

MULTIMARGINAL OPTIMAL TRANSPORT WITH A TREE-STRUCTURED COST AND THE SCHRÖDINGER BRIDGE PROBLEM*

ISABEL HAASLER[†], AXEL RINGH[‡], YONGXIN CHEN[§], AND JOHAN KARLSSON[†]

Abstract. The optimal transport problem has recently developed into a powerful framework for various applications in estimation and control. Many of the recent advances in the theory and application of optimal transport are based on regularizing the problem with an entropy term, which connects it to the Schrödinger bridge problem and thus to stochastic optimal control. Moreover, the entropy regularization makes the otherwise computationally demanding optimal transport problem feasible even for large scale settings. This has led to an accelerated development of optimal transport based methods in a broad range of fields. Many of these applications have an underlying graph structure, for instance, information fusion and tracking problems can be described by trees. In this work we consider multimarginal optimal transport problems with a cost function that decouples according to a tree structure. The entropy regularized multimarginal optimal transport problem can be viewed as a generalization of the Schrödinger bridge problem with the same tree-structure, and by utilizing these connections we extend the computational methods for the classical optimal transport problem in order to solve structured multimarginal optimal transport problems in an efficient manner. In particular, the algorithm requires only matrix-vector multiplications of relatively small dimensions. We show that the multimarginal regularization introduces less diffusion, compared to the commonly used pairwise regularization, and is therefore more suitable for many applications. Numerical examples illustrate this, and we finally apply the proposed framework for the tracking of an ensemble of indistinguishable agents.

Key words. multimarginal optimal transport, Schrödinger bridge, hidden Markov chain, graph signal processing, ensemble estimation

AMS subject classifications. 93E10, 60J10, 93E11, 60J22, 90C35

DOI. 10.1137/20M1320195

1. Introduction. An optimal transport problem is to find a transport plan that minimizes the cost of moving the mass of one distribution to another distribution [62]. Historically this problem has been important in economics and operations research, but as a result of recent progress in the area it has become a popular tool in a wide range of fields such as control theory [1, 6, 15, 35, 51, 64], signal processing [30, 43], computer vision [27, 59], and machine learning [2, 4, 52].

An extension to the standard optimal transport framework is multimarginal optimal transport [53], which seeks a transport plan between not only two, but several distributions. Early works on multimarginal optimal transport include [33, 55, 56]. In this work we consider multimarginal optimal transport problems with cost functions that decouple according to a tree structure. We refer to such a problem as a tree-

*Received by the editors February 19, 2020; accepted for publication (in revised form) March 12, 2021; published electronically July 6, 2021.

<https://doi.org/10.1137/20M1320195>

Funding: This work was supported by the Swedish Research Council (VR), grant 2014-5870, KTH Digital Futures, SJTU-KTH cooperation grant, the NSF under grants 1901599 and 1942523, and the Knut and Alice Wallenberg foundation under grant KAW 2018.0349.

[†]Division of Optimization and Systems Theory, Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden, SE—100 44 (haasler@kth.se, johan.karlsson@math.kth.se).

[‡]Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, China (eeringh@ust.hk).

[§]School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (yongchen@gatech.edu).

structured multimarginal optimal transport problem. This should not be confused with optimal transport problems on graphs as in [16, 18, 19], where the distributions are defined over the nodes of the graphs. Tree-structured cost functions generalize many structures that are commonly used in applications of optimal transport. For example, a path tree naturally appears in tracking and interpolation applications [9, 17, 59]. Similarly, star trees are used in barycenter problems, which occur, for instance, in information fusion applications [23, 29].

The optimal transport problem can be formulated as a linear program. However, in many practical applications the problem is too large to be solved directly. For the bimarginal case, these computational limitations have recently been alleviated by regularizing the problem with an entropy term [22]. The optimal solution to the regularized optimal transport problem can then be expressed in terms of dual variables, which can be efficiently found by an iterative scheme, called Sinkhorn iterations [58]. For the multimarginal case, although the Sinkhorn iterations can be generalized in a straightforward fashion [7], the complexity of the scheme increases dramatically with the number of marginals [47]. Thus Sinkhorn iterations alone are not sufficient to address many multimarginal problems. However, in specific settings, structures in the cost function can be exploited in order to derive computationally feasible methods, e.g., for Euler flows [7] and in tracking and information fusion applications [30].

Unregularized multimarginal optimal transport problems with transport cost that decouples according to a tree structure can be equivalently formulated as a sum of coupled pairwise optimal transport problems, as, for instance, tracking and interpolation problems [9, 17, 29] and barycenter problems [3, 7, 23, 59]. Typically these problems are solved using pairwise regularization (see, e.g., [7, section 3.2] and [8, 44, 46]). However, we have empirically observed in some applications that the multimarginal formulation yields favorable solutions compared to a corresponding pairwise optimal transport estimate [30]. One main contribution of this work is to develop a framework for solving tree-structured optimal transport problems with a multimarginal regularization. For these problems, we show that the Sinkhorn algorithm can be performed in an efficient way, requiring only successive matrix-vector multiplications of relatively small size compared to that of the original multimarginal problem. Thus we extend the computational results from [30] to general trees.

The entropy regularized formulation of optimal transport is connected to another classical topic, the Schrödinger bridge problem [13, 14, 45, 48, 50]. Schrödinger was interested in determining the most likely evolution of a particle cloud observed at two time instances, where the particle dynamics have deviated from the expected Brownian motion [57]. Schrödinger showed that this particle evolution can be characterized as the one, out of all the theoretically possible ones, that minimizes the relative entropy to the Wiener measure. This optimal solution may be found by solving a so-called Schrödinger system, which turns out to be tightly connected to the Sinkhorn iterations for the entropy regularized optimal transport problem [13, 48]. This framework has been used in robust and stochastic control problems [15, 63]. A version of the Schrödinger bridge problem that is discrete in both time and space can be formulated by modeling the evolutions of a number of particles as a Markov chain [34, 54]. In the infinite particle limit, the maximum likelihood solution of the Markov process can then be approximated by solving a relative entropy problem [16]. An analogous approach has recently been used to develop a framework for modelling ensemble flows on hidden Markov chains [37]. Other optimal transport based state estimation problems for a continuum of agents and in continuous time have been con-

sidered in [11, 12, 17]. For further work on ensemble controllability and observability see, e.g., [10, 67].

Based on [37], we extend the discrete Schrödinger bridge in [54] to trees. In particular, we derive a maximum likelihood estimate for Markov processes defined on the edges of a rooted directed tree. This leads to our second main result: Interestingly, it turns out that the solution to the entropy regularized tree-structured multimarginal optimal transport problem corresponds to the solution of the generalized Schrödinger bridge with the same tree-structure. This generalizes the established equivalence of the classical bimarginal entropy regularized optimal transport problem and the Schrödinger bridge problem [48]. Moreover, this also gives an additional motivation for using multimarginal optimal transport instead of the often used pairwise optimal transport for tree structured problems.

This paper is organized as follows: section 2 is an introduction to the problems of optimal transport and Schrödinger bridges. The main results are presented in section 3 and 4. In section 3 we define the tree-structured multimarginal optimal transport problem and provide an algorithm for efficiently solving its entropy regularized version. In section 4 we generalize the Schrödinger bridge problem to trees and show that it is equivalent to an entropy regularized multimarginal optimal transport problem on the same tree. Section 5 is on the discrepancy between the multimarginal and pairwise optimal transport formulations with tree structure. A larger numerical simulation is detailed in section 6, where the proposed framework is used to estimate ensemble flows from aggregate and incomplete measurements. Most of the proofs can be found in the appendix, but due to space limitations two of the proofs are deferred to the supplementary material in [38].

2. Background. In this section we summarize some background material on optimal transport and Schrödinger bridges. This is also used to set up the notation. To this end, first note that throughout we let $\exp(\cdot)$, $\log(\cdot)$, \odot , and $./$ denote the elementwise exponential, logarithm, multiplication, and division of vectors, matrices, and tensors, respectively. Moreover, \otimes denotes the outer product. By $\mathbf{1}$ we denote a column vector of ones, the size of which will be clear from the context.

2.1. Optimal transport. In this work we consider the discrete optimal transport problem. For its continuous counterpart see, e.g., [62]. Let the vectors $\mu_1, \mu_2 \in \mathbb{R}_+^n$ describe two nonnegative distributions with equal mass. The optimal transport problem is to find a mapping that moves the mass from μ_1 to μ_2 , while minimizing the total transport cost. Here the transport cost is defined in terms of a underlying cost matrix $C \in \mathbb{R}_+^{n \times n}$, where C_{i_1, i_2} denotes the cost of moving a unit mass from location i_1 to i_2 . Analogously, define a transport plan $M \in \mathbb{R}_+^{n \times n}$, where M_{i_1, i_2} describes the amount of mass that is moved from location i_1 to i_2 . An optimal transport plan from μ_1 to μ_2 is then a minimizing solution of

$$(2.1) \quad \begin{aligned} T(\mu_1, \mu_2) := & \underset{M \in \mathbb{R}_+^{n \times n}}{\text{minimize}} \text{trace}(C^T M) \\ & \text{subject to } M\mathbf{1} = \mu_1, \quad M^T \mathbf{1} = \mu_2. \end{aligned}$$

Multimarginal optimal transport extends the concept of the classical optimal transport problem (2.1) to the setting with a set of marginals μ_1, \dots, μ_J , where $J \geq 2$ [7, 30, 53]. In this setting, the transport cost and transport plan in (2.1) are described by J -mode tensors $\mathbf{C}, \mathbf{M} \in \mathbb{R}_+^{n \times n \times \dots \times n}$. For a tuple (i_1, \dots, i_J) , the value $\mathbf{C}_{i_1, \dots, i_J}$ denotes the transport cost for a unit mass corresponding to the tuple, and similarly

$\mathbf{M}_{i_1, \dots, i_J}$ represents the amount of transported mass associated with this tuple. The multimarginal optimal transport problem then reads

$$(2.2) \quad \begin{aligned} & \underset{\mathbf{M} \in \mathbb{R}_+^{n \times \dots \times n}}{\text{minimize}} \langle \mathbf{C}, \mathbf{M} \rangle \\ & \text{subject to } P_j(\mathbf{M}) = \mu_j \text{ for } j \in \Gamma, \end{aligned}$$

where $\langle \mathbf{C}, \mathbf{M} \rangle = \sum_{i_1, \dots, i_J} \mathbf{C}_{i_1, \dots, i_J} \mathbf{M}_{i_1, \dots, i_J}$ is the standard inner product, the projection on the j th marginal of \mathbf{M} , denoted by $P_j(\mathbf{M})$, is defined as

$$(2.3) \quad P_j(\mathbf{M})_{i_j} = \sum_{i_1, \dots, i_{j-1}, i_{j+1}, \dots, i_J} \mathbf{M}_{i_1, \dots, i_{j-1}, i_j, i_{j+1}, \dots, i_J},$$

and Γ denotes an index set corresponding to the given marginals. In the original multimarginal optimal transport formulation, constraints are typically given on all marginals, i.e., for the index set $\Gamma = \{1, 2, \dots, J\}$. However, in this work we consider the case where constraints typically are only imposed on a subset of marginals, i.e., $\Gamma \subset \{1, 2, \dots, J\}$. Note that the standard bimarginal optimal transport problem (2.1) is a special case of the multimarginal optimal transport problem (2.2), where $J = 2$ and $\Gamma = \{1, 2\}$.

2.2. Entropy regularized optimal transport. Although linear, the number of variables in the multimarginal optimal transport problem (2.2) is often too large to be solved directly. A popular approach for the bimarginal setting to bypass the size of the problem has been to add a regularizing entropy term to the objective. In theory the same approach can be used also for the multimarginal case.

DEFINITION 2.1 ([36, Chapter 4]). *Let p and q be two nonnegative vectors, matrices, or tensors of the same dimension. The normalized Kullback–Leibler (KL) divergence of p from q is defined as $H(p|q) := \sum_i (p_i \log(p_i/q_i) - p_i + q_i)$, where $0 \log 0$ is defined to be 0. Similarly, define $H(p) := H(p|\mathbf{1}) = \sum_i (p_i \log(p_i) - p_i + 1)$, which is effectively the negative of the entropy of p .*

Note that $H(p|q)$ is jointly convex over p, q . For a detailed description of the KL divergence see, e.g., [20, 21]. The entropy regularized multimarginal optimal transport problem is the convex problem

$$(2.4) \quad \begin{aligned} & \underset{\mathbf{M} \in \mathbb{R}_+^{n \times \dots \times n}}{\text{minimize}} \langle \mathbf{C}, \mathbf{M} \rangle + \epsilon H(\mathbf{M}) \\ & \text{subject to } P_j(\mathbf{M}) = \mu_j \text{ for } j \in \Gamma, \end{aligned}$$

where $\epsilon > 0$ is a regularization parameter. For the bimarginal case (2.1), where the cost and mass transport tensors are matrices, the entropy regularized problem reads

$$(2.5) \quad \begin{aligned} T_\epsilon(\mu_1, \mu_2) := & \underset{\mathbf{M} \in \mathbb{R}_+^{n \times n}}{\text{minimize}} \text{trace}(\mathbf{C}^T \mathbf{M}) + \epsilon H(\mathbf{M}) \\ & \text{subject to } \mathbf{M}\mathbf{1} = \mu_1, \quad \mathbf{M}^T \mathbf{1} = \mu_2. \end{aligned}$$

It is well established that the entropy regularized bimarginal optimal transport problem is connected to the Schrödinger bridge problem [13, 14, 48, 50], which is introduced in section 2.3. More importantly, from a computational perspective, the introduction of the entropy term in problem (2.5) allows for expressing the optimal solution \mathbf{M}

in terms of the Lagrange dual variables, which may be computed by Sinkhorn iterations [22]. This procedure can be generalized to the setting of multimarginal optimal transport [7].

In particular, for the multimarginal entropy regularized optimal transport problem (2.4) it can be shown that the optimal solution is of the form [30]

$$(2.6) \quad \mathbf{M} = \mathbf{K} \odot \mathbf{U},$$

where $\mathbf{K} = \exp(-\mathbf{C}/\epsilon)$ and where \mathbf{U} can be decomposed as

$$(2.7) \quad \mathbf{U} = u_1 \otimes u_2 \otimes \cdots \otimes u_J \quad \text{with } u_j = \begin{cases} \exp(\lambda_j/\epsilon) & \text{if } j \in \Gamma, \\ 1 & \text{else.} \end{cases}$$

Here $\lambda_j \in \mathbb{R}^n$ for $j \in \Gamma$ are optimal dual variables in the dual problem of (2.4):

$$(2.8) \quad \underset{\lambda_j \in \mathbb{R}^n, j \in \Gamma}{\text{maximize}} \quad -\epsilon \langle \mathbf{K}, \mathbf{U} \rangle + \sum_{j \in \Gamma} \lambda_j^T \mu_j,$$

where \mathbf{U} depends on the variables λ_j as specified by (2.7). Note that the dual variable λ_j corresponds to the constraint on the j th marginal. For details the reader is referred to, e.g., [7, 30].

The Sinkhorn scheme for finding \mathbf{U} in (2.7) is to iteratively update u_j as

$$(2.9) \quad u_j \leftarrow u_j \odot \mu_j ./ P_j(\mathbf{K} \odot \mathbf{U})$$

for all $j \in \Gamma$. This scheme may, for instance, be derived as Bregman projections [7] or a block coordinate ascent in the dual (2.8) [30, 42, 61]. As a result, global convergence of the Sinkhorn scheme (2.9) is guaranteed [5]. For the sake of completeness, we also provide a result on the linear convergence rate in our presentation (see Theorem 3.5). We also note that (2.9) reduces to standard Sinkhorn iterations,

$$(2.10) \quad u_1 \leftarrow \mu_1 ./ (K u_2), \quad u_2 \leftarrow \mu_2 ./ (K^T u_1)$$

for the two-marginal case (2.5). The iterations (2.10) converge linearly to an optimal solution u_1, u_2 , which is unique up to multiplication/division with a constant [13, 32].

The computational bottleneck of the Sinkhorn iterations (2.9) is computing the projections $P_j(\mathbf{M})$ for $j \in \Gamma$, which in general scales exponentially in J . In fact, even storing the tensor \mathbf{M} is a challenge as it consists of n^J elements. However, in many cases of interest, structures in the cost tensors can be exploited to make the computation of the projections feasible. In [30] this is shown on the example of cost functions that decouple sequentially, centrally, or a combination of both. Computing the projections requires then only repeated matrix vector multiplications. In this work, we show that these efficient methods can be generalized to the setting where the cost tensor decouples according to a tree structure, that is, when the marginals of the optimal transport problem are associated with the nodes of a tree, and cost matrices are defined on its edges (see section 3).

2.3. Schrödinger bridge problem. The Schrödinger bridge problem is to determine the most likely evolution of a particle cloud observed at two time instances [57]. In the case that the second observation cannot be explained as a Brownian motion of the initially observed particle cloud, Schrödinger aimed to find the most likely particle evolution connecting, hence bridging, the two distributions.

A useful mathematical framework to solve this problem, however not yet developed in Schrödinger's time, is the theory of large deviations, which studies so-called rare events, meaning deviations from the law of large numbers [24, 25]. The probability of these rare events approaches zero, as the number of trials goes to infinity, and large deviation theory analyzes the rate of this decay, which can often be expressed as the exponential of a so-called rate function [25, 28].

For a large deviation interpretation of the Schrödinger bridge [31, section II.1.3], the particle evolutions are modelled as independent identically distributed random variables on path space, and the Schrödinger bridge is the probability measure \mathcal{P} on path space that is most likely to describe the rare event of observing the two particle distributions. Let \mathcal{W} be the Wiener measure, which corresponds to the probability law of the Brownian motion. Then the Schrödinger bridge \mathcal{P} is found by minimizing the corresponding rate function $\int \log(d\mathcal{P}/d\mathcal{W})d\mathcal{P}$ over all probability measures that are absolutely continuous with respect to \mathcal{W} and have the given particle cloud distributions as marginals. A space and time discrete Schrödinger bridge problem for Markov chains is treated in [34, 54].

We follow the exposition in [37], where the space and time discrete Schrödinger bridge has been derived as a maximum likelihood approximation for Markov chains. Consider a cloud of N particles and assume that each particle evolves according to a Markov chain. Denote the states of the Markov chain by $X = \{X_1, X_2, \dots, X_n\}$ and let the transition probability matrix be $A^j \in \mathbb{R}_+^{n \times n}$ with elements $A_{k\ell}^j = P(q_{j+1} = X_\ell | q_j = X_k)$, where q_j denotes the state at time j . Let the vector μ_j describe the particle distribution over the discrete state space X at time j for $j = 1, \dots, J$. As in the optimal transport framework we define the mass transport matrix M^j where element $M_{k\ell}^j$ describes the number of particles transitioning from state k at time j to state ℓ at time $j + 1$.

Given the distribution μ_j and transition probability matrix A^j , a large deviation argument similar to the continuous setting can be used to approximate the most likely transfer matrix M^j by the minimizer of the rate function $H(\cdot | \text{diag}(\mu_j)A^j)$. Assuming that the initial and final marginal, μ_1 and μ_J , are given, these large deviation theoretic considerations motivate the formulation of an optimization problem to find the most likely mass transfer matrices M^j for $j = 1, \dots, J - 1$ and intermediate distributions μ_j for $j = 2, \dots, J - 1$, as

$$(2.11) \quad \begin{aligned} & \underset{M_{[1:J-1]}, \mu_{[2:J-1]}}{\text{minimize}} && \sum_{j=1}^{J-1} H(M^j | \text{diag}(\mu_j)A^j) \\ & \text{subject to} && M^j \mathbf{1} = \mu_j, \quad (M^j)^T \mathbf{1} = \mu_{j+1} \quad \text{for } j = 1, \dots, J-1. \end{aligned}$$

Indeed, this problem is equivalent to the discrete time and space Schrödinger bridge problem in [54] (for details see [37] or section 5.1).

An extension to a hidden Markov model formulation with aggregate indirect observations of the distributions was described in [37]. In this article the framework is extended to general tree structures, where each vertex is associated with a distribution, and each edge with a Markov process.

3. Tree-structured multimarginal optimal transport. In this section we introduce tree-structured multimarginal optimal transport problems. For the entropy regularized version of these problems, the projections (2.3) can be computed by only matrix-vector multiplications. This yields an efficient Sinkhorn algorithm, which we present in this section. Moreover, some properties of tree-structured multimarginal optimal transport problems are discussed.

DEFINITION 3.1. A graph $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, with vertices \mathcal{V} and edges \mathcal{E} , is a tree if it is acyclic and connected [26]. The vertices with degree 1 are called leaves, and we denote the set of leaves by \mathcal{L} . For a vertex $j \in \mathcal{V}$, the set of neighbors \mathcal{N}_j is defined as the set of vertices that have a common edge with j . The path between two vertices j_1 and j_L is a sequence of edges that connect j_1 and j_L , such that all edges are distinct. We also denote a path by the set of intermediate vertices.

Let $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ be a tree. In this section we consider a multimarginal optimal transport problem where the marginals correspond to the nodes of the tree, i.e., $\mathcal{V} = \{1, 2, \dots, J\}$. Assume that the cost tensor decouples as

$$(3.1) \quad \mathbf{C}_{i_1, \dots, i_J} = \sum_{(j_1, j_2) \in \mathcal{E}} C_{i_{j_1}, i_{j_2}}^{(j_1, j_2)},$$

where $C^{(j_1, j_2)} \in \mathbb{R}_+^{n \times n}$ is a cost matrix characterizing the cost of transportation between marginal j_1 and j_2 for all $(j_1, j_2) \in \mathcal{E}$. Since we consider an undirected tree, the expression (3.1) should not depend on the ordering of the indices (j_1, j_2) in the edges. This is achieved by simply letting $C^{(j_2, j_1)} = (C^{(j_1, j_2)})^T$ for $(j_1, j_2) \in \mathcal{E}$. We refer to problem (2.2) with a cost of the form (3.1) as a multimarginal optimal transport problem on the tree \mathcal{T} .

In the following, we consider discretized and entropy regularized multimarginal optimal transport with tree structure. The transport plan can then be expressed as $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$, where $\mathbf{K} = \exp(-\mathbf{C}/\epsilon)$ (cf. (2.6)). Due to the structured cost (3.1), this tensor decouples as

$$(3.2) \quad \mathbf{K}_{i_1, \dots, i_J} = \prod_{(j_1, j_2) \in \mathcal{E}} K_{i_{j_1}, i_{j_2}}^{(j_1, j_2)}$$

with the matrices defined as $K_{i_{j_1}, i_{j_2}}^{(j_1, j_2)} = \exp(-C_{i_{j_1}, i_{j_2}}^{(j_1, j_2)}/\epsilon)$ for $(j_1, j_2) \in \mathcal{E}$. Thus, it holds that $K^{(j_2, j_1)} = (K^{(j_1, j_2)})^T$ for $(j_1, j_2) \in \mathcal{E}$. We show that for these problems the projections (2.3) can therefore be computed by successive matrix vector multiplications. As the computation of the projections accounts for the bottleneck of the Sinkhorn iterations (2.9), this yields an efficient algorithm for solving entropy regularized multimarginal optimal transport problems with tree structure.

We first illustrate the computation of the projections on a small example.

Example 1. Consider the tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ with vertices $\mathcal{V} = \{1, 2, 3, 4\}$ and edges $\mathcal{E} = \{(1, 2), (2, 3), (1, 4)\}$ as depicted in Figure 1. Assume that the cost in problem (2.4) decouples as $\mathbf{C}_{i_1, i_2, i_3, i_4} = C_{i_1, i_2}^{(1, 2)} + C_{i_2, i_3}^{(2, 3)} + C_{i_1, i_4}^{(1, 4)}$, where $C^{(j_1, j_2)}$, for $(j_1, j_2) \in \mathcal{E}$, are cost matrices defined on the respective edges. The transport tensor is then of the form $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$, where $\mathbf{K} = \exp(-\mathbf{C}/\epsilon)$ and $\mathbf{U} = u_1 \otimes u_2 \otimes u_3 \otimes u_4$. Thus, denoting

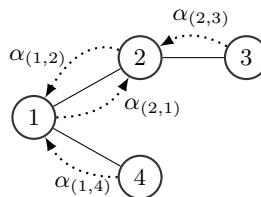


FIG. 1. Illustration of the tree in Example 1.

$K^{(j_1, j_2)} = \exp(-C^{(j_1, j_2)}/\epsilon)$ it holds that $\mathbf{K}_{i_1, i_2, i_3, i_4} = K_{i_1, i_2}^{(1,2)} K_{i_2, i_3}^{(2,3)} K_{i_1, i_4}^{(1,4)}$. Hence, the projection on the first marginal of $\mathbf{K} \odot \mathbf{U}$ can be computed as

$$P_1(\mathbf{K} \odot \mathbf{U})_{i_1} = \sum_{i_2, i_3, i_4} \mathbf{K}_{i_1, i_2, i_3, i_4} \mathbf{U}_{i_1, i_2, i_3, i_4} = (u_1)_{i_1} \sum_{i_2, i_3, i_4} K_{i_1, i_2}^{(1,2)} K_{i_2, i_3}^{(2,3)} K_{i_1, i_4}^{(1,4)} (u_2)_{i_2} (u_3)_{i_3} (u_4)_{i_4}.$$

The sum can be computed as successive matrix-vector multiplications, starting from the leaves of the tree. In particular, denote the products starting in the two leaves

$$\sum_{i_3} K_{i_2, i_3}^{(2,3)} (u_3)_{i_3} = \left(K^{(2,3)} u_3 \right)_{i_2} =: (\alpha_{(2,3)})_{i_2}, \quad \sum_{i_4} K_{i_1, i_4}^{(1,4)} (u_4)_{i_4} = \left(K^{(1,4)} u_4 \right)_{i_1} =: (\alpha_{(1,4)})_{i_1}.$$

On the lower branch of the tree, we have thereby brought the expression to a vector indexed by i_1 . On the upper branch, another multiplication, corresponding to the edge $(1, 2)$, is required. This leads to defining

$$\sum_{i_2} K_{i_1, i_2}^{(1,2)} (u_2 \odot \alpha_{(2,3)})_{i_2} = \left(K^{(1,2)} (u_2 \odot \alpha_{(2,3)}) \right)_{i_1} =: (\alpha_{(1,2)})_{i_1}.$$

The full expression for the projection thus reads

$$P_1(\mathbf{K} \odot \mathbf{U}) = u_1 \odot K^{(1,2)} (u_2 \odot K^{(2,3)} u_3) \odot K^{(1,4)} u_4 = u_1 \odot \alpha_{(1,2)} \odot \alpha_{(1,4)}.$$

This thus allows for decomposing the computation of $P_1(\mathbf{K} \odot \mathbf{U})$ into the three parts $u_1, \alpha_{(1,2)}, \alpha_{(1,4)}$, which represent the contributions from node 1; from nodes 2 and 3; and from node 4, respectively. Analogously, the projection on the second marginal is computed as $P_2(\mathbf{K} \odot \mathbf{U}) = u_2 \odot \alpha_{(2,3)} \odot \alpha_{(2,1)}$, where $\alpha_{(2,1)} := K^{(2,1)}(u_1 \odot \alpha_{(1,4)})$. In particular, note here that the tuple (j, k) in $\alpha_{(j,k)}$ is ordered. For instance, it holds that $\alpha_{(1,2)} \neq \alpha_{(2,1)}$. The other marginals can be computed similarly by performing the corresponding matrix vector products, starting from the leaves of the tree.

Illustratively, the procedure in Example 1 for projecting the information of the full tensor down to one marginal can be understood as passing down the information from all branches of the tree to the desired vertex, where we interpret the vectors u_j as local information in each node. Starting from the leaves of the tree, this is done according to the following rules:

1. Information is passed down the edge $(j_1, j_2) \in \mathcal{E}$ by multiplication with the matrix $K^{(j_1, j_2)}$.
2. Information is collected in the node $j \in \mathcal{V}$ by elementwise multiplication of the vector u_j , which contains local information, with information passed down from the connected branches.

In this interpretation the vector $\alpha_{(j,k)}$ can be understood as the collected information that is sent from node k to node j . The following theorem shows that these rules hold for any tensor of the form $\mathbf{K} \odot \mathbf{U}$ with $\mathbf{U} = u_1 \otimes \cdots \otimes u_j$ and \mathbf{K} decoupling according to an arbitrary tree.

THEOREM 3.2. *Let $\mathbf{K} = \exp(-\mathbf{C}/\epsilon)$ with \mathbf{C} as in (3.1) for the tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, and let $\mathbf{U} = u_1 \otimes u_2 \otimes \cdots \otimes u_J$. Define $K^{(j_1, j_2)} = \exp(-C^{(j_1, j_2)}/\epsilon)$ for $(j_1, j_2) \in \mathcal{E}$. Then the projection on the j th marginal of $\mathbf{K} \odot \mathbf{U}$ is of the form*

$$P_j(\mathbf{K} \odot \mathbf{U}) = u_j \odot \bigodot_{k \in \mathcal{N}_j} \alpha_{(j,k)}.$$

The vectors $\alpha_{(j,k)}$, for all ordered tuples $(j, k) \in \mathcal{E}$, can be computed recursively starting

in the leaves of the tree according to

$$(3.3) \quad \begin{aligned} \alpha_{(j,k)} &= K^{(j,k)} u_k && \text{for } k \in \mathcal{L}, \\ \alpha_{(j,k)} &= K^{(j,k)} \left(u_k \odot \bigodot_{\ell \in \mathcal{N}_k \setminus \{j\}} \alpha_{(k,\ell)} \right) && \text{for } k \notin \mathcal{L}. \end{aligned}$$

Proof. See the appendix. \square

Analogously to the marginal projections (2.3), let $P_{j_1,j_2}(\mathbf{M})$ denote the bimarginal projections on marginals j_1 and j_2 as $P_{j_1,j_2}(\mathbf{M})_{i_{j_1},i_{j_2}} = \sum_{i_1, \dots, i_J \setminus \{i_{j_1}, i_{j_2}\}} \mathbf{M}_{i_1, \dots, i_J}$. These pairwise projections can be expressed similarly to the marginal projections in Theorem 3.2.

PROPOSITION 3.3. *Let the assumptions in Theorem 3.2 hold, let $j_1, j_L \in \mathcal{V}$, and let j_1, j_2, \dots, j_L denote the path between j_1 and j_L . Then the pairwise projection of $\mathbf{K} \odot \mathbf{U}$ on the marginals j_1 and j_L , denoted by $P_{j_1,j_L}(\mathbf{K} \odot \mathbf{U})$, is*

$$(3.4) \quad \left(\prod_{\ell=1}^{L-1} \text{diag} \left(u_{j_\ell} \odot \bigodot_{\substack{k \in \mathcal{N}_{j_\ell} \\ k \notin \{j_1, \dots, j_L\}}} \alpha_{(j_\ell,k)} \right) K^{(j_\ell, j_{\ell+1})} \right) \text{diag} \left(u_{j_L} \odot \bigodot_{k \in \mathcal{N}_{j_L} \setminus \{j_{L-1}\}} \alpha_{(j_L,k)} \right),$$

where $\alpha_{(j_1,j_2)}$, for $(j_1, j_2) \in \mathcal{E}$, are defined as in Theorem 3.2.

Proof. See the supplementary material [38]. \square

Multimarginal optimal transport problems on some specific trees have previously been introduced in [30]. It is worth noting that the expressions for the projections of the mass tensor \mathbf{M} for the examples in [30] are special cases of the results in Theorem 3.2 and Proposition 3.3. For instance, in tracking problems each marginal in the optimal transport problem is associated with a time instance, and the corresponding graph is a path graph, as in [30, sections 3.1 and 5.1]. Sensor fusion applications can be cast as barycenter problems, which are described by star-shaped graphs, as in [30, sections 3.2 and 5.2]. Moreover, a combination of these two applications is the tracking of barycenters over time, which is described by a graph as in Figure 6(a) (see [30, sections 3.3 and 5.3]).

The expressions for the projections in Theorem 3.2 can be used to solve a multimarginal optimal transport problem, with cost structure according to the tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, by a Sinkhorn method as in (2.9). To this end, without loss of generality, we consider tree-structured multimarginal optimal transport problems (2.2), where the constraints are given on the set of leaves, i.e., $\Gamma = \mathcal{L}$. The following proposition shows that if a marginal is given for a node that is not a leaf, then the problem can be decomposed into smaller problems on subtrees, where the marginals are known exactly on the set of leaves of the subtrees.

PROPOSITION 3.4. *Consider problem (2.4) with constraint set Γ and cost structure according to the tree \mathcal{T} with leaves \mathcal{L} . The following holds:*

1. *If there is a marginal $k \in \Gamma$, but $k \notin \mathcal{L}$, then the solution to (2.4) can be found by solving problems of form (2.4) on the subtrees of \mathcal{T} , which are obtained by cutting \mathcal{T} in node k .*
2. *If there is a marginal $\ell \in \mathcal{L}$, but $\ell \notin \Gamma$, then the solution to (2.4) can be found by solving the problem on the subtree of \mathcal{T} that is obtained by removing the node ℓ and its adjoining edge from \mathcal{T} .*

Proof. See the appendix. \square

Algorithm 3.1 Sinkhorn method for the tree-structured multimarginal optimal transport problem.

Given: Tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ with leaves $\mathcal{L} = \{1, 2, \dots, |\mathcal{L}|\}$;
 Initial guess u_j for $j \in \mathcal{L}$
for all ordered tuples $(j_1, j_2) \in \mathcal{E}$ **do**
 Initialize $\alpha_{(j_1, j_2)}$ according to (3.3)
end for
 Initialize $j \in \mathcal{L}$
while Sinkhorn not converged **do**
 $u_j \leftarrow \mu_j / \bigodot_{k \in \mathcal{N}_j} \alpha_{(j, k)}$
 for $(j_1, j_2) \in \mathcal{E}$ on the path from j to $(j + 1 \bmod |\mathcal{L}|)$ **do**
 Update $\alpha_{(j_1, j_2)}$ according to (3.3)
 end for
 $j \leftarrow j + 1 \bmod |\mathcal{L}|$
end while
return u_j for $j \in \mathcal{L}$

Note that since $\Gamma = \mathcal{L}$, the factors u_j , for $j \in \mathcal{V} \setminus \mathcal{L}$, can be neglected in Theorem 3.2 and Proposition 3.3. Moreover, when applying the iterative scheme (2.9), in each iteration step some of the factors $\alpha_{(j_1, j_2)}$ do not change. Specifically, between the update of u_{j_1} and u_{j_2} only the factors on all edges that lie on the path between nodes j_1 and j_2 need to be updated. The full method for solving tree-structured multimarginal optimal transport problems is summarized in Algorithm 3.1. Moreover, the algorithm has in fact linear convergence speed.

THEOREM 3.5. *Let $\{u_j^k\}_{j \in \mathcal{L}}$ be the set of vectors after the k th iteration in the while-loop in Algorithm 3.1. Then the sequence $\{u_j^k\}_{j \in \mathcal{L}}$ converges at least linearly to an optimal solution of (2.8), as $k \rightarrow \infty$.*

Proof. Algorithm 3.1 implements the Sinkhorn iterations (2.9), which are a block coordinate ascent in the dual multimarginal optimal transport problem (2.8) [30]. Thus, the result follows from [49, Application 5.3 and Theorem 2.1]. \square

It should be noted that computing the marginals $P_j(\mathbf{K} \odot \mathbf{U})$ by summing over the indices as defined in (2.3) scales exponentially in the number of marginals J of the tensor $\mathbf{K} \odot \mathbf{U}$. Thus, the complexity of one Sinkhorn update in (2.9) is in general of the order $\mathcal{O}(n^J)$. Algorithm 3.1 utilizes the result in Theorem 3.2 in order to exploit the graph-structure in the tensor $\mathbf{K} \odot \mathbf{U}$ to compute the projections. This decreases the computational complexity of the Sinkhorn iterations substantially. In particular, if the direct path between node $j - 1$ and node j is of length p , then the update of the vector u_j requires p matrix-vector multiplications. The complexity of one update is thus of the order $\mathcal{O}(pn^2)$. As a rule of thumb the leaves in the underlying tree \mathcal{T} should thus be labeled such that the paths between any two nodes that are updated successively are short. In case there are additional structures in the cost matrices $C^{(j_1, j_2)}$, for $(j_1, j_2) \in \mathcal{E}$, e.g., they describe the squared Euclidean distance, the cost of the matrix-vector multiplications can be decreased further (see, e.g., [30, Remark 4], [42, Remark 3.11]).

Note that Proposition 3.4.1 implies that if the tensor \mathbf{K} is normalized to be a discrete probability distribution, it defines a Markov random field (see also [39]). Every Markov process is also a reciprocal process [40], which were first introduced in the con-

text of studying Schrödinger bridges [41]. Recall that the maximum likelihood estimation problem for a Markov chain in section 2.3 is equivalent to a discrete Schrödinger bridge. This suggests a connection between the multimarginal optimal transport problem and the maximum likelihood estimation problem in section 2.3. In fact, under certain conditions a generalization of the Markov chain problem (2.11) to trees, which is introduced in section 4.1, yields the same solution as the regularized multimarginal optimal transport problem (2.4) on the same tree, as we will see in section 4.2.

4. Multimarginal optimal transport vs. Schrödinger bridge with tree-structure. In this section, we extend the Schrödinger bridge problem, which is defined on a path tree, to a formulation valid on an arbitrary tree graph. We then show that under certain conditions this problem yields a solution which is equivalent to the solution of the entropy regularized multimarginal optimal transport problem (2.4) on the same tree, thus extending the well-known equivalence from the optimal transport problem and the Schrödinger bridge problem on a path tree [13, 14, 48, 50].

4.1. Tree-structured Schrödinger bridges. The notion of time of the Markov chain introduces a direction to the Schrödinger bridge problem, which can be seen directly in the optimization problem (2.11). This directionality needs to be taken into account when extending the problem to a general tree. To this end, we introduce the following notation: Consider a tree $\mathcal{T}_r = (\mathcal{V}, \mathcal{E}_r)$ that is rooted in a vertex $r \in \mathcal{L}$, i.e., one of the leaves is defined to be the root of the tree. This defines a partial ordering on the tree, and we write $k < j$ if node k lies on the path between r and node j . To formulate the following results, we assume that all edges are directed according to this partial ordering, i.e., $j_1 < j_2$ for all $(j_1, j_2) \in \mathcal{E}_r$. For a vertex $j \in \mathcal{V} \setminus \{r\}$, its parent is then defined as the (unique) vertex $p(j)$ such that $(p(j), j) \in \mathcal{E}_r$. In the following, without loss of generality we denote the root vertex by $r = 1$.

Let $A^{(j_1, j_2)}$ be the probability transition matrix on edge $(j_1, j_2) \in \mathcal{E}_r$. Then, the Schrödinger bridge problem in (2.11) may be naturally extended to the tree structure as

$$(4.1) \quad \begin{aligned} & \underset{\substack{M^{(j_1, j_2)}, (j_1, j_2) \in \mathcal{E}_r, \\ \mu_j, j \in \mathcal{V} \setminus \Gamma}}{\text{minimize}} && \sum_{(j_1, j_2) \in \mathcal{E}_r} H\left(M^{(j_1, j_2)} \mid \text{diag}(\mu_{j_1}) A^{(j_1, j_2)}\right) \\ & \text{subject to} && M^{(j_1, j_2)} \mathbf{1} = \mu_{j_1}, \quad (M^{(j_1, j_2)})^T \mathbf{1} = \mu_{j_2} \quad \text{for } (j_1, j_2) \in \mathcal{E}_r. \end{aligned}$$

Moreover, we can assume without loss of generality that $\Gamma = \mathcal{L}$, by similar reasoning as in section 3. The extension of (2.11) to (4.1) builds on the large deviation principle in [37, Proposition 1], which requires that an initial distribution is known. This is why we restrict our analysis to the case when the directed tree $\mathcal{T}_r(\mathcal{V}, \mathcal{E}_r)$ is rooted in a leaf.

Similarly to the multimarginal optimal transport problem, the optimal solution to (4.1) can be expressed in terms of its dual variables. In fact this is a natural generalization of the forward and backward propagation factors in the Schrödinger system [34, 54] to tree graphs.

THEOREM 4.1. *Let $A^{(j_1, j_2)}$ be probability transition matrices for $(j_1, j_2) \in \mathcal{E}_r$ with strictly positive elements and assume that $\mathbf{1}^T \mu_{j_1} = \mathbf{1}^T \mu_{j_2}$ for all leaves $j_1, j_2 \in \mathcal{L}$. Then an optimal solution to (4.1) can be written as*

$$\begin{aligned} \mu_j &= \varphi_j \odot \hat{\varphi}_j, \\ M^{(j_1, j_2)} &= \text{diag}(\hat{\varphi}_{j_1} \odot \varphi_{j_1 \setminus j_2}) A^{(j_1, j_2)} \text{diag}(\varphi_{j_2}) \end{aligned}$$

for a set of vectors $\varphi_j, \hat{\varphi}_j$, for $j \in \mathcal{V}$, and $\varphi_{j_1 \setminus j_2}$ for $(j_1, j_2) \in \mathcal{E}_r$.

Moreover, there exists a set of vectors v_j for $j \in \Gamma$ such that the vectors φ_j , $\hat{\varphi}_j$, and $\varphi_{j_1 \setminus j_2}$ can be written as

$$(4.2) \quad \varphi_j = \begin{cases} v_j & \text{if } j \in \mathcal{L} \setminus \{1\} \\ \bigodot_{k:(j,k) \in \mathcal{E}_r} A^{(j,k)} \varphi_k & \text{else,} \end{cases}$$

$$(4.3) \quad \hat{\varphi}_j = \begin{cases} \mathbf{1}./v_1 & \text{if } j = 1 \\ (A^{(p(j),j)})^T (\hat{\varphi}_{p(j)} \odot \varphi_{p(j) \setminus j}) & \text{else,} \end{cases}$$

and $\varphi_{j_1 \setminus j_2} = \varphi_{j_1} ./ (A^{(j_1,j_2)} \varphi_{j_2})$.

Proof. See the appendix. \square

Note that the vectors v_j for $j \in \mathcal{L}$ in Theorem 4.1 satisfy the nonlinear equations $\varphi_1./v_1 = \mu_1$ and $v_j \odot \hat{\varphi}_j = \mu_j$ for $j \in \mathcal{L} \setminus \{1\}$. When considering the Schrödinger bridge, which is the special case with two leaves, these equations correspond to the boundary conditions in the Schrödinger system [54, Theorem 4.1]. The factors φ_j and $\hat{\varphi}_j$ can be seen as propagating information from the leaves to the other nodes in a similar manner as $\alpha_{(j,k)}$ in Theorem 3.2.

Next, we show that the tree-structured Schrödinger bridge problem (4.1) is independent of the choice of root $r \in \mathcal{L}$. This follows as a corollary to the next proposition.

PROPOSITION 4.2. *Let $\mathcal{T}_r = (\mathcal{V}, \mathcal{E}_r)$ be a rooted directed tree with root $r \in \mathcal{L}$. Then problem (4.1) is equivalent to the problem*

$$(4.4) \quad \begin{aligned} & \underset{\substack{M^{(j_1,j_2)}, (j_1,j_2) \in \mathcal{E}_r, \\ \mu_j, j \in \mathcal{V} \setminus \Gamma}}{\text{minimize}} \quad \sum_{(j_1,j_2) \in \mathcal{E}_r} H\left(M^{(j_1,j_2)} \mid A^{(j_1,j_2)}\right) - \sum_{j \in \mathcal{V} \setminus \mathcal{L}} (\deg(j) - 1) H(\mu_j) \\ & \text{subject to} \quad M^{(j_1,j_2)} \mathbf{1} = \mu_{j_1}, \quad (M^{(j_1,j_2)})^T \mathbf{1} = \mu_{j_2} \quad \text{for } (j_1, j_2) \in \mathcal{E}_r. \end{aligned}$$

Proof. See the appendix. \square

COROLLARY 4.3. *Let $\mathcal{T}_r = (\mathcal{V}, \mathcal{E}_r)$ be a directed tree with root $r \in \mathcal{L}$ and $\mathcal{T}_{\hat{r}} = (\mathcal{V}, \mathcal{E}_{\hat{r}})$ be a directed tree with root $\hat{r} \in \mathcal{L}$ with the same structure as \mathcal{T}_r and the edges on the path between r and \hat{r} reversed. Then, the solution to (4.1) on \mathcal{T}_r and on $\mathcal{T}_{\hat{r}}$ are equivalent in the sense that the marginal distributions μ_j for $j \in \mathcal{V}$ and transport plans $M^{(j_1,j_2)}$ for $(j_1, j_2) \in \mathcal{E}_r \cap \mathcal{E}_{\hat{r}}$ are the same, and on the reversed edges it holds that $M^{(j_2,j_1)} = (M^{(j_1,j_2)})^T$.*

Proof. See the supplementary material [38]. \square

4.2. Equivalence of problems (2.4) and (4.1) for tree-structures. We will now verify that the generalized Schrödinger bridge (4.1) on a rooted tree \mathcal{T}_r is equivalent to an entropy regularized multimarginal optimal transport problem (2.4) on an undirected tree \mathcal{T} with the same structure as \mathcal{T}_r . In particular, given positive transition probability matrices $A^{(j_1,j_2)}$ for $(j_1, j_2) \in \mathcal{E}_r$ and a regularization parameter ϵ , there is a natural choice of cost matrices $C^{(j_1,j_2)}$ for $(j_1, j_2) \in \mathcal{E}$, so that the minimizers to both problems represent the same solution. For the original bimarginal optimal transport problem (2.1) the equivalence between entropy regularized optimal transport and the Schrödinger bridge problem has been studied extensively [13, 14, 48, 50].

Remark 1 (cf. [37, Remark 1]). In the discrete setting, it is easy to see that with a cost matrix defined as $C = -\epsilon \log(A)$, the entropy regularized optimal transport

problem (2.5) is equivalent to the bimarginal Schrödinger bridge (2.11), by noting that the objective can then be written as

$$\text{trace}(C^T M) + \epsilon H(M) = \sum_{i,j=1}^n \epsilon \left(M_{ij} \log \left(\frac{M_{ij}}{A_{ij}} \right) - M_{ij} + 1 \right) = \epsilon H(M | A).$$

Both bimarginal problems thus yield the same optimal solution M .

We will show that this equivalence holds true even for the multimarginal case with tree-structured cost.

THEOREM 4.4. *Let $\mathcal{T}_r = (\mathcal{V}, \mathcal{E}_r)$ be a rooted directed tree with root $r = 1 \in \mathcal{L}$, let $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ be its undirected counterpart, and let μ_j for $j \in \mathcal{L}$ be given marginals. Let $\epsilon > 0$ and $C^{(j_1, j_2)}$ be such that*

$$K^{(j_1, j_2)} = A^{(j_1, j_2)} \quad \text{for } (j_1, j_2) \in \mathcal{E}_r.$$

Let $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$ be the solution to the entropy regularized multimarginal optimal transport problem (2.4) on \mathcal{T} , where \mathbf{K} and \mathbf{U} are defined as in Theorem 3.2. Moreover, let μ_j for $j \in \mathcal{V}$, and $M^{(j_1, j_2)}$ for $(j_1, j_2) \in \mathcal{E}_r$, be the solution to the generalized Schrödinger bridge (4.1) on \mathcal{T}_r . Then,

$$\begin{aligned} P_j(\mathbf{M}) &= \mu_j && \text{for } j \in \mathcal{V}, \\ P_{(j_1, j_2)}(\mathbf{M}) &= M^{(j_1, j_2)} && \text{for } (j_1, j_2) \in \mathcal{E}_r. \end{aligned}$$

Furthermore, if v_j for $j \in \mathcal{L}$ are the corresponding vectors in Theorem 4.1, then it holds that

$$u_j = \begin{cases} \mathbf{1}./v_1 & \text{if } j = 1, \\ v_j & \text{otherwise.} \end{cases}$$

Proof. See the appendix. □

In particular, given a maximum likelihood problem (4.1) structured according to the tree $\mathcal{T}_r = (\mathcal{V}, \mathcal{E}_r)$, a corresponding multimarginal optimal transport problem can be formulated by defining the cost matrices as $C^{(j_1, j_2)} = -\epsilon \log(A^{(j_1, j_2)})$ for all $(j_1, j_2) \in \mathcal{E}_r$. Vice versa, given an entropy regularized multimarginal optimal transport problem (2.4) with cost structured according to the tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, a corresponding maximum likelihood problem can be formulated by transforming the undirected tree into a rooted tree $\mathcal{T}_r = (\mathcal{V}, \mathcal{E}_r)$, where $r \in \mathcal{L}$, and defining the matrices $A^{(j_1, j_2)} = \exp(-C^{(j_1, j_2)}/\epsilon)$ for all $(j_1, j_2) \in \mathcal{E}_r$. However, if these matrices are not row-stochastic, they cannot be interpreted as transition probability matrices of a Markov process.

Remark 2. For a nonnegative matrix A with at least one positive element in each row, define the vector $b = \mathbf{1}./A\mathbf{1}$. Then the matrix $\hat{A} = \text{diag}(b)A$ is row stochastic. Note that we can write $H(M | \text{diag}(\mu)A) = H(M | \hat{A}) + H(\mu | b)$. For a given rooted tree $\mathcal{T}_r = (\mathcal{V}, \mathcal{E}_r)$ assume that there is a set of matrices $A^{(j_1, j_2)}$ for $(j_1, j_2) \in \mathcal{E}_r$, which are not necessarily row-stochastic, but are such that for every $j_1 \in \mathcal{E}_r$ there is a vector b_{j_1} such that it holds that $b_{j_1} = \mathbf{1}./A^{(j_1, j_2)}\mathbf{1}$ for all j_2 such that $(j_1, j_2) \in \mathcal{E}_r$. Then, according to Proposition 4.2, the objective function of the maximum likelihood

estimation problem of a Markov process on \mathcal{T}_r in (4.1) may be written as

$$(4.5) \quad \sum_{(j_1, j_2) \in \mathcal{E}_r} H(M^{(j_1, j_2)} | \hat{A}^{(j_1, j_2)}) - \sum_{j \in \mathcal{V} \setminus \mathcal{L}} (\deg(j) - 1) H(\mu_j | b_j).$$

The Sinkhorn scheme (2.9) can equivalently be formulated in terms of the notation from the maximum likelihood estimation problem in Theorem 4.1. Given an initial set of positive vectors v_j for $j \in \mathcal{L}$, the updates are then expressed as

$$(4.6) \quad \begin{aligned} v_1 &= \varphi_1 \cdot / \mu_1, \\ v_j &= \mu_j \cdot / \hat{\varphi}_j \text{ for } j \in \mathcal{L} \setminus 1, \end{aligned}$$

where φ_1 is computed as in (4.2) before each update of v_1 , and $\hat{\varphi}_j$ is computed as in (4.3) before updating v_j for each $j \in \mathcal{L} \setminus 1$. This can be done efficiently, and by storing and reusing intermediate results appropriately, this algorithm is equivalent to Algorithm 3.1. It was shown that the classical Sinkhorn iterations are a block coordinate ascent in a Lagrange dual problem [42, 61]. In this sense, the scheme (4.6) can be understood as Sinkhorn iterations even for the problem (4.1).

PROPOSITION 4.5. *Let v_j , for $j \in \mathcal{L}$, be an initial set of positive vectors. Iteratively performing the updates (4.6) corresponds to a block coordinate ascent in a Lagrange dual problem to (4.1).*

In the light of Proposition 4.5 it is not surprising that the algorithm presented in [37], for the special case where the tree represents a hidden Markov chain, is of the form (4.6), where the vectors φ_j and $\hat{\varphi}_j$ are constructed according to the special tree structure.

Remark 3. The tree-structured optimal transport problem, i.e., problem (2.2) where the cost is of the form (3.1) for a tree-graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, can be formulated as a sum of pairwise optimal transport costs

$$(4.7) \quad \underset{\mu_j, j \in \mathcal{V} \setminus \Gamma}{\text{minimize}} \sum_{(j_1, j_2) \in \mathcal{E}} T^{(j_1, j_2)}(\mu_{j_1}, \mu_{j_2}),$$

where $T^{(j_1, j_2)}(\cdot, \cdot)$ is the optimal transport problem (2.1) with cost matrix $C^{(j_1, j_2)}$ for $(j_1, j_2) \in \mathcal{E}$. In particular, both problems have common optimal solutions which are identical in the sense that the projections of an optimal tensor \mathbf{M} in (2.2) coincide with a set of optimal transport plans in (4.7) in the sense that $P_{(j_1, j_2)}(\mathbf{M}) = M^{(j_1, j_2)}$ for $(j_1, j_2) \in \mathcal{E}$.

This, however, cannot be generalized to graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ that contain cycles. A counterexample can be found for a complete graph with three nodes (Figure 2). Let the cost matrices be

$$C^{(1,2)} = C^{(2,3)} = J = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad C^{(1,3)} = I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

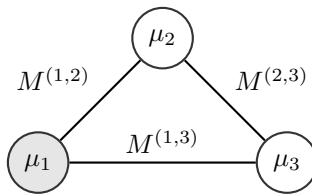


FIG. 2. Illustration of the cyclic graph in Remark 3.

where $\mu_1 = [1, 1]^T$ and $\Gamma = \{1\}$. The unique optimal solution to (4.7) with these parameters is $\mu_2 = \mu_3 = [1 \ 1]^T$, $M^{(1,2)} = M^{(2,3)} = I$, $M^{(1,3)} = J$, and the objective value is 0. However, note that the corresponding multimarginal cost tensor given by $\mathbf{C}_{i_1, i_2, i_3} = C_{i_1, i_2}^{(1,2)} + C_{i_1, i_3}^{(1,3)} + C_{i_2, i_3}^{(2,3)}$ is elementwise strictly positive, and thus the multimarginal problem cannot attain the value 0. In fact, there is no tensor $\mathbf{M} \in \mathbb{R}_+^{2 \times 2 \times 2}$ that is consistent with these projections $P_{1,2}(\mathbf{M}) = P_{2,3}(\mathbf{M}) = I$ and $P_{1,3}(\mathbf{M}) = J$.

Here we have illustrated a fundamental difficulty that must be handled when extending tree-graphs to graphs with cycles. Thus, naively extending the Schrödinger bridge from trees to general graphs by defining transition probability matrices $A^{(j_1, j_2)} = \exp(-C^{(j_1, j_2)}/\epsilon)$ for all $(j_1, j_2) \in \mathcal{E}$ and solving (4.1) does not yield an equivalence result with a multimarginal optimal transport problem as in Theorem 4.4.

5. Multimarginal vs. pairwise regularization. Another natural way to define an optimal transport problem on the tree \mathcal{T} is to minimize the sum of all bimarginal transport costs on the edges of \mathcal{T} , as in (4.7). In fact, the unregularized multimarginal optimal transport problem (2.2) structured according to $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ is equivalent to this pairwise problem. An alternative computational approach to the one taken in this paper is thus to regularize each term in the sum of (4.7) by an entropy term. The entropy regularized pairwise optimal transport problem on \mathcal{T} is then

$$(5.1) \quad \underset{\mu_j, j \in \mathcal{V} \setminus \Gamma}{\text{minimize}} \sum_{(j_1, j_2) \in \mathcal{E}} T_\epsilon^{(j_1, j_2)}(\mu_{j_1}, \mu_{j_2}),$$

where $T_\epsilon^{(j_1, j_2)}(\cdot, \cdot)$ is defined as the regularized bimarginal optimal transport problem (2.5) with cost matrix $C^{(j_1, j_2)}$. In particular, it is common to formulate barycenter problems in this pairwise regularized manner [7, 8, 23, 44, 46, 59]. Recently, we have empirically observed that in some applications the barycenter problem with multimarginal regularization as in (2.4) gives better results compared to the pairwise regularization in (5.1) (see [30, section 6.3]). Therein solutions using the multimarginal regularization are less smoothed out compared to the pairwise regularization. In the following we investigate the difference between the two problems. In particular, we show that the latter is not equivalent to the generalized Schrödinger bridge (4.1).

Consider a rooted version of the tree, denoted $\mathcal{T}_r = (\mathcal{V}, \mathcal{E}_r)$, where $r \in \mathcal{L}$. In the case that $A^{(j_1, j_2)} = K^{(j_1, j_2)}$ for all $(j_1, j_2) \in \mathcal{E}_r$, the multimarginal optimal transport problem has the same solution as problem (4.4). Now note that using Proposition 4.2 this problem can be written as

$$(5.2) \quad \underset{\mu_j, j \in \mathcal{V} \setminus \Gamma}{\text{minimize}} \sum_{(j_1, j_2) \in \mathcal{E}} T_\epsilon^{(j_1, j_2)}(\mu_{j_1}, \mu_{j_2}) - \sum_{j \in \mathcal{V} \setminus \mathcal{L}} (\deg(j) - 1)H(\mu_j).$$

Thus, the generalized Schrödinger bridge problem (4.1), or equivalently (5.2), is not equivalent to the pairwise regularized optimal transport problem (5.1), unless for the trivial case of a tree with only two vertices. Moreover, we can see from (5.2) that the multimarginal optimal transport problem penalizes not only the transport cost between the marginals, but in addition favors marginal distributions with low entropy.¹ One can thus expect less smoothed out distributions when solving the multimarginal optimal transport problem, which is desirable in many applications, such as localization problems [30] and computer vision applications [59]. Solving tree-structured

¹Recall that the definition of $H(\mu_j)$ essentially corresponds to the negative of the entropy of μ_j .

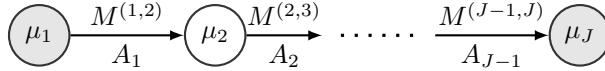


FIG. 3. Illustration of the linear path tree in section 5.1.

problems by using multimarginal regularization thus has some advantages as compared to pairwise regularization. First, it yields less smoothed out solutions, and second it preserves the connections to the Schrödinger bridge problem. Finally, empirical study suggests that the multimarginal problem is better conditioned compared to the pairwise problem, which allows for smaller values of the regularization parameter ϵ , while still yielding a numerically stable algorithm. We give empirical evidence for this behavior in section 5.2.

The solution of the pairwise optimal transport problem (5.1) is compared to the solution of the two equivalent problems (2.4) and (4.1) on the example of a path graph in section 5.1 and on a more complex tree in section 5.2.

5.1. The discrete time Schrödinger bridge problem. Consider a path tree $\mathcal{T}_r = (\mathcal{V}, \mathcal{E}_r)$, where $\mathcal{V} = \{1, 2, \dots, J\}$ and $\mathcal{E}_r = \{(j, j+1) | j = 1, \dots, J-1\}$, as sketched in Figure 3. Let A^j denote the probability transition matrix on $(j, j+1) \in \mathcal{E}_r$. This model corresponds to a Markov chain of length J . Assume that the distributions on the leaves $j = 1$ and $j = J$ are known. The most likely particle evolutions between them are then found by solving (2.11). Following [37], we see that this problem is equivalent to the discrete time and discrete space Schrödinger bridge in [54]. Assume that the distributions μ_j , for $j = 1, \dots, J$, are strictly positive, and define the row stochastic matrices $\bar{A}^j = \text{diag}(\mu_j)^{-1} M^j$ for $j = 1, \dots, J$. In terms of these matrices problem (2.11) reads

$$(5.3) \quad \begin{aligned} & \underset{\bar{A}_{[1:J-1]}, \mu_{[2:J-1]}}{\text{minimize}} && \sum_{j=1}^{J-1} \sum_{i=1}^n (\mu_j)_{iH} \left(\bar{A}_i^j \mid A_i^j \right) \\ & \text{subject to} && \bar{A}^j \mathbf{1} = \mathbf{1}, \quad \mu_{j+1} = (\bar{A}^j)^T \mu_j \quad \text{for } j = 1, \dots, J-1. \end{aligned}$$

Here $A_{i \cdot}$ denotes the i th row of A . Problem (5.3) is exactly the formulation of a Schrödinger bridge over a Markov chain from [54, equation (24)]. In [54] it is shown that a unique solution to a corresponding Schrödinger system exists if μ_J is a strictly positive distribution and the matrix $\prod_{j=1}^{J-1} A^j$ has only positive elements. The solution to the Schrödinger system may be obtained as a fixed point iteration [34], which is linked to the Sinkhorn iterations for entropy regularized optimal transport problems. We recall from [37] that the optimization problem (5.3) is nonconvex, whereas the equivalent formulation (2.11) is convex.

We note that with cost matrices defined as $C^j = -\epsilon \log(A^j)$ for $j = 1, \dots, J-1$, the discrete Schrödinger bridge problem may be written as

$$\underset{\mu_2, \dots, \mu_{J-1}}{\text{minimize}} \sum_{j=1}^{J-1} T_\epsilon^{(j,j+1)}(\mu_j, \mu_{j+1}) - \sum_{j=2}^{J-1} H(\mu_j).$$

In comparison to the pairwise optimal transport problem on the path graph \mathcal{T} , the Schrödinger bridge thus favors intermediate distributions with lower entropy, resulting in less smoothed out solutions. We illustrate this behavior for a path tree $\mathcal{T}_r = (\mathcal{V}, \mathcal{E}_r)$ with vertices $\mathcal{V} = \{1, 2, \dots, 6\}$, and where the initial and final distribution

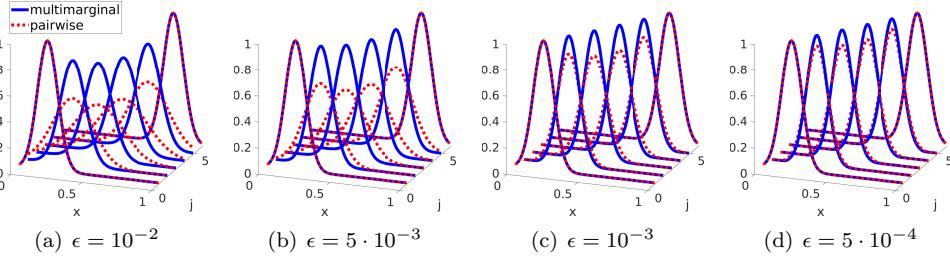


FIG. 4. Solutions to the multimarginal and pairwise entropy regularized optimal transport problem on a path graph for varying regularization parameter ϵ .

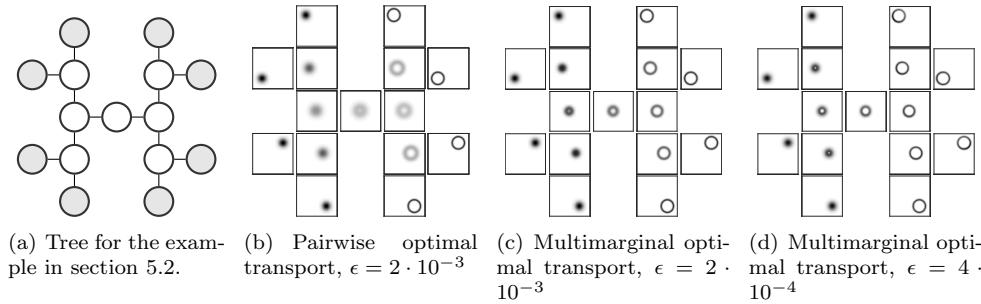


FIG. 5. Estimated marginals of the pairwise (b) and multimarginal (c), (d) optimal transport solutions on the tree in (a).

are given as $\mu_1(x) = \exp(-(\frac{x-0.2}{10})^2)$ and $\mu_J(x) = \exp(-(\frac{x-0.8}{10})^2)$, and the cost function is defined as the Euclidean distance. The results for problem (5.3) and problem (5.1) are compared for different values of the regularization parameter ϵ in Figure 4. Both optimal transport solutions describe a smooth way of shifting the mass from distribution μ_1 to μ_J . The larger the regularization parameter ϵ is chosen, the more smoothed out are the intermediate solutions to both problems. However, for each value of ϵ the multimarginal optimal transport solution is substantially less smoothed than the corresponding pairwise optimal transport solution.

5.2. Optimal transport with tree-structure. In this section we compare the entropy regularized multimarginal and pairwise optimal transport solutions on a more general tree. Consider the tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ illustrated in Figure 5(a) with 15 nodes, each representing a 50×50 pixel image. The marginal images on the 8 leaves, colored in gray, are known. Each edge on the tree \mathcal{T} is associated with a cost function defined by the Euclidean distance between any two pixels. Using this choice of cost function in the corresponding optimal transport problems yields smooth translations in power for the intermediate marginals.

We solve the entropy regularized pairwise optimal transport problem (5.1) with regularization parameter $\epsilon = 2 \cdot 10^{-3}$ on \mathcal{T} . The solution can be seen in Figure 5(b). Compared to the pairwise optimal transport estimate, the solution to the entropy regularized multimarginal optimal transport problem (2.4) on the same tree \mathcal{T} and with the same regularization parameter ϵ is significantly sharper and less smoothed out (see Figure 5(c)). For the pairwise optimal transport problem, the method diverges with a smaller regularization parameter, e.g., $\epsilon = 10^{-3}$. In contrast, for the multimarginal formulation the regularization parameter can be decreased further, still

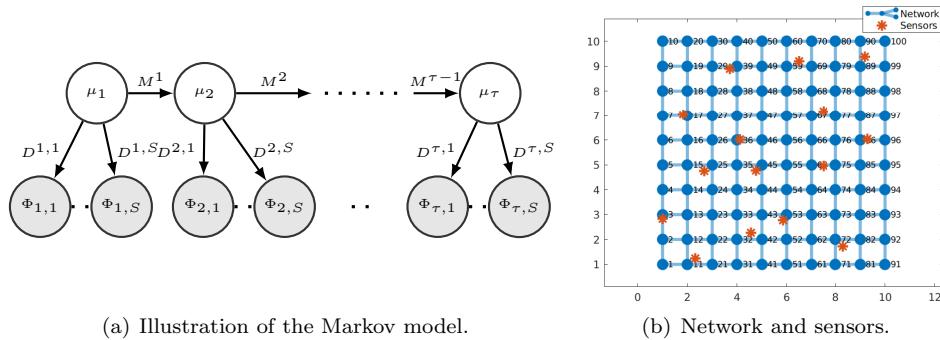


FIG. 6. Ensemble flow estimation example in section 6.

yielding a numerically stable algorithm. We have found that the method is still stable for a regularization parameter of $\epsilon = 4 \cdot 10^{-4}$, which results in very clear estimates on the intermediate nodes.

6. Estimating ensemble flows on a hidden Markov chain. We consider the problem of tracking an ensemble of agents on a network based on aggregate measurements from sensors distributed around the network. This is similar to [37], where ensemble flows of indistinguishable agents have been estimated as the maximum likelihood solution on a hidden Markov chain. The present work generalizes this method and provides an algorithm for solving the problem introduced therein. In particular, the framework in [37] is a special case of the method in Theorem 4.1 and can therefore be solved with Algorithm 3.1. Herein, we study different observation models and robustness of the estimates with respect to the number of agents.

Consider a graph $\mathcal{G} = (\mathcal{V}_{\mathcal{G}}, \mathcal{E}_{\mathcal{G}})$ with vertices $\mathcal{V}_{\mathcal{G}}$ and edges $\mathcal{E}_{\mathcal{G}}$. Let the set of nodes $\mathcal{V}_{\mathcal{G}}$ be the set of states of a Markov model with transition probability matrix $A^t \in \mathbb{R}^{n \times n}$, where $n = |\mathcal{V}_{\mathcal{G}}|$, for $t = 1, \dots, \tau-1$. We simulate a finite number of agents to evolve according to this Markov model for a number of time steps $t = 1, \dots, \tau$. Let $\mu_t \in \mathbb{R}^n$ denote the distribution of agents over the set of states for the times $t = 1, \dots, \tau$. An observation model is represented by the detection probability matrices $B^s \in \mathbb{R}^{n \times m}$ for $s = 1, \dots, S$, where m is the size of the observation space and S denotes the number of uncoupled observations at each time point. We then estimate the ensemble evolution as the solution of

$$\begin{aligned} & \underset{M_{[1:\tau-1]}, D_{[1:\tau], [1:S]}, \mu_{[1:\tau]}}{\text{minimize}} \quad \sum_{t=1}^{\tau-1} H(M^t \mid \text{diag}(\mu_t) A^t) + \sum_{t=1}^{\tau} \sum_{s=1}^S H(D^{s,t} \mid \text{diag}(\mu_t) B^s) \\ & \text{subject to} \quad M^t \mathbf{1} = \mu_{t-1}, \quad (M^t)^T \mathbf{1} = \mu_t, \quad D^{t,s} \mathbf{1} = \mu_t, \quad (D^{t,s})^T \mathbf{1} = \Phi_{t,s} \\ & \quad \text{for } t = 1, \dots, \tau, \text{ and } s = 1, \dots, S. \end{aligned}$$

The tree corresponding to the Markov model (6.1) is sketched in Figure 6(a). This optimization problem is similar to the one in [37], but therein the initial distribution of agents, i.e., the marginal μ_1 , is assumed to be known.

In the following we consider a number of N agents evolving on the network displayed in Figure 6(b) with $n = 100$ nodes and 180 edges. Figure 6(b) also shows the location of the $N_S = 15$ sensors. In a first step, we compute the discrete Schrödinger bridge (2.11) between the distribution μ_1 , where all agents are in node 1, and the

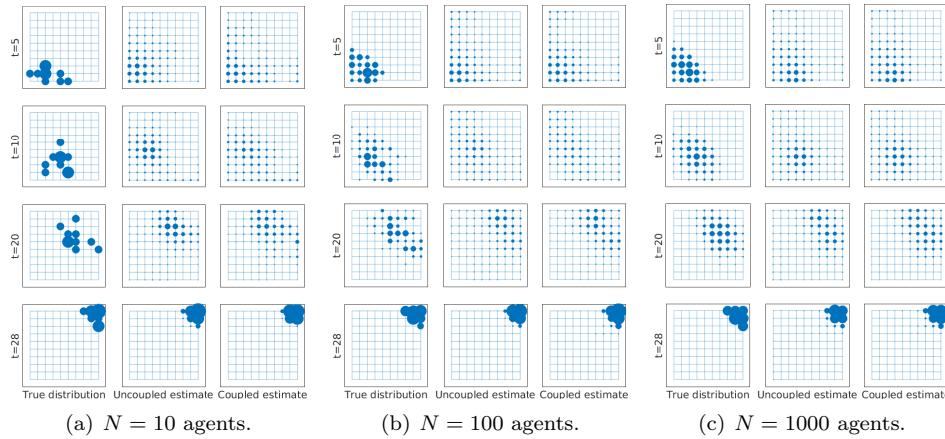


FIG. 7. *True ensemble flow and estimates with the two observation models for a varying number of agents.*

distribution μ_τ , where all agents are in node 100, with the random walk on \mathcal{G} as a prior. The resulting particle evolutions M^t , for $t = 1, \dots, \tau - 1$, in (2.11) define the transition probability matrices $\bar{A}^t = \text{diag}(\mu_t)^{-1}M^t$, for $t = 1, \dots, \tau - 1$, as in (5.3), used to simulate a number of N agents with initial distribution μ_1 . By construction, the final distribution is then μ_τ .

We consider two observation models, one where the sensors make uncoupled measurements, similar to the example in [37], and one where the sensors are coupled and form a joint measurement. In the uncoupled setting, we consider an observation space of two states for each sensor, where one state denotes that an agent is detected and the other one that the agent is undetected by the sensor. Hence, we define an observation probability matrix $B^s \in \mathbb{R}^{n \times 2}$ for each sensor $s = 1, \dots, S = N_S$, where the probability for an agent on node i to be detected by the sensor s is defined as $B_{i1}^s = \min(0.99, 2e^{-d_{s,i}})$, where $d_{s,i}$ denotes the Euclidean distance between the location of sensor s and the node i . Consequently the probability of not being detected is $B_{i2}^s = 1 - B_{i1}^s$. For the coupled observation model, the observation space consists of all possible sets of sensors that detect a given agent, i.e., all subsets of the set $\{1, \dots, N_S\}$. The size of the observation space is thus 2^{N_S} , and there is only one observation at each time instance, i.e., $S = 1$. Hence, we define an observation probability matrix $B^{\text{joint}} \in \mathbb{R}^{n \times 2^S}$, where the probability for an agent in node i to be detected by exactly the set \mathfrak{S} of sensors is given by $(\prod_{s \in \mathfrak{S}} B_{i1}^s)(\prod_{s \notin \mathfrak{S}} B_{i2}^s)$.

We solve the multimarginal optimal transport problem corresponding to (6.1) with assumed probability transition matrix A describing a random walk on \mathcal{G} and observation probability matrices B^s for the two described observation models. The results for both observation models are compared for $N \in \{10, 100, 1000\}$ agents in Figure 7, where the estimated number of agents in each node is plotted as a circle with size corresponding to the log-scaled weighted number of agents in that node. In Figure 7(a), one can see that for $N = 10$ agents the coupled estimate localizes the position of the agents slightly better than the uncoupled estimate. However, this effect decreases with an increasing number of agents, as the evolution of a single agent has less influence on the empirical distribution of agents. Already from a number of $N = 100$ agents the uncoupled estimate is competitive with the coupled estimate, albeit relying on significantly less information (cf. Figure 7(b)). For an ensemble of

$N = 1000$ agents one can hardly see any difference between the two estimates, as shown in Figure 7(c).

Note that the number of observations at each time instance for the uncoupled estimate is only of the size $2N_S = 30$, whereas the number of coupled observations is $2^{N_S} = 32768$. The improved observation model thus comes at an increased computational cost, which is exponential in the number of sensors, and becomes infeasible for larger numbers of sensors.

7. Conclusion. In this work we consider multimarginal optimal transport problems with cost functions that decouple according to a tree, and we show that the entropy regularized formulation of this problem is equivalent to a generalization of a time and space discrete Schrödinger bridge defined on the same tree. Moreover, we derive an efficient algorithm for solving this problem. We also compare the multimarginally regularized optimal transport problem to a commonly used pairwise regularized optimal transport problem and illustrate the benefits in theory and practice. Finally, we describe how to apply the framework to the problem of tracking an ensemble of indistinguishable agents.

Interestingly, the construction of the marginals in Theorem 3.2 is of the same form as the belief propagation algorithm [60, 65] for inference in graphical models. In this interpretation, the vectors u_j correspond to the local evidence in node $j \in \mathcal{V}$ and $\alpha_{(j_1, j_2)}$ is the message passed from node j_2 to node j_1 . Moreover, the objective function in (4.5) can be interpreted as the Bethe free energy [66], which is connected to belief propagation [65]. These similarities of our framework to belief propagation algorithms are investigated in [39] and could be a stepping stone to extending our framework to graphs with cycles. Another direction of interest is the extension to continuous state models.

Appendix A. Proofs. For an index set Γ , we denote the double sum $\sum_{j \in \Gamma} \sum_{i_j}$ by $\sum_{i_j: j \in \Gamma}$. For simplicity of notation, we let \mathcal{M} denote the set of matrices $M^{(j_1, j_2)}$ for $(j_1, j_2) \in \mathcal{E}$ and similarly write $\boldsymbol{\mu}$ and $\boldsymbol{\lambda}$ for the respective sets of optimization variables. The proof of Theorem 3.2 is based on the following lemma.

LEMMA A.1 ([30, Lemmas 1 and 2]). *Let $\mathbf{U} = u_1 \otimes u_2 \otimes \cdots \otimes u_J$ for a set of vectors u_1, u_2, \dots, u_J and \mathbf{K} be a tensor of the same size. If $\langle \mathbf{K}, \mathbf{U} \rangle = w^T u_j$ for a vector w that does not depend on u_j , then it holds that $P_j(\mathbf{K} \odot \mathbf{U}) = w \odot u_j$. Similarly, if $\langle \mathbf{K}, \mathbf{U} \rangle = w^T \text{diag}(u_{j_1}) W \text{diag}(u_{j_2}) \hat{w}$ for two vectors w, \hat{w} , and a matrix W that does not depend on u_{j_1} and u_{j_2} , then it holds that $P_{j_1, j_2}(\mathbf{K} \odot \mathbf{U}) = \text{diag}(w \odot u_{j_1}) W \text{diag}(u_{j_2} \odot \hat{w})$.*

Proof of Theorem 3.2. Due to the decoupling of the cost tensor \mathbf{C} as in (3.1), the tensor $\mathbf{K} = \exp(-\mathbf{C}/\epsilon)$ decouples according to (3.2) with the matrices $K^{(j_1, j_2)} = \exp(C^{(j_1, j_2)}/\epsilon)$ for $(j_1, j_2) \in \mathcal{E}$. We can therefore write

$$\langle \mathbf{K}, \mathbf{U} \rangle = \sum_{i_j: j \in \mathcal{V}} \left(\prod_{(j_1, j_2) \in \mathcal{E}} K_{i_{j_1}, i_{j_2}}^{(j_1, j_2)} \right) \left(\prod_{k \in \mathcal{V}} (u_k)_{i_k} \right) = \sum_{i_j} (u_j)_{i_j} (w_j)_{i_j},$$

where the vector w_j is of the form

$$(w_j)_{i_\ell} = \sum_{i_k: k \in \mathcal{V} \setminus j} \left(\prod_{(j_1, j_2) \in \mathcal{E}} K_{i_{j_1}, i_{j_2}}^{(j_1, j_2)} \right) \left(\prod_{j \in \mathcal{V}, j \neq \ell} (u_j)_{i_j} \right).$$

Consider the underlying tree to be rooted in node j . Then this can be written as

$(w_j)_{i_\ell} = \prod_{k: k \in \mathcal{N}_j} (\alpha_{(j,k)})_{i_j}$, where $\alpha_{(p(k),k)} \in \mathbb{R}^n$ is defined by

$$(\alpha_{(p(k),k)})_{i_{p(k)}} = \sum_{i_\ell: \ell \geq k} \prod_{m > k} K_{i_{p(m)}, i_m}^{(p(m), m)} (u_m)_{i_m}.$$

If $k \in \mathcal{L}$, then

$$(\alpha_{(p(k),k)})_{i_{p(k)}} = \sum_{i_k} K_{i_{p(k)}, i_k}^{(p(k), k)} (u_k)_{i_k} = \left(K^{(p(k), k)} u_k \right)_{i_{p(k)}}.$$

Otherwise, it holds that

$$\begin{aligned} (\alpha_{(p(k),k)})_{i_{p(k)}} &= \sum_{i_k} \left(K_{i_{p(k)}, i_k}^{(p(k), k)} (u_k)_{i_k} \prod_{\substack{\ell \in \mathcal{N}_k \\ \ell \neq p(k)}} \left(\sum_{i_m: m \geq \ell} \prod_{j > \ell} K_{i_{p(j)}, i_j}^{(p(j), j)} (u_j)_{i_j} \right) \right) \\ &= \sum_{i_k} \left(K_{i_{p(k)}, i_k}^{(p(k), k)} (u_k)_{i_k} \prod_{\substack{\ell \in \mathcal{N}_k \\ \ell \neq p(k)}} (\alpha_{(p(\ell), \ell)})_{i_{p(\ell)}} \right). \end{aligned}$$

This inductively defines the vectors $\alpha_{(j,k)}$ as in (3.3). The expression for the projection follows from Lemma A.1 with $w_j = \bigodot_{k \in \mathcal{N}_j} \alpha_{(j,k)}$. \square

Proof of Proposition 3.4. Note that the optimal solution to (2.4) with cost structured according to the tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ is of the form $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$, where \mathbf{K} is defined by (3.2) and \mathbf{U} is defined by (2.7). To prove the first claim, let $k \in \mathcal{V}$ lie on the direct path between $j \in \mathcal{V}$ and $\ell \in \mathcal{V}$. Then straightforward computation implies that

$$(P_{jk\ell}(\mathbf{K} \odot \mathbf{U}))_{i_j, i_k, i_\ell} (P_k(\mathbf{K} \odot \mathbf{U}))_{i_k} = (P_{jk}(\mathbf{K} \odot \mathbf{U}))_{i_j i_k} (P_{k\ell}(\mathbf{K} \odot \mathbf{U}))_{i_k, i_\ell}.$$

In particular, for any fixed i_k , the matrix $(P_{jk\ell}(\mathbf{K} \odot \mathbf{U}))_{\cdot, i_k, \cdot}$ is of rank 1. Hence, given $P_k(\mathbf{K} \odot \mathbf{U})$, the marginals $P_j(\mathbf{K} \odot \mathbf{U})$ and $P_\ell(\mathbf{K} \odot \mathbf{U})$ have no influence on each other. Thus, the optimal transport problem (2.4) on \mathcal{T} can be decoupled into smaller problems, by cutting \mathcal{T} in $k \in \mathcal{V}$, and solving an optimal transport problem (2.4) on each of the resulting subtrees, where $k \in \Gamma$ and $k \in \mathcal{L}$ for each of the subtrees.

To prove the second claim, recall from the definition of the tensor \mathbf{U} in (2.7) that $u_j = \mathbf{1}$ for all $j \in \mathcal{V} \setminus \Gamma$. Thus, for k such that $(k, \ell) \in \mathcal{E}$, the corresponding vector (3.3) is $\alpha_{(k,\ell)} = K^{(k,\ell)} \mathbf{1}$, which is constant and does not need to be updated when recomputing the projections (3.4). It thus suffices to solve (2.4) on the subtree of \mathcal{T} that is obtained by removing $\ell \in \mathcal{V}$ and $(k, \ell) \in \mathcal{E}$ from \mathcal{T} . \square

Proof of Theorem 4.1. Assume that the only edge with node $j = 1$ is denoted $(1, 2)$. We add the trivial constraints $M^{(j_1, j_2)} \mathbf{1} = (M^{(p(j_1), j_1)})^T \mathbf{1}$ for $(j_1, j_2) \in \mathcal{E}_r \setminus (1, 2)$. We relax this constraint, together with the constraint on μ_1 , and let $\lambda_{(j_1, j_2)}$, for $(j_1, j_2) \in \mathcal{E}_r$, denote the corresponding dual variables. Furthermore, we relax the constraints $(M^{(j_p, j)})^T \mathbf{1} = \mu_j$ with dual variables λ_j for the leaves $j \in \mathcal{L} \setminus 1$. A Lagrangian for (4.1) is then

$$\begin{aligned} L(\mathcal{M}, \boldsymbol{\mu}, \boldsymbol{\lambda}) &= \sum_{(j_1, j_2) \in \mathcal{E}} H \left(M^{(j_1, j_2)} \mid \text{diag}(\mu_{j_1}) A^{(j_1, j_2)} \right) + \lambda_{(1,2)}^T (M^{(1,2)} \mathbf{1} - \mu_1) \\ &\quad + \sum_{j_2 \in \mathcal{L}} \lambda_{j_2}^T (\mu_{j_2} - (M^{(j_1, j_2)})^T \mathbf{1}) \\ &\quad + \sum_{\substack{(j_1, j_2) \in \mathcal{E} \\ j_1 \neq 1}} \lambda_{(j_1, j_2)}^T (M^{(j_1, j_2)} \mathbf{1} - (M^{(j_1, j_1)})^T \mathbf{1}). \end{aligned}$$

When j_2 is an inner node on the tree, i.e., $j_2 \notin \mathcal{L}$, the derivative with respect to the entries $M_{i_1 i_2}^{(j_1, j_2)}$, for all $i_1, i_2 = 1, \dots, n$, and $(j_1, j_2) \in \mathcal{E}_r$ is

$$(A.1) \quad \log \left(\frac{M_{i_1 i_2}^{(j_1, j_2)}}{(\mu_{j_1})_{i_1} A_{i_1 i_2}^{(j_1, j_2)}} \right) + (\lambda_{(j_1, j_2)})_{i_2} - \sum_{k: (j_2, k) \in \mathcal{E}} (\lambda_{(j_2, k)})_{i_2}.$$

Since (4.1) is convex, a mass transport plan is optimal if this gradient vanishes, yielding the expression in terms of the other variables,

$$(A.2) \quad M^{(j_1, j_2)} = \text{diag}(\mu_{j_1} \cdot / v_{(j_1, j_2)}) A^{(j_1, j_2)} \text{diag} \left(\bigodot_{k: (j_2, k) \in \mathcal{E}} v_{(j_2, k)} \right),$$

where $v_{(j_1, j_2)} = \exp(\lambda_{(j_1, j_2)})$ for all $(j_1, j_2) \in \mathcal{E}_r$. In the case $j_2 \in \mathcal{L}$, the last sum in the derivative (A.1) is replaced by $(\lambda_{j_2})_l$. Defining $v_{j_2} = \exp(\lambda_{j_2})$, the optimal mass transport plan is thus of the form

$$(A.3) \quad M^{(j_1, j_2)} = \text{diag}(\mu_{j_1} \cdot / v_{(j_1, j_2)}) A^{(j_1, j_2)} \text{diag}(v_{j_2}).$$

Next, note that the marginal of the optimal transport plan satisfies

$$(A.4) \quad \mu_{j_2} = M_{(j_1, j_2)}^T \mathbf{1} = \hat{\varphi}_{j_2} \odot \varphi_{j_2}$$

$$(A.5) \quad \text{with } \varphi_{j_2} = \bigodot_{k: (j_2, k) \in \mathcal{E}} v_{(j_2, k)} \quad \text{and} \quad \hat{\varphi}_{j_2} = (A^{(j_1, j_2)})^T (\mu_{j_1} \cdot / v_{(j_1, j_2)}).$$

Since for all k such that $(j_2, k) \in \mathcal{E}_r$ it holds that $M_{(j_2, k)} \mathbf{1} = \mu_{j_2}$, we get

$$v_{(j_2, k)} = A^{(j_2, k)} \left(\bigodot_{\ell: (k, \ell) \in \mathcal{E}} v_{(k, \ell)} \right) = A^{(j_2, k)} \varphi_k,$$

which completes the definition of the vectors $(\varphi_j)_{j \in \mathcal{V}}$. Similarly to (A.4) it holds that

$$(A.6) \quad \mu_{j_1} = \left(\bigodot_{k: (j_1, k) \in \mathcal{E}} v_{(j_1, k)} \right) \odot A \left(\mu_{j_{p(1)}} \cdot / v_{(j_{p(1)}, j_1)} \right).$$

Plugging (A.6) into $\hat{\varphi}_{j_2}$ in (A.5) yields the recursive definition of the vectors $(\hat{\varphi}_j)_{j \in \mathcal{V}}$. The expression for the mass transport plans $M^{(j_1, j_2)}$, for $(j_1, j_2) \in \mathcal{E}$, in terms of the vectors $\hat{\varphi}_{j_1}$, φ_{j_2} and $\varphi_{j_1 \setminus j_2}$ follows by identifying them in expressions (A.2) and (A.3). \square

Proof of Proposition 4.2. For any matrices $M, A \in \mathbb{R}_+^{n \times n}$ and $\mu \in \mathbb{R}_+^n$ one can write the term $H(M | \text{diag}(\mu)A)$ as

$$(A.7) \quad \sum_{i, j=1}^n \left(M_{ij} \log \left(\frac{M_{ij}}{A_{ij}} \right) - M_{ij} \right) + \sum_{i, j=1}^n (\mu_i A_{ij} - M_{ij} \log(\mu_i)).$$

Since $M\mathbf{1} = \mu$ and $A\mathbf{1} = \mathbf{1}$, the second sum can be simplified to $\sum_{i=1}^n (\mu_i - \mu_i \log(\mu_i))$. Furthermore, adding the term $\sum_{i, j=1}^n A_{ij} - \sum_{i=1}^n 1 = 0$ to (A.7), the expression can be written as $H(M | A) - H(\mu)$. Due to the underlying tree structure of problem (4.1), the number of outgoing edges from the root node j_r is $\deg(j_r)$, and for all other vertices $j \in \mathcal{V} \setminus \{j_r\}$ the number of outgoing edges is $\deg(j) - 1$. In the case that $j_r \in \mathcal{L}$, since the marginal μ_{j_r} is known, the term $H(\mu_{j_r})$ is constant and can be removed from the objective without changing the optimal solution. \square

Proof of Theorem 4.4. Note that for a rooted directed tree, in each node $j \in \mathcal{V} \setminus \mathcal{L}$, there is one incoming edge and the rest of the connected edges are outgoing. Its neighboring nodes are therefore given by the set $\mathcal{N}_j = p(j) \cup \{k : (j, k) \in \mathcal{E}_r\}$. Thus,

$$(A.8) \quad \bigodot_{k \in \mathcal{N}_j} \alpha_{(j,k)} = (K^{(p(j),j)})^T \alpha_{(j,p(j))} \odot \bigodot_{k:(j,k) \in \mathcal{E}_r} K^{(j,k)} \alpha_{(j,k)}.$$

For any edge $(j, k) \in \mathcal{E}$ and for the reverse edge (j, j_p) it holds that

$$\begin{aligned} \alpha_{(j,k)} &= \bigodot_{\ell \in \mathcal{N}_k \setminus \{j\}} K^{(k,\ell)} \alpha_{(k,\ell)} = \bigodot_{\ell:(k,\ell) \in \mathcal{E}} K^{(k,\ell)} \alpha_{(k,\ell)}, \\ \alpha_{(j,p(j))} &= \bigodot_{\ell \in \mathcal{N}_{p(j)} \setminus \{j\}} K^{(p(j),\ell)} \alpha_{(j,\ell)} = (K^{(p(p(j)),p(j))})^T \alpha_{(p(j),p(p(j)))} \odot \bigodot_{\substack{\ell:(p(j),\ell) \in \mathcal{E}, \\ \ell \neq j}} K^{(p(j),\ell)} \alpha_{(p(j),\ell)}. \end{aligned}$$

By associating the first term in (A.8) with $\hat{\varphi}_j$ and the second term with φ_j , we see that the tensor structure of \mathbf{K} gives rise to the construction of φ_j and $\hat{\varphi}_j$ as in Theorem 4.1. Hence, we can define a tensor in analogy to \mathbf{U} of the form $\mathbf{V} = (\mathbf{1}./v_1) \otimes v_2 \otimes \cdots \otimes v_J$, where v_j are the vectors from Theorem 4.1 if $j \in \mathcal{L}$, and $v_j = \mathbf{1}$ if $j \in \mathcal{V} \setminus \mathcal{L}$. Then, the tensor $\mathbf{K} \odot \mathbf{V}$ has the same marginals as the tensor $\mathbf{K} \odot \mathbf{U}$. Due to an extension of Sinkhorn's theorem to tensors [32], it follows that $\mathbf{U} = \mathbf{V}$ (up to scaling with a factor and its inverse in the vectors u_1, \dots, u_J) and $P_j(\mathbf{M}) = \mu_j$ for all $j \in \mathcal{V}$. \square

Proof of Proposition 4.5. Based on the proof to Theorem 4.1, a Lagrange dual to problem (4.1) can be formulated as to maximize

$$(A.9) \quad \begin{aligned} & - \sum_{(j_1,j_2) \in \mathcal{E}} (\mu_{j_1} \odot \exp(-\lambda_{(j_1,j_2)})) A^{(j_1,j_2)} \left(\bigodot_{k:(j_2,k) \in \mathcal{E}} \exp(\lambda_{(j_2,k)}) \right) \\ & - \sum_{j \in \mathcal{L}} (\mu_{j_p} \odot \exp(-\lambda_{(j_p,j)})) A^{(j_1,j_2)} (\exp(\lambda_j)) - \lambda_{(1,2)}^T \mu_1 + \sum_{j \in \mathcal{L}} \lambda_j^T \mu_j \end{aligned}$$

with respect to μ_j for all inner nodes j , and the dual variables $\lambda_{(j_1,j_2)}$ for $(j_1, j_2) \in \mathcal{E}$, and λ_j for $j \in \mathcal{L}$. A block coordinate ascent in the dual is then to iteratively maximize (A.9) with respect to one of the dual variable vectors, while keeping the others fixed. Denote $v_{(j_1,j_2)} = \exp(\lambda_{(j_1,j_2)})$ for $(j_1, j_2) \in \mathcal{E}$, and $v_j = \exp(\lambda_j)$ for $j \in \mathcal{L}$. The gradient of (A.9) with respect to $\lambda_{(1,2)}$ vanishes if it holds that

$$v_{(1,2)} = A^{(j_1,j_2)} \left(\bigodot_{k:(2,k) \in \mathcal{E}} v_{(2,k)} \right),$$

and the gradient with respect to λ_j vanishes if $v_j = \mu_j ./ ((A^{(j_p,j)})^T (\mu_{j_p} ./ v_{(j_p,j)}))$ for $j \in \mathcal{L}$, and where $j_p \in \mathcal{V}$ is the parent of node j . Finally, the gradient of (A.9) with respect to $\lambda_{(j_1,j_2)}$, where $j_1 \neq 1$, vanishes if $M_{(j_1,j_2)} \mathbf{1} = M_{(j_p,j_1)}^T \mathbf{1}$, where j_p is the parent of node j_1 , which leads to the recursive definition of φ_j and $\hat{\varphi}_j$ for $j \in \mathcal{V}$, as in Theorem 4.1. Hence, given an initial set of positive vectors $v_{(1,2)}$ and v_j for $j \in \mathcal{L}$, the scheme (4.6) is a block coordinate ascent in a Lagrange dual of problem (4.1). \square

REFERENCES

- [1] B. ACCIAIO, J. BACKHOFF-VERAGUAS, AND R. CARMONA, *Extended mean field control problems: Stochastic maximum principle and transport perspective*, SIAM J. Control Optim., 57 (2019), pp. 3666–3693.
- [2] J. ADLER, A. RINGH, O. ÖKTEM, AND J. KARLSSON, *Learning to Solve Inverse Problems Using Wasserstein Loss*, preprint, arXiv:1710.10898, 2017.
- [3] M. AGUEH AND G. CARLIER, *Barycenters in the Wasserstein space*, SIAM J. Math. Anal., 43 (2011), pp. 904–924.
- [4] M. ARJOVSKY, S. CHINTALA, AND L. BOTTOU, *Wasserstein generative adversarial networks*, in Proceedings of the International Conference on Machine Learning, 2017, pp. 214–223.
- [5] H. BAUSCHKE AND A. LEWIS, *Dykstras algorithm with Bregman projections: A convergence proof*, Optimization, 48 (2000), pp. 409–427.
- [6] E. BAYRAKTAR, A. M. COX, AND Y. STOEV, *Martingale optimal transport with stopping*, SIAM J. Control Optim., 56 (2018), pp. 417–433.
- [7] J.-D. BENAMOU, G. CARLIER, M. CUTURI, L. NENNA, AND G. PEYRÉ, *Iterative Bregman projections for regularized transportation problems*, SIAM J. Sci. Comput., 37 (2015), pp. A1111–A1138.
- [8] N. BONNEEL, G. PEYRÉ, AND M. CUTURI, *Wasserstein barycentric coordinates: Histogram regression using optimal transport*, ACM Trans. Graph., 35 (2016), pp. 71:1–71:10.
- [9] N. BONNEEL, M. VAN DE PANNE, S. PARIS, AND W. HEIDRICH, *Displacement interpolation using Lagrangian mass transport*, ACM Trans. Graph., 30 (2011), pp. 158:1–158:12.
- [10] X. CHEN, *Structure theory for ensemble controllability, observability, and duality*, Math. Control Signals Systems, 31 (2019), pp. 1–40.
- [11] Y. CHEN, G. CONFORTI, AND T. GEORGIOU, *Measure-valued spline curves: An optimal transport viewpoint*, SIAM J. Math. Anal., 50 (2018), pp. 5947–5968.
- [12] Y. CHEN, G. CONFORTI, T. GEORGIOU, AND L. RIPANI, *Multi-Marginal Schrödinger Bridges*, in Proceedings of the International Conference on Geometric Science of Information, 2019, pp. 725–732.
- [13] Y. CHEN, T. GEORGIOU, AND M. PAVON, *Entropic and displacement interpolation: A computational approach using the Hilbert metric*, SIAM J. Appl. Math., 76 (2016), pp. 2375–2396.
- [14] Y. CHEN, T. GEORGIOU, AND M. PAVON, *On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint*, J. Optim. Theory Appl., 169 (2016), pp. 671–691.
- [15] Y. CHEN, T. GEORGIOU, AND M. PAVON, *Optimal steering of a linear stochastic system to a final probability distribution, Part I*, IEEE Trans. Automat. Control, 61 (2016), pp. 1158–1169.
- [16] Y. CHEN, T. GEORGIOU, M. PAVON, AND A. TANNENBAUM, *Robust transport over networks*, IEEE Trans. Automat. Control, 62 (2016), pp. 4675–4682.
- [17] Y. CHEN AND J. KARLSSON, *State tracking of linear ensembles via optimal mass transport*, IEEE Control Syst. Lett., 2 (2018), pp. 260–265.
- [18] S.-N. CHOW, W. HUANG, Y. LI, AND H. ZHOU, *Fokker–Planck equations for a free energy functional or Markov process on a graph*, Arch. Ration. Mech. Anal., 203 (2012), pp. 969–1008.
- [19] G. CONFORTI AND C. LÉONARD, *Reciprocal classes of random walks on graphs*, Stochastic Process. Appl., 127 (2017), pp. 1870–1896.
- [20] T. COVER AND J. THOMAS, *Elements of Information Theory*, John Wiley & Sons, Hoboken, NJ, 2006.
- [21] I. CSISZAR, *Why least squares and maximum entropy? An axiomatic approach to inference for linear inverse problems*, Ann. Statist., 19 (1991), pp. 2032–2066.
- [22] M. CUTURI, *Sinkhorn distances: Lightspeed computation of optimal transport*, in Proceedings of Advances in Neural Information Processing Systems, 2013, pp. 2292–2300.
- [23] M. CUTURI AND A. DOUCET, *Fast computation of Wasserstein barycenters*, in Proceedings of the International Conference on Machine Learning, 2014, pp. 685–693.
- [24] D. DAWSON, L. GOROSTIZA, AND A. WAKOLBINGER, *Schrödinger processes and large deviations*, J. Math. Phys., 31 (1990), pp. 2385–2388.
- [25] A. DEMBO AND O. ZEITOUNI, *Large Deviations Techniques and Applications*, Springer, Berlin, 2009.
- [26] R. DIESTEL, *Graph Theory*, Springer, Berlin, 2010.
- [27] A. DOMINITZ AND A. TANNENBAUM, *Texture mapping via optimal mass transport*, IEEE Trans. Vis. Comput. Graph., 16 (2010), pp. 419–433.

- [28] R. ELLIS, *Entropy, Large Deviations, and Statistical Mechanics*, Springer, Berlin, 2006.
- [29] F. ELVANDER, I. HAASLER, A. JAKOBSSON, AND J. KARLSSON, *Tracking and sensor fusion in direction of arrival estimation using optimal mass transport*, in Proceedings of the 26th European Signal Processing Conference (EUSIPCO), IEEE, 2018, pp. 1617–1621.
- [30] F. ELVANDER, I. HAASLER, A. JAKOBSSON, AND J. KARLSSON, *Multi-marginal optimal mass transport using partial information with applications in robust localization and sensor fusion*, Signal Process., 171 (2020).
- [31] H. FÖLLMER, *Random fields and diffusion processes*, in École d'Été de Probabilités de Saint-Flour XV–XVII, 1985–87, P.-L. Hennequin, ed., Lecture Notes in Math. 1362, Springer, Berlin, 1988, pp. 101–203.
- [32] J. FRANKLIN AND J. LORENZ, *On the scaling of multidimensional matrices*, Linear Algebra Appl., 114–115 (1989), pp. 717–735.
- [33] W. GANGBO AND A. ŚWIĘCH, *Optimal maps for the multidimensional Monge–Kantorovich problem*, Comm. Pure Appl. Math., 51 (1998), pp. 23–45.
- [34] T. GEORGIU AND M. PAVON, *Positive contraction mappings for classical and quantum Schrödinger systems*, J. Math. Phys., 56 (2015), 033301.
- [35] N. GHOUSSOUB, Y.-H. KIM, AND A. Z. PALMER, *Optimal transport with controlled dynamics and free end times*, SIAM J. Control Optim., 56 (2018), pp. 3239–3259.
- [36] H. GZYL, *The Method of Maximum Entropy*, World Scientific, River Edge, NJ, 1995.
- [37] I. HAASLER, A. RINGH, Y. CHEN, AND J. KARLSSON, *Estimating ensemble flows on a hidden Markov chain*, in Proceedings of the IEEE Conference on Decision and Control, IEEE, 2019, pp. 1331–1338.
- [38] I. HAASLER, A. RINGH, Y. CHEN, AND J. KARLSSON, *Multi-Marginal Optimal Transport with a Tree-Structured Cost and the Schrödinger Bridge Problem*, 2020, <https://arxiv.org/abs/2004.06909v3>.
- [39] I. HAASLER, R. SINGH, Q. ZHANG, J. KARLSSON, AND Y. CHEN, *Multi-Marginal Optimal Transport and Probabilistic Graphical Models*, IEEE Trans. Inform. Theory, 67 (2021), pp. 4647–4668, preprint, arXiv:2006.14113, 2020.
- [40] B. JAMISON, *Reciprocal processes*, Probab. Theory Related Fields, 30 (1974), pp. 65–86.
- [41] B. JAMISON, *The Markov processes of Schrödinger*, Probab. Theory Related Fields, 32 (1975), pp. 323–331.
- [42] J. KARLSSON AND A. RINGH, *Generalized Sinkhorn iterations for regularizing inverse problems using optimal mass transport*, SIAM J. Imaging Sci., 10 (2017), pp. 1935–1962.
- [43] S. KOLOURI, S. R. PARK, M. THORPE, D. SLEPCEV, AND G. K. ROHDE, *Optimal mass transport: Signal processing and machine-learning applications*, IEEE Signal Process. Mag., 34 (2017), pp. 43–59.
- [44] A. KROSHNIN, N. TUPITSA, D. DVINSKIKH, P. DVURECHENSKY, A. GASNIKOV, AND C. URIBE, *On the complexity of approximating Wasserstein barycenters*, in Proceedings of the International Conference on Machine Learning, 2019, pp. 3530–3540.
- [45] C. LÉONARD, *From the Schrödinger problem to the Monge–Kantorovich problem*, J. Funct. Anal., 262 (2012), pp. 1879–1920.
- [46] T. LIN, N. HO, X. CHEN, M. CUTURI, AND M. JORDAN, *Revisiting Fixed Support Wasserstein Barycenter*, preprint, arXiv:2002.04783, 2020.
- [47] T. LIN, N. HO, M. CUTURI, AND M. JORDAN, *On the Complexity of Approximating Multi-marginal Optimal Transport*, preprint, arXiv:1910.00152, 2019.
- [48] C. LÉONARD, *A survey of the Schrödinger problem and some of its connections with optimal transport*, Discrete Contin. Dyn. Syst. A, 34 (2014), pp. 1533–1574.
- [49] Z.-Q. LUO AND P. TSENG, *On the convergence of the coordinate descent method for convex differentiable minimization*, J. Optim. Theory Appl., 72 (1992), pp. 7–35.
- [50] T. MIKAMI, *Monge's problem with a quadratic cost by the zero-noise limit of h-path processes*, Probab. Theory Related Fields, 129 (2004), pp. 245–260.
- [51] T. MIKAMI AND M. THIEULLEN, *Optimal transportation problem by stochastic optimal control*, SIAM J. Control Optim., 47 (2008), pp. 1127–1139.
- [52] G. MONTAVON, K.-R. MÜLLER, AND M. CUTURI, *Wasserstein training of restricted Boltzmann machines*, in Proceedings of Advances in Neural Information Processing Systems, 2016, pp. 3718–3726.
- [53] B. PASS, *Multi-marginal optimal transport: Theory and applications*, ESAIM Math. Model. Numer. Anal., 49 (2015), pp. 1771–1790.
- [54] M. PAVON AND F. TICOZZI, *Discrete-time classical and quantum Markovian evolutions: Maximum entropy problems on path space*, J. Math. Phys., 51 (2010), 042104.
- [55] L. RÜSCHENDORF, *Optimal solutions of multivariate coupling problems*, Appl. Math. (Warsaw), 23 (1995), pp. 325–338.

- [56] L. RÜSCHENDORF AND L. UCKELMANN, *On the n -coupling problem*, J. Multivariate Anal., 81 (2002), pp. 242–258.
- [57] E. SCHRÖDINGER, *Über die Umkehrung der Naturgesetze*, Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse, 1931, pp. 144–153.
- [58] R. SINKHORN, *Diagonal equivalence to matrices with prescribed row and column sums*, Amer. Math. Monthly, 74 (1967), pp. 402–405.
- [59] J. SOLOMON, F. DE GOES, G. PEYRÉ, M. CUTURI, A. BUTSCHER, A. NGUYEN, T. DU, AND L. GUIBAS, *Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains*, ACM Trans. Graph., 34 (2015), pp. 66:1–66:11.
- [60] Y. TEH AND M. WELLING, *The unified propagation and scaling algorithm*, in Proceedings of Advances in Neural Information Processing Systems, 2002, pp. 953–960.
- [61] P. TSENG, *Dual ascent methods for problems with strictly convex costs and linear constraints: A unified approach*, SIAM J. Control Optim., 28 (1990), pp. 214–242.
- [62] C. VILLANI, *Optimal Transport: Old and New*, Springer, Berlin, 2008.
- [63] I. G. VLADIMIROV AND I. R. PETERSON, *State distributions and minimum relative entropy noise sequences in uncertain stochastic systems: The discrete-time case*, SIAM J. Control Optim., 53 (2015), pp. 1107–1153.
- [64] I. YANG, *A convex optimization approach to distributionally robust Markov decision processes with Wasserstein distance*, IEEE Control Syst. Lett., 1 (2017), pp. 2475–1456.
- [65] J. S. YEDIDIA, W. T. FREEMAN, AND Y. WEISS, *Understanding belief propagation and its generalizations*, in Exploring Artificial Intelligence in the New Millennium, Morgan Kaufmann Publishers, San Francisco, 2003, pp. 239–269.
- [66] J. S. YEDIDIA, W. T. FREEMAN, AND Y. WEISS, *Constructing free-energy approximations and generalized belief propagation algorithms*, IEEE Trans. Inform. Theory, 51 (2005), pp. 2282–2312.
- [67] S. ZENG, *Sample-based population observers*, Automatica, 101 (2019), pp. 166–174.