AlignOT: An optimal transport based algorithm for fast 3D alignment with applications to cryogenic electron microscopy density maps

Aryan Tajmir Riahi[1], Geoffrey Woollard[1], Frédéric Poitevin[2], Anne Condon[1], and Khanh Dao Duc[1,3,*]

[1]Department of Computer Science, University of British Columbia, Vancouver, BC V6T 1Z4, Canada
[2] SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA
[3]Department of Mathematics, University of British Columbia, Vancouver, BC V6T 1Z4, Canada

Abstract

Aligning electron density maps from Cryogenic electron microscopy (cryo-EM) is a first key step for studying multiple conformations of a biomolecule. As this step remains costly and challenging, with standard alignment tools being potentially stuck in local minima, we propose here a new procedure, called *AlignOT*, which relies on the use of computational optimal transport (OT) to align EM maps in 3D space. By embedding a fast estimation of OT maps within a stochastic gradient descent algorithm, our method searches for a rotation that minimizes the Wasserstein distance between two maps, represented as point clouds. We quantify the impact of various parameters on the precision and accuracy of the alignment, and show that *AlignOT* can outperform the standard local alignment methods, with an increased range of rotation angles leading to proper alignment. We further benchmark *AlignOT* on various pairs of experimental maps, which account for different types of conformational heterogeneities and geometric properties. As our experiments show good performance, we anticipate that our method can be broadly applied to align 3D EM maps.

Introduction

Solving the 3D structures of biomolecules is key to their function and the mechanisms underlying biological processes. For this purpose, cryogenic electron microscopy (cryo-EM) has become in recent years the most used technique to solve structures [1]. One main advantage of this technique, in contrast with X-ray crystallography, is that it potentially allows various conformations (or 3D configurations) of the same molecule to be solved [2]. Once different conformations are obtained as 3D EM density maps (i.e., large 3D grids of voxels with different levels of intensities), aligning these maps in 3D space is needed to further compare them.

Efficient methods have been developed to align two protein structures [3, 4], assuming their atomic composition is known. In this case, aligning two conformational structures is tantamount to finding an optimal rigid body transformation (i.e. a combination of 3D translation and rotation) that can align homologous atoms. However, when density maps are only given, one cannot directly establish such a homology correspondence from voxel to voxel, so the same problem becomes more challenging as the grid size increases and with the computational cost of searching over all possible rigid body transformations.

To solve the rigid body alignment problem for 3D cryo-EM density maps, standard approaches use various algorithms to maximize correlation [5, 6, 7]. More recently, Han et al. introduced a new method, which relies on representing the maps as sets of unit vectors before performing alignment [8]. Overall, both the choice of the metric to optimize, as well as the representation of the maps, can play important roles in getting a successful alignment. In this paper, we introduce a novel approach, called *AlignOT*, that uses a point-cloud representation of 3D maps, and minimizes the so-called Wasserstein distance between two maps with a stochastic gradient algorithm. This non-Euclidean distance is associated with the theory of Optimal Transport (OT) [9], with recent advances that make tractable the computation of transport-based distances [10, 11]. After describing the procedure of *AlignOT* in detail, we present the results of our experiments that quantify the precision and accuracy of this method, and benchmark it on a set of representative experimental

---

* corresponding author: kdd@math.ubc.ca

maps. Overall, the good performance of *AlignOT* with respect to standard local alignment methods suggests that our method can be broadly applied, as an alternative for aligning 3D EM maps with a non-Euclidean metric. We finally discuss the potential limitations of *AlignOT*, and connections with other recent methods and problems in optimal transport that would help to further improve and generalize it.

Material and Methods

Point cloud representation of EM maps

To represent voxelized cryo-EM maps, we use the topology representing network algorithm (TRN) [12], which reduces a map of d^3 voxels to a point cloud, i.e., a set of n points $\in \mathbb{R}^3$. Briefly, the voxelized maps (typically $\sim 100^3$ to $\sim 500^3$ voxels in the 3D grid) are first thresholded with a noise floor, to set low intensity regions to zero, and normalized to a probability mass function $\mathcal{P}$ over the grid points. n points $(\mathbf{r}_i(0))_{i=1\ldots n}$ are initially sampled from $\mathcal{P}$, and then updated in parallel for t_f rounds by taking weighted steps, according to the following equations [12]

$$\mathbf{r}_i(t+1) = \mathbf{r}_i(t) + \varepsilon(t)\exp[-k_i/\lambda(t)](\mathbf{r}_t - \mathbf{r}_i(t)), \tag{1}$$

$$\lambda(t) = \lambda_0\Big(\frac{\lambda_f}{\lambda_0}\Big)^{t/t_f}, \tag{2}$$

$$\varepsilon(t) = \varepsilon_0\Big(\frac{\varepsilon_f}{\varepsilon_0}\Big)^{t/t_f}, \tag{3}$$

where $\mathbf{r}_t$ is a single grid point sampled from $\mathcal{P}$, k_i denotes the rank of point $\mathbf{r}_i(t)$, by its Euclidean distance to $\mathbf{r}_t$, and $\lambda_0, \varepsilon_0, \lambda_f, \varepsilon_f, t_f$ are hyperparameters. In practice, we used $\varepsilon_0 = 0.3$ and $\varepsilon_f = 0.05$ for the initial and final step sizes, $\lambda_0 = 0.005 \times n$ and $\lambda_f = 0.5$ for the initial and final scaling ranks, and $t_f = 8 \times n$ for the total number of steps (adapted from [13]).

Optimal Transport and Wasserstein distance

To compare the point cloud representations of EM maps, we use a non-Euclidean metric that derives from the theory of Optimal Transport [9]. For two given point clouds, $\mathbf{A} = \{a_1, \ldots, a_n\}$ and $\mathbf{B} = \{b_1, \ldots, b_n\}$, we define a cost matrix $C_{i,j} = d(a_i, b_j)^2$, where d is the Euclidean distance. The *entropy regularized* 2-*Wasserstein distance* between $\mathbf{A}$ and $\mathbf{B}$, denoted by $\mathcal{W}_{2,\epsilon}(\mathbf{A}, \mathbf{B})$, is then defined as

$$\mathcal{W}_{2,\epsilon}(\mathbf{A}, \mathbf{B}) = [\min_{P \in \mathbb{R}_+^{n \times n}} \sum_{i,j=1}^{n} C_{i,j} P_{i,j} + \epsilon H(P)]^{1/2}, \tag{4}$$

$$\text{s.t.} \quad P.\mathbb{1} = P^T.\mathbb{1} = \mathbb{1}/n$$

where $\epsilon \in \mathbb{R}_+$ is the *regularization parameter* and the entropy $H(P)$ is given by

$$H(P) = \sum_{i,j=1}^{n} P_{i,j} \log P_{i,j}. \tag{5}$$

The minimizer of equation (4) is called the *transport plan*. For the rest of the Methods section, we will simply denote the Wasserstein distance as $\mathcal{W}_{2,\epsilon}$, and $P_{i(a),i(b)}$ as $P_{a,b}$, where $i(a)$ and $i(b)$ are the indices of the two points a and b in $\mathbf{A}$ and $\mathbf{B}$, respectively.

AlignOT: Algorithm for 3D map alignment

To align two 3D maps from their point cloud representations $\mathbf{A}$ and $\mathbf{B}$, we solve the optimization problem

$$q_{\text{opt}} = \operatorname*{argmin}_{q \in \mathbb{H}} \mathcal{W}_{2,\epsilon}(R_q(\mathbf{A}), \mathbf{B}), \tag{6}$$

where q is a quaternion (defined over the quaternion space $\mathbb{H}$), that we identify to a 3D rotation R_q in $SO(3)$, so $R_q(\mathbf{A}) = \{R_q(a_i)|a_i \in \mathbf{A}\}$. We explain in the Results section and Appendix A why we can only consider rotations and ignore translations to solve the general alignment problem, and provide more details on the identification of q to R_q in Appendix B. Our stochastic gradient descent procedure to solve (6), called *AlignOT*, is detailed in Algorithm 1. At each iteration, the algorithm updates q from the transport plan P between $R_q(\mathbf{A})$ and $\mathbf{B}$ as follows: After sampling one point $a \in R_q(\mathbf{A})$, we evaluate $\pi(a) = \mathrm{argmax}_{b \in \mathbf{B}} P_{a,b}$, and compute the gradient in q associated with $d(\pi(a), a)^2$, where d is the Euclidean distance, to update q. To compute the transport plan, we apply the Sinkhorn algorithm [10], with the initial vectors set as the outputs of the previous iteration. In practice, we also set the convergence condition $\|d(\pi(a), a)^2\| < \delta$ (where $\delta > 0$), that stops the algorithm before the maximum number of iterations. The hyperparameters of this procedure are the learning rate α associated with gradient descent, the regularization parameter ϵ associated with the Wasserstein distance, and a threshold δ associated with the number of iterations. In all our experiments, we set $\epsilon = 100$, $\delta = 10^{-10}$, and the maximum number of iterations equal to 500.

Algorithm 1 *AlignOT*: 3D density maps alignment with SGD using unit quaternions and Wasserstein distance

---

    Input two 3D density maps $\mathcal{A}, \mathcal{B}$, number of sampled points $n \in \mathbb{R}$, learning rate $\alpha \in \mathbb{R}$, regularization parameter $\epsilon \in \mathbb{R}$, maximum number of iterations $L \in \mathbb{N}$, and gradient threshold $\delta \in \mathbb{R}$

1: Sample two sets of n points $\mathbf{A}, \mathbf{B} \subset \mathbb{R}^3$ from $\mathcal{A}, \mathcal{B}$ respectively, using TRNs
2: $q = 1 + 0i + 0j + 0k$
3: $G = \alpha^2$
4: **while** not converged **and** the number of iterations is at most L **do**
5:     Compute $R_q(\mathbf{A})$
6:     Compute P to be the OT plan matrix between $R_q(\mathbf{A})$ and $\mathbf{B}$
7:     Randomly select $a \in R_q(\mathbf{A})$
8:     $b = \pi(a)$
9:     $f(q) = d(R_q(a), b)^2$ (where d is Euclidean distance in $\mathbb{R}^3$)
10:     $G = G + \|\nabla f(q)\|^2$
11:     $q = q - \frac{\alpha}{\sqrt{G}} \times \nabla f(q)$
12:     $q = \frac{q}{\|q\|}$
13: **end while**
14: **return** q

Implementation

We implemented *AlignOT* in Python 3.6.4. To sample a point cloud representation of an EM map using TRN, we adapted code from ProDy [14]. We used the NumPy package for matrix operations and POT's implementation of the Sinkhorn algorithm, which was modified to set the initial vectors (instead of initializing with uniform vectors). Our code is available in this GitHub repository.

Datasets

We tested *AlignOT* on various publicly available EM maps available from the EMDB, as listed in Table 1, and shown in Figure 1. For experiments that involved aligning a pair of distinct maps, we also used the corresponding pair of structures taken from the PDB [15], and used *MM-align* [3] to set a ground truth alignment, before converting the structures into EM maps using the function `molmap` in Chimera X [16]. All the datasets used in this study are available at this OSF page.

Results

Formulation of the alignment problem for EM maps with point cloud representation

We present here a new procedure, called *AlignOT*, that aligns 3D density maps from cryo-EM. Assuming that the EM maps are represented by two 3D point clouds $\mathbf{A} = \{a_1, \ldots, a_n\}$ and $\mathbf{B} = \{b_1, \ldots, b_n\}$, and for a given distance function d defined over the space of point clouds, the problem of aligning these maps consists of finding a rigid body transformation that minimizes the objective function

$$\mathcal{L}_d(R, T) = d(\text{move}_{R,T}(\mathbf{A}), \mathbf{B})^2, \tag{7}$$

where $\mathcal{L}_d(R, T)$ is defined over rotation matrices $R \in SO(3)$ and translation vectors $T \in \mathbb{R}^3$, and the operator $\text{move}_{R,T}(\mathbf{A})$ is defined as

$$\text{move}_{R,T}(\mathbf{A}) = \{Ra_i + T | a_i \in \mathbf{A}\}. \tag{8}$$

As the choice of d influences both the accuracy and the computational cost of the solution to the rigid body alignment problem, we here use the 2-Wasserstein distance, associated with the theory of Optimal Transport [9]. This distance can be used to compute distances between probability distributions, and is applied here more specifically for two distributions of 3D point clouds of same size. To efficiently evaluate this distance, we consider a regularized version (see equation (4) in the Methods section), denoted $\mathcal{W}_{2,\epsilon}$. Besides, given the centers of mass $\bar{a} = \frac{1}{n}\sum_{i=1}^{n} a_i$ and $\bar{b} = \frac{1}{n}\sum_{i=1}^{n} b_i$, and the centered point clouds $\mathbf{A_c} = \{a_{c_i} = a_i - \bar{a} | a_i \in \mathbf{A}\}$ and $\mathbf{B_c} = \{b_{c_i} = b_i - \bar{b} | b_i \in \mathbf{B}\}$, we can show that the optimal translation of the objective function (7) is

$$T_{\text{opt}} = \bar{b} - R_{\text{opt}}\bar{a}, \tag{9}$$

where

$$R_{\text{opt}} = \underset{R \in SO(3)}{\text{argmin}} \; \mathcal{W}_{2,\epsilon}(R(\mathbf{A_c}), \mathbf{B_c}). \tag{10}$$

Thus, the search for an optimal rigid body transformation in (10) can be simplified to rotations after matching the centers of mass of $\mathbf{A}, \mathbf{B}$, leading to equation (6) of the Methods section . We provide a detailed proof of this result in Appendix A.

AlignOT optimization procedure and complexity

To search for an optimal rotation that minimizes the Wasserstein distance between two 3D point clouds, we use an iterative Stochastic Gradient Descent algorithm, detailed in Algorithm 1. Basically, the algorithm aims to improve the alignment at each iteration, by rotating a map according to an assignment provided by the evaluation of the Wasserstein distance, with quaternions used to represent the 3D rotations (see Appendix B for a full description). Since EM maps are defined on a 3D voxelized grid, we also first convert the maps into point clouds, using the TRN algorithm [12]. We cover the application of the TRN application, the Wasserstein distance and the optimization procedure, as well as its implementation in the Material and Methods sections .

The complexity of our method (summarized in Algorithm 1) can be evaluated as follows: Assume that L is the maximum number of iterations, n the size of the point cloud, and ϵ the regularization parameter of the $\mathcal{W}_{2,\epsilon}$ distance. Each iteration of this algorithm consists of three steps. First, rotating one point cloud, second, computing the OT plan matrix, and last, sampling a random point and computing the gradient. Rotating a point cloud requires computing the coordinates of each point after rotation and takes $O(n)$ time, where n represents the number of points. To compute the OT plan matrix we use the Sinkhorn algorithm [11] which solves this problem in $O(n^2 \log n\epsilon^{-3})$ time, where ϵ is the regularization parameter. Finally, the gradient step takes $O(1)$ time. Overall, the most time-consuming part of this algorithm is computing the OT plan matrix at each iteration. Therefore, the overall time complexity of the algorithm is $O(n^2 L \log n\epsilon^{-3})$.

4

Alignment between maps can be obtained by minimizing the Wasserstein distance

To evaluate our method, we first tested *AlignOT* on aligning two point clouds that differ by a rotation only. To do so, we used an experimental map of a ribosome from EMDB 1717, shown in Figure 1a (ribosome structures are broadly studied in cryo-EM [2, 17]). First, we sampled two clouds of 500 points, and applied a rotation defined in its axis-angle representation by an arbitrary axis, and an angle $\theta = 20°$. Figure 2a illustrates how the moving point cloud gets closer to the targeted one over the iterations of the algorithm, until the convergence criterion is reached. To confirm this visual impression, we repeated the procedure with different initial angles $\theta \in \{10°, 30°, 50°, 70°\}$. The corresponding Wasserstein distance obtained across the iterations is shown in Figure 2b, with all the four trajectories converging to the same value and resulting in a successful alignment. However, we also observed that as θ increases, it takes more iterations for the algorithm to converge, with longer periods of slow variations at the beginning of the procedure, suggesting that this alignment can only be achieved within a certain range of θ. Aside from the initial angle, we also studied in Figure 2c how the size of the point cloud can affect the convergence plot. For $\theta = 50°$, decreasing the point cloud size from 500 to 250 leads to a potential loss of precision (with the Wasserstein distance between two aligned point clouds increasing from ~ 115 to 130), but with faster convergence from the algorithm (from $\sim$ 200 to 50 iterations to converge), indicating some trade-off between accuracy and speed.

To interpret these results, we further plotted in Figure 2d how the Wasserstein distance varies on average (after sampling different point clouds of same size 250 and 500), as a function of θ (and same rotation axis). In addition to the global minimum achieved for $\theta = 0$, another local minimum was detected at $\theta = 180°$, separated by a peak around 90°. While the sharp decrease observed towards the global minimum suggests that the optimization procedure can converge well for $\theta \leq 60°$, it is also possible that the algorithm does not converge to the global minimum above this range. Besides, the higher variability obtained from sampling point clouds of size 250, compared with 500, suggests that the final alignment may be less accurate in this case. While the same fixed rotation axis was considered in the previous experiments (with different values of θ), we finally evaluated the probability to successfully align the maps for initial rotations of fixed angle θ (45, 60, 75 and 90°), and across different axes covering half of the sphere $\mathcal{S}^2$. Upon mapping the axes on the planar disk in Figure 2e, we found local regions of poorer alignment that match with local minima of the Wasserstein distance. These results also confirm the existence of a limiting range within which the method can align two maps. While a successful alignment is overall obtained for $\theta = 45°$, the maps get partially aligned in different regions of the disks for 60° (see Figure 2e), with the performance worsening as θ increases (see Figure 3).

Benchmarking performance and accuracy of *AlignOT* with a pair of identical maps

We next focused on point cloud size and the range of convergence, and quantified their impact on the accuracy and computational cost of *AlignOT*. We first considered the same map and alignment task with $\theta = 20°$ as previously, with different point cloud sizes n (from 50 to 1000). The results obtained with a standard workstation are shown in Table 2, and confirm the existence of a trade-off between accuracy and speed, with the alignment improving as n increases, but with a larger runtime that goes from a few seconds for $n = 50, 100, 200$ to approximately a minute for $n = 1000$. While the accuracy of the alignment remains poor for $n = 50$ and 100 (with an average error of 12.6° and 7.72° respectively), it significantly improves for $n = 500$ with $2.35 \pm 1.23°$ error observed, with a runtime suggesting that *AlignOT* can be used in practice on standard density maps. We also noted that *AlignOT* runs slower in comparison with Chimera and Chimera X's alignment function `fitmap`, which performs a steepest ascent optimization to align maps according to their overlapping score [5, 16] (0.3 s on average). On the other hand, we show in our next experiments that *AlignOT* outperforms `fitmap` in accuracy as θ increases.

More precisely, we determined the range of θ within which the method converges, and compared our method with `fitmap`. As shown in Figure 2c, we found that *AlignOT* can cover a wider range that extends to 75°, while `fitmap` starts failing at approximately 45° (see also Figure 6 for a visualization of some representative alignments obtained). Beyond this value, *AlignOT* leads to some variable results up until 100°, due to the stochasticity of the algorithm. It then converges towards the other local minimum found

for the Wasserstein distance observed in Figure 2d, with a difference of $\sim 175°$ from the true alignment. In practice, Chimera also provides a global search option that randomly generates different initial instances of the rotated map, and keeps the best alignment to the target map using `fitmap`. To estimate the potential gain from using $AlignOT$, we generated 500 random initial placements, and computed the rates of successful alignments from $AlignOT$ and `fitmap`'s local search, to be 15.6% and 2.2%, respectively. In this context, $AlignOT$ thus reduces the number of initial placements needed in the global search by a factor of 7.6 (as $(1 - 0.022)^{7.6} \approx 1 - 0.156$).

Benchmarking *AlignOT* with conformationally heterogeneous pairs

We finally tested $AlignOT$ on pairs of distinct maps, as this reflects how the method should be applied in practice. More precisely, we considered six pairs, listed in Table 1 and shown in Figure 1, that all account for different conformations of a protein or complex structure, and that were recently used to evaluate another alignment method [8]. As illustrated in Figure 5a, we first used the molecular structures associated with the maps to define a ground truth alignment using the structure-based alignment function *MM-align* [3] (see also the Datasets section). This ground truth was used to evaluate the performance of $AlignOT$, as well as Chimera's `fitmap`, which we used for benchmarking against our method.

Upon testing for different values of point cloud size $n \in \{250, 500, 1000\}$ and initial rotation angle $\theta \in \{45°, 60°, 90°\}$, we recorded the angle differences between the ground truth alignment and the output using both `fitmap` and $AlignOT$, as reported in Table 3. Over the 18 cases tested, $AlignOT$ outperformed `fitmap` 13 times (72%), with some significant improvement observed in the majority of them (8 cases with an average improvement of more than $20°$). For example, the improvement obtained in the pair ID 4 ($\theta = 45°$) is visualized in Figure 5b, and shows how the presence of some symmetries (observed in the figure at the bottom of the molecule) can lead `fitmap` to misalign maps. Our experiments also confirm the improvement of the range of convergence using $AlignOT$, which we illustrate in Figure 5c with the pair ID 5 (with good alignment obtained at $60°$, contrary to `fitmap`). Besides, over the 5 cases where `fitmap` outperformed $AlignOT$, either the difference between the two methods is marginal (less than $2°$, for ID 1, 3, 6, and $\theta = 45°$), or both methods perform poorly (ID 2 with $\theta = 60°, 90°$). With the exception of the pair ID 2 (with poor performance from both methods), we found that the best alignment yields a Wasserstein distance that is the closest to that of the ground truth alignment, regardless of the method associated with it. These results thus suggest that the Wasserstein distance, which determines the objective function of our method, is a more appropriate metric to use than the Euclidean norm. Finally, our experiments highlight that increasing point size cloud does not necessarily lead to the best alignment from $AlignOT$, as a result of the trade-off between performance and speed that we also previously observed. Such a phenomenon was observed for the pair ID 6 ($\theta = 45°$), with a better alignment for $n = 250$ than $n = 1000$, as illustrated in in Figure 5d. As shown by the Wasserstein distance plots, the convergence is in this case significantly slower for $n = 1000$, which leads the algorithm to stop before the convergence condition gets achieved.

Discussion

In this paper, we present $AlignOT$, a new method for aligning cryo-EM density maps that relies on minimizing the Wasserstein distance between sampled point clouds. As shown in our experiments, $AlignOT$ is scalable to the typical size of density maps, with a good compromise between accuracy and speed that can be achieved upon tuning the point cloud size. Our method can thus be used to quickly align maps that come from different conformations of the same protein or complex. In particular, optimizing for a transport-based metric, instead of other common metrics (e.g. overlap, correlation), allows $AlignOT$ to generally outperform the standard local optimization method implemented in Chimera. Interestingly, the Wasserstein (or Earth-mover) distance was used in other applications in cryo-EM and tomographic projections (e.g. in interpolation or clustering [18, 19, 20, 21]), as its natural interpretation as the cost of displacing mass between two distributions makes it appropriate to compare volume-objects. While the choice of the TRN algorithm to generate point clouds is justified by its previous use to represent molecular structures [22], it can also be replaced by any other point cloud generation method, and it would be interesting to explore how to possibly improve our method on this aspect (in particular, we could for example use Vector Quantization, as it also has been used for approximating EM maps [23, 24]).

In the more general context of solving a rigid body alignment problem, *AlignOT* can be seen as a variant of the Iterative Closest Point method (ICP) framework [25], that consists of iteratively moving the point clouds according to the best way to match them. Among the many different variants of the ICP, Grave *et al.* similarly employed the Wasserstein distance to align language models [26]. Compared with their approach, our method differs with the use of quaternions in the gradient evaluation, as they make it possible to efficiently represent and manipulate 3D rotations. In addition, our modified version of the Sinkhorn algorithm, which evaluates the transport plan between the point clouds, provides a more rapid computation of the transport plan at each iteration, by simply reusing the vectors estimated at the previous iteration.

While Chimera's `fitmap` or *VESPER* [8] support a global search option (with multiple initializations of the local search function), we have not at this stage implemented a similar option for *AlignOT*, which also limits the direct comparison with these global methods. In addition, it also be interesting to study how to further optimize the learning rate and cloud size to improve the method. Yet, as our comparative experiments indicate an improvement of the range of convergence, as well as the number of random placements leading to a successful alignment, they suggest that *AlignOT* can potentially provide a useful alternative from the aforementioned current tools, that could also in the future be implemented as a plug-in for Chimera (as in [18]). Another current limitation is that the optimization problem assumes that the maps to align represent different conformations of a molecule, and thus carry approximately the same total density. However, it is also important in the context of Cryo-EM to consider the case (which *VESPER* does), of fitting two maps of different sizes, with one representing only a part of the other [6, 27]. In this case, since we cannot simplify, as in Appendix A , the rigid body alignment problem into an optimization over the rotations only, our method does not apply directly. However, as one can naturally formulate this problem in our framework as a problem of *unbalanced*, or *partial* Optimal Transport [28], and with the recent development of computational methods to solve it [29, 30], it makes such a generalization of *AlignOT* a promising future direction to pursue.

Acknowledgments

References

[1] Bai XC, McMullan G, Scheres SH. How Cryo-EM is revolutionizing structural biology. Trends in Biochemical Sciences. 2015;40(1):49–57.

[2] Poitevin F, Kushner A, Li X, Dao Duc K. Structural heterogeneities of the ribosome: new frontiers and opportunities for cryo-EM. Molecules. 2020;25(18):4262.

[3] Mukherjee S, Zhang Y. MM-align: a quick algorithm for aligning multiple-chain protein complex structures using iterative dynamic programming. Nucleic Acids Research. 2009;37(11):e83–e83.

[4] Moreta LS, Al-Sibahi AS, Theobald D, Bullock W, Rommes BN, Manoukian A, et al. A probabilistic programming approach to protein structure superposition. In: 2019 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB). IEEE; 2019. p. 1–5.

[5] Pettersen EF, Goddard TD, Huang CC, Couch GS, Greenblatt DM, Meng EC, et al. UCSF Chimera—a visualization system for exploratory research and analysis. Journal of Computational Chemistry. 2004;25(13):1605–1612.

[6] Evans P, McCoy A. An introduction to molecular replacement. Acta Crystallographica Section D: Biological Crystallography. 2008;64(1):1–10.

[7] Kawabata T. Multiple subunit fitting into a low-resolution density map of a macromolecular complex using a gaussian mixture model. Biophysical journal. 2008;95(10):4643–4658.

[8] Han X, Terashi G, Christoffer C, Chen S, Kihara D. VESPER: global and local cryo-EM map alignment using local density vectors. Nature communications. 2021;12(1):1–12.

[9] Peyré G, Cuturi M, et al. Computational optimal transport: With applications to data science. Foundations and Trends® in Machine Learning. 2019;11(5-6):355–607.

[10] Cuturi M. Sinkhorn distances: Lightspeed computation of optimal transport. Advances in neural information processing systems. 2013;26.

[11] Altschuler J, Weed J, Rigollet P. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. arXiv preprint arXiv:170509634. 2017;.

[12] Martinetz T, Schulten K. Topology representing networks. Neural Networks. 1994;7(3):507–522.

[13] Zhang Y, Krieger J, Mikulska-Ruminska K, Kaynak B, Sorzano COS, Carazo JM, et al. State-dependent sequential allostery exhibited by chaperonin TRiC/CCT revealed by network analysis of Cryo-EM maps. Progress in Biophysics and Molecular Biology. 2021;160:104–120. doi:10.1016/j.pbiomolbio.2020.08.006.

[14] Zhang S, Krieger JM, Zhang Y, Kaya C, Kaynak B, Mikulska-Ruminska K, et al. ProDy 2.0: increased scale and scope after 10 years of protein dynamics modelling with Python. Bioinformatics;(20):3657–3659. doi:10.1093/bioinformatics/btab187.

[15] Berman H, Henrick K, Nakamura H, Markley JL. The worldwide Protein Data Bank (wwPDB): ensuring a single, uniform archive of PDB data. Nucleic acids research. 2007;35(suppl_1):D301–D303.

[16] Goddard TD, Huang CC, Meng EC, Pettersen EF, Couch GS, Morris JH, et al. UCSF ChimeraX: Meeting modern challenges in visualization and analysis. Protein Science. 2018;27(1):14–25.

[17] Kushner A, Petrov AS, Dao Duc K. RiboXYZ: A comprehensive database for visualizing and analyzing ribosome structures. bioRxiv. 2022;.

[18] Ecoffet A, Poitevin F, Dao Duc K. MorphOT: Transport-based interpolation between EM maps with UCSF ChimeraX. Bioinformatics. 2020;36(22-23):5528–5529. doi:10.1093/bioinformatics/btaa1019.

[19] Ecoffet A, Woollard G, Kushner A, Poitevin F, Dao Duc K. Application of transport-based metric for continuous interpolation between cryo-EM density maps. AIMS Mathematics. 2022;7(1):986–999.

[20] Zelesko N, Moscovich A, Kileel J, Singer A. Earthmover-based manifold learning for analyzing molecular conformation spaces. In: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI). IEEE; 2020. p. 1715–1719.

[21] Rao R, Moscovich A, Singer A. Wasserstein k-means for clustering tomographic projections. arXiv preprint arXiv:201009989. 2020;.

[22] Zhang Y, Krieger J, Mikulska-Ruminska K, Kaynak B, Sorzano COS, Carazo JM, et al. State-dependent sequential allostery exhibited by chaperonin TRiC/CCT revealed by network analysis of Cryo-EM maps. Progress in Biophysics and Molecular Biology. 2021;160:104–120.

[23] De-Alarcón PA, Pascual-Montano A, Gupta A, Carazo JM. Modeling shape and topology of low-resolution density maps of biological macromolecules. Biophysical Journal. 2002;83(2):619–632.

[24] Jonić S, Sorzano COS. Versatility of approximating single-particle electron microscopy density maps using pseudoatoms and approximation-accuracy control. BioMed Research International. 2016;2016.

[25] Besl PJ, McKay ND. Method for registration of 3-D shapes. In: Sensor fusion IV: control paradigms and data structures. vol. 1611. International Society for Optics and Photonics; 1992. p. 586–606.

[26] Grave E, Joulin A, Berthet Q. Unsupervised alignment of embeddings with Wasserstein procrustes. In: The 22nd International Conference on Artificial Intelligence and Statistics. PMLR; 2019. p. 1880–1890.

[27] Kawabata T. Rigid-body fitting of atomic models on 3D density maps of electron microscopy. Integrative Structural Biology with Hybrid Methods. 2018; p. 219–235.

[28] Chizat L, Peyré G, Schmitzer B, Vialard FX. Unbalanced optimal transport: Dynamic and Kantorovich formulations. Journal of Functional Analysis. 2018;274(11):3090–3123.

[29] Koehl P, Delarue M, Orland H. Physics approach to the variable-mass optimal-transport problem. Physical Review E. 2021;103(1):012113.

[30] Le K, Nguyen H, Nguyen K, Pham T, Ho N. On multimarginal partial optimal transport: Equivalent forms and computational complexity. In: International Conference on Artificial Intelligence and Statistics. PMLR; 2022. p. 4397–4413.

[31] Fischer N, Konevega AL, Wintermeyer W, Rodnina MV, Stark H. Ribosome dynamics and tRNA movement by time-resolved electron cryomicroscopy. Nature. 2010;466(7304):329–333.

[32] Bai Xc, Rajendra E, Yang G, Shi Y, Scheres SH. Sampling the conformational space of the catalytic subunit of human γ-secretase. elife. 2015;4:e11182.

[33] Bai Xc, Yan C, Yang G, Lu P, Ma D, Sun L, et al. An atomic structure of human γ-secretase. Nature. 2015;525(7568):212–217.

[34] Chen Q, She J, Zeng W, Guo J, Xu H, Bai Xc, et al. Structure of mammalian endolysosomal TRPML1 channel in nanodiscs. Nature. 2017;550(7676):415–418.

[35] Hirschi M, Herzik Jr MA, Wie J, Suo Y, Borschel WF, Ren D, et al. Cryo-electron microscopy structure of the lysosomal calcium-permeable channel TRPML3. Nature. 2017;550(7676):411–414.

[36] Wu J, Yan Z, Li Z, Qian X, Lu S, Dong M, et al. Structure of the voltage-gated calcium channel Ca v 1.1 at 3.6 Å resolution. Nature. 2016;537(7619):191–196.

[37] Wu J, Yan Z, Li Z, Yan C, Lu S, Dong M, et al. Structure of the voltage-gated calcium channel Cav1. 1 complex. Science. 2015;350(6267).

[38] Zhao J, Benlekbir S, Rubinstein JL. Electron cryomicroscopy observation of rotational states in a eukaryotic V-ATPase. Nature. 2015;521(7551):241–245.

[39] Zhao J, Beyrakhova K, Liu Y, Alvarez CP, Bueler SA, Xu L, et al. Molecular basis for the binding and modulation of V-ATPase by a bacterial effector protein. PLoS pathogens. 2017;13(6):e1006394.

[40] Verba KA, Wang RYR, Arakawa A, Liu Y, Shirouzu M, Yokoyama S, et al. Atomic structure of Hsp90-Cdc37-Cdk4 reveals that Hsp90 traps and stabilizes an unfolded kinase. Science. 2016;352(6293):1542–1547.

[41] Gordiyenko Y, Llácer JL, Ramakrishnan V. Structural basis for the inhibition of translation through eIF2α phosphorylation. Nature communications. 2019;10(1):1–11.

Figures

a

1717

b

5fn5

5a63

c

5w3s

5wpq

d

5gjw

3jbr

e

5vox

3j9t

f

5fwm

5fwl

g

6qg6

6qg5

Figure 1: 3D maps used in our experiments, visualized with Chimera [5]. **(a):** Density map used in our first experiment (EMDB:1717). **(b-g):** Pairs of maps used in the second set of experiments, representing different conformational states of a given complex/molecule (see also Table 1).
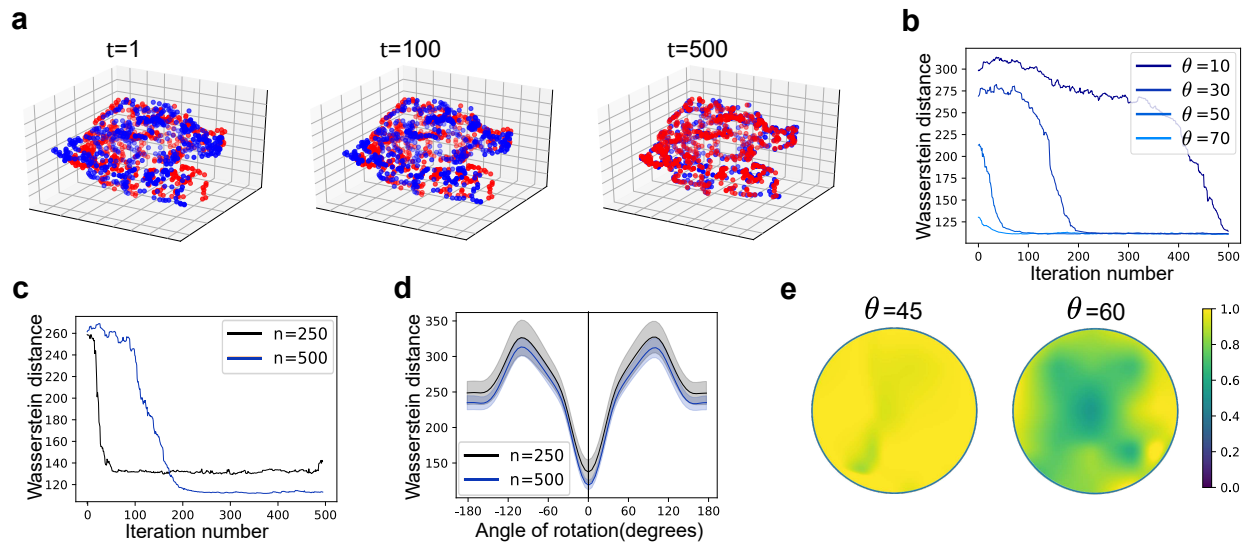
Figure 2: Alignment of two copies of a single map (from EMDB 1717) using *AlignOT* **(a):** We generated two copies of the map that differ from a 3D rotation, and applied *AlignOT* (see Algorithm 1). The figure shows the result of the procedure at different iterations ($t = 1, 100, 500$). The blue and red dots represent the rotated and target point clouds, respectively. **(b):** For different initial values θ ($= 10°, 30°, 50°, 70°$) of the angle difference between the two maps and same rotation axis, we plot the Wasserstein distance between the two target and rotated point clouds across the iterations of the algorithm, showing that they all converge to the same limit. **(c):** For the same initial rotation (with $\theta = 50°$), we plot the variations of the Wasserstein distance across iterations as in **(b)**, for point clouds of size $n = 250$ (black) and 500 (blue). **(d):** For the same rotation axis as in **(b)** and **(c)**, we plot the average Wasserstein distance between two sampled point clouds as a function of θ, for point clouds of size 250 (black) and 500 (blue). Error regions show the standard deviation from sampling 100 different point clouds for each angle. **(e):** Alignment success rate of *AlignOT* at fixed angles $\theta = 45$, and $60°$, and over the rotation axes that cover the upper hemishpere of S^2. Heatmaps show the percentage of outcomes that result in an alignment with error $\leq 5°$, where each point of the disk is the projection of the axis considered in S^2 (89 in total, with 20 runs for each).

Figure 3: Alignment success rate of $AlignOT$ at fixed angles $\theta = 45, 60, 75$, and $90°$, and over the rotation axes that cover the upper hemishpere of S^2. Heatmaps show the percentage of outcomes that result in an alignment with error $\leq 5°$, where each point of the disk is the projection of the axis considered in S^2 (89 in total, with 20 runs for each).
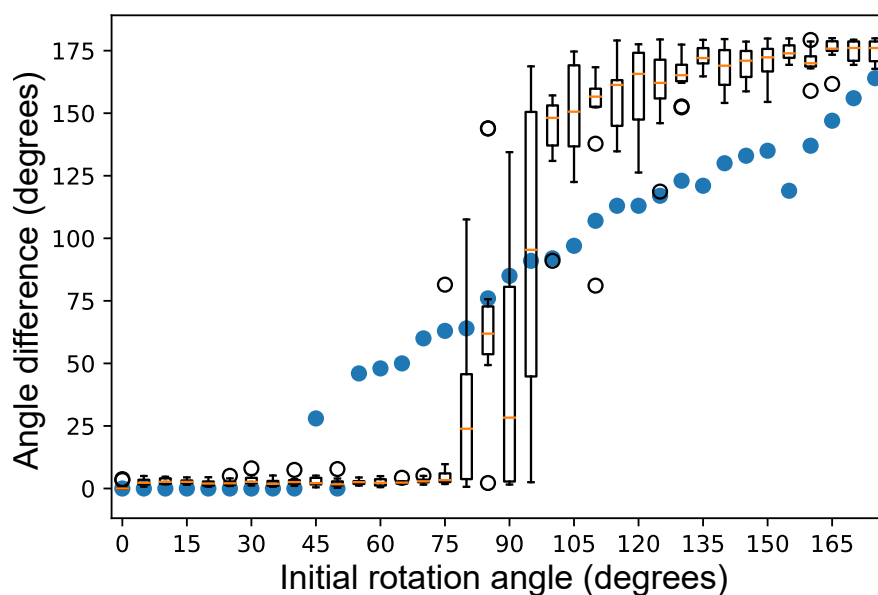
Figure 4: Comparison between *AlignOT* and Chimera's `fitmap` function: With a fixed axis and for a rotation angle difference θ between 0 and 180 degrees (x-axis), we set up a couple of maps from EMDB:1717 and run *AlignOT* and Chimera's `fitmap` function local search. The blue dots represent the error obtain in the alignment using `fitmap`, while the box plot (showing the minimum, first quartile, median, third quartile, and maximum) indicates the error of *AlignOT* over 10 runs.
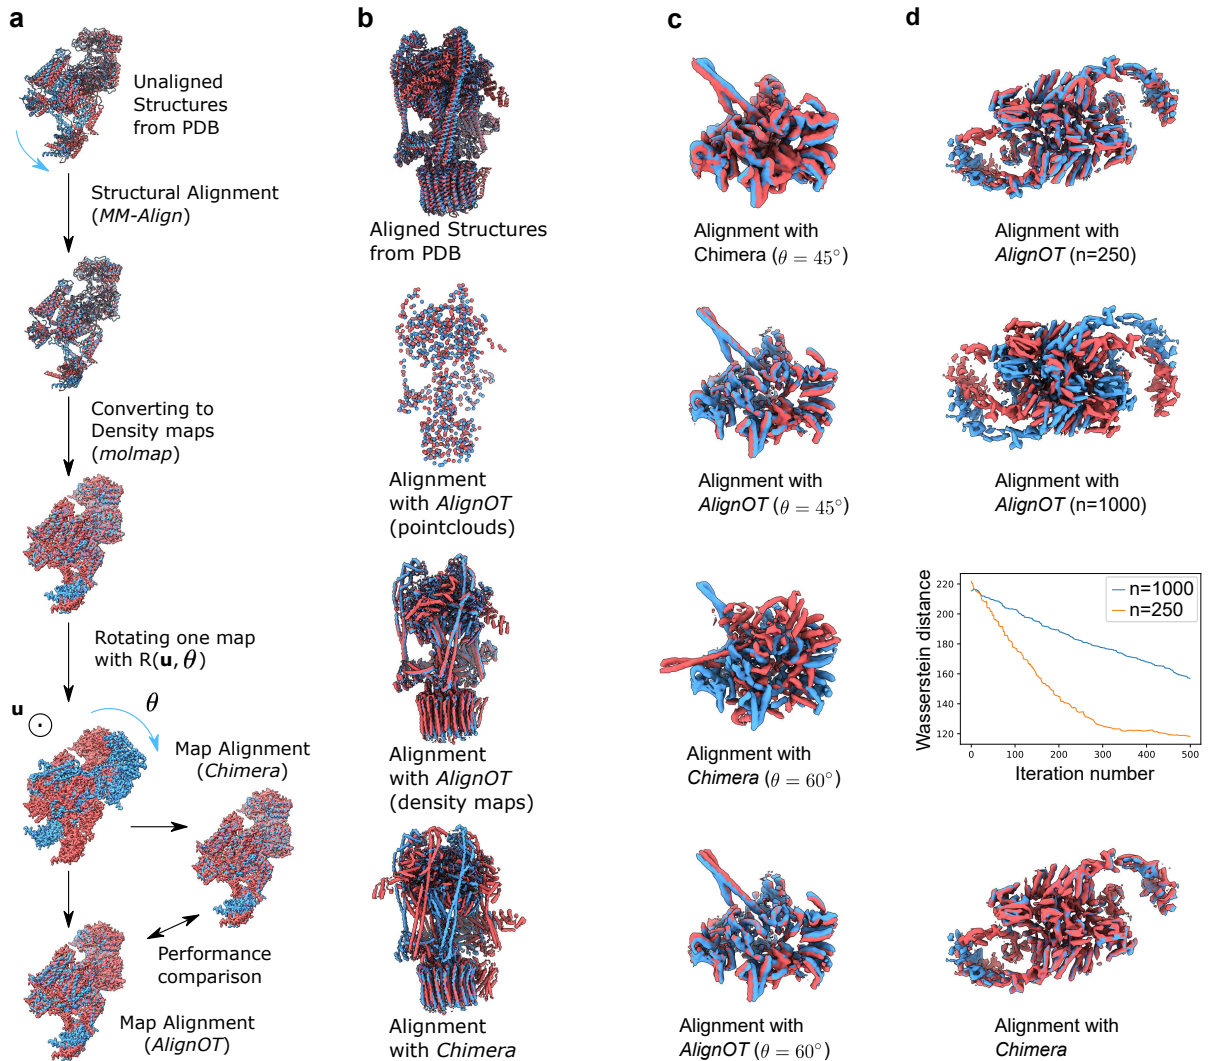
Figure 5: Performance of *AlignOT* with conformationally heterogeneous pairs **(a):** To quantify the performance of *AlignOT* on aligning different maps, we first aligned the corresponding PDB structures using *MM-align* [3] to set a ground truth alignment, and converted the aligned structures into cryo-EM density maps using the `molmap` command in Chimera [16], which were then used to benchmark Chimera's `fitmap` function and *AlignOT* with different initial rotations (see also Table 3). **(b):** Alignments obtained for the pair ID 4 (see Table 1) and initial rotation angle $\theta = 45°$. From top to bottom, we compare the alignment obtained with (1) *MM-Align* from PDB structures, (2) *AlignOT* using a point cloud size $n = 500$, with both point clouds and resulting density maps shown, and (3) Chimera's `fitmap` function (from the density maps). **(c):** Alignments obtained with the pair ID 5 show the increased range of convergence of *AlignOT* compared with `fitmap`. While *AlignOT* successfully align the maps for $\theta = 45$ and $60°$ (with the figure showing the result for $n = 1000$), it is not the case for `fitmap`, with a misalignment observed at $60°$. **(d):** From top to bottom, aligning pair ID 6 with $\theta = 45°$ shows a decrease in accuracy for *AlignOT* using a point cloud size $n = 1000$, compared with $n = 250$. We plot the corresponding the Wasserstein distance over the iterations of the algorithm, showing a larger final value for $n = 1000$, as the algorithm converges more slowly. As a result, this is one of the few cases where we found that Chimera produces a better alignment.

Figure 6: Visual comparison between *AlignOT* and Chimera's `fitmap` function. We set up a couple of maps from EMDB:1717 that differ from a rotation of fixed axis and angles $\theta = 60$ and $120°$, and ran *AlignOT* and Chimera's `fitmap` function local search. The figure shows the result of the procedure, with the final rotation (axis, angle) between the two maps.

Tables

Table 1: EM maps used in our experiments, also shown in Figure 1. The map from EMDB:1717 is used for aligning two copies of the same map. The other maps are grouped and aligned in our experiments as pair of different conformations of the same molecule, with the PDB structures used to obtain a ground truth (see Datasets and Results sections).

EMDB ID	Protein or Complex Name	Figure
1717 [31]	Ribosome	1 a

EMDB/PDB ID	Protein or Complex Name	Figure/ID
3240 / 5fn5 [32]	Human γ-secretase	1 b / 1
2677 / 5a63 [33]		
8881 / 5wpq [34]	TRPML	1 c / 2
8764 / 5w3s [35]		
9515 / 5gjw [36]	Voltage-gated calcium channel	1 d / 3
6475 / 3jb [37]		
6284 / 3j9t [38]	Yeast V-ATPase	1 e / 4
8724 / 5vox [39]		
3342 / 5fwm [40]	Hsp90-Cdc37-Cdk4	1 f / 5
3341 / 5fwl [40]		
4547 / 6qg5 [41]	eIF2B-eIF2	1 g / 6
4548 / 6qg6 [41]		

Table 2: Computational performance of *AlignOT*. For point clouds of size $n \in \{50, 100, 200, 500, 1000\}$, and for a fixed rotation of $\theta = 20°$, we ran *AlignOT* on EMDB:1717 50 times. We evaluate the accuracy of the method by computing the angle difference with the ground truth, and the associated runtime in seconds (mean and std). Experiments were run on an Intel(R) Core(TM) workstation with i7-7700HQ CPU @ 2.80GHz 2.81 GHz with 16.0 GB RAM.

n	Angle difference	Runtime (s)
50	12.60°±4.51°	2.36±0.42
100	7.72°±3.84°	4.66±0.57
200	3.85°±1.53°	4.14±0.37
500	2.35°±1.23°	16.36±1.37
1000	2.20°±0.96°	54.75±2.77

Table 3: Benchmarking of *AlignOT* on pairs of conformationally heterogeneous maps, as listed in Table 1. For each pair, we ran *AlignOT* and `fitmap` and reported the angle difference between the resulting algorithm and the ground truth (see Figure 5a). For each pair, we ran our test with initial rotation angles $\theta \in \{45, 60, 90°\}$, across 89 different axes covering half of the sphere $\mathcal{S}^2$ (with 20 runs for each), and point cloud sizes $n \in \{250, 500, 1000\}$ in *AlignOT*. Each table entry reports the mean and std observed (with sample size 20), with the best results for each pair highlighted in bold. To compute the regularized Wasserstein distance of final output for different methods, we sampled 500 points from each map and used the same regularization parameter $\epsilon = 100$.

ID	Angle	MMAlign $\mathcal{W}^2_{2,\epsilon}$	fitmap Angle diff. (°)	fitmap $\mathcal{W}^2_{2,\epsilon}$	AlignOT $n=250$ Angle diff. (°)	AlignOT $n=250$ $\mathcal{W}^2_{2,\epsilon}$	AlignOT $n=500$ Angle diff. (°)	AlignOT $n=500$ $\mathcal{W}^2_{2,\epsilon}$	AlignOT $n=1000$ Angle diff. (°)	AlignOT $n=1000$ $\mathcal{W}^2_{2,\epsilon}$
1	45	125.14 ±3.79	**0.97 ±0.71**	**125.36 ±4.15**	2.82 ±1.34	130.73 ±8.36	2.30 ±1.08	125.65 ±5.26	5.73 ±3.84	127.51 ±7.27
1	60	125.14 ±3.79	6.39 ±14.35	133.71 ±38.01	**2.79 ±1.32**	130.82 ±8.42	2.99 ±1.95	**126.19 ±5.22**	10.78 ±8.21	140.99 ±23.44
1	90	125.14 ±3.79	46.93 ±30.70	267.78 ±129.91	**9.22 ±25.99**	**140.72 ±40.58**	16.77 ±22.09	169.43 ±74.51	32.48 ±25.48	223.20 ±94.98
2	45	152.69 ±9.67	37.91 ±8.38	292.44 ±30.47	20.01 ±34.43	188.17 ±29.67	14.99 ±30.83	163.25 ±25.32	**11.06 ±26.68**	**149.06 ±21.79**
2	60	152.69 ±9.67	**55.33 ±5.79**	317.66 ±42.78	62.32 ±42.37	214.83 ±42.83	62.88 ±41.47	195.23 ±42.07	62.58 ±39.83	**188.42 ±45.09**
2	90	152.69 ±9.67	**85.21 ±5.42**	295.43 ±49.50	98.20 ±21.54	245.20 ±40.92	99.69 ±14.84	228.26 ±39.47	99.66 ±13.48	**221.37 ±40.41**
3	45	237.78 ±31.97	**3.56 ±10.48**	268.36 ±147.22	6.71 ±10.13	263.32 ±46.31	5.41 ±1.66	221.49 ±22.01	5.14 ±1.20	**201.35 ±11.38**
3	60	237.78 ±31.97	38.35 ±20.65	823.87 ±415.12	9.49 ±24.21	267.61 ±64.31	7.00 ±16.86	224.56 ±44.38	**5.13 ±1.20**	**201.86 ±12.05**
3	90	237.78 ±31.97	80.53 ±7.75	1623.24 ±521.41	37.25 ±66.03	330.59 ±157.46	30.94 ±61.25	281.16 ±144.28	**25.27 ±55.53**	**248.83 ±130.86**
4	45	413.84 ±35.57	31.65 ±10.33	**325.64 ±83.43**	10.45 ±2.67	422.41 ±44.29	10.05 ±1.94	408.48 ±30.38	**9.65 ±1.49**	401.14 ±21.32
4	60	413.84 ±35.57	47.21 ±11.27	435.61 ±126.95	10.75 ±6.04	422.46 ±42.70	10.07 ±4.09	407.52 ±30.02	**9.56 ±1.51**	**401.28 ±21.67**
4	90	413.84 ±35.57	81.79 ±6.69	655.52 ±170.82	33.51 ±57.08	440.99 ±65.55	29.54 ±52.75	**425.56 ±60.30**	**26.65 ±46.27**	432.05 ±85.33
5	45	104.93 ±1.53	2.02 ±5.41	111.62 ±15.84	5.94 ±15.79	106.87 ±3.95	3.77 ±11.68	104.93 ±2.66	**2.02 ±1.22**	**103.98 ±1.32**
5	60	104.93 ±1.53	26.10 ±19.97	138.80 ±30.50	9.37 ±28.47	107.42 ±5.51	6.64 ±24.80	105.38 ±4.53	**3.37 ±14.33**	**104.35 ±3.80**
5	90	104.93 ±1.53	60.52 ±16.78	162.30 ±23.35	17.454 ±45.22	108.72 ±7.94	26.27 ±58.88	108.56 ±9.91	**16.71 ±47.56**	**106.79 ±9.03**
6	45	119.92 ±2.91	**0.91 ±4.98**	**113.25 ±16.73**	4.26 ±1.58	120.77 ±5.34	7.54 ±4.20	122.15 ±5.49	14.53 ±6.93	138.26 ±11.07
6	60	119.92 ±2.91	14.73 ±19.64	155.34 ±58.90	**5.17 ±2.33**	**121.43 ±5.52**	11.50 ±7.26	128.46 ±11.18	22.12 ±10.46	161.75 ±20.83
6	90	119.92 ±2.91	43.43 ±27.87	204.40 ±57.11	**18.37 ±26.30**	**143.74 ±45.67**	31.68 ±25.94	178.16 ±60.72	46.43 ±22.89	258.30 ±45.43

Appendix

Appendix A: Optimal translation for the rigid body alignment problem

To minimize the loss function over rigid body transformations (translation and rotations), we can simplify the problem and restrict the search to rotations only as follows: For two point clouds $\mathbf{A} = \{a_1, \ldots, a_n\}$ and $\mathbf{B} = \{b_1, \ldots, b_n\}$, and upon introducing the centers of mass $\bar{a} = \frac{1}{n} \sum_{i=1}^{n} a_i$ and $\bar{b} = \frac{1}{n} \sum_{i=1}^{n} b_i$, and the centered coordinates $a_{c_i} = a_i - \bar{a}$ and $b_{c_i} = b_i - \bar{b}$, equation (7) yields

$$
\begin{aligned}
\mathcal{L}(R, T) &= \mathcal{W}_{2,\epsilon}(\text{move}_{R,T}(\mathbf{A}), \mathbf{B})^2 \\
&= \min_{P \in \mathbb{R}_+^{n \times n}} \sum_{i,j=1}^{n} \|Ra_i + T - b_j\|_2^2 P_{i,j} + \epsilon H(P) \quad (\text{s.t.} \quad \forall 1 \le j \le n : \sum_{i=1}^{n} P_{i,j} = \sum_{i=1}^{n} P_{j,i} = \frac{1}{n}) \\
&= \min_{P \in \mathbb{R}_+^{n \times n}} \sum_{i,j=1}^{n} \|Ra_{c_i} + R\bar{a} + T - b_{c_j} - \bar{b}\|_2^2 P_{i,j} + \epsilon H(P),
\end{aligned}
$$

where we used definition of the entropy regularized Wasserstein distance from equation (4). This further simplifies as

$$
\begin{aligned}
\mathcal{L}(R, T) &= \min_{P \in \mathbb{R}_+^{n \times n}} \sum_{i,j=1}^{n} \|Ra_{c_i} - b_{c_j}\|_2^2 P_{i,j} + \sum_{i,j=1}^{n} \|R\bar{a} + T - \bar{b}\|_2^2 P_{i,j} \\
&\quad + \sum_{i,j=1}^{n} (Ra_{c_i} - b_{c_j}).(R\bar{a} + T - \bar{b}) P_{i,j} + \epsilon H(P) \quad (11) \\
&= \min_{P \in \mathbb{R}_+^{n \times n}} \sum_{i,j=1}^{n} \|Ra_{c_i} - b_{c_j}\|_2^2 P_{i,j} + \sum_{i,j=1}^{n} \|R\bar{a} + T - \bar{b}\|_2^2 P_{i,j} + \epsilon H(P) \quad (12) \\
&= \min_{P \in \mathbb{R}_+^{n \times n}} [\sum_{i,j=1}^{n} \|Ra_{c_i} - b_{c_j}\|_2^2 P_{i,j} + \epsilon H(P)] + \|R\bar{a} + T - \bar{b}\|_2^2 \quad (13)
\end{aligned}
$$

where we used in (11) that $\sum_{i,j=1}^{n} P_{i,j} a_{c_i} = \frac{1}{n} \sum_{i=1}^{n} a_{c_i} = 0$ and $\sum_{i,j=1}^{n} P_{i,j} b_{c_j} = \frac{1}{n} \sum_{j=1}^{n} b_{c_j} = 0$. Also, we used the fact that $\sum_{i,j=1}^{n} P_{i,j} = 1$ in (12). The second term in eq. (13) is minimized for $T = \bar{b} - R\bar{a}$, i.e. the translation that aligns the two centers of mass. Therefore, to solve the rigid body alignment problem, we can assume that the distributions of points $\mathbf{A}$ and $\mathbf{B}$ are both centered at the origin, and that we only need to solve the optimization problem over the rotations in $SO(3)$, as stated in equation (10).

Appendix B: Quaternion Representation of 3D rotations

Quaternion Representation of 3D rotations and gradient formula

To formalize 3D rotations in $AlignOT$ we use the quaternion representation ($\mathbb{H}$). In this representation, given a point $a = (x, y, z) \in \mathbb{R}^3$ and a rotation with angle θ around axis $\vec{u} = (u_x, u_y, u_z)$, we form quaternions

$$
\begin{aligned}
q &= \cos \theta/2 + u_x \sin \theta/2 i + u_y \sin \theta/2 j + u_z \sin \theta/2 k, \\
p &= xi + yj + zk,
\end{aligned}
$$

where i, j, k are the basic quaternions such that

$$
i^2 = j^2 = k^2 = ijk = -1. \quad (14)
$$

Using equation (14) we compute the following term

$$x_{\mathrm{res}}i + y_{\mathrm{res}}j + z_{\mathrm{res}}k = qpq^*,$$

where $R_q(a) = (x_{\mathrm{res}}, y_{\mathrm{res}}, z_{\mathrm{res}}) \in \mathbb{R}^3$ is the coordinates of a after rotation, and q^* is the conjugate of q and is defined as

$$(q_0 + q_1i + q_2j + q_3k)^* = q_0 - q_1i - q_2j - q_3k.$$

We then define the absolute norm of q as

$$\|q\|^2 = qq^* = q_0^2 + q_1^2 + q_2^2 + q_3^2 \in \mathbb{R}.$$

In *AlignOT*, we also compute the gradient associated with the function $f(q) = f(q_0, q_1, q_2, q_3) = \|R_q(a) - b\|_2^2$, where $a, b \in \mathbb{R}^3$ as

$$\nabla f = \frac{\partial f}{\partial q_0} + \frac{\partial f}{\partial q_1}i + \frac{\partial f}{\partial q_2}j + \frac{\partial f}{\partial q_3}k,$$

where the explicit analytical formula of ∇f is provided in the Appendix B.

Analytical formula for the gradient update

In the previous section we defined gradient of a general function $f : \mathbb{H} \to \mathbb{R}$. In this sectionuse this definition to compute an explicit formula for the gradient of the function used in Algorithm 1 i.e. $f(q) = d(R_q(a), b)^2$, where d is Euclidean distance in $\mathbb{R}^3$, $a = (a_x, a_y, a_z), b = (b_x, b_y, b_z)$ are two points in $\mathbb{R}^3$, and $q = q_0 + q_1i + q_2j + q_3k$ is a quaternion. Using the definition of $R_q(a)$ in Appendix B to compute the coordinates of $R_q(a) = (a_{x_{res}}, a_{y_{res}}, a_{z_{res}})$, we have

$$a_{x_{res}}i + a_{y_{res}}j + a_{z_{res}}k = q(a_xi + a_yj + a_zk)q^* = f_i(q)i + f_j(q)j + f_k(q)k.$$

After some simplification, we can write f_i, f_j and f_k as

$$a_{x_{res}} = f_i(q) = q_1^2 a_x + 2q_1q_2a_y + 2q_1q_3a_z + q_0^2a_x + 2q_0q_2a_z - 2q_0q_3a_y - q_3^2a_x - q_2^2a_x,$$
$$a_{y_{res}} = f_j(q) = 2q_1q_2a_x + q_2^2a_y + 2q_2q_3a_z + q_0^2a_y + 2q_0q_3a_x - 2q_0q_1a_z - q_1^2a_y - q_3^2a_y,$$
$$a_{z_{res}} = f_k(q) = 2q_1q_3a_x + 2q_2q_3a_y + q_3^2a_z + q_0^2a_z + 2q_0q_1a_y - 2q_0q_2a_x - q_2^2a_z - q_1^2a_z.$$

To compute partial derivatives of f we use that

$$\frac{\partial f}{\partial q_l} = \frac{(a_{x_{res}} - b_x)^2 + (a_{y_{res}} - b_y)^2 + (a_{z_{res}} - b_z)^2}{\partial q_l}$$
$$= 2(f_i(q) - b_x)\frac{\partial f_i(q)}{\partial q_l} + 2(f_j(q) - b_y)\frac{\partial f_j(q)}{\partial q_l} + 2(f_k(q) - b_z)\frac{\partial f_k(q)}{\partial q_l}$$

for $l \in \{0, 1, 2, 3\}$, which yields

$$\frac{\partial f}{\partial q_0} = 2(f_i(q) - b_x)\frac{\partial f_i(q)}{\partial q_0} + 2(f_j(q) - b_y)\frac{\partial f_j(q)}{\partial q_0} + 2(f_k(q) - b_z)\frac{\partial f_k(q)}{\partial q_0}$$

$$= 4(f_i(q) - b_x)(q_0 a_x + q_2 a_z - q_3 a_y) + 4(f_j(q) - b_y)(q_0 a_y + q_3 a_x - q_1 a_z)$$

$$+ 4(f_k(q) - b_z)(q_0 a_z + q_1 a_y - q_2 a_x),$$

$$\frac{\partial f}{\partial q_1} = 2(f_i(q) - b_x)\frac{\partial f_i(q)}{\partial q_1} + 2(f_j(q) - b_y)\frac{\partial f_j(q)}{\partial q_1} + 2(f_k(q) - b_z)\frac{\partial f_k(q)}{\partial q_1}$$

$$= 4(f_i(q) - b_x)(q_1 a_x + q_2 a_y + q_3 a_z) + 4(f_j(q) - b_y)(q_2 a_x - q_0 a_z - q_1 a_y)$$

$$+ 4(f_k(q) - b_z)(q_3 a_x + q_0 a_y - q_1 a_z),$$

$$\frac{\partial f}{\partial q_2} = 2(f_i(q) - b_x)\frac{\partial f_i(q)}{\partial q_2} + 2(f_j(q) - b_y)\frac{\partial f_j(q)}{\partial q_2} + 2(f_k(q) - b_z)\frac{\partial f_k(q)}{\partial q_2}$$

$$= 4(f_i(q) - b_x)(q_1 a_y + q_0 a_z - q_2 a_x) + 4(f_j(q) - b_y)(q_1 a_x + q_2 a_y + q_3 a_z)$$

$$+ 4(f_k(q) - b_z)(q_3 a_y - q_0 a_x - q_2 a_z),$$

$$\frac{\partial f}{\partial q_3} = 2(f_i(q) - b_x)\frac{\partial f_i(q)}{\partial q_3} + 2(f_j(q) - b_y)\frac{\partial f_j(q)}{\partial q_3} + 2(f_k(q) - b_z)\frac{\partial f_k(q)}{\partial q_3}$$

$$= 4(f_i(q) - b_x)(q_1 a_z - q_0 a_y - q_3 a_x) + 4(f_j(q) - b_y)(q_2 a_z + q_0 a_x - q_3 a_y)$$

$$+ 4(f_k(q) - b_z)(q_1 a_x + q_2 a_y + q_3 a_z),$$

$$\nabla f = \frac{\partial f}{q_0} + \frac{\partial f}{q_1}i + \frac{\partial f}{q_2}j + \frac{\partial f}{q_3}k.$$