An Optimal Transport Approach to Robust Reconstruction and Simplification of 2D Shapes

Fernando de Goes, David Cohen-Steiner, Pierre Alliez, Mathieu Desbrun

March 22, 2018

Problem

Given a point set $\mathcal{S}$ and considering $\mathcal{S}$ as a measure μ consisting of Dirac masses, find a coarse simplicial complex $\mathcal{T}$ such that μ is well approximated by a linear combination of uniform measures on the edges and vertices of $\mathcal{T}$.

Contributions

Optimal transport formulation allows for a unified treatment of noise, outliers, boundaries and sharp features. None of the previous work could handle all of these concurrently.

Preliminaries

Delaunay Triangulation(DT) - Given a point set P, $DT(P)$ is a triangulation of P such that no point in P lies inside the circumcircle of any triangle in $DT(P)$. Such a triangulation can be computed in time $O(n \log n)$.

Half-edge - A half-edge is a directed line segment described by an origin vertex and a destination vertex.

One-ring of a point - The one-ring of a point x in a triangulation $\mathcal{T}$ is the set of all vertices adjacent to x in $\mathcal{T}$.

Kernel of a polygon P - This is a non-empty set K of points in the interior of P such that there exists a line segment from every point in K to every other point in P lying entirely inside P.

Flippable edge - An edge e in a triangulation $\mathcal{T}$ is called flippable if its end points and its two opposite vertices form a convex quadrilateral.

Shape reconstruction algorithm

1. **Input** - point set $\mathcal{S} = \{p_1, \ldots, p_n\}$
2. Construct Delaunay triangulation $\mathcal{T}_0$ of $\mathcal{S}$.
3. Compute initial transport plan π_0 from $\mathcal{S}$ to $\mathcal{T}_0$.
4. Set $k = 0$.
5. **Repeat** steps 6-11 **Until** desired vertex count is obtained.
6. Pick best half-edge $e = (x_i, x_j)$ to collapse (simplification).
7. Create $\mathcal{T}_{k+1}$ by merging x_i onto x_j.
8. $\pi'_{k+1} := \pi_k$ with local reassignments (update transport).
9. Optimize position of vertices in the one-ring of x_i (vertex relocation).
10. $\pi_{k+1} := \pi'_{k+1}$ with local reassignments (update transport).
11. $k \rightarrow k + 1$.
12. Filter edges based on relevance (optional).

Optimal Transport Formulation

- Let $\mathcal{S} = \{p_i\}_{i \in I}$ denote the input point set. Every point p_i is seen as a Dirac measure μ_i centered at p_i and of mass m_i. The point set is thus considered as a measure $\mu = \sum_i \mu_i$.
- Assume that we are given a triangulation $\mathcal{T}$ and a *point to simplex assignment* which maps every point p_i to either an edge e or a vertex v of $\mathcal{T}$.
- Each vertex v of $\mathcal{T}$ is seen as a Dirac measure and every edge e is a uniform 1D measure defined over the edge e.

Optimal Transport Formulation

- For every vertex v of $\mathcal{T}$, let $\mathcal{S}_v$ denote the set of points of $\mathcal{S}$ assigned to the vertex v.
- For every edge e of $\mathcal{T}$, let $\mathcal{S}_e$ denote the set of points of $\mathcal{S}$ assigned to the edge e.
- Assume these sets are disjoint with $\cup_{v \in \mathcal{T}} \mathcal{S}_v \cup \cup_{e \in \mathcal{T}} \mathcal{S}_e = \mathcal{S}$.
- Let M_v denote the total mass of $\mathcal{S}_v$ and M_e denote the total mass of $\mathcal{S}_e$. Then, $\sum_{e \in \mathcal{T}} M_e + \sum_{v \in \mathcal{T}} M_v = \sum_i m_i$.
- Let π denote the transport plan satisfying the *point to simplex assignment* and $W_2(\pi)$ its transport cost.

Optimal Transport Cost

- **Points to Vertex** - For a vertex $v \in \mathcal{T}$, the cost to transport the measure $\mathcal{S}_v$ to the Dirac measure centered on v with mass M_v is given by

$$W_2(v, \mathcal{S}_v) = \sqrt{\sum_{p_i \in \mathcal{S}_v} m_i ||p_i - v||^2}.$$

- **Points to Edge** - For an edge e, the transport plan is decomposed into a normal and a tangential component to e. For every $p_i \in \mathcal{S}_e$, let q_i denote the orthogonal projection of p_i onto e. The transport cost N of the normal plan is given by

$$N(e, \mathcal{S}_e) = \sqrt{\sum_{p_i \in \mathcal{S}_e} m_i ||p_i - q_i||^2}.$$

Optimal Transport Cost

The tangential plan is obtained as follows:

- The projected points $\{q_i\}$ are sorted along e and the edge is partitioned into $|\mathcal{S}_e|$ segment bins, with the i-th bin having length $l_i = (m_i/M_e) len(e)$. Here, $len(e)$ denoted the length of edge e.
- Consider a point p_i of mass m_i that projects onto q_i on edge e. Set a 1D coordinate axis along the edge with origin at the center of the i-th bin and let c_i be the coordinate of q_i in this coordinate axis. The tangential cost t_i of p_i is given by

$$t_i = \frac{M_e}{len(e)} \int_{-l_i/2}^{l_i/2} (x - c_i)^2 dx = m_i \left(\frac{l_i^2}{12} + c_i^2 \right).$$

Optimal Transport Cost

The tangential component of the optimal transport cost for an edge e is given by

$$T(e, \mathcal{S}_e) = \sqrt{\sum_{p_i \in \mathcal{S}_e} m_i \left(\frac{l_i^2}{12} + c_i^2 \right)}.$$

Note that the above definition of tangential cost ensures that the boundaries and features are preserved.
The total cost to transport $\mathcal{S}$ to $\mathcal{T}$ through the transport plan π is therefore given by

$$W_2(\pi) = \sqrt{\sum_{e \in \mathcal{T}} [N(e, \mathcal{S}_e)^2 + T(e, \mathcal{S}_e)^2] + \sum_{v \in \mathcal{T}} W_2(v, \mathcal{S}_v)^2}.$$

Point to simplex assignment

Given a triangulation $\mathcal{T}$, an assignment of the point set $\mathcal{S}$ to the vertices and edges of $\mathcal{T}$ is given as follows:

- Each point p_i is first temporarily assigned to the closest edge of the simplicial complex.
- This results into a partition of $\mathcal{S}$ into subsets $\mathcal{S}_e$.
- For every edge e, the points in $\mathcal{S}_e$ are either kept assigned to e or every point of $\mathcal{S}_e$ is assigned to its closest endpoint of e.
- The assignment that minimizes the optimal transport cost is chosen.

Shape reconstruction algorithm

1. **Input** - point set $\mathcal{S} = \{p_1, \ldots, p_n\}$
2. Construct Delaunay triangulation $\mathcal{T}_0$ of $\mathcal{S}$.
3. Compute initial transport plan π_0 from $\mathcal{S}$ to $\mathcal{T}_0$.
4. Set $k = 0$.
5. **Repeat** steps 6-11 **Until** desired vertex count is obtained.
6. Pick best half-edge $e = (x_i, x_j)$ to collapse (simplification).
7. Create $\mathcal{T}_{k+1}$ by merging x_i onto x_j.
8. $\pi'_{k+1} := \pi_k$ with local reassignments (update transport).
9. Optimize position of vertices in the one-ring of x_i (vertex relocation).
10. $\pi_{k+1} := \pi'_{k+1}$ with local reassignments (update transport).
11. $k \rightarrow k + 1$.
12. Filter edges based on relevance (optional).

Simplification

- Collapsing an edge changes a triangulation $\mathcal{T}_k$ to a triangulation $\mathcal{T}_{k+1}$ and thus changes the cost by $\delta_k = W_2(\pi_{k+1}) - W_2(\pi_k)$.
- Since the goal is to minimize the increase in total cost, edge collapses are applied in increasing order of δ.
- Therefore, all feasible collapses are initially simulated and their associated δ is added to a dynamic priority queue sorted in increasing order.
- Edge collapse is done by repeatedly popping from the queue the next edge to collapse, performing the collapse, updating the transport plan and cost and updating the priority queue.
- Note that updating the transport involves only the edges in the one-ring of the removed vertex and updating the priority queue is required for edges incident to the modified one ring.

Collapsing edges

- A half-edge is called **collapsible** if its collpase creates neither overlaps nor fold-overs in the triangulation.
- Every edge is made collapsible by the following procedure: let (x_i, x_j) denote the edge we want to collapse. Let P_{x_i} denote the counter-clockwise oriented polygon formed by the one-ring of x_i and let K_{x_i} denote its kernel. An edge $(a, b) \in P_{x_i}$ is *blocking* x_j if the triangle (x_j, a, b) has clockwise orientation.
- The edge (a, b) is removed from P_{x_i} by flipping either (a, x_i) or (b, x_i). Note that one of these edges is flippable.
- The flipping of blocking edges is continued until there are no blocking edges in P_{x_i}.

Vertex Relocation

- The triangulations obtained by edge collapses have their vertices on the input points. However, the presence of noise and missing data make interpolated triangulations poorly adapted to recover features.
- To overcome this problem, vertex relocation is performed after every edge collapse.
- The square of the normal part of the W_2 cost associated with a vertex v of $\mathcal{T}$ is given by

$$\sum_{p_i \in \mathcal{S}_v} m_i ||p_i - v||^2 + \sum_{b \in \mathcal{N}_1(v)} \sum_{p_i \in \mathcal{S}_{(v,b)}} m_i ||p_i - q_i||^2.$$

The optimal position v^* of v is computed by equating the gradient of the above expression to zero.

Edge filtering

- The presence of noise and outliers can lead to a few undesirable solid edges in the triangulation.
- Therefore, the solid edges are eliminated based on a notion of relevance r_e, given by

$$r_e = \frac{M_e len(e)^2}{N(e, \mathcal{S}_e)^2 + T(e, \mathcal{S}_e)^2}.$$