



Polynomial-time algorithms for multimarginal optimal transport problems with structure

Jason M. Altschuler¹ · Enric Boix-Adserà¹

Received: 13 December 2020 / Accepted: 29 June 2022 / Published online: 16 August 2022
© The Author(s) 2022

Abstract

Multimarginal Optimal Transport (MOT) has attracted significant interest due to applications in machine learning, statistics, and the sciences. However, in most applications, the success of MOT is severely limited by a lack of efficient algorithms. Indeed, MOT in general requires exponential time in the number of marginals k and their support sizes n . This paper develops a general theory about what “structure” makes MOT solvable in $\text{poly}(n, k)$ time. We develop a unified algorithmic framework for solving MOT in $\text{poly}(n, k)$ time by characterizing the structure that different algorithms require in terms of simple variants of the dual feasibility oracle. This framework has several benefits. First, it enables us to show that the Sinkhorn algorithm, which is currently the most popular MOT algorithm, requires strictly more structure than other algorithms do to solve MOT in $\text{poly}(n, k)$ time. Second, our framework makes it much simpler to develop $\text{poly}(n, k)$ time algorithms for a given MOT problem. In particular, it is necessary and sufficient to (approximately) solve the dual feasibility oracle—which is much more amenable to standard algorithmic techniques. We illustrate this ease-of-use by developing $\text{poly}(n, k)$ -time algorithms for three general classes of MOT cost structures: (1) graphical structure; (2) set-optimization structure; and (3) low-rank plus sparse structure. For structure (1), we recover the known result that Sinkhorn has $\text{poly}(n, k)$ runtime; moreover, we provide the first $\text{poly}(n, k)$ time algorithms for computing solutions that are exact and sparse. For structures (2)–(3), we give the first $\text{poly}(n, k)$ time algorithms, even for approximate computation. Together, these three structures encompass many—if not most—current applications of MOT.

Keywords Multimarginal optimal transport · Polynomial-time algorithms · Implicit linear programming · Structured linear programs

Mathematics Subject Classification 90C08 · 90C06

Work partially supported by NSF GRFP 1122374, a TwoSigma PhD fellowship, and a Siebel Scholarship.

✉ Jason M. Altschuler
jasonalt@mit.edu

¹ Laboratory for Information and Decision Systems (LIDS), MIT, Cambridge, MA 02139, US

Contents

1	Introduction	1109
1.1	Contribution 1: unified algorithmic framework for MOT	1110
1.1.1	Answer to Q1: candidate $\text{poly}(n, k)$ -time algorithms	1111
1.1.2	Answer to Q2: structure necessary to run candidate algorithms	1111
1.1.3	Answer to Q3: characterizing what MOT problems each algorithm can solve	1112
1.1.4	Answer to Q4: ease-of-use for checking if MOT is solvable in polynomial time	1113
1.1.5	Practical algorithmic tradeoffs beyond polynomial-time solvability	1113
1.2	Contribution 2: applications to general classes of structured MOT problems	1114
1.2.1	Graphical structure	1115
1.2.2	Set-optimization structure	1116
1.2.3	Low-rank and sparse structure	1117
1.3	Related work	1119
1.3.1	MOT algorithms	1119
1.3.2	Graphically structured MOT problems with constant treewidth	1120
1.3.3	Tractable MOT problems beyond graphically structured costs	1121
1.3.4	Intractable MOT problems	1122
1.3.5	Variants of MOT	1122
1.3.6	Optimization over joint distributions	1123
1.4	Organization	1123
2	Preliminaries	1123
2.1	Multimarginal optimal transport	1124
2.2	Regularization	1125
3	Oracles	1126
4	Algorithms to oracles	1129
4.1	The Ellipsoid algorithm and the MIN oracle	1129
4.1.1	Algorithm	1130
4.1.2	Equivalence of bottleneck to MIN	1131
4.1.3	Practical implementation via Column Generation	1133
4.2	The Multiplicative Weights Update and the AMIN oracle	1134
4.2.1	Algorithm	1135
4.2.2	Equivalence of bottleneck to AMIN	1137
4.3	The Sinkhorn algorithm and the SMIN oracle	1140
4.3.1	Algorithm	1141
4.3.2	Equivalence of bottleneck to SMIN	1143
5	Application: MOT problems with graphical structure	1145
5.1	Setup	1145
5.2	Polynomial-time algorithms	1147
5.3	Application vignette: Fluid dynamics	1149
6	Application: MOT problems with set-optimization structure	1151
6.1	Setup	1151
6.2	Polynomial-time algorithms	1153
6.3	Application vignette: Network reliability with correlations	1156
7	Application: MOT problems with low-rank plus sparse structure	1159
7.1	Setup	1160
7.2	Polynomial-time algorithms	1161
7.2.1	Technical ingredients	1162
7.2.2	Proof of Theorem 7.4	1166
7.3	Application vignette: Risk estimation	1167
7.4	Application vignette: Projection to the transportation polytope	1169
8	Discussion	1170
A	Deferred proof details	1171
A.1	Proof of Lemma 4.11	1171
A.2	Proof of Lemma 4.16	1173
B	Additional numerical experiments	1174
	References	1175

1 Introduction

Multimarginal Optimal Transport (MOT) is the problem of linear programming over joint probability distributions with fixed marginal distributions. In this way, MOT generalizes the classical Kantorovich formulation of Optimal Transport from 2 marginal distributions to an arbitrary number $k \geq 2$ of them.

More precisely, an MOT problem is specified by a cost tensor C in the k -fold tensor product space $(\mathbb{R}^n)^{\otimes k} = \mathbb{R}^n \otimes \dots \otimes \mathbb{R}^n$, and k marginal distributions μ_1, \dots, μ_k in the simplex $\Delta_n = \{v \in \mathbb{R}_{\geq 0}^n : \sum_{i=1}^n v_i = 1\}$.¹ The MOT problem is to compute

$$\min_{P \in \mathcal{M}(\mu_1, \dots, \mu_k)} \langle P, C \rangle \tag{MOT}$$

where $\mathcal{M}(\mu_1, \dots, \mu_k)$ is the “transportation polytope” consisting of all entrywise non-negative tensors $P \in (\mathbb{R}^n)^{\otimes k}$ satisfying the marginal constraints $\sum_{j_1, \dots, j_{i-1}, j_{i+1}, \dots, j_k} P_{j_1, \dots, j_{i-1}, j, j_{i+1}, \dots, j_k} = [\mu_i]_j$ for all $i \in \{1, \dots, k\}$ and $j \in \{1, \dots, n\}$.

This MOT problem has many applications throughout machine learning, computer science, and the natural sciences since it arises in tasks that require “stitching” together aggregate measurements. For instance, applications of MOT include inference from collective dynamics [38, 49], information fusion for Bayesian learning [80], averaging point clouds [2, 36], the n -coupling problem [76], quantile aggregation [60, 75], matching for teams [28, 31], image processing [74, 79], random combinatorial optimization [1, 50, 61, 65, 68, 89, 93], Distributionally Robust Optimization [30, 62, 66], simulation of incompressible fluids [17, 26], and Density Functional Theory [16, 27, 32].

However, in most applications, the success of MOT is severely limited by the lack of efficient algorithms. Indeed, in general, MOT requires *exponential time* in the number of marginals k and their support sizes n . For instance, applying a linear program solver out-of-the-box takes $n^{\Theta(k)}$ time because MOT is a linear program with n^k variables, n^k non-negativity constraints, and nk equality constraints. Specialized algorithms in the literature such as the Sinkhorn algorithm yield similar $n^{\Theta(k)}$ runtimes. Such runtimes currently limit the applicability of MOT to tiny-scale problems (e.g., $n = k = 10$).

Polynomial-time algorithms for MOT. This paper develops *polynomial-time* algorithms for MOT, where here and henceforth “polynomial” means in the number of marginals k and their support sizes n —and possibly also C_{\max}/ε for ε -additive approximation, where C_{\max} is a bound on the entries of C .

At first glance, this may seem impossible for at least two “trivial” reasons. One is that it takes exponential time to read the input cost C since it has n^k entries. We circumvent this issue by considering costs C with $\text{poly}(n, k)$ -size implicit representations, which

¹ For simplicity, all μ_i are assumed to have the same support size n . Everything in this paper extends in a straightforward way to non-uniform sizes n_i where n^k is replaced by $\prod_{i=1}^k n_i$, and $\text{poly}(n, k)$ is replaced by $\text{poly}(\max_i n_i, k)$.

encompasses essentially all MOT applications.² A second obvious issue is that it takes exponential time to write the output variable P since it has n^k entries. We circumvent this issue by returning solutions P with $\text{poly}(n, k)$ -size implicit representations, for instance sparse solutions.

But, of course, circumventing these issues of input/output size is not enough to actually solve MOT in polynomial time. See [6] for examples of NP-hard MOT problems with costs that have $\text{poly}(n, k)$ -size implicit representations.

Remarkably, for several MOT problems, there are specially-tailored algorithms that run in polynomial time—notably, for MOT problems with graphically-structured costs of constant treewidth [47, 49, 82], variational mean-field games [15], computing generalized Euler flows [14], computing low-dimensional Wasserstein barycenters [7, 14, 29], and filtering and estimation tasks based on target tracking [38, 46, 47, 49, 77]. However, the number of MOT problems that are known to be solvable in polynomial time is small, and it is unknown if these techniques can be extended to the many other MOT problems arising in applications. This motivates the central question driving this paper:

Are there general “structural properties” that make MOT solvable in $\text{poly}(n, k)$ time?

This paper is conceptually divided into two parts. In the first part of the paper, we develop a unified algorithmic framework for MOT that characterizes the structure required for different algorithms to solve MOT in $\text{poly}(n, k)$ time, in terms of simple variants of the dual feasibility oracle. This enables us to prove that some algorithms can solve MOT problems in polynomial time whenever any algorithm can; whereas the popular Sinkhorn algorithm cannot. Moreover, this algorithmic framework makes it significantly easier to design a $\text{poly}(n, k)$ time algorithm for a given MOT problem (when possible) because it now suffices to solve the dual feasibility oracle—and this is much more amenable to standard algorithmic techniques. In the second part of the paper, we demonstrate the ease-of-use of our algorithmic framework by applying it to three general classes of MOT cost structures.

Below, we detail these two parts of the paper in Sects. 1.1 and 1.2, respectively.

1.1 Contribution 1: unified algorithmic framework for MOT

In order to understand what structural properties make MOT solvable in polynomial time, we first lay a more general groundwork. The purpose of this is to understand the following fundamental questions:

- Q1 What are reasonable candidate algorithms for solving structured MOT problems in polynomial time?
- Q2 What structure must an MOT problem have for these algorithms to have polynomial runtimes?

² E.g., in the MOT problems of Wasserstein barycenters, generalized Euler flows, and Density Functional Theory, C has entries $C_{j_1, \dots, j_k} = \sum_{i, i'=1}^k g_{i, i'}(j_i, j_{i'})$ and thus can be implicitly input via the k^2 functions $g_{i, i'} : \{1, \dots, n\}^2 \rightarrow \mathbb{R}$.

Table 1 These MOT algorithms have polynomial runtime except for a bottleneck “oracle”

Algorithm	Oracle	Runtime	Always applicable?	Exact solution?	Sparse solution?	Practical?
ELLIPSOID	MIN	Theorem 4.1	Yes	Yes	Yes	No
MWU	AMIN	Theorem 4.7	Yes	No	Yes	Yes
SINKHORN	SMIN	Theorem 4.18	No	No	No	Yes

Each oracle is a simple variant of the dual feasibility oracle for MOT. The number of oracle computations is $\text{poly}(n, k)$ for ELLIPSOID, and $\text{poly}(n, k, C_{\max}/\epsilon)$ for both MWU and SINKHORN. From a theoretical perspective, the most important aspect of an algorithm is whether it can solve MOT in polynomial time if and only if any algorithm can. We show that ELLIPSOID and MWU satisfy this (Theorem 1.1), but SINKHORN does not (Theorem 1.3). From a practical perspective, SINKHORN is the most scalable when applicable.³ Code for implementing these algorithms and reproducing all numerical simulations in this paper is provided at <https://github.com/eboix/mot>

Q3 Is the structure required by a given algorithm more restrictive than the structure required by a different algorithm (or *any* algorithm)?

Q4 How to check if this structure occurs for a given MOT problem?

We detail our answers to these four questions below in Sects. 1.1.1 to 1.1.4, and then briefly discuss practical tradeoffs beyond polynomial-time solvability in Sect. 1.1.5; see Table 1 for a summary. We expect that this general groundwork will prove useful in future investigations of tractable MOT problems.

1.1.1 Answer to Q1: candidate $\text{poly}(n, k)$ -time algorithms

We consider three algorithms for MOT whose exponential runtimes can be isolated into a single bottleneck—and thus can be implemented in polynomial time whenever that bottleneck can. These algorithms are the Ellipsoid algorithm ELLIPSOID [44], the Multiplicative Weights Update algorithm MWU [91], and the natural multidimensional analog of Sinkhorn’s scaling algorithm SINKHORN [14, 70]. SINKHORN is specially tailored to MOT and is currently the predominant algorithm for it. To foreshadow our answer to Q3, the reason that we restrict to these candidate algorithms is: we show that ELLIPSOID and MWU can solve an MOT problem in polynomial time if and only if any algorithm can.

1.1.2 Answer to Q2: structure necessary to run candidate algorithms

These three algorithms only access the cost tensor C through polynomially many calls of their respective bottlenecks. Thus the structure required to implement these candidate algorithms in polynomial time is equivalent to the structure required to implement their respective bottlenecks in polynomial time.

In Sect. 4, we show that the bottlenecks of these three algorithms are polynomial-time equivalent to natural analogs of the feasibility oracle for the dual LP to MOT.

Namely, given weights $p_1, \dots, p_k \in \mathbb{R}^n$, compute

$$\min_{(j_1, \dots, j_k) \in \{1, \dots, n\}^k} C_{j_1, \dots, j_k} - \sum_{i=1}^k [p_i]_{j_i} \quad (1.1)$$

either exactly for ELLIPSOID, approximately for MWU, or with the “min” replaced by a “softmin” for SINKHORN. We call these three tasks the MIN, AMIN, and SMIN oracles, respectively. See Remark 3.4 for the interpretation of these oracles as variants of the dual feasibility oracle.

These three oracles take n^k time to implement in general. However, for a wide range of structured cost tensors C they can be implemented in $\text{poly}(n, k)$ time, see Sect. 1.2 below. For such structured costs C , our oracle abstraction immediately implies that the MOT problem with cost C and any input marginals μ_1, \dots, μ_k can be (approximately) solved in polynomial time by any of the three respective algorithms.

Our characterization of the algorithms’ bottlenecks as variations of the dual feasibility oracle has two key benefits—which are the answers to Q3 and Q4, described below.

1.1.3 Answer to Q3: characterizing what MOT problems each algorithm can solve

A key benefit of our characterization of the algorithms’ bottlenecks as variations of the dual feasibility oracles is that it enables us to establish whether the structure required by a given MOT algorithm is more restrictive than the structure required by a different algorithm (or by *any* algorithm).

In particular, this enables us to answer the natural question: why restrict to just the three algorithms described above? Can other algorithms solve MOT in $\text{poly}(n, k)$ time in situations when these algorithms cannot? Critically, the answer is no: restricting ourselves to the ELLIPSOID and MWU algorithms is at no loss of generality.

Theorem 1.1 (Informal statement of part of Theorems 4.1 and 4.7) *For any family of costs $C \in (\mathbb{R}^n)^{\otimes k}$:*

- ELLIPSOID computes an exact solution for MOT in $\text{poly}(n, k)$ time if and only if any algorithm can.
- MWU computes an ε -approximate solution for MOT in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time if and only if any algorithm can.

The statement for ELLIPSOID is implicit from classical results about LP [44] combined with arguments from [7], see the previous work section Sect. 1.3. The statement for MWU is new to this paper.

The oracle abstraction helps us show Theorem 1.1 because it reduces this question of what structure is needed for the algorithms to solve MOT in polynomial time, to the question of what structure is needed to solve their respective bottlenecks in polynomial time. Thus Theorem 1.1 is a consequence of the following result. (The “if” part of this result is a contribution of this paper; the “only if” part was shown in [6].)

Theorem 1.2 (Informal statement of part of Theorems 4.1 and 4.7) *For any family of costs $C \in (\mathbb{R}^n)^{\otimes k}$:*

- *MOT can be exactly solved in $\text{poly}(n, k)$ time if and only if MIN can.*
- *MOT can be ε -approximately solved in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time if and only if AMIN can.*

Interestingly, a further consequence of our unified algorithm-to-oracle abstraction is that it enables us to show that SINKHORN—which is currently the most popular algorithm for MOT by far—requires strictly more structure to solve an MOT problem than other algorithms require. This is in sharp contrast to the complete generality of the other two algorithms shown in Theorem 1.1.

Theorem 1.3 (Informal statement of Theorem 4.19) *Under standard complexity-theoretic assumptions, there exists a family of MOT problems that can be solved exactly in $\text{poly}(n, k)$ time using ELLIPSOID, however it is impossible to implement a single iteration of SINKHORN (even approximately) in $\text{poly}(n, k)$ time.*

The reason that our unified algorithm-to-oracle abstraction helps us show Theorem 1.3 is that it puts SINKHORN on equal footing with the other two classical algorithms in terms of their reliance on variants of the dual feasibility oracle. This reduces proving Theorem 1.3 to showing the following separation between the SMIN oracle and the other two oracles.

Theorem 1.4 (Informal statement of Lemma 3.7) *Under standard complexity-theoretic assumptions, there exists a family of cost tensors $C \in (\mathbb{R}^n)^{\otimes k}$ such that there are $\text{poly}(n, k)$ -time algorithms for MIN and AMIN, however it is impossible to solve SMIN (even approximately) in $\text{poly}(n, k)$ time.*

1.1.4 Answer to Q4: ease-of-use for checking if MOT is solvable in polynomial time

The second key benefit of this oracle abstraction is that it is helpful for showing that a given MOT problem (whose cost C is input implicitly through some concise representation) is solvable in polynomial time as it without loss of generality reduces MOT to solving any of the three corresponding oracles in polynomial time. The upshot is that these oracles are more directly amenable to standard algorithmic techniques since they are phrased as more conventional combinatorial-optimization problems. In the second part of the paper, we illustrate this ease-of-use via applications to three general classes of structured MOT problems; for an overview see Sect. 1.2.

1.1.5 Practical algorithmic tradeoffs beyond polynomial-time solvability

From a theoretical perspective, the most important aspect of an algorithm is whether it can solve MOT in polynomial time if and only if any algorithm can. As we have discussed, this is true for ELLIPSOID and MWU (Theorem 1.1) but not for SINKHORN (Theorem 1.3). Nevertheless, for a wide range of MOT cost structures, all three oracles can be implemented in polynomial time, which means that all three algorithms

ELLIPSOID, MWU, and SINKHORN can be implemented in polynomial time. Which algorithm is best in practice depends on the relative importance of the following considerations for the particular application.

- *Error.* ELLIPSOID computes exact solutions, whereas MWU and SINKHORN only compute low-precision solutions due to $\text{poly}(1/\varepsilon)$ runtime dependence.
- *Solution sparsity.* ELLIPSOID and MWU output solutions with polynomially many non-zero entries (roughly nk), whereas SINKHORN outputs fully dense solutions with n^k non-zero entries (through a polynomial-size implicit representation, see Sect. 4.3). Solution sparsity enables interpretability, visualization, and efficient downstream computation—benefits which are helpful in diverse applications, for example ranging from computer graphics [20, 71, 79] to facility location problems [10] to machine learning [7, 33] to ecological inference [64] to fluid dynamics (see Sect. 5.3), and more. Furthermore, in Sect. 7.4, we show that sparse solutions for MOT (a.k.a. linear optimization over the transportation polytope) enable efficiently solving certain non-linear optimization problems over the transportation polytope.
- *Practical runtime.* Although all three algorithms enjoy polynomial runtime guarantees, the polynomials are smaller for some algorithms than for others. In particular, SINKHORN has remarkably good scalability in practice as long the error ε is not too small and its bottleneck oracle SMIN is practically implementable. By Theorems 1.1 and 1.3, MWU can solve strictly more MOT problems in polynomial time than SINKHORN; however, it is less scalable in practice when both MWU and SINKHORN can be implemented. ELLIPSOID is not practical and is used solely as a proof of concept that problems are tractable to solve exactly; in practice, we use Column Generation (see, e.g., [18, §6.1]) rather than ELLIPSOID as it has better empirical performance, yet still has the same bottleneck oracle MIN, see Sect. 4.1.3. Column Generation is not as practically scalable as SINKHORN in n and k but has the benefit of computing exact, sparse solutions.

To summarize: which algorithm is best in practice depends on the application. For example, Column Generation produces the qualitatively best solutions for the fluid dynamics application in Sect. 5.3, SINKHORN is the most scalable for the risk estimation application in Sect. 7.3, and MWU is the most scalable for the network reliability application in Sect. 6.3 (for that application there is no known implementation of SINKHORN that is practically efficient).

1.2 Contribution 2: applications to general classes of structured MOT problems

In the second part of the paper, we illustrate the algorithmic framework developed in the first part of the paper by applying it to three general classes of MOT cost structures:

1. Graphical structure (in Sect. 5).
2. Set-optimization structure (in Sect. 6).
3. Low-rank plus sparse structure (in Sect. 7).

Specifically, if the cost C is structured in any of these three ways, then MOT can be (approximately) solved in $\text{poly}(n, k)$ time for any input marginals μ_1, \dots, μ_k .

Table 2 In the second part of the paper, we illustrate the ease-of-use of our algorithmic framework by applying it to three general classes of MOT cost structures. These structures encompass many—if not most—current applications of MOT

Structure	Definition	Complexity measure	Polynomial-time algorithm?	
			Approximate	Exact
Graphical (Sect. 5)	$C_{\vec{j}} = \sum_{S \in \mathcal{S}} f_S(\vec{j}_S)$	Treewidth	Known [49, 82]	Corollary 5.6
Set- optimization (Sect. 6)	$C_{\vec{j}} = \mathbb{1}[\vec{j} \notin S]$	Optimization oracle over S	Corollary 6.9	Corollary 6.9
Low-rank + sparse (Sect. 7)	$C = R + S$	Rank of R , sparsity of S	Corollary 7.5	Unknown

Previously, it was known how to solve MOT problems with structure (1) using SINKHORN [49, 82], but this only computes solutions that are dense (with n^k non-zero entries) and low-precision (due to $\text{poly}(1/\epsilon)$ runtime dependence). We therefore provide the first solutions that are sparse and exact for structure (1). For structures (2) and (3), we provide the first polynomial-time algorithms, even for approximate computation. These three structures are incomparable: it is in general not possible to model a problem falling under any of the three structures in a non-trivial way using any of the others, for details see Remarks 6.7 and 7.3. This means that the new structures (2) and (3) enable capturing a wide range of new applications.

Below, we detail these structures individually in Sects. 1.2.1, 1.2.2, and 1.2.3. See Table 2 for a summary.

1.2.1 Graphical structure

In Sect. 5, we apply our algorithmic framework to MOT problems with graphical structure, a broad class of MOT problems that have been previously studied [47, 49, 82]. Briefly, an MOT problem has graphical structure if its cost tensor C decomposes as

$$C_{j_1, \dots, j_k} = \sum_{S \in \mathcal{S}} f_S(\vec{j}_S),$$

where $f_S(\vec{j}_S)$ are arbitrary “local interactions” that depend only on tuples $\vec{j}_S := \{j_i\}_{i \in S}$ of the k variables.

In order to derive efficient algorithms, it is necessary to restrict how local the interactions are because otherwise MOT is NP-hard (even if all interaction sets $S \in \mathcal{S}$ have size 2) [6]. We measure the locality of the interactions via the standard complexity measure of the “treewidth” of the associated graphical model. See Sect. 5.1 for formal definitions. While the runtimes of our algorithms (and all previous algorithms) depend exponentially on the treewidth, we emphasize that the treewidth is a very small constant

(either 1 or 2) in all current applications of MOT falling under this framework; see the related work section.

We show that for MOT cost tensors that have graphical structure of constant treewidth, all three oracles can be implemented in $\text{poly}(n, k)$ time. We accomplish this by leveraging the known connection between graphically structured MOT and graphical models shown in [49]. In particular, the MIN, AMIN, and SMIN oracles are respectively equivalent to the mode, approximate mode, and log-partition function of an associated graphical model. Thus we can implement all oracles in $\text{poly}(n, k)$ time by simply applying classical algorithms from the graphical models community [55, 88].

Theorem 1.5 (Informal statement of Theorem 5.5) *Let $C \in (\mathbb{R}^n)^{\otimes k}$ have graphical structure of constant treewidth. Then the MIN, AMIN, and SMIN oracles can be computed in $\text{poly}(n, k)$ time.*

It is an immediate corollary of Theorem 1.5 and our algorithms-to-oracles reduction described in Sect. 1.1 that one can implement ELLIPSOID, MWU, and SINKHORN in polynomial time. Below, we record the theoretical guarantee of ELLIPSOID since it is the best of the three algorithms as it computes exact, sparse solutions.

Theorem 1.6 (Informal statement of Corollary 5.6) *Let $C \in (\mathbb{R}^n)^{\otimes k}$ have graphical structure of constant treewidth. Then an exact, sparse solution for MOT can be computed in $\text{poly}(n, k)$ time.*

Previously, it was known how to solve such MOT problems [49, 82] using SINKHORN, but this only computes a solution that is fully dense (with n^k non-zero entries) and low-precision (due to $\text{poly}(1/\varepsilon)$ runtime dependence). Details in the related work section. Our result improves over this state-of-the-art algorithm by producing solutions that are *exact* and *sparse* in $\text{poly}(n, k)$ time.

In Sect. 5.3, we demonstrate the benefit of Theorem 1.6 on the application of computing generalized Euler flows, which was historically the motivation of MOT and has received significant attention, e.g., [14, 17, 23–26]. While there is a specially-tailored version of the SINKHORN algorithm for this problem that runs in polynomial time [14, 17], it produces solutions that are approximate and fully dense. Our algorithm produces exact, sparse solutions which lead to sharp visualizations rather than blurry ones (see Fig. 4).

1.2.2 Set-optimization structure

In Sect. 6, we apply our algorithmic framework to MOT problems whose cost tensors C take value 0 or 1 in each entry. That is, costs C of the form

$$C_{j_1, \dots, j_k} = \mathbb{1}[(j_1, \dots, j_k) \notin S],$$

for some subset $S \subseteq [n]^k$. Such MOT problems arise naturally in applications where one seeks to minimize the probability that some event S occurs, given marginal probabilities on each variable j_i , see Example 6.1.

In order to derive efficient algorithms, it is necessary to restrict the (otherwise arbitrary) set S . We parametrize the complexity of such MOT problems via the complexity of finding the minimum-weight object in S . This opens the door to combinatorial applications of MOT because finding the minimum-weight object in S is well-known to be polynomial-time solvable for many “combinatorially-structured” sets S of interest—e.g., the set S of cuts in a graph, or the set S of independent sets in a matroid.

We show that for MOT cost tensors with this structure, all three oracles can be implemented efficiently.

Theorem 1.7 (Informal statement of Theorem 6.8) *Let $C \in (\mathbb{R}^n)^{\otimes k}$ have set-optimization structure. Then the MIN, AMIN, and SMIN oracles can be computed in $\text{poly}(n, k)$ time.*

It is an immediate corollary of Theorem 1.7 and our algorithms-to-oracles reduction described in Sect. 1.1 that one can implement ELLIPSOID, MWU, and SINKHORN in polynomial time. Below, we record the theoretical guarantee for ELLIPSOID since it is the best of these three algorithms as it computes exact, sparse solutions.

Theorem 1.8 (Informal statement of Corollary 6.9) *Let $C \in (\mathbb{R}^n)^{\otimes k}$ have set-optimization structure. Then an exact, sparse solution for MOT can be computed in $\text{poly}(n, k)$ time.*

This is the first polynomial-time algorithm for this class of MOT problems. We note that a more restrictive class of MOT problems was studied in [93] under the additional restriction that S is upwards-closed.

In Sect. 6.3, we show how this general class of set-optimization structure captures, for example, the classical application of computing the extremal reliability of a network with stochastic edge failures. Network reliability is a fundamental topic in network science and engineering [12, 13, 41] which is often studied in an average-case setting where each edge fails independently with some given probability [52, 63, 72, 85]. The application investigated here is a robust notion of network reliability in which edge failures may be maximally correlated (e.g., by an adversary) or minimally correlated (e.g., by a network maintainer) subject to a marginal constraint on each edge’s failure probability, a setting that dates back to the 1980s [89, 93]. We show how to express both the minimally and maximally correlated network reliability problems as MOT problems with set-optimization structure, recovering as a special case of our general framework the known polynomial-time algorithms in [89, 93] as well as more practical polynomial-time algorithms that scale to input sizes that are an order-of-magnitude larger.

1.2.3 Low-rank and sparse structure

In Sect. 7, we apply our algorithmic framework to MOT problems whose cost tensors C decompose as

$$C = R + S,$$

where R is a constant-rank tensor, and S is a polynomially-sparse tensor. We assume that R is represented in factored form, and that S is represented through its non-zero entries, which overall yields a $\text{poly}(n, k)$ -size representation of C .

We show that for MOT cost tensors with low-rank plus sparse structure, the AMIN and SMIN oracles can be implemented in polynomial time.³ This may be of independent interest because, by taking all oracle inputs $p_i = 0$ in (1.1), this generalizes the previously open problem of approximately computing the smallest entry of a constant-rank tensor with n^k entries in $\text{poly}(n, k)$ time.

Theorem 1.9 (Informal statement of Theorem 7.4) *Let $C \in (\mathbb{R}^n)^{\otimes k}$ have low-rank plus sparse structure. Then the AMIN and SMIN oracles can be computed in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time.*

It is an immediate corollary of Theorem 1.9 and our algorithms-to-oracles reduction described in Sect. 1.1 that one can implement MWU and SINKHORN in polynomial time. Of these two algorithms, MWU computes sparse solutions, yielding the following theorem.

Theorem 1.10 (Informal statement of Corollary 7.5) *Let $C \in (\mathbb{R}^n)^{\otimes k}$ have low-rank plus sparse structure. Then a sparse, ε -approximate solution for MOT can be computed in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time.*

This is the first polynomial-time result for this class of MOT problems. We note that the runtime of our MOT algorithm depends exponentially on the rank r of R , hence why we take r to be constant. Nevertheless, such a restriction on the rank is unavoidable since unless $\text{P} = \text{NP}$, there does not exist an algorithm with runtime that is jointly polynomial in n, k , and the rank r [6].

We demonstrate this polynomial-time algorithm concretely on two applications. First, in Sect. 7.3 we consider the risk estimation problem of computing an investor's expected profit in the worst-case over all future prices that are consistent with given marginal distributions. We show that this is equivalent to an MOT problem with a low-rank tensor and thereby provide the first efficient algorithm for it.

Second, in Sect. 7.4, we consider the fundamental problem of projecting a joint distribution Q onto the transportation polytope. We provide the first polynomial-time algorithm for solving this when Q decomposes into a constant-rank and sparse component, which models mixtures of product distributions with polynomially many corruptions. This application illustrates the versatility of our algorithmic results beyond polynomial-time solvability of MOT, since this projection problem is a *quadratic* optimization over the transportation polytope rather than linear optimization (a.k.a. MOT). In order to achieve this, we develop a simple quadratic-to-linear reduction tailored to this problem that crucially exploits the sparsity of the MOT solutions enabled by the MWU algorithm.

³ It is an interesting open question if the MIN oracle can similarly be implemented in $\text{poly}(n, k)$ time. This would enable implementing ELLIPSOID in $\text{poly}(n, k)$ time by our algorithms-to-oracles reduction, and thus would enable computing exact solutions for this class of MOT problems (cf., Theorem 1.10).

1.3 Related work

1.3.1 MOT algorithms

MOT algorithms fall into two categories. One category consists of general-purpose algorithms that do not depend on the specific MOT cost. For example, this includes running an LP solver out-of-the-box, or running the Sinkhorn algorithm where in each iteration one sums over all n^k entries of the cost tensor to implement the marginalization bottleneck [40, 58, 84]. These approaches are robust in the sense that they do not need to be changed based on the specific MOT problem. However, they are impractical beyond tiny input sizes (e.g., $n = k = 10$) because their runtimes scale as $n^{\Omega(k)}$.

The second category consists of algorithms that are much more scalable but require extra structure of the MOT problem. Specifically, these are algorithms that somehow exploit the structure of the relevant cost tensor C in order to (approximately) solve an MOT problem in $\text{poly}(n, k)$ time [1, 7, 14, 15, 17, 29, 30, 38, 46, 47, 49, 50, 61, 62, 65–68, 77, 82, 89, 93]. Such a $\text{poly}(n, k)$ runtime is far more tractable—but it is not well understood for which MOT problems such a runtime is possible. The purpose of this paper is to clarify this question.

To contextualize our answer to this question with the rapidly growing literature requires further splitting this second category of algorithms.

Sinkhorn algorithm. Currently, the predominant approach in the second category is to solve an entropically regularized version of MOT with the Sinkhorn algorithm, a.k.a. Iterative Proportional Fitting or Iterative Bregman Projections or RAS algorithm or Iterative Scaling algorithm, see e.g., [15–17, 47, 49, 67, 82]. Recent work has shown that a polynomial number of iterations of this algorithm suffices [40, 58, 84]. However, the bottleneck is that each iteration requires n^k operations in general because it requires marginalizing a tensor with n^k entries. The critical question is therefore: what structure of an MOT problem enables implementing this marginalization bottleneck in polynomial time.

This paper makes two contributions to this question. First, we identify new broad classes of MOT problems for which this bottleneck can be implemented in polynomial time, and thus SINKHORN can be implemented in polynomial time (see Sect. 1.2). Second, we propose other algorithms that require strictly less structure than SINKHORN does in order to solve an MOT problem in polynomial time (Theorem 4.19).

Ellipsoid algorithm. The Ellipsoid algorithm is among the most classical algorithms for implicit LP [43, 44, 53], however it has taken a back seat to the SINKHORN algorithm in the vast majority of the MOT literature.

In Sect. 4.1, we make explicit the fact that the variant of ELLIPSOID from [7] can solve MOT exactly in $\text{poly}(n, k)$ time if and only if any algorithm can (Theorem 4.1). This is implicit from combining several known results [6, 7, 44]. In the process of making this result explicit, we exploit the special structure of the MOT LP to significantly simplify the reduction from the dual violation oracle to the dual feasibility oracle. The previously known reduction is highly impractical as it requires an indirect “back-and-forth” use of the Ellipsoid algorithm [44, page 107]. In contrast, our reduction is direct and simple; this is critical for implementing our practical alternative to ELLIPSOID, namely COLGEN, with the dual feasibility oracle.

Multiplicative Weights Update algorithm. This algorithm, first introduced by [91], has been studied in the context of optimal transport when $k = 2$ [19, 73], in which case implicit LP is not necessary for a polynomial runtime. MWU lends itself to implicit LP [91], but is notably absent from the MOT literature.

In Sect. 4.2, we show that MWU can be applied to MOT in polynomial time if and only if the approximate dual feasibility oracle can be solved in polynomial time. To do this, we show that in the special case of MOT, the well-known “softmax-derivative” bottleneck of MWU is polynomial-time equivalent to the approximate dual feasibility oracle. Since it is known that the approximate dual feasibility oracle is polynomial-time reducible to approximate MOT [6], we therefore establish that MWU can solve MOT approximately in polynomial time if and only if any algorithm can (Theorem 4.7).

1.3.2 Graphically structured MOT problems with constant treewidth

We isolate here graphically structured costs with constant treewidth because this framework encompasses all MOT problems that were previously known to be tractable in polynomial time [49, 82], with the exceptions of the fixed-dimensional Wasserstein barycenter problem and MOT problems related to combinatorial optimization—both of which are described below in Sect. 1.3.3. This family of graphical structured costs with treewidth 1 (a.k.a. “tree-structured costs” [47]) includes applications in economics such as variational mean-field games [15], interpolating histograms on trees [3], matching for teams [29, 67]; as well as encompasses applications in filtering and estimation for collective dynamics such as target tracking [38, 46, 47, 49, 77] and Wasserstein barycenters in the case of fixed support [14, 29, 38, 67]. With treewidth 2, this family of costs also includes dynamic multi-commodity flow problems [48], as well as the application of computing generalized Euler flows in fluid dynamics [14, 17, 67], which was historically the original motivation of MOT [23–26].

Previous polynomial-time algorithms for graphically structured MOT compute approximate, dense solutions. Implementing SINKHORN for graphically structured MOT problems by using belief propagation to efficiently implement the marginalization bottleneck was first proposed twenty years ago in [82]. There have been recent advancements in understanding connections of this algorithm to the Schrödinger bridge problem in the case of trees [47], as well as developing more practically efficient single-loop variations [49].

All of these works prove theoretical runtime guarantees only in the case of tree structure (i.e., treewidth 1). However, this graphical model perspective for efficiently implementing SINKHORN readily extends to any constant treewidth: simply implement the marginalization bottleneck using junction trees. This, combined with the iteration complexity of SINKHORN which is known to be polynomial [40, 58, 84], immediately yields an overall polynomial runtime. This is why we cite [49, 82] throughout this paper regarding the fact that SINKHORN can be implemented in polynomial time for graphical structure with any constant treewidth.

While the use of SINKHORN for graphically structured MOT is mathematically elegant and can be impressively scalable in practice, it has two drawbacks. The first drawback of this algorithm is that it computes (implicit representations of)

solutions that are fully dense with n^k non-zero entries. Indeed, it is well-known that SINKHORN finds the unique optimal solution to the entropically regularized MOT problem $\min_{P \in \mathcal{M}(\mu_1, \dots, \mu_k)} \langle P, C \rangle - \eta^{-1} H(P)$, and that this solution is fully dense [70]. For example, in the simple case of cost $C = 0$, uniform marginals μ_i , and any strictly positive regularization parameter $\eta > 0$, this solution P has value $1/n^k$ in each entry.

The second drawback of this algorithm is that it only computes solutions that are low-precision due to $\text{poly}(1/\varepsilon)$ runtime dependence on the accuracy ε . This is because the number of SINKHORN iterations is known to scale polynomially in the entropic regularization parameter η even in the matrix case $k = 2$ [59, §1.2], and it is known that $\eta = \Omega(\varepsilon^{-1} k \log n)$ is necessary for the converged solution of SINKHORN to be an ε -approximate solution to the (unregularized) original MOT problem [58].

Improved algorithms for graphically structured MOT problems. The contribution of this paper to the study of graphically structured MOT problems is that we give the first $\text{poly}(n, k)$ time algorithms that can compute solutions which are exact and sparse (Corollary 5.6). Our framework also directly recovers all known results about SINKHORN for graphically structured MOT problems—namely that it can be implemented in polynomial time for trees [47, 82] and for constant treewidth [49, 82].

1.3.3 Tractable MOT problems beyond graphically structured costs

The two new classes of MOT problems studied in this paper—namely, set-optimization structure and low-rank plus sparse structure—are incomparable to each other as well as to graphical structure. Details in Remarks 6.7 and 7.3. This lets us handle a wide range of new MOT problems that could not be handled before.

There are two other classes of MOT problems studied in the literature which do not fall under the three structures studied in this paper. We elaborate on both below.

Remark 1.11 (Low-dimensional Wasserstein barycenter) This MOT problem has cost $C_{j_1, \dots, j_k} = \sum_{i, i'=1}^k \|x_{i, j_i} - x_{i', j_{i'}}\|^2$ where $x_{i, j} \in \mathbb{R}^d$ denotes the j -th atom in the distribution μ_i . Clearly this cost is not a graphically structured cost of constant treewidth—indeed, representing it through the lens of graphical structure requires the complete graph of interactions, which means a maximal treewidth of $k - 1$.⁴ This problem also does not fall under the set-optimization or constant-rank structures. Nevertheless, this MOT problem can be solved in $\text{poly}(n, k)$ time for any fixed dimension d by exploiting the low-dimensional geometric structure of the points $\{x_{i, j}\}$ that implicitly define the cost [7].

Remark 1.12 (Random combinatorial optimization) MOT problems also appear in the random combinatorial optimization literature since the 1970s, see e.g., [50, 61, 65, 89,

⁴ We remark that the related but different problem of *fixed-support* Wasserstein barycenters has graphical structure with treewidth 1 [14, 29, 38, 67]. However, it should be emphasized that the fixed-support Wasserstein barycenter problem is different from the Wasserstein barycenter problem: it only approximates the latter to ε accuracy if the fixed support is restricted to an $O(\varepsilon)$ -net which requires $n = 1/\varepsilon^{\Omega(d)}$ discretization size for the barycenter's support, and thus (i) even in constant dimension, does not lead to high-precision algorithms due to $\text{poly}(1/\varepsilon)$ runtime; and (ii) scales exponentially in the dimension d . See [8, §1.3] for further details about the complexity of Wasserstein barycenters.

93], although under a different name and in a different community. These papers consider MOT problems with costs of the form $C(x) = \min_{v \in V} \langle x, v \rangle$ for polytopes $V \subseteq \{0, 1\}^k$ given through a list of their extreme points. Applications include PERT (Program Evaluation and Review Technique), extremal network reliability, and scheduling. Recently, applications to Distributionally Robust Optimization were investigated in [30, 62, 66] which considered general polytopes $V \subset \mathbb{R}^k$, as well as in [68] which considered MOT costs of the related form $C(x) = \mathbb{1}[\min_{v \in V} \langle x, v \rangle \geq t]$, and in [1] which considers other combinatorial costs C such as sub/supermodular functions. These papers show that these random combinatorial optimization problems are in general intractable, and give sufficient conditions on when they can be solved in polynomial time. In general, these families of MOT problems are different from the three structures studied in this paper, although some MOT applications fall under multiple umbrellas (e.g., extremal network reliability). It is an interesting question to understand to what extent these structures can be reconciled (as well as the algorithms, which sometimes use extended formulations in these papers).

1.3.4 Intractable MOT problems

These algorithmic results beg the question: what are the fundamental limitations of this line of work on polynomial-time algorithms for structured MOT problems? To this end, the recent paper [6] provides a systematic investigation of NP-hardness results for structured MOT problems, including converses to several results in this paper. In particular, [6, Propositions 4.1 and 4.2] justify the constant-rank regime studied in Sect. 7 by showing that unless $P = NP$, there does not exist an algorithm with runtime that is jointly polynomially in the rank r and the input parameters n and k . Similarly, [6, Propositions 5.1 and 5.2] justify the constant-treewidth regime for graphically structured costs studied in Sect. 5 and all previous work by showing that unless $P = NP$, there does not exist an algorithm with polynomial runtime even in the seemingly simple class of MOT costs that decompose into pairwise interactions $C_{j_1, \dots, j_k} = \sum_{i \neq i' \in [k]} c_{i, i'}(j_i, j_{i'})$. The paper [6] also shows NP-hardness for several MOT problems with repulsive costs, including for example the MOT formulation of Density Functional Theory with Coulomb-Buckingham potential. It is an problem whether the Coulomb potential, studied in [16, 27, 32], also leads to an NP-hard MOT problem [6, Conjecture 6.4].

1.3.5 Variants of MOT

The literature has studied several other variants of the MOT problem, notably with entropic regularization and/or with constraints on a subset of the k marginals, see, e.g., [14–17, 38, 46–49, 58, 77]. Our techniques readily apply with little change. Briefly, to handle entropic regularization, simply use the SMIN oracle and SINKHORN algorithm with fixed regularization parameter $1/\eta > 0$ (rather than $1/\eta$ of vanishing size $\Theta(\varepsilon/\log n)$) as described in Sect. 4.3. And to handle partial marginal constraints, essentially the only change is that in the MIN, AMIN, and SMIN oracles, the potentials p_i are zero for all indices $i \in [k]$ corresponding to unconstrained marginals $m_i(P)$.

Full details are omitted for brevity since they are straightforward modifications of our main results.

1.3.6 Optimization over joint distributions

Optimization problems over exponential-size joint distributions appear in many domains. For instance, they arise in game theory when computing correlated equilibria [69]; however, in that case the optimization has different constraints which lead to different algorithms. Such problems also arise in variational inference [87]; however, the optimization there typically constrains this distribution to ensure tractability (e.g., mean-field approximation restricts to product distributions). The different constraints in these optimization problems over joint distributions versus MOT lead to significant differences in computational complexity, and thus also necessitate different algorithmic techniques.

1.4 Organization

In Sect. 2 we recall preliminaries about MOT and establish notation. The first part of the paper then establishes our unified algorithmic framework for MOT. Specifically, in Sect. 3 we define and compare three variants of the dual feasibility oracle; and in Sect. 4 we characterize the structure that MOT algorithms require for polynomial-time implementation in terms of these three oracles. For an overview of these results, see Sect. 1.1. The second part of the paper applies this algorithmic framework to three general classes of MOT cost structures: graphical structure (Sect. 5), set-optimization structure (Sect. 6), and low-rank plus sparse structure (Sect. 7). For an overview of these results, see Sect. 1.2. These three application sections are independent of each other and can be read separately. We conclude in Sect. 8.

2 Preliminaries

General notation. The set $\{1, \dots, n\}$ is denoted by $[n]$. For shorthand, we write $\text{poly}(t_1, \dots, t_m)$ to denote a function that grows at most polynomially fast in those parameters. Throughout, we assume for simplicity of exposition that all entries of the input C and μ_1, \dots, μ_k have bit complexity at most $\text{poly}(n, k)$, and same with the components defining C in structured settings. As such, throughout runtimes refer to the number of arithmetic operations. The set $\mathbb{R} \cup \{-\infty\}$ is denoted by $\tilde{\mathbb{R}}$, and note that the value $-\infty$ can be represented efficiently by adding a single flag bit. We use the standard $O(\cdot)$ and $\Omega(\cdot)$ notation, and use $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$ to denote that polylogarithmic factors may be omitted.

Tensor notation. The k -fold tensor product space $\mathbb{R}^n \otimes \dots \otimes \mathbb{R}^n$ is denoted by $(\mathbb{R}^n)^{\otimes k}$, and similarly for $(\mathbb{R}_{\geq 0}^n)^{\otimes k}$. Let $P \in (\mathbb{R}^n)^{\otimes k}$. Its i -th marginal, $i \in [k]$, is denoted by $m_i(P) \in \mathbb{R}^n$ and has entries $[m_i(P)]_j := \sum_{j_1, \dots, j_{i-1}, j_{i+1}, \dots, j_k} P_{j_1, \dots, j_{i-1}, j, j_{i+1}, \dots, j_k}$. For shorthand, we often denote an index (j_1, \dots, j_k) by \vec{j} . The sum of P 's entries is

denoted by $m(P) = \sum_{\vec{j}} P_{\vec{j}}$. The maximum absolute value of P 's entries is denoted by $\|P\|_{\max} := \max_{\vec{j}} |P_{\vec{j}}|$, or simply P_{\max} for short. For $\vec{j} \in [n]^k$, we write $\delta_{\vec{j}}$ to denote the tensor with value 1 at entry \vec{j} , and 0 elsewhere. The operations \odot and \otimes respectively denote the entrywise product and the Kronecker product. The notation $\otimes_{i=1}^k d_i$ is shorthand for $d_1 \otimes \cdots \otimes d_k$. A non-standard notation we use throughout is that $f[P]$ denotes a function $f : \mathbb{R} \rightarrow \mathbb{R}$ (typically \exp , \log , or a polynomial) applied *entrywise* to a tensor P .

2.1 Multimarginal optimal transport

The transportation polytope between measures $\mu_1, \dots, \mu_k \in \Delta_n$ is

$$\mathcal{M}(\mu_1, \dots, \mu_k) := \left\{ P \in (\mathbb{R}_{\geq 0}^n)^{\otimes k} : m_i(P) = \mu_i, \forall i \in [k] \right\}. \tag{2.1}$$

For a fixed cost $C \in (\mathbb{R}^n)^{\otimes k}$, the MOT_C problem is to solve the following linear program, given input measures $\mu = (\mu_1, \dots, \mu_k) \in (\Delta_n)^k$:

$$\min_{P \in \mathcal{M}(\mu_1, \dots, \mu_k)} \langle P, C \rangle. \tag{MOT}$$

In the $k = 2$ matrix case, (MOT) is the Kantorovich formulation of OT [86]. Its dual LP is

$$\max_{p_1, \dots, p_k \in \mathbb{R}^n} \sum_{i=1}^k \langle p_i, \mu_i \rangle \quad \text{subject to} \quad C_{j_1, \dots, j_k} - \sum_{i=1}^k [p_i]_{j_i} \geq 0, \forall (j_1, \dots, j_k) \in [n]^k. \tag{MOT-D}$$

A basic, folklore fact about MOT is that it always has a sparse optimal solution (e.g., [10, Lemma 3]). This follows from elementary facts about standard-form LP; we provide a short proof for completeness.

Lemma 2.1 (Sparse solutions for MOT) *For any cost $C \in (\mathbb{R}^n)^{\otimes k}$ and any marginals $\mu_1, \dots, \mu_k \in \Delta_n$, there exists an optimal solution P to $\text{MOT}_C(\mu)$ that has at most $nk - k + 1$ non-zero entries.*

Proof Since (MOT) is an LP over a compact domain, it has an optimal solution at a vertex [18, Theorem 2.7]. Since (MOT) is a standard-form LP, these vertices are in correspondence with basic solutions, thus their sparsity is bounded by the number of linearly dependent constraints defining $\mathcal{M}(\mu_1, \dots, \mu_k)$ [18, Theorem 2.4]. We bound this quantity by $nk - k + 1$ via two observations. First, $\mathcal{M}(\mu_1, \dots, \mu_k)$ is defined by nk equality constraints $[m_i(P)]_j = [\mu_i]_j$ in (2.1), one for each coordinate $j \in [n]$ of each marginal constraint $i \in [k]$. Second, at least $k - 1$ of these constraints are linearly dependent because we can construct k distinct linear combinations of them, namely $\sum_{j \in [n]} [m_i(P)]_j = \sum_{j \in [n]} [\mu_i]_j$ for each marginal $i \in [k]$, which all simplify to the same constraint $m(P) = 1$, and thus are redundant with each other. \square

Definition 2.2 (ε -approximate MOT solution) P is an ε -approximate solution to $\text{MOT}_C(\mu)$ if P is feasible (i.e., $P \in \mathcal{M}(\mu_1, \dots, \mu_k)$) and $\langle C, P \rangle$ is at most ε more than the optimal value.

2.2 Regularization

We introduce two standard regularization operators. First is the Shannon entropy $H(P) := -\sum_j P_j \log P_j$ of a tensor $P \in (\mathbb{R}_{\geq 0}^n)^{\otimes k}$ with entries summing to $m(P) = 1$. We adopt the standard notational convention that $0 \log 0 = 0$. Second is the softmin operator, which is defined for parameter $\eta > 0$ as

$$\text{smin}_{i \in [m]} a_i := -\frac{1}{\eta} \log \left(\sum_{i=1}^m e^{-\eta a_i} \right). \tag{2.2}$$

This softmin operator naturally extends to $a_i \in \mathbb{R} \cup \{\infty\}$ by adopting the standard notational conventions that $e^{-\infty} = 0$ and $\log 0 = -\infty$.

We make use of the following folklore fact, which bounds the error between the min and smin operators based on the regularization and the number of points. For completeness, we provide a short proof.

Lemma 2.3 (Softmin approximation bound) *For any $a_1, \dots, a_m \in \mathbb{R} \cup \{\infty\}$ and $\eta > 0$,*

$$\min_{i \in [m]} a_i \geq \text{smin}_{i \in [m]} a_i \geq \min_{i \in [m]} a_i - \frac{\log m}{\eta}.$$

Proof Assume without loss of generality that all a_i are finite, else a_i can be dropped (if all $a_i = \infty$ then the claim is trivial). For shorthand, denote $\min_{i \in [m]} a_i$ by a_{\min} . For the first inequality, use the non-negativity of the exponential function to bound

$$\text{smin}_{i \in [m]} a_i = -\frac{1}{\eta} \log \left(\sum_{i=1}^m e^{-\eta a_i} \right) \leq -\frac{1}{\eta} \log (e^{-\eta a_{\min}}) = a_{\min}.$$

For the second inequality, use the fact that each $a_i \geq a_{\min}$ to bound

$$\text{smin}_{i \in [m]} a_i = -\frac{1}{\eta} \log \left(\sum_{i=1}^m e^{-\eta a_i} \right) \geq -\frac{1}{\eta} \log (m e^{-\eta a_{\min}}) = a_{\min} - \frac{\log m}{\eta}.$$

□

The entropically regularized MOT problem (RMOT for short) is the convex optimization problem

$$\min_{P \in \mathcal{M}(\mu_1, \dots, \mu_k)} \langle P, C \rangle - \eta^{-1} H(P). \tag{RMOT}$$

This is the natural multidimensional analog of entropically regularized OT, which has a rich literature in statistics [57] and transportation theory [90], and has recently attracted significant interest in machine learning [35, 70]. The convex dual of (RMOT) is the convex optimization problem

$$\max_{p_1, \dots, p_k \in \mathbb{R}^n} \sum_{i=1}^k \langle p_i, \mu_i \rangle + \operatorname{smin}_{\vec{j} \in [n]^k} \left(C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i} \right). \tag{RMOT-D}$$

In contrast to MOT, there is no analog of Lemma 2.1 for RMOT: the unique optimal solution to RMOT is dense. Further, this solution may not even be “approximately” sparse. For example, when $C = 0$, all μ_i are uniform, and $\eta > 0$ is any positive number, the solution is fully dense with all entries equal to $1/n^k$.

We define P to be an ε -approximate RMOT solution in the analogous way as in Definition 2.2. A basic, folklore fact about RMOT is that if the regularization η is sufficiently large, then RMOT and MOT are equivalent in terms of approximate solutions.

Lemma 2.4 (MOT and RMOT are close for large regularization η) *Let $P \in \mathcal{M}(\mu_1, \dots, \mu_k)$, $\varepsilon > 0$, and $\eta \geq \varepsilon^{-1}k \log n$. If P is an ε -approximate solution to (RMOT), then P is also a (2ε) -approximate solution to (MOT); and vice versa.*

Proof Since a discrete distribution supported on n^k atoms has entropy at most $k \log n$ [34], the objectives of (MOT) and (RMOT) differ pointwise by at most $\eta^{-1}k \log n \leq \varepsilon$. Since (MOT) and (RMOT) also have the same feasible sets, their optimal values therefore differ by at most ε . □

3 Oracles

Here we define the three oracle variants described in the introduction and discuss their relations. In the below definitions, let $C \in (\mathbb{R}^n)^{\otimes k}$ be a cost tensor.

Definition 3.1 (*MIN oracle*) For weights $p = (p_1, \dots, p_k) \in \mathbb{R}^{n \times k}$, $\operatorname{MIN}_C(p)$ returns

$$\min_{\vec{j} \in [n]^k} C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i}.$$

Definition 3.2 (*AMIN oracle*) For weights $p = (p_1, \dots, p_k) \in \mathbb{R}^{n \times k}$ and accuracy $\varepsilon > 0$, $\operatorname{AMIN}_C(p, \varepsilon)$ returns $\operatorname{MIN}_C(p)$ up to additive error ε .

Definition 3.3 (*SMIN oracle*) For weights $p = (p_1, \dots, p_k) \in \bar{\mathbb{R}}^{n \times k}$ and regularization parameter $\eta > 0$, $\operatorname{SMIN}_C(p, \eta)$ returns

$$\operatorname{smin}_{\vec{j} \in [n]^k} C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i}.$$

An algorithm is said to “solve” or “implement” MIN_C if given input p , it outputs $\text{MIN}_C(p)$. Similarly for AMIN_C and SMIN_C . Note that the weights p that are input to SMIN have values inside $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty\}$; this simplifies the notation in the treatment of the SINKHORN algorithm below and does not increase the bit-complexity by more than 1 bit by adding a flag for the value $-\infty$.

Remark 3.4 (Interpretation as variants of the dual feasibility oracle) These three oracles can be viewed as variants of the feasibility oracle for (MOT-D). For $\text{MIN}_C(p)$, this relationship is exact: $p \in \mathbb{R}^{n \times k}$ is feasible for (MOT-D) if and only if $\text{MIN}_C(p)$ is non-negative. For AMIN_C and SMIN_C , this relationship is approximate, with the approximation depending on how small ε is and how large η is, respectively.

Since these oracles form the respective bottlenecks of all algorithms from the MOT and implicit linear programming literatures (see the overview in the introduction Sect. 1.1), an important question is: if one oracle can be implemented in $\text{poly}(n, k)$ time, does this imply that the other can be too?

Two reductions are straightforward: the AMIN oracle can be implemented in $\text{poly}(n, k)$ time whenever either the MIN oracle or the SMIN oracle can be implemented in $\text{poly}(n, k)$ time. We record these simple observations in remarks for easy recall.

Remark 3.5 (MIN implies AMIN) For any accuracy $\varepsilon > 0$, the $\text{MIN}_C(p)$ oracle provides a valid answer to the $\text{AMIN}_C(p, \varepsilon)$ oracle by definition.

Remark 3.6 (SMIN implies AMIN) For any $p \in \mathbb{R}^{n \times k}$ and regularization $\eta \geq \varepsilon^{-1}k \log n$, the $\text{SMIN}_C(p, \eta)$ oracle provides a valid answer to the $\text{AMIN}_C(p, \varepsilon)$ oracle due to the approximation property of the smin operator (Lemma 2.3).

In the remainder of this section, we show a separation between the SMIN oracle and both the MIN and AMIN oracles by exhibiting a family of cost tensors C for which there exist polynomial-time algorithms for MIN and AMIN, however there is no polynomial-time algorithm for SMIN. The non-existence of a polynomial-time algorithm of course requires a complexity theoretic assumption; our result holds under #BIS-hardness—which is a by-now standard complexity theory assumption introduced in [37], and in words is the statement that there does not exist a polynomial-time algorithm for counting the number of independent sets in a bipartite graph.

Lemma 3.7 (Restrictiveness of the SMIN oracle) *There exists a family of costs $C \in (\mathbb{R}^n)^{\otimes k}$ for which MIN_C and AMIN_C can be solved in $\text{poly}(n, k)$ time, however SMIN_C is #BIS-hard.*

Proof In order to prove hardness for general n , it suffices to exhibit such a family of cost tensors when $n = 2$. Since $n = 2$, it is convenient to abuse notation slightly by indexing a cost tensor $C \in (\mathbb{R}^n)^{\otimes k}$ by $\vec{j} \in \{-1, 1\}^k$ rather than by $\vec{j} \in \{1, 2\}^k$. The family we exhibit is $\{C(A, b) : A \in \mathbb{R}_{\geq 0}^{k \times k}, b \in \mathbb{R}^k\}$, where the cost tensors $C(A, b)$ are parameterized by a non-negative square matrix A and a vector b , and have entries of the form

$$C_{\vec{j}}(A, b) := -\langle \vec{j}, A\vec{j} \rangle - \langle b, \vec{j} \rangle, \quad \vec{j} \in \{\pm 1\}^k.$$

Polynomial-time algorithm for MIN and AMIN. We show that given a matrix $A \in \mathbb{R}_{\geq 0}^{k \times k}$, vector $b \in \mathbb{R}^k$, and weights $p \in \mathbb{R}^{2 \times k}$, it is possible to compute $\text{MIN}_C(p)$ on the cost tensor $C(A, b)$ in $\text{poly}(k)$ time. Clearly this also implies a $\text{poly}(k)$ time algorithm for $\text{AMIN}_C(p, \varepsilon)$ for any $\varepsilon > 0$, see Remark 3.5.

To this end, we first re-write the $\text{MIN}_C(p)$ problem on $C(A, b)$ in a more convenient form that enables us to “ignore” the weights p . Recall that $\text{MIN}_C(p)$ is the problem of

$$\text{MIN}_C(p) = \min_{\vec{j} \in \{\pm 1\}^k} -\langle \vec{j}, A\vec{j} \rangle - \langle b, \vec{j} \rangle - \sum_{i=1}^k [p_i]_{j_i}.$$

Note that the linear part of the cost is equal to

$$\langle b, \vec{j} \rangle + \sum_{i=1}^k [p_i]_{j_i} = \langle \ell, \vec{j} \rangle + d, \tag{3.1}$$

where $\ell \in \mathbb{R}^k$ is the vector with entries $\ell_i = b_i + \frac{1}{2}((p_i)_1 - (p_i)_{-1})$, and d is the scalar $d = \frac{1}{2} \sum_{i=1}^k ([p_i]_1 + [p_i]_{-1})$. Thus, since d is clearly computable in $O(k)$ time, the MIN_C problem is equivalent to solving

$$\min_{\vec{j} \in \{\pm 1\}^k} -\langle \vec{j}, A\vec{j} \rangle - \langle \ell, \vec{j} \rangle, \tag{3.2}$$

when given as input a non-negative matrix $A \in \mathbb{R}_{\geq 0}^{k \times k}$ and a vector $\ell \in \mathbb{R}^k$.

To show that this task is solvable in $\text{poly}(k)$ time, note that the objective in (3.2) is a submodular function because it is a quadratic whose Hessian $-A$ has non-positive off-diagonal terms [11, Proposition 6.3]. Therefore (3.2) is a submodular optimization problem, and thus can be solved in $\text{poly}(k)$ time using classical algorithms from combinatorial optimization [44, Chapter 10.2].

SMIN oracle is #BIS-hard. On the other hand, by using the definition of the SMIN oracle, the re-parameterization (3.1), and then the definition of the softmin operator,

$$\begin{aligned} \text{SMIN}_C(p, \eta) &= \text{smin}_{\vec{j} \in \{\pm 1\}^k} -\langle \vec{j}, A\vec{j} \rangle - \langle b, \vec{j} \rangle - \sum_{i=1}^k [p_i]_{j_i} \\ &= \text{smin}_{\vec{j} \in \{\pm 1\}^k} -\langle \vec{j}, A\vec{j} \rangle - \langle \ell, \vec{j} \rangle - d = -\frac{\log Z}{\eta} - d, \end{aligned}$$

where $Z = \sum_{\vec{j} \in \{\pm 1\}^k} Q(\vec{j})$ is the partition function of the ferromagnetic Ising model with inconsistent external fields given by

$$Q(\vec{j}) = \exp\left(\eta \langle \vec{j}, A\vec{j} \rangle + \eta \langle \ell, \vec{j} \rangle\right).$$

Because it is #BIS hard to compute the partition function Z of a ferromagnetic Ising model with inconsistent external fields [42], it is #BIS hard to compute the value $-\eta^{-1} \log Z - d$ of the oracle $\text{SMIN}_C(p, \eta)$. \square

Remark 3.8 (The restrictiveness of SMIN extends to approximate computation) The separation between the oracles shown in Lemma 3.7 further extends to approximate computation of the SMIN oracle under the assumption that #BIS is hard to approximate, since under this assumption it is hard to approximate the partition function of a ferromagnetic Ising model with inconsistent external fields [42].

4 Algorithms to oracles

In this section, we consider three algorithms for MOT. Each is iterative and requires only polynomially many iterations. The key issue for each algorithm is the per-iteration runtime, which is in general exponential (roughly n^k). We isolate the respective bottlenecks of these three algorithms into the three variants of the dual feasibility oracle defined in Sect. 3. See Sect. 1.1 and Table 1 for a high-level overview of this section's results.

4.1 The Ellipsoid algorithm and the MIN oracle

Among the most classical algorithms for implicit LP is the Ellipsoid algorithm [43, 44, 53]. However it has taken a back seat to the SINKHORN algorithm in the vast majority of the MOT literature. The very recent paper [7], which focuses on the specific MOT application of computing low-dimensional Wasserstein barycenters, develops a variant of the classical Ellipsoid algorithm specialized to MOT; henceforth this is called ELLIPSOID , see Sect. 4.1.1 for a description of this algorithm. The objective of this section is to analyze ELLIPSOID in the context of general MOT problems in order to prove the following.

Theorem 4.1 *For any family of cost tensors $C \in (\mathbb{R}^n)^{\otimes k}$, the following are equivalent:*

- (i) *ELLIPSOID takes $\text{poly}(n, k)$ time to solve the MOT_C problem. (Moreover, it outputs a vertex solution represented as a sparse tensor with at most $nk - k + 1$ non-zeros.)*
- (ii) *There exists a $\text{poly}(n, k)$ time algorithm that solves the MOT_C problem.*
- (iii) *There exists a $\text{poly}(n, k)$ time algorithm that solves the MIN_C problem.*

Interpretation of results. In words, the equivalence “(i) \iff (ii)” establishes that ELLIPSOID can solve any MOT problem in polynomial time that any other algorithm can. Thus from a theoretical perspective, this paper's restriction to ELLIPSOID is at no loss of generality for developing polynomial-time algorithms that exactly solve MOT. In words, the equivalence “(ii) \iff (iii)” establishes that the MOT and MIN problems are polynomial-time equivalent. Thus we may investigate when MOT is tractable by instead investigating the more amenable question of when MIN is tractable (see Sect. 1.1.4) at no loss of generality.

As stated in the related work section, Theorem 4.1 is implicit from combining several known results [6, 7, 44]. Our contribution here is to make this result explicit, since this allows us to unify algorithms from the implicit LP literature with the SINKHORN algorithm. We also significantly simplify part of the implication “(iii) \implies (i)”, which is crucial for making an algorithm that relies on the MIN oracle practical—namely, the Column Generation algorithm discussed below.

Organization of Sect. 4.1. In Sect. 4.1.1, we recall this ELLIPSOID algorithm and how it depends on the violation oracle for (MOT-D). In Sect. 4.1.2, we give a significantly simpler proof that the violation and feasibility oracles are polynomial-time equivalent in the case of (MOT-D), and use this to prove Theorem 4.1. In Sect. 4.1.3, we describe a practical implementation that replaces the ELLIPSOID outer loop with Column Generation.

4.1.1 Algorithm

A key component of the proof of Theorem 4.1 is the ELLIPSOID algorithm introduced in [7] for MOT, which we describe below. In order to present this, we first define a variant of the MIN oracle that returns a minimizing tuple rather than the minimizing value.

Definition 4.2 (*Violation oracle for (MOT-D)*) Given weights $p = (p_1, \dots, p_k) \in \mathbb{R}^{n \times k}$, ARGMIN_C returns the minimizing solution \vec{j} and value of $\min_{\vec{j} \in [n]^k} C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i}$.

ARGMIN_C can be viewed as a violation oracle⁵ for the decision set to (MOT-D). This is because, given $p = (p_1, \dots, p_k) \in \mathbb{R}^{n \times k}$, the tuple \vec{j} output by $\text{ARGMIN}_C(p)$ either provides a violated constraint if $C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i} < 0$, or otherwise certifies p is feasible. In [7] it is proved that MOT can be solved with polynomially many calls to the ARGMIN_C oracle.

Theorem 4.3 (ELLIPSOID guarantee; Proposition 12 of [7]) Algorithm 1 finds an optimal vertex solution for $\text{MOT}_C(\mu)$ using $\text{poly}(n, k)$ calls to the ARGMIN_C oracle and $\text{poly}(n, k)$ additional time. The solution is returned as a sparse tensor with at most $nk - k + 1$ non-zero entries.

Sketch of algorithm. Full details and a proof are in [7]. We give a brief overview here for completeness. First, recall from the implicit LP literature that the classical Ellipsoid algorithm can be implemented in polynomial time for an LP with arbitrarily many constraints so long as it has polynomially many variables and the violation oracle for its decision set is solvable in polynomial time [44]. This does not directly apply to the LP (MOT) because that LP has n^k variables. However, it can apply to the dual LP (MOT-D) because that LP only has nk variables.

This suggests a natural two-step algorithm for MOT. First, compute an optimal dual solution by directly applying the Ellipsoid algorithm to (MOT-D). Second, use

⁵ Recall that a violation oracle for a polytope $K = \{x : \langle a_i, x \rangle \leq b_i, \forall i \in [N]\}$ is an algorithm that given a point p , either asserts p is in K , or otherwise outputs the index i of a violated constraint $\langle a_i, p \rangle > b_i$.

Algorithm 1 ELLIPSOID: specialization of the classical Ellipsoid algorithm to MOT

Input: Cost $C \in (\mathbb{R}^n)^{\otimes k}$, marginals $\mu_1, \dots, \mu_k \in \Delta_n$

Output: Vertex solution to $\text{MOT}_C(\mu)$

\ \ Solve dual

1: Solve (MOT-D) using the Ellipsoid algorithm, with ARGMIN_C as the violation oracle. Let S denote the set of tuples returned by all calls to ARGMIN_C .

\ \ Solve primal

2: Solve (4.1) using the Ellipsoid algorithm.

this dual solution to construct a sparse primal solution. Although this dual-to-primal conversion does not extend to arbitrary LP [18, Exercise 4.17], the paper [7] provides a solution by exploiting the standard-form structure of MOT. The procedure is to solve

$$\begin{aligned} \min \quad & \langle C, P \rangle \\ \text{s.t. } & P_{\vec{j}} = 0, \forall \vec{j} \notin S \end{aligned} \tag{4.1}$$

which is the MOT problem restricted to sparsity pattern S , where S is the set of tuples \vec{j} returned by the violation oracle during the execution of step one of Algorithm 1. This second step takes $\text{poly}(n, k)$ time using a standard LP solver, because running the Ellipsoid algorithm in the first step only calls the violation oracle $\text{poly}(n, k)$ times, and thus S has $\text{poly}(n, k)$ size, and therefore the LP (4.1) has $\text{poly}(n, k)$ variables and constraints. In [7] it is proved that this produces a primal vertex solution to the original MOT problem.

4.1.2 Equivalence of bottleneck to MIN

Although Theorem 4.3 shows that ELLIPSOID can solve MOT in $\text{poly}(n, k)$ time using the ARGMIN oracle, this is not sufficient to prove the implication “(iii) \implies (i)” in Theorem 4.1. In order to prove that implication requires showing the polynomial-time equivalence between MIN and ARGMIN.

Lemma 4.4 (Equivalence of MIN and ARGMIN) *Each of the oracles MIN_C and ARGMIN_C can be implemented using $\text{poly}(n, k)$ calls of the other oracle and $\text{poly}(n, k)$ additional time.*

This equivalence follows from classical results about the equivalence of violation and feasibility oracles [92]. However, the known proof of that general result requires an involved and indirect argument based on “back-and-forth” applications of the Ellipsoid algorithm [44, §4.3]. Here we exploit the special structure of MOT to give a direct and elementary proof. This is essential to practical implementations (see Sect. 4.1.3).

Proof It is obvious how the MIN_C oracle can be implemented via a single call of the ARGMIN_C oracle; we now show the converse. Specifically, given $p_1, \dots, p_k \in \mathbb{R}^n$, we show how to compute a solution $\vec{j} = (j_1, \dots, j_k) \in [n]^k$ for $\text{ARGMIN}_C([p_1, \dots, p_k])$ using nk calls to the MIN_C oracle and polynomial additional time. We use the first n

calls to compute the first index j_1 of the solution, the next n calls to compute the next index j_2 , and so on.

Formally, for $s \in [k]$, let us say that $(j_1^*, \dots, j_s^*) \in [n]^s$ is a “partial solution” of size s if there exists a solution $j \in [n]^k$ for $\text{ARGMIN}_C([p_1, \dots, p_k])$ that satisfies $j_i = j_i^*$ for all $i \in [s]$. Then it suffices to show that for every $s \in [k]$, it is possible to compute a partial solution (j_1^*, \dots, j_s^*) of size s from a partial solution $(j_1^*, \dots, j_{s-1}^*)$ of size $s - 1$ using n calls to the MIN_C oracle and polynomial additional time.

The simple but key observation enabling this is the following. Below, for $i \in [k]$ and $j \in [n]$, define $q_{i,j}$ to be the vector in \mathbb{R}^n with value $[p_i]_j$ on entry j , and value $-M$ on all other entries. In words, the following observation states that if the constant M is sufficiently large, then for any indices j'_i , replacing the vectors p_i with the vectors q_{i,j'_i} in a MIN oracle query effectively performs a MIN oracle query on the original input p_1, \dots, p_k except that now the minimization is only over $\vec{j} \in [n]^k$ satisfying $j_i = j'_i$.

Observation 4.5 *Set $M := 2C_{\max} + 2 \sum_{i=1}^k \|p_i\|_{\max} + 1$. Then for any $s \in [k]$ and any $(j'_1, \dots, j'_s) \in [n]^s$,*

$$\text{MIN}_C([q_{1,j'_1}, \dots, q_{s,j'_s}, p_{s+1}, \dots, p_k]) = \min_{\substack{\vec{j} \in [n]^k \\ \text{s.t. } j_1=j'_1, \dots, j_s=j'_s}} C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i}.$$

Proof By definition of the MIN oracle,

$$\text{MIN}_C([q_{1,j'_1}, \dots, q_{s,j'_s}, p_{s+1}, \dots, p_k]) = \min_{\vec{j} \in [n]^k} C_{\vec{j}} - \sum_{i=1}^s [q_{i,j'_i}]_{j_i} - \sum_{i=s+1}^k [p_i]_{j_i}$$

It suffices to prove that every minimizing tuple $\vec{j} \in [n]^k$ for the right hand side satisfies $j_i = j'_i$ for all $i \in [s]$. Suppose not for sake of contradiction. Then there exists a minimizing tuple $\vec{j} \in [n]^k$ for which $j_\ell \neq j'_\ell$ for some $\ell \in [s]$. But then $[q_{\ell,j'_\ell}]_{j_\ell} = -M$, so the objective value of \vec{j} is at least

$$C_{\vec{j}} - \sum_{i=1}^s [q_{i,j'_i}]_{j_i} - \sum_{i=s+1}^k [p_i]_{j_i} \geq M - C_{\max} - \sum_{i=1}^k \|p_i\|_{\max} = C_{\max} + \sum_{i=1}^k \|p_i\|_{\max} + 1.$$

But this is strictly larger (by at least 1) than the value of any tuple with prefix (j'_1, \dots, j'_s) , contradicting the optimality of \vec{j} . \square

Thus, given a partial solution $(j_1^*, \dots, j_{s-1}^*)$ of length $s - 1$, we construct a partial solution (j_1^*, \dots, j_s^*) of length s by setting j_s^* to be a minimizer of

$$\min_{j'_s \in [n]} \text{MIN}_C([q_{1,j_1^*}, \dots, q_{s-1,j_{s-1}^*}, q_{s,j'_s}, p_{s+1}, \dots, p_k]). \tag{4.2}$$

The runtime bound is clear; it remains to show correctness. To this end, note that

$$\begin{aligned}
 \min_{\substack{\bar{j} \in [n]^k \\ \text{s.t. } j_1=j_1^*, \dots, j_s=j_s^*}} C_{\bar{j}} - \sum_{i=1}^k [p_i]_{j_i} &= \text{MIN}_C([q_1, j_1^*, \dots, q_s, j_s^*, p_{s+1}, \dots, p_k]) \\
 &= \min_{\bar{j}'_s \in [n]} \text{MIN}_C([q_1, j_1^*, \dots, q_{s-1}, j_{s-1}^*, q_s, j'_s, p_{s+1}, \dots, p_k]) \\
 &= \min_{\bar{j}'_s \in [n]} \min_{\substack{\bar{j} \in [n]^k \\ \text{s.t. } j_1=j_1^*, \dots, j_{s-1}=j_{s-1}^*, j_s=j'_s}} C_{\bar{j}} - \sum_{i=1}^k [p_i]_{j_i} \\
 &= \min_{\substack{\bar{j} \in [n]^k \\ \text{s.t. } j_1=j_1^*, \dots, j_{s-1}=j_{s-1}^*}} C_{\bar{j}} - \sum_{i=1}^k [p_i]_{j_i} \\
 &= \text{MIN}_C([p_1, \dots, p_k]),
 \end{aligned}$$

where above the first and third steps are by Observation 4.5, the second step is by construction of j_s^* , the fourth step is by simplifying, and the final step is by definition of $(j_1^*, \dots, j_{s-1}^*)$ being a partial solution of size $s - 1$. We conclude that (j_1^*, \dots, j_s^*) is a partial solution of size s , as desired. \square

We can now conclude the proof of the main result of Sect. 4.1.

Proof of Theorem 4.1 The implication “(i) \implies (ii)” is trivial, and the implication “(ii) \implies (iii)” is shown in [6]. It therefore suffices to show the implication “(iii) \implies (i)”. This follows from combining the fact that ELLIPSOID solves MOT_C in polynomial time given an efficient implementation of ARGMIN_C (Theorem 4.3), with the fact that the MIN_C and ARGMIN_C oracles are polynomial-time equivalent (Lemma 4.4). \square

4.1.3 Practical implementation via Column Generation

Although ELLIPSOID enjoys powerful theoretical runtime guarantees, it is slow in practice because the classical Ellipsoid algorithm is. Nevertheless, whenever ELLIPSOID is applicable (i.e., whenever the MIN_C oracle can be efficiently implemented), we can use an alternative practical algorithm, namely the delayed Column Generation method COLGEN, to compute exact, sparse solutions to MOT.

For completeness, we briefly recall the idea behind COLGEN; for further details see the standard textbook [18, §6.1]. COLGEN runs the Simplex method, keeping only basic variables in the tableau. Each time that COLGEN needs to find a Simplex variable on which to pivot, it solves the “pricing problem” of finding a variable with negative reduced cost. This is the key subroutine in COLGEN. In the present context of the MOT LP, this pricing problem is equivalent to a call to the ARGMIN violation oracle (see [18, Definition 3.2] for the definition of reduced costs). By the polynomial-time equivalence of the ARGMIN and MIN oracles shown in Lemma 4.4, this bottleneck subroutine in COLGEN can be computed in polynomial time whenever the MIN oracle can. For easy recall, we summarize this discussion as follows.

Theorem 4.6 (Standard guarantee for COLGEN; Section 6.1 of [18]) *For any $T > 0$, one can implement T iterations of COLGEN in $\text{poly}(n, k, T)$ time and calls to the*

MIN_C oracle. When COLGEN terminates, it returns an optimal vertex solution, which is given as a sparse tensor with at most $nk - k + 1$ non-zero entries.

Note that COLGEN does not have a theoretical guarantee stating that it terminates after a polynomial number of iterations. But it often performs well in practice and terminates after a small number of iterations, leading to much better empirical performance than ELLIPSOID.

4.2 The Multiplicative Weights Update and the AMIN oracle

The second classical algorithm for solving implicitly-structured LPs that we study in the context of MOT is the Multiplicative Weights Update algorithm MWU [91]. The objective of this section is to prove the following guarantees for its specialization to MOT.

Theorem 4.7 *For any family of cost tensors $C \in (\mathbb{R}^n)^{\otimes k}$, the following are equivalent:*

- (i) *For any $\varepsilon > 0$, MWU takes $\text{poly}(n, k, C_{\max}/\varepsilon)$ time to solve the MOT_C problem ε -approximately. (Moreover, it outputs a sparse solution with at most $\text{poly}(n, k, C_{\max}/\varepsilon)$ non-zero entries.)*
- (ii) *There exists a $\text{poly}(n, k, C_{\max}/\varepsilon)$ -time algorithm that solves the MOT_C problem ε -approximately for any $\varepsilon > 0$.*
- (iii) *There exists a $\text{poly}(n, k, C_{\max}/\varepsilon)$ -time algorithm that solves the AMIN_C problem ε -approximately for any $\varepsilon > 0$.*

Interpretation of results. Similarly to the analogous Theorem 4.1 for ELLIPSOID, the equivalence “(i) \iff (ii)” establishes that MWU can approximately solve any MOT problem in polynomial time that any other algorithm can. Thus, from a theoretical perspective, restricting to MWU for approximately solving MOT problems is at no loss of generality. In words, the equivalence “(ii) \iff (iii)” establishes that approximating MOT and approximating MIN are polynomial-time equivalent. Thus we may investigate when MOT is tractable to approximate by instead investigating the more amenable question of when MIN is tractable (see Sect. 1.1.4) at no loss of generality.

Theorem 4.7 is new to this work. In particular, equivalences between problems with polynomially small error do not fall under the purview of classical LP theory, which deals with exponentially small error [44]. Our use of the MWU algorithm exploits a simple reduction of MOT to a mixed packing-covering LP that has appeared in the $k = 2$ matrix case of Optimal Transport in [19, 73], where implicit LP is not necessary for polynomial runtime.

Organization of Sect. 4.2. In Sect. 4.2.1 we present the specialization of Multiplicative Weights Update to MOT, and recall how it runs in polynomial time and calls to a certain bottleneck oracle. In Sect. 4.2.2, we show that this bottleneck oracle is equivalent to the AMIN oracle, and then use this to prove Theorem 4.7.

4.2.1 Algorithm

Here we present the MWU algorithm, which combines the generic Multiplicative Weights Update algorithm of [91] specialized to MOT, along with a final rounding step that ensures feasibility of the solution.

In order to present MWU , it is convenient to assume that the cost C has entries in the range $[1, 2] \subset \mathbb{R}$, which is at no loss of generality by simply translating and rescaling the cost (see Sect. 4.2.2), and can be done implicitly given a bound on C_{\max} . This is why in the rest of this subsection, every runtime dependence on ε is polynomial in $1/\varepsilon$ for costs in the range $[1, 2]$; after transformation back to $[-C_{\max}, C_{\max}]$, this is polynomial dependence in the standard scale-invariant quantity C_{\max}/ε .

Since the cost C is assumed to have non-negative entries, for any $\lambda \in [1, 2]$, the polytope

$$K(\lambda) = \{P \in \mathcal{M}(\mu_1, \dots, \mu_k) : \langle C, P \rangle \leq \lambda\}$$

of couplings with cost at most λ is a mixed packing-covering polytope (i.e., all variables are non-negative and all constraints have non-negative coefficients). Note that $K(\lambda)$ is non-empty if and only if $MOT_C(\mu)$ has value at most λ . Thus, modulo a binary search on λ , this reduces computing the value of $MOT_C(\mu)$ to the task of detecting whether $K(\lambda)$ is empty. Since $K(\lambda)$ is a mixed packing-covering polytope, the Multiplicative Weights Update algorithm of [91] determines whether $K(\lambda)$ is empty, and runs in polynomial time apart from one bottleneck, which we now define.

In order to define the bottleneck, we first define a potential function. For this, we define the softmax analogously to the softmin as

$$\text{smax}(a_1, \dots, a_t) = -\text{smin}(-a_1, \dots, -a_t) = \log \left(\sum_{i=1}^t e^{a_i} \right).$$

Here we use regularization parameter $\eta = 1$ for simplicity, since this suffices for analyzing MWU , and thus we have dropped this index η for shorthand.

Definition 4.8 (*Potential function for MWU*) Fix a cost $C \in (\mathbb{R}^n)^{\otimes k}$, target marginals $\mu \in (\Delta_n)^k$, and target value $\lambda \in \mathbb{R}$. Define the potential function $\Phi := \Phi_{C, \mu, \lambda} : (\mathbb{R}_{\geq 0}^n)^{\otimes k} \rightarrow \mathbb{R}$ by

$$\Phi(P) = \text{smax} \left(\frac{\langle C, P \rangle}{\lambda}, \frac{m_1(P)}{\mu_1}, \dots, \frac{m_k(P)}{\mu_k} \right).$$

The softmax in the above expression is interpreted as a softmax over the $nk + 1$ values in the concatenation of vectors and scalars in its input. (This slight abuse of notation significantly reduces notational overhead.)

Given this potential function, we now define the bottleneck operation for MWU : find a direction $\vec{j} \in [n]^k$ in which P can be increased such that the potential is increased as little as possible.

Definition 4.9 (*Bottleneck oracle for MWU*) Given iterate $P \in (\mathbb{R}_{\geq 0}^n)^{\otimes k}$, target marginals $\mu \in (\Delta_n)^k$, target value $\lambda \in \mathbb{R}$, and accuracy $\varepsilon > 0$, MWU_BOTTLENECK_C($P, \mu, \lambda, \varepsilon$) either:

- Outputs “null”, certifying that $\min_{\vec{j} \in [n]^k} \frac{\partial}{\partial h} \Phi(P + h \cdot \delta_{\vec{j}}) |_{h=0} > 1$, or
- Outputs $\vec{j} \in [n]^k$ such that $\frac{\partial}{\partial h} \Phi(P + h \cdot \delta_{\vec{j}}) |_{h=0} \leq 1 + \varepsilon$.

(If $\min_{\vec{j} \in [n]^k} \frac{\partial}{\partial h} \Phi(P + h \cdot \delta_{\vec{j}}) |_{h=0}$ is within $(1, 1 + \varepsilon)$, then either return behavior is a valid output.)

Pseudocode for the MWU algorithm is given in Algorithm 2. We prove that MWU runs in polynomial time given access to this bottleneck oracle.

Theorem 4.10 *Let the entries of the cost C lie in the range $[1, 2]$. Given $\lambda \in \mathbb{R}$ and accuracy parameter $\varepsilon > 0$, MWU either certifies that $MOT_C(\mu) \leq \lambda$, or returns a poly($n, k, 1/\varepsilon$)-sparse $P \in \mathcal{M}(\mu_1, \dots, \mu_k)$ satisfying $\langle C, P \rangle \leq \lambda + 8\varepsilon$.*

Furthermore, the loop in Step 1 runs in $\tilde{O}(nk/\varepsilon^2)$ iterations, and Step 2 runs in poly($n, k, 1/\varepsilon$) time.

The MWU algorithm can be used to output a $O(\varepsilon)$ -approximate solution for MOT time via an outer loop that performs binary search over λ ; this only incurs $O(\log(1/\varepsilon))$ -multiplicative overhead in runtime.

Algorithm 2 MWU: specialization of Multiplicative Weights Update [91] to MOT

Input: Accuracy $\varepsilon > 0$, marginals $\mu_1, \dots, \mu_k \in \Delta_n$, target value $\lambda > 0$

Output: Either certifies $MOT_C(\mu) > \lambda$ by returning “infeasible”, or returns a solution P with $\langle C, P \rangle \leq \lambda + 8\varepsilon$

```

    \ Assume cost  $C$  satisfies  $C_{\vec{j}} \in [1, 2]$  for all  $\vec{j} \in [n]^k$  (without loss of generality by rescaling)
    \ Step 1: Multiplicative Weights Update
1:  $P \leftarrow 0 \in (\mathbb{R}_{\geq 0}^n)^{\otimes k}$ ,  $\eta \leftarrow 2(\log(nk + 1))/\varepsilon$ 
2: while  $m(P) < \eta$  do                                     ▷ While total mass is small
3:    $\vec{j} \leftarrow \text{MWU\_BOTTLENECK}_C(P, \mu, \lambda, \varepsilon)$    ▷ Bottleneck: find direction with small potential increase
4:   if  $\vec{j} = \text{“null”}$  then return “infeasible”           ▷ Infeasible if no good direction
5:   else  $P \leftarrow P + \delta_{\vec{j}} \cdot \varepsilon \cdot \min(\lambda/C_{\vec{j}}, \min_i[\mu_i]_{j_i})$    ▷ Else, increase in good direction
6:  $P \leftarrow P/(\eta(1 + \varepsilon)^4)$                                ▷ Rescale

    \ Step 2: round to transportation polytope
7: while  $m(P) < 1$  do                                       ▷ Until all constraints are satisfied
8:    $j_i \leftarrow \operatorname{argmax}_{j \in [n]} ([\mu_i]_j - [m_i(P)]_j)$  for each  $i \in [k]$    ▷ For each marginal, find unsatisfied constraint
9:    $\alpha \leftarrow \min_{i \in [k]} ([\mu_i]_{j_i} - [m_i(P)]_{j_i})$    ▷ Maximal mass to add
10:   $P \leftarrow P + \alpha \cdot \delta_{\vec{j}}$                                ▷ Add mass to saturate at least one constraint

```

Proof We analyze Step 1 (Multiplicative Weights Update) and Step 2 (rounding) of MWU separately. □

Lemma 4.11 (Correctness of Step 1) *Step 1 of Algorithm 2 runs in $\tilde{O}(nk/\varepsilon^2)$ iterations. It either returns (i) “infeasible”, certifying that $K(\lambda)$ is empty; or (ii) finds*

a poly($n, k, 1/\varepsilon$)-sparse tensor $P \in (\mathbb{R}_{\geq 0}^n)^{\otimes k}$ that is approximately in $K(\lambda)$, i.e., P satisfies:

$$m(P) \geq 1 - 4\varepsilon, \quad \langle C, P \rangle \leq \lambda, \quad \text{and} \quad m_i(P) \leq \mu_i \text{ for all } i \in [k]$$

Step 1 is the Multiplicative Weights Update algorithm of [91] applied to the polytope $K(\lambda)$, so correctness follows from the analysis of [91]. We briefly recall the main idea behind this algorithm for the convenience of the reader, and provide a full proof in Appendix A.1 for completeness. The main idea behind the algorithm is that on each iteration, $\vec{j} \in [n]^k$ is chosen so that the increase in the potential $\Phi(P)$ is approximately bounded by the increase in the total mass $m(P)$. If this is impossible, then the bottleneck oracle returns null, which means $K(\lambda)$ is empty. So assume otherwise. Then once the total mass has increased to $m(P) = \eta + O(\varepsilon)$, the potential $\Phi(P)$ must be bounded by $\eta(1 + O(\varepsilon))$. By exploiting the inequality between the max and the softmax, this means that $\max(\langle C, P \rangle/\lambda, \max_{i \in [n], j \in [k]} [m_i(P)]_j / [\mu_i]_j) \leq \Phi(P) \leq \eta(1 + O(\varepsilon))$ as well. Thus, rescaling P by $1/(\eta(1 + O(\varepsilon)))$ in Line 6 satisfies $m(P) \geq 1 - O(\varepsilon)$, $\langle C, P \rangle/\lambda \leq 1$, and $m_i(P)/\mu_i \leq 1$. See Appendix A.1 for full details and a proof of the runtime and sparsity claims.

Lemma 4.12 (Correctness of Step 2) *Step 2 of Algorithm 2 runs in poly($n, k, 1/\varepsilon$) time and returns $P \in \mathcal{M}(\mu_1, \dots, \mu_k)$ satisfying $\langle C, P \rangle \leq \lambda + 8\varepsilon$. Furthermore, P only has poly($n, k, 1/\varepsilon$) non-zero entries.*

Proof of Lemma 4.12 By Lemma 4.11, P satisfies $m_i(P) \leq \mu_i$ for all $i \in [k]$. Observe that this is an invariant that holds throughout the execution of Step 2. This, along with the fact that $\sum_{j=1}^n [m_i(P)]_j = m(P)$ is equal for all i , implies that the indices (j_1, \dots, j_k) found in Line 8 satisfy $[\mu_i]_{j_i} - [m_i(P)]_{j_i} > 0$ for each $i \in [k]$. Thus in particular $\alpha > 0$ in Line 9. It follows that Line 10 makes at least one more constraint satisfied (in particular the constraint “ $[\mu_i]_{j_i} = [m_i(P)]_{j_i}$ ” where i is the minimizer in Line 9). Since there are nk constraints total to be satisfied, Step 2 terminates in at most nk iterations. Each iteration increases the number of non-zero entries in P by at most one, thus P is poly($n, k, 1/\varepsilon$) sparse throughout. That P is sparse also implies that each iteration can be performed in poly($n, k, 1/\varepsilon$) time, thus Step 2 takes poly($n, k, 1/\varepsilon$) time overall.

Finally, we establish the quality guarantee on $\langle C, P \rangle$. By Lemma 4.11, this is at most λ before starting Step 2. During Step 2, the total mass added to P is equal to $1 - m(P)$. This is upper bounded by 4ε by Lemma 4.11. Since $C_{\max} \leq 2$, we conclude that the value of $\langle C, P \rangle$ is increased by at most 8ε in Step 2. □

Combining Lemmas 4.11 and 4.12 concludes the proof of Theorem 4.10. □

4.2.2 Equivalence of bottleneck to AMIN

In order to prove Theorem 4.7, we show that the MWU algorithm can be implemented in polynomial time and calls to the AMIN oracle. First, we prove this fact for the ARGAMIN oracle, which differs from the AMIN oracle in that it also returns a tuple $\vec{j} \in [n]^k$ that is an approximate minimizer.

Definition 4.13 (*Approximate violation oracle for (MOT-D)*) Given weights $p = (p_1, \dots, p_k) \in \mathbb{R}^{n \times k}$ and accuracy $\varepsilon > 0$, ARGAMIN_C returns $\vec{j} \in [n]^k$ that minimizes $\min_{\vec{j} \in [n]^k} C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i}$ up to additive error ε , and its value up to additive error ε .

Lemma 4.14 *Let the entries of the cost C lie in the range $[1, 2]$. The MWU algorithm (Algorithm 2), can be implemented by $\text{poly}(n, k, 1/\varepsilon)$ time and calls to the ARGAMIN_C oracle with accuracy parameter $\varepsilon' = \Theta(\varepsilon^2/(nk))$.*

Proof We show that on each iteration of Step 1 of Algorithm 2 we can emulate the call to the MWU_BOTTLENECK oracle with one call to the ARGAMIN oracle. Recall that $\text{MWU_BOTTLENECK}_C(P, \mu, \lambda, \varepsilon)$ seeks to find $\vec{j} \in [n]^k$ such that

$$V_{\vec{j}} := \left. \frac{\partial}{\partial h} \Phi(P + h\delta_{\vec{j}}) \right|_{h=0}$$

is at most $1 + \varepsilon$, or to certify that for all \vec{j} it is greater than 1. By explicit computation,

$$\begin{aligned} V_{\vec{j}} &= \left. \frac{\partial}{\partial h} \log \left(\exp \left(\frac{\langle C, P \rangle + hC_{\vec{j}}}{\lambda} + \sum_{s=1}^k \sum_{t=1}^n \exp \left(\frac{[m_s(P)]_t + \delta_{t,j_s}}{[\mu_s]_t} \right) \right) \right) \right|_{h=0} \\ &= \left(C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i} \right) \frac{\exp(\langle C, P \rangle / \lambda)}{\exp(\langle C, P \rangle / \lambda) + \sum_{s=1}^k \sum_{t=1}^n \exp([m_s(P)]_t / [\mu_s]_t)}, \end{aligned} \tag{4.3}$$

where the weights $p = (p_1, \dots, p_k) \in \mathbb{R}^{k \times n}$ in the last line are defined as

$$[p_i]_j = - \frac{\lambda}{\exp(\langle C, P \rangle / \lambda)} \cdot \frac{\exp([m_i(P)]_j / [\mu_i]_j)}{[\mu_i]_j}, \quad \forall i \in [k], j \in [n].$$

Note that the second term in the product in (4.3) is positive and does not depend on \vec{j} . This suggests that in order to minimize (4.3), it suffices to compute $\vec{j} \leftarrow \text{ARGAMIN}_C(p, \varepsilon')$ for some accuracy parameter $\varepsilon' > 0$. □

The main technical difficulty with formalizing this intuitive approach is that the weights p are not necessarily efficiently computable. Nevertheless, using $\text{poly}(n, k)$ extra time on each iteration, we can compute the marginals $m_1(P), \dots, m_k(P)$. Since the ARGAMIN oracle returns an ε' -additive approximation of the cost, we can also compute a running estimate \tilde{c} of the cost such that, on iteration T ,

$$\tilde{c} - T\varepsilon' \leq \langle C, P \rangle \leq \tilde{c} + T\varepsilon'.$$

Therefore, we define weights $\tilde{p} \in \mathbb{R}^{n \times k}$, which approximate p and which can be computed in $\text{poly}(n, k)$ time on each iteration:

$$[\tilde{p}_i]_j = - \frac{\lambda}{\exp(\tilde{c}/\lambda)} \cdot \frac{\exp([m_i(P)]_j / [\mu_i]_j)}{[\mu_i]_j}, \quad \forall i \in [k], j \in [n].$$

We also define the approximate value for any $\vec{j} \in [n]^k$:

$$\tilde{V}_{\vec{j}} := \left(C_{\vec{j}} - \sum_{i=1}^k [\tilde{p}_i]_{j_i} \right) \frac{\exp(\tilde{c}/\lambda)/\lambda}{\exp(\tilde{c}/\lambda) + \sum_{s=1}^k \sum_{t=1}^n \exp([m_s(P)]_t / [\mu_s]_t)}$$

It holds that $\text{ARGAMIN}_C(\tilde{p}, \varepsilon')$ returns a $\vec{j} \in [n]^k$ that minimizes $C_{\vec{j}} - \sum_{i=1}^k [\tilde{p}_i]_{j_i}$ up to multiplicative error $1/(1 - \varepsilon')$, because the entries of the cost C are lower-bounded by 1, and $[\tilde{p}_i]_j \leq 0$ for all $i \in [n], j \in [k]$. In particular, $\text{ARGAMIN}_C(\tilde{p}, \varepsilon')$ minimizes $\tilde{V}_{\vec{j}}$ up to multiplicative error $1/(1 - \varepsilon')$. We prove the following claim relating $V_{\vec{j}}$ and $\tilde{V}_{\vec{j}}$:

Claim 4.15 *For any $\vec{j} \in [n]^k$, on iteration T , we have $V_{\vec{j}}/\tilde{V}_{\vec{j}} \in [\exp(-2T\varepsilon'/\lambda), \exp(2T\varepsilon'/\lambda)]$.*

By the above claim, therefore $\text{ARGAMIN}_C(\tilde{p}, \varepsilon')$ minimizes $V_{\vec{j}}$ up to multiplicative error $\exp(4T\varepsilon'/\lambda)/(1 - \varepsilon') \leq (1 + \varepsilon/3)$ if we choose $\varepsilon' = \Omega(\lambda\varepsilon/T)$. Thus one can implement $\text{MWU_BOTTLENECK}_C(p, \mu, \lambda, \varepsilon)$ by returning the value of $\text{ARGAMIN}_C(\tilde{p}, \varepsilon')$ if its value is estimated to be at most $1 + \varepsilon/3$, and returning “null” otherwise. The bound on the accuracy $\varepsilon' = \tilde{\Omega}(\varepsilon^2/(nk))$ follows since $\lambda \in [1, 2]$ follows since $\lambda \in [1, 2]$ and $T = \tilde{O}(nk/\varepsilon^2)$ by Theorem 4.10.

Proof of Claim We compare the expressions for $V_{\vec{j}}$ and $\tilde{V}_{\vec{j}}$. Each of these is a product of two terms. Since $C_{\vec{j}} \geq 0$, and $[\tilde{p}_i]_{j_i}, [p_i]_{j_i} \leq 0$ for all i , the ratio of the first terms is

$$\frac{C_{\vec{j}} - \sum_{i=1}^k [\tilde{p}_i]_{j_i}}{C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i}} \in \left[\min_i [\tilde{p}_i]_{j_i} / [p_i]_{j_i}, \max_i [\tilde{p}_i]_{j_i} / [p_i]_{j_i} \right] \subset [\exp(-T\varepsilon'/\lambda), \exp(T\varepsilon'/\lambda)],$$

where we have used that, for all $i \in [k]$,

$$[\tilde{p}_i]_{j_i} / [p_i]_{j_i} = \exp((C, P)/\lambda) / \exp(\tilde{c}/\lambda) \in [\exp(-T\varepsilon'/\lambda), \exp(T\varepsilon'/\lambda)].$$

Similarly the ratio of the second terms in the expression for $V_{\vec{j}}$ and $\tilde{V}_{\vec{j}}$ is also in the range $[\exp(-T\varepsilon'/\lambda), \exp(T\varepsilon'/\lambda)]$. This concludes the proof of the claim. \square

Finally, we show that the ARGAMIN oracle can be reduced to the AMIN oracle, which completes the proof that MWU can be run with AMIN.

Lemma 4.16 (Equivalence of AMIN and ARGAMIN) *Each of the oracles AMIN_C and ARGAMIN_C with accuracy parameter $\varepsilon > 0$ can be implemented using $\text{poly}(n, k)$ calls of the other oracle with accuracy parameter $\Theta(\varepsilon/k)$ and $\text{poly}(n, k)$ additional time.*

It is worth remarking that the equivalence that we show between AMIN and ARGAMIN is *not* known to hold for the feasibility and separation oracles of general LPs, since the known result for general LPs requires exponentially small error in nk [44, §4.3]. However, in the case of MOT the equivalence follows from a direct

and practical reduction, similar to the proof of the equivalence of the exact oracles (Lemma 4.4). The main difference is that some care is needed to bound the propagation of the errors of the approximate oracles. For completeness, we provide the full proof of Lemma 4.16 in Appendix A.

We conclude by proving Theorem 4.7.

Proof of Theorem 4.7 The implication “(i) \implies (ii)” is trivial, and the implication “(ii) \implies (iii)” is shown in [6]. It therefore suffices to show the implication “(iii) \implies (i)”. For costs C with entries in the range $[1, 2]$, this follows from combining the fact that MWU can be implemented to solve MOT_C in $\text{poly}(n, k, 1/\varepsilon)$ time given an efficient implementation of ARGAMIN_C with polynomially-sized accuracy parameter $\varepsilon' = \text{poly}(1/n, 1/k, \varepsilon)$ (Lemma 4.14), along with the fact that the AMIN_C and ARGAMIN_C oracles are polynomially-time equivalent with polynomial-sized accuracy parameter (Lemma 4.16).

The assumption that C has entries within the range $[1, 2]$ can be removed with no loss by translating and rescaling the original cost $C' \leftarrow (C + 3C_{\max})/(2C_{\max})$ and running Algorithm 2 on C' with approximation parameter $\varepsilon' \leftarrow \varepsilon/(2C_{\max})$. Each τ' -approximate query to the $\text{AMIN}_{C'}$ oracle can be simulated by a τ -approximate query to the AMIN_C oracle, where $\tau = 2C_{\max}\tau'$. \square

Remark 4.17 (Practical optimizations) Our numerical implementation of MWU has two modifications that provide practical speedups. One is maintaining a cached list of the tuples $\vec{j} \in [n]^k$ previously returned by calls to MWU_BOTTLENECK . Whenever MWU_BOTTLENECK is called, we first check whether any tuple \vec{j} in the cache satisfies the desiderata $\frac{\partial}{\partial h} \Phi(P + h \cdot \delta_{\vec{j}})|_{h=0} \leq 1 + \varepsilon$, in which case we use this \vec{j} to answer the oracle query. Otherwise, we answer the oracle query using AMIN as explained above. In practice, this cache allows us to avoid many calls to the potentially expensive AMIN bottleneck. Our second optimization is that, at each iteration of MWU , we check whether the current iterate P can be rescaled in order to satisfy the guarantees in Lemma 4.11 required from Step 1. If so, we stop Step 1 early and use this rescaled P .

4.3 The Sinkhorn algorithm and the SMIN oracle

The Sinkhorn algorithm (SINKHORN) is specially tailored to MOT , and does not apply to general exponential-size LP. Currently it is by far the most popular algorithm in the MOT literature (see Sect. 1.3). However, in general each iteration of SINKHORN takes exponential time $n^{\Theta(k)}$, and it is not well-understood when it can be implemented in polynomial-time. The objective of this section is to show that this bottleneck is polynomial-time equivalent to the SMIN oracle, and in doing so put SINKHORN on equal footing with classical implicit LP algorithms in terms of their reliance on variants of the dual feasibility oracle for MOT . Concretely, this lets us establish the following two results.

First, SINKHORN can solve MOT in polynomial time whenever SMIN can be solved in polynomial time.

Theorem 4.18 *For any family of cost tensors $C \in (\mathbb{R}^n)^{\otimes k}$ and accuracy $\varepsilon > 0$, SINKHORN solves MOT_C to ε accuracy in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time and*

$\text{poly}(n, k, C_{\max}/\varepsilon)$ calls to the SMIN_C oracle with regularization $\eta = (2k \log n)/\varepsilon$. (The solution is output through a polynomial-size implicit representation, see Sect. 4.3.1.)

Second, we show that SINKHORN requires strictly more structure than other algorithms do to solve an MOT problem. This is why the results about ELLIPSOID (Theorem 4.1) and MWU (Theorem 4.7) state that those algorithms solve an MOT problem whenever possible, whereas Theorem 4.18 cannot be analogously extended to such an “if and only if” characterization.

Theorem 4.19 *There is a family of cost tensors $C \in (\mathbb{R}^n)^{\otimes k}$ for which ELLIPSOID solves MOT_C exactly in $\text{poly}(n, k)$ time, however it is $\#BIS$ -hard to implement a single iteration of SINKHORN in $\text{poly}(n, k)$ time.*

Organization of Sect. 4.3. In Sect. 4.3.1, we recall this SINKHORN algorithm and how it depends on a certain marginalization oracle. In Sect. 4.3.2, we show that this marginalization oracle is polynomial-time equivalent to the SMIN oracle, and use this to prove Theorems 4.18 and 4.19.

4.3.1 Algorithm

Here we recall SINKHORN and its known guarantees. To do this, we first define the following oracle. While this oracle does not have an interpretation as a dual feasibility oracle, we show below that it is polynomial-time equivalent to SMIN , which is a specific type of approximate dual feasibility oracle (Remark 3.6).

Definition 4.20 (*MARG*) Given scalings $d = (d_1, \dots, d_k) \in \mathbb{R}_{\geq 0}^{n \times k}$, regularization $\eta > 0$, and an index $i \in [k]$, the marginalization oracle $\text{MARG}_C(d, \eta, i)$ returns the vector $m_i((\otimes_{i' \neq i}^k d_{i'}) \odot \exp[-\eta C]) \in \mathbb{R}_{\geq 0}^n$.

It is known that SINKHORN can solve MOT with only polynomially many calls to this oracle [58]. The approximate solution that SINKHORN computes is a fully dense tensor with n^k non-zero entries, but it is output implicitly in $O(nk)$ space through “scaling vectors” and “rounding vectors”, described below.

Theorem 4.21 (*SINKHORN* guarantee, [58]) *Algorithm 3 computes an ε -approximate solution to $\text{MOT}_C(\mu)$ using $\text{poly}(n, k, C_{\max}/\varepsilon)$ calls to the MARG_C oracle with parameter $\eta = (2k \log n)/\varepsilon$, and $\text{poly}(n, k, C_{\max}/\varepsilon)$ additional time. The solution is of the form*

$$P = \left(\otimes_{i=1}^k d_i \right) \odot \exp[-\eta C] + \left(\otimes_{i=1}^k v_i \right), \tag{4.4}$$

and is output implicitly via the scaling vectors $d_1, \dots, d_k \in \mathbb{R}_{\geq 0}^n$ and rounding vectors $v_1, \dots, v_k \in \mathbb{R}_{\geq 0}^n$.

Sketch of algorithm. Full details and a proof are in [58]. We give a brief overview here for completeness. The main idea of SINKHORN is to solve RMOT , the entropically regularized variant of MOT described in Sect. 2.2. On one hand, this provides

Algorithm 3 SINKHORN: multidimensional analog of classical Sinkhorn scaling

Input: Cost $C \in (\mathbb{R}^n)^{\otimes k}$, marginals $\mu_1, \dots, \mu_k \in \Delta_n$
Output: Implicit representation of tensor (4.4) that is an ε -approximate solution to $\text{MOT}_C(\mu)$

\ \ Step 1: scale

- 1: $d_1, \dots, d_k \leftarrow \mathbf{1}$ and $\eta \leftarrow (2k \log n)/\varepsilon$ ▷ Initialize (no scaling)
- 2: **for** poly($n, k, C_{\max}/\varepsilon$) iterations **do**
- 3: Choose $i \in [k]$ ▷ Round-robin, greedily, or randomly
- 4: $\tilde{\mu}_i \leftarrow \text{MARG}_C(d, \eta, i)$ ▷ Bottleneck: compute i -th marginal
- 5: $d_i \leftarrow d_i \odot (\mu_i/\tilde{\mu}_i)$ ▷ Rescale i -th marginal (division is entrywise)

\ \ Step 2: round to transportation polytope

- 6: **for** $i = 1, \dots, k$ **do** ▷ Rescale each marginal to be below marginal constraints
- 7: $\tilde{\mu}_i \leftarrow \text{MARG}_C(d, \eta, i)$ ▷ Bottleneck: compute i -th marginal
- 8: $d_i \leftarrow d_i \odot \min[\mathbf{1}, \mu_i/\tilde{\mu}_i]$ ▷ Rescale i -th marginal (operations are entrywise)
- 9: $v_i \leftarrow \mu_i - \text{MARG}_C(d, \eta, i)$ for each $i \in [k]$ ▷ Add back mass
- 10: $v_1 \leftarrow v_1/\|v_1\|_1^{k-1}$ ▷ Rescale so that (4.4) satisfies marginal constraints
- 11: **Return** d_1, \dots, d_k and v_1, \dots, v_k ▷ Implicit representation of solution (4.4)

an ε -approximate solution to MOT by taking the regularization parameter $\eta = \Theta(\varepsilon^{-1}k \log n)$ sufficiently high (Lemma 2.4). On the other hand, solving RMOT rather than MOT enables exploiting the first-order optimality conditions of RMOT, which imply that the unique solution to RMOT is the unique tensor in $\mathcal{M}(\mu_1, \dots, \mu_k)$ of the form

$$P^* = (\otimes_{i=1}^k d_i^*) \odot K, \quad (4.5)$$

where K denotes the entrywise exponentiated tensor $\exp[-\eta C]$, and $d_1^*, \dots, d_k^* \in \mathbb{R}_{\geq 0}^n$ are non-negative vectors. The SINKHORN algorithm approximately computes this solution in two steps.

The first and main step of Algorithm 3 is the natural multimarginal analog of the Sinkhorn scaling algorithm [78]. It computes an approximate solution $P = (\otimes_{i=1}^k d_i) \odot K$ by finding d_1, \dots, d_k such that P is nearly feasible in the sense that $m_i(P) \approx \mu_i$ for each $i \in [k]$. Briefly, it does this via alternating optimization: initialize d_i to the all-ones vector $\mathbf{1} \in \mathbb{R}^n$, and then iteratively update one d_i so that the i -th marginal $m_i(P)$ of the current scaled iterate $P = (\otimes_{i=1}^k d_i) \odot K$ is μ_i . Although correcting one marginal can detrimentally affect the others, this algorithm nevertheless converges—in fact, in a polynomial number of iterations [58].

The second step of Algorithm 3 is the natural multimarginal analog of the rounding algorithm [5, Algorithm 2]. It rounds the solution $P = (\otimes_{i=1}^k d_i) \odot K$ found in step one to the transportation polytope $\mathcal{M}(\mu_1, \dots, \mu_k)$. Briefly, it performs this by scaling each marginal $m_i(P)$ to be entrywise less than the desired μ_i , and then adding mass back to P so that all marginal constraints are exactly satisfied. The former adjustment is done by adjusting the diagonal scalings d_1, \dots, d_k , and the latter adjustment is done by adding a rank-1 term $\otimes_{i=1}^k v_i$.

Critically, note that Algorithm 3 takes polynomial time except for possibly the calls to the MARG_C oracle. In the absence of structure in the cost tensor C , evaluating this

MARG_C oracle takes exponential time because it requires computing marginals of a tensor with n^k entries.

We conclude this discussion with several remarks about SINKHORN.

Remark 4.22 (Choice of update index in SINKHORN) In line 3 there are several ways to choose update indices, all of which lead to the polynomial iteration complexity we desire. Iteration-complexity bounds are shown for a greedy choice in [40, 58]. Similar bounds can be shown for random and round-robin choices by adapting the techniques of [9]. These latter two choices do not incur the overhead of k MARG computations per iteration required by the greedy choice, which is helpful in practice. Empirically, we observe that round-robin works quite well, and we use this in our experiments.

Remark 4.23 (Alternative implementations of SINKHORN) For simplicity, Algorithm 3 provides pseudocode for the “vanilla” version of SINKHORN as it performs well in practice and it achieves the polynomial iteration complexity we desire. There are several variants in the literature, including accelerated versions and first rounding small entries of the marginals—these variants have iteration-complexity bounds with better polynomial dependence on ϵ and k , albeit sometimes at the expense of larger polynomial factors in n [58, 84].

Remark 4.24 (Output of SINKHORN and efficient downstream tasks) While the output P of SINKHORN is fully dense with n^k non-zero entries, its specific form (4.4) enables performing downstream tasks in polynomial time. This is conditional on a polynomial-time MARG_C oracle, which is at no loss of generality since that is required for running SINKHORN in polynomial time in the first place. The basic idea is that P is a mixture of two simple distributions (modulo normalization). The first is $(\otimes_{i=1}^k d_i) \odot \exp[-\eta C]$, which is marginalizable using MARG_C. The second is $\otimes_{i=1}^k v_i$, which is easily marginalizable since it is a product distribution (as the v_i are non-negative). Together, this enables efficient marginalization of P . By recursively marginalizing on conditional distributions, this enables efficiently sampling from P . This in turn enables efficient estimation of bounded statistics of P (e.g., the cost $\langle C, P \rangle$) by Hoeffding’s inequality.

4.3.2 Equivalence of bottleneck to SMIN

Although Theorem 4.21 shows that SINKHORN solves MOT in polynomial time using the MARG oracle, this is neither sufficient to prove the implication “(ii) \implies (i)” in Theorem 4.18, nor to prove Theorem 4.19. In order to prove these results, we show that SMIN and MARG are polynomial-time equivalent.

Lemma 4.25 (Equivalence of MARG and SMIN) *For any regularization $\eta > 0$, each of the oracles MARG_C and SMIN_C can be implemented using poly(n) calls of the other oracle and poly(n, k) additional time.*

Proof *Reduction from SMIN to MARG.* First, we show how to compute SMIN_C(p, η) using one call to the marginalization oracle and $O(n)$ additional time. Consider the entrywise exponentiated matrix $d = \exp[\eta p] \in \mathbb{R}_{\geq 0}^{n \times k}$, and let $\mu_1 = m_1((\otimes_{i=1}^k d_i) \odot$

$\exp[-\eta C])$ be the answer to $\text{MARG}_C(d, \eta, 1)$. Observe that

$$\begin{aligned} -\eta^{-1} \log \left(\sum_{j_1=1}^n [\mu_1]_{j_1} \right) &= -\eta^{-1} \log \left(\sum_{j_1=1}^n \sum_{j_2, \dots, j_k \in [n]} \prod_{i=1}^k [d_i]_{j_i} e^{-\eta C_{\vec{j}}} \right) \\ &= -\eta^{-1} \log \left(\sum_{\vec{j} \in [n]^k} e^{-\eta(C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i})} \right) \\ &= \text{smin}_{\vec{j} \in [n]^k} \left(C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i} \right), \end{aligned}$$

where above the first step is by definition of μ_1 , the second step is by definition of d and combining the sums, and the third step is by definition of smin . We conclude that $-\eta^{-1} \log \sum_{j_1=1}^n [\mu_1]_{j_1}$ is a valid answer to $\text{SMIN}_C(p, \eta)$. Since this is clearly computable from μ_1 in $O(n)$ time, this establishes the claimed reduction.

Reduction from MARG to SMIN. Next, we show that for any marginalization index $i \in [k]$ and entry $\ell \in [n]$, it is possible to compute the ℓ -th entry of the vector $\text{MARG}_C(d, \eta, i)$ using one call to the SMIN_C oracle and $\text{poly}(n, k)$ additional time. Define $v \in \mathbb{R}^n$ to be the vector with ℓ -th entry equal to $[d_i]_\ell$, and all other entries 0. Define the matrix $p = \eta^{-1} \log[d_1, \dots, d_{i-1}, v, d_{i+1}, \dots, d_k] \in \mathbb{R}^{n \times k}$, where recall that $\log 0 = -\infty$ (see Sect. 2). Let $s \in \mathbb{R}$ denote the answer to $\text{SMIN}_C(p, \eta)$. Observe that

$$\begin{aligned} e^{-\eta s} &= \sum_{\vec{j} \in [n]^k} e^{-\eta(C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i})} = \sum_{\vec{j} \in [n]^k : j_i = \ell} \prod_{i=1}^k [d_i]_{j_i} e^{-\eta C_{\vec{j}}} \\ &= \left[m_i \left((\otimes_{i=1}^k d_i) \odot \exp[-\eta C] \right) \right]_\ell, \end{aligned}$$

where above the first step is by definition of s , the second step is by definition of p and v , and the third step is by definition of the marginalization notation $m_i(\cdot)$. We conclude that $\exp(-\eta s)$ is a valid answer for the ℓ -th entry of the vector $\text{MARG}_C(d, \eta, i)$. This establishes the claimed reduction since we may repeat this procedure n times to compute all n entries of the the vector $\text{MARG}_C(d, \eta, i)$. \square

We can now conclude the proofs of the main results of Sect. 4.3.

Proof of Theorem 4.18 This follows from the fact that SINKHORN approximates MOT_C in polynomial time given a efficient implementation of MARG_C (Theorem 4.21), combined with the fact that the MARG_C and SMIN_C oracles are polynomial-time equivalent (Lemma 4.25).

Proof of Theorem 4.19 Consider the family of cost tensors in Lemma 3.7 for which the MIN_C oracle admits a polynomial-time algorithm, but for which the SMIN_C oracle is $\#$ BIS-hard. Then on one hand, the ELLIPSOID algorithm solves MOT_C in polynomial time by Theorem 4.1. And on the other hand, it is $\#$ BIS-hard to implement a single

iteration of SINKHORN because that requires implementing the $MARG_C$ oracle, which is polynomial-time equivalent to the $SMIN_C$ oracle by Lemma 4.25. \square

5 Application: MOT problems with graphical structure

In this section, we illustrate our algorithmic framework on MOT problems with graphical structure. Although a polynomial-time algorithm is already known for this particular structure [49, 82], that algorithm computes solutions that are approximate and dense; see the related work section for details. By combining our algorithmic framework developed above with classical facts about graphical models, we show that it is possible to compute solutions that are exact and sparse in polynomial time.

The section is organized as follows. In Sect. 5.1, we recall the definition of graphical structure. In Sect. 5.2, we show that the MIN, AMIN, and SMIN oracles can be implemented in polynomial time for cost tensors with graphical structure; from this it immediately follows that all of the MOT algorithms discussed in part 1 of this paper can be implemented in polynomial time. Finally, in Sect. 5.3, we demonstrate our results on the popular application of computing generalized Euler flows, which was the original motivation of MOT. Numerical simulations demonstrate how the exact, sparse solutions produced by our new algorithms provide qualitatively better solutions than previously possible in polynomial time.

5.1 Setup

We begin by recalling preliminaries about undirected graphical models, a.k.a., Markov Random Fields. We recall only the relevant background; for further details we refer the reader to the textbooks [55, 88].

In words, graphical models provide a way of encoding the independence structure of a collection of random variables in terms of a graph. The formal definition is as follows. Below, all graphs are undirected, and the notation 2^V means the power set of V (i.e., the set of all subsets of V).

Definition 5.1 (*Graphical model structure*) Let $\mathcal{S} \subset 2^{[k]}$. The graphical model structure corresponding to \mathcal{S} is the graph $G_{\mathcal{S}} = (V, E)$ with vertices $V = [k]$ and edges $E = \{(i, j) : i, j \in S, \text{ for some } S \in \mathcal{S}\}$.

Definition 5.2 (*Graphical model*) Let $\mathcal{S} \subset 2^{[k]}$. A probability distribution P over $\{X_i\}_{i \in [k]}$ is a graphical model with structure \mathcal{S} if there exist functions $\{\psi_S\}_{S \in \mathcal{S}}$ and normalizing constant Z such that

$$P(\{x_i\}_{i \in [k]}) = \frac{1}{Z} \prod_{S \in \mathcal{S}} \psi_S(\{x_i\}_{i \in S}).$$

A standard measure of complexity for graphical models is the treewidth of the underlying graphical model structure $G_{\mathcal{S}}$ because this captures not just the storage complexity, but also the algorithmic complexity of performing fundamental tasks such




Fig. 1 The path graph (left) has treewidth 1 because the corresponding junction tree (right) has bags of size at most 2




Fig. 2 The graph that has an edges between all vertices of distance at most two when ordered sequentially (left) has treewidth 2 because the corresponding junction tree (right) has bags of size at most 3




Fig. 3 The cycle graph (left) has treewidth 2 because the corresponding junction tree (right) has bags of size at most 3

as computing the mode, log-partition function, and marginal distributions [55, 88]. There are a number of equivalent definitions of treewidth [22]. Each requires defining intermediate combinatorial concepts. We recall here the definition that is based on the concept of a junction tree because this is perhaps the most standard definition in the graphical models community.

Definition 5.3 (*Junction tree, treewidth*) A junction tree $T = (V_T, E_T, \{B_u\}_{u \in V_T})$ for a graph $G = (V, E)$ consists of a tree (V_T, E_T) and a set of bags $\{B_u \subseteq V\}_{u \in V_T}$ satisfying:

- For each variable $i \in V$, the set of nodes $U_i = \{u \in V_T : i \in B_u\}$ induces a subtree of T .
- For each edge $e \in E$, there is some bag B_u containing both endpoints of e .

The width of the junction tree is one less than the size of the largest bag, i.e., is $\max_{u \in V_T} |B_u| - 1$. The treewidth of a graph is the width of its minimum-width junction tree.

See Figs. 1, 2, and 3 for illustrated examples.

We now formally recall the definition of graphical structure for MOT.

Definition 5.4 (*Graphical structure for MOT*) An MOT cost tensor $C \in (\mathbb{R}^n)^{\otimes k}$ has graphical structure with treewidth ω if there is a graphical model structure $\mathcal{S} \subset 2^{[k]}$ and functions $\{f_S\}_{S \in \mathcal{S}}$ such that

$$C_{\vec{j}} = \sum_{S \in \mathcal{S}} f_S(\{j_i\}_{i \in S}), \quad \forall \vec{j} := (j_1, \dots, j_k) \in [n]^k, \quad (5.1)$$

and such that the graph G_S has treewidth ω .

We make three remarks about this structure. First, note that the functions $\{f_S\}_{S \in \mathcal{S}}$ can be arbitrary so long as the corresponding graphical model structure has treewidth at most ω .

Second, if Definition 5.4 did not constrain the treewidth ω , then every tensor C would trivially have graphical structure with maximal treewidth $\omega = k - 1$ (take \mathcal{S} to be the singleton containing $[k]$, G_S to be the complete graph, and $f_{[k]}$ to be C). Just like all previous algorithms, our algorithms have runtimes that depend exponentially (only) on the treewidth of G_S . This is optimal in the sense that unless $P = NP$, there is no algorithm with jointly polynomial runtime in the input size and treewidth [6]. We also point out that in all current applications of graphically structured MOT, the treewidth is either 1 or 2, see Sect. 1.3.

Third, as in all previous work on graphically structured MOT, we make the natural assumptions that the cost C is input implicitly through the functions $\{f_S\}_{S \in \mathcal{S}}$, and that each function f_S can be evaluated in polynomial time, since otherwise graphical structure is useless for designing polynomial-time algorithms. In all applications in the literature, these two basic assumptions are always satisfied. Note also that if the treewidth of the graphical structure is constant, then there is a linear-time algorithm to compute the treewidth and a corresponding minimum-width junction tree [21].

5.2 Polynomial-time algorithms

By our oracle reductions in Sect. 4, in order to design polynomial-time algorithms for MOT with graphical structure, it suffices to design polynomial-time algorithms for the MIN, AMIN, or SMIN oracles. This follows directly from classical algorithmic results in the graphical models literature [55].

Theorem 5.5 (Polynomial-time algorithms for the MIN, AMIN, and SMIN oracles for costs with graphical structure) *Let $C \in (\mathbb{R}^n)^{\otimes k}$ be a cost tensor that has graphical structure with constant treewidth ω (see Definition 5.4). Then the MIN_C , $AMIN_C$, and $SMIN_C$ oracles can be computed in $\text{poly}(n, k)$ time.*

Algorithm 4 Polynomial-time algorithm for MIN for graphically structured costs (Theorem 5.5).

Input: Cost C with graphical structure, matrix $p \in \mathbb{R}^{n \times k}$

Output: Solution to $MIN_C(p)$

- 1: $\vec{j} \leftarrow$ mode of the graphical model P in (5.2) ▷ Using the classical max-product algorithm [55, §13.3]
 - 2: Return $C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i}$ ▷ Value of the $MIN_C(p)$ oracle
-

Algorithm 5 Polynomial-time algorithm for SMIN for graphically structured costs (Theorem 5.5).

Input: Cost C with graphical structure, matrix $p \in \mathbb{R}^{n \times k}$, regularization $\eta > 0$	
Output: Solution to $\text{SMIN}_C(p, \eta)$	
1: $Z \leftarrow$ partition function of the graphical model P in (5.2)	▷ Using the classical sum-product algorithm [55, §10.2]
2: Return $-\eta^{-1} \log Z$	▷ Value of the $\text{SMIN}_C(p, \eta)$ oracle

Proof Consider input p for the oracles. Let P denote the probability distribution on $[n]^k$ given by

$$P(\vec{j}) = \frac{1}{Z} \exp\left(-\eta\left(C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i}\right)\right), \quad \forall \vec{j} \in [n]^k, \tag{5.2}$$

where $Z = \sum_{\vec{j} \in [n]^k} \exp(-\eta(C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i}))$ ensures P is normalized. Observe that the MIN_C oracle amounts⁶ to computing the mode of the distribution P because $\text{MIN}_C(p) = C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i}$, where $\vec{j} \in [n]^k$ is a maximizer of $P_{\vec{j}}$. Also, the SMIN_C oracle amounts to computing the partition function Z because $\text{SMIN}_C(p) = -\eta^{-1} \log Z$. Thus it suffices to compute the mode and partition function of P in polynomial time. (The AMIN_C oracle follows from the MIN_C oracle by Remark 3.5).

To this end, observe that by assumption on C , there is a graphical model structure $\mathcal{S} \in 2^{[k]}$ and functions $\{f_S\}_{S \in \mathcal{S}}$ such that the corresponding graph G_S has treewidth ω and the distribution P factors as

$$P(\vec{j}) = \exp\left(-\eta\left(\sum_{S \in \mathcal{S}} f_S(\{j_i\}_{i \in S}) - \sum_{i=1}^k [p_i]_{j_i}\right)\right).$$

It follows that P is a graphical model with respect to the same graphical model structure \mathcal{S} because the “vertex potentials” $\exp(\eta[p_i]_{j_i})$ do not affect the underlying graphical model structure. Thus P is a graphical model with constant treewidth ω , so we may compute the mode and partition function of P in $\text{poly}(n, k)$ time using, respectively, the classical max-product and sum-product algorithms [55, Chapters 13.3 and 10.2]. For convenience, pseudocode summarizing this discussion is provided in Algorithms 4 and 5. □

An immediate consequence of Theorem 5.5 combined with our oracle reductions is that all candidate MOT algorithms in Sect. 4 can be efficiently implemented for MOT problems with graphical structure. From a theoretical perspective, **ELLIPSOID** gives the best guarantee since it produces an exact, sparse solution.

Corollary 5.6 (Polynomial-time algorithms for MOT problems with graphical structure) *Let $C \in (\mathbb{R}^n)^{\otimes k}$ be a cost tensor that has graphical structure with constant treewidth ω (see Definition 5.4). Then:*

⁶ In fact, for the purpose of computing MIN_C , the distribution $P(\vec{j})$ can be defined using any $\eta > 0$.

- The *ELLIPSOID* algorithm in Sect. 4.1 computes an exact solution to MOT_C in $\text{poly}(n, k)$ time.
- The *MWU* algorithm in Sect. 4.2 computes an ε -approximate solution to MOT_C in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time.
- The *SINKHORN* algorithm in Sect. 4.3 computes an ε -approximate solution to MOT_C in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time.
- The *COLGEN* algorithm in Sect. 4.1.3 can be run for T iterations in $\text{poly}(n, k, T)$ time.

Moreover, *ELLIPSOID*, *MWU*, and *COLGEN* output a polynomially sparse tensor; whereas *SINKHORN* outputs a fully dense tensor through the implicit representation described in Sect. 4.3.1.

Proof Combine the polynomial-time implementations of the oracles in Theorem 5.5 with the polynomial-time algorithm-to-oracle reductions in Theorems 4.1, 4.7, 4.18, and 4.6, respectively. □

5.3 Application vignette: Fluid dynamics

In this section, we numerically demonstrate our new results for graphically structured MOT—namely the ability to compute exact, sparse solutions in polynomial time (Corollary 5.6). We illustrate this on the problem of computing generalized Euler flows—an MOT application which has received significant interest and which was historically the motivation of MOT, see e.g., [14, 16, 23–26]. This MOT problem is already known to be tractable via a popular, specially-tailored modification of *SINKHORN* [14]—which can be interpreted as implementing *SINKHORN* using graphical structure [49, 82]. However, that algorithm is based on *SINKHORN* and thus unavoidably produces solutions that are low-precision (due to $\text{poly}(1/\varepsilon)$ runtime dependence), fully dense (with n^k non-zero entries), and have well-documented numerical precision issues. We offer the first polynomial-time algorithm for computing exact and/or sparse solutions.

We briefly recall the premise of this MOT problem; for further background see [14, 26]. An incompressible fluid (e.g., water) is modeled by n particles which are uniformly distributed in space (due to incompressibility) at all times $t \in \{1, \dots, k + 1\}$. We observe each particle’s location at initial time $t = 1$ and final time $t = k + 1$. The task is to infer the particles’ locations at all intermediate times $t \in \{2, \dots, k\}$, and this is modeled by an MOT problem as follows.

Specifically, the locations of the fluid particles are discretized to points $\{x_j\}_{j \in [n]} \subset \mathbb{R}^d$, and σ is a known permutation on this set that encodes the relation between each initial location x_j at time $t = 1$ and final location $\sigma(x_j)$ at time $t = k + 1$. The total movement of a particle that takes the trajectory $x_{j_1}, x_{j_2}, \dots, x_{j_k}, \sigma(x_{j_1})$ is given by

$$C_{j_1, \dots, j_k} = \|\sigma(x_{j_1}) - x_{j_k}\|^2 + \sum_{t=1}^{k-1} \|x_{j_{t+1}} - x_{j_t}\|^2, \tag{5.3}$$




Fig. 4 Transport maps computed by the fast implementation of SINKHORN [14] (left) and our COLGEN implementation (right) on a standard fluid dynamics benchmark problem in dimension $d = 1$ [26]. The pairwise transport maps between successive timesteps are plotted with opacity proportional to the mass. The SINKHORN algorithm is run at the highest precision (i.e., smallest regularization) before serious numerical precision issues (NaNs). It returns a dense, approximate solution in 2.25 seconds. (All experiments in this paper are run on a standard-issue Apple MacBook Pro 2020 laptop with an M1 Chip.) COLGEN returns an exact, sparse solution in 9.52 seconds. Furthermore, in this particular problem instance, the COLGEN method returns a Monge solution, i.e., the sparsity is n so that the particles never split in the computed trajectories

By the principle of least action, the generalized Euler flow problem of inferring the most likely trajectories of the fluid particles is given by the solution to the MOT problem with this cost C and uniform marginals $\mu_t = \mathbf{1}_n/n \in \Delta_n$ which impose the constraint that the fluid is incompressible.

Corollary 5.7 (Exact, sparse solutions for generalized Euler flows) *The MOT problem with cost (5.3) can be solved in $d \cdot \text{poly}(n, k)$ time. The solution is returned as a sparse tensor with at most $nk - k + 1$ non-zeros.*

Proof This cost tensor C can be expressed in graphical form $C_{\vec{j}} = \sum_{S \in \mathcal{S}} f_S(\{j_i\})$ where \mathcal{S} consists of the sets $\{1, 2\}, \dots, \{k-1, k\}$ of adjacent time points as well as the set $\{1, k\}$. Moreover, each function $f_S : [n]^2 \rightarrow \mathbb{R}$ can be computed in $O(dn^2)$ time since this simply requires computing $\|x_j - x_{j'}\|^2$ for n^2 pairs of points $x_j, x_{j'} \in \mathbb{R}^d$. Once this graphical representation is computed, Corollary 5.6 implies a $\text{poly}(n, k)$ time algorithm for this MOT problem because the graphical model structure \mathcal{S} is a cycle graph and thus has treewidth 2 (cf., Fig. 3). \square

Figure 4 illustrates how the exact, sparse solutions found by our new algorithm provide visually sharper estimates than the popular modification of SINKHORN in [14], which blurs the trajectories. The latter is the state-of-the-art algorithm in the literature and in particular is the only previously known non-heuristic algorithm that has polynomial-time guarantees. Note that this algorithm is identical to implementing SINKHORN by exploiting the graphical structure to perform exact marginalization efficiently [49, 82].

The numerical simulation is on a standard benchmark problem used in the literature (see e.g., [14, Figure 9] and [26, Figure 2]) in which the particle at initial location $x \in [0, 1]$ moves to final location $\sigma(x) = x + \frac{1}{2} \pmod{1}$. This is run with $k = 6$ and marginals $\mu_1 = \dots = \mu_k$ uniformly supported on $n = 51$ positions in $[0, 1]$. See Appendix B for numerics on other standard benchmark instances. Note that this

amounts to solving an MOT LP with $n^k = 51^6 \approx 1.8 \times 10^{10}$ variables, which is infeasible for standard LP solvers. Our algorithm is the first to compute exact solutions for problem instances of this scale.

Two important remarks. First, since this MOT problem is a discretization of the underlying PDE, an exact solution is of course not necessary; however, there is an important—even qualitative—difference between low-precision solutions (computable with $\text{poly}(1/\varepsilon)$ runtime) and high-precision solutions (computable with $\text{polylog}(1/\varepsilon)$ runtime) for the discretized problem. Second, a desirable feature of SINKHORN that should be emphasized is its practical scalability, which might make it advantageous for problems where very fine discretization is required. It is an interesting direction of practical relevance to develop algorithms that can compute high-precision solutions at a similarly large scale in practice (see the discussion in Sect. 8).

6 Application: MOT problems with set-optimization structure

In this section, we consider MOT problems whose cost tensors C take values 0 and 1—or more generally any two values, by a straightforward reduction⁷. Such MOT problems arise naturally in applications where one seeks to minimize or maximize the probability that some event occurs given marginal probabilities on each variable (see Example 6.1). We establish that this general class of MOT problems can be solved in polynomial time under a condition on the sparsity pattern of C that is often simple to check due its connection to classical combinatorial optimization problems.

The section is organized as follows. In Sect. 6.1 we formally describe this setup and discuss why it is incomparable to all other structures discussed in this paper. In Sect. 6.2, we show that for costs with this structure, the MIN, AMIN, and SMIN oracles can be implemented in polynomial time; from this it immediately follows that the ELLIPSOID, MWU, SINKHORN, and COLGEN algorithms discussed in part 1 of this paper can be implemented in polynomial time. In Sect. 6.3, we illustrate our results via a case study on network reliability.

6.1 Setup

Example 6.1 (Motivation for binary-valued MOT costs: minimizing/maximizing probability of an event) Let $S \subset [n]^k$. If $C_{\vec{j}} = \mathbb{1}[\vec{j} \in S]$, then the MOT_C problem amounts to minimizing the probability that event S occurs, given marginals on each variable. On the other hand, if $C_{\vec{j}} = \mathbb{1}[\vec{j} \notin S]$, then the MOT_C problem amounts to maximizing the probability that event S occurs since

$$\text{MOT}_C(\mu_1, \dots, \mu_k) = \min_{P \in \mathcal{M}(\mu_1, \dots, \mu_k)} \mathbb{P}_{\vec{j} \sim P}[\vec{j} \notin S] = 1 - \max_{P \in \mathcal{M}(\mu_1, \dots, \mu_k)} \mathbb{P}_{\vec{j} \sim P}[\vec{j} \in S].$$

⁷ If C takes two values $a < b$, then define the tensor \tilde{C} with $\{0, 1\}$ -entries by $\tilde{C}_{\vec{j}} = (C_{\vec{j}} - a)/(b - a)$. It is straightforward to see that the MOT problems with costs C and \tilde{C} have identical solutions.

Even if the cost is binary-valued, there is no hope to solve MOT in polynomial time without further assumptions—essentially because in the worst case, any algorithm must query all n^k entries if C is a completely arbitrary $\{0, 1\}$ -valued tensor.

We show that MOT is polynomial-time solvable under the general and often simple-to-check condition that the MIN, AMIN, and SMIN oracles introduced in Sect. 3 are polynomial-time solvable when restricted to the set S of indices $\vec{j} \in [n]^k$ for which $C_{\vec{j}} = 0$. For simplicity, our definition of these set oracles removes the cost $C_{\vec{j}}$ as it is constant on S . Of course it is also possible to remove the negative sign in $-p$ by re-parameterizing the inputs as $w = -p$; however, we keep this notation in order to parallel the original oracles.

Definition 6.2 (*MIN set oracle*) Let $S \subset [n]^k$. For weights $p = (p_1, \dots, p_k) \in \mathbb{R}^{n \times k}$, $\text{MIN}_{C,S}(p)$ returns

$$\min_{\vec{j} \in S} - \sum_{i=1}^k [p_i]_{j_i}.$$

Definition 6.3 (*AMIN set oracle*) Let $S \subset [n]^k$. For weights $p = (p_1, \dots, p_k) \in \mathbb{R}^{n \times k}$ and accuracy $\varepsilon > 0$, $\text{AMIN}_{C,S}(p, \varepsilon)$ returns $\text{MIN}_{C,S}(p)$ up to additive error ε .

Definition 6.4 (*SMIN set oracle*) Let $S \subset [n]^k$. For weights $p = (p_1, \dots, p_k) \in \mathbb{R}^{n \times k}$ and regularization parameter $\eta > 0$, $\text{SMIN}_{C,S}(p, \eta)$ returns

$$\text{smin}_{\vec{j} \in S} - \sum_{i=1}^k [p_i]_{j_i}.$$

The key motivation behind these set oracle definitions (aside from the syntactic similarity to the original oracles) is that they encode the problem of (approximately) finding the min-weight object in S . This opens the door to combinatorial applications of MOT because finding the min-weight object in S is well-known to be polynomial-time solvable for many “combinatorial-structured” sets S of interest—e.g., the set S of cuts in a graph, or the set S of independent sets in a matroid. See Sect. 6.3 for fully-detailed applications.

Definition 6.5 (*Set-optimization structure for MOT*) An MOT cost tensor $C \in (\mathbb{R}^n)^{\otimes k}$ has exact, approximate, or soft set-optimization structure of complexity β if

$$C_{\vec{j}} = \mathbb{1}[\vec{j} \notin S]$$

for a set $S \subset [n]^k$ for which there is an algorithm solving $\text{MIN}_{C,S}$, $\text{AMIN}_{C,S}$, or $\text{SMIN}_{C,S}$, respectively, in β time.

We make two remarks about this structure.

Remark 6.6 (Only require set oracle for $C^{-1}(0)$, not for $C^{-1}(1)$) Note that Definition 6.5 only requires the set oracles for the set S of entries where C is 0, and does *not*

need the set oracles for the set $[n]^k \setminus S$ where C is 1. The fact that both set oracles are not needed makes set-optimization structure easier to check than the original oracles in Sect. 3, because those effectively require optimization over both S and $[n]^k \setminus S$.

Remark 6.7 (Set-optimization structure is incomparable to graphical and low-rank plus sparse structure) Costs C that satisfy Definition 6.5 in general do not have non-trivial graphical structure or low-rank plus sparse structure. Specifically, there are costs C that satisfy Definition 6.5, yet require maximal $k - 1$ treewidth to model via graphical structure, and super-constant rank or exponential sparsity to model via low-rank plus sparse structure. (A concrete example is the network reliability application in Sect. 6.3.) Because of the NP-hardness of MOT problems with $(k - 1)$ -treewidth graphical structure or super-constant rank [6], simply modeling such problems with graphical structure or low-rank plus rank structure is therefore useless for the purpose of designing polynomial-time MOT algorithms.

6.2 Polynomial-time algorithms

By our oracle reductions in part 1 of this paper, in order to design polynomial-time algorithms for MOT with set-optimization structure, it suffices to design polynomial-time algorithms for the MIN, AMIN, or SMIN oracles. We show how to do this for all three oracles in a straightforward way by exploiting the set-optimization structure.

Theorem 6.8 (Polynomial-time algorithms for the MIN, AMIN, and SMIN oracles for costs with set-optimization structure) *If $C \in (\mathbb{R}^n)^{\otimes k}$ is a cost tensor with exact, approximate, or soft set-optimization structure of complexity β (see Definition 6.5), then the MIN_C , $AMIN_C$, and $SMIN_C$ oracles, respectively, can be computed in $\beta + \text{poly}(n, k)$ time.*

Algorithm 6 Polynomial-time algorithm for MIN for costs with exact set-optimization structure (Theorem 6.8).

Input: Access to C via $MIN_{C,S}$ oracle, matrix $p \in \mathbb{R}^{n \times k}$

Output: Solution to $MIN_C(p)$

- | | |
|---|----------------------|
| 1: $a \leftarrow MIN_{C,S}(p)$ | ▷ One oracle call |
| 2: $x \leftarrow - \sum_{i=1}^k \max_{j \in [n]} [p_i]_j$ | ▷ Takes $O(nk)$ time |
| 3: Return a if $a \leq x$, or $\min(a, 1 + x)$ otherwise | ▷ Takes $O(1)$ time |
-

Proof *Polynomial-time algorithm for MIN.* We first claim that Algorithm 6 implements the $MIN_C(p)$ oracle. To this end, define

$$a := MIN_{C,S}(p) = \min_{\substack{\vec{j} \in [n]^k \\ \text{s.t. } C_{\vec{j}}=0}} - \sum_{i=1}^k [p_i]_{j_i} \quad \text{and} \quad b := \min_{\substack{\vec{j} \in [n]^k \\ \text{s.t. } C_{\vec{j}}=1}} - \sum_{i=1}^k [p_i]_{j_i}. \quad (6.1)$$

Algorithm 7 Polynomial-time algorithm for SMIN for costs with soft set-optimization structure (Theorem 6.8).

Input: Access to C via $\text{SMIN}_{C,S}$ oracle, matrix $p \in \mathbb{R}^{n \times k}$, regularization η
Output: Solution to $\text{SMIN}_C(p, \eta)$

1: $a \leftarrow \exp(-\eta \cdot \text{SMIN}_{C,S}(p))$ ▷ One oracle call
2: $x \leftarrow \prod_{i=1}^k \sum_{j=1}^n \exp(\eta [p_i]_{j_i})$ ▷ Takes $O(nk)$ time
3: Return $-\eta^{-1} \log(e^{-\eta} x + (1 - e^{-\eta})a)$ ▷ Takes $O(1)$ time

By re-arranging the sum and max, it follows that

$$x := - \sum_{i=1}^k \max_{j \in [n]} [p_i]_j = - \max_{\vec{j} \in [n]^k} \sum_{i=1}^k [p_i]_{j_i} = \min_{\vec{j} \in [n]^k} \sum_{i=1}^k -[p_i]_{j_i} = \min(a, b). \tag{6.2}$$

Therefore

$$\begin{aligned} \text{MIN}_C(p) &= \min_{\vec{j} \in [n]^k} C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i} = \min(a, 1 + b) = \begin{cases} a & \text{if } a \leq b \\ \min(a, 1 + \min(a, b)) & \text{if } a > b \end{cases} \\ &= \begin{cases} a & \text{if } a \leq x \\ \min(a, 1 + x) & \text{if } a > x \end{cases}, \end{aligned} \tag{6.3}$$

where above the first step is by definition of MIN_C ; the second step is by partitioning the minimization over $\vec{j} \in [n]^k$ into \vec{j} such that $C_{\vec{j}} = 0$ or $C_{\vec{j}} = 1$, and then plugging in the definitions of a and b ; the third step is by manipulating $\min(a, 1 + b)$ in both cases; and the last step is because $x = \min(a, b)$ as shown above. We conclude that Algorithm 6 correctly outputs $\text{MIN}_C(p)$. Since the algorithm uses one call to the $\text{MIN}_{C,S}$ oracle and $O(nk)$ additional time, the claim is proven.

Polynomial-time algorithm for AMIN. Next, we claim that the same Algorithm 6, now run with the approximate oracle $\text{AMIN}_{C,S}(p, \varepsilon)$ in the first step instead of the exact oracle $\text{MIN}_{C,S}(p)$, computes a valid solution to $\text{AMIN}_C(p, \varepsilon)$. To prove this, let a, b , and x be as defined in (6.1) and (6.2) for the MIN analysis, and let $\tilde{a} = \text{AMIN}_{C,S}(p, \varepsilon)$. By the same logic as in (6.3), except now reversed, the output

$$\begin{cases} \tilde{a} & \text{if } \tilde{a} \leq x \\ \min(\tilde{a}, 1 + x) & \text{if } \tilde{a} > x \end{cases}$$

is equal to $\min(\tilde{a}, 1 + b)$. Now because \tilde{a} is within additive ε error of a (by definition of the $\text{AMIN}_{C,S}$ oracle), it follows that the above output is within ε additive error of

$$\min(a, 1 + b) = \min_{\vec{j} \in [n]^k} C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i} = \text{MIN}_C(p).$$

Thus the output is a valid answer to $\text{AMIN}_C(p, \varepsilon)$, establishing correctness. The run-time claim is obvious.

Polynomial-time algorithm for SMIN. Finally, we claim that Algorithm 7 implements the $SMIN_C(p, \eta)$ oracle. To this end, define

$$a := e^{-\eta \cdot SMIN_{C,S}(p,\eta)} = \sum_{\substack{\vec{j} \in [n]^k \\ \text{s.t. } C_{\vec{j}}=0}} e^{\eta \sum_{i=1}^k [p_i]_{j_i}} \quad \text{and} \quad b := \sum_{\substack{\vec{j} \in [n]^k \\ \text{s.t. } C_{\vec{j}}=1}} e^{\eta \sum_{i=1}^k [p_i]_{j_i}}.$$

By re-arranging products and sums, it follows that

$$x := \prod_{i=1}^k \sum_{j=1}^n e^{\eta [p_i]_{j_i}} = \sum_{\vec{j} \in [n]^k} \prod_{i=1}^k e^{\eta [p_i]_{j_i}} = a + b.$$

Therefore

$$\begin{aligned} SMIN_C(p, \eta) &= -\frac{1}{\eta} \log \left(\sum_{\vec{j} \in [n]^k} e^{-\eta(C_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i})} \right) = -\frac{1}{\eta} \log (a + e^{-\eta} b) \\ &= -\frac{1}{\eta} \log (e^{-\eta} x + (1 - e^{-\eta}) a), \end{aligned}$$

where above the first step is by definition of $SMIN_C$; the second step is by partitioning the sum over $\vec{j} \in [n]^k$ into \vec{j} such that $C_{\vec{j}} = 0$ or $C_{\vec{j}} = 1$, and then plugging in the definitions of a and b ; and the third step is because $x = a + b$ as shown above. We conclude that Algorithm 7 correctly outputs $SMIN_C(p, \eta)$. Since the algorithm uses one call to the $SMIN_{C,S}$ oracle and $O(nk)$ additional time, the claim is proven. \square

An immediate consequence of Theorem 6.8 combined with our oracle reductions is that all of the candidate MOT algorithms described in Sect. 4 can be efficiently implemented for MOT problems with set-optimization structure. From a theoretical perspective, the ELLIPSOID algorithm gives the best guarantee since it produces an exact, sparse solution in polynomial time.

Corollary 6.9 (Polynomial-time algorithms for MOT problems with set-optimization structure) *Let $C \in (\mathbb{R}^n)^{\otimes k}$ be a cost tensor that has set-optimization structure with $\text{poly}(n, k)$ complexity (see Definition 6.5).*

- *Exact set-optimization structure.* The ELLIPSOID algorithm in Sect. 4.1 computes an exact solution to MOT_C in $\text{poly}(n, k)$ time. Also, the COLGEN algorithm in Sect. 4.1.3 can be run for T iterations in $\text{poly}(n, k, T)$ time.
- *Approximate set-optimization structure.* The MWU algorithm in Sect. 4.2 computes an ε -approximate solution to MOT_C in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time.
- *Soft set-optimization structure.* The SINKHORN algorithm in Sect. 4.3 computes an ε -approximate solution to MOT_C in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time.

Moreover, ELLIPSOID, MWU, and COLGEN output a polynomially sparse tensor, whereas SINKHORN outputs a fully dense tensor through the implicit representation described in Sect. 4.3.1.

Proof Combine the polynomial-time implementations of the oracles in Theorem 7.4 with the polynomial-time algorithm-to-oracle reductions in Theorems 4.1, 4.6, 4.7, and 4.18, respectively. \square

6.3 Application vignette: Network reliability with correlations

In this section, we illustrate this class of MOT structures via an application to network reliability, a central topic in network science, engineering, and operations research, see e.g., the textbooks [12, 13, 41]. The basic network reliability question is: given an undirected graph $G = (V, E)$ where each edge $e \in E$ is reliable with some probability q_e and fails with probability $1 - q_e$, what is the probability that all vertices are reachable from all others? This connectivity is desirable in applications, e.g., if G is a computer cluster, the vertices are the machines, and the edges are communication links, then connectivity corresponds to the reachability of all machines. See the aforementioned textbooks for many other applications.

Of course, the above network reliability question is not yet well-defined since the edge failures are only prescribed up to their marginal distributions. *Which joint distribution* greatly impacts the answer.

The most classical setup posits that edge failures are independent [63]. Denote the network reliability probability for this setting by ρ^{ind} . This quantity ρ^{ind} is #P-complete [72, 85] and thus NP-hard to compute, but there exist fully polynomial randomized approximation schemes (a.k.a. FPRAS) for multiplicatively approximating both the connection probability ρ^{ind} [52] and the failure probability $1 - \rho^{\text{ind}}$ [45].

Here we investigate the setting of *coordinated* edge failures, which dates back to the 1980s [89, 93]. This coordination may optimize for disconnection (e.g., by an adversary), or for connection (e.g., maximize the time a network is connected while performing maintenance on each edge e during $1 - q_e$ fraction of the time). We define these notions below; see also Fig. 5 for an illustration. Below, $\text{Ber}(q_e)$ denotes the Bernoulli distribution with parameter q_e .

Definition 6.10 (*Network reliability with correlations*) For an undirected graph $G = (V, E)$ and edge reliability probabilities $\{q_e\}_{e \in E}$:

- The *worst-case network reliability* is


$$\rho^{\min} := \min_{P \in \mathcal{M}(\{\text{Ber}(q_e)\}_{e \in E})} \mathbb{P}_{H \sim P} [H \text{ is a connected subgraph of } G].$$

- The *best-case network reliability* is

$$\rho^{\max} := \max_{P \in \mathcal{M}(\{\text{Ber}(q_e)\}_{e \in E})} \mathbb{P}_{H \sim P} [H \text{ is a connected subgraph of } G].$$

Clearly $\rho^{\min} \leq \rho^{\text{ind}} \leq \rho^{\max}$. These gaps can be large (e.g., see Fig. 5), which promises large opportunities for applications in which coordination is possible. However, in order to realize such an opportunity requires being able to compute ρ^{\min} and ρ^{\max} , and both of these problems require solving an exponentially large LP with

Worst-case reliability



Best-case reliability




Fig. 5 Optimal decompositions for the the worst-case (top) and best-case (bottom) reliability problems on the same graph G and edge reliability probabilities q_e (left). Coordinating edge failures yields significantly different connection probabilities: $\rho^{\min} = 40\%$, $\rho^{\text{ind}} \approx 60\%$, and $\rho^{\max} = 90\%$

$2^{|E|}$ variables. Below we show how to use set-optimization structure to compute these quantities in $\text{poly}(|E|)$ time, thereby recovering as a special case of our general framework the known polynomial-time algorithms for this particular problem in [89, 93], as well as more practical polynomial-time algorithms that scale to input sizes that are an order-of-magnitude larger.

Corollary 6.11 (Polynomial-time algorithm for network reliability with correlations) *The worst-case and best-case network reliability can both be computed in $\text{poly}(|E|)$ time.*

Proof By the observation in Example 6.1, the optimization problems defining ρ^{\min} and $1 - \rho^{\max}$ are instances of MOT in which $k = |E|$, $n = 2$, $\mu_e = \text{Ber}(q_e)$, and each entry of the cost $C \in \{0, 1\}^{|E|}$ is the indicator of whether that subset of edges is a connected or disconnected subgraph of G , respectively. It therefore suffices to show that both of these MOT cost tensors satisfy exact set-optimization structure (Definition 6.5) since that implies a polynomial-time algorithm for exactly solving MOT (Corollary 6.9).

Set-optimization structure for $1 - \rho^{\max}$. In this case, S is the set of connected subgraphs of G . Thus the $\text{MIN}_{C,S}$ problem is: given weights $p \in \mathbb{R}^{2 \times |E|}$, compute

$$\min_{\text{connected subgraph } H \text{ of } G} - \sum_{e \in H} p_{2,e} - \sum_{e \notin H} p_{1,e}.$$

Note that this objective is equal to $\sum_{e \in H} x_e - \sum_{e \in E} p_{1,e}$ where $x_e := p_{1,e} - p_{2,e}$. Since the latter sum is independent of H , the $\text{MIN}_{C,S}$ problem therefore reduces to the problem of finding a minimum-weight connected subgraph in G ; that is, given edge weights $x \in \mathbb{R}^{|E|}$, compute

$$\min_{\text{connected subgraph } H \text{ of } G} \sum_{e \in H} x_e. \tag{6.4}$$

We first show how to solve this in polynomial time in the case that all edge weights x_e are positive. In this case, the optimal solution H is a minimum-weight spanning tree of G . This can be found by Kruskal's algorithm in $O(|E| \log |E|)$ time [56].

For the general case of arbitrary edge weights, note that the edges e with non-positive weight $x \leq 0$ can be added to any solution without worsening the cost or feasibility. Thus these edges are without loss of generality in every solution H , and so it suffices to solve the same problem (6.4) on the graph G' obtained by contracting these non-positively-weighted edges in G . This reduces (6.4) to the same problem of finding a minimum-weight connected subgraph, except now in the special case that all edge weights are positive. Since we have already shown how to solve this case in polynomial time, the proof is complete.

Set-optimization structure for ρ^{\min} . In this case, S is the set of disconnected subgraphs of G . We may simplify the $\text{MIN}_{C,S}$ problem for ρ^{\min} by re-parameterizing the input $p \in \mathbb{R}^{2 \times |E|}$ to edge weights $x \in \mathbb{R}^{|E|}$ as done above in (6.4) for $1 - \rho^{\max}$. Thus the $\text{MIN}_{C,S}$ problem for ρ^{\min} is: given weights $x \in \mathbb{R}^{|E|}$, compute

$$\min_{\text{disconnected subgraph } H \text{ of } G} \sum_{e \in H} x_e. \quad (6.5)$$

We first show how to solve this in the case that all edge weights x_e are negative. In that case, the optimal solution is of the form $H = E \setminus C$, where C is a maximum-weight cut of the graph G with weights x_e . Equivalently, by negating all edge weights, C is a minimum-weight cut of the graph G with weights $-x_e$. Since a minimum-weight cut of a graph with positive weights can be found in polynomial time [81], the problem (6.5) can be solved in polynomial time when all x_e are negative.

Now in the general case of arbitrary edge weights, note that the edges e with non-negative weight $x \geq 0$ can be removed from any solution without worsening the cost or feasibility. Thus these edges are without loss of generality not in every solution H , and so it suffices to solve the same problem (6.5) on the graph G' obtained by deleting these non-negatively-weighted edges in G . This reduces (6.5) to the same problem of finding a minimum-weight disconnected subgraph, except now in the special case that all edge weights are negative. Since we have already shown how to solve this case in polynomial time, the proof is complete. \square

In Fig. 6, we compare the numerical performance of the algorithms in Corollary 6.11—COLGEN and MWU with polynomial-time implementation of their bottlenecks—with the fastest previous algorithms for both best-case and worst-case network reliability. Previously, the fastest algorithms that apply to this problem are (1) out-of-the-box LP solvers run on MOT, (2) the brute-force implementation of SINKHORN which marginalizes over all $n^k = 2^{|E|}$ entries in each iteration, and (3) this COLGEN algorithm that we recover [89, 93]. It is unknown if there is a practically efficient implementation of the $\text{SMIN}_{C,S}$ oracle (and thus of SINKHORN) for both best-case or worst-case reliability. Since the previous algorithms (1) and (2) have exponential runtime that scales as $n^{\Omega(k)} = 2^{\Omega(|E|)}$, they do not scale past tiny input sizes. In contrast, the algorithms in Corollary 6.11 scale to much larger inputs. Indeed, the COLGEN algorithm that our framework recovers can compute exact solutions roughly

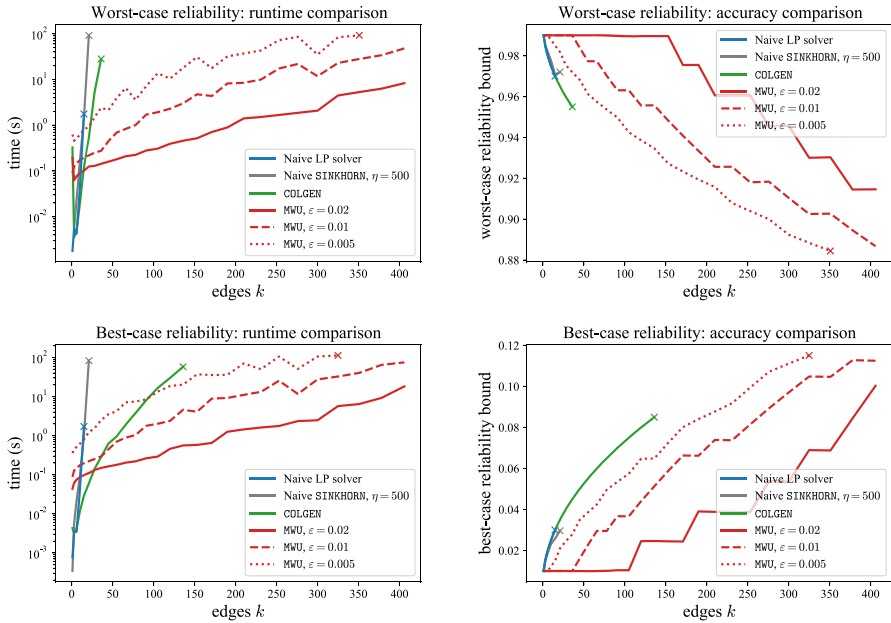


Fig. 6 Top: comparison of the runtime (left) and accuracy (right) of the algorithms described in the main text, for the worst-case reliability of a clique graph on t vertices and $k = \binom{t}{2}$ edges with reliability probabilities $q_e = 0.99$. Bottom: same, but for best-case reliability and reliability probabilities $q_e = 0.01$. For worst-case reliability, the algorithms compute an upper bound, so a smaller value is better; reverse for best-case reliability. The algorithms are cut off at 2 minutes, denoted by an “x”. SINKHORN is run at the highest precision (i.e., highest η) before numerical precision issues. The COLGEN algorithm that our framework recovers computes exact solutions an order-of-magnitude faster than the other algorithms, and the new MWU algorithm computes reasonably approximate solutions for $k = 400$, which amounts to an MOT LP with $n^k = 2^{400} \approx 2.6 \times 10^{120}$ variables

an order-of-magnitude faster than the other algorithms, and the new MWU algorithm computes reasonably approximate solutions beyond $k = 400$, which amounts to an MOT LP with $n^k = 2^{400} \approx 2.6 \times 10^{120}$ variables.

7 Application: MOT problems with low-rank plus sparse structure

In this section, we consider MOT problems whose cost tensors C decompose into low-rank and sparse components. We propose the first polynomial-time algorithms for this general class of MOT problems.

The section is organized as follows. In Sect. 7.1 we formally describe this setup and discuss why it is incomparable to all other structures discussed in this paper. In Sect. 7.2, we show that for costs with this structure, the AMIN and SMIN oracles can be implemented in polynomial time; from this it immediately follows that MWU and SINKHORN can be implemented in polynomial time. Finally, in Sect. 7.3 and Sect. 7.4, we provide two illustrative applications of these algorithms. The former regards portfolio risk management and is a direct application of our result for MOT

with low-rank cost tensors. The latter regards projecting mixture distributions to the transportation polytope and illustrates the versatility of our algorithmic results since this problem is quadratic optimization over the transportation polytope rather than linear (a.k.a. MOT).

7.1 Setup

We begin by recalling the definition of tensor rank. It is the direct analog of the standard concept of matrix rank. See the survey [54] for further background.

Definition 7.1 (*Tensor rank*) A rank- r factorization of a tensor $R \in (\mathbb{R}^n)^{\otimes k}$ is a collection of rk vectors $\{u_{i,\ell}\}_{i \in [k], \ell \in [r]} \subset \mathbb{R}^n$ satisfying

$$R = \sum_{\ell=1}^r \bigotimes_{i=1}^k u_{i,\ell}.$$

The rank of a tensor is the minimal r for which there exists a rank- r factorization.

In this section we consider MOT problems with the following “low-rank plus sparse” structure.

Definition 7.2 (*Low-rank plus sparse structure for MOT*) An MOT cost tensor $C \in (\mathbb{R}^n)^{\otimes k}$ has low-rank plus sparse structure of rank r and sparsity s if it decomposes as

$$C = R + S, \tag{7.1}$$

where R is a rank- r tensor and S is an s -sparse tensor.

Throughout, we make the natural assumption that S is input through its s non-zero entries, and that R is input through a rank- r factorization. We also make the natural assumption that the entries of both R and S are of size $O(C_{\max})$ —this rules out the case of having extremely large entries of R and S , one positive and one negative, which cancel to yield a small entry of $C = R + S$.

Remark 7.3 (Neither low-rank structure nor sparse structure can be modeled by graphical structure or set-optimization structure) In general, both rank-1 costs and polynomially sparse costs do not have non-trivial graphical structure. Specifically, modeling these costs with graphical structure requires the complete graph (a.k.a., maximal treewidth of $k - 1$)—and because MOT problems with graphical structure of treewidth $k - 1$ are NP-hard to solve in the absence of further structure [6], modeling such problems with graphical structure is useless for the purpose of designing polynomial-time MOT algorithms. It is also clear that neither low-rank structure nor sparse structure can be modeled by set-optimization structure because in general, neither R nor S nor $R + S$ has binary-valued entries.

7.2 Polynomial-time algorithms

From a technical perspective, the main result of this section is that there is a polynomial-time algorithm for approximating the minimum entry of a tensor that decomposes into constant-rank and sparse components. Previously, this was not known even for constant-rank tensors. This result may be of independent interest. We remark that this result is optimal in the sense that unless $P = NP$, there does not exist an algorithm with runtime that is *jointly* polynomial in the input size and the rank r [6].

Theorem 7.4 (Polynomial-time algorithm solving AMIN and SMIN for low-rank + sparse costs) *Consider cost tensors $C \in (\mathbb{R}^n)^{\otimes k}$ that have low-rank plus sparse structure of rank r and sparsity s (see Definition 7.2). For any fixed r , Algorithm 8 runs in $\text{poly}(n, k, s, C_{\max}/\varepsilon)$ time and solves the ε -approximate AMIN_C oracle. Furthermore, it also solves the $\text{SMIN}_{\tilde{C}}$ oracle for $\eta = (2k \log n)/\varepsilon$ on some cost tensor $\tilde{C} \in (\mathbb{R}^n)^{\otimes k}$ satisfying $\|C - \tilde{C}\|_{\max} \leq \varepsilon/2$.*

We make three remarks about Theorem 7.4. First, we are unaware of any polynomial-time implementation of SMIN_C for the cost C . Instead, Theorem 7.4 solves the $\text{SMIN}_{\tilde{C}}$ oracle for an $O(\varepsilon)$ -approximate cost tensor \tilde{C} since this is sufficient for implementing SINKHORN on the original cost tensor C (see Corollary 7.5 below). Second, it is an interesting open question if the $\text{poly}(n, k, C_{\max}/\varepsilon)$ runtime for the ε -approximate AMIN_C oracle can be improved to $\text{poly}(n, k, \log(C_{\max}/\varepsilon))$, as this would imply a $\text{poly}(n, k)$ runtime for the MIN_C oracle and thus for this class of MOT problems (see also Footnote 4 in the introduction). Third, we remark about practical efficiency: the runtime of Algorithm 8 is not just polynomially small in s and n , but in fact linear in s and near-linear in n . However, since this improved runtime is not needed for the theoretical results in the sequel, we do not pursue this further.

Combining the efficient oracle implementations in Theorem 7.4 with our algorithm-to-oracles reductions in Sect. 4 implies the first polynomial-time algorithms for MOT problems with costs that have constant-rank plus sparse structure. This is optimal in the sense that unless $P = NP$, there does not exist an algorithm with runtime that is *jointly* polynomial in the input size and the rank r [6].

Corollary 7.5 (Polynomial-time algorithms solving MOT for low-rank + sparse costs) *Consider cost tensors $C \in (\mathbb{R}^n)^{\otimes k}$ that have low-rank plus sparse structure of constant rank r and $\text{poly}(n, k)$ sparsity s (see Definition 7.2). For any $\varepsilon > 0$:*

- The MWU algorithm in Sect. 4.2 computes an ε -approximate solution to MOT_C in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time.
- The SINKHORN algorithm in Sect. 4.3 computes an ε -approximate solution to MOT_C in $\text{poly}(n, k, C_{\max}/\varepsilon)$ time.

Moreover, MWU outputs a polynomially sparse tensor, whereas SINKHORN outputs a fully dense tensor through the implicit representation described in Sect. 4.3.1.

Proof For MWU , simply combine the polynomial-time reduction to the AMIN_C oracle (Theorem 4.7) with the polynomial-time algorithm for the AMIN oracle (Theorem 7.4). For SINKHORN , combining the polynomial-time reduction to the $\text{SMIN}_{\tilde{C}}$ oracle (Theorem 4.18) with the polynomial-time algorithm for the $\text{SMIN}_{\tilde{C}}$ oracle (Theorem 7.4)

yields a poly($n, k, C_{\max}/\varepsilon$) algorithm for $\varepsilon/2$ -approximating the MOT problem with cost tensor \tilde{C} . It therefore suffices to show that the values of the MOT problems with cost tensors C and \tilde{C} differ by at most $\varepsilon/2$, that is,

$$\left| \min_{P \in \mathcal{M}(\mu_1, \dots, \mu_k)} \langle P, C \rangle - \min_{P \in \mathcal{M}(\mu_1, \dots, \mu_k)} \langle P, \tilde{C} \rangle \right| \leq \varepsilon/2.$$

But this holds because both MOT problems have the same feasible set, and for any feasible $P \in \mathcal{M}(\mu_1, \dots, \mu_k)$ it follows from Hölder's inequality that the objectives of the two MOT problems differ by at most

$$\left| \langle P, C \rangle - \langle P, \tilde{C} \rangle \right| \leq \|P\|_1 \|C - \tilde{C}\|_{\max} \leq \varepsilon/2.$$

□

Below, we describe the algorithm in Theorem 7.4. Specifically, in Sect. 7.2.1, we give four helper lemmas which form the technical core of our algorithm; and then in Sect. 7.2.2, we combine these ingredients to design the algorithm and prove its correctness. Throughout, recall that we use the bracket notation $f[A]$ to denote the *entrywise* application of a univariate function f (e.g., exp, log, or a polynomial) to A .

7.2.1 Technical ingredients

At a high level, our approach to designing the algorithm in Theorem 7.4 is to approximately compute the SMIN oracle in polynomial time by synthesizing four facts:

1. By expanding the softmax and performing simple operations, it suffices to compute the total sum of all n^k entries of the entrywise exponentiated tensor $\exp[-\eta R]$ (modulo simple transforms).
2. Although $\exp[-\eta R]$ is in general a full-rank tensor, we can exploit the fact that R is a low-rank tensor in order to approximate $\exp[-\eta R]$ by a low-rank tensor L . (Moreover, we can efficiently compute a low-rank factorization of L in closed form.)
3. There is a simple algorithm for computing the sum of all n^k entries of L in polynomial time because L is low-rank. (And thus we may approximate the sum of all n^k entries of $\exp[-\eta R]$ as desired in step 1.)
4. This approximation is sufficient for computing both the AMIN and SMIN oracle in Theorem 7.4.

Of these four steps, the main technical step is the low-rank approximation in step two. Below, we formalize these four steps individually in Lemmas 7.6, 7.7, 7.8, and 7.9. Further detail on how to synthesize these four steps is then provided afterwards, in the proof of Theorem 7.4.

It is convenient to write the first lemma in terms of an approximate tensor $\tilde{C} = \tilde{R} + S$ rather than the original cost $C = R + S$.

Lemma 7.6 (Softmin for cost with sparse component) *Let $\tilde{C} = \tilde{R} + S$ and $p_1, \dots, p_k \in \mathbb{R}^n$. Then*

$$\operatorname{smin}_{\vec{j} \in [n]^k} \tilde{C}_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i} = -\eta^{-1} \log(a + b),$$

where $d_i := \exp[\eta p_i] \in \mathbb{R}_{>0}^n$,

$$a := \sum_{\substack{\vec{j} \in [n]^k \\ \text{s.t. } S_{\vec{j}} \neq 0}} \prod_{i=1}^k [d_i]_{j_i} \cdot e^{-\eta \tilde{R}_{\vec{j}}} \cdot (e^{-\eta S_{\vec{j}}} - 1) \tag{7.2}$$

and

$$b := \sum_{\vec{j} \in [n]^k} \prod_{i=1}^k [d_i]_{j_i} \cdot e^{-\eta \tilde{R}_{\vec{j}}}. \tag{7.3}$$

Proof By expanding the definition of softmin, and then substituting p_i with d_i and \tilde{C} with $\tilde{R} + S$,

$$\begin{aligned} \operatorname{smin}_{\vec{j} \in [n]^k} \tilde{C}_{\vec{j}} - \sum_{i=1}^k [p_i]_{j_i} &= -\frac{1}{\eta} \log \left(\sum_{\vec{j} \in [n]^k} e^{\eta \sum_{i=1}^k [p_i]_{j_i}} e^{-\eta \tilde{C}_{\vec{j}}} \right) \\ &= -\frac{1}{\eta} \log \left(\sum_{\vec{j} \in [n]^k} \prod_{i=1}^k [d_i]_{j_i} \cdot e^{-\eta \tilde{R}_{\vec{j}}} e^{-\eta S_{\vec{j}}} \right). \end{aligned}$$

By simple manipulations, we conclude that the above quantity is equal to the desired quantity:

$$\begin{aligned} \dots &= -\frac{1}{\eta} \log \left(\sum_{\substack{\vec{j} \in [n]^k \\ \text{s.t. } S_{\vec{j}} \neq 0}} \prod_{i=1}^k [d_i]_{j_i} \cdot e^{-\eta \tilde{R}_{\vec{j}}} e^{-\eta S_{\vec{j}}} + \sum_{\substack{\vec{j} \in [n]^k \\ \text{s.t. } S_{\vec{j}} = 0}} \prod_{i=1}^k [d_i]_{j_i} \cdot e^{-\eta \tilde{R}_{\vec{j}}} \right) \\ &= -\frac{1}{\eta} \log \left(\sum_{\substack{\vec{j} \in [n]^k \\ \text{s.t. } S_{\vec{j}} \neq 0}} \prod_{i=1}^k [d_i]_{j_i} \cdot e^{-\eta \tilde{R}_{\vec{j}}} (e^{-\eta S_{\vec{j}}} - 1) + \sum_{\vec{j} \in [n]^k} \prod_{i=1}^k [d_i]_{j_i} \cdot e^{-\eta \tilde{R}_{\vec{j}}} \right) \\ &= -\frac{1}{\eta} \log(a + b). \end{aligned}$$

Above, the first step is by partitioning the sum over $\vec{j} \in [n]^k$ based on if $S_{\vec{j}} = 0$, the second step is by adding and subtracting $\sum_{\vec{j} \in [n]^k \text{ s.t. } S_{\vec{j}} \neq 0} \prod_{i=1}^k [d_i]_{j_i} \cdot e^{-\eta \tilde{R}_{\vec{j}}}$, and the last step is by definition of a and b . □

Lemma 7.7 (Low-rank approximation of the exponential of a low-rank tensor) *There is an algorithm that given $R \in (\mathbb{R}^n)^{\otimes k}$ in rank- r factored form, $\eta > 0$, and a precision $\tilde{\epsilon} < e^{-\eta R_{\max}}$, takes $n \cdot \text{poly}(k, \tilde{r})$ time to compute a rank- \tilde{r} tensor $L \in (\mathbb{R}^n)^{\otimes k}$ in factored form satisfying $\|L - \exp[-\eta R]\|_{\max} \leq \tilde{\epsilon}$, where*

$$\tilde{r} \leq \binom{r + O(\log \frac{1}{\tilde{\epsilon}})}{r}. \tag{7.4}$$

Proof By classical results from approximation theory (see, e.g., [83]), there exists a polynomial q of degree $m = O(\log 1/\tilde{\epsilon})$ satisfying

$$|\exp(-\eta x) - q(x)| \leq \tilde{\epsilon}, \quad \forall x \in [-R_{\max}, R_{\max}].$$

For instance, the Taylor or Chebyshev expansion of $x \mapsto \exp(-\eta x)$ suffices. Thus the tensor L with entries

$$L_{\vec{j}} = q(R_{\vec{j}})$$

approximates $\exp[-\eta R]$ to error

$$\|L - \exp[-\eta R]\|_{\max} \leq \tilde{\epsilon}.$$

We now show that L has rank $\tilde{r} \leq \binom{r+m}{r}$, and moreover that a rank- \tilde{r} factorization can be computed in $n \cdot \text{poly}(k, \tilde{r})$ time. Denote $q(x) = \sum_{t=0}^m a_t x^t$ and $R = \sum_{\ell=1}^r \otimes_{i=1}^k u_{i,\ell}$. By definition of L , definition of q and R , and then the Multinomial Theorem,

$$L_{\vec{j}} = q(R_{\vec{j}}) = \sum_{t=0}^m a_t \left(\sum_{\ell=1}^r \prod_{i=1}^k [u_{i,\ell}]_{j_i} \right)^t = \sum_{\alpha \in \mathbb{N}_0^r : |\alpha| \leq m} \binom{|\alpha|}{\alpha} a_{|\alpha|} \prod_{\ell=1}^r \prod_{i=1}^k [u_{i,\ell}]_{j_i}^{\alpha_\ell},$$

where the sum is over r -tuples α with non-negative entries summing to at most m . Thus

$$L = \sum_{\alpha \in \mathbb{N}_0^r : |\alpha| \leq m} \bigotimes_{i=1}^k v_{i,\alpha},$$

where $v_{i,\alpha} \in \mathbb{R}^n$ denotes the vector with j -th entry $\binom{|\alpha|}{|\alpha|} a_{|\alpha|} \prod_{\ell=1}^r [u_{i,\ell}]_j^{\alpha_i}$ for $i = 1$, and $\prod_{\ell=1}^r [u_{i,\ell}]_j^{\alpha_i}$ for $i > 1$. This yields the desired low-rank factorization of L because

$$\tilde{r} \leq \#\{\alpha \in \mathbb{N}_0^r : |\alpha| \leq m\} = \binom{r+m}{r}.$$

Finally, since each of the $k\tilde{r}$ vectors $v_{i,\alpha}$ in the factorization of L can be computed efficiently from the closed-form expression above, the desired runtime follows. \square

Lemma 7.8 (Marginalizing a scaled low-rank tensor) *Given vectors $d_1, \dots, d_k \in \mathbb{R}^n$ and a tensor $L \in (\mathbb{R}^n)^{\otimes k}$ through a rank \tilde{r} factorization, we can compute $m((\otimes_{i=1}^k d_i) \odot L)$ in $O(nk\tilde{r})$ time.*

Proof Denote the factorization of L by $L = \sum_{\ell=1}^{\tilde{r}} \otimes_{i=1}^k v_{i,\ell}$. Then

$$\begin{aligned} m((\otimes_{i=1}^k d_i) \odot L) &= \sum_{\tilde{j} \in [n]^k} [(\otimes_{i=1}^k d_i) \odot L]_{\tilde{j}} = \sum_{\tilde{j} \in [n]^k} \sum_{\ell=1}^{\tilde{r}} \prod_{i=1}^k [d_i]_{j_i} [v_{i,\ell}]_{j_i} \\ &= \sum_{\ell=1}^{\tilde{r}} \prod_{i=1}^k \sum_{j=1}^n [d_i]_j [v_{i,\ell}]_j = \sum_{\ell=1}^{\tilde{r}} \prod_{i=1}^k \langle d_i, v_{i,\ell} \rangle, \end{aligned}$$

where the first step is by definition of the $m(\cdot)$ operation that sums over all entries, the second step is by definition of L , and the third step is by swapping products and sums. Thus computing the desired quantity amounts to computing $\tilde{r}k$ inner products of n -dimensional vectors. This can be done in $O(nr\tilde{k})$ time. \square

Lemma 7.9 (Precision of the low-rank approximation) *Let $\varepsilon \leq 1$. Suppose $L \in (\mathbb{R}^n)^{\otimes k}$ satisfies $\|L - \exp[-\eta R]\|_{\max} \leq \frac{\varepsilon}{3} e^{-\eta R_{\max}}$. Then the matrix $\tilde{C} := -\frac{1}{\eta} \log[L] + S$ satisfies*

$$\|\tilde{C} - C\|_{\max} \leq \frac{\varepsilon}{2}. \tag{7.5}$$

Proof Observe that the minimum entry of L is at least

$$e^{-\eta R_{\max}} - \frac{\varepsilon}{3} e^{-\eta R_{\max}} \geq \frac{2}{3} e^{-\eta R_{\max}}. \tag{7.6}$$

Since this is strictly positive, the tensor $\tilde{R} := -\eta^{-1} \log[L]$ is well defined. Furthermore,

$$\|\eta \tilde{R} - \eta R\|_{\max} = \max_{\tilde{j} \in [n]^k} \left| \eta \tilde{R}_{\tilde{j}} - \eta R_{\tilde{j}} \right| \leq \max_{\tilde{j} \in [n]^k} \frac{|L_{\tilde{j}} - e^{-\eta R_{\tilde{j}}}|}{\min(L_{\tilde{j}}, e^{-\eta R_{\tilde{j}}})} \leq \frac{\frac{\varepsilon}{3} e^{-\eta R_{\max}}}{\frac{2}{3} e^{-\eta R_{\max}}} = \frac{\varepsilon}{2},$$

where above the first step is by definition of the max norm; the second step is by the elementary inequality $|\log x - \log y| \leq |x - y| / \min(x, y)$ which holds for positive

scalars x and y [4, Lemma K]; and the third step is by (7.6) and the approximation bound of L . Since $\eta \geq 1$, we therefore conclude that $\|\tilde{R} - R\|_{\max} \leq \varepsilon/2$. By adding and subtracting S , this implies $\|\tilde{C} - C\|_{\max} = \|\tilde{R} - R\|_{\max} \leq \varepsilon/2$. \square

7.2.2 Proof of Theorem 7.4

We are now ready to state the algorithm in Theorem 7.4. Pseudocode is in Algorithm 8. Note that $\tilde{R} = -\eta^{-1} \log[L]$ and $\tilde{C} = \tilde{R} + S$ are never explicitly computed because in both Lines 3 and 4, the algorithm performs the relevant operations only through the low-rank tensor L and the sparse tensor S .

Algorithm 8 Polynomial-time algorithm for AMIN and SMIN for low-rank + sparse costs (Theorem 7.4).

Input: Low-rank tensor R , sparse tensor S , matrix $p \in \mathbb{R}^{n \times k}$, accuracy $\varepsilon > 0$

Output: Solution to both $\text{AMIN}_C(p, \varepsilon)$ on cost tensor $C = R + S$, and also $\text{SMIN}_{\tilde{C}}(p, (2k \log n)/\varepsilon)$ on some approximate cost tensor \tilde{C} satisfying $\|C - \tilde{C}\|_{\max} \leq \varepsilon/2$

- 1: $\eta \leftarrow (2k \log n)/\varepsilon$
 - 2: Compute low-rank approximation L of $\exp[-\eta R]$ via Lemma 7.7, for precision $\tilde{\varepsilon} = \frac{\varepsilon}{3} e^{-\eta R_{\max}}$
 - 3: Compute a in (7.2) directly by enumerating over the polynomially many non-zero entries of S , where $\tilde{R} = -\eta^{-1} \log[L]$
 - 4: Compute b in (7.3) via Lemma 7.8, where $\tilde{R} = -\eta^{-1} \log[L]$
 - 5: Return $-\eta^{-1} \log(a + b)$
-

Proof of Theorem 7.4 Proof of correctness for SMIN. Consider any oracle inputs $p = (p_1, \dots, p_k) \in \mathbb{R}^{n \times k}$. By Lemma 7.9, the tensor $\tilde{C} = \tilde{R} + S = -\eta^{-1} \log L + S$ satisfies $\|\tilde{C} - C\|_{\max} \leq \varepsilon/2$. Therefore it suffices to show that Algorithm 8 correctly computes $\text{SMIN}_{\tilde{C}}(p, \eta)$. This is true because that quantity is equal to $-\eta^{-1} \log(a + b)$ by Lemma 7.6.

Proof of correctness for AMIN. We have just established that Algorithm 8 computes $\text{SMIN}_{\tilde{C}}(p, \eta)$. Because $\eta = (2k \log n)/\varepsilon$ and the fact that SMIN is a special case of AMIN (Remark 3.6), it follows that $\text{SMIN}_{\tilde{C}}(p, \eta)$ is within additive accuracy $\varepsilon/2$ of $\text{MIN}_{\tilde{C}}(p, \eta)$. Therefore, by the triangle inequality, it suffices to show that $\text{MIN}_{\tilde{C}}(p)$ is within $\varepsilon/2$ additive accuracy of $\text{MIN}_C(p)$. That is, it suffices to show that

$$\left| \min_{\tilde{j} \in [n]^k} C_{\tilde{j}} - \sum_{i=1}^k [p_i]_{j_i} - \min_{\tilde{j} \in [n]^k} \tilde{C}_{\tilde{j}} - \sum_{i=1}^k [p_i]_{j_i} \right| \leq \varepsilon/2.$$

But this is true because $\|C - \tilde{C}\|_{\max} \leq \varepsilon/2$ by Lemma 7.9, and thus the quantities $C_{\tilde{j}} - \sum_{i=1}^k [p_i]_{j_i}$ and $\tilde{C}_{\tilde{j}} - \sum_{i=1}^k [p_i]_{j_i}$ are within additive accuracy $\varepsilon/2$ for each $\tilde{j} \in [n]^k$.

Proof of runtime. We prove the claimed runtime bound simultaneously for the AMIN and SMIN computation because we use the same algorithm for both. To this end, we first bound the rank \tilde{r} of the low-rank approximation L computed in Lemma 7.7. Note

that since $\tilde{\varepsilon} = \frac{\varepsilon}{3} e^{-\eta R_{\max}}$ and since it is assumed that $R_{\max} = O(C_{\max})$, we have $\log 1/\tilde{\varepsilon} = O(\frac{C_{\max}}{\varepsilon} k \log n)$. Therefore

$$\tilde{r} \leq \binom{r + O(\log 1/\tilde{\varepsilon})}{r} = O(\log 1/\tilde{\varepsilon})^r = O(\frac{C_{\max}}{\varepsilon} k \log n)^r = \text{poly}(\log n, k, C_{\max}/\varepsilon).$$

Above, the first step is by Lemma 7.7, and the final step is because r is assumed constant.

Therefore Line 2 in Algorithm 8 takes polynomial time by Lemma 7.7, Line 3 takes polynomial time by simply enumerating over the s non-zero entries of S , and Line 4 takes polynomial time by Lemma 7.8. □

7.3 Application vignette: Risk estimation

Here we consider an application to portfolio risk management. For simplicity of exposition, let us first describe the setting of 1 financial instrument (“stock”). Consider investing in one unit of a stock for k years. For $i \in \{0, \dots, k\}$, let X_i denote the price of the stock at year i . Suppose that the return $\rho_i = X_i/X_{i-1}$ of the stock between years $i - 1$ and i is believed to follow some distribution $\rho_i \sim \mu_i$. A fundamental question about the riskiness of this stock is to compute the investor’s expected profit in the worst-case over all joint probability distributions on future returns (ρ_1, \dots, ρ_k) that are consistent with the modeled marginal distributions (μ_1, \dots, μ_k) . This is an MOT problem with cost C given by

$$C(\rho_1, \dots, \rho_k) = \prod_{i \in [k]} \rho_i,$$

where here we view C as a function rather than a tensor for notational simplicity. If each return ρ_i has n possible values (e.g., after quantization), then the cost C is equivalently represented as a rank-1 tensor in $(\mathbb{R}^n)^{\otimes k}$ (by assigning an index to each of the n possible values of each ρ_i). Therefore our result Corollary 7.5 provides a polynomial-time algorithm for solving this MOT problem defining the investor’s worst-case profit.

Rather than formalize this proof for 1 stock, we directly generalize to the general case of investing in r stocks, $r \geq 1$. This is essentially identical to the simple case of $r = 1$ stock, modulo additional notation.

Corollary 7.10 (Polynomial-time algorithm for expected profit given marginals on the returns) *Suppose an investor holds 1 unit of r stocks for k years. For each stock $\ell \in [r]$ and each year $i \in [k]$, let $\rho_{i,\ell}$ denote the relative price of stock ℓ between years i and $i - 1$. Suppose $\rho_{i,\ell}$ has distribution $\mu_{i,\ell}$, and that each $\mu_{i,\ell}$ has at most n atoms. Let $R_{\max} = \max_{\{\rho_{i,\ell}\}} \sum_{\ell=1}^r \prod_{i=1}^k \rho_{i,\ell}$ denote the maximal possible return. For any constant number of stocks r , there is a $\text{poly}(n, k, R_{\max}/\varepsilon)$ time algorithm for ε -approximating the expected profit in the worst-case over all futures that are consistent with the returns’ marginal distributions.*

Proof This is the optimization problem

$$\min_{P \in \mathcal{M}(\{\mu_{i,\ell}\}_{i \in [k], \ell \in [r]})} \mathbb{E}_{\{\rho_{i,\ell}\}_{i \in [k], \ell \in [r]} \sim P} \left[\sum_{\ell=1}^r \prod_{i=1}^k \rho_{i,\ell} \right]$$

over all joint distributions P on the returns $\{\rho_{i,\ell}\}_{i \in [k], \ell \in [k]}$ that are consistent with the marginal distributions $\{\mu_{i,\ell}\}_{i \in [k], \ell \in [k]}$. This is an MOT problem with $k' = rk$ marginals, each over n atoms, with cost function

$$C(\{\rho_{i,\ell}\}_{i \in [k], \ell \in [r]}) = \sum_{\ell' \in [r]} \prod_{(i,\ell) \in [k] \times [r] \cong [k']} (\rho_{i,\ell} \cdot \mathbb{1}[\ell = \ell'] + \mathbb{1}[\ell \neq \ell'])). \quad (7.7)$$

By viewing this cost function C as a cost tensor in the natural way (i.e., assigning an index to each of the n possible values of $\rho_{i,\ell}$), this representation (7.7) shows that the corresponding cost tensor $C \in (\mathbb{R}^n)^{\otimes k'}$ has rank r . Moreover, observe that the maximum entry of the cost is R_{\max} . Therefore we may appeal to our polynomial-time MOT algorithms in Corollary 7.5 for costs with constant rank. \square

The algorithm is readily generalized, e.g., if the investor has different units of a stock, or if a stock is held for a different number of years. The former is modeled simply by adding an extra year in which the return of stock ℓ is equal to the number of units, with probability 1. The latter is modeled simply by setting the return of stock ℓ to be 1 for all years after it is held, with probability 1.

In Fig. 7, we provide a numerical illustration comparing our new polynomial-time algorithms for this risk estimation task with the previous fastest algorithms. Previously, the fastest algorithms that apply to this problem are out-of-the-box LP solvers run on MOT, and the brute-force implementation of SINKHORN which marginalizes over all n^k entries in each iteration. Since both of these previous algorithms have exponential runtime that scales as $n^{\Omega(k)}$, they do not scale beyond tiny input sizes of $n = 10$ and $k = 8$ even with two minutes of computation time. In contrast, our new polynomial-time algorithms compute high-quality solutions for problems that are orders-of-magnitude larger. For example, our polynomial-time implementation of SINKHORN takes less than a second to solve a MOT LP with $n^k = 10^{30}$ variables.

Details for this numerical experiment: we consider $r = 1$ stock over k timesteps, where each marginal distribution μ_i is uniform on $[1, 1 + 1/k]$, discretized with $n = 10$. We implement the AMIN and SMIN oracle efficiently by using our above algorithm to exploit the rank-one structure of the cost tensor. In particular, the polynomial approximation we use here to approximate $\exp[-\eta C]$ is the degree-5 Taylor approximation (cf., Lemma 7.7). This lets us run SINKHORN and MWU in polynomial time, as described above. In the numerical experiment, we also implement an approximate version of COLGEN using our polynomial-time implementation of the approximate violation oracle AMIN. Since the algorithms compute an upper bound, lower value is better in the right plot of Fig. 7. We observe that MWU yields the loosest approximation for this application, whereas our implementations of SINKHORN and COLGEN produce high-quality approximations, as is evident by comparing to the exact LP solver in the regime that the latter is tractable to run.

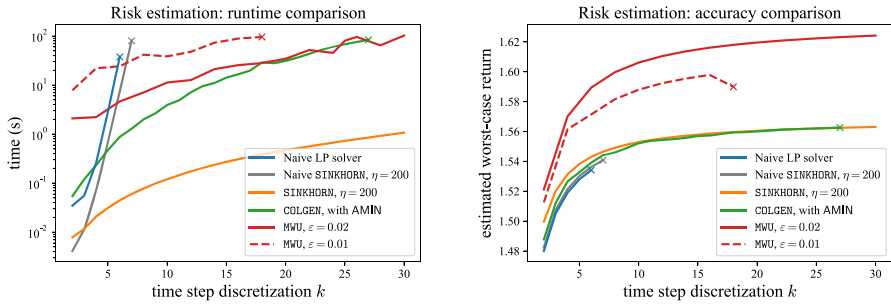


Fig. 7 Comparison of the runtime (left) and accuracy (right) of the fastest existing algorithms (naive LP solver and naive SINKHORN which both have exponential runtimes that scale as $n^{\Omega(k)}$) with our algorithms (SINKHORN, MWU, and COLGEN and MWU with polynomial-time implementations of their bottlenecks) for the risk estimation problem described in the main text. The algorithms are cut off at 2 minutes, denoted by an “x”. Our new polynomial-time implementation of SINKHORN returns high-quality solutions for problems that are orders-of-magnitude larger than previously possible: e.g., it takes less than a second to solve the problem for $k = 30$, which amounts to an MOT LP with 10^{30} variables

7.4 Application vignette: Projection to the transportation polytope

Here we consider the fundamental problem of projecting a joint probability distribution Q onto the transportation polytope $\mathcal{M}(\mu_1, \dots, \mu_k)$, i.e.,

$$\operatorname{argmin}_{P \in \mathcal{M}(\mu_1, \dots, \mu_k)} \sum_j (P_j - Q_j)^2. \tag{7.8}$$

We provide the first polynomial-time algorithm for solving this problem in the case where Q is a distribution that decomposes into a low-rank component plus a sparse component. The low-rank component enables modeling mixtures of product distributions (e.g., mixtures of isotropic Gaussians), which arise frequently in statistics and machine learning; see, e.g., [39]. In such applications, the number of product distributions in the mixture corresponds to the tensor rank. The sparse component further enables modeling arbitrary corruptions to the distribution in polynomially many entries.

We emphasize that this projection problem (7.8) is *not* an MOT problem since the objective is quadratic rather than linear. This illustrates the versatility of our algorithmic results. Our algorithm is based on a reduction from quadratic optimization to linear optimization over $\mathcal{M}(\mu_1, \dots, \mu_k)$ that is tailored to this problem. Crucial to this reduction is the fact that the MOT algorithms in Sect. 4 can compute *sparse* solutions. In particular, this reduction does not work with SINKHORN because SINKHORN cannot compute sparse solutions.

Corollary 7.11 (Efficient projection to the transportation polytope) *Let $Q = R + S \in (\mathbb{R}_{\geq 0}^n)^{\otimes k}$, where R has constant rank and S is polynomially sparse. Suppose that R_{\max} and S_{\max} are $O(1)$. Given R in factored form, S through its non-zero entries, measures $\mu_1, \dots, \mu_k \in \Delta_n$, and accuracy $\varepsilon > 0$, we can compute in $\operatorname{poly}(n, k, 1/\varepsilon)$ time a feasible $P \in \mathcal{M}(\mu_1, \dots, \mu_k)$ that has ε -suboptimal cost for the projection*

problem (7.8). This solution P is a sparse tensor output through its $\text{poly}(n, k, 1/\varepsilon)$ non-zero entries.

Proof We apply the Frank-Wolfe algorithm (a.k.a., Conditional Gradient Descent) to solve (7.8), specifically using approximate LP solutions for the descent direction as in [51, Algorithm 2]. By the known convergence guarantee of this algorithm [51, Theorem 1.1], if each LP is solved to $\varepsilon' = O(\varepsilon)$ accuracy, then $T = O(1/\varepsilon)$ Frank-Wolfe iterations suffice to obtain an ε -suboptimal solution to (7.8).

The crux, therefore, is to show that each Frank-Wolfe iteration can be computed efficiently, and that the final solution is sparse. Initialize $P^{(0)}$ to be an arbitrary vertex of $\mathcal{M}(\mu_1, \dots, \mu_k)$. Then $P^{(0)}$ is feasible and is polynomially sparse (see Sect. 2.1). Let $P^{(t)} \in (\mathbb{R}_{\geq 0}^n)^{\otimes k}$ denote the t -th Frank-Wolfe iterate. Performing the next iteration requires two computations:

1. Approximately solve the following LP to ε' accuracy:

$$D^{(t)} \leftarrow \min_{P \in \mathcal{M}(\mu_1, \dots, \mu_k)} \langle P, P^{(t)} - Q \rangle. \quad (7.9)$$

2. Update $P^{(t+1)} \leftarrow (1 - \gamma_t)P^{(t)} + \gamma_t D^{(t)}$, where $\gamma_t = 2/(t + 2)$ is the current stepsize.

For the first iteration $t = 0$, note that the LP (7.9) is an MOT problem with cost $C^{(0)} = P^{(0)} - Q = P^{(0)} - R - S$ which decomposes into a polynomially sparse tensor $P^{(0)} - S$ plus a constant-rank tensor $-R$. Therefore the algorithm in Corollary 7.5 can solve the LP (7.9) to $\varepsilon' = O(\varepsilon)$ additive accuracy in $\text{poly}(n, k, 1/\varepsilon)$ time, and it outputs a solution $D^{(0)}$ that is $\text{poly}(n, k, 1/\varepsilon)$ sparse. It follows that $P^{(1)}$ can be computed in $\text{poly}(n, k, 1/\varepsilon)$ time and moreover is $\text{poly}(n, k, 1/\varepsilon)$ sparse since it is a convex combination of the similarly sparse tensors $P^{(0)}$ and $D^{(0)}$. By repeating this argument identically for $T = O(1/\varepsilon)$ iterations, it follows that each iteration takes $\text{poly}(n, k, 1/\varepsilon)$ time, and that each iterate $P^{(t)}$ is $\text{poly}(n, k, 1/\varepsilon)$ sparse. \square

8 Discussion

In this paper, we investigated what structure enables MOT—an LP with n^k variables—to be solved in $\text{poly}(n, k)$ time. We developed a unified algorithmic framework for MOT by characterizing what “structure” is required to solve MOT in polynomial time by different algorithms in terms of simple variants of the dual feasibility oracle. On one hand, this enabled us to show that ELLIPSOID and MWU solve MOT in polynomial time whenever any algorithm can, whereas SINKHORN requires strictly more structure. And on the other hand, this made the design of polynomial-time algorithms for MOT much simpler, as we illustrated on three general classes of MOT cost structures.

Our results suggest several natural directions for future research. One exciting direction is to identify further tractable classes of MOT cost structures beyond the three studied in this paper, since this may enable new applications of MOT. Our results help guide this search because they make it significantly easier to identify if an MOT problem is polynomial-time solvable (see Sect. 1.1.4).

Another important direction is practicality. While the focus of this paper is to characterize when MOT problems can be solved in polynomial time, in practice there is of course a difference between small and large polynomial runtimes. It is therefore a question of practical significance to improve our “proof of concept” polynomial-time algorithms by designing algorithms with smaller polynomial runtimes. Our theoretical results help guide this search for practical algorithms because they make it significantly easier to identify if an MOT problem is polynomial-time solvable in the first place.

In order to develop more practical algorithms, recall that, roughly speaking, our approach for designing MOT algorithms consisted of three parts:

- An “outer loop” algorithm such as ELLIPSOID, MWU, or SINKHORN that solves MOT in polynomial time conditionally on a polynomial-time implementation of a certain bottleneck oracle.
- An “intermediate” algorithm that reduces this bottleneck oracle to polynomial calls of a variant of the dual feasibility oracle.
- An “inner loop” algorithm that solves the relevant variant of the dual feasibility oracle for the structured MOT problem at hand.

Obtaining a smaller polynomial runtime for any of these three parts immediately implies smaller polynomial runtimes for the overall MOT algorithm. Another approach is to design altogether different algorithms that avoid the polynomial blow-up of the runtime that arises from composing these three parts. Understanding how to solve an MOT problem more “directly” in this way is an interesting question.

Acknowledgements We are grateful to Jonathan Niles-Weed, Pablo Parrilo, and Philippe Rigollet for insightful conversations; to Frederic Koehler for suggesting a simpler proof of Lemma 3.7; to Ben Edelman and Siddhartha Jayanti who were involved in the brainstorming stages and provided helpful references; and to Karthik Natarajan for references to the random combinatorial optimization literature.

Funding Open Access funding provided by the MIT Libraries

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Deferred proof details

A.1 Proof of Lemma 4.11

Our proof is based on two helper claims.

Claim A.1 (Lemma 3 of [91]) *If $MWU_BOTTLENECK(P, \mu, \lambda, \varepsilon)$ returns “null”, then $MOT_C(\mu) > \lambda$.*

Proof We prove the contrapositive. Let $P^* \in \mathcal{M}(\mu_1, \dots, \mu_k)$ such that $\langle C, P^* \rangle \leq \lambda$. Then

$$\begin{aligned} & \sum_{\vec{j} \in [n]^k} P_j^* \frac{\partial}{\partial h} \Phi(P + h\delta_{\vec{j}}) \Big|_{h=0} \\ &= \frac{\sum_{\vec{j} \in [n]^k} P_j^* \left(\frac{C_{\vec{j}}}{\lambda} \exp(\langle C, P \rangle / \lambda) + \sum_{i=1}^k \frac{1}{[\mu_i]_{j_i}} \exp([m_i(P)]_{j_i} / [\mu_i]_{j_i}) \right)}{\exp(\langle C, P \rangle / \lambda) + \sum_{s=1}^k \sum_{t=1}^n \exp([m_s(P)]_t / [\mu_s]_t)} \\ &= \frac{\frac{\langle C, P^* \rangle}{\lambda} \exp(\langle C, P \rangle / \lambda) + \sum_{s=1}^k \sum_{t=1}^n \frac{[m_s(P^*)]_t}{[\mu_s]_t} \exp([m_s(P)]_t / [\mu_s]_t)}{\exp(\langle C, P \rangle / \lambda) + \sum_{s=1}^k \sum_{t=1}^n \exp([m_s(P)]_t / [\mu_s]_t)} \\ &\leq \frac{\exp(\langle C, P \rangle / \lambda) + \sum_{s=1}^k \sum_{t=1}^n \exp([m_s(P)]_t / [\mu_s]_t)}{\exp(\langle C, P \rangle / \lambda) + \sum_{s=1}^k \sum_{t=1}^n \exp([m_s(P)]_t / [\mu_s]_t)} \\ &= 1. \end{aligned}$$

Since P^* is non-negative and $\sum_{\vec{j} \in [n]^k} P_j^* = m(P^*) = 1$, there must exist a $\vec{j} \in [n]^k$ satisfying $\frac{\partial}{\partial h} \Phi(P + h\delta_{\vec{j}}) \Big|_{h=0} \leq 1$. □

This first claim ensures that if the algorithm returns “infeasible”, then indeed $\text{MOT}_C(\mu) > \lambda$. This proves the first part of the lemma. We now prove a second claim, useful for bounding the running time and the quality of the returned solution when the algorithm does not return “infeasible”.

Claim A.2 (Lemmas 1 and 4 of [91]) *In Lines 1 to 5, we maintain the invariant that $\Phi(P) - (1 + \varepsilon)^2 m(P) \leq \log(nk + 1)$.*

Proof Note that initially $P = 0$ and $\Phi(0) - (1 + \varepsilon)^2 m(0) = \Phi(0) = \log(nk + 1)$. So it suffices to prove that $\Phi(P) - (1 + \varepsilon)^2 m(P)$ does not increase on each iteration. Indeed, if \vec{j} is returned by `MWU_BOTTLENECK`, then $\frac{\partial}{\partial h} \Phi(P + h\delta_{\vec{j}}) \Big|_{h=0} \leq (1 + \varepsilon)$. Furthermore, let $\varepsilon' = \varepsilon \cdot \min(\lambda / C_{\vec{j}}, \min_i [\mu_i]_{j_i})$. By the smoothness of the softmax, it holds that

$$\Phi(P + \varepsilon' \delta_{\vec{j}}) \leq \Phi(P) + \varepsilon' (1 + \varepsilon) \frac{\partial}{\partial h} \Phi(P + h\delta_{\vec{j}}) \Big|_{h=0} \leq \varepsilon' (1 + \varepsilon)^2.$$

So, in Line 5, $(1 + \varepsilon)^2 m(P)$ increases by $(1 + \varepsilon)^2 \varepsilon'$, and $\Phi(P)$ increases by at most $(1 + \varepsilon)^2 \varepsilon'$. This proves the claim. □

We use this second claim to bound the running time of Step 1. Each iteration of the loop from Line 1 to Line 5 of Algorithm 2, increases the value of

$$\Psi(P) := \langle C, P \rangle / \lambda + \sum_i \|m_i(P) / \mu_i\|_1$$

by at least ε . So after T iterations we must have

$$\Phi(P) \geq \Psi(P) / (nk + 1) \geq T\varepsilon / (nk + 1),$$

where we have used Jensen’s inequality to relate Φ and Ψ . By the claim, this means that on any iteration $T \geq (\eta + \log(nk + 1))(nk + 1)(1 + \varepsilon)^2/\varepsilon = \tilde{O}(nk/\varepsilon^2)$, we must have

$$m(P) \geq (\Phi(P) - \log(nk + 1))/(1 + \varepsilon)^2 \geq \eta,$$

so the loop must terminate after at most $\tilde{O}(nk/\varepsilon^2)$ iterations. Since each iteration can increase the number of non-zero entries by at most 1, this also proves the sparsity bound on P .

We now prove the bounds on the marginals and cost, again using the claim. When Line 6 is reached, we must have $m(P) \in [\eta, \eta + \varepsilon]$, because each iteration increases $m(P)$ by at most ε . Therefore,

$$\begin{aligned} \max((C, P)/\lambda, m_1(P)/\mu_1, \dots, m_k(P)/\mu_k) &\leq \Phi(P) && \text{by Lemma 2.3 for softmax} \\ &\leq \log(nk + 1) + (1 + \varepsilon)^2 m(P) && \text{by the claim} \\ &\leq \log(nk + 1) + (1 + \varepsilon)^2 (\eta + \varepsilon) \\ &\leq (1 + \varepsilon)^4 \eta && \text{by } \eta \geq 2 \log(nk + 1) \geq 1 \end{aligned}$$

Therefore, the rescaling in Line 6 yields P satisfying the guarantees of Lemma 4.11.

A.2 Proof of Lemma 4.16

It is obvious how the AMIN_C oracle can be implemented via a single call of the ARGAMIN_C oracle; we now show the converse. Specifically, given $p_1, \dots, p_k \in \mathbb{R}^n$, we show how to compute a solution $\vec{j} = (j_1, \dots, j_k) \in [n]^k$ for $\text{ARGAMIN}_C([p_1, \dots, p_k], \varepsilon)$ using nk calls to the AMIN_C oracle with accuracy $\varepsilon/(2k)$. As in the proof of Lemma 4.4, we use the first n calls to compute the first index j_1 of the solution, the next n calls to compute the next index j_2 , and so on.

Formally, for $s \in [k]$, let us say that $(j_1^*, \dots, j_s^*) \in [n]^s$ is a δ -approximate “partial solution” of size s if there exists a solution $j \in [n]^k$ for $\text{ARGAMIN}_C([p_1, \dots, p_k], \delta)$ that satisfies $j_i = j_i^*$ for all $i \in [s]$. Then it suffices to show that for every $s \in [k]$, it is possible to compute an $(s\varepsilon/k)$ -approximate partial solution (j_1^*, \dots, j_s^*) of size s from an $((s - 1)\varepsilon/k)$ -approximate partial solution $(j_1^*, \dots, j_{s-1}^*)$ of size $s - 1$ using n calls to AMIN_C and polynomial additional time.

Do this by setting j_s^* to be a minimizer of

$$\min_{j'_s \in [n]} \text{AMIN}_C \left(\left[q_{1, j_1^*}, \dots, q_{s-1, j_{s-1}^*}, q_{s, j'_s}, p_{s+1}, \dots, p_k \right], \frac{\varepsilon}{2k} \right), \tag{A.1}$$

where the q vectors are defined as in the proof of Lemma 4.4. The runtime claim is obvious; it suffices to prove correctness. To this end, observe that

$$\begin{aligned}
 & \min_{\substack{\bar{j} \in [n]^k \\ \text{s.t. } j_1=j_1^*, \dots, j_s=j_s^*}} C_{\bar{j}} - \sum_{i=1}^k [p_i]_{j_i} \\
 &= \text{MIN}_C \left(\left[q_{1, j_1^*}, \dots, q_{s, j_s^*}, p_{s+1}, \dots, p_k \right] \right) \\
 &\leq \frac{\varepsilon}{2k} + \text{AMIN}_C \left(\left[q_{1, j_1^*}, \dots, q_{s, j_s^*}, p_{s+1}, \dots, p_k \right], \frac{\varepsilon}{2k} \right) \\
 &= \frac{\varepsilon}{2k} + \min_{j'_s \in [n]} \text{AMIN}_C \left(\left[q_{1, j_1^*}, \dots, q_{s-1, j_{s-1}^*}, q_{s, j'_s}, p_{s+1}, \dots, p_k \right], \frac{\varepsilon}{2k} \right) \\
 &\leq \frac{\varepsilon}{k} + \min_{j'_s \in [n]} \text{MIN}_C \left(\left[q_{1, j_1^*}, \dots, q_{s-1, j_{s-1}^*}, q_{s, j'_s}, p_{s+1}, \dots, p_k \right] \right) \\
 &= \frac{\varepsilon}{k} + \min_{j'_s \in [n]} \min_{\substack{\bar{j} \in [n]^k \\ \text{s.t. } j_1=j_1^*, \dots, j_{s-1}=j_{s-1}^*, j_s=j'_s}} C_{\bar{j}} - \sum_{i=1}^k [p_i]_{j_i} \\
 &= \frac{\varepsilon}{k} + \min_{\substack{\bar{j} \in [n]^k \\ \text{s.t. } j_1=j_1^*, \dots, j_{s-1}=j_{s-1}^*}} C_{\bar{j}} - \sum_{i=1}^k [p_i]_{j_i} \\
 &= \frac{s\varepsilon}{k} + \text{MIN}_C \left(\left[p_1, \dots, p_k \right] \right).
 \end{aligned}$$

Above, the first and fifth steps are by Observation 4.5, the second and fourth steps are by definition of the AMIN oracle, the third step is by construction of j_s^* , the penultimate step is by simplifying, and the final step is by definition of $(j_1^*, \dots, j_{s-1}^*)$ being an $((s-1)\varepsilon/k)$ -approximate partial solution of size $s-1$. We conclude that (j_1^*, \dots, j_s^*) is an $(s\varepsilon/k)$ -approximate partial solution of size s , as desired.

B Additional numerical experiments

Here, we provide additional numerics for the generalized Euler flow application in Sect. 5.3 in order to demonstrate that similar behavior is observed on other standard benchmark inputs in the literature [26]. See Figs. 8 and 9. These instances are identical to Fig. 4, except with different input permutations σ between the initial and final positions of the particles. Note that our algorithm COLGEN computes an exact, sparse solution with at most $nk - k + 1$ non-zero entries (Theorem 4.6). In contrast, the SINKHORN algorithm of [14] computes approximate, fully dense solutions with n^k non-zero entries, which leads to blurry visualizations.




Fig. 8 Same as Fig. 4, but now with the permutation σ that sends the particle at initial location $x \in [0, 1]$ to final location $\sigma(x) = \min(2x, 2-2x)$. COLGEN runs in 7.88 seconds, while SINKHORN with regularization $\eta = 2000$ runs in 6.97 seconds. In the COLGEN solution, roughly half the particles have trajectories that never split




Fig. 9 Same as Fig. 4, but now with the permutation σ that sends the particle at initial location $x \in [0, 1]$ to final location $\sigma(x) = 1 - x$. COLGEN runs in 10.53 seconds, while SINKHORN with regularization $\eta = 1500$ runs in 2.10 seconds

References

1. Agrawal, S., Ding, Y., Saberi, A., Ye, Y.: Price of correlations in stochastic optimization. *Oper. Res.* **60**(1), 150–162 (2012)
2. Agueh, M., Carlier, G.: Barycenters in the Wasserstein space. *SIAM J. on Math. Anal.* **43**(2), 904–924 (2011)
3. Akagi, Y., Tanaka, Y., Iwata, T., Kurashima, T., Toda, H.: Probabilistic optimal transport based on collective graphical models. (2020). Preprint at [arXiv:2006.08866](https://arxiv.org/abs/2006.08866)
4. Altschuler, J., Bach, F., Rudi, A., Niles-Weed, J.: Massively scalable Sinkhorn distances via the Nyström method. *Adv. in Neural Inf. Process. Syst.* **32**, 4429–4439 (2019)
5. Altschuler, J., Weed, J., Rigollet, P.: Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. *Advances in Neural Information Processing Systems* (2017)
6. Altschuler, J.M., Boix-Adserà, E.: Hardness results for Multimarginal Optimal Transport problems. *Discrete Optim.* **42**, 100669 (2021)
7. Altschuler, J.M., Boix-Adserà, E.: Wasserstein barycenters can be computed in polynomial time in fixed dimension. *J. of Machine Learning Res.* **22**, 44–1 (2021)
8. Altschuler, J.M., Boix-Adserà, E.: Wasserstein barycenters are NP-hard to compute. *SIAM Journal on Mathematics of Data Science* (2022)
9. Altschuler, J.M., Parrilo, P.A.: Near-linear convergence of the Random Osborne algorithm for Matrix Balancing. *Mathematical Programming*, 1–35 (2022)
10. Anderes, E., Borgwardt, S., Miller, J.: Discrete Wasserstein barycenters: Optimal transport for discrete data. *Math. Methods of Oper. Res.* **84**(2), 389–409 (2016)
11. Bach, F.: Learning with submodular functions: A convex optimization perspective. *Found. and Trends in Mach. Learning* **6**(2–3), 145–373 (2013)

12. Ball, M.O.: Computational complexity of network reliability analysis: An overview. *IEEE Trans. on Reliab.* **35**(3), 230–239 (1986)
13. Ball, M.O., Colbourn, C.J., Provan, J.S.: Network reliability. *Handbooks in Oper. Res. and Management Sci.* **7**, 673–762 (1995)
14. Benamou, J.-D., Carlier, G., Cuturi, M., Nenna, L., Peyré, G.: Iterative Bregman projections for regularized transportation problems. *SIAM J. on Scientific Comput.* **37**(2), A1111–A1138 (2015)
15. Benamou, J.-D., Carlier, G., Di Marino, S., Nenna, L.: An entropy minimization approach to second-order variational mean-field games. *Math. Models and Methods in Appl. Sci.* **29**(08), 1553–1583 (2019)
16. Benamou, J.-D., Carlier, G., Nenna, L.: A numerical method to solve multi-marginal optimal transport problems with Coulomb cost. In: *Splitting Methods in Communication, Imaging, Science, and Engineering*, pp. 577–601. Springer (2016)
17. Benamou, J.-D., Carlier, G., Nenna, L.: Generalized incompressible flows, multi-marginal transport and Sinkhorn algorithm. *Numerische Mathematik* **142**(1), 33–54 (2019)
18. Bertsimas, D., Tsitsiklis, J.N.: Introduction to linear optimization, vol. 6. Athena Scientific Belmont, MA (1997)
19. Blanchet, J., Jambulapati, A., Kent, C., Sidford, A.: Towards optimal running times for optimal transport. (2018). Preprint at [arXiv:1810.07717](https://arxiv.org/abs/1810.07717)
20. Blondel, M., Seguy, V., Rolet, A.: Smooth and sparse optimal transport. In: *International conference on Artificial Intelligence and Statistics*, pp. 880–889. PMLR (2018)
21. Bodlaender, H.L.: A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. on comput.* **25**(6), 1305–1317 (1996)
22. Bodlaender, H.L.: Treewidth: Structure and algorithms. In: *International Colloquium on Structural Information and Communication Complexity*, pp. 11–25. Springer (2007)
23. Brenier, Y.: The least action principle and the related concept of generalized flows for incompressible perfect fluids. *J. of the Am. Math. Soc.* **2**(2), 225–255 (1989)
24. Brenier, Y.: The dual least action problem for an ideal, incompressible fluid. *Archive for Rational Mech. and Anal.* **122**(4), 323–351 (1993)
25. Brenier, Y.: Minimal geodesics on groups of volume-preserving maps and generalized solutions of the Euler equations. *Commun. on Pure and Appl. Math.* **52**(4), 411–452 (1999)
26. Brenier, Y.: Generalized solutions and hydrostatic approximation of the Euler equations. *Physica D: Nonlinear Phenomena* **237**(14–17), 1982–1988 (2008)
27. Buttazzo, G., De Pascale, L., Gori-Giorgi, P.: Optimal-transport formulation of electronic density-functional theory. *Phys. Review A* **85**(6), 062502 (2012)
28. Carlier, G., Ekeland, I.: Matching for teams. *Econ. Theory* **42**(2), 397–418 (2010)
29. Carlier, G., Oberman, A., Oudet, E.: Numerical methods for matching for teams and Wasserstein barycenters. *ESAIM: Math. Modelling and Numerical Anal.* **49**(6), 1621–1642 (2015)
30. Chen, L., Ma, W., Natarajan, K., Simchi-Levi, D., Yan, Z.: Distributionally robust linear and discrete optimization with marginals. *Operations Research* (2022)
31. Chiappori, P.-A., McCann, R.J., Nesheim, L.P.: Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness. *Econ. Theory* **42**(2), 317–354 (2010)
32. Cotar, C., Friesecke, G., Klüppelberg, C.: Density functional theory and optimal transportation with Coulomb cost. *Commun. on Pure and Appl. Math.* **66**(4), 548–599 (2013)
33. Courty, N., Flamary, R., Tuia, D., Rakotomamonjy, A.: Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **1**, (2016)
34. Cover, T.M., Thomas, J.A.: *Elements of information theory*. John Wiley & Sons, Hoboken, New Jersey, US (2012)
35. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in Neural Information Processing Systems* (2013)
36. Cuturi, M., Doucet, A.: Fast computation of Wasserstein barycenters. *International Conference on Machine Learning* (2014)
37. Dyer, M., Goldberg, L.A., Greenhill, C., Jerrum, M.: On the relative complexity of approximate counting problems. In: *International Workshop on Approximation Algorithms for Combinatorial Optimization*, pp. 108–119. Springer (2000)
38. Elvander, F., Haasler, I., Jakobsson, A., Karlsson, J.: Multi-marginal optimal transport using partial information with applications in robust localization and sensor fusion. *Signal Process.* **171**, 107474 (2020)

39. Feldman, J., O'Donnell, R., Servedio, R.A.: Learning mixtures of product distributions over discrete domains. *SIAM J. on Comput.* **37**(5), 1536–1564 (2008)
40. Friedland, S.: Optimal transport, distance between sets of measures and tensor scaling. (2020). Preprint at [arXiv:2005.00945](https://arxiv.org/abs/2005.00945)
41. Gertsbakh, I., Shpungin, Y.: Network reliability and resilience. Springer Science & Business Media, Heidelberg, Germany (2011)
42. Goldberg, L.A., Jerrum, M.: The complexity of ferromagnetic Ising with local fields. *Comb., Probab. and Comput.* **16**(1), 43–61 (2007)
43. Grötschel, M., Lovász, L., Schrijver, A.: The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* **1**(2), 169–197 (1981)
44. Grötschel, M., Lovász, L., Schrijver, A.: Geometric algorithms and combinatorial optimization, vol. 2. Springer Science & Business Media, Heidelberg, Germany (2012)
45. Guo, H., Jerrum, M.: A polynomial-time approximation algorithm for all-terminal network reliability. *SIAM J. on Comput.* **48**(3), 964–978 (2019)
46. Haasler, I., Ringh, A., Chen, Y., Karlsson, J.: Estimating ensemble flows on a hidden Markov chain. In: IEEE Conference on Decision and Control, pp. 1331–1338. IEEE (2019)
47. Haasler, I., Ringh, A., Chen, Y., Karlsson, J.: Multi-marginal optimal transport and Schrödinger bridges on trees. (2020). Preprint at [arXiv:2004.06909](https://arxiv.org/abs/2004.06909)
48. Haasler, I., Ringh, A., Chen, Y., Karlsson, J.: Scalable computation of dynamic flow problems via multi-marginal graph-structured optimal transport. (2021). Preprint at [arXiv:2106.14485](https://arxiv.org/abs/2106.14485)
49. Haasler, I., Singh, R., Zhang, Q., Karlsson, J., Chen, Y.: Multi-marginal optimal transport and probabilistic graphical models. *IEEE Transactions on Information Theory* (2021)
50. Haneveld, W.K.: Robustness against dependence in PERT: An application of duality and distributions with known marginals. In: *Stochastic Programming*, pp. 153–182. Springer (1986)
51. Jaggi, M.: Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In: *International Conference on Machine Learning*, pp. 427–435 (2013)
52. Karger, D.R.: A randomized fully polynomial time approximation scheme for the all-terminal network reliability problem. *SIAM Rev.* **43**(3), 499–522 (2001)
53. Khachiyan, L.G.: Polynomial algorithms in linear programming. *USSR Comput. Math. and Math. Phys.* **20**(1), 53–72 (1980)
54. Kolda, T.G., Bader, B.W.: Tensor decompositions and applications. *SIAM Rev.* **51**(3), 455–500 (2009)
55. Koller, D., Friedman, N.: Probabilistic graphical models: principles and techniques. MIT Press, Cambridge, Massachusetts, US (2009)
56. Kruskal, J.B.: On the shortest spanning subtree of a graph and the traveling salesman problem. *Proc. of the Am. Math. soc.* **7**(1), 48–50 (1956)
57. Léonard, C.: A survey of the Schrödinger problem and some of its connections with optimal transport (2013). Preprint at [arXiv:1308.0215](https://arxiv.org/abs/1308.0215)
58. Lin, T., Ho, N., Cuturi, M., Jordan, M.I.: On the complexity of approximating multimarginal optimal transport (2019). Preprint at [arXiv:1910.00152](https://arxiv.org/abs/1910.00152)
59. Linial, N., Samorodnitsky, A., Wigderson, A.: A deterministic strongly polynomial algorithm for matrix scaling and approximate permanents. In: *Symposium on Theory of Computing*, pp. 644–652 (1998)
60. Makarov, G.: Estimates for the distribution function of a sum of two random variables when the marginal distributions are fixed. *Theory of Probab. and its Appl.* **26**(4), 803–806 (1982)
61. Meilijson, I., Nádas, A.: Convex majorization with an application to the length of critical paths. *J. of Appl. Probab.* **16**(3), 671–677 (1979)
62. Mishra, V.K., Natarajan, K., Padmanabhan, D., Teo, C.-P., Li, X.: On theoretical and empirical aspects of marginal distribution choice models. *Management Sci.* **60**(6), 1511–1531 (2014)
63. Moore, E.F., Shannon, C.E.: Reliable circuits using less reliable relays. *J. of the Franklin Institute* **262**(3), 191–208 (1956)
64. Muzellec, B., Nock, R., Patrini, G., Nielsen, F.: Tsallis regularized optimal transport and ecological inference. In: *AAAI Conference on Artificial Intelligence*, volume 31 (2017)
65. Nadas, A.: Probabilistic PERT. *IBM J. of Res. and Dev.* **23**(3), 339–347 (1979)
66. Natarajan, K., Song, M., Teo, C.-P.: Persistency model and its applications in choice modeling. *Management Sci.* **55**(3), 453–469 (2009)
67. Nenna, L.: Numerical methods for multi-marginal optimal transportation. PhD thesis (2016)
68. Padmanabhan, D., Ahipasaoglu, S.D., Ramachandra, A., Natarajan, K.: Extremal probability bounds in combinatorial optimization (2021). Preprint at [arXiv:2109.01591](https://arxiv.org/abs/2109.01591)

69. Papadimitriou, C.H., Roughgarden, T.: Computing correlated equilibria in multi-player games. *J. of the ACM* **55**(3), 1–29 (2008)
70. Peyré, G., Cuturi, M.: Computational optimal transport. *Foundations and Trends in Machine Learning* (2017)
71. Pitié, F., Kokaram, A.C., Dahyot, R.: Automated colour grading using colour distribution transfer. *Comput. Vision and Image Underst.* **107**(1–2), 123–137 (2007)
72. Provan, J.S., Ball, M.O.: The complexity of counting cuts and of computing the probability that a graph is connected. *SIAM J. on Comput.* **12**(4), 777–788 (1983)
73. Quanrud, K.: Approximating optimal transport with linear programs. In: *Symposium on Simplicity in Algorithms* (2018)
74. Rabin, J., Peyré, G., Delon, J., Bernot, M.: Wasserstein barycenter and its application to texture mixing. In: *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 435–446. Springer (2011)
75. Rüschendorf, L.: Random variables with maximum sums. *Advances in Applied Probability*, pp. 623–632 (1982)
76. Rüschendorf, L., Uckelmann, L.: On the n-coupling problem. *J. of Multivar. Anal.* **81**(2), 242–258 (2002)
77. Singh, R., Haasler, I., Zhang, Q., Karlsson, J., Chen, Y.: Incremental inference of collective graphical models. *IEEE Control Systems Letters* (2020)
78. Sinkhorn, R.: Diagonal equivalence to matrices with prescribed row and column sums. *The Am. Math. Monthly* **74**(4), 402–405 (1967)
79. Solomon, J., De Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., Guibas, L.: Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Trans. on Graphics* **34**(4), 1–11 (2015)
80. Srivastava, S., Li, C., Dunson, D.B.: Scalable Bayes via barycenter in Wasserstein space. *J. of Mach. Learning Res.* **19**(1), 312–346 (2018)
81. Stoer, M., Wagner, F.: A simple min-cut algorithm. *J. of the ACM* **44**(4), 585–591 (1997)
82. Teh, Y.W., Welling, M.: The unified propagation and scaling algorithm. *Advances in Neural Information Processing Systems*, 953–960 (2002)
83. Trefethen, L.N.: Approximation theory and approximation practice, volume 164. SIAM (2019)
84. Tupitsa, N., Dvurechensky, P., Gasnikov, A., Uribe, C.A.: Multimarginal optimal transport by accelerated alternating minimization. In: *2020 IEEE Conference on Decision and Control*, pp. 6132–6137. IEEE (2020)
85. Valiant, L.G.: The complexity of enumeration and reliability problems. *SIAM J. on Comput.* **8**(3), 410–421 (1979)
86. Villani, C.: Topics in optimal transportation. Number 58. American Mathematical Society, Providence, Rhode Island, US (2003)
87. Wainwright, M.J., Jordan, M.I.: Variational inference in graphical models: The view from the marginal polytope. In *Allerton Conf. on Commun., Control, and Comput.* **41**, 961–971 (2003)
88. Wainwright, M.J., Jordan, M.I.: Graphical models, exponential families, and variational inference. Now Publishers Inc, Hanover, Massachusetts, US (2008)
89. Weiss, G.: Stochastic bounds on distributions of optimal value functions with applications to PERT, network flows and reliability. *Oper. Res.* **34**(4), 595–605 (1986)
90. Wilson, A.G.: The use of entropy maximising models, in the theory of trip distribution, mode split and route split. *Journal of Transport Economics and Policy*, 108–126 (1969)
91. Young, N.E.: Sequential and parallel algorithms for mixed packing and covering. In: *Symposium on Foundations of Computer Science*, pp. 538–546. IEEE (2001)
92. Yudin, D.B., Nemirovskii, A.S.: Informational complexity and efficient methods for the solution of convex extremal problems. *Matekon* **13**(2), 22–45 (1976)
93. Zemel, E.: Polynomial algorithms for estimating network reliability. *Netw.* **12**(4), 439–452 (1982)