

—Supplementary Material—
**Multiview Regenerative Morphing
with Dual Flows**

Chih-Jung Tsai¹, Cheng Sun^{1,2}, and Hwann-Tzong Chen^{1,3}

¹ National Tsing Hua University

² ASUS AICS Department

³ Aeolus Robotics

In this supplementary document, we provide *i*) derivations of the optimization on rigid transformation flow and optimal transport flow using Sinkhorn divergence, and *ii*) additional morphing results rendered in different views.

1 Sinkhorn Divergence and Morphing Flows

Sinkhorn divergence [2, 3] is a popular metric for comparing distributions in the form of regularized optimal transport. Given two positive, unit-mass measures $\gamma, \beta \in \mathcal{M}^+(\mathcal{X})$ on some feature space \mathcal{X} , the Sinkhorn divergence with parameter ε is expressed as

$$\text{SD}_\varepsilon(\gamma, \beta) = \text{OT}_\varepsilon(\gamma, \beta) - \frac{1}{2}\text{OT}_\varepsilon(\gamma, \gamma) - \frac{1}{2}\text{OT}_\varepsilon(\beta, \beta), \quad (1)$$

where OT_ε is an entropy-regularized optimal transport [1, 4] that can be solved using the famous Sinkhorn algorithm. The optimal transport OT_ε is defined as

$$\text{OT}_\varepsilon(\gamma, \beta) = \min_{\pi_1=\gamma, \pi_2=\beta} \int_{\mathcal{X}^2} C \, d\pi + \varepsilon \text{KL}(\pi \mid \gamma \otimes \beta), \quad (2)$$

where KL is the Kullback-Leibler divergence. The coupled measure $\pi \in \mathcal{M}^+(\mathcal{X}^2)$ has two marginals (π_1, π_2) , and the optimization is to find the transport plan π that moves all the mass of γ toward β , subject to the constraints on the two marginals that $\pi_1 = \gamma$ and $\pi_2 = \beta$. The distance function C is defined as $C(x, y) = \|x - y\|^p$ and often referred to as the *p-Wasserstein* distance. Eq. (2) can be optimized by taking the Sinkhorn iterations with its dual form. For discretely sampled measures, such as $\gamma = \sum_{i=1}^N \gamma_i \Delta_{\mathbf{x}_i}$ and $\beta = \sum_{j=1}^M \beta_j \Delta_{\mathbf{y}_j}$, the optimized OT_ε can be expressed as

$$\text{OT}_\varepsilon(\{(\gamma_i, \mathbf{x}_i)\}, \{(\beta_j, \mathbf{y}_j)\}) = \sum_{i=1}^N \gamma_i \mathbf{u}_i + \sum_{j=1}^M \beta_j \mathbf{v}_j, \quad (3)$$

where (\mathbf{u}, \mathbf{v}) are the optimal dual vectors with

$$\begin{aligned}\mathbf{u}_i &= -\varepsilon \log \sum_{k=1}^M \exp \left(\log(\beta_k) + \frac{1}{\varepsilon} \mathbf{v}_k - \frac{1}{\varepsilon} C(\mathbf{x}_i, \mathbf{y}_k) \right), \\ \mathbf{v}_j &= -\varepsilon \log \sum_{k=1}^N \exp \left(\log(\gamma_k) + \frac{1}{\varepsilon} \mathbf{u}_k - \frac{1}{\varepsilon} C(\mathbf{x}_k, \mathbf{y}_j) \right).\end{aligned}\quad (4)$$

Therefore, the Sinkhorn divergence between γ and β in the discrete form can be computed as

$$\text{SD}_\varepsilon(\{(\gamma_i, \mathbf{x}_i)\}, \{(\beta_j, \mathbf{y}_j)\}) = \sum_{i=1}^N \gamma_i (\mathbf{u}_i - \mathbf{a}_i) + \sum_{j=1}^M \beta_j (\mathbf{v}_j - \mathbf{b}_j), \quad (5)$$

where $\mathbf{a}_i, \mathbf{b}_j$ are the optimal dual vectors that yield optimal $\text{OT}_\varepsilon(\gamma, \gamma)$ and $\text{OT}_\varepsilon(\beta, \beta)$, the regularization terms in Sinkhorn divergence. We can minimize Sinkhorn divergence by moving mass from \mathbf{x}_i to \mathbf{y}_j . This is achieved by taking the derivative $\partial_{\mathbf{x}_i} \text{SD}_\varepsilon = \gamma_i \nabla (\mathbf{u}_i - \mathbf{a}_i) = \gamma_i \nabla \Phi(\mathbf{x}_i)$, where Φ is defined as

$$\begin{aligned}\Phi(\mathbf{x}) &= -\varepsilon \log \sum_{j=1}^M \exp \left[\log(\beta_j) + \frac{1}{\varepsilon} \mathbf{v}_j - \frac{1}{\varepsilon} C(\mathbf{x}, \mathbf{y}_j) \right] \\ &\quad + \varepsilon \log \sum_{i=1}^N \exp \left[\log(\gamma_i) + \frac{1}{\varepsilon} \mathbf{a}_i - \frac{1}{\varepsilon} C(\mathbf{x}, \mathbf{x}_i) \right].\end{aligned}\quad (6)$$

Based on Eqs. (4) and (5), if we plug \mathbf{x}_i into Eq. (6) we can get $\Phi(\mathbf{x}_i) = \mathbf{u}_i - \mathbf{a}_i$ as expected.

In our experiment, the source and target shape representations \mathcal{S} and \mathcal{T} are the collections of their opacity information and the coordinates of the respective voxels. The representations can be expressed as sums of weighted Dirac masses

$$\mathcal{S} = \sum_{i=1}^{N^{\mathcal{S}}} \omega_i^{\mathcal{S}} \Delta_{\mathbf{x}_i^{\mathcal{S}}}, \quad \mathcal{T} = \sum_{j=1}^{N^{\mathcal{T}}} \omega_j^{\mathcal{T}} \Delta_{\mathbf{x}_j^{\mathcal{T}}}, \quad (7)$$

where $\omega_i^{\mathcal{S}}$ and $\omega_j^{\mathcal{T}}$ represent the normalized opacity collections α derived from $\mathcal{V}_\alpha^{\mathcal{S}}$ and $\mathcal{V}_\alpha^{\mathcal{T}}$. We aim to minimize Sinkhorn divergence $\text{SD}_\varepsilon(\{(\omega_i^{\mathcal{S}}, \mathbf{x}_i^{\mathcal{S}})\}, \{(\omega_j^{\mathcal{T}}, \mathbf{x}_j^{\mathcal{T}})\})$ between \mathcal{S} and \mathcal{T} using the rigid transformation flow and the optimal transport flow.

1.1 Rigid transformation flow

The rigid transformation flow produces an as-rigid-as-possible visual effect for morphing. Without a definite pose between two unrelated objects, the flow

estimates an optimal 6D transformation $\Psi \in \text{SE}(3)$, with rotation \mathbf{R} and translation \mathbf{z} , that best aligns the shapes of the two objects. The optimization can be formulated as

$$\hat{\Psi} = \arg \min_{\Psi} \text{SD}_{\varepsilon} (\{(\omega_i^S, \Psi(\mathbf{x}_i^S))\}, \{(\omega_j^T, \mathbf{x}_j^T)\}), \quad (8)$$

which is optimized through gradients descent. Considering Eq. (5), the differentiation of SD_{ε} with respect to Ψ is

$$\frac{1}{\omega_i^S} \partial_{\Psi} \text{SD}_{\varepsilon} (\{(\omega_i^S, \Psi(\mathbf{x}_i^S))\}, \{(\omega_j^T, \mathbf{x}_j^T)\}) = \nabla \Phi (\Psi(\mathbf{x}_i^S)) \cdot \frac{\partial \Psi(\mathbf{x}_i^S)}{\partial \Psi}, \quad (9)$$

where $\Phi(\cdot)$ is defined in Eq. (6) with specialization on $\{(\omega_i^S, \Psi(\mathbf{x}_i^S))\}$ and $\{(\omega_j^T, \mathbf{x}_j^T)\}$. Using the above equations, we can update every component of Ψ . In our experiment, we use *GeomLoss* [2] to calculate the Sinkhorn divergence, and *Pytorch* to compute the gradient and the partial differentiation in Eq. (9). Since the update on the rotation matrix \mathbf{R} does not guarantee to be a rotation matrix, we use singular value decomposition to normalize the updated \mathbf{R} in every iteration. After K iterations, we obtain an estimate $\hat{\Psi}$ and the rigid transformation flow \mathbf{f} :

$$\mathbf{f}(t) = t \cdot (\hat{\Psi}(\mathcal{S}) - \mathcal{S}). \quad (10)$$

$\hat{\Psi}(\mathcal{S})$ is also used for computing the optimal transport flow.

1.2 Optimal transport flow

The optimal transport flow deforms the source shape to the target shape using the gradient of Sinkhorn divergence. Consider a Dirac mass located at \mathbf{x}_i , with weight ω_i^S and belonging to the transformed source shape $\hat{\Psi}(\mathcal{S})$, the optimal transport flow related to \mathbf{x}_i can be written as

$$\mathbf{g}_i(t) = -t \cdot \partial_{\mathbf{x}_i} \text{SD}(\hat{\Psi}(\mathcal{S}), \mathcal{T}), \quad (11)$$

for some blending weight $t \in [0, 1]$. In our experiment, the flow deforms from rigidly-transformed source shape to target shape. The setting constrains the deformation to happen between aligned objects, instead of deforming directly from source to target that leads to a shattered shape. Therefore, we can formulate Eq. (11) as

$$\begin{aligned} \mathbf{g}_i(t) &= -t \cdot \partial_{\mathbf{x}_i} \text{SD}_{\varepsilon} (\{(\omega_i^S, \hat{\Psi}(\mathbf{x}_i^S))\}, \{(\omega_j^T, \mathbf{x}_j^T)\}) \\ &= -t \cdot \nabla \Phi (\hat{\Psi}(\mathbf{x}_i^S)), \end{aligned} \quad (12)$$

which can be computed by Eq. (6) with specialization on $\{(\omega_i^S, \hat{\Psi}(\mathbf{x}_i^S))\}$ and $\{(\omega_j^T, \mathbf{x}_j^T)\}$. The flow component \mathbf{g}_i moves \mathbf{x}_i^S from $\hat{\Psi}(\mathcal{S})$ to close to \mathcal{T} .

2 Morphing Results

We show in Figs. 1 to 4 more morphing results on different combinations of sources and targets rendered in different viewing angles. For more results and multiview-consistent visualizations please see the video. Fig. 1 shows the transition from ‘Chair’ to ‘Mic’, rendered in three different views. Fig. 2 shows morphs between ‘Lego’ and ‘Materials’ in three different views. Fig. 3 renders more views of morphing with Synthetic-NSVF shown in the main paper. Fig. 4 demonstrates more morphing results of scenes from Synthetic-NSVF.



Fig. 1. The morphing shows the transition from ‘Chair’ to ‘Mic’. From left to right, the morphs have blending weights $t = \{0, .2, .4, .6, .8, 1\}$. From top to down, the images are rendered with azimuth $\theta = \{30^\circ, 150^\circ, 270^\circ\}$ and elevation $\phi = 30^\circ$. Note that the weight t , the azimuth θ , and the elevation ϕ are chosen randomly and can be replaced with different values.

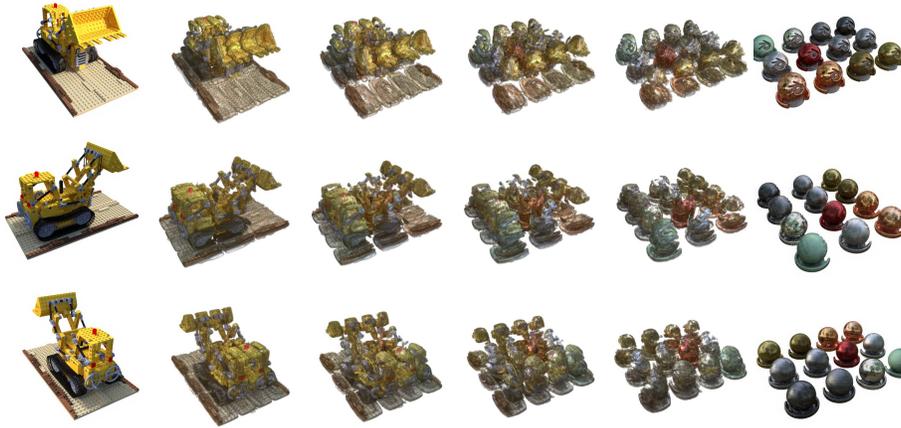


Fig. 2. The morphing shows the transition from ‘Lego’ to ‘Materials’. From left to right, the morphs have blending weights $t = \{0, .2, .4, .6, .8, 1\}$. From top to down, the images are rendered with azimuth $\theta = \{30^\circ, 150^\circ, 270^\circ\}$ and elevation $\phi = 30^\circ$. Note that the weight t , the azimuth θ , and the elevation ϕ are chosen randomly and can be replaced with different values.



Fig. 3. Rendering in more views of morphing results of Synthetic-NSVF shown in the main paper. Each row presents the morphing process as in Fig. 1, but with different azimuths. In the first and the second canvases, the rows are rendered with azimuth $\theta = \{30^\circ, 150^\circ, 270^\circ\}$, from top to bottom. While in the third canvas, the rows are rendered with $\theta = 60^\circ$ and 240° .



Fig. 4. More morphing results between scenes from Synthetic-NSVF. Each row presents the morphing process as in Fig. 1, but with different azimuths. In the first and the second canvases, the rows are rendered with azimuth $\theta = \{30^\circ, 150^\circ, 270^\circ\}$, from top to bottom. While in the third canvas, the rows are rendered with $\theta = 60^\circ$ and 240° .

References

1. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems* **26**, 2292–2300 (2013) [1](#)
2. Feydy, J., Séjourné, T., Vialard, F.X., Amari, S.i., Trounev, A., Peyré, G.: Interpolating between optimal transport and mmd using sinkhorn divergences. In: *The 22nd International Conference on Artificial Intelligence and Statistics*. pp. 2681–2690. PMLR (2019) [1](#), [3](#)
3. Genevay, A., Peyré, G., Cuturi, M.: Learning generative models with sinkhorn divergences. In: *International Conference on Artificial Intelligence and Statistics*. pp. 1608–1617. PMLR (2018) [1](#)
4. Peyré, G., Cuturi, M., et al.: Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning* **11**(5-6), 355–607 (2019) [1](#)