CSIAM Trans. Appl. Math.
doi: 10.4208/csiam-am.SO-2024-0025

Vol. 7, No. 1, pp. 86-119
March 2026

A Numerical Algorithm with Linear Complexity for Multi-Marginal Optimal Transport with L^1 Cost

Chunhui Chen[1], Jing Chen[2], Baojia Luo[3], Shi Jin[4] and Hao Wu[1,*]

[1] *Department of Mathematical Sciences, Tsinghua University, Beijing 100084, China.*
[2] *School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 639798.*
[3] *Theory Lab, Central Research Institute, 2012 Labs, Huawei Technologies Co. Ltd., Hong Kong, SAR, China.*
[4] *School of Mathematical Sciences, Institute of Natural Sciences, and MOE-LSC Shanghai Jiao Tong University, Shanghai 200240, China.*

Received 15 June 2024; Accepted 22 April 2025

Abstract. Numerically solving multi-marginal optimal transport (MMOT) problems is computationally prohibitive, even for moderate-scale instances involving $l \geq 4$ marginals with support sizes of $N \geq 1000$. The cost in MMOT is represented as a tensor with N^l elements. Even accessing each element once incurs a significant computational burden. In fact, many algorithms require direct computation of tensor-vector products, leading to a computational complexity of $\mathcal{O}(N^l)$ or beyond. In this paper, inspired by our previous work [Comm. Math. Sci., 20, 2022], we observe that the costly tensor-vector products in the Sinkhorn Algorithm can be computed with a recursive process by separating summations and dynamic programming. Based on this idea, we propose a fast tensor-vector product algorithm to solve the MMOT problem with L^1 cost, achieving a miraculous reduction in the computational cost of the entropy regularized solution to $\mathcal{O}(N)$. Numerical experiment results confirm such high performance of this novel method which can be several orders of magnitude faster than the original Sinkhorn algorithm.

AMS subject classifications: 49M25, 65K10

Key words: Multi-marginal optimal transport, Sinkhorn algorithm, linear complexity, entropy regularization.

*Corresponding author. *Email addresses:* hwu@tsinghua.edu.cn (H. Wu), cch21@mails.tsinghua.edu.cn (C. Chen), jing.chen@ntu.edu.sg (J. Chen), luobaojia2@huawei.com (B. Luo), shijin-m@sjtu.edu.cn (S. Jin)

1 Introduction

Multi-marginal optimal transport, first proposed by Gangbo and Święch [18], is an extension of the classical optimal transport problem [7,36]. It aims to find an optimal transport plan that minimizes total cost while fitting multiple marginal distributions. MMOT problems naturally arise in various fields, such as machine learning [2, 5, 11, 20, 32], incompressible fluid dynamics [4,8,42], density function theory [10,14,22,23], Schrödinger bridge [12, 19], and tomographic reconstruction [1], and thus have attracted wide attention in recent years.

However, the heavy computational burden of solving general MMOT problems limits its broad application. For l-marginal distributions with support sizes of N, the l-th order cost tensor in MMOT problems contains N^l elements. To fully obtain the information of the cost tensor, it is inevitable to repeatedly access all elements in the tensor, leading to a significant computational cost. For example, the generalized Sinkhorn algorithms [6, 7, 30, 34, 40] require repeated computation of tensor-vector products, resulting in a computational complexity of $\mathcal{O}(N^l)$. More severely, directly solving linear programming problems has a computational complexity of $N^{\mathcal{O}(l)}$. Therefore, these methods remain computationally prohibitive even for moderate-scale MMOT problems. Some modified algorithms with lower computational complexity have been proposed for specific MMOT problems, such as the MMOT problem with a tree structure [19,41] and the Wasserstein barycenter [3,6,39].

In this work, we propose a novel implementation of the Sinkhorn algorithm for solving the entropy regularized MMOT problem with L_1 cost applications in image processing [38], computer vision [35] and seismic tomography [16,33,43], which has linear computational complexity relative to support size N. This work is a follow-up work of the fast Sinkhorn algorithms [28, 29] , which observe the special structure of the kernel matrix with Wasserstein-1 metric and utilize dynamic programming techniques [21, 24–26], achieving linear computational complexity for solving up to 2-marginal optimal transport problem. Unlike the previous situation, the kernel matrix evolves into an l-th order tensor in the l-marginal optimal transport problem. The computational burden of the Sinkhorn algorithm becomes prohibitive due to the $\mathcal{O}(N^l)$ operations required by the tensor-vector products. To address this problem, we observe a similar special structure of the kernel tensor and accelerate the tensor-vector products using the series rearrangement and dynamic programming techniques [25], which results in a fast Sinkhorn algorithm with $\mathcal{O}(N)$ computational complexity.

The rest of the paper is organized as follows. In Section 2, we review the MMOT problem and the generalized Sinkhorn algorithm for the 1-dimensional (1D) and 3-marginal case. Then, we introduce the key fast tensor-vector product technique and provide a detailed implementation of our algorithm in Section 3. This algorithm can be conducted in more general scenarios, such as high-dimensional and l-marginal optimal transport problems, which are presented in Section 4. In Section 5, numerical experiments are carried