor, such that only one index x_n is varying, and the other $N - 1$ indices are frozen. This approach is inspired by alternating least square optimization algorithms used in statistics [8, 33, 46, 69] and quantum physics [16, 38, 58]. On each step, positions of the next fiber(s) are chosen adaptively in order to maximize the volume (absolute value of the determinant) [25, 29] of the matrix of tensor elements formed by the intersection in each cross. Maximum-volume submatrices are empirically known to contain large elements, hence the entries selected for cross interpolation can also be good candidates for the globally maximal element within the tensor. Hence, a natural idea is to approximate a given tensor using tensor cross interpolation algorithm, and look for the maximum (in modulus) value of tensor among the entries chosen by the maximum-volume principle. Crucially, the number of samples per iteration is linear in N and $\#\mathbb{X}$, hence the outlined procedure is computationally feasible and allows a direct search for the maximum among the evaluated samples. The concrete way to generate fibers may vary. One option is to apply the so-called *maxvol* iteration algorithm [25] to reshaped factors (*cores*) of the TT decomposition directly [53]. Optimization based on this method was used successfully to speed up the search for molecular docking configurations [61, 62, 72, 73, 74] and reinforcement learning [59].

In this paper we apply a slightly different tensor cross interpolation with a greedy search for fibers [17, 56] to global discrete optimization which arises in Bayesian inference of a contact network from the observed dynamics of epidemic evolving on this network. The problem boils down to finding the network with the largest a posteriori estimate, which in absence of prior knowledge is equivalent to maximization of the likelihood function. Evaluating likelihood for each contact network is rather numerically expensive, since it requires either solving a high-dimensional system of ODEs, known as Markovian master equation, or accurately estimating the likelihood by extensive application of the Gillespie's stochastic simulation algorithm (SSA) [22]. Thus, we prefer the greedy cross interpolation instead of the *maxvol* TT cross algorithm [53], because in our numerical experience the former tends to require fewer samples than the latter.

Although the main contribution of this paper consists in testing the efficiency of the existing method applied to a relatively novel problem, we also suggest, test and evaluate a number of technical improvements to the greedy tensor cross interpolation algorithm [56], namely tempering and caching. Tempering introduces a parameter called temperature, which modifies the likelihood function and changes the high-dimensional landscape for interpolation and optimization. Caching simply means keeping all evaluated elements of the tensor, even those that are tested but not chosen by the adaptive tensor cross interpolation algorithm, as they may be tested again (and even selected) at later steps of the method. We demonstrate numerically that both enhancements considerably speed up the global optimization of the likelihood.

The rest of the paper is organised as follows. In Sec. 2 we describe the general framework for the Bayesian inference of the contact network from epidemiological data. In Sec. 3 we present the specific epidemiological model, denoted ε -SIS, that we use for the experiments in this paper. In Sec. 4 we recall the tensor cross interpolation algorithm from [17, 56] and state it using notation related to the specific problem considered in this paper to simplify the presentation. In Sec. 5 we describe tempering and caching. In Sec. 6 we apply the proposed method to Bayesian network inference of a simple contact network. We demonstrate that the tensor cross interpolation can be efficient as a heuristic algorithm for global discrete optimization, and that caching and tempering additionally improve its performance.

2. Bayesian network identification model

Following from our previous work [18], we consider the problem of identifying the network of contacts between nodes (species, individuals, groups, etc.) from observed discrete states of these nodes at given time points of a continuous-time Markov process. The network is an unweighted undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, consisting of a set of nodes $\mathcal{V} = \{1, 2, \dots, N\}$ and links (or edges) $\mathcal{E} = \{(m, n) : m \in \mathcal{V}, n \in \mathcal{V}, m \neq n\}$, such that $(m, n) \in \mathcal{E} \Leftrightarrow (n, m) \in \mathcal{E}$. We will say that nodes $m, n \in \mathcal{V}$ are connected, $m \sim n$, iff $(m, n) \in \mathcal{E}$. Each node can take one of discrete states, $x_n \in \mathbb{X}$, e.g. susceptible, infected, recovered. The state of the whole network is therefore a vector $\mathbf{x} = (x_1 \ x_2 \ \dots \ x_N)^T \in \mathbb{X}^N$.

We consider a continuous-time Markov jump process on the state space $\Omega = \mathbb{X}^N$, assuming that the probability to switch from a state \mathbf{x} to a state \mathbf{y} in a infinitesimal time interval dt is $p_{\mathbf{x} \rightarrow \mathbf{y}}^{\mathcal{G}} dt$ with some known transition rate $p_{\mathbf{x} \rightarrow \mathbf{y}}^{\mathcal{G}}$ that depends on the connectivity of the given network \mathcal{G} . For example, the probability of infection spreading between nodes n and m can be nonzero only if $m \sim n$. The system dynamics is

described by the probabilities of network states $p(x, t|\mathcal{G}) = \mathbb{P}(\text{system on a network } \mathcal{G} \text{ is in state } x \text{ at time } t)$. This probability function satisfies a system of ordinary differential equations (ODEs), known as *Markovian master equation* [9, 66], *Chapman* or *forward Kolmogorov equations*:

$$p'(x, t|\mathcal{G}) = \sum_{y \in \mathbb{X}^N} (p_{y \rightarrow x}^{\mathcal{G}} \cdot p(y, t|\mathcal{G}) - p_{x \rightarrow y}^{\mathcal{G}} \cdot p(x, t|\mathcal{G})), \quad x \in \mathbb{X}^N. \quad (1)$$

In practice, we can measure only particular states at particular points in time, which we collect into a dataset $X = \{t_k, x(t_k)\}_{k=0}^K$ of samples of the random variable $X(t)$ with the probability distribution function $p(x, t|\mathcal{G})$. We seek the maximum likelihood estimator (MLE) of the network configuration,

$$\mathcal{G}_{\text{opt}} = \arg \max_{\mathcal{G} \in \mathbb{G}} \mathcal{L}(\mathcal{G}), \quad (2)$$

where the likelihood of the dataset X is defined as

$$\mathcal{L}(\mathcal{G}) = \mathbb{P}(X|\mathcal{G}) = \mathbb{P}(X(t_1) = x_1, \dots, X(t_K) = x_K|\mathcal{G}), \quad (3)$$

and \mathbb{G} is some prior set of admissible networks, for example, all networks with unweighted edges. If a uniform prior is used, \mathcal{G}_{opt} is also the maximum a posteriori (MAP) estimator. Since the system dynamics is Markovian, the likelihood can be expanded as follows

$$\begin{aligned} \mathcal{L}(\mathcal{G}) &= \mathbb{P}(X(t_1) = x_1|X(t_0) = x_0, \mathcal{G}) \cdots \mathbb{P}(X(t_K) = x_K|X(t_{K-1}) = x_{K-1}, \mathcal{G}), \\ &= \prod_{k=1}^K \underbrace{\mathbb{P}(X(t_k) = x_k|X(t_{k-1}) = x_{k-1}, \mathcal{G})}_{\mathbb{P}(x_{k-1} \rightarrow x_k|\mathcal{G})}. \end{aligned} \quad (4)$$

The latter conditional probability is the solution to the master equation (1) started from 100% probability assigned to x_{k-1} as the initial state, and integrated over $[t_{k-1}, t_k]$,

$$\mathbb{P}(x_{k-1} \rightarrow x_k|\mathcal{G}) = p(x_k, t_k|\mathcal{G}), \quad \text{given} \quad p(x, t_{k-1}|\mathcal{G}) = \begin{cases} 1, & x = x_{k-1}, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

The master equation (1) can be solved by stochastic simulation algorithms [2, 22, 23, 36, 48], or directly, approximating the probability function $p(x, t|\mathcal{G})$ for example by sparse grids [34], adaptive finite state projections [7, 40, 50], radial basis functions [47], neural networks [30, 60], and tensor product approximations, such as canonical polyadic (CP) format [1, 35, 41], and TT format [11, 14, 15, 16, 19, 21, 39, 45, 68]. In this paper we use the latter approach [15, 19] since the likelihoods require both breaking the curse of dimensionality for large N , and estimating even small probabilities accurately, such that the error is below 10^{-3} – 10^{-4} .

3. ε -SIS model process and network parametrization

As a particular model, we consider the so-called ε -SIS process, as shown in Fig. 1. Each node can take only one of two states: $x_n \in \mathbb{X} = \{\text{susceptible, infected}\} = \{0, 1\}$, i.e. there is no recovered state. The system exhibits only two types of transitions:

$$p_{x \rightarrow y}^{\mathcal{G}} = \begin{cases} p_{x \rightarrow y}^{(\text{inf})}, & \text{if } \exists n \in \mathcal{V}, \text{ such that } y = x + e_n & (\text{infection}); \\ p_{x \rightarrow y}^{(\text{rec})}, & \text{if } \exists n \in \mathcal{V}, \text{ such that } y = x - e_n & (\text{recovery}); \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where $e_n \in \mathbb{R}^N$ is the n -th unit vector. For the simplicity sake, recovery rates $p_{x \rightarrow y}^{(\text{rec})} = \gamma$ are considered to be the same for all infected nodes. The infection rate for a node n (nonzero only if n is susceptible, i.e. $x_n = 0$) is set to $p_{x \rightarrow y}^{(\text{inf})} = I_n(x)\beta + \varepsilon$. Note that it grows with the number of infected neighbouring nodes, $I_n(x) = \#\{m \in \mathcal{V} : m \sim n, x_m = 1\}$, at the constant per-contact rate β , on top of the baseline infection rate ε introduced from outside the network.

With known number of nodes and constant rates $\beta, \gamma, \varepsilon$ of the only possible transitions in the ε -SIS process, the last bit of information comes from a contact network \mathcal{G} . The unweighted network can be




Figure 1. Markov chain and transitions between network states for the ε -SIS infection process. (a) In this example, the contact network is a chain of $N = 4$ people. (b) Transition rates in the Markov process depend not just on the size of infected and susceptible group, but on the exact state of the network. Note that the infection rate for node n is $I_n\beta + \varepsilon$, where I_n is the number of infected neighbours of node n , β is per contact infection rate, and ε is the self-infection rate, responsible for contacts with potential threats outside the network. (c) All network states, with recovery transitions shown with green arrows, and infection transition shown with red arrows with a circled number k indicating the infection rate $k\beta + \varepsilon$.

represented by adjacency matrix of edges, $G \in \{0, 1\}^{n \times n}$, where $g_{m,n} = 1$ indicates that m and n are connected, and $g_{m,n} = 0$ means they are not. Since the graph \mathcal{G} is undirected, its adjacency matrix is symmetric, $G = G^T$. Without loss of generality we take $g_{n,n} = 0$. Under these assumptions the graph \mathcal{G} is defined uniquely by the elements of the adjacency matrix G above the diagonal. We will collect the $\frac{1}{2}N(N-1)$ elements from the upper triangular part of G and stretch them in a single *adjacency vector* $\hat{\mathbf{g}} \in \{0, 1\}^d$, where $d := \frac{1}{2}N(N-1)$, and the elements are arranged as follows

$$\hat{\mathbf{g}} = (\hat{g}_1 \ \dots \ \hat{g}_d)^T := (g_{1,2} \ g_{1,3} \ g_{2,3} \ \dots \ g_{N-1,N})^T \in \{0, 1\}^d. \quad (7)$$

Since $\mathcal{G} \Leftrightarrow \hat{\mathbf{g}}$, the likelihood (3) can also be seen as a function of the adjacency vector, $L(\mathcal{G}) = L(\hat{\mathbf{g}}) : \{0, 1\}^d \rightarrow \mathbb{R}$, and the array of all likelihoods $[L(\mathcal{G})]_{\mathcal{G} \in \mathbb{G}}$ can be seen as a $2 \times 2 \times \dots \times 2$ tensor with d binary indices.

4. Tensor Train cross approximation and optimization

All possible likelihood values $L(\hat{\mathbf{g}})$ for all possible combinations of $\hat{\mathbf{g}}$ can be collected into a tensor of order d , with elements $L(\hat{g}_1, \dots, \hat{g}_d) = L(\hat{\mathbf{g}})$. In this section we consider low-rank cross approximation methods for approximation and optimization of this tensor.

4.1. Matrix cross interpolation

Consider first only two variables $\hat{\mathbf{g}} = (\hat{g}_1, \hat{g}_2)$, such that the likelihood tensor is a matrix with elements $L(\hat{g}_1, \hat{g}_2)$. To introduce a nontrivial low-rank matrix approximation (which will be necessary in higher-dimensional case), let us relax the assumption $\hat{\mathbf{g}} \in \{0, 1\}^2$ for the moment, and consider $\hat{\mathbf{g}} \in \{1, \dots, \hat{M}\} \times \{1, \dots, \hat{N}\}$. If the matrix $L \in \mathbb{R}^{\hat{M} \times \hat{N}}$ has rank $r \leq \min(\hat{M}, \hat{N})$, it can be interpolated exactly via any r of its linearly independent rows and columns,

$$\begin{aligned} L(\hat{g}_1, \hat{g}_2) &= \tilde{L}(\hat{g}_1, \hat{g}_2) := \sum_{s=1}^r \sum_{t=1}^r L(\hat{g}_1, \mathcal{I}_2^{(t)}) [L(\mathcal{I}_1, \mathcal{I}_2)]_{t,s}^{-1} L(\mathcal{I}_1^{(s)}, \hat{g}_2) \\ &= L(\hat{g}_1, \mathcal{I}_2) [L(\mathcal{I}_1, \mathcal{I}_2)]^{-1} L(\mathcal{I}_1, \hat{g}_2), \end{aligned} \quad (8)$$

where $\mathcal{I}_1 \subset \{1, \dots, \hat{M}\}$ is a set of row indices and $\mathcal{I}_2 \subset \{1, \dots, \hat{N}\}$ is a set of column indices. In general, one can aim at approximating $\tilde{L} \approx L$. The approximation (8) is still exact on the positions of the computed rows \mathcal{I}_1 and columns \mathcal{I}_2 , which together form a cross of elements of L , hence the name *cross interpolation*. However, to minimize the error in the other elements it becomes crucial to select $\mathcal{I}_1, \mathcal{I}_2$ that draw not just

Algorithm 1 One step of the matrix cross interpolation algorithm

Require: Current row and column sets $\mathcal{I}_1, \mathcal{I}_2$.

- 1: Pick a random set of samples $\mathcal{R} = \{(\hat{g}_1, \hat{g}_2)\}$ and choose the one with the largest error,
 $(\hat{g}_1^*, \hat{g}_2^*) \leftarrow \arg \max_{(\hat{g}_1, \hat{g}_2) \in \mathcal{R}} |\mathcal{L}(\hat{g}_1, \hat{g}_2) - \tilde{\mathcal{L}}(\hat{g}_1, \hat{g}_2)|$
- 2: **repeat**
- 3: $\hat{g}_1^* \leftarrow \arg \max_{\hat{g}_1 \in \{1, \dots, \hat{M}\} \setminus \mathcal{I}_1} |\mathcal{L}(\hat{g}_1, \hat{g}_2^*) - \tilde{\mathcal{L}}(\hat{g}_1, \hat{g}_2^*)|$
- 4: $\hat{g}_2^* \leftarrow \arg \max_{\hat{g}_2 \in \{1, \dots, \hat{N}\} \setminus \mathcal{I}_2} |\mathcal{L}(\hat{g}_1^*, \hat{g}_2) - \tilde{\mathcal{L}}(\hat{g}_1^*, \hat{g}_2)|$
- 5: **until** rook condition (9) holds **or** maximal number of iterations reached

Ensure: Expanded sets $\mathcal{I}_1 \leftarrow \mathcal{I}_1 \cup \{\hat{g}_1^*\}$, $\mathcal{I}_2 \leftarrow \mathcal{I}_2 \cup \{\hat{g}_2^*\}$, maximal error $\mathcal{E} = |\mathcal{L}(\hat{g}_1^*, \hat{g}_2^*) - \tilde{\mathcal{L}}(\hat{g}_1^*, \hat{g}_2^*)|$.




Figure 2. One step of the matrix cross interpolation algorithm

linearly independent, but as well-conditioned as possible rows and columns. Theoretical guarantees are available if $\mathcal{I}_1, \mathcal{I}_2$ are chosen such that the intersection matrix has maximum *volume*, $|\det \mathcal{L}(\mathcal{I}_1, \mathcal{I}_2)| \rightarrow \max$ over all possible $\mathcal{I}_1 \subset \{1, \dots, \hat{M}\}$ and $\mathcal{I}_2 \subset \{1, \dots, \hat{N}\}$ of size r [25, 27, 28, 29, 57]. By replacing the inverse of $\mathcal{L}(\mathcal{I}_1, \mathcal{I}_2)$ by the Moore–Penrose pseudoinverse, the results have been generalized to rectangular submatrices [49, 71]. However, to guarantee that a $r \times r$ submatrix of a $\hat{M} \times \hat{N}$ matrix with *both* $\hat{M}, \hat{N} > r$ has the maximum volume is NP-hard [4]. In practice, cheaper heuristic algorithms are used. For example, the *incomplete cross approximation* [53, 65] is based on fixing e.g. \mathcal{I}_2 to some initial guess, and searching only \mathcal{I}_1 as a maximum-volume set of rows in a smaller $\hat{M} \times r$ matrix, and vice versa. This so-called *maxvol* algorithm works well for many approximation and optimization problems; see e.g. [3, 10, 12, 13, 55, 59, 72, 73] for a very incomplete list of examples. However, for our needs it is somewhat inconvenient: the corresponding interpolation formula interpolates only after the full iteration over all variables is completed, which may require quite a few of expensive samples from \mathcal{L} .

The greedy algorithm for the matrix cross interpolation that we use in this paper can be seen as a version of the *adaptive cross approximation* (ACA) algorithm [5], or as the Gaussian elimination with partial pivoting [24]. The algorithm builds up the subsets of rows and columns for the interpolation by subsequently adding one row and column, which is characteristic for *greedy* algorithms. Specifically, we use the *rook* pivoting [51], which searches for a leading index $(\hat{g}_1^*, \hat{g}_2^*)$ in both row and column of the residual spawned from \hat{g}_1^* and \hat{g}_2^* ,

$$|\mathcal{L}(\hat{g}_1^*, \hat{g}_2^*) - \tilde{\mathcal{L}}(\hat{g}_1^*, \hat{g}_2^*)| \geq |\mathcal{L}(\hat{g}_1, \hat{g}_2) - \tilde{\mathcal{L}}(\hat{g}_1, \hat{g}_2)|, \quad \text{for all } (\hat{g}_1, \hat{g}_2) \text{ such that } \hat{g}_1 = \hat{g}_1^* \text{ or } \hat{g}_2 = \hat{g}_2^*. \quad (9)$$

The rook pivoting still needs an initial guess for \hat{g}_1^*, \hat{g}_2^* . We search for one among $\min(\hat{N}, \hat{M})$ indices \hat{g}_1, \hat{g}_2 sampled uniformly at random, as explained in Alg. 1 and visualised in Fig. 2. Note that one iteration of Alg. 1 needs at most $\hat{M} + \hat{N}$ evaluations of \mathcal{L} , and both the approximation error and candidate maxima can be updated frequently.

4.2. TT cross interpolation algorithm

Matrix methods can be applied to tensors reshaped to matrices. We introduce *multi-indices*

$$\hat{g}_{\leq k} = \hat{g}_1 \hat{g}_2 \dots \hat{g}_k, \quad \hat{g}_{> k} = \hat{g}_{k+1} \dots \hat{g}_d, \quad \forall k = 1, \dots, d-1.$$

Those allow us to reshape a tensor, e.g. $\mathbf{L} \in \mathbb{R}^{2 \times \dots \times 2}$, to a matrix $\mathbf{L}^{[k]} \in \mathbb{R}^{2^k \times 2^{d-k}}$, with

$$\mathbf{L}^{[k]}(\hat{g}_{\leq k}, \hat{g}_{> k}) = \mathbf{L}^{[k]}(\hat{g}_1 \hat{g}_2 \dots \hat{g}_k; \hat{g}_{k+1} \dots \hat{g}_d) = \mathbf{L}(\hat{g}_1, \hat{g}_2, \dots, \hat{g}_d),$$

The matrix interpolation formula (8) can be written for $\mathbf{L}^{[k]}$,

$$\mathbf{L}^{[k]}(\hat{g}_{\leq k}, \hat{g}_{> k}) \approx \tilde{\mathbf{L}}^{[k]}(\hat{g}_{\leq k}, \hat{g}_{> k}) = \mathbf{L}^{[k]}(\hat{g}_{\leq k}, \mathcal{I}_{> k})[\mathbf{L}^{[k]}(\mathcal{I}_{\leq k}, \mathcal{I}_{> k})]^{-1} \mathbf{L}^{[k]}(\mathcal{I}_{\leq k}, \hat{g}_{> k}). \quad (10)$$

However, since the ranges of $\hat{g}_{\leq k}$ and $\hat{g}_{> k}$ still suffer from the curse of dimensionality for intermediate k , direct usage of (10) is not practical.

Nonetheless, one of the factors of (10) is feasible for borderline k . For example, for $k = 1$, we can assume that some small enough index set $\mathcal{I}_{> 1}$ is given. Without a better initial guess, we can initialise it at random. The left factor $\mathbf{L}^{[1]}(\hat{g}_{\leq 1}, \mathcal{I}_{> 1}) = \mathbf{L}(\hat{g}_1, \mathcal{I}_{> 1})$ is sufficiently small, and can be assembled. Similarly, the left index set $\mathcal{I}_{\leq 1}$ can be assembled and/or updated by the information in $\mathbf{L}(\hat{g}_1, \mathcal{I}_{> 1})$. This sets the base for a recursive factorization of $\mathbf{L}(\mathcal{I}_{\leq 1}, \hat{g}_2, \dots, \hat{g}_d)$.

For a general $k \in \{1, \dots, d-1\}$, we can assume that both $\mathcal{I}_{\leq k-1} = \{\hat{g}_{\leq k-1}\}$ of size r_{k-1} and $\mathcal{I}_{> k+1} = \{\hat{g}_{> k+1}\}$ of size r_{k+1} are available. For $k = 1$ or $k = d$ we can let $\mathcal{I}_{\leq 0} = \mathcal{I}_{> d} = \emptyset$. The corresponding subtensor defined by $\mathbf{L}(\mathcal{I}_{\leq k-1}, \hat{g}_k, \hat{g}_{k+1}, \mathcal{I}_{> k+1})$ contains only $4r_{k-1}r_{k+1}$ elements. We can reshape it into a $2r_{k-1} \times 2r_{k+1}$ matrix with elements $\mathbf{L}(\mathcal{I}_{\leq k-1} \hat{g}_k; \hat{g}_{k+1} \mathcal{I}_{> k+1})$, and factorize it using the matrix cross interpolation algorithm. Moreover, if some $\mathcal{I}_{\leq k}$ and $\mathcal{I}_{> k}$ are already available too, we don't need to assemble $\mathbf{L}(\mathcal{I}_{\leq k-1} \hat{g}_k; \hat{g}_{k+1} \mathcal{I}_{> k+1})$ in full, but we can apply Algorithm 1 starting with $\mathcal{I}_{\leq k}$ and $\mathcal{I}_{> k}$ to find only one new pivot at a time. This gives us a cross interpolation

$$\begin{aligned} \mathbf{L}(\mathcal{I}_{\leq k-1} \hat{g}_k; \hat{g}_{k+1} \mathcal{I}_{> k+1}) &\approx \tilde{\mathbf{L}}(\mathcal{I}_{\leq k-1} \hat{g}_k; \hat{g}_{k+1} \mathcal{I}_{> k+1}) \\ &:= \mathbf{L}(\mathcal{I}_{\leq k-1}, \hat{g}_k, \mathcal{I}_{> k})[\mathbf{L}(\mathcal{I}_{\leq k}, \mathcal{I}_{> k})]^{-1} \mathbf{L}(\mathcal{I}_{\leq k}, \hat{g}_{k+1}, \mathcal{I}_{> k+1}). \end{aligned} \quad (11)$$

Recall that in the base of the recursion we left behind the left factor $\mathbf{L}(\hat{g}_1, \mathcal{I}_{> 1})[\mathbf{L}(\mathcal{I}_{\leq 1}, \mathcal{I}_{> 1})]^{-1}$. Completing the recursion, we arrive at the following high-dimensional cross interpolation:

$$\begin{aligned} \mathbf{L}(\hat{g}_1, \dots, \hat{g}_d) &\approx \tilde{\mathbf{L}}(\hat{g}_1, \dots, \hat{g}_d) \\ &:= \mathbf{L}(\hat{g}_1, \mathcal{I}_{> 1})[\mathbf{L}(\mathcal{I}_{\leq 1}, \mathcal{I}_{> 1})]^{-1} \mathbf{L}(\mathcal{I}_{\leq 1}, \hat{g}_2, \mathcal{I}_{> 2})[\mathbf{L}(\mathcal{I}_{\leq 2}, \mathcal{I}_{> 2})]^{-1} \dots \mathbf{L}(\mathcal{I}_{\leq d-1}, \hat{g}_d). \end{aligned} \quad (12)$$

Merging the inverse matrices with either left or right factors, we see that (12) resembles the traditional TT decomposition [52]

$$\tilde{\mathbf{L}}(\hat{g}_1, \dots, \hat{g}_d) = \mathbf{L}^{(1)}(1, \hat{g}_1, :) \mathbf{L}^{(2)}(:, \hat{g}_2, :) \dots \mathbf{L}^{(d)}(:, \hat{g}_d, 1), \quad (13)$$

where $\mathbf{L}^{(k)} \in \mathbb{R}^{r_{k-1} \times 2 \times r_k}$ are the *TT cores*, such that $\mathbf{L}^{(k)}(:, \hat{g}_k, :) \in \mathbb{R}^{r_{k-1} \times r_k}$ for $k = 1, \dots, d$, with $r_0 = r_d = 1$, and the sizes of TT cores and index sets r_1, \dots, r_{d-1} are called *TT ranks*.

Note that in the subtensor interpolation (11), the set $\mathcal{I}_{\leq k}$ contains indices, each of which consists of an element from $\mathcal{I}_{\leq k-1}$ and some sample of \hat{g}_k . Similarly, $\mathcal{I}_{> k}$ samples from the union of $\mathcal{I}_{> k+1}$ and the range of \hat{g}_{k+1} . This defines *nestedness* relations between the index sets,

$$\mathcal{I}_{\leq k+1} \subset \mathcal{I}_{\leq k} \times \{0, 1\}, \quad \mathcal{I}_{> k} \subset \{0, 1\} \times \mathcal{I}_{> k+1}, \quad k = 1, \dots, d-1, \quad (14)$$

which is maintained as long as the indices are only added, never removed. This allows us to update the interpolations (11) for any $k = 1, \dots, d-1$, in any order, which enabled efficient parallelisation of this algorithm [17]. For simplicity in this paper, we update the interpolations sequentially [56], as explained in Alg. 2 and shown in Fig. 3.

It remains to iterate the whole process over the TT ranks until an error or complexity stopping criterion is met, recording the maximal element of $\mathbf{L}(\hat{\mathbf{g}})$ among those seen in the course of iterations. This is summarized in Algorithm 3.

Algorithm 2 Left-to-right sweep of the greedy TT cross interpolation algorithm [56]

Require: Sets $(\mathcal{I}_{\leq k}, \mathcal{I}_{> k})$ of the interpolation (12)

- 1: **for** $k = 1, \dots, d - 1$ **do**
- 2: Apply Alg. 1 to the subtensor $\mathcal{L}(\mathcal{I}_{\leq k-1}^* \hat{g}_k, \hat{g}_{k+1} \mathcal{I}_{> k+1})$ seen as a $(r_{k-1} \cdot 2) \times (2 \cdot r_{k+1})$ matrix. Find the new pivot $(\hat{g}_{\leq k}^*, \hat{g}_{> k}^*)$ and the maximal error $\mathcal{E}_k = |\mathcal{L}(\hat{g}_{\leq k}^*, \hat{g}_{> k}^*) - \tilde{\mathcal{L}}(\hat{g}_{\leq k}^*, \hat{g}_{> k}^*)|$
- 3: $\mathcal{I}_{\leq k}^* \leftarrow \mathcal{I}_{\leq k} \cup \{\hat{g}_{\leq k}^*\}; \mathcal{I}_{> k}^* \leftarrow \mathcal{I}_{> k} \cup \{\hat{g}_{> k}^*\}$
- 4: **end for**

Ensure: Updated index sets $(\mathcal{I}_{\leq k}, \mathcal{I}_{> k}) \leftarrow (\mathcal{I}_{\leq k}^*, \mathcal{I}_{> k}^*), k = 1, \dots, d - 1$, maximal error $\mathcal{E} = \max_k \mathcal{E}_k$.




Figure 3. Tensor cross interpolation algorithm visualised. Top: one step of tensor cross interpolation updates searches for a new pivot $(\hat{g}_{\leq k}^*, \hat{g}_{> k}^*)$ in the submatrix $\mathcal{L}(\mathcal{I}_{\leq k-1} \hat{g}_k, \hat{g}_{k+1} \mathcal{I}_{> k+1})$ and adds a new cross to the index set $(\mathcal{I}_{\leq k}, \mathcal{I}_{> k})$. Bottom: the process repeats for $k = 1, \dots, d - 1$ (left-to-right sweep) and back for $k = d - 1, \dots, 1$ (right-to-left sweep).

5. Tempering and caching of the objective function

Efficiency of the TT approximation depends on the decay of singular values of the matricisations $\mathcal{L}^{[k]}$ of the tensor. In turn, those depend on complexity of the likelihood function. Although, as we will see in the numerical experiments, the candidate maximizer produced by Alg. 3 can be much more accurate than the approximate tensor itself, it is still beneficial to be able to vary the complexity of the likelihood for the same data.

One non-intrusive way to achieve this is *tempering*. Instead of approximating (and maximizing) $\mathcal{L}(\hat{\mathbf{g}})$ directly, we compute the pointwise power of the likelihood, $\mathcal{L}(\hat{\mathbf{g}})^{1/\tau}$. The parameter $\tau > 0$ is similar to temperature in the Boltzmann distribution, hence the name. To avoid over- and under-flow errors in practical computations, we assemble the tempered likelihood from a sum of logarithms of likelihoods in (4),

$$\mathcal{L}(\hat{\mathbf{g}})^{1/\tau} = \exp \left(\frac{\sum_{k=1}^K \log \mathcal{P}(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k | \mathcal{G})}{\tau} \right). \quad (15)$$

Note that $\mathcal{L}(\hat{\mathbf{g}})^{1/\tau} \rightarrow 1$ as $\tau \rightarrow \infty$. As all entries of the tensor approach 1, the whole tensor approaches an array of ones, which has a trivial rank-1 TT decomposition:

$$\lim_{\tau \rightarrow \infty} [\mathcal{L}(\hat{\mathbf{g}})^{1/\tau}]_{\hat{\mathbf{g}} \in \{0,1\}^d} = [1]_{\hat{\mathbf{g}} \in \{0,1\}^d} = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \times \begin{pmatrix} 1 \\ 1 \end{pmatrix} \times \dots \times \begin{pmatrix} 1 \\ 1 \end{pmatrix}. \quad (16)$$

This example shows that by increasing the temperature τ we reduce the *contrast* of the high-dimensional

Algorithm 3 Full greedy TT cross interpolation and maximization

Require: A procedure to evaluate $L(\hat{\mathbf{g}})$ for any $\hat{\mathbf{g}}$, initial index $\hat{\mathbf{g}}^0 = (\hat{g}_1^0, \dots, \hat{g}_d^0)$, stopping thresholds for error $\delta \geq 0$, TT rank r_{\max} and computational budget N_{\max} .

- 1: Initialise $\mathcal{I}_{\leq k} = \{\hat{g}_1^0 \dots \hat{g}_k^0\}$ and $\mathcal{I}_{> k} = \{\hat{g}_{k+1}^0 \dots \hat{g}_d^0\}$ for $k = 1, \dots, d-1$, $\hat{\mathbf{g}}^{\max} = \hat{\mathbf{g}}^0$, $r_0 = \dots = r_d = 1$, $N_{\text{eval}} = 0$.
- 2: **repeat**
- 3: Apply Algorithm 2 to $L(\hat{\mathbf{g}})$ starting with $\mathcal{I}_{\leq k}, \mathcal{I}_{> k}$. Find the new index sets $\mathcal{I}_{\leq k}, \mathcal{I}_{> k}$ for $k = 1, \dots, d-1$, and the maximal error \mathcal{E} .
- 4: Update $\hat{\mathbf{g}}^{\max} = \arg \max L(\hat{\mathbf{g}})$ over $\hat{\mathbf{g}} \in \{\hat{\mathbf{g}}^{\max}, (\mathcal{I}_{\leq 1}, \mathcal{I}_{> 1}), \dots, (\mathcal{I}_{\leq d-1}, \mathcal{I}_{> d-1})\}$.
- 5: Update $r_k := r_k + 1$ for $k = 1, \dots, d-1$.
- 6: Update $N_{\text{eval}} := N_{\text{eval}} + \text{number of samples of } L(\hat{\mathbf{g}}) \text{ drawn in Alg. 2.}$
- 7: **until** $\mathcal{E} \leq \delta \cdot L(\hat{\mathbf{g}}^{\max})$, $r_k \geq r_{\max} \forall k = 1, \dots, d-1$ or $N_{\text{eval}} \geq N_{\max}$.

Ensure: Expanded index sets $(\mathcal{I}_{\leq k}, \mathcal{I}_{> k})$, candidate maximizer $\hat{\mathbf{g}}^{\max}$.

landscape of the likelihoods, which also reduces the complexity of representation manifested by low TT ranks.

Increasing the temperature and lowering the contrast generally makes it easier for the tensor cross interpolation algorithm to approximate the tempered high-dimensional likelihood. However, extremely high temperatures will wipe all contrast and lose all spacial information. In the extreme example above, Alg. 3 (through Alg. 1) is able to achieve good approximation of (16) by choosing any interpolation indices, not necessarily those that maximize $L(\hat{\mathbf{g}})$.

In practice, we aim to find a suitable value of τ that maintains reasonable contrast to allow us to perform optimization, while reducing the TT ranks of the representation and hence the number of computations of the likelihoods required to locate the maximum. In our numerical examples we find that it is sufficient to find an optimal τ up to an order of magnitude, which is feasible for trial and error over only a few values.

The TT approximation of $P(x_{k-1} \rightarrow x_k | \mathcal{G}) = p(x_k, t_k | \mathcal{G})$ notwithstanding, solving many master equations for many data points can take a lot of computing time, e.g. tens to hundreds of seconds per total likelihood evaluation. Fortunately, the cross structure of samples in Alg. 3 contains a lot of repeating indices. To reduce the computing time, we store a *cache* of all unique indices appearing in Alg. 3 together with their (tempered) likelihood values. Whenever a tensor element is requested at the index which was already considered in the previous iterations, it is taken from the cache instead of evaluating the likelihood from scratch.

6. Numerical experiments

We test the accuracy of the cross interpolation applied for the global optimization of the likelihood by comparing the inferred network \mathcal{G} to the ‘ground truth’ network \mathcal{G}_\star , which in this experiment is a linear chain with $N = 9$ nodes as shown in Fig. 4(a). We assume that the initial state $x_0 = (1 \ 0 \ \dots \ 0)$ is the same for all data samples. We sample random walks of ε -SIS process with parameters $\beta = 1$, $\gamma = 0.5$ and $\varepsilon = 0.01$ every $\Delta t = 0.1$ time units for the duration of $T = 200$ time units. This produces $K = 2000$ data points $t_k = \Delta t k$, $k = 1, \dots, K$, which form the likelihood as shown in (3).

The TT cross interpolation algorithm starts from the index corresponding to the score-based approximate inference as described in [18], and runs for 4 full iterations, thereby achieving the maximal TT rank of 5. The master equation (1) is solved on each time interval $t \in [t_{k-1}, t_k]$ using the tAMEn algorithm [15] with Chebyshev polynomials of degree 12 in time and a relative stopping and TT truncation threshold $\epsilon_{\text{tAMEn}} = 10^{-6}$.

In Fig. 4(b) and (c) we show the convergence of the inferred maximum-likelihood network \mathcal{G} towards the ground truth network \mathcal{G}_\star , as the TT cross interpolation algorithm progresses. The error in the network is computed as the number of incorrectly inferred links,

$$\|\mathcal{G} - \mathcal{G}_\star\|_1 = \#\{m, n \in \mathcal{V}, m > n : g_{m,n} \neq g_{m,n}^\star\}, \quad (17)$$

relative to the total number of possible links, $\frac{1}{2}N(N-1)$. Since the error (17) depends on the particular data in addition to the algorithm, we repeat each experiment (both the data generation and the network inference) 42 times (the number of available cores on the Intel Xeon Platinum 8168 CPU node used), and plot the average error over those runs, and the interval 1 standard deviation above and below the average.

Fig. 4(b) shows the convergence with respect to the number of elements the TT cross algorithm draws from the original tensor. Many of them are repetitive, and hence take almost zero time compared to the likelihood computation at the new samples. Fig. 4(b) shows the convergence with respect to the actual CPU time. For this example, about 2/3 of the samples are retrieved from the cache. This is a reasonable balance between sparing the computational costs and exploration of the likelihood landscape.

We see that the lowest temperature seems to be most efficient, so potentially one would try to reduce it even further to e.g. 0.1. However, for most of the datasets this experiment could not complete, since some likelihood values exceeded the overflow threshold, and the algorithm has stopped. This indicates that the tempering should be chosen reasonably. Fortunately, the algorithm is quite insensitive to the particular temperature value within an order of magnitude of the optimal one, so one needs to try only a few values (such as the orders of 10).

The maximal TT rank of 5 seems too low to approximate a complicated likelihood accurately. However, the cross algorithm can find a good candidate maximizer even with such small TT ranks, far before a good approximation is reached. In Fig. 4(d) and (e) we plot the histogram of $\|\mathcal{G} - \mathcal{G}_\star\|_1$ and the relative mean square error of the whole tensor, $\|\mathcal{L} - \tilde{\mathcal{L}}\|_F / \|\mathcal{L}\|_F$ for $\tau = 1$ over $N_s = 39$ experiments (the remaining 3 experiments did not complete due to overflow in particularly high likelihood values). For such a high-contrast likelihood we see that the tensor approximation error is actually larger than 1 for 70% of the data realisations, while 1 or more incorrect links are found in only 30% of the cases.

7. Conclusion

In this work, we have shown that the TT cross interpolation algorithm can be used to infer the network structure of the epidemic process from observational data on the states of this process. The approach is based on the tendency of the rook pivoting in the TT cross interpolation algorithm to find tensor elements close to the global maximum among local maxima of the residual. It was observed that maximum-volume submatrices tend to contain the maximal element of the matrix [26], and here we demonstrated the same effect for the greedy rook pivoting cross interpolation. The latter requires fewer evaluations of tensor elements after which a consistent approximation or maximization can be produced: the rook pivoting samples only one tensor fiber at a time, while the maximum-volume based methods sample $O(r)$ fibers at once. A significant further reduction of the computing time (66% in our example) is achieved by caching of tensor elements, thanks to many cross indices inherited from iteration to iteration in the greedy cross algorithm.

However, the efficiency of maximization can be limited by the restricted nested indices, maintaining the low-rank TT structure. Moreover, a local maximum of the residual does not have to be the global maximum of the tensor, especially if the tensor contains many nearby elements. In the latter case we can say that the global maximum has low contrast. Specifically for the likelihood maximization, we considered tempering as a trade-off between the low-rankness of the tensor of values of the objective function and the contrast of the global maximum. Intermediate values of the tempering parameter perform best, but it does not have to be very precise, since the cross interpolation algorithm can reveal a good candidate maximizer even at low TT ranks, that would be insufficient for an accurate approximation of the whole tensor. It was shown that high-dimensional low-rank canonical tensor decomposition can implicitly denoise the tensor [54]. The TT cross interpolation seems to be recovering most important fibers (and maxima therein) at low TT ranks too. This paves the way for using the proposed tempered cross interpolation for global maximization of other functions, especially in discrete high-dimensional spaces. Sharp estimation of TT ranks and optimization performance for tempered functions remains to be a future research. A necessary condition seems to be that the function is strictly uniformly positive, for example, of exponential family, such that the tempered function can be made arbitrarily close to 1.

References

- [1] A. Ammar, E. Cueto, and F. Chinesta. Reduction of the chemical master equation for gene regulatory networks using proper generalized decompositions. *Int. J. Numer. Meth. Biomed. Engng.*, 28(9): 960–973, 2012.

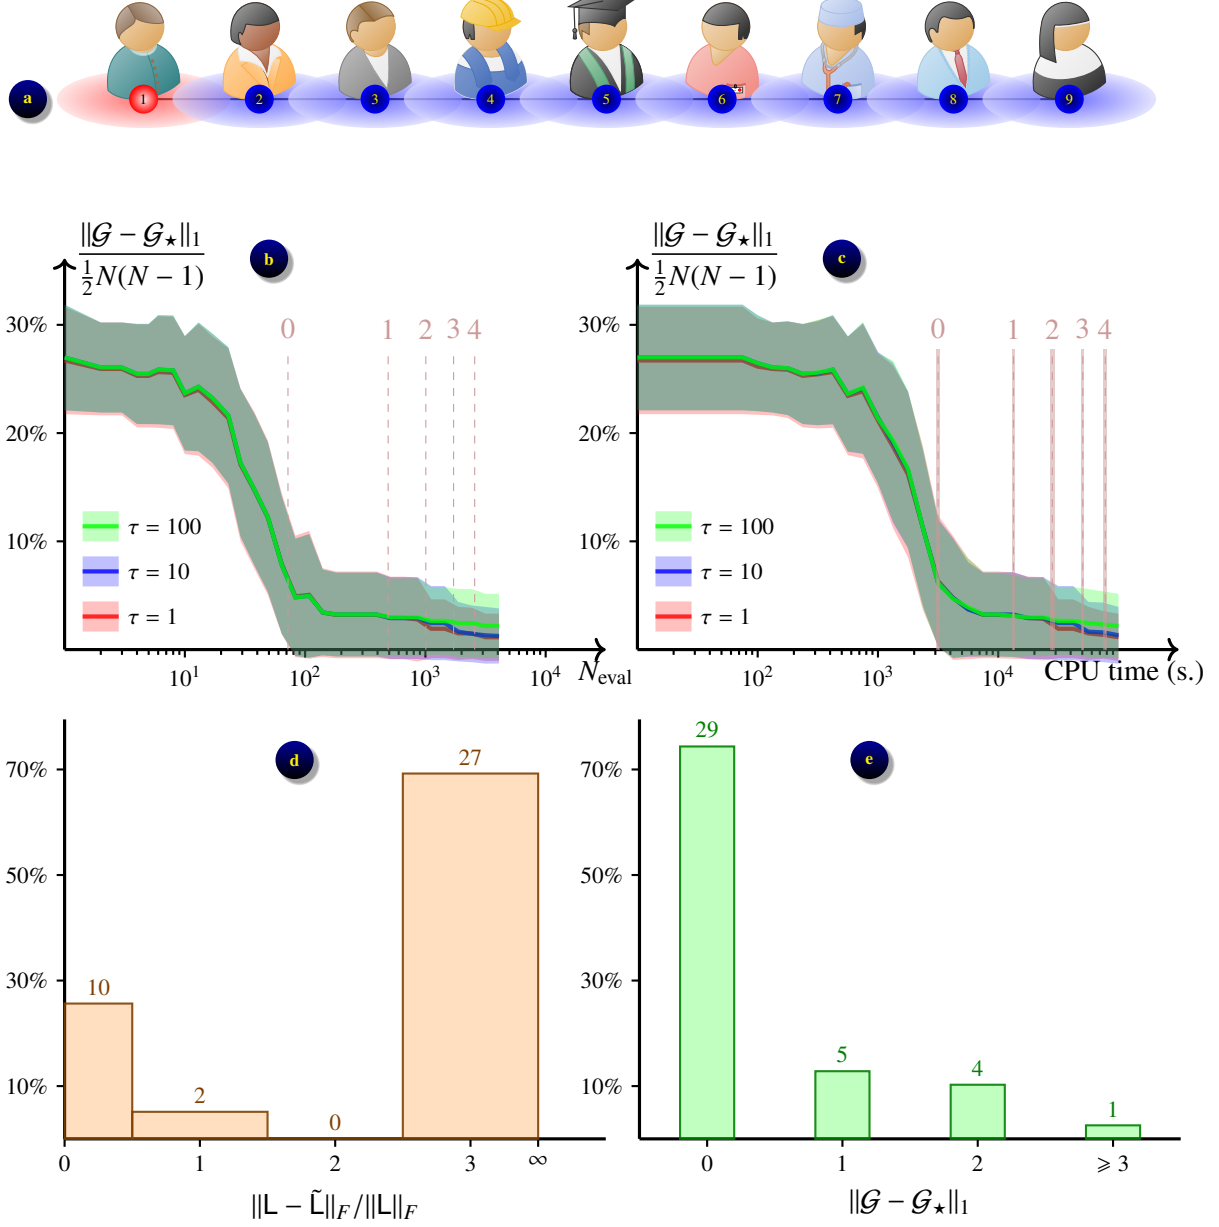


Figure 4. Inferring linear chain network with $N = 9$ people from ε -SIS epidemic process with $\beta = 1$, $\gamma = 0.5$ and $\varepsilon = 0.01$: (a) the ground truth network \mathcal{G}_\star in its initial state; (b) and (c) convergence of the network \mathcal{G} towards \mathcal{G}_\star in the TT cross interpolation algorithm; average (solid lines) \pm one standard deviation (shaded areas) over $N_s = 42$ datasets; shown for temperatures $\tau = 1, 10, 100$. Vertical dashed lines in (b) and (c) indicate full iterations of the TT cross interpolation algorithm. Intervals around the dashed lines in (c) indicate the average time ± 1 standard deviation. (d) distribution of relative errors for the tensor approximation, based on $N_s = 39$ datasets; (e) distribution of errors for the network inference, based on $N_s = 39$ datasets.

- [2] David F. Anderson and Desmond J. Higham. Multilevel Monte Carlo for continuous time Markov chains, with applications in biochemical kinetics. *Multiscale Modeling & Simulation*, 10(1):146–179, 2012.
- [3] Harbir Antil, Sergey Dolgov, and Akwum Onwunta. TTRISK: Tensor train decomposition algorithm for risk averse optimization. *Numerical Linear Algebra with Applications*, 30(3):e2481, 2023.
- [4] John J Bartholdi III. A good submatrix is hard to find. *Operations Research Lett.*, 1(5):190–193, 1982.
- [5] M. Bebendorf. Approximation of boundary element matrices. *Numer. Math*, 86(4):565–589, 2000.
- [6] Giacomo Borghi, Michael Herty, and Lorenzo Pareschi. Constrained consensus-based optimization. *SIAM J Optimization*, 33(1):211–236, 2023.
- [7] Youfang Cao, Anna Terebus, and Jie Liang. State space truncation with quantified errors for accurate solutions to discrete chemical master equation. *Bulletin of Mathematical Biology*, 78(4):617–661, 2016.
- [8] J. D. Carroll and J. J Chang. Analysis of individual differences in multidimensional scaling via n-way generalization of Eckart–Young decomposition. *Psychometrika*, 35:283–319, 1970.
- [9] Wei-Yin Chen and Sankar Bokka. Stochastic modeling of nonlinear epidemiology. *Journal of Theoretical Biology*, 234(4):455–470, 2005.
- [10] T. Cui and S. Dolgov. Deep composition of Tensor Trains using squared inverse Rosenblatt transports. arXiv preprint 2007.06968, 2020.
- [11] Trang Dinh and Roger B Sidje. An adaptive solution to the chemical master equation using quantized tensor trains with sliding windows. *Physical Biology*, 17(6):065014, 2020.
- [12] S. Dolgov, K. Anaya-Izquierdo, C. Fox, and R. Scheichl. Approximation and sampling of multivariate probability distributions in the tensor train decomposition. *Statistics and Computing*, 30:603–625, 2020.
- [13] S. Dolgov, D. Kalise, and L. Saluzzi. Data-driven Tensor Train gradient cross approximation for Hamilton–Jacobi–Bellman equations. *SIAM Journal on Scientific Computing*, 45(5):A2153–A2184, 2023.
- [14] S. Dolgov and B. Khoromskij. Simultaneous state-time approximation of the chemical master equation using tensor product formats. *Numer. Linear Algebra Appl.*, 22(2):197–219, 2015.
- [15] S. V. Dolgov. A tensor decomposition algorithm for large ODEs with conservation laws. *Computational Methods in Applied Mathematics*, 19:23–38, 2019.
- [16] S. V. Dolgov and D. V. Savostyanov. Alternating minimal energy methods for linear systems in higher dimensions. *SIAM J. Sci. Comput.*, 36(5):A2248–A2271, 2014.
- [17] Sergey Dolgov and Dmitry Savostyanov. Parallel cross interpolation for high-precision calculation of high-dimensional integrals. *Comp. Phys. Comm.*, 246:106869, 2020.
- [18] Sergey Dolgov and Dmitry Savostyanov. Tensor product algorithms for inference of contact network from epidemiological data. *BMC Bioinformatics*, 25, 2024. article no. 285.
- [19] Sergey Dolgov and Dmitry Savostyanov. Tensor product approach to modelling epidemics on networks. *Applied Mathematics and Computation*, 460:128290, 2024.
- [20] Massimo Fornasier, Timo Klock, and Konstantin Riedl. Consensus-based optimization methods converge globally. *SIAM J Optimization*, 34(3):2973–3004, 2024.
- [21] Patrick Gelß, Sebastian Matera, and Christof Schütte. Solving the master equation without kinetic Monte Carlo: Tensor train approximations for a CO oxidation model. *Journal of Computational Physics*, 314:489–502, 2016.
- [22] Daniel T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J Comput. Phys.*, 22(4):403–434, 1976.
- [23] Daniel T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 115(4):1716–1733, 2001.
- [24] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, 2013.
- [25] S. A. Goreinov, I. V. Oseledets, D. V. Savostyanov, E. E. Tyrtyshnikov, and N. L. Zamarashkin. How to find a good submatrix. In *Matrix Methods: Theory, Algorithms, Applications*, V. Olshevsky and E. Tyrtyshnikov, editors, World Scientific, Hackensack, NY, 2010, pages 247–256.
- [26] S. A. Goreinov and E. E. Tyrtyshnikov. The maximal-volume concept in approximation by low-rank matrices. *Contemporary Mathematics*, 280:47–51, 2001.
- [27] S. A. Goreinov and E. E. Tyrtyshnikov. Quasioptimality of skeleton approximation of a matrix in the Chebyshev norm. *Doklady Math.*, 83(3):374–375, 2011.
- [28] S. A. Goreinov, E. E. Tyrtyshnikov, and N. L. Zamarashkin. A theory of pseudo-skeleton approximations. *Linear Algebra Appl.*, 261:1–21, 1997.
- [29] S. A. Goreinov, N. L. Zamarashkin, and E. E. Tyrtyshnikov. Pseudo-skeleton approximations by

- matrices of maximum volume. *Mathematical Notes*, 62(4):515–519, 1997.
- [30] Ankit Gupta, Christoph Schwab, and Mustafa Khammash. DeepCME: A deep learning framework for computing solution statistics of the chemical master equation. *PLoS Comput Biol*, 17(12):e1009623, 2021.
 - [31] Seung-Yeal Ha, Shi Jin, and Doheon Kim. Convergence and error estimates for time-discrete consensus-based optimization algorithms. *Numer. Math.*, 147:255–282, 2021.
 - [32] W. Hackbusch. *Tensor Spaces And Numerical Tensor Calculus*. Springer-Verlag, Berlin, 2012. ISBN 978-3642280269.
 - [33] R. A. Harshman. Foundations of the PARAFAC procedure: models and conditions for an explanatory multimodal factor analysis. *UCLA Working Papers in Phonetics*, 16:1–84, 1970.
 - [34] M. Hegland, C. Burden, L. Santoso, S. MacNamara, and H. Booth. A solver for the stochastic master equation applied to gene regulatory networks. *Journal of Computational and Applied Mathematics*, 205(2):708 – 724, 2007.
 - [35] M. Hegland and J. Garcke. On the numerical solution of the chemical master equation with sums of rank one tensors. *ANZIAM*, 52:C628–C643, 2011.
 - [36] M. Hemberg and M. Barahona. Perfect sampling of the master equation for gene regulatory networks. *Biophysical journal*, 93(2):401–410, 2007.
 - [37] Hui Huang, Jinniao Qiu, and Konstantin Riedl. Consensus-based optimization for saddle point problems. *SIAM J Control and Optimization*, 62(2):1093–1121, 2024.
 - [38] C. Hubig, I. P. McCulloch, U. Schollwöck, and F. A. Wolf. Strictly single-site DMRG algorithm with subspace expansion. *Phys. Rev. B*, 91(15):155115, 2015.
 - [39] Ion Gabriel Ion, Christian Wildner, Dimitrios Loukrezis, Heinz Koepl, and Herbert De Gersem. Tensor-train approximation of the chemical master equation and its application for parameter inference. *The Journal of Chemical Physics*, 155(3):034102, 2021.
 - [40] T. Jahnke. An adaptive wavelet method for the chemical master equation. *SIAM J. Sci. Comput.*, 31(6):4373, 2010.
 - [41] T. Jahnke and W. Huisinga. A dynamical low-rank approach to the chemical master equation. *Bulletin of Mathematical Biology*, 70:2283–2302, 2008.
 - [42] H. Jensen, D. Jerez, and M. Beer. A general two-phase Markov chain Monte Carlo approach for constrained design optimization: Application to stochastic structural optimization. *Computer Methods in Applied Mechanics and Engineering*, 373:113487, 2021.
 - [43] H.A. Jensen, D.J. Jerez, and M. Valdebenito. An adaptive scheme for reliability-based global design optimization: A Markov chain Monte Carlo approach. *Mechanical Systems and Signal Processing*, 143:106836, 2020.
 - [44] Dante Kalise, Akash Sharma, and Michael V. Tretyakov. Consensus-based optimization via jump-diffusion stochastic differential equations. *Mathematical Models and Methods in Applied Sciences*, 33(02):289–339, 2023.
 - [45] Vladimir Kazeev, Mustafa Khammash, Michael Nip, and Christoph Schwab. Direct solution of the Chemical Master Equation using Quantized Tensor Trains. *PLOS Computational Biology*, 10(3): e100359, 2014.
 - [46] P.M. Kroonenberg and J. de Leeuw. Principal component analysis of three-mode data by means of alternating least squares algorithms. *Psychometrika*, 45(1):69–97, 1980.
 - [47] Ivan Kryven, Susanna Röblitz, and Christof Schütte. Solution of the chemical master equation by radial basis functions approximation with interface tracking. *BMC Systems Biology*, 9(1):67, 2015.
 - [48] Christopher Lester, Ruth E. Baker, Michael B. Giles, and Christian A. Yates. Extending the multi-level method for the simulation of stochastic biological systems. *Bulletin of Mathematical Biology*, 78(8): 1640–1677, 2016.
 - [49] A. Yu. Mikhalev and I. V. Oseledets. Rectangular maximum-volume submatrices and their applications. *Linear Algebra Appl.*, 538:187–211, 2018.
 - [50] B. Munsky and M. Khammash. The finite state projection algorithm for the solution of the chemical master equation. *The Journal of chemical physics*, 124:044104, 2006.
 - [51] Larry Neal and George Poole. A geometric analysis of Gaussian elimination. II. *Linear Alg. Appl.*, 173:239–264, 1992.
 - [52] I. V. Oseledets. Tensor-train decomposition. *SIAM J. Sci. Comput.*, 33(5):2295–2317, 2011.
 - [53] I. V. Oseledets and E. E. Tyrtyshnikov. TT-cross approximation for multidimensional arrays. *Linear Algebra Appl.*, 432(1):70–88, 2010.
 - [54] Sergey Petrov and Nikolai Zamarashkin. Sufficient recovery conditions for noise-buried low rank tensors. Technical Report arXiv:2312.02088, arXiv, 2023.
 - [55] M. V. Rakhuba and I. V. Oseledets. Grid-based electronic structure calculations: the tensor

- decomposition approach. *J. Comput. Phys.*, 2016.
- [56] D. V. Savostyanov. Quasioptimality of maximum-volume cross interpolation of tensors. *Linear Algebra Appl.*, 458:217–244, 2014.
 - [57] J. Schneider. Error estimates for two-dimensional cross approximation. *J. Approx. Theory*, 162: 1685–1700, 2010.
 - [58] U. Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96–192, 2011.
 - [59] Konstantin Sozykin, Andrei Chertkov, Roman Schutski, Anh-Huy Phan, Andrzej S Cichocki, and Ivan Oseledets. TTOpt: A maximum volume quantized tensor train-based optimization and its application to reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 35, 2022, pages 26052–26065.
 - [60] Augustinas Sukys, Kaan Öcal, and Ramon Grima. Approximating solutions of the chemical master equation using neural networks. *iScience*, 25:105010, 2022.
 - [61] A. Sulimov, D. Kutov, I. Ilin, D. Zheltkov, E. Tyrtshnikov, and V. Sulimov. Supercomputer docking with a large number of degrees of freedom. *SAR and QSAR in Environmental Research*, 30(10): 733–749, 2019.
 - [62] Alexey V. Sulimov, Dmitry A. Zheltkov, Igor V. Oferkin, Danil C. Kutov, Ekaterina V. Katkova, Eugene E. Tyrtshnikov, and Vladimir B. Sulimov. Evaluation of the novel algorithm of flexible ligand docking with moveable target-protein atoms. *Computational and Structural Biotechnology Journal*, 15:275–285, 2017.
 - [63] Claudia Totzeck. Trends in consensus-based optimization. In *Active Particles, Volume 3: Advances in Theory, Models, and Applications*, Nicola Bellomo, José Antonio Carrillo, and Eitan Tadmor, editors, Springer, 2022, pages 201–226.
 - [64] Zhenguo Tu and Yong Lu. A robust stochastic genetic algorithm (StGA) for global numerical optimization. *IEEE Trans Evolutionary Computation*, 8(5):456–470, 2004.
 - [65] E. E. Tyrtshnikov. Incomplete cross approximation in the mosaic-skeleton method. *Computing*, 64 (4):367–380, 2000.
 - [66] N. G. van Kampen. *Stochastic processes in physics and chemistry*. North Holland, Amsterdam, 1981.
 - [67] S. Venkatraman and G.G. Yen. A generic framework for constrained optimization using genetic algorithms. *IEEE Trans Evolutionary Computation*, 9(4):424–435, 2005.
 - [68] Huy D. Vo and Roger B. Sidje. An adaptive solution to the chemical master equation using tensors. *The Journal of Chemical Physics*, 147(4):044102, 2017.
 - [69] J.-H. Wang, P. K. Hopke, T. M. Hanczewicz, and S. L. Zhang. Application of modified least squares regression to spectroscopic image analysis. *Analys Chim. Acta.*, 476:93–109, 2003.
 - [70] Thomas Weise. *Global optimization algorithms — theory and application*. Self-Published Thomas Weise, 2009.
 - [71] N. L. Zamarashkin and A. I. Osinsky. New accuracy estimates for pseudoskeleton approximations of matrices. *Doklady Mathematics*, 94(3):643–645, 2016.
 - [72] D. A. Zheltkov, I. V. Oferkin, E. V. Katkova, A. V. Sulimov, V. B. Sulimov, and E. E. Tyrtshnikov. TTDock: a docking method based on tensor train decompositions. *Numerical Methods and Programming*, 14(3):279–291, 2013.
 - [73] Dmitry Zheltkov and Eugene Tyrtshnikov. Global optimization based on TT-decomposition. *Russian Journal of Numerical Analysis and Mathematical Modelling*, 35(4):247–261, 2020.
 - [74] Dmitry A. Zheltkov and Alexander Osinsky. Global optimization algorithms using tensor trains. In *Large-Scale Scientific Computing*, Ivan Lirkov and Svetozar Margenov, editors, Cham, 2020, pages 197–202. Springer International Publishing.