

BSP-OT: Sparse transport plans between discrete measures in loglinear time

Supplementary material

BAPTISTE GENEST, Université Claude Bernard Lyon 1, CNRS, INSA Lyon, France
NICOLAS BONNEEL, CNRS, Université Claude Bernard Lyon 1, INSA Lyon, France
VINCENT NIVOLIER, Université Claude Bernard Lyon 1, CNRS, INSA Lyon, France
DAVID COEURJOLLY, CNRS, Université Claude Bernard Lyon 1, INSA Lyon, France

ACM Reference Format:

Baptiste Genest, Nicolas Bonneel, Vincent Nivoliers, and David Coeurjolly. 2025. BSP-OT: Sparse transport plans between discrete measures in loglinear time Supplementary material. *ACM Trans. Graph.* 44, 6 (December 2025), 2 pages. <https://doi.org/10.1145/3763281>

A Statistics of connected components sizes

In Fig. 1-a, we evaluate the sorted connected components' size, averaged on 1000 runs while (monotone) merging the current bijection T and the k -th bijection T_k (for k from 1 to 64). We first observe a few large connected components and many small ones in this experiment. As k increases, the number of connected components gets smaller but with larger components.

B Accuracy when considering other transport metrics

The merging strategies depend on a ground metric that defines the transport cost we aim to minimize. In Fig. 1-b, we plot the relative error, as k increases, for various metrics of the form :

$$\mathcal{W}_{p,q}(T) = \left(\sum_{i=1}^n d_p(x_i, T(x_i))^q \right)^{\frac{1}{q}}, \quad (1)$$

where d_p is the Euclidean p -norm $d_p(x, y) = (\sum |x_i - y_i|^p)^{\frac{1}{p}}$. We observe that our pipeline reasonably approximates optimal transport regardless of the considered metric as k increases.

C Orthogonal slicing strategies

Another strategy for building BSP matchings is to use an orthogonal slicing strategy, i.e. using alternating orthogonal directions, (as in kd-tree matchings of Nurbekyan et al. [2020]). In order to randomize this strategy, we perform random rotations to μ and ν , with either ρ fixed to $\frac{1}{2}$, or random as-well. While cheaper in practice than Gaussian slicing, this strategy also has a $O(d^3)$ complexity due to rotation sampling. As illustrated in Fig. 2, this strategy gives comparable results to random slicing (left). It is also an order of magnitude less efficient than Gaussian slicing (right).

References

Levon Nurbekyan, Alexander Iannantuono, and Adam M Oberman. 2020. No-collision transportation maps. *Journal of Scientific Computing* 82, 2 (2020), 45.

Authors' Contact Information: Baptiste Genest, Université Claude Bernard Lyon 1, CNRS, INSA Lyon, France, baptiste.genest@liris.cnrs.fr; Nicolas Bonneel, CNRS, Université Claude Bernard Lyon 1, INSA Lyon, France, nicolas.bonneel@liris.cnrs.fr; Vincent Nivoliers, Université Claude Bernard Lyon 1, CNRS, INSA Lyon, France, vincent.nivoliers@liris.cnrs.fr; David Coeurjolly, CNRS, Université Claude Bernard Lyon 1, INSA Lyon, France, david.coeurjolly@cnrs.fr.



This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2025 Copyright held by the owner/author(s).
ACM 1557-7368/2025/12-ART
<https://doi.org/10.1145/3763281>




Fig. 1. (a) When considering a pair of 1000 random points, we plot the average size of the connected components when using the monotone merging (averaged over 1000 runs, one color per component). (b) Comparison of relative error with number of bijection merged across various transport metrics, using random slicing between the armadillo and the ball in 3D, over 100 batches.




Fig. 2. Comparisons of the evolution of the relative error as a function of number of bijections merged using orthogonal slicing: Comparison with random slicing (left), and Gaussian Slicing (right). Curves with long dashes for orthogonal slices with ρ random and short dashes for orthogonal slices with $\rho = 0.5$.