HAL
open science

A survey of Optimal Transport for Computer Graphics and Computer Vision

Nicolas Bonneel, Julie Digne

To cite this version:

Nicolas Bonneel, Julie Digne. A survey of Optimal Transport for Computer Graphics and Computer Vision. Computer Graphics Forum, 2023, 42 (2). hal-04029613

HAL Id: hal-04029613

https://hal.science/hal-04029613v1

Submitted on 4 Apr 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

EUROGRAPHICS 2023
A. Bousseau and C. Theobalt
(Guest Editors)

Volume 42 (2023), Number 2
STAR – State of The Art Report

A survey of Optimal Transport for Computer Graphics and Computer Vision

Nicolas Bonneel[1] and Julie Digne[1]

[1]CNRS, Univ. Lyon, France

Abstract

Optimal transport is a long-standing theory that has been studied in depth from both theoretical and numerical point of views. Starting from the 50s this theory has also found a lot of applications in operational research. Over the last 30 years it has spread to computer vision and computer graphics and is now becoming hard to ignore. Still, its mathematical complexity can make it difficult to comprehend, and as such, computer vision and computer graphics researchers may find it hard to follow recent developments in their field related to optimal transport. This survey first briefly introduces the theory of optimal transport in layman's terms as well as most common numerical techniques to solve it. More importantly, it presents applications of these numerical techniques to solve various computer graphics and vision related problems. This involves applications ranging from image processing, geometry processing, rendering, fluid simulation, to computational optics, and many more. It is aimed at computer graphics researchers desiring to follow optimal transport research in their field as well as optimal transport researchers willing to find applications for their numerical algorithms.

CCS Concepts

• ***Computing methodologies*** → ***Computer graphics; Computer vision;***

1. Introduction

Optimal transport is a theory that has been mainly used to compare probability distributions and interpolate between them. The optimal transport problem was first stated by Monge in 1781 [Mon81] as the problem of moving a pile of earth from one location to another with minimum effort. While Monge did not succeed in solving this problem, Kantorovich did more than 150 years later, with his new linear programming formalism, in 1942 [Kan42]. While the problem was formally solved, fast algorithmical solutions remained limited. For decades this problem has thus remained computationally intractable even for moderately coarse discretizations. With recent advances in numerical methods and solvers, optimal transport has finally found its way to solve realistic computer graphics and vision related problems. It is now routinely used to solve problems in image processing, geometry processing, rendering, and animation, as well as many other fields – problems we will cover in breadth throughout this survey.

While the mathematical theory behind optimal transport and numerical techniques are largely presented in recent books [Vil09, San15, PC17] and the state-of-the-art in numerical solvers is extensively developed in a recent survey [PC19], our state-of-the-art report focuses on computer graphics and vision applications. Applications in image processing are more exhaustively covered in the Habilitation of Papadakis [Pap15]. The aim of this survey is twofold. First, it gives a brief and gentle introduction to optimal transport for computer graphics and computer vision researchers who are not necessarily comfortable with heavy mathematical developments but desiring to follow optimal transport research in their own field. Second, it gives optimal transport researchers a set of potentially exciting applications to their research. To our knowledge, these aspects are not covered by existing surveys. Researchers familiar with optimal transport can directly skip to the application section they are interested in. These applications are organized into sections corresponding to different fields of computer graphics and computer vision: image processing (sec. 4), rendering (sec. 5), geometry processing (sec. 6), topological data analysis (sec. 7), and simulation and animation (sec. 8). Table 6 is a good starter for an overview of the possible applications of optimal transport.

2. Principles of optimal transport

While this report does not aim at fully exposing optimal transport theory and numerical techniques, we briefly present the problem in layman's terms as well as broad classes of solvers. This will involve sloppy notations and definitions that might hurt mathematically inclined researchers – proper definitions and introductions to optimal transport can be found in other work [San15, PC19]. For instance, our exposition will not require knowledge of measure theory, and we will merely refer to the optimal transport of nonnegative functions or, at worst, set of Dirac distributions, seen as point sets.

© 2023 The Author(s).
Computer Graphics Forum published by Eurographics - The European Association for Computer Graphics and John Wiley & Sons Ltd.
This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

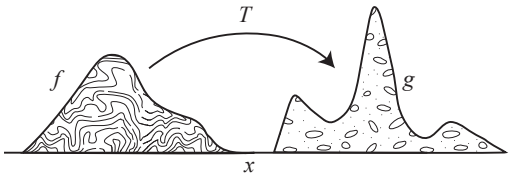

Figure 1: *The Monge problem seeks the minimum cost transport map T allowing to reshape a function f to another function g of the same total integral (mass). This figure illustrates the 1-d case, but the theory extends to functions over arbitrary dimensional spaces, and even Riemannian manifolds.*

2.1. The Wasserstein cost

2.1.1. The Monge problem

Monge was interested in moving a pile of earth from one location to another with minimum amount of effort (Fig. 1). He formulated the problem as an energy minimization problem, considering the cost of moving the pile of earth (or of sand) as the sum of the costs of moving each of its individual particles, each particle of earth being moved from location $x \in X$ to location $T(x) \in X$. The cost of moving each particle from a location x to a location y is denoted $c(x,y)$, and is also called *ground distance*. Monge initially used $c(x,y) = \|x-y\|$, the distance traveled by the particle of sand, though this renders the problem much more complicated. For many practical applications in computer graphics, the *quadratic* cost $c(x,y) = \|x-y\|^2$ is used, i.e., the squared distance, which has nice properties and physical interpretation. We assume that the pile of earth to move has a shape that can be represented by a function f over domain X, and we want to give it a new shape g elsewhere in the domain X. The Monge problem defines the transport cost $W(f,g)$ between f and g as follows:

$$W(f,g) = \min_T \int_X c(x,T(x)) f(x) \mathrm{d}x \tag{1}$$

$$\text{s.t.} \quad f(x) = g(x)\,|\det J_T(x)| \;\forall x \in X \tag{2}$$

The optimization is performed over all possible ways to reshape the function f using all possible warps T. Here, X is here an n-dimensional Euclidean space, and the formulation extends to surfaces or Riemannian manifolds, replacing distances with geodesic distances). Denoting $x = (x_0, x_1, \ldots, x_n)$ the space variable, the optimization constraint 2 involves the Jacobian J_T of the transformation $T = (T_1, \ldots, T_n)$ defined at x as:

$$J_T(x) = \left[\frac{\partial T_i(x)}{\partial x_j}\right]_{i,j} \tag{3}$$

and enforces "mass preservation", i.e., that after reshaping f by the warp T, the pile of sand matches g. One can see $y = T(x)$ as a change of variable, and constraint 2 as the well-known "change of variable formula" for integration.

Of course, way beyond moving earth, this problem allows to reshape actual functions. The *transport cost* $W(f,g)$ gives the optimal cost of moving f towards g, assuming that the total mass (i.e., the integral of) f equals that of g, and that these functions are never negative. A particular case of "non-negative functions having the same integral" is non-negative functions of integral 1. This arises notably when dealing with probability density functions. This framework has thus met with large success in the probabilistic setting, to transport probability density functions and compare them with this transport cost $W(f,g)$ (it also generalizes to probability distributions).

If $c(x,y) = c(y,x)$, we have $W(f,g) = W(g,f)$, and if $c(x,y)$ is a distance between points x and y, then $W(f,g)$ is also a distance between functions f and g. In this case, this guarantees that $W(f,g) = W(g,f)$, that $W(f,g) = 0$ iff $f = g$, and that W respects the triangle inequality $W(f,g) \geq W(f,h) + W(h,g) \quad \forall h$. In fact, when $c(x,y) = \|x-y\|^p$ for $p \geq 1$, the function $W_p(f,g) = W(f,g)^{1/p}$ is called *Wasserstein-p distance* and is also a distance between functions f and g. This distance is also called *Earth Mover's Distance* (EMD), in reference to the Monge problem.

2.1.2. The Kantorovich problem

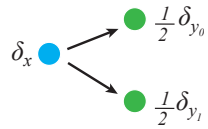

The Monge problem is non-linear, extremely difficult to solve in general, and may not have solutions in contexts useful to the computer scientist. For instance, replacing the function f by a Dirac distribution δ_x and function g by a sum of two half Dirac distributions $\frac{1}{2}\delta_{y_0} + \frac{1}{2}\delta_{y_1}$, it becomes impossible to find a one-to-one warp T that sends the mass at a single x to two different locations y_0 and y_1. Dirac distributions are often encountered when dealing with discrete problems, e.g., attempting to discretize a continuous problem or when sampling from a probability distribution. Fortunately, Kantorovich came up with a solution that both simplifies the resolution of the optimal transportation problem, and generalizes the formulation to extend it to a broader class of objects. Now, one does not attempt to find a *warp* (more often called *transport map*, see Fig. 2 right) $T : X \to X$ that maps each particle to its new location, but a *transport plan* $P : X \times X \to \mathbb{R}$ that indicates how much mass is transported between two locations. This led to the (sloppy) formulation[†]:

$$W(f,g) = \min_P \int_{X\times X} c(x,y) P(x,y) \mathrm{d}x\,\mathrm{d}y \tag{4}$$

$$\text{s.t.} \quad \int_X P(x,y)\mathrm{d}y = f(x) \quad \forall x \in X \tag{5}$$

$$\int_X P(x,y)\mathrm{d}x = g(y) \quad \forall y \in X \tag{6}$$

$$P(x,y) \geq 0 \quad \forall (x,y) \in X \times X \tag{7}$$

where the optimization is now performed over all transport plans. The constraints on the transport plans express the fact that the total mass moved from each location x corresponds to the mass $f(x)$ present at location x (eq. 5), that the total mass received at location y corresponds to what is needed at that location, that is $g(y)$ (eq. 6),

† Beware that more formally, P is a measure and requires more care in writing the integrals which should read $\int_{X\times X} c(x,y) dP(x,y)$. Similarly, constraints 5,6 would require the notion of push-forward, which is beyond the scope of our survey.

© 2023 The Author(s)

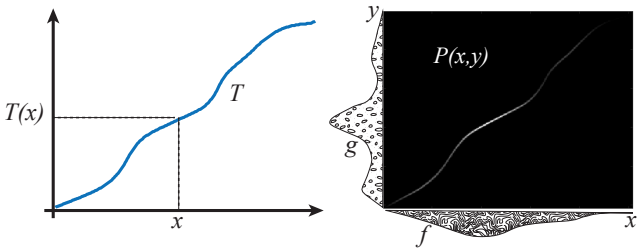

Figure 2: *Solution to the problem of Fig. 1.* ***Left****. The Monge solution yields a transport map T that tells where the mass of f located at x should move.* ***Right.*** *The Kantorovich solution yields a transport plan P that tells how much mass is moved between all pairs (x,y) of locations. This transport plan is very sparse. When there is no Dirac or oddities, both solutions coincide, and the transport plan is the graph of the transport map. In 1-d like here, this graph is monotonically increasing. More generally, in higher dimension, it is the gradient of a convex function.*

and that the mass transport from each location and received by another location cannot be negative (eq. 7), all these constraints enforce mass conservation.

One advantage of this formulation is that it is more flexible, and when Monge's solution T exists, it coincides with that of Kantorovich by taking $P = (Id, T)$ with Id the identity function (Fig. 2). Another advantage is that it yields a linear program. This is more easily seen, from a computer scientist perspective, if one discretizes the function f at m discrete locations $(x_i)_{i=1..m} \in X^m$ and function g at n locations $(y_j)_{j=1..n} \in X^n$:

$$W(f,g) = \min_P \sum_{1\le i,j\le m,n} c(x_i,y_j)P_{i,j} \tag{8}$$

$$\text{s.t.} \quad \sum_{j=1}^{n} P_{i,j} = f(x_i) \quad \forall i \in \{1\dots m\} \tag{9}$$

$$\sum_{i=1}^{m} P_{i,j} = g(y_j) \quad \forall j \in \{1\dots n\} \tag{10}$$

$$P_{i,j} \ge 0 \quad 1 \le i,j \le m,n \tag{11}$$

where the optimization is now performed over all $m \times n$ matrices $P \in \mathbb{R}^{m\times n}$. This indeed corresponds to minimizing a linear combination of variables, under both $m+n$ linear equality constraints and mn positivity constraints.

Since it corresponds to a linear program, it admits a dual formulation with the same optimal cost, which, in its continuous formulation, gives:

$$W(f,g) = \sup_{\phi,\psi} \int_X \phi(y)g(y)\mathrm{d}y + \int_X \psi(x)f(x)\mathrm{d}x \tag{12}$$

$$\text{s.t.} \quad \phi(y) + \psi(x) \le c(x,y) \quad \forall (x,y) \in X \times X \tag{13}$$

where $\psi : X \to \mathbb{R}$ and $\phi : X \to \mathbb{R}$ are dual variables (the plus sign for ψ depends on conventions, you may find it with a minus). Duality can probably again be better understood from a computer scientist

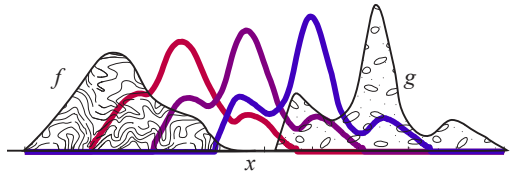

Figure 3: *Using the map $T_t = (1-t)Id + tT$, one can smoothly interpolate between the two functions shown in Fig. 1. This is called displacement interpolation (here with $t = 0.25, 0.5, 0.75$, from red to blue.*

point of view in the discrete setting:

$$W(f,g) = \max_{\phi,\psi} \sum_{j=1}^{n} \phi_j g(y_j) + \sum_{i=1}^{m} \psi_i f(x_i) \tag{14}$$

$$\text{s.t.} \quad \phi_j + \psi_i \le c(x_i,y_j) \quad 1 \le i,j \le m,n \tag{15}$$

which now optimizes vectors $\phi \in \mathbb{R}^n$ and $\psi \in \mathbb{R}^m$ under mn inequality constraints. Intuitively, this corresponds to the problem of an transport company who charges the consumer an amount of money $\psi(x)$ for loading their truck at location x with 1 unit of mass and charging again $\phi(y)$ for unloading their truck at location y of 1 unit of mass. The company's goal is to maximize their profit, under the constraint that the cost of loading plus unloading one unit of mass is not greater than the cost $c(x,y)$ defined in the primal problem. The dual formulation is sometimes used in linear program solvers, and for solving the semi-discrete optimal transport problem (see sec. 3.2).

2.2. Wasserstein barycenters

As mentioned previously, the optimal transport framework produces a distance between non-negative functions that have the same integral when the cost of moving a particle from x to y is $c(x,y) = \|x-y\|^p$. This notably allows to produce interpolations between functions. The case $p = 2$ is particularly interesting. When displacing particles of mass from x to their target $T(x)$ (e.g., assuming the Monge formulation), one may stop the motion of all particles simultaneously, i.e., replacing the map T by $T_t = (1-t)Id + tT$ for some time parameter t. Doing so continuously warps f towards g and the resulting interpolation, $f_t(x) = f(T_t(x))|\det J_{T_t}(x)|$, is called *displacement interpolation* (Fig. 3). Displacement interpolation can also be interpreted as moving along a geodesic between the two functions with respect to the Wasserstein metric: any point on this geodesic corresponds to an intermediate function during the interpolation.

This interpolation can also be obtained by solving the following optimization problem that resembles the one used for obtaining a linear interpolation between any two points (replacing W_2 by $\|.\|$):

$$f_t = \mathrm{argmin}_h (1-t)W_2^2(h,f) + tW_2^2(h,g). \tag{16}$$

For instance, for $t = 0.5$, this amounts to finding a function h that is *as equidistant as possible* to the functions f and g in term of Wasserstein-2 distance.

The formulation as an energy minimization problem allows to

© 2023 The Author(s)

generalize the notion of displacement interpolation to the case where one wants to interpolate between more than two non-negative functions of the same integral. This leads to the notion of *Wasserstein barycenter* of a set of functions $\{f_k\}_k$ with their corresponding barycentric weights $\{\lambda_k\}_k$ as:

$$b = \text{argmin}_h \sum_k \lambda_k W_2^2(h, f_k) \tag{17}$$

For instance, the *Wasserstein centroid* of three functions f_0, f_1, f_2 would be the Wasserstein barycenter of $\{(f_0, \frac{1}{3}), (f_1, \frac{1}{3}), (f_2, \frac{1}{3})\}$, i.e. $b = \text{argmin}_h W_2^2(h, f_1) + W_2^2(h, f_2) + W_2^2(h, f_3)$.

2.3. Particular cases

There are simple cases which can be trivially solved.

2.3.1. The 1-d case

Computing the optimal transport map T between two 1-d functions f and g can be easily achieved when the cost $c(x,y)$ is a convex function of the distance $\|x-y\|$ (e.g., when computing W_p for $p > 1$). We consider the cumulative sums:

$$F(x) = \int_{t=-\infty}^{x} f(t)\text{d}t \tag{18}$$

$$G(x) = \int_{t=-\infty}^{x} g(t)\text{d}t \tag{19}$$

The optimal transport is then:

$$T(x) = G^{-1}(F(x)) \tag{20}$$

Here G^{-1} denotes the inverse function. However, since g may be sometimes zero, it means $G(x)$ may be constant for a certain range of x values, thus, G may not be injective, and G^{-1} may be ill defined. It is thus necessary to define a more general inverse function:

$$G^{-1}(x) = \inf\{y \in \mathbb{R}; G(y) > x\} \tag{21}$$

In the discrete setting, F and G are obtained using discrete cumulative sums (e.g., Matlab's *cumsum*) or higher order integration schemes (e.g., Newton-Cotes), and the inverse function can be obtained using a line search. Computing the map for consecutive discrete values of x allows to start the line search at the previously found location, producing an $\mathcal{O}(m+n)$ cost algorithm.

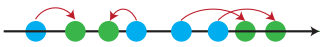

In the case where f and g actually consist in a sum of the same number of uniformly weighted Diracs, a transport map merely assign each Dirac of f to a Dirac of g. It can be shown that in the convex cost setting above, one can get the optimal transport map simply by sorting the Diracs along the line from left to right, and associating the kth Dirac of f with the kth Dirac of g.

In 1-d with convex cost, Wasserstein barycenters can also be easily achieved by considering:

$$B^{-1} = \sum_k \lambda_k F_k^{-1} \tag{22}$$

This allows to obtain B^{-1}, the inverse cumulative sum of the desired barycenter b, as a function of all $\{F_k^{-1}\}_k$, the inverse cumulative sums of each input distribution. In the discrete setting, once the inverse cumulative sum of the result is numerically computed, one can obtain the Wasserstein barycenter b by numerical inversion followed by numerical differentiation.

Costs that are concave functions of the ℓ^2 distance between particles leads to much more complex algorithms [DSS12], and other 1-d extensions have been proposed (e.g., transporting on the circle [RDG11] or transporting Diracs towards a larger number of Diracs [BC19]).

To simply compute the optimal transport cost, for more general costs of the form $c(x,y) = c(x-y)$, we can use [Vil09]:

$$W(f,g) = \int_0^1 c(F^{-1}(x) - G^{-1}(x))\text{d}x \tag{23}$$

One dimensional transport is the basis for higher-dimensional techniques that rely on 1-d projections.

2.3.2. Transporting Gaussians

In some cases, and in particular high-dimensional datasets, the distribution of data is well approximated by a Gaussian distribution, and one needs to merely transport Gaussians. This can be easily achieved. Interestingly, one can show that Wasserstein barycenters of Gaussians are also Gaussians.

For the quadratic cost, the optimal transport cost between two d-dimensional Gaussians $\mathcal{N}_0 = \mathcal{N}(\mu_0, \Sigma_0)$ and $\mathcal{N}_1 = \mathcal{N}(\mu_1, \Sigma_1)$ of means $\mu_0 \in \mathbb{R}^d$ and $\mu_1 \in \mathbb{R}^d$ and covariance matrices $\Sigma_0 \in \mathbb{R}^{d\times d}$ and $\Sigma_1 \in \mathbb{R}^{d\times d}$ is obtained by:

$$W_2(\mathcal{N}_0, \mathcal{N}_1) = \text{tr}(\Sigma_0 + \Sigma_1 - 2\Sigma_{0,1}) + \|\mu_0 - \mu_1\|^2 \tag{24}$$

with $\Sigma_{0,1} = (\Sigma_0^{1/2}\Sigma_1\Sigma_0^{1/2})^{1/2}$. The square roots $^{1/2}$ are understood as matrix square roots (involving square roots of eigenvalues).

The displacement interpolation between two Gaussians $\mathcal{N}_0$ and $\mathcal{N}_1$ is the Gaussian $\mathcal{N}_t = \mathcal{N}((1-t)\mu_0 + t\mu_1, \Sigma_t)$, with

$$\Sigma_t = [(1-t)\text{Id} + tC]\ \Sigma_0\ [(1-t)\text{Id} + tC] \tag{25}$$

with $C = \Sigma_1^{1/2}\Sigma_{0,1}^{+}\Sigma_1^{1/2}$, denoting $\Sigma_{0,1}^{+}$ the Moore-Penrose pseudo-inverse of $\Sigma_{0,1}$.

When considering more than two Gaussians $\{\mathcal{N}_k = \mathcal{N}(\mu_k, \Sigma_k)\}_k$ with weights $\{\lambda_k\}_k$, their Wasserstein barycenter is $b = \mathcal{N}(\sum_k \lambda_k \mu_k, \Sigma)$ with Σ obtained by solving a fixed point equation:

$$\Sigma^{(n+1)} = \sum_k \lambda_k \left(\sqrt{\Sigma^{(n+1)}}\Sigma_k\sqrt{\Sigma^{(n+1)}}\right)^{1/2} \tag{26}$$

3. Numerical methods

Again, exhaustive coverage of numerical solvers can be found in the recent survey of Peyré and Cuturi [PC19], and this section is aimed at briefly exposing broad classes of solvers.

© 2023 The Author(s)

3.1. LP solvers

Linear programming (LP) solvers provide exact solutions to the discrete optimal transport problem. The optimal transport linear program can be seen as a graph, where each of the m and n discrete locations of the input and target functions are nodes, and moving mass from location i to location j (i.e., when the entry (i, j) of the transport plan matrix P is non-zero) amounts to connecting node i with node j with an arc. An interesting property of this LP is that it possess at most $m+n-1$ non-zero variables, i.e., the graph of the optimal transport plan has at most $m+n-1$ arcs with non-zero flow, among mn possible arcs, which makes it suitable for sparse graph representations. Displacement interpolation between two distributions can be obtained in this way by advecting particles along arcs [BvdPPH11]. When the particles are uniformly weighted and $m = n$, the optimal transport problems amounts to a Linear Assignment Problem (LAP), i.e., the transport plan is actually a one-to-one mapping, and there are only m non-zero arcs. Note that while an LP formulation of Wasserstein barycenters (for more than two distributions) is possible, existing solvers are limited to very small problems [ABM16].

For the most general problem, with arbitrary cost functions, the network simplex algorithm has shown to work fastest [BvdPPH11]. However, memory storage can still become a bottleneck for values of $m = n \approx 40,000$ (problems that typically take minutes to solve, though this heavily depends on the cost function).

For problems structured on a grid with quadratic ground distance, it becomes possible to benefit from the separability of the quadratic cost and split the optimal transport problem along each axis, leading to faster solutions [ABGV18] (about 50x faster than default network simplex for 128x128 grids, and works with 512x512 grids). Multiscale approaches may also be used in the Cartesian grid setting [Sch16, OR20].

The special case of ℓ^1 ground distance has also received a lot of attention. To alleviate the computational complexity of LP solvers, in the case of 2-d histogram to 2-d histogram mass transportation, Ling and Okada [LO07] proposed to split the particle flow into sequences of motions of magnitude 1. This drastically reduces the number of variables and allows to define a network flow interpretation to solve the problem efficiently. They further reduce the complexity by introducing a tree-based algorithm extending the transportation simplex to the ℓ^1 case yielding the EMD-ℓ^1 algorithm.

A state-of-the-art general purpose solver is the Network Simplex of the LEMON C++ library [DJK11] adapted by Bonneel et al. [BvdPPH11] and found at `https://github.com/nbonneel/network_simplex`.

3.2. Semi-discrete solvers

In the special case where one is interested in transporting a continuous function towards a set of Dirac distributions (of the same total mass), the Monge optimal transport problem leads to a solution that can be expressed with a geometric construction called Laguerre's diagrams [AHA92]. If in addition the ground distance is quadratic, this corresponds to a power diagram, i.e., a Voronoï diagram for which the size of each cell can be controlled with a weighting parameter. An intuitive example of the semi-discrete optimal transport problem consists in a density of population desiring to get their morning bread from a set of bakeries at specific locations. If the bread supplies of all bakeries is unlimited, the optimal choice for each inhabitant would be to go to the nearest bakery to buy its bread – this leads to a Voronoï diagram partition of the space. However, if the sum of the bread supplies of all bakeries exactly equals the sum of the bread demands of all inhabitants, each inhabitant would not necessarily go to its nearest bakery as its bread supply might be exhausted by other customers. Instead, each bakery would attract customers from a connected and convex polygonal neighborhood (in the case of the quadratic ground distance), that may not necessarily contains the bakery location (Fig. 4). The diagram produced by other costs can be much more complex – e.g., for the ground distance $c(x, y) = \|x - y\|$, this leads to the Apollonius diagram, that has curved edges.

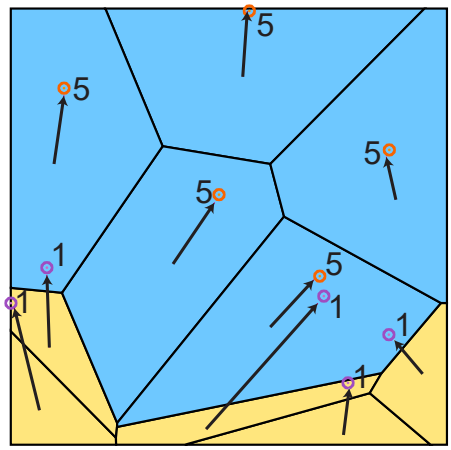

Figure 4: *Semi-discrete optimal transport leads to power diagrams. Here, a set of 10 bakeries can produce either 5 units of bread (orange dots) or 1 unit of bread (pink dots). Inhabitants of a uniform density population are divided among these bakeries such that the total transportation effort is minimized* and *their need for bread is exactly satisfied. This leads to a cell decomposition where each inhabitant in one cell will go to its associated bakery (arrows) – bakeries producing more bread lead to bigger cells (in blue). Note that the bakery sites are not necessarily inside their corresponding cells (here, all bakeries producing 1 unit of bread are located outside of their cell).*

This lends itself to efficient implementations, as power diagrams are well studied geometric constructions that can be obtained by computing a Voronoï diagram in a higher dimensional space. Typical semi-discrete solvers optimize an energy designed so that, at optimality, the area of each power cell (i.e., the polygonal part of the city going towards its optimal bakery, in the above example) matches the mass of each Dirac (i.e., the amount of bread available at that bakery). This leads to fast, multiscale, implementations in 2-d [Mér11] and 3-d [Lév15], or using second order optimization schemes [Lév15, SWS*15, KMT19], and extends to optimal transport on the sphere [CQW*19] or when the masses in the two distributions are different [Lév22]. This is also relatively easy to implement – e.g., as a student project or lab. This however suffers

© 2023 The Author(s)

from the difficulty in implementing these constructions in higher than 3-dimensional spaces.

Typical sizes of distributions handled by this setting by a state-of-the-art solver involves up to millions of particles, in 3-d, on the GPU [vHLM22]. Insights on semi-discrete optimal transport solvers and their implementation can be found in the paper of Lévy and Schwindt [LS18].

A fast semi-discrete C++ implementation in 3-d [Lév15] is available within the WarpDrive plugin of Levy's Graphite software at https://github.com/BrunoLevy/GraphiteThree.

3.3. Entropy regularized solver

By adding the entropy of the transport plan as a regularization term in the Kantorovich minimization problem (eq. 4 and 8), it turns out that the discrete optimal transport problem can be solved with the help of iterative matrix multiplications and elementwise divisions. The resulting algorithm is called the *Sinkhorn algorithm*. It is extremely easy to implement (a couple of lines of Matlab or Python code), can be very fast, works with arbitrary costs and geometric settings, and extends to the construction of Wasserstein barycenters [CD14, SdGP*15]. In the special case of grids and quadratic ground distance, these matrix multiplications can advantageously be replaced by fast Gaussian convolutions [SdGP*15], and the implementation can be even faster and simpler to implement (also lending itself well to labs or projects for students).

The downside is that the entropic regularization adds a small amount of blur to the transport plans: the amount of mass leaving location i will spread a little bit over its target location j (technically, it is spread infinitely, but the mass traveling far away is negligible – see Fig. 5). It will spread even more as the amount of added entropy increases. While it is tempting to make the amount of added entropy tend to zero to recover sharp transport plans or Wasserstein barycenters, this approach leads to numerical instabilities as the amount entropy decreases too much. Performing computations in the log-domain stabilizes these computations, but also makes them much slower [CPSV18].

Another downside of entropy regularized distances, denoted $W_\varepsilon(f,g)$, is that the resulting optimal energy is not a distance anymore, while the corresponding non-regularized problem leads to a distance. Specifically, with the entropic term, we do not have $W_\varepsilon(f,f)=0$ and the triangle inequality also largely breaks. An interesting approach is to debias the result energy, and instead consider $\tilde{W}_\varepsilon(f,g) = W_\varepsilon(f,g) - \frac{1}{2}W_\varepsilon(f,f) - \frac{1}{2}W_\varepsilon(g,g)$. This enforces that $\tilde{W}_\varepsilon(f,f)=0$ by construction, and in practice, reduces the number of cases where triangle inequality is not satisfied. The resulting energy is called the *Sinkhorn divergence* [FSV*19]. This energy can be used in an energy optimization routine to produce (approximate) Wasserstein barycenters. In particular, using a single descent step already produces plausible results, and amounts to linearizing the problem [Fey20].

Nowadays, a state-of-the-art Python solver for Sinkhorn divergences for scattered data, implemented on the GPU, with an additional multiscale strategy, is that of GeomLoss [FSV*19]. It supports discrete problems of about a million of data points within seconds and is currently available at https://www.kernel-operations.io/geomloss/.

Replacing entropy by other types of regularization is possible, such as regularizing by the quadratic norm of the transport plan [ES18], but are far less used in practice for computer vision and graphics applications.

3.4. Sliced optimal transport

An alternative functional for high-dimensional optimal transport relies on optimal transport over projections on random one-dimensional lines [RPDB12, Bon13]. The corresponding distance between two functions is thus merely the average over all 1-d projections of 1-d optimal transport problems over these projections (the projection of a function onto a straight line is the integral of that function over the hyperplane orthogonal to that line passing through each point on that line). This can be computed efficiently using Radon transform in the case of densities stored on a grid, or by projecting points in the case of sums of Dirac distributions [RPDB12, BRPP15]. They lead to either Radon or Sliced optimal transport problems, and can be used to compute Sliced and Radon Wasserstein barycenters. They are also extremely easy to implement and very fast.

Contrary to entropic regularization or Sinkhorn divergences, this does not rely on a parameter that makes the problem converge to the "true" optimal transport problem as it tends to zero. Projecting on 1-d lines makes the problem fundamentally easier to solve, but it does not produce the same result, and does not converge to the classical optimal transport problem, even as the number of 1-d projection increases. However, when the cost writes as a power of the 1-d distance, $c(x,y) = |x-y|^p$, it also produces a distance between functions, similarly to the Wasserstein-p distance. In this case, a (loose) error bounds exist. Denoting $SW_p(f,g)$ the corresponding sliced Wasserstein-p distance in d dimension, one has [Bon13]:

$$SW_p(f,g)^p \leq C_{d,p} W_p(f,g)^p \leq C_{d,p} R^{p-1/(d+1)} SW_p(f,g)^{1/(d+1)} \tag{27}$$

for f and g supported in a ball of radius R and some constant $C_{d,p}$.

While sliced transport is easily implemented, reference Matlab code for sliced and Radon barycenters can be found from the work of Bonneel et al. [BRPP15] at https://github.com/gpeyre/2014-JMIV-SlicedTransport.

3.5. Dynamic formulation

A quantity f_t is preserved when it is advected by a vector field v if it satisfies the conservation PDE:

$$\frac{\partial f_t}{\partial t} + \operatorname{div}(f_t v) = 0 \tag{28}$$

It can be shown that for the quadratic ground distance, a function f_t such that $f_0 = f$ and $f_1 = g$ satisfying this conservation PDE for a vector field $v(x,t)$ minimizing:

$$\min \int_0^1 \int_X \|v(x,t)\|^2 f_t(x,t) \mathrm{d}x\mathrm{d}t \tag{29}$$

is the displacement interpolation between f and g. Intuitively, one is looking for the motion of a gas shaped as f at $t=0$ and g at

© 2023 The Author(s)

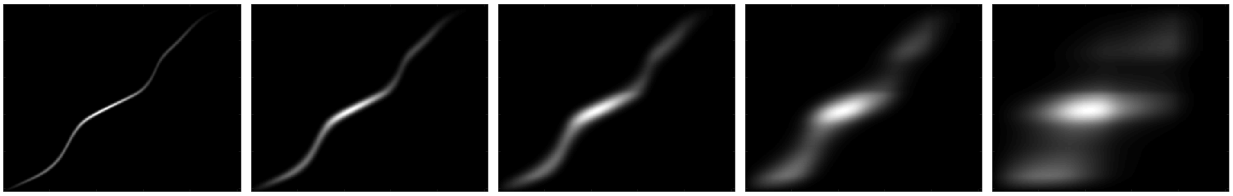

Figure 5: *Solution to the problem of Fig. 1 with increasing amount of entropy. From left to right, the Gaussian standard deviation used in the convolutional approach of Solomon et al. [SdGP*15] is* $\sigma = 2, 5, 10, 20, 50$. *As* σ *increases, the transport plan converges towards the trivial transport plan* $P(x,y) = f(x)g(y)$.

$t = 1$ minimizing its kinetic energy. Rewriting these expressions with the variables $(f_t, f_t v)$, the problem becomes convex and has led to the popular Benamou and Brenier numerical technique [BB00] and subsequent works [PPO14].

In computer graphics, the dynamic formulation has been used for various problems such as computing geodesics on surfaces [SRGB14a], or for other 2-d problems [NG18], using dedicated solvers.

The C++ implementation of the fast 2-d solver of Nader and Guennebaud [NG18] is available at `https://github.com/ggael/otmap`.

3.6. Other solvers

Another recent trend is to solve optimal transport problems using deep learning. It includes learning the optimal transport plan [CDC22], the Wasserstein embedding [CFD18], or Wasserstein barycenters [LDCB22].

Generalizations of the optimal transport problems have also been investigated. Notably, when the two functions being compared do not have the same mass, this leads to the unbalanced and partial optimal transportation problems [CPSV18, BC19, Lév22], with similar solutions (entropy regularized, sliced, semi-discrete or dynamical approaches). Similarly, when the two functions live on two different spaces and one only has access to pairwise distances within each of these spaces that need to be matched, the corresponding problem amounts to the Gromov-Wasserstein problem [Mém11].

Many different solvers are also available within the Python Optimal Transport library "POT" [FCG*21], available at `https://pythonot.github.io/`. Among others, it includes linear program, entropy-regularized, and sliced solvers. In Fig. 6, we summarize the various solvers and problem settings encountered in computer graphics and vision applications, which we detail next.

4. Applications to Image processing

4.1. Image retrieval

Probably the earliest use of optimal transport within the computer graphics and vision literature consists in using the CIE-Lab color distributions of images as a feature for image retrieval [RTG00], and matching it using EMD. Due to low efficiency of optimal transport solvers in the years 2000s (their LP solver took less than a second for $m = n = 100$, and minutes for $m = n = 1000$), image colors were first quantified. On average they used less than 9 color clusters and the clusters spanned less than 25 units in any of the L, a and b axes. This allows to retrieve images of the same color distribution within a database – though they also proposed extensions for querying other features with spatial and textural components. To alleviate the linear search complexity of retrieval algorithms and general speed issue, Indyk and Thaper [IT03] proposed an approximate embedding of the optimal transport cost into $\mathbb{R}^k$ (with relatively large k) equipped with the Manhattan distance in the case of the cost $c(x,y) = \|x-y\|$. This allows for the use of Locality-Sensitive Hashing and thus, sublinear query times. Ling and Okada [LO07] also applied their EMD-ℓ^1 formulation for shape retrieval and shape recognition using histograms of SIFT descriptors [Low04], shape contexts [BMP02] and spin images [JH99].

For image retrieval, via histogram matching, a key feature is that histograms might be non-normalized. Pele and Werman [PW08] proposed to modify the original EMD formulation into $\hat{EMD}$ by adding a term accounting for the difference in total mass, which under mild assumption is a metric even for non-normalized histograms. By further thresholding ground distances above 2, and using their modified $\hat{EMD}$, they turn the LP formulation into a linear-time algorithm, since the maxflow iterations in the simplex algorithm are efficiently reduced through ground distance thresholding. This proved very efficient to match SIFT features [Low04] for which the histogram magnitude conveys critical information. They later extended this method to any threshold value, leading to flow network with drastically reduced number of edges and faster computations [PW09] with applications to SIFT matching, Lab histogram matching, and a combination of both.

4.2. Color grading

Altering colors of images is certainly one of the most widely spread application to illustrate one's optimal transport algorithm. Specifically, the goal is to transform the color distribution of an input photograph or video frame so that the color distribution of the result matches that of an exemplar image (Fig. 7).

To achieve this effect, the color distribution of an input image in a given color space is first stored, discarding any spatial information. The distribution can be stored either as a density on a grid, or directly considering individual pixels as Diracs in that color space, or using color clusters, or even considering the color distribution as a Gaussian distribution. The color distribution of the exemplar image is stored similarly. An optimal transport map between the

© 2023 The Author(s)

Application	LP	Semi-discrete	Entropy-regularized	Sliced	PDE	Others
Image retrieval	[RTG00, LO07, PW08, PW09]					[IT03] (embedding)
Color grading	[MS03, BvdPPH11, RFP14, FSDH14, TPK*17]		[SdGP*15, CPSV18, TPK*17, BPC16]	[PKD05, PKD07, BRPP15, BC19]		[PK07, BSPP13] (gaussian)
Style transfer		[LR21]		[SHS*17]		[KSS19] (no mass preservation)
Image interpolation		[Mér11]	[SA20]		[HZTA04, PPO14, HMP15]	[CFTR17] (embedding, deep learning), [LDCB22] (deep learning)
Optical flow and image matching			[XLY*22, LZYY20]		[KKTH10]	
Image segmentation			[PR17]	[PFR12]		[CEN07, NBCE09] (1d)
Generative Models and texture synthesis		[GLR18, LR21, HLPR22, LSC*19]	[PCVS19]	[RPDB12, TPG16, HVCB21, BRPP15]		[HB95, MZD05] (1d), [HN18, FXPA13, VDKCC20] (Gaussian), [ACB17, GAA*17] (deep learning)
Other image proc. appl.				[TPG16] (denoising, super-resolution)		[GRK*22] (deep learning for super-resolution)
Image stippling and sampling		[XLC*16, dGBOD12, LdGKW19]	[QCHC17]	[SGSS22, PBC*20]		
Reflectance manipulation	[BvdPPH11, GK18, WKB14]		[SdGP*15, BPC16, LBFS21]			
Computational optics and imaging	[AAMT15]	[STTP14, MMT18b]				
Shape comparison and retrieval	[LD09, LD11]			[RPC10]		
Shape interpolation		[Lév15]	[SdGP*15]	[RPDB12]		
Shape registration		[MMT18a]	[SFL*21, MCSK*17, ETLTC20, PRM*21, LLX21]	[BC19]		
Shape reconstruction	[GCSAD11, DCSA*14]					[GCSAD11] (heuristics)
Shape parameterization		[ZSG*13, SWS*15]			[DT09]	
Transport on surfaces	[SV19]		[SdGP*15, PCVS19]		[SRGB14a, LCCS18]	
Topological Data Analysis	[TMMH14, VBT19]		[LCO18]			
Fluid simulation		[GM18, dGWH*15, Lév22, QLDGJ22]	[QLDGJ22]			
Animation	[WCGG18]		[ZSS22, GKK*20, MDZ*21]			

Figure 6: *We summarize the solvers or problem settings used for the applications detailed in this STAR. Note that a solver can be used in different settings – e.g., an entropy-regularized or sliced solver for a semi-discrete OT problem. This classification is thus only indicative.*

input and exemplar distributions is then computed. Finally, the color of each pixel of the input image is matched to its target color using the optimal transport map or plan (in case of a plan, one color can be matched with multiple targets – in that case, a weighted average can be used to obtain a single target [SdGP*15]).

Sliced [PKD05, PKD07, BRPP15] and unbalanced sliced [BC19] formulations are efficient for solving the optimal transport problem in that case, but entropy regularization can be used as well [SdGP*15] or even using a simplex-based solver on color clusters [MS03, BvdPPH11, RFP14, FSDH14]. Even faster is to consider the closed-form expression obtained for Gaussian distributions in the color space [PK07]. To illustrate 2-d optimal transport solvers, the color distribution is sometimes stored in the CIE-Lab color space, and the color transfer problem is separated into matching a 1-d luminance channel and 2-d chrominance channels [SdGP*15, BSPP13].

Directly using the optimal transport result to match colors may lead to quantization artifacts, since similar neighboring colors maybe be stretched to visibly different colors. To alleviate these artefacts, Rabin et al. proposed to regularize the transport plan directly by filtering it using a nonlocal means-like filter [BCM05] adapted to the space of transport maps [RDG10]. Later, Rabin et al. [RFP14] directly add regularization within the linear program by relaxing the mass preservation constraints. Similar regularization for color transfer can be obtained via unbalanced or partial optimal transport [CPSV18, BC19]. Pitié et al. instead propose a variational formulation that enforces the resulting image's gradient to be regularized towards the original image gradients [PKD07] in order to reduce grain. Solving an optimal transport between the graphs of two functions leads to the transportation L^p-distance, that can be used to match colors accounting for spatial correlations [TPK*17] and can more generally be used to compare general functions (that can be negative, and do not have the same mass).

Note that these methods often produce convincing results (see Fig. 7 where we compared standard optimal transport based approaches to direct Adobe Photoshop color matching), but quantitative perceptual evaluation of the color transformed images is challenging, and in practice, to our knowledge, not done. In this context, it is hard to assess which technique in general works best.

Instead of matching one input image color distribution to one exemplar, another goal can be to harmonize the colors of a set of input images, e.g., to make them all look as if they had been shot with the same camera (or with the same camera settings). In that case, it can be useful to transfer the color of each input image to the Wasserstein barycenter of the set of images [BRPP15]. A single image can also be matched to multiple exemplars. In that case, its color distribution can be projected to the nearest Wasserstein barycenter of the exemplar images [BPC16], which is formulated as an optimization problem, or Wasserstein barycenters can be used to obtain

© 2023 The Author(s)

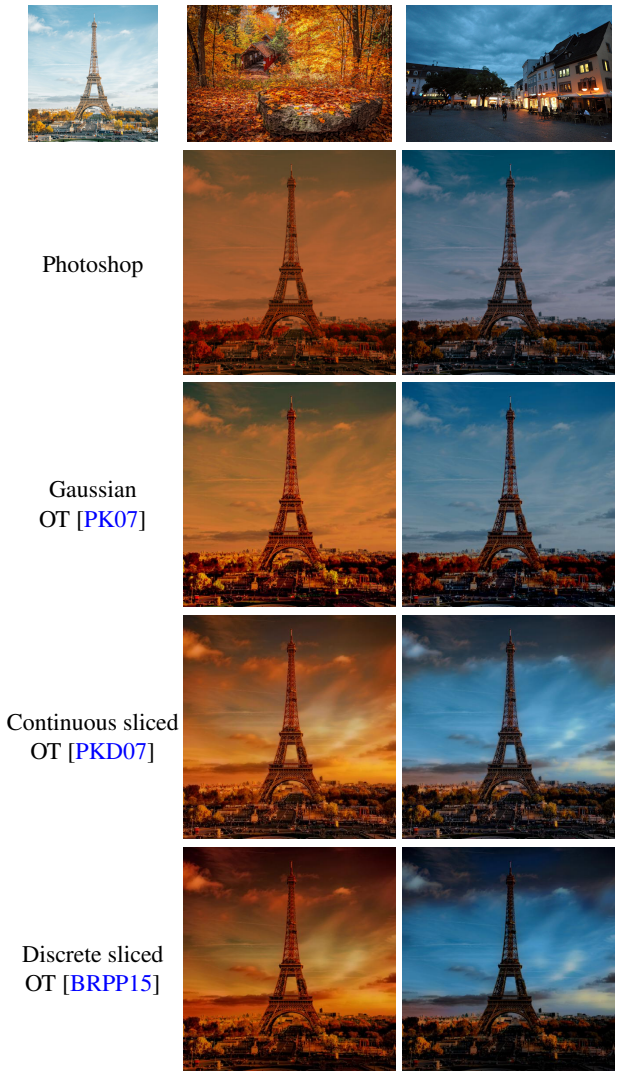

Figure 7: ***From top to bottom***. *Example-based color grading using Photoshop "Match Color", optimal transport-based color grading using a Gaussian approximation [PK07], sliced optimal transport using a continuous cdf-based transfer plus TV-based reconstruction to reduce grain [PKD07] and sliced optimal transport using a discrete 1-d matching [BRPP15] plus filtering to reduce grain [RDG10]*

interpolations between different color styles [BRPP15]. Finally, the problem of the temporal consistency of per-frame color transfer in the context of video color grading has been studied by exploiting the differential geometry of Gaussian optimal transport [BSPP13].

4.3. Style Transfer

Optimal transport is not only useful for transferring color histograms, it can transport all sorts of distributions. As such, it lends itself well to the transport of illumination between image and examples. Shu et al. [SHS*17] tackle this problem for portrait relighting by transporting histograms of normals, positions, colors (hence histograms of 8-dimensional vectors, normals being identified as 2-d quantities) between 3-d templates fitted to 2-d portraits. Interestingly, instead of using an LP formulation, they returned to the original Monge formulation, in order to avoid discontinuities in uniform color regions. This high dimensional transport problem is solved efficiently by using the sliced approximation.

In 2016, Gatys et al. [GEB16] defined the style of an image as the Gram matrix of its CNN features, and used a gradient descent to transfer style from an image to another. The core idea to match the distribution of features of two images can be seen as an optimal transport problem, and many subsequent methods have explored this path. Kolkin et al. [KSS19] proposed to use relaxed optimal transport for their style loss. This relaxation called Word Mover Distance [KSKW15] drops the target mass conservation constraint. Hence, an optimal solution moves the total mass of each point to its nearest point in the target distribution. By inverting source/target roles in this computation, and taking the maximum of the two approximations one gets the final relaxed EMD. The result is more akin to a closest-point assignment than an optimal transportation, but it has been presented and used as an optimal transport surrogate. The relighting problem is finally solved by a gradient descent minimizing this relaxed EMD combined with a content loss and a moment matching loss leading to an efficient content-preserving style transfer.

Leclaire and Rabin [LR21] addressed style transfer using a stochastic approximation of semi-discrete optimal transport in patch space. For the continuous measure, they consider the Gaussian random field synthesized patches distribution, and for the discrete measure they consider the discrete empirical distribution of the sample patches. It works by computing multiple restricted optimal transport problems between the source measure and a hierarchical simplification of the target measure, given by multi-resolution clustering. A blending step mixes the style feature with geometric features of the source image.

Style transfer is intrinsically linked with texture synthesis. This is exemplified by the seminal works of Gatys et al. [GEB15, GEB16], who addressed both in a similar gradient descent on extracted features and Gram matrix of the features. Many methods tackling one of the problems can be adapted to the other [HVCB21, LR21], and the reader should refer to Section 4.7 for more details.

4.4. Image interpolation

Optimal transport has also been considered as a tool to morph images, notably with images of low structures such as fluids

© 2023 The Author(s)

in which classical feature detection+matching algorithms might fail [HZTA04]. This however supposes a notion of "mass" to be matched with pixel color or intensity, and that this mass should somewhat be preserved along the morphing, which is not intuitive in the context of image morphing. Optimal transport may also split mass (e.g., one eye on a face could be divided into two eyes during the interpolation!), but solutions to enforce rigidity have been proposed [HMP15]. To interpolate between sketches Courty et al. proposed to learn a Wasserstein embedding of sketches, allowing then to interpolate in the embedding space [CFTR17], however the results remain blurry. For the same problem, Lacombe et al. proposed to learn directly the barycenter prediction without building an embedding [LDCB22]. Acknowledging the non-natural artefacts caused by Wasserstein-based interpolation, Simon and Aberdam [SA20] proposed to constrain the barycenters to follow an image prior given by a GAN or a sparse prior. In practice this is done through an ADMM optimization, by computing the Wasserstein barycenter and projecting it on the image prior manifold. The resulting interpolation avoids unnatural artefacts (mass splitting...) but is mostly demonstrated on background-less pictures of objects.

Nowadays, many other tools – notably based on deep learning – achieve much better results, and optimal-transport based image interpolation is mostly used as illustrative of optimal transport algorithms performance or behavior [Mér11, PPO14].

4.5. Optical flow and image matching

Similarly to image interpolation, optimal transport has been used to estimate optical flow and match images. Within a differentiable renderer, Xing et al. [XLY*22] uses a differentiable Sinkhorn divergence to match pixels between the current rendering estimate and the reference rendering, where each pixel has a unit mass and the ground distance accounts for color and positional distances. Liu et al. [LZYY20] uses the Sinkhorn algorithm with a cosine similarity ground distance between feature maps to establish semantic correspondences between images. In a more specific setting, Kolesov et al. [KKTH10] computes an optical flow of image sequences of fire and smoke by considering that fire and smoke approximately conserve their pixel intensity, using a dynamic formulation.

4.6. Image segmentation

Optimal transport has also been applied to the image segmentation problem. In the regional active contour framework Chan et al. [CEN07, NBCE09] proposed to combine shape features with an optimal transport approach to model regional statistics. It strives to find an optimal image segmentation into 2 regions Σ, Σ^c, such that the local histograms inside Σ (resp. Σ^c) can be well approximated by a single histogram representing Σ (resp. Σ^c). This approximation is captured by a W_1 distance. The segmentation functional also includes a more traditional term aiming at reducing the region perimeter. Following this trend, Peyré et al. [PFR12] proposed Wasserstein Active Contours, based on similar ideas but with an explicit derivation of the active contour velocity. In addition, to make the computation tractable they relied on sliced Wasserstein distances. Papadakis [PR17] proposed a segmentation energy based on the minimization of the Wasserstein-p optimal transport distance between the histograms of the 2 segmented regions and two example histograms. While the segmentation results are not on par with other image segmentation techniques (even old graph cuts based methods), the mathematical derivation for these algorithms is very interesting and provides new insight on optimal transport driven functionals.

4.7. Generative Models and texture synthesis

Texture synthesis is a widely spread application of optimal transport approaches, for both synthesizing new textures from one exemplar image (either patch-based or using statistical models), or for interpolating between different texture models. A related application is the synthesis of natural, non-stochastic, images.

Synthesizing textures. Multiple approaches exist to synthesize new textures from an exemplar image. This includes pixel-based, patch-based, or stochastic models.

The widely used Heeger and Bergen texture synthesis algorithm [HB95] decomposes an exemplar image into steerable pyramid coefficients, and then, from a Gaussian white noise image, iteratively matches each sub-band of its own steerable pyramid decomposition to the exemplar's one using cdf-based 1-d optimal transport. Tartavel et al. [TPG16] instead replace this matching with an LBFGS-based optimization of a sliced optimal transport functional.

Optimal transport is also instrumental for spot noise-based texture synthesis. Discrete spot noise is defined as the addition of n random translations of an input image – as n grows this can serve to synthesize textures. It can also be seen as convolving a random point set with the input image. This process yields a Gaussian Random Field, a $\mathbb{R}^2$ field such that any vector of field values anywhere in $\mathbb{R}^2$ are distributed as a multivariate Gaussian. Galerne et al. [GLR18] proposed to use optimal transport to constrain a spot-noise texture synthesis method to match an exemplar texture. It starts by sampling a spot noise texture, and then solves a semi-discrete optimal transport problem between the set of possible 3×3 generated patches (as a continuous distribution) and the set of exemplar patches (as a discrete distribution). Each spot noise patch can then be transformed to match a texture patch and the final image is obtained by averaging. In a way it is mixing parametric models given by spot noise, generating a texture as a stationary Gaussian Random Field and nonparametric methods, such as patch-based texture synthesis. Leclaire and Rabin's fast semi-discrete optimal transport approximation using hierarchical decomposition of the target measure allowed to improve this method numerically, allowing to use larger image patches [LR21] (see Section 4.3).

Another way to include optimal transport into texture synthesis is to use it as a loss for deep learning models. Heitz et al [HVCB21] proposed to replace the features Gram matrix loss in texture synthesis task [GEB15] with a sliced Wasserstein loss between features (Fig. 8). Houdard et al. [HLPR22] later leveraged the semi-discrete idea in a deep learning synthesis framework. Their method uses the semi-discrete optimal transport map between the texture features continuous distribution synthesized by a generator and the target texture features discrete distribution as a loss to learn to synthesize a stationary texture from a single example.

Optimal transport was also used as a preprocessing for texture

© 2023 The Author(s)

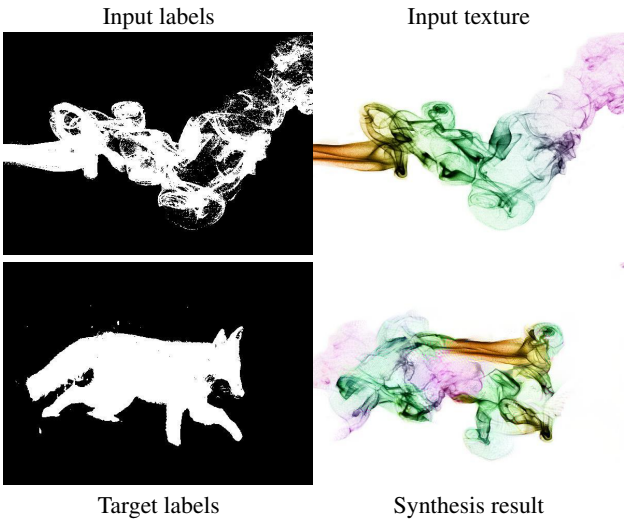

Figure 8: *Replacing the Gram matrix distance by a sliced Wasserstein loss in the approach of Gatys et al. [GEB15] produces better texture synthesis and transfer results ; here, a texture synthesis guided by a label mask [HVCB21].*

patches. Heitz and Neyret perform patch-based texture synthesis by blending patches taken from an exemplar image [HN18]. The color distribution within each patch is considered Gaussian, and a Gaussian distribution blending formula is used (without resorting to optimal transport). However, to conform to this Gaussian distribution hypothesis, in a precomputation step, patch color distributions are *made* Gaussian by computing an optimal transport map between the actual color distribution and a sampled Gaussian distribution, using a linear program solver.

Interpolating between textures. Matusik et al. [MZD05] interpolate between multiple textures by first computing a (non-transport based) warping between pairs of textures, warping them halfway, and blending them using 1-d optimal transport (Fig. 9). More precisely, after alignment using the precomputed warp, textures are blended by decomposing them into steerable pyramid subbands [HB95], and adjusting the bands' histograms by using the 1-d cdf-based formula for Wasserstein barycenters. Rabin et al. [RPDB12] proposed to interpolate between textures by computing sliced Wasserstein barycenters between the distributions of the texture statistical textures. This method can be seen as a way to generalize [HB95]: instead of considering only 1d distribution it considers higher order distributions. Ferradans et al. [FXPA13] proposed an optimal transport-based texture interpolation method in the spot noise model. Static (resp. dynamic) textures are represented as 2-d (resp. 3-d) stationary Gaussian processes. Optimal transport is computed between these Gaussian random fields, and texture mixing is achieved by sampling along the geodesics. To synthesize textures that interpolate between more than two input textures, Bonneel et al. use the Radon Wasserstein barycenter of input texture power spectra, and apply a random phase [BRPP15]. In a similar setting, Peyré et al. [PCVS19] better account for correlations between colors by interpolating 3x3 Gaussian covariance matrices instead of scalar values (see Sec. 6.6).

By considering the mean and covariance of a texture CNN features, and using closed-form formulas for the Wasserstein-2 distance between the estimated Gaussians, Vacher et al. [VDKCC20] interpolate textures along geodesics defined by the Wasserstein metric. The results are reported to be better in terms of visual perception.

Synthesizing natural images. Optimal transport was also used in the more general setting of natural image synthesis, in the context of Generative Adversarial Networks (GAN). Arjovsky et al. [ACB17] introduced Wasserstein GAN, which replaces the Jensen-Shannon divergence with a Wasserstein-1 distance. Given a generator G, a discriminator (called critic in the WGAN setting) D, μ_G the generated distribution, and μ_{ref} the reference distribution one wants to learn, the WGAN loss writes:

$$\mathcal{L}_{WGAN}(G,D) = \mathbb{E}_{x\sim\mu_G}[D(x)] - \mathbb{E}_{x\sim\mu_{ref}}[D(x)]. \tag{30}$$

D tries to maximize this loss which G tries to minimize in turn. By using the Kantorovich-Rubinstein dual formulation, one gets that for a given generator G, the optimal D can be constrained to produce a measurable Lispchitz function using weight clipping so that $\mathcal{L}_{WGAN} \propto W_1(\mu_G, \mu_{ref})$. Then G in turn tries to minimize this Wasserstein-1 Distance, leading to more realistic samples. Wasserstein GAN perform better than regular GANs, avoiding vanishing gradients and mode collapse problems, but they are still unstable and slow to converge. Later improvements have been proposed to stabilize the training, replacing weight clipping by a discriminator gradient penalty [GAA*17]. Wasserstein GANs are a broad image generation technique, which can be used in a wide variety of applications. Lei et al. [LSC*19] revisit the Wasserstein GAN architecture with a semi-discrete optimal transport perspective. They consider that G optimizes a transport map from the latent space to the true data manifold, while D estimates the Kantorovich potential permitting to compute the Wasserstein-1 distance between the synthesized continuous distribution and the discrete target distribution. Using this geometric view, they show that if the cost $c(x,y) = h(x-y)$ is a strictly convex function, the generator can be written in closed form from the optimal D. While this is shown mainly theoretically, some promising experiments are also shown. GAN (and WGAN) are now superseded by Diffusion Model approaches [HJA20], for which Optimal Transportation can also provide interesting insights. In that context, De Bortoli [DB22] provided bounds on the Wasserstein-1 distance between the distribution synthesized by a diffusion model and the true distribution. More generally, since the goal of any generative model is to learn a distribution of samples (images, shapes, or data in general), then Optimal Transportation is a good way of measuring the performance of the model, by computing the Wasserstein-p distance between the generated distribution and the target distribution.

4.8. Other image processing applications

Optimal transport has been also applied to other image processing tasks, more marginally. Tartavel et al. use a sliced Wasserstein term in a Total Variation (TV) based image restoration framework, and applies it for denoising and super-resolution tasks [TPG16], with the

© 2023 The Author(s)

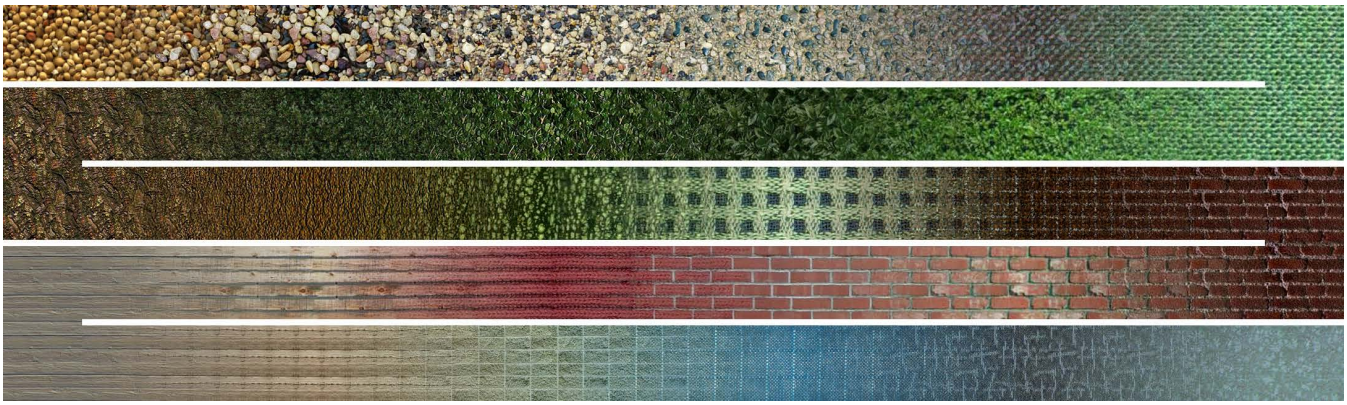

Figure 9: *A path in texture space interpolating between 26 different textures [MZD05]. After spatial alignment, textures are blended after matching steerable pyramid sub-bands 1-d histograms [HB95] to their 1-d Wasserstein barycenters.*

goal to recover details lost by the TV term by matching a distribution of features to a prior distribution. Gazdieva et al. perform super-resolution by learning an approximate optimal transport map from a set of low resolution images to a set of (unpaired) high resolution images [GRK*22].

5. Applications in rendering

5.1. Image stippling and sampling

Semi-discrete optimal transport provides a way to measure the distance between a set of Diracs and a density. As such, it easily found its place for problems related to the approximation of a density with a set of Diracs, that is, the energy optimization problem consisting in minimizing the optimal transport distance between a density and a sum of Diracs. This approach finds applications for two related problems. First, that of stippling an image, for instance in the context of artistic image stylization or for dithering an image in the context of printing grayscale (or color) images with black (or CMYK) ink droplets. Second, that of sampling a density (most often the uniform distribution on a unit hypercube domain) for Monte Carlo integration, and in particular for physically-based rendering. While the first application mainly looks for visually pleasing point distributions, the second application is more interested in reducing the error or variance in integral estimators. Optimal transport offers benefits for both, as it produces interesting blue noise properties [dGBOD12, PSC*15].

The general idea behind these approaches is to initialize with a random point set, and to iterate the computation of a semi-discrete optimal transport computation step between the current point set and a density (uniform or not), and a centering step that moves each sample towards or at the location of the centroid of its power cell, in a spirit similar to Lloyd's algorithm [Llo82].

Stippling an image has been proposed as an application for several semi-discrete approaches [XLC*16, dGBOD12] and state-of-the-art results were obtained by the "BNOT" approach [dGBOD12] (Fig. 10, left, only recently outperformed by a non optimal transport based approach [ARW22]). A different, much faster but approximate, solver on 2-d grids has been proposed [NG18] and allows to generate optimal transport stippling patterns visually similar to that of de Goes et al. [dGBOD12]. Qin et al. instead uses entropy-regularized optimal transport for sampling multiple classes (e.g., for color stippling), but also to sample meshes [QCHC17]. Their algorithm should be general but is demonstrated in two dimensions (or on surfaces embedded in 3-d). A sliced approach also allows for multi-class color stippling [SGSS22]. An interesting generalization relates to the optimal transport approximation of a density by other non punctual measures [MM99], such as a single (long) continuous curve [LdGKW19] (Fig. 10, right).

Sampling for rendering involves different constraints. While image stippling usually restricts the problem to the approximation of non-uniform 2-d densities, sampling for rendering generally requires higher-dimensional densities, albeit generally uniform (they are usually non-linearly transformed to match the desired distribution for importance sampling afterwards). This adds difficulties as the semi-discrete optimal transport problem is much easier in the low-dimensional setting. Even in the case of 2-d or 3-d integration problems, semi-discrete optimal transport sampling has shown benefits in term of integration error because the resulting point distribution exhibits blue noise properties [PSC*15]. In the higher dimensional setting (up to 20-d), a sliced approach has notably been proposed for physically-based rendering applications [PBC*20] by enforcing blue noise properties both in the high-dimensional space and in its lower dimensional projections. This approach also works for the task of multi-class blue noise sampling. Sliced multi-class sampling has also been proposed by Salaün et al. [SGSS22] with a custom energy that tends to produce blue noise spatial error distribution in the image plane, by producing one point set per pixel such that neighboring pixel samples are well interleaved in a blue noise fashion.

5.2. Reflectance manipulation

Displacement interpolation can be used to obtain intermediate Bidirectional Reflectance Distribution Functions (BRDFs) between two input BRDFs. While for parameterized BRDF models, it is often more intuitive to directly interpolate BRDF parameters, this is not the case for BRDFs that are captured by gonioreflectometers and

© 2023 The Author(s)

Figure 10: *Semi-discrete optimal transport allows to solve the problem of approximating a density with discrete distributions. This finds applications in image stippling when the target distribution is a set of Diracs [dGBOD12], or stylization when the target distribution is a continuous curve [LdGKW19]*

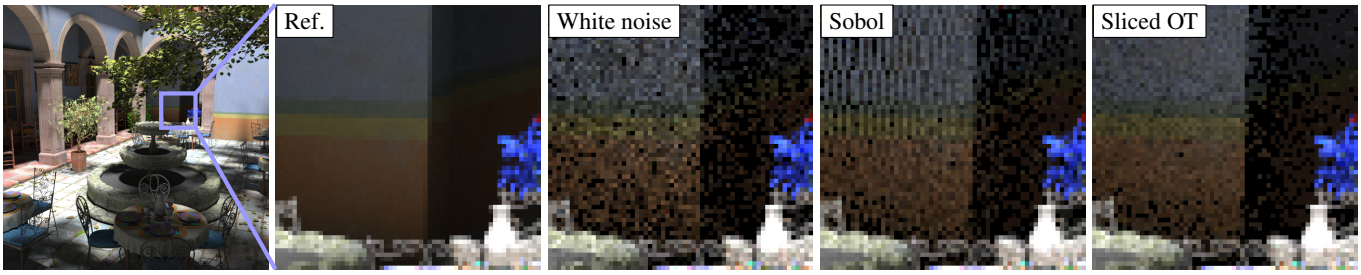

Figure 11: *Sliced Optimal Transport Sampling reduces variance in Monte Carlo physically-based rendering applications by optimizing sample location in high dimension to uniformly cover the unit hypercube [PBC*20].*

thus only accessible by their measurements at discrete locations. To achieve that, Bonneel et al. [BvdPPH11] interpolate each 2-d slice of two 4-d measured BRDF, per color channel (or wavelength), by advecting gaussian particles along a path on the hemisphere obtained by solving the optimal transport linear program. Solomon et al. [SdGP*15] adopt a similar slicing strategy, but uses an entropy regularized solver to produce Wasserstein barycenters, using fast approximate spherical Gaussian convolutions. For car paint interpolation, Golla and Klein [GK18] use both a parametric model, and displacement interpolation of color clusters in the CIE-Lab color space for metallic flakes using a linear program solver. Optimal transport can be used to densify sparse BRDF measurements, by triangulating the set of incoming directions, and interpolating the corresponding 2d slices between different incoming directions [WKB14], or by projecting a sparse BRDF onto a database of denser BRDFs using optimal transport [BPC16]. The Sinkhorn divergence between BRDFs is also shown to correlate well with perceptual dissimilarity measurements [LBFS21].

5.3. Computational optics and imaging

Optimal transport has a relatively long history in solving non-imaging optic problems [GO03]. Specifically, one problem is to redirect some incoming light distribution via a mirror or lens so as to produce a specific reflected or refracted light distribution, e.g., an image. Intuitively, mass conservation translates into light energy preservation, and a transport plan would dictate the shape of the mirror or lens, which makes optimal transport an appropriate tool for solving this kind of problem.

In the work of Schwartzburg et al. [STTP14] a lens is designed by first computing a set of normals using semi-discrete optimal transport. Here, a flat light emitting surface is discretized into centroidal Voronoi cells, and the semi-discrete optimal transport plan between this discretization and the target light distribution on the flat receiver surface produces one-to-one correspondences between the Voronoi cells of the emitter and the power diagram cells produced by the semi-discrete optimal transport map on the receiver. Using Snell's law, the distribution of normals over the lens surface is then computed, and an optimization reconstructs the corresponding surface (Fig. 12, left).

In the far-field reflector design problem of Andre et al. [AAMT15], a light distribution at infinity (on the sphere of directions, as opposed to a near-field receiver surface) is prescribed and should be produced by a point light reflecting off a mirror whose shape needs to be found. The reflector is designed as a set of intersecting paraboloids whose focal point correspond to the point light location and the focal distances need to be determined, and the paraboloids axis cover the set of desired directions. They solve a linear program formulation of optimal transport with cost

© 2023 The Author(s)

$c(x,y) = -\log(1 - x.y)$ for an incoming light direction x and a paraboloid of axis y, between a discretization of the point light directional intensity distribution and a discretization of the desired light distribution on the sphere. The resulting transport plan indicates how much light from a incoming direction reaches any outgoing direction after being reflected. It also turns out that the solution to the dual linear program directly gives the desired focal distances, and thus solves the problem.

Meyron et al. [MMT18b] generalize the non-imaging optic problem (Fig. 12, right). They show that all the problems of obtaining a target near-field or far-field image from a given collimated or punctual light source whose light is refracted in a lens or reflected over a mirror that can be concave or convex, can all be solved via a semi-discrete optimal transport problem with quadratic cost, up to a change of variable in the power diagram weights. They solve the semi-discrete optimal transport using a damped Newton scheme, with automatic differentiation.

6. Applications in Geometry Processing

6.1. Shape Comparison and retrieval

Mesh surface comparison can be performed using optimal transport [LD09,LD11], by mapping simply connected Riemann surfaces to unit disks and transporting the domains onto each other using a ground distance invariant to Möbius coordinate changes. By discretizing the two domains, the optimal transport problem can be solved using a linear program solver. This method is limited to simply connected surfaces and inherits linear program solver limits, i.e., the sampling of the unit disks should remain relatively small for the transport to be computed.

Similarly to image retrieval, 2-d and 3-d shape retrieval can also benefit from optimal transport tools, by computing shape signatures and matching these signatures using optimal transport. Rabin et al. [RPC10] compute a shape signature by sampling points on the surface and computing quantiles of geodesic distances distributions. They then match these descriptors using the sliced Wasserstein-1 distance.

6.2. Shape Interpolation

Optimal transport based image interpolation is only partially satisfying because features can get split in a non semantic way, while other methods would match image features and yield much more natural interpolations. For 2-d or 3-d shapes, the problem is however slightly different. Take for instance two spheres or two icosahedra. Matching features would be impossible (a sphere has no meaningful feature) or error-prone (an icosahedron has repeated features). By considering the shapes as boundaries of indicator functions, and normalizing the mass, one can match the shapes as 2-d or 3-d densities using optimal transport, and perform displacement interpolation.

Rabin et al. [RPDB12] already proposed interpolating between shapes using sliced Wasserstein barycenters. The volume of the shape is first sampled uniformly and the sliced barycenter is computed between these discrete distributions. Lévy [Lév15] proposed to use semi-discrete optimal transport for interpolating between two tetrahedral meshes M and M', given an additional number of desired vertices k. The algorithm is strikingly simple: one starts by sampling M' with a set $Y = (y_i)_{i=1\dots k}$ of points. The semi-discrete transport between Y and M is then computed as a power diagram restricted to M, from which one can extract the dual regular triangulation. A correspondence between the triangles of this triangulation and the ones of the restricted Delaunay of Y on M' can be established since they are related to the same y_i, and the common topology can be found by looking for the triangles that exist in both (i.e., the corresponding cells do not degenerate). The vertices positions are linearly interpolated between the barycenters of the power cells on M and the y_i on M'. This provides an efficient way to interpolate between meshes without resorting to feature point computation and matching (Fig. 13).

Instead of sampling the shape volume, Solomon et al. [SdGP*15] considered shapes a indicator functions discretized on 2-d or 3-d grids and, proposed to compute interpolation between several voxelized shapes using fast convolutional Wasserstein distances.

Although problems are less obvious than for image interpolation, optimal transport-based shape interpolation still faces issues: the interpolated shapes might split and have more connected components than what would be expected, which is due to optimal transport being topology agnostic.

6.3. Shape Registration

The optimal transportation problem consists in matching data. As such, a natural application is to look for reliable and consistent pointwise correspondences between shapes.

Point Cloud to Point Cloud Discrete formulation of optimal transport naturally leads to solutions of point clouds (i.e., sets of Dirac) processing problems. In the work of Bonneel and Coeurjolly [BC19], a gradient flow of a sliced partial optimal transport energy allows to advect a set of points in 3-d towards a larger set of points. They introduce a 1-d partial optimal transport algorithm and use it in a sliced context to propose a variant of the Iterative Closest Point (ICP) point registration algorithm [BM92]. In classical ICP, a rigid transformation between the input point cloud and the nearest neighbor of each input point in the target point cloud is estimated – a possible issue is that many input points can be matched to the same nearest neighbor. Using partial optimal transport, injectivity is guaranteed and this issue cannot happen. In the proposed approach, the input point cloud is advected towards the target point cloud via the sliced partial optimal transport gradient flow, and the rigid (or similarity) transform is computed between the input point cloud and the advected point cloud (Fig. 14).

Shen et al. 2021 [SFL*21] add an entropy-regularized unbalanced optimal transport [CPSV18] layer to a deep learning architecture as a drop-in replacement of closest-point matching, along with confidence weights. Specifically, an optimal transport plan leads to a one-to-many point matching that is not suitable for registration problems, but one-to-one matching can be approximately recovered by a weighted average of the many matches of soft transport plans [SdGP*15]. They apply this to rigid, affine or deformable transformation estimation to register point sets, and use a deep learning architecture to directly predict model parameters.

© 2023 The Author(s)

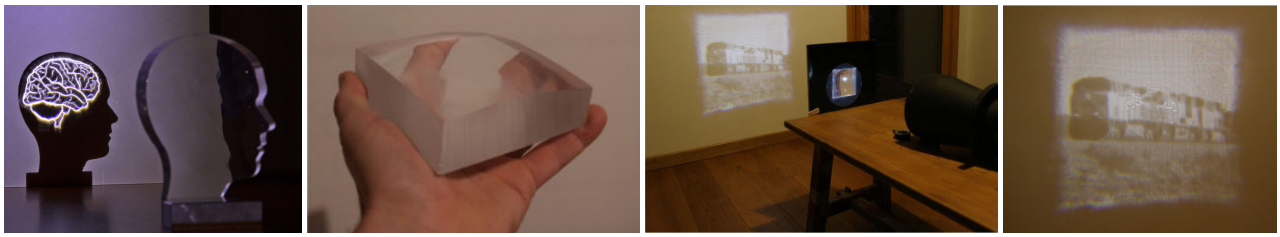

Figure 12: ***Left.*** *A brain caustic obtained from a light panel and optimized lens by Schwartzburg et al. [STTP14].* ***Right three images.*** *A lens producing an image of a train from a collimated light by Meyron et al. [MMT18b].*

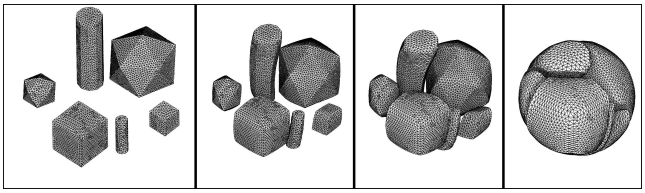

Figure 13: *3-d mesh interpolation between 6 volume shapes and a ball, performed by Lévy [Lév15] using semi-discrete optimal transport. Optimal transport allows to interpolate between shapes of different topology, and for which feature points would be hard to compute.*

Optimal transport has also been used for scene flow estimation. Li et al. [LLX21] uses the Sinkhorn algorithm to match consecutive point clouds in time using a ground distance accounting for differences in point position, colors, and estimated normals.

Point Cloud to Mesh. Mérigot et al. [MMT18a] introduced semi-discrete optimal transport to compute distances between a point cloud and a mesh. As done traditionally [AHA98], the problem is cast into a power diagram weight optimization, by considering the intersection of the power diagram with the mesh. A damped Newton algorithm is proposed to for finding the power diagram weights, with proven convergence in linear time under mild conditions. They propose two applications to this algorithm, remeshing and Iterative Closest Point. For remeshing, they use the continuous measure on the mesh as a target triangle density measure and the mesh vertices as the point set, which produces a remeshing as the dual of the computed Laguerre diagram. Optimal transport-based Iterative Closest Point is another interesting application of this algorithm, since it allows to drop the costly assignment step and replace it by a semi-discrete optimal transport computation between a constant measure on the mesh and a constant discrete measure on the point set. The alorithm then associates each point of the point set with the barycenter of the corresponding power cell, yielding faster empirical convergence. Dropping the nearest neighbor search also permits to better handle far-away initial poses.

Mesh to Mesh To register a mesh onto another mesh, Mandad et al. [MCSK*17] relied on Sinkhorn iterations to compute a transport map and deducing a variance-minimizing correspondence. To do so, they rely on a surface discretization and aim at minimizing the variance of transported neighborhoods, by alternating between relocating samples to local barycenters and solving an optimal transport problem. Optimal transport is computed via Sinkhorn iterations, relaxed to allow for small changes of mass, and a coarse-to-fine strategy. Eisenberger [ETLTC20] introduced Deep Shells extending the Smooth Shells [ELC20] correspondence framework. It replaces the non-differentiable costly correspondence search step with an entropy-regularized optimal transport computation. This change makes all steps of the smooth shells approach differentiable, and thus trainable in an end-to-end manner. This approach interestingly mixes optimal transport with the functional map approach. The development of entropy-regularized tools for optimal transportation led to other breakthrough for functional maps shape registration via Sinkhorn filters [PRM*21]. While this is not directly an optimal transport approach, it benefited from advances in this field.

6.4. Shape Reconstruction

The shape reconstruction problem can be seen in a similar light as the point cloud to mesh registration problem.

De Goes et al. [GCSAD11] proposed to compute the distance between a 2-d mesh and a (noisy) 2-d pointset where points are located on curves. In this setting, optimal transport is computed by first assigning each point to the nearest mesh edge and distributing its mass regularly on the edge. The corresponding cost can be computed in closed form. This principle is applied to curve reconstruction: starting from an initial Delaunay triangulation of the point set, vertices are iteratively removed by choosing the point whose removal increases the optimal transport distance least, and the triangulation is adapted accordingly. Final edges are filtered out by removing edges without any mass, and vertex positions are optimized to minimize the transport cost (with frozen transport plan). Digne et al. [DCSA*14] extended the problem to the 3-d setting. They discretize a constant measure on regularly sampled bins on the simplices, and formulate the optimal transport problem as a linear program between points and simplices, with additional constraints enforcing the piecewise constant measure condition. To make the problem tractable, local problems are considered iteratively. Applications to mesh reconstruction, similary to de Goes et al. [GCSAD11], and sharp edge recovery (e.g., after a Marching Cube-based reconstruction) are proposed by alternating between the relocation of vertices to minimize transport cost with frozen transport plan, and transport plan update. Arguably the setting of both methods is not a "true" optimal transport since the continuous

© 2023 The Author(s)

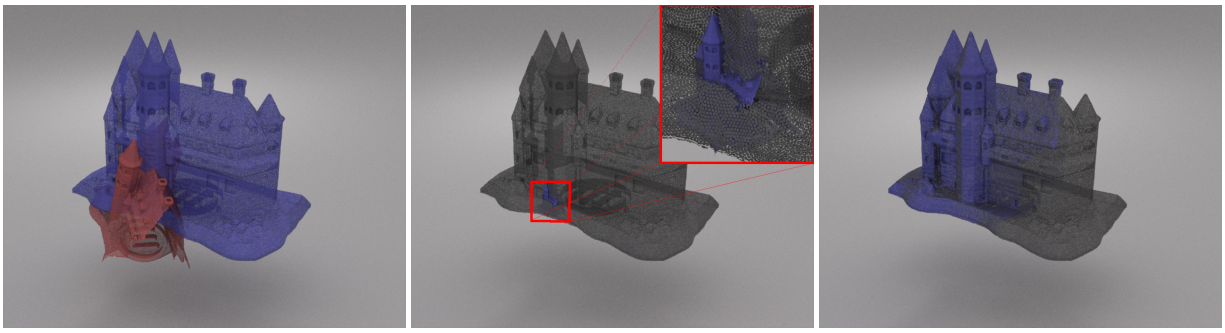

Figure 14: *A sliced partial optimal transport gradient flow can be used instead of nearest neighbor queries in the ICP algorithm to register point sets.* ***Left.*** *Initial configuration.* ***Middle.*** *Classical ICP to register the point sets using rotation, translation and scaling. Since many points are matched to the same nearest neighbor, this leads to a degenerate configuration in the scaling factor.* ***Right.*** *The sliced optimal transport solution enforces injectivity and leads to a correct registration.*

measure on the simplicial complex has to be optimized as well: we only know its total mass and that it is piecewise constant on the simplices.

6.5. Shape Parameterization

Shape parameterization corresponds to the task of unfolding a mesh on a planar or spherical domain, and has also attracted optimal transport applications. Parameterization methods usually target some quantity (areas, angles, distances) preservation. The most common mapping is conformal mapping which preserves angles but can exhibit large area distortions (Fig. 15). To reduce this distortion, Dominitz and Tannenbaum [DT09] first compute a conformal parameterization of a mesh onto a sphere, and then compute the optimal transport map between the conformal factor and a uniform density over the sphere using a PDE-based gradient flow formulation and multiresolution solver using Discrete Exterior Calculus (DEC). The UV coordinates are then advected along this flow. This allows to produce area-preserving mapping of low angle distortion. Zhao et al. [ZSG*13] use a similar intuition but rely on a semi-discrete optimal transport solver to transport the conformal factor of the embedding onto a disk to the uniform density over that disk. UV coordinates are then moved towards the centroid of their power cells.

Instead of transporting the conformal factor, Su et al. [SWS*15] directly transports the area of neighboring triangles in a conformal map parameterization onto the disk or sphere towards the uniform measure using semi-discrete optimal transport. They again move UVs coordinates towards the centroid of their corresponding power cells to obtain an area-preserving map. This further serves to compute a Wasserstein-2 distance between surfaces by comparing their mapped domains and then be used for mesh retrieval [SWS*15].

6.6. Transport on Surfaces

Instead of transporting mass on Euclidean domains, it is also possible to transport densities defined on Riemann surfaces. Solomon et al. [SRGB14a] considered the problem of computing the optimal transport between densities on a manifold surface M represented as

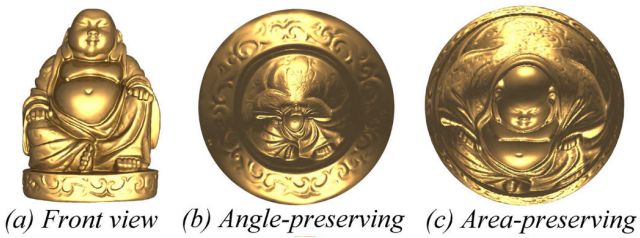

Figure 15: *Semi-discrete optimal transport allows to reduce area distortion of conformal parameterizations [ZSG*13]. These parameterizations can be used for shape matching [SWS*15].*

a mesh and linked it with geodesics computation. It uses a variational formulation casting the Wasserstein-1 distance between two densities defined on M as solving a flow with constraints on the boundaries of M and divergence of the flow. More precisely, the method looks for a tangential vector field V on M:

$$\begin{aligned} W_1(\mu_0,\mu_1) = \inf_J \int_M \|V(x)\| dx \\ s.t. \nabla \cdot V(x) = \mu_1(x) - \mu_0(x) \\ V(x) \cdot n(x) = 0 \quad \forall x \in \partial M \end{aligned} \tag{31}$$

where $n(x)$ is the normal to the boundary ∂M at point $x \in \partial M$.

The flow lines of J are known to be geodesics on M. The vector field can be easily decomposed using the Helmoltz-Hodge decomposition, and projected on a spectral basis which allows for an efficient Finite Elements discretization. This formulation is both fast and accurate, and yields a family of geodesic distances, depending on the size of the spectral basis, ranging from a purely spectral distance to the "canonical" geodesic distance. However this method cannot be applied for quadratic costs, and does not support smooth interpolation.

An alternative to using a dynamical optimal transport formulation is to rely on entropic regularization, and use a heat kernel instead of a distance [SdGP*15], since both are related by Varadhan's for-

© 2023 The Author(s)

mula [Var67]. However the resulting interpolations remain blurry, as expected when using entropy regularization.

Lavenant et al. [LCCS18] proposed a dynamic formulation for quadratic costs. It is based on a similar discretization of the Benamou-Brenier definition on a manifold, but it is strictly convex and more amenable to interpolation computation than the Wasserstein 1 method [SRGB14a]. It is also less blurry than entropy-regularized variants or convolutional Wasserstein distances [SdGP*15].

Solomon and Vaxman [SV19] proposed to interpolate between two tangent vector fields defined on a mesh. It uses optimal transport as a way of matching the vector field singularity in an interesting way. Singularities have – possibly negative – indices (i.e. classification) that sum to the same value for two smooth tangent vector fields on the surface. Therefore these singularity indices can be considered as mass to be transported. A variant of discrete optimal transport is presented, allowing for negative mass, and it is solved using a linear program solver, which is completely tractable, since the number of singularities is low. The rest of the vector field is constructed using other dedicated optimization schemes. Going higher in dimension, Peyré et al. [PCVS19] considered the transport of tensor-valued distributions on a mesh surface using quantum-entropy regularized optimization, which in practice leads to a modified Sinkhorn algorithm. They applied their algorithm to the interpolation of orientation fields (as opposed to vector fields) and metric tensors for anisotropic remeshing.

7. Applications in Topological Data Analysis

In topological data analysis, data can be analyzed via their persistence diagram, an object accounting for birth and death of features along a filtration [EH08]. A natural question when dealing with such objects is to compute distances and persistence diagrams, and optimal transport is particularly well suited for these tasks. Turner et al. [TMMH14] proposed a Wasserstein barycenter computation based on the auction algorithm for point matching, but it was extremely slow. Lacombe et al. [LCO18] proposed to use the Sinkhorn algorithm to compute distances between persistence diagrams discretized on a grid. To avoid the discretization step and the rasterization effects it can cause, Vidal et al. [VBT19] proposed a progressive Wasserstein barycenter computation technique. It iteratively decreases the Fréchet energy while increasing the accuracy of the computation, leading to more accurate results and better computation times.

8. Applications to simulation and animation

8.1. Fluid simulation

Brenier's polar factorization theorem [Bre91] indicates that optimal transport provides a projection of any map to the closest volume-preserving map. By considering the incompressible Euler equations as a flow under a volume-preserving map (i.e., one starts with the identity map, and at each time step, this map is transformed by a divergence-free velocity field), Gallouët and Mérigot proposed a Lagrangian fluid simulation scheme based on semi-discrete optimal transport [GM18] inspired by the discrete optimal transport scheme of Brenier [Bre00], which they illustrate in 2-d. The idea is to start with a set of particles covering the unit square along with initial velocities. Then, at each time step, one advects the particles according to their velocities, computes a semi-discrete optimal transport map between a uniform density on the unit square and the advected particles thus defining power cells (see Sec. 3.2), and finally pushes the particles towards the barycenter of their power cells in order to recover incompressibility. This formalizes the optimal-transport based fluid simulation heuristic presented earlier by de Goes et al. [dGWH*15]. Then, Lévy extended the approach to handle free-boundary fluids [Lév22], by considering a partial optimal transport problem, and illustrate it in 3-d (Fig. 16), with viscosity and surface tension added. Considering the weights of the power cell of each fluid particle as a Lagrange multiplier of mass preservation constraint, an additional Lagrange multiplier is introduced to receive the entire mass of the empty space unoccupied by fluid. The fluid interface is shown to be the intersection between power cells and a balls, which makes the modeling of the interface accurate. To make the fluid simulation faster, Qu et al. [QLDGJ22] leverage semi-discrete entropy-regularized optimal transport to avoid computing costly power diagrams, by applying Sinkhorn algorithm between particles and a uniform density grid.

8.2. Animation

Shape animation through optimal transport is a difficult task since optimal transport can break shapes and does not have any topological preservation guarantees. However it is extremely powerful at modeling unstructured data evolution. Zhang et al. [ZSS22] propose to interpolate unstructured freeform objects between user-provided keyframes by combining optimal transport with Continuous Normalizing Flows (CNF) [CRBD18]. Here, the Sinkhorn divergence is used to replace the Kullback-Leibler divergence in the CNF framework. It is also used to ensure that the trajectories adhere to the keyframes by computing Wasserstein barycenters in between ODE-advected keyframes. Webanck et al. [WCGG18] proposed to interpolate between clouds keyframes using pointwise density matching. In a different context, optimal transport was used to interpolate between point clouds captured from plants at several timesteps [GKK*20]: plant scans are segmented and the segments are matched across time. A Sinkhorn-regularized optimal transport is then computed in-between the matched segments, yielding interpolated point cloud sequences. Interestingly this method copes with the necessary topology preservation by matching segments, hence ensuring that optimal transport does not break the plant structure.

Another way to use optimal transport is for parameterizing shapes inside a shape space. Ma et al. [MDZ*21] used Wasserstein barycenters to model soft underwater swimmers, that are optimized via differentiable swimming simulation, yielding better swimming performances. Here optimal transport serves to parameterize the shape space and optimize easily the shape geometry, since it is fully differentiable using the formulation of Bonneel et al. [BPC16].

9. Other applications

Optimal transportation has been applied to a wide variety of other problems far beyond computer graphics and vision, from economics [Gal18] to operations research. We shall give a few words

© 2023 The Author(s)

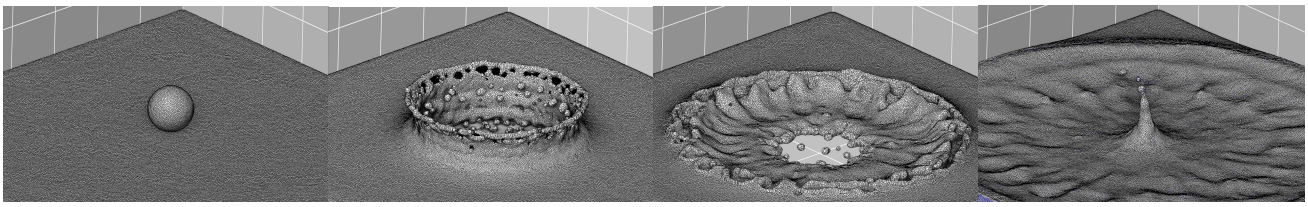

Figure 16: *Free-boundary 3d fluid simulation including viscosity and surface tension, by projecting velocities on divergence-free velocity fields using semi-discrete optimal transport, with 500k power cells [Lév22]*

on notable applications, either related to the computer graphics and vision communities, or widely adopted by their communities.

For music interpretation, optimal transport has been applied to music transcription, i.e. recovering the notes from a music spectrum [FFCE16]. Henderson and Solomon [HS19] used 1-d optimal transportation to match pitches between two audio signals yielding a natural *portamento* (i.e., a transition) even between different polyphonic instruments.

In cosmology, optimal transport has been used to reconstruct accurate acoustic baryon variations, using a dedicated algorithm based on semi-discrete optimal transport [vHLM22]. The idea is to compute a displacement interpolation between a uniform density (the universe at time zero) and the current distribution of mass in the universe.

In text retrieval and analysis, document distances can be computed by computing words embedding and comparing the words distribution in the Euclidean space, using Relaxed EMD [KSKW15].

Optimal transport has also been used extensively in the broader machine learning literature, as a way to cope with semi-supervision [SRGB14b] or to perform transfer learning [CFTR17].

In genomics, genes or gene signatures can be extracted from different cells at several times, and unbalanced optimal transport can be used to recover developmental trajectories by coupling these data over time [SST*19]. The Gene Mover's Distance [BCG*21] computes the similarity between two cells by computing an earth mover's distance with a ground distance consisting of the ℓ^2 distance in a 200-dimensional euclidean embedding of genes.

10. Discussion and conclusion

In this survey of optimal transport applications, we encountered various uses of optimal transport, notably for interpolation purposes. For such applications, care must be taken as optimal transport is agnostic to topology changes and semantics. For instance, interpolating between two sketches will produce a result which is well defined from an optimal transport perspective, but might not be interesting for the application: connected components might break and merge throughout the process, and semantics may be lacking. For unorganized matter, such as clouds, fluids, smoke, this is an acceptable behavior, but it is not the case for structured data such as human, animals or object shapes. When the input data do not directly represent probability distributions, one issue for this type of application is the interpretation of the input data (e.g., pixel intensities) as mass that ought to be preserved during interpolation.

This survey explored the many facets of optimal transport applications, from archetypal applications such as image color grading to more advanced or niche albeit interesting applications, or even applications whose main purpose is to illustrate one's optimal transport algorithm behavior. While optimal transport's main downside is its computational cost, many fast relaxations or approximations have been proposed and continue to be developed. Still, scaling to millions of variables remain to this day a challenge in time critical applications, preventing its wide integration in many computational problems.

Acknowledgments

We thank authors for their permission to use their images in our survey. This work was funded in part by ANR MoCaMed (ANR 20-CE45-0025) and TOPACS (ANR-19-CE45-0015).

References

[AAMT15] André J., Attali D., Mérigot Q., Thibert B.: Far-field reflector problem under design constraints. *International Journal of Computational Geometry & Applications 25*, 02 (2015), 143–162. 8, 13

[ABGV18] Auricchio G., Bassetti F., Gualandi S., Veneroni M.: Computing kantorovich-wasserstein distances on d-dimensional histograms using $(d+1)$-partite graphs. *Advances in Neural Information Processing Systems 31* (2018). 5

[ABM16] Anderes E., Borgwardt S., Miller J.: Discrete wasserstein barycenters: Optimal transport for discrete data. *Mathematical Methods of Operations Research 84*, 2 (2016), 389–409. 5

[ACB17] Arjovsky M., Chintala S., Bottou L.: Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning* (06–11 Aug 2017), vol. 70 of *Proceedings of Machine Learning Research*, PMLR, pp. 214–223. 8, 11

[AHA92] Aurenhammer F., Hoffmann F., Aronov B.: Minkowski-type theorems and least-squares partitioning. In *Proceedings of the eighth annual symposium on Computational geometry* (1992), pp. 350–357. 5

[AHA98] Aurenhammer F., Hoffmann F., Aronov B.: Minkowski-type theorems and least-squares clustering. *Algorithmica 20*, 1 (1998), 61–76. 15

[ARW22] Ahmed A. G., Ren J., Wonka P.: Gaussian blue noise. *arXiv preprint arXiv:2206.07798* (2022). 12

[BB00] Benamou J.-D., Brenier Y.: A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. *Numerische Mathematik 84*, 3 (2000), 375–393. 7

[BC19] Bonneel N., Coeurjolly D.: SPOT: Sliced Partial Optimal Transport. *ACM Transactions on Graphics (SIGGRAPH) 38*, 4 (July 2019). 4, 7, 8, 14

© 2023 The Author(s)

[BCG*21] BELLAZZI R., CODEGONI A., GUALANDI S., NICORA G., VERCESI E.: The gene mover's distance: Single-cell similarity via optimal transport. *arXiv preprint arXiv:2102.01218* (2021). 18

[BCM05] BUADES A., COLL B., MOREL J.-M.: A non-local algorithm for image denoising. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)* (2005), vol. 2, pp. 60–65 vol. 2. 8

[BM92] BESL P., MCKAY N. D.: A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence 14*, 2 (1992), 239–256. 14

[BMP02] BELONGIE S., MALIK J., PUZICHA J.: Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence 24*, 4 (2002), 509–522. 7

[Bon13] BONNOTTE N.: *Unidimensional and evolution methods for optimal transportation*. PhD thesis, Paris 11, 2013. 6

[BPC16] BONNEEL N., PEYRÉ G., CUTURI M.: Wasserstein barycentric coordinates: histogram regression using optimal transport. *ACM Trans. Graph. 35*, 4 (2016), 71–1. 8, 13, 17

[Bre91] BRENIER Y.: Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics 44*, 4 (1991), 375–417. 17

[Bre00] BRENIER Y.: Derivation of the euler equations from a caricature of coulomb interaction. *Communications in Mathematical Physics 212*, 1 (2000), 93–104. 17

[BRPP15] BONNEEL N., RABIN J., PEYRÉ G., PFISTER H.: Sliced and Radon Wasserstein Barycenters of Measures. *Journal of Mathematical Imaging and Vision 1*, 51 (2015), 22–45. 6, 8, 9, 11

[BSPP13] BONNEEL N., SUNKAVALLI K., PARIS S., PFISTER H.: Example-based video color grading. *ACM Trans. Graph. 32*, 4 (2013), 39–1. 8, 9

[BvdPPH11] BONNEEL N., VAN DE PANNE M., PARIS S., HEIDRICH W.: Displacement Interpolation Using Lagrangian Mass Transport. *ACM Transactions on Graphics (SIGGRAPH ASIA 2011) 30*, 6 (2011). 5, 8, 13

[CD14] CUTURI M., DOUCET A.: Fast computation of wasserstein barycenters. In *International conference on machine learning* (2014), PMLR, pp. 685–693. 6

[CDC22] COEURDOUX F., DOBIGEON N., CHAINAIS P.: Learning optimal transport between two empirical distributions with normalizing flows. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD 2022)* (Grenoble, France, Sept. 2022). 7

[CEN07] CHAN T., ESEDOGLU S., NI K.: Histogram based segmentation using wasserstein distances. In *Scale Space and Variational Methods in Computer Vision* (Berlin, Heidelberg, 2007), Sgallari F., Murli A., Paragios N., (Eds.), Springer Berlin Heidelberg, pp. 697–708. 8, 10

[CFD18] COURTY N., FLAMARY R., DUCOFFE M.: Learning Wasserstein Embeddings. In *ICLR 2018 - 6th International Conference on Learning Representations* (Vancouver, Canada, Apr. 2018), pp. 1–13. 7

[CFTR17] COURTY N., FLAMARY R., TUIA D., RAKOTOMAMONJY A.: Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence 39*, 9 (2017), 1853–1865. 8, 10, 18

[CPSV18] CHIZAT L., PEYRÉ G., SCHMITZER B., VIALARD F.-X.: Scaling algorithms for unbalanced optimal transport problems. *Mathematics of Computation 87*, 314 (2018), 2563–2609. 6, 7, 8, 14

[CQW*19] CUI L., QI X., WEN C., LEI N., LI X., ZHANG M., GU X.: Spherical optimal transportation. *Computer-Aided Design 115* (2019), 181–193. 5

[CRBD18] CHEN R. T. Q., RUBANOVA Y., BETTENCOURT J., DUVENAUD D. K.: Neural ordinary differential equations. In *Advances in Neural Information Processing Systems* (2018), Bengio S., Wallach H., Larochelle H., Grauman K., Cesa-Bianchi N., Garnett R., (Eds.), vol. 31, Curran Associates, Inc. 17

[DB22] DE BORTOLI V.: Convergence of denoising diffusion models under the manifold hypothesis. *Transactions on Machine Learning Research (TMLR)* (2022). 11

[DCSA*14] DIGNE J., COHEN-STEINER D., ALLIEZ P., DE GOES F., DESBRUN M.: Feature-preserving surface reconstruction and simplification from defect-laden point sets. *Journal of Mathematical Imaging and Vision 48*, 2 (2014), 369–382. 8, 15

[dGBOD12] DE GOES F., BREEDEN K., OSTROMOUKHOV V., DESBRUN M.: Blue noise through optimal transport. *ACM Trans. Graph. 31*, 6 (nov 2012). 8, 12, 13

[dGWH*15] DE GOES F., WALLEZ C., HUANG J., PAVLOV D., DESBRUN M.: Power particles: An incompressible fluid solver based on power diagrams. *ACM Trans. Graph. 34*, 4 (jul 2015). 8, 17

[DJK11] DEZSŐ B., JÜTTNER A., KOVÁCS P.: Lemon–an open source c++ graph template library. *Electronic notes in theoretical computer science 264*, 5 (2011), 23–45. 5

[DSS12] DELON J., SALOMON J., SOBOLEVSKI A.: Local matching indicators for transport problems with concave costs. *SIAM Journal on Discrete Mathematics 26*, 2 (2012), 801–827. 4

[DT09] DOMINITZ A., TANNENBAUM A.: Texture mapping via optimal mass transport. *IEEE transactions on visualization and computer graphics 16*, 3 (2009), 419–433. 8, 16

[EH08] EDELSBRUNNER H., HARER J.: Persistent homology —– a survey. In *Surveys on discrete and computational geometry* (2008), vol. 453, Amer Mathematical Society, p. 257. 17

[ELC20] EISENBERGER M., LAHNER Z., CREMERS D.: Smooth shells: Multi-scale shape registration with functional maps. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (jun 2020), IEEE Computer Society, pp. 12262–12271. 15

[ES18] ESSID M., SOLOMON J.: Quadratically regularized optimal transport on graphs. *SIAM Journal on Scientific Computing 40*, 4 (2018), A1961–A1986. 6

[ETLTC20] EISENBERGER M., TOKER A., LEAL-TAIXÉ L., CREMERS D.: Deep shells: Unsupervised shape correspondence with optimal transport. In *Advances in Neural Information Processing Systems* (2020), Larochelle H., Ranzato M., Hadsell R., Balcan M., Lin H., (Eds.), vol. 33, Curran Associates, Inc., pp. 10491–10502. 8, 15

[FCG*21] FLAMARY R., COURTY N., GRAMFORT A., ALAYA M. Z., BOISBUNON A., CHAMBON S., CHAPEL L., CORENFLOS A., FATRAS K., FOURNIER N., GAUTHERON L., GAYRAUD N. T., JANATI H., RAKOTOMAMONJY A., REDKO I., ROLET A., SCHUTZ A., SEGUY V., SUTHERLAND D. J., TAVENARD R., TONG A., VAYER T.: Pot: Python optimal transport. *Journal of Machine Learning Research 22*, 78 (2021), 1–8. 7

[Fey20] FEYDY J.: *Geometric data analysis, beyond convolutions*. PhD thesis, Université Paris-Saclay Gif-sur-Yvette, France, 2020. 6

[FFCE16] FLAMARY R., FÉVOTTE C., COURTY N., EMIYA V.: Optimal spectral transportation with application to music transcription. In *Advances in Neural Information Processing Systems (NIPS)* (2016). 18

[FSDH14] FRIGO O., SABATER N., DEMOULIN V., HELLIER P.: Optimal transportation for example-guided color transfer. In *Asian Conference on Computer Vision* (2014), Springer, pp. 655–670. 8

[FSV*19] FEYDY J., SÉJOURNÉ T., VIALARD F.-X., AMARI S.-I., TROUVÉ A., PEYRÉ G.: Interpolating between optimal transport and mmd using sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics* (2019), PMLR, pp. 2681–2690. 6

[FXPA13] FERRADANS S., XIA G.-S., PEYRÉ G., AUJOL J.-F.: Static and dynamic texture mixing using optimal transport. In *Scale Space and Variational Methods in Computer Vision* (Berlin, Heidelberg, 2013), Kuijper A., Bredies K., Pock T., Bischof H., (Eds.), Springer Berlin Heidelberg, pp. 137–148. 8, 11

[GAA*17] GULRAJANI I., AHMED F., ARJOVSKY M., DUMOULIN V., COURVILLE A. C.: Improved training of wasserstein gans. In *Advances*

© 2023 The Author(s)

in Neural Information Processing Systems (2017), Guyon I., Luxburg U. V., Bengio S., Wallach H., Fergus R., Vishwanathan S., Garnett R., (Eds.), vol. 30, Curran Associates, Inc. 8, 11

[Gal18] GALICHON A.: *Optimal transport methods in economics*. Princeton University Press, 2018. 17

[GCSAD11] GOES F. D., COHEN-STEINER D., ALLIEZ P., DESBRUN M.: An optimal transport approach to robust reconstruction and simplification of 2d shapes. *Computer Graphics Forum* (2011). 8, 15

[GEB15] GATYS L., ECKER A. S., BETHGE M.: Texture synthesis using convolutional neural networks. In *Advances in Neural Information Processing Systems* (2015), Cortes C., Lawrence N., Lee D., Sugiyama M., Garnett R., (Eds.), vol. 28. 9, 10, 11

[GEB16] GATYS L. A., ECKER A. S., BETHGE M.: Image style transfer using convolutional neural networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016), pp. 2414–2423. 9

[GK18] GOLLA T., KLEIN R.: Interactive interpolation of metallic effect car paints. In *Proceedings of the Conference on Vision, Modeling, and Visualization* (2018), pp. 11–20. 8, 13

[GKK*20] GOLLA T., KNEIPHOF T., KUHLMANN H., WEINMANN M., KLEIN R.: Temporal upsampling of point cloud sequences by optimal transport for plant growth visualization. *Computer Graphics Forum 39*, 6 (2020), 167–179. 8, 17

[GLR18] GALERNE B., LECLAIRE A., RABIN J.: A texture synthesis model based on semi-discrete optimal transport in patch space. *SIAM Journal on Imaging Sciences 11*, 4 (2018), 2456–2493. 8, 10

[GM18] GALLOUËT T. O., MÉRIGOT Q.: A lagrangian scheme à la brenier for the incompressible euler equations. *Foundations of Computational Mathematics 18*, 4 (2018), 835–865. 8, 17

[GO03] GLIMM T., OLIKER V.: Optical design of single reflector systems and the monge–kantorovich mass transfer problem. *Journal of Mathematical Sciences 117*, 3 (2003), 4096–4108. 13

[GRK*22] GAZDIEVA M., ROUT L., KOROTIN A., FILIPPOV A., BURNAEV E.: Unpaired image super-resolution with optimal transport maps. *arXiv preprint arXiv:2202.01116* (2022). 8, 12

[HB95] HEEGER D. J., BERGEN J. R.: Pyramid-based texture analysis/synthesis. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques* (1995), pp. 229–238. 8, 10, 11, 12

[HJA20] HO J., JAIN A., ABBEEL P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems 33* (2020), 6840–6851. 11

[HLPR22] HOUDARD A., LECLAIRE A., PAPADAKIS N., RABIN J.: A generative model for texture synthesis based on optimal transport between feature distributions. *Journal of Mathematical Imaging and Vision* (2022), 1–25. 8, 10

[HMP15] HUG, ROMAIN, MAITRE, EMMANUEL, PAPADAKIS, NICOLAS: Multi-physics optimal transportation and image interpolation. *ESAIM: M2AN 49*, 6 (2015), 1671–1692. 8, 10

[HN18] HEITZ E., NEYRET F.: High-performance by-example noise using a histogram-preserving blending operator. *Proceedings of the ACM on Computer Graphics and Interactive Techniques 1*, 2 (2018), 1–25. 8, 11

[HS19] HENDERSON T., SOLOMON J.: Audio transport: A generalized portamento via optimal transport. In *International Conference on Digital Audio Effects 2019* (2019). 18

[HVCB21] HEITZ E., VANHOEY K., CHAMBON T., BELCOUR L.: A sliced wasserstein loss for neural texture synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2021). 8, 9, 10, 11

[HZTA04] HAKER S., ZHU L., TANNENBAUM A., ANGENENT S.: Optimal mass transport for registration and warping. *International Journal of computer vision 60*, 3 (2004), 225–240. 8, 10

[IT03] INDYK P., THAPER N.: Fast image retrieval via embeddings. In *3rd international workshop on statistical and computational theories of vision* (2003), vol. 2, Nice, France, p. 5. 7, 8

[JH99] JOHNSON A., HEBERT M.: Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence 21*, 5 (1999), 433–449. 7

[Kan42] KANTOROVICH L. V.: On the translocation of masses. In *Dokl. Akad. Nauk. USSR (NS)* (1942), vol. 37, pp. 199–201. 1

[KKTH10] KOLESOV I., KARASEV P., TANNENBAUM A., HABER E.: Fire and smoke detection in video with optimal mass transport based optical flow and neural networks. In *2010 IEEE International Conference on Image Processing* (2010), IEEE, pp. 761–764. 8, 10

[KMT19] KITAGAWA J., MÉRIGOT Q., THIBERT B.: Convergence of a newton algorithm for semi-discrete optimal transport. *Journal of the European Mathematical Society 21*, 9 (2019), 2603–2651. 5

[KSKW15] KUSNER M. J., SUN Y., KOLKIN N. I., WEINBERGER K. Q.: From word embeddings to document distances. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37* (2015), ICML'15, JMLR.org, p. 957–966. 9, 18

[KSS19] KOLKIN N. I., SALAVON J., SHAKHNAROVICH G.: Style transfer by relaxed optimal transport and self-similarity. *IEEE CVPR 2018* (2019). 8, 9

[LBFS21] LAVOUÉ G., BONNEEL N., FARRUGIA J.-P., SOLER C.: Perceptual quality of brdf approximations: dataset and metrics. In *Computer Graphics Forum* (2021), vol. 40, Wiley Online Library, pp. 327–338. 8, 13

[LCCS18] LAVENANT H., CLAICI S., CHIEN E., SOLOMON J.: Dynamical optimal transport on discrete surfaces. *ACM Trans. Graph. 37*, 6 (dec 2018). 8, 17

[LCO18] LACOMBE T., CUTURI M., OUDOT S.: Large scale computation of means and clusters for persistence diagrams using optimal transport. In *Advances in Neural Information Processing Systems* (2018), Bengio S., Wallach H., Larochelle H., Grauman K., Cesa-Bianchi N., Garnett R., (Eds.), vol. 31, Curran Associates, Inc. 8, 17

[LD09] LIPMAN Y., DAUBECHIES I.: Surface comparison with mass transportation, 2009. 8, 14

[LD11] LIPMAN Y., DAUBECHIES I.: Conformal wasserstein distances: Comparing surfaces in polynomial time. *Advances in Mathematics 227*, 3 (jun 2011), 1047–1077. 8, 14

[LDCB22] LACOMBE J., DIGNE J., COURTY N., BONNEEL N.: Learning to Generate Wasserstein Barycenters. *Journal of Mathematical Imaging and Vision* (2022). 7, 8, 10

[LdGKW19] LEBRAT L., DE GOURNAY F., KAHN J., WEISS P.: Optimal transport approximation of 2-dimensional measures. *SIAM Journal on Imaging Sciences 12*, 2 (2019), 762–787. 8, 12, 13

[Lév15] LÉVY B.: A numerical algorithm for l2 semi-discrete optimal transport in 3d. *ESAIM: M2AN 49*, 6 (2015), 1693–1715. 5, 6, 8, 14, 15

[Lév22] LÉVY B.: Partial optimal transport for a constant-volume lagrangian mesh with free boundaries. *Journal of Computational Physics 451* (2022), 110838. 5, 7, 8, 17, 18

[Llo82] LLOYD S.: Least squares quantization in pcm. *IEEE Transactions on Information Theory 28*, 2 (1982), 129–137. 12

[LLX21] LI R., LIN G., XIE L.: Self-point-flow: Self-supervised scene flow estimation from point clouds with optimal transport and random walk. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2021), pp. 15577–15586. 8, 15

[LO07] LING H., OKADA K.: An efficient earth mover's distance algorithm for robust histogram comparison. *IEEE Trans. Pattern Anal. Mach. Intell. 29*, 5 (may 2007), 840–853. 5, 7, 8

[Low04] LOWE D. G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision 60*, 2 (2004), 91–110. 7

© 2023 The Author(s)

[LR21] LECLAIRE A., RABIN J.: A stochastic multi-layer algorithm for semi-discrete optimal transport with applications to texture synthesis and style transfer. *Journal of Mathematical Imaging and Vision 63*, 2 (2021), 282–308. 8, 9, 10

[LS18] LÉVY B., SCHWINDT E. L.: Notions of optimal transport theory and how to implement them on a computer. *Computers & Graphics 72* (2018), 135–148. 6

[LSC*19] LEI N., SU K., CUI L., YAU S.-T., GU X. D.: A geometric view of optimal transportation and generative model. *Computer Aided Geometric Design 68* (2019), 1–21. 8, 11

[LZYY20] LIU Y., ZHU L., YAMADA M., YANG Y.: Semantic correspondence as an optimal transport problem. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 4463–4472. 8, 10

[MCSK*17] MANDAD M., COHEN-STEINER D., KOBBELT L., ALLIEZ P., DESBRUN M.: Variance-minimizing transport plans for inter-surface mapping. *ACM Trans. Graph. 36*, 4 (jul 2017). 8, 15

[MDZ*21] MA P., DU T., ZHANG J. Z., WU K., SPIELBERG A., KATZSCHMANN R. K., MATUSIK W.: Diffaqua: A differentiable computational design pipeline for soft underwater swimmers with shape interpolation. *ACM Transactions on Graphics (TOG) 40*, 4 (2021), 132. 8, 17

[Mém11] MÉMOLI F.: Gromov–wasserstein distances and the metric approach to object matching. *Found. Comput. Math. 11*, 4 (2011), 417–487. 7

[Mér11] MÉRIGOT Q.: A multiscale approach to optimal transport. *Computer Graphics Forum 30*, 5 (Aug. 2011), 1584–1592. 18 pages. 5, 8, 10

[MM99] MCASEY M., MOU L.: Optimal locations and the mass transport problem. *Contemporary Mathematics 226* (1999), 131–148. 12

[MMT18a] MÉRIGOT Q., MEYRON J., THIBERT B.: An algorithm for optimal transport between a simplex soup and a point cloud. *SIAM Journal on Imaging Sciences 11*, 2 (2018), 1363–1389. 8, 15

[MMT18b] MEYRON J., MÉRIGOT Q., THIBERT B.: Light in power: A general and parameter-free algorithm for caustic design. *ACM Trans. Graph. 37*, 6 (dec 2018). 8, 14, 15

[Mon81] MONGE G.: Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.* (1781), 666–704. 1

[MS03] MOROVIC J., SUN P.-L.: Accurate 3d image colour histogram transformation. *Pattern Recognition Letters 24*, 11 (2003), 1725–1735. 8

[MZD05] MATUSIK W., ZWICKER M., DURAND F.: Texture design using a simplicial complex of morphable textures. *ACM Transactions on Graphics (TOG) 24*, 3 (2005), 787–794. 8, 11, 12

[NBCE09] NI K., BRESSON X., CHAN T., ESEDOGLU S.: Local histogram based segmentation using the wasserstein distance. *Int. J. Comput. Vision 84*, 1 (aug 2009), 97–111. 8, 10

[NG18] NADER G., GUENNEBAUD G.: Instant Transport Maps on 2D Grids. *ACM Transactions on Graphics 37*, 6 (Nov. 2018), 13. 7, 12

[OR20] OBERMAN A. M., RUAN Y.: Solution of optimal transportation problems using a multigrid linear programming approach. *Journal of Computational Mathematics 38*, 6 (2020), 933. 5

[Pap15] PAPADAKIS N.: *Optimal transport for image processing*. PhD thesis, Université de Bordeaux, 2015. 1

[PBC*20] PAULIN L., BONNEEL N., COEURJOLLY D., IEHL J.-C., WEBANCK A., DESBRUN M., OSTROMOUKHOV V.: Sliced optimal transport sampling. *ACM Trans. Graph. 39*, 4 (2020), 99. 8, 12, 13

[PC17] PEYRÉ G., CUTURI M.: Computational optimal transport. *Center for Research in Economics and Statistics Working Papers*, 2017-86 (2017). 1

[PC19] PEYRÉ G., CUTURI M.: Computational optimal transport. *Foundations and Trends in Machine Learning 11*, 5-6 (2019), 355–607. 1, 4

[PCVS19] PEYRÉ G., CHIZAT L., VIALARD F.-X., SOLOMON J.: Quantum entropic regularization of matrix-valued optimal transport. *European Journal of Applied Mathematics 30*, 6 (Dec. 2019), 1079–1102. 8, 11, 17

[PFR12] PEYRÉ G., FADILI J. M., RABIN J.: Wasserstein Active Contours. In *ICIP'12* (2012), pp. 2541 – 2544. 8, 10

[PK07] PITIÉ F., KOKARAM A.: The linear monge-kantorovitch linear colour mapping for example-based colour transfer. 8, 9

[PKD05] PITIE F., KOKARAM A. C., DAHYOT R.: N-dimensional probability density function transfer and its application to color transfer. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1* (2005), vol. 2, IEEE, pp. 1434–1439. 8

[PKD07] PITIÉ F., KOKARAM A. C., DAHYOT R.: Automated colour grading using colour distribution transfer. *Computer Vision and Image Understanding 107*, 1-2 (2007), 123–137. 8, 9

[PPO14] PAPADAKIS N., PEYRÉ G., OUDET E.: Optimal transport with proximal splitting. *SIAM Journal on Imaging Sciences 7*, 1 (2014), 212–238. 7, 8, 10

[PR17] PAPADAKIS N., RABIN J.: Convex histogram-based joint image segmentation with regularized optimal transport cost. *Journal of Mathematical Imaging and Vision* (2017). 8, 10

[PRM*21] PAI G., REN J., MELZI S., WONKA P., OVSJANIKOV M.: Fast sinkhorn filters: Using matrix scaling for non-rigid shape correspondence with functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2021), pp. 384–393. 8, 15

[PSC*15] PILLEBOUE A., SINGH G., COEURJOLLY D., KAZHDAN M., OSTROMOUKHOV V.: Variance analysis for monte carlo integration. *ACM Transactions on Graphics (TOG) 34*, 4 (2015), 1–14. 12

[PW08] PELE O., WERMAN M.: A linear time histogram metric for improved sift matching. In *Computer Vision – ECCV 2008* (2008), Springer Berlin Heidelberg, pp. 495–508. 7, 8

[PW09] PELE O., WERMAN M.: Fast and robust earth mover's distances. In *2009 IEEE 12th International Conference on Computer Vision* (2009), pp. 460–467. 7, 8

[QCHC17] QIN H., CHEN Y., HE J., CHEN B.: Wasserstein blue noise sampling. *ACM Transactions on Graphics (TOG) 36*, 5 (2017), 1–13. 8, 12

[QLDGJ22] QU Z., LI M., DE GOES F., JIANG C.: The power particle-in-cell method. *ACM Trans. Graph. 41*, 4 (jul 2022). 8, 17

[RDG10] RABIN J., DELON J., GOUSSEAU Y.: Regularization of Transportation Maps for Color and Contrast Transfer. In *ICIP* (Sept. 2010), IEEE, pp. 1933–1936. 8, 9

[RDG11] RABIN J., DELON J., GOUSSEAU Y.: Transportation distances on the circle. *J. Math. Imaging Vis. 41*, 1–2 (2011), 147–167. 4

[RFP14] RABIN J., FERRADANS S., PAPADAKIS N.: Adaptive color transfer with relaxed optimal transport. In *2014 IEEE international conference on image processing (ICIP)* (2014), IEEE, pp. 4852–4856. 8

[RPC10] RABIN J., PEYRÉ G., COHEN L. D.: Geodesic shape retrieval via optimal mass transport. In *European Conference on Computer Vision* (2010), Springer, pp. 771–784. 8, 14

[RPDB12] RABIN J., PEYRÉ G., DELON J., BERNOT M.: Wasserstein barycenter and its application to texture mixing. In *Scale Space and Variational Methods in Computer Vision* (2012), Springer Berlin Heidelberg. 6, 8, 11, 14

[RTG00] RUBNER Y., TOMASI C., GUIBAS L. J.: The earth mover's distance as a metric for image retrieval. *International journal of computer vision 40*, 2 (2000), 99–121. 7, 8

[SA20] SIMON D., ABERDAM A.: Barycenters of natural images constrained wasserstein barycenters for image morphing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 7910–7919. 8, 10

© 2023 The Author(s)

[San15] SANTAMBROGIO F.: Optimal transport for applied mathematicians. *Birkäuser, NY 55*, 58-63 (2015), 94. 1

[Sch16] SCHMITZER B.: A sparse multiscale algorithm for dense optimal transport. *Journal of Mathematical Imaging and Vision 56*, 2 (2016), 238–259. 5

[SdGP*15] 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. Graph. 34*, 4 (jul 2015). 6, 7, 8, 13, 14, 16, 17

[SFL*21] SHEN Z., FEYDY J., LIU P., CURIALE A. H., SAN JOSE ESTEPAR R., SAN JOSE ESTEPAR R., NIETHAMMER M.: Accurate point cloud registration with robust optimal transport. In *Advances in Neural Information Processing Systems* (2021), Ranzato M., Beygelzimer A., Dauphin Y., Liang P., Vaughan J. W., (Eds.), vol. 34, Curran Associates, Inc., pp. 5373–5389. 8, 14

[SGSS22] SALAÜN C., GEORGIEV I., SEIDEL H.-P., SINGH G.: Scalable multi-class sampling via filtered sliced optimal transport. *ACM Transactions on Graphics (SIGGRAPH ASIA 2022) 41*, 6 (2022). 8, 12

[SHS*17] SHU Z., HADAP S., SHECHTMAN E., SUNKAVALLI K., PARIS S., SAMARAS D.: Portrait lighting transfer using a mass transport approach. *ACM Trans. Graph. 37*, 1 (2017). 8, 9

[SRGB14a] SOLOMON J., RUSTAMOV R., GUIBAS L., BUTSCHER A.: Earth mover's distances on discrete surfaces. *ACM Trans. Graph. 33*, 4 (jul 2014). 7, 8, 16, 17

[SRGB14b] SOLOMON J., RUSTAMOV R., GUIBAS L., BUTSCHER A.: Wasserstein propagation for semi-supervised learning. In *Proceedings of the 31st International Conference on Machine Learning* (Bejing, China, 22–24 Jun 2014), Xing E. P., Jebara T., (Eds.), vol. 32 of *Proceedings of Machine Learning Research*, PMLR, pp. 306–314. 18

[SST*19] SCHIEBINGER G., SHU J., TABAKA M., CLEARY B., SUBRAMANIAN V., SOLOMON A., GOULD J., LIU S., LIN S., BERUBE P., ET AL.: Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell 176*, 4 (2019), 928–943. 18

[STTP14] SCHWARTZBURG Y., TESTUZ R., TAGLIASACCHI A., PAULY M.: High-contrast computational caustic design. *ACM Transactions on Graphics (TOG) 33*, 4 (2014), 1–11. 8, 13, 15

[SV19] SOLOMON J., VAXMAN A.: Optimal transport-based polar interpolation of directional fields. *ACM Trans. Graph. 38*, 4 (jul 2019). 8, 17

[SWS*15] SU Z., WANG Y., SHI R., ZENG W., SUN J., LUO F., GU X.: Optimal mass transport for shape matching and comparison. *IEEE Transactions on Pattern Analysis and Machine Intelligence 37*, 11 (2015), 2246–2259. 5, 8, 16

[TMMH14] TURNER K., MILEYKO Y., MUKHERJEE S., HARER J.: Fréchet means for distributions of persistence diagrams. *Discrete Comput. Geom. 52*, 1 (jul 2014), 44–70. 8, 17

[TPG16] TARTAVEL G., PEYRÉ G., GOUSSEAU Y.: Wasserstein loss for image synthesis and restoration. *SIAM Journal on Imaging Sciences 9*, 4 (2016), 1726–1755. 8, 10, 11

[TPK*17] THORPE M., PARK S., KOLOURI S., ROHDE G. K., SLEPČEV D.: A transportation lp distance for signal analysis. *Journal of mathematical imaging and vision 59*, 2 (2017), 187–210. 8

[Var67] VARADHAN S. R. S.: Diffusion processes in a small time interval. *Communications on Pure and Applied Mathematics 20*, 4 (1967), 659–685. 17

[VBT19] VIDAL J., BUDIN J., TIERNY J.: Progressive Wasserstein barycenters of persistence diagrams. *IEEE transactions on visualization and computer graphics 26*, 1 (2019), 151–161. 8, 17

[VDKCC20] VACHER J., DAVILA A., KOHN A., COEN-CAGLI R.: Texture interpolation for probing visual perception. *Advances in Neural Information Processing Systems* (2020). 8, 11

[vHLM22] VON HAUSEGGER S., LÉVY B., MOHAYAEE R.: Accurate baryon acoustic oscillations reconstruction via semidiscrete optimal transport. *Phys. Rev. Lett. 128* (May 2022), 201302. 6, 18

[Vil09] VILLANI C.: *Optimal transport: old and new*, vol. 338. Springer, 2009. 1, 4

[WCGG18] WEBANCK A., CORTIAL Y., GUÉRIN E., GALIN E.: Procedural Cloudscapes. *Computer Graphics Forum* (2018). 8, 17

[WKB14] WARD G., KURT M., BONNEEL N.: Reducing anisotropic bsdf measurement to common practice. In *Material Appearance Modeling* (2014), pp. 5–8. 8, 13

[XLC*16] XIN S.-Q., LÉVY B., CHEN Z., CHU L., YU Y., TU C., WANG W.: Centroidal power diagrams with capacity constraints: Computation, applications, and extension. *ACM Transactions on Graphics (TOG) 35*, 6 (2016), 1–12. 8, 12

[XLY*22] XING J., LUAN F., YAN L.-Q., HU X., QIAN H., XU K.: Differentiable rendering using rgbxy derivatives and optimal transport. *ACM Transactions on Graphics (TOG) 41*, 6 (2022), 1–13. 8, 10

[ZSG*13] ZHAO X., SU Z., GU X. D., KAUFMAN A., SUN J., GAO J., LUO F.: Area-preservation mapping using optimal mass transport. *IEEE transactions on visualization and computer graphics 19*, 12 (2013), 2838–2847. 8, 16

[ZSS22] ZHANG P., SMIRNOV D., SOLOMON J.: Wassersplines for stylized neural animation, 2022. 8, 17

© 2023 The Author(s)