Eurographics Symposium on Geometry Processing 2025
M. Attene and S. Sellán
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 44 (2025), Number 5

Volume Preserving Neural Shape Morphing - Supplementary

Camille Buonomo[1] and Julie Digne [1] and Raphaëlle Chaine[1]

[1]CNRS, Université Claude Bernard Lyon 1, INSA Lyon, LIRIS UMR5205

1. Resulting fields

To better visualize the learned velocity field, we trained our model on a 2D toy experiment morphing two parallel rectangles into a hollowed disk (Figure 1). The velocity field appears smooth and well aligned with the expected deformation.

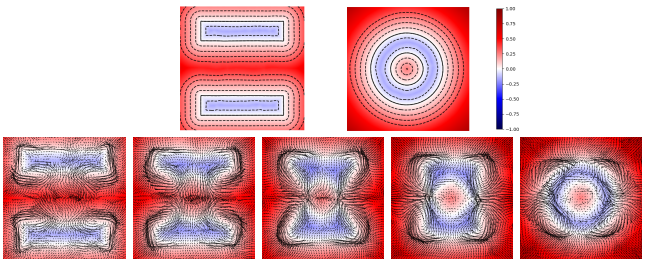

Figure 1: *Evolution of the velocity field (black arrows) during shape interpolation. The value of f_θ is reflected by a color scale. The source and target SDF are displayed in the first row.*

Figure 2 shows the different components combined to create our adaptive divergence vector field on a 2D experiment for easier visualization. $\mathcal{H}$ is the 2D equivalent of **D**, it is the scalar field from which we derive the divergence free component of **V**. One can see that, by construction, the level lines of $\mathcal{H}$ are tangent to **W**. The second component $\mathrm{div}(\beta(f)\nabla u)$ is divergence free on the shape boundary. Near the medial axis, $\mathrm{div}(\beta(f)\nabla u)$ is high and some level sets are absorbed (change of volume in those areas).

2. Behavior with respect to rotation and trivial motions

As a sanity check, we interpolate between a torus and itself, for which we expect the transform to be the identity. The results are reported on Figure 3. As expected, our method is able to reproduce this trivial motion. The second test is to interpolate between a torus and its 90° rotated version, for which we expect to recover a rotation. Figure 3 (bottom row) shows that the inferred deformation is consistent with a rotation, but adds a small nonrigid deformation.

3. Robustness to noise

We consider morphing between source and target shapes with increasing levels of noise. In practice, our method inherits the robustness of the source and target SDF training method and produces a consistent transition between very noisy shapes (Figure 4). Input point clouds are degraded through a Gaussian Noise, normal directions are recomputed while keeping the overall orientation consistent, the auxiliary source and target SDF are trained on these noisy inputs, and our method is applied to these data. This also tends to show that our method is relatively stable with respect to small pertubations of the input data.

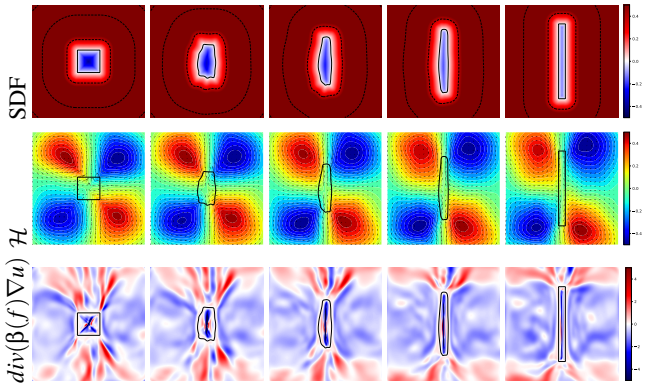

Figure 2: *Visualization of the components used for building our adaptive divergence vector field in 2D. Top row: resulting SDF approximation, for reference; middle row: scalar potential $\mathcal{H}$, with black arrows highlighting the divergence free field constructed over $\mathcal{H}$; bottom row: adaptive divergence component $div(\beta(f)\nabla u)$.*

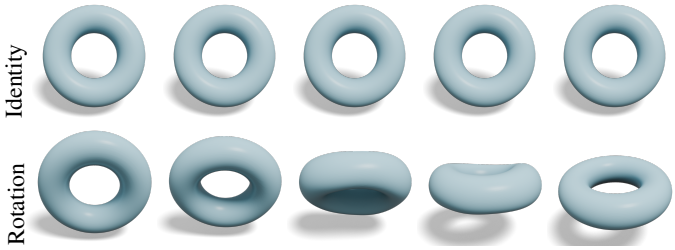

Figure 3: *Interpolations produced for an identity groundtruth transform and a rotation groundtruth transform.*

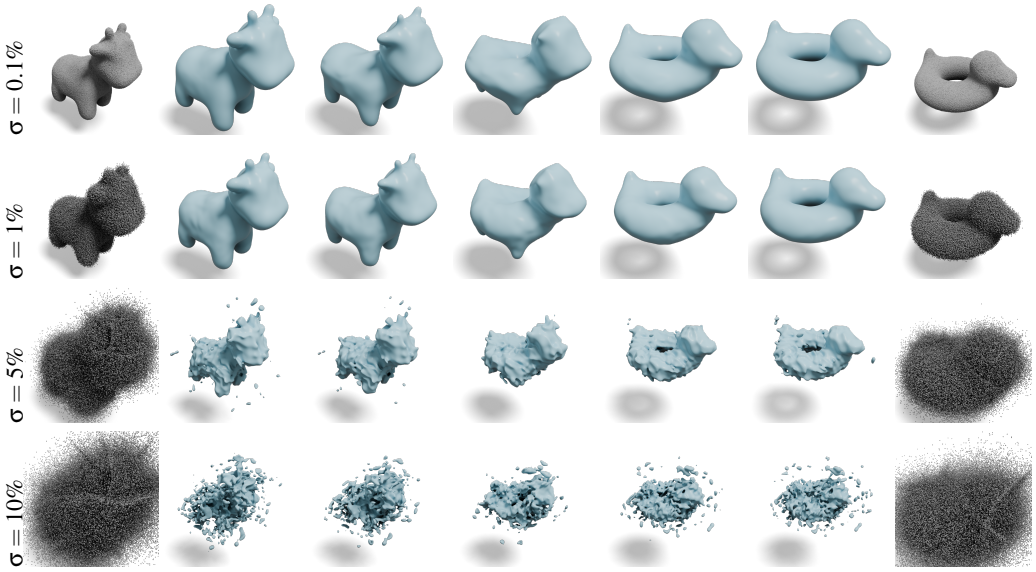

Figure 4: *Results obtained for various noise levels (standard deviation σ expressed as a percentage of the shape diagonal). Our method can handle moderate amounts of noise, but fails for large noise (σ = 10%diagonal).*

4. Ablation study - Eikonal-LSE Balance

We study the balance between Eikonal and Level Set equation losses by modifying the weights λ_{LSE} and $\lambda_{Eikonal}$. Figure 5 shows the distribution of f_θ's gradient norm depending on the weight of the Eikonal term in the loss. Compared to NISE and INSD, we can see that enforcing the Eikonal at any time yields a gradient distribution which is much more centered around 1 (a signed distance field gradient histogram should have a single peak at 1). Figure 6 shows the impact of the weight of the LSE loss on the volume preservation. When λ_{LSE} increases, the volume is better preserved throughout the deformation.

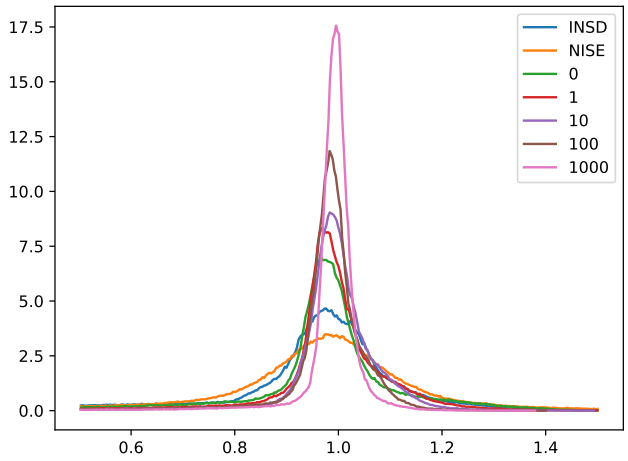

Figure 5: *f_θ's gradient norm distribution for different Eikonal weights in the loss for the cylinders to torus experiment, compared to NISE and INSD.*

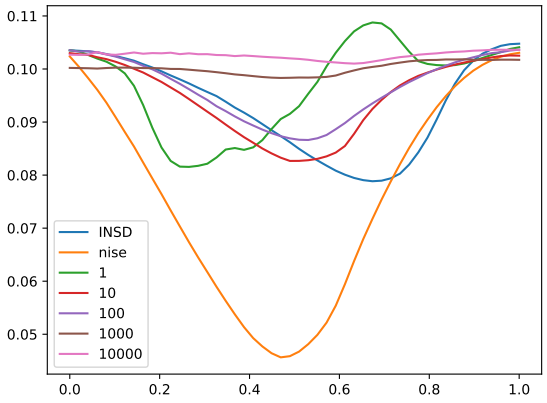

Figure 6: *Volume variation for different LSE weights in the loss, compared to NISE and INSD on the cylinders to torus experiment.*

Figure 8 and 7 illustrate the distribution of the LSE error term and the Eikonal error term at different time steps on a slice of the cylinders-torus experiment.

All methods yield a vanishing gradient of f_θ on the medial axes of the shapes, since the gradient is undefined on this part of the space (Figure 7). NISE does not have any constraint on the gradient of the intermediate shapes, resulting in a poor distribution of the gradient norm. Our method allows to easily tune the penalization of gradient by adjusting λ_{Eik} to better preserve the gradient norm, even for intermediate shapes. INSD uses a modified level set equation to preserve the gradient norm only along the shape contour for

Figure 7: *Distribution of $\|\nabla f\|$ for different time steps on a slice of the cylinders-torus experiment.*

any time t, but the authors did not provide any insight on how the value of the hyper parameter of the right hand side of their modified LSE could impact the distribution of the gradient. We chose to use $\lambda_{Eik} = 1$ in our experiments, which gave good results. The high computation time of our implementation of INSD did not allow to do a full grid search on the hyper parameters.

On Figure 8, we can see that the compliance with the LSE increases with λ_{LSE}. Interestingly, the highest errors occur at topological changes where and when surface shape parts merge (see $\lambda_{LSE} = 1e4$). Without topological changes, the LSE error is low everywhere. For low λ_{LSE} we can see that the LSE error distribution is spread over spacetime. ($\lambda_{LSE} = 1$ for example), compared to NISE. This can be explained by the fact that NISE explicitely uses the scalar potential in the handcrafted $\mathbf{V}$, while it appears only through its space and time derivatives in our case. Since it is important for volume preservation that the LSE equation is fulfilled on surfaces, we chose $\lambda_{LSE} = 1e4$ (Figure 6). We recommend that the weight of the LSE loss is at least equivalent to the weight $\lambda_{Dirichlet}$ of the Dirichlet loss. In the case of INSD the approach does not intend to satisfy the original expression of the LSE away from the

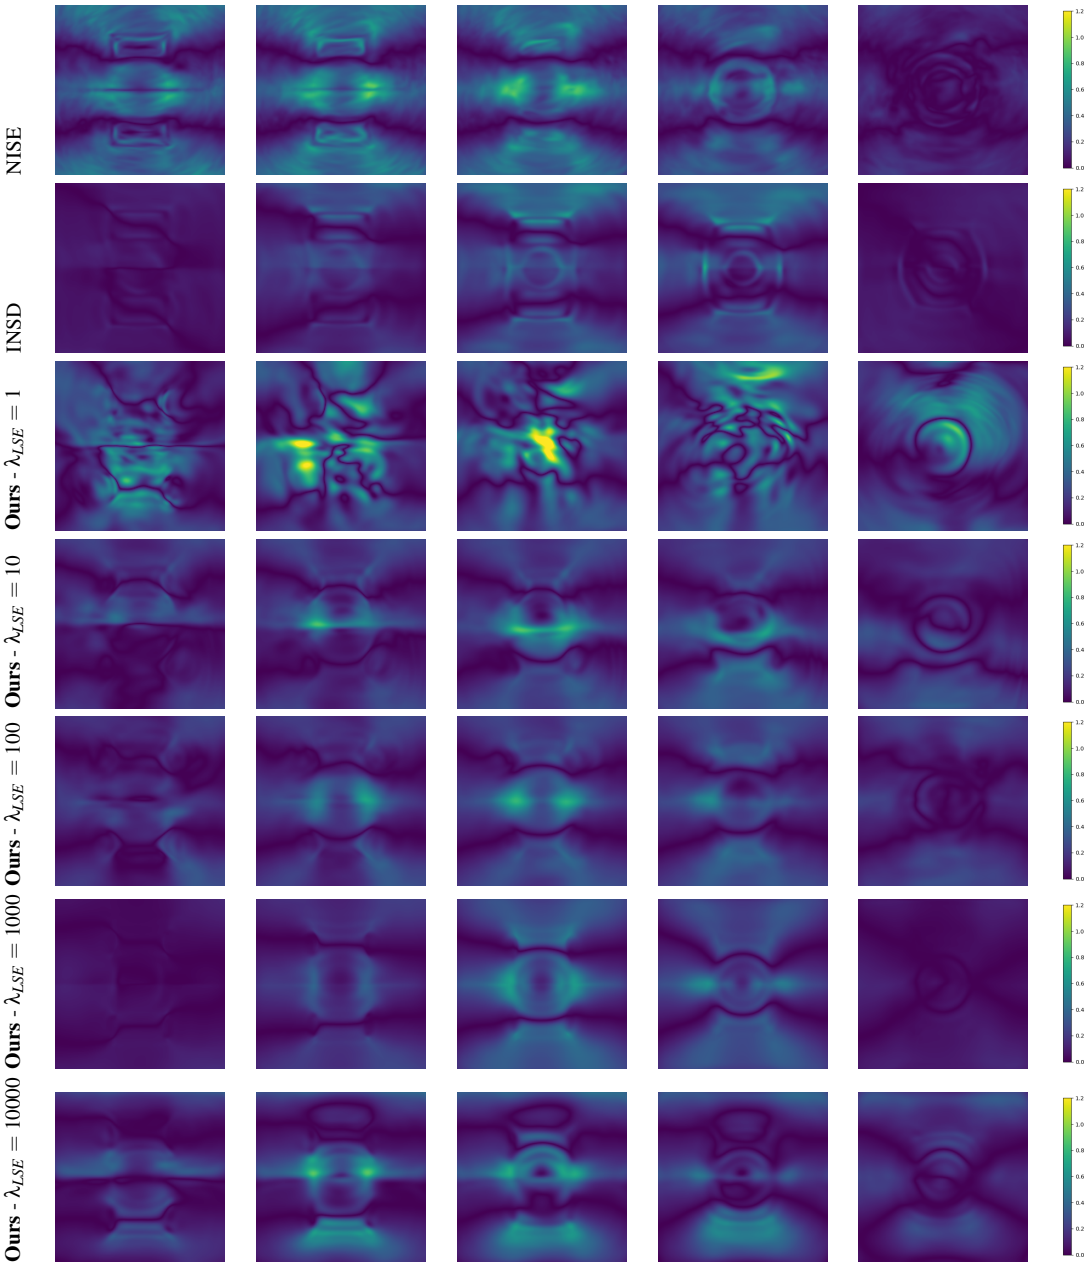

Figure 8: *Distribution of the LSE loss value for different time steps on a slice of the cylinders-torus experiment.*

surface. Hence the error is higher avay from the surface. However the modified LSE they introduce prevents the error to explode.

Changing the LSE loss weight has an impact on the detail preservation as well: lowering it allow the method to focus more on the source and target details at the cost of a blobby morphing, while increasing it yields less detailed shapes, but a morphing with less spurious blobby artifacts. Figure 9 shows the different results obtained by changing the LSE weight.

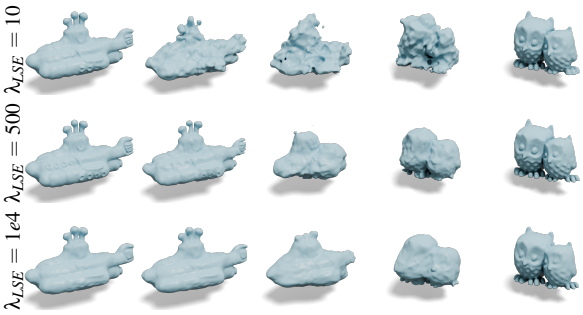

Figure 9: *Impact on the detail preservation and morphing quality of the LSE loss weight change.*

5. Comparison to baselines methods

We compare our method to simple neural baselines in Figure 10. The first method (w/out LSE) consists in learning directly $f(x,t)$ from the source and target shapes only with an Eikonal constraint. The second baseline is to compute an explicit flow of the form $x + t\phi(x,t)$ to get an exact solution of the LSE with $g_0(x+t\phi(x,t))$. Our experiments show that both baselines do not capture meaningful interpolations, and are obviously not volume preserving.

6. Incorporating landmark correspondences

Although our method targets morphing between shapes without semantic correspondences, landmarks can be incorporated in our method. To do so, we use the method proposed in [SCC*25]. Let $(y_i^0, y_i^1)_i$ be a pair of correspondences on shape $\mathcal{S}_0$ and $\mathcal{S}_1$, each pair must verify:

$$y_i^1 = y_i^0 + \int_0^1 \mathbf{V}(y_i^t, t) \tag{1}$$

This induces an additional loss term:

$$l_{Corr} = \sum_i \|y_i^1 - \tilde{y}_i^1\| \tag{2}$$

Where $\tilde{y}_i^1$ is obtained via the forward Euler scheme using our learned vector field $\mathbf{V}_\theta$: $\tilde{y}_i^{t_{n+1}} = \tilde{y}^{t_n} + \Delta t \mathbf{V}_\theta(\tilde{y}^{t_n}, t_n)$, $\tilde{y}_i^0 = y_i^0$. Additionally, every point remains on the surface through the integration, their implicit value must be 0, which is mirrored in another loss:

$$l_{time-Dirichlet} = \sum_i \sum_n \|f_\theta(\tilde{y}_i^{t_n}, t_n)\| \tag{3}$$

Figure 11 shows the results of using landmarks for our method and INSD. Although, the twist motion is captured by both methods, clearly, INSD yields much better intermediate steps, we hypothesize that this is due to an additional Laplacian regularization on the vector field in the INSD method, a constraint which we did not use in our method.

7. Comparison to an explicit morphing method

We compare our method to [ELC18] (DFSCD), an explicit method that iteratively computes keypoints on the shapes, matches them and computes a divergence free vector field, before updating the correspondences. We use a pair of shapes with the same genus. The result shows that the method fails to match the keypoints meaningfully, which can be due to the fact that they are repeated. [ELC18] is much more meaningful to compute matches between semantically close shapes, such as human bodies for example. Figure 12 shows a morphing example.

8. Stability wrt network initialization

We analyze the impact of the random seed. Changing the random initialization of the network weights, leads to a convergence to a different local minimum and consequently different intermediate shapes (Figure 13). This is an expected result since the problem is under-constrained and an infinite number of solutions exist. But we can observe that the solutions are always comparable and the difference are limited.

9. Computation time analysis

Our method is computation-intensive, we show on Figure 14 that, with a lower number of epochs, the obtained morphing is still reasonable but the intermediate steps are less relevant. A strategy to mitigate the computation time would be to use our strategy for morphing coarse shapes in a meaningful way and then transfer the details along, a promising future research direction. We also compare our method with NISE and LF-INSD with a fixed training time of one hour. This time corresponds to the convergence time of NISE. Figure 15 shows that our result is already meaningful even if details are smooth (see the teaser). This shows that our method is a good starting point for morphing coarse shapes before applying detail transfer for sharper details.

10. Additional comparisons and results

We show additional comparisons in Figure 17, 18, 19 and corresponding volume variations on Figure 16.

When interpolating between a feline and a rotated horse, one can see that the motion obtained by most methods is irrelevant. Without landmarks, our method still manages to use the wings for the front legs and regrow the back legs. Other methods either tear the surface (OT) or produce less significant blobby intermediate steps. To get a better deformation, we can estimate a rigid transform alignment through PCA and compute the nonrigid deformation in these aligned poses (Figure 18). This alignment allows our method to better leverage shape parts with similar geometry, such as the legs for example. Finally Figure 20 shows some additional morphing results. It should be noted that the quality of the interpolation depends also on the quality of the source and target SDF. For example, the falcon to witch morphing used a low-detail representation.

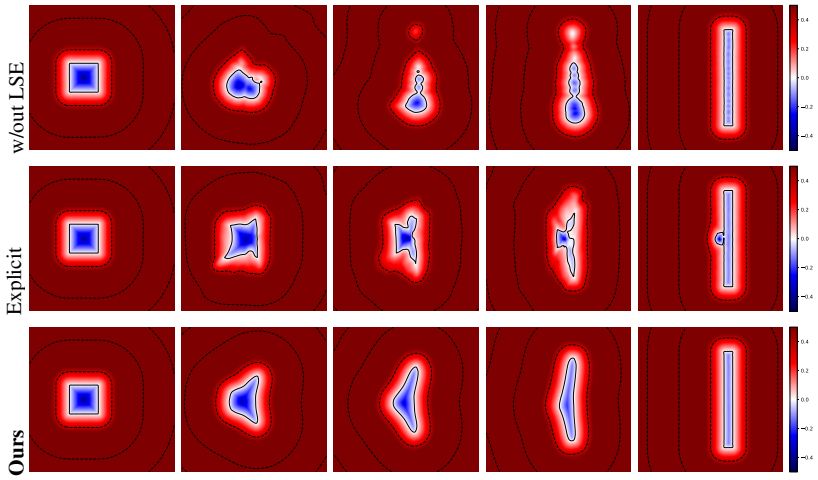

Figure 10: *Comparison of the square to rectangle morphing using various baselines methods (without LSE, with explicit flows, ours). Note that the baseline without LSE still efficiently learns the boundary shapes, but fails to produce smooth and meaningful intermediary shapes. For the explicit flow baseline, the first shape is exact (in terms of implicit representation), but the baseline fails to learn both the final shape and smooth intermediated shapes.*

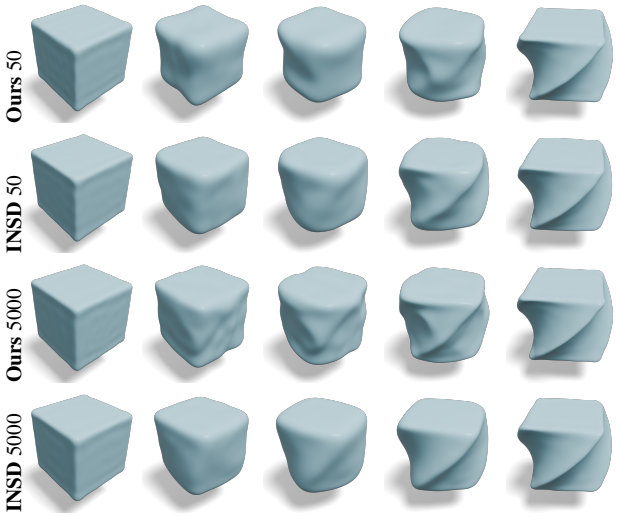

Figure 11: *Comparison of the effect of sparse and dense landmark correspondences on our method and INSD [SCC*25].*

Acknowledgements This work was partially funded by ANR-23-PEIA-0004 (PDE-AI). Work under CC-BY 4.0 license https://creativecommons.org/licenses/by/4.0/.

References

[ELC18] EISENBERGER M., LÄHNER Z., CREMERS D.: Divergence-free shape interpolation and correspondence, 2018. 5, 7

[SCC*25] SANG L., CANFES Z., CAO D., BERNARD F., CREMERS D.: Implicit neural surface deformation with explicit velocity fields. In *The Thirteenth International Conference on Learning Representations* (2025). 5, 6

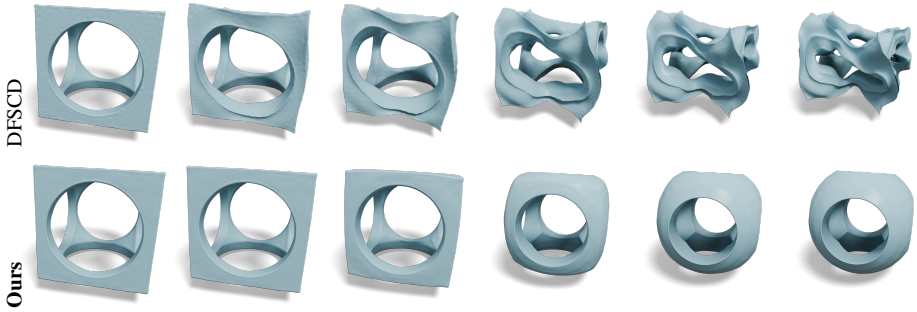

Figure 12: *Comparison with the method of Eisenberger et al. [ELC18] on a hollow cube to a hollow sphere.*

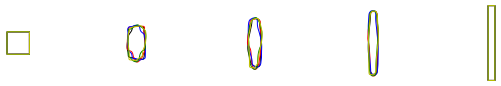

Figure 13: *Different seeds for the network weights random initialization yield different intermediate shapes. Contours are drawn in different colors for each seed (square to rectangle experiment, Figure 10).*

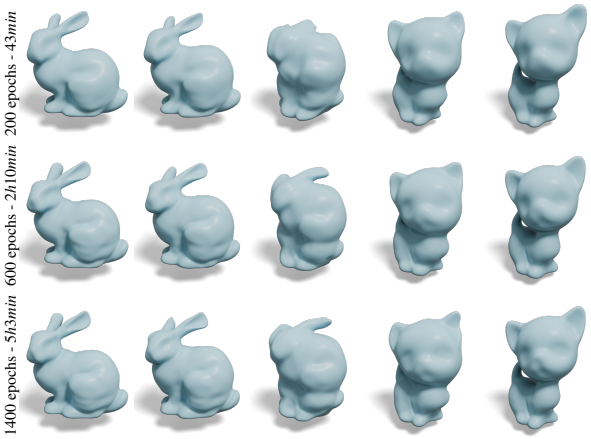

Figure 14: *Comparisons of different morphing obtained for various numbers of epochs.*

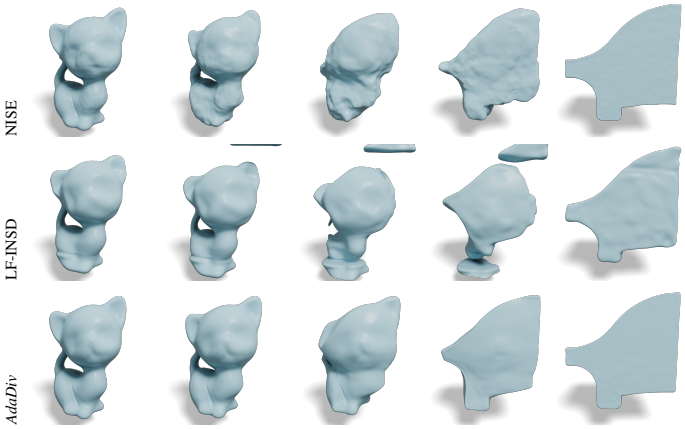

Figure 15: *Comparisons of the results obtained after one hour for various methods.*

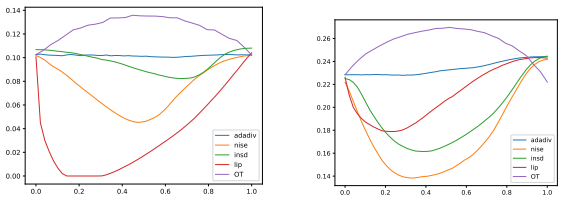

Figure 16: *Volume evolution when morphing cylinders to a torus (left) and feline to horse in aligned poses (right) using our method compared to INSD, NISE, LipMLP and OT. See Figures 19 and 18 for the corresponding interpolations.*

Figure 17: *Comparison on Max-Bust. On this simple shape all methods produce a reasonable morphing. With NISE, intermediate steps are a little more blobby.*

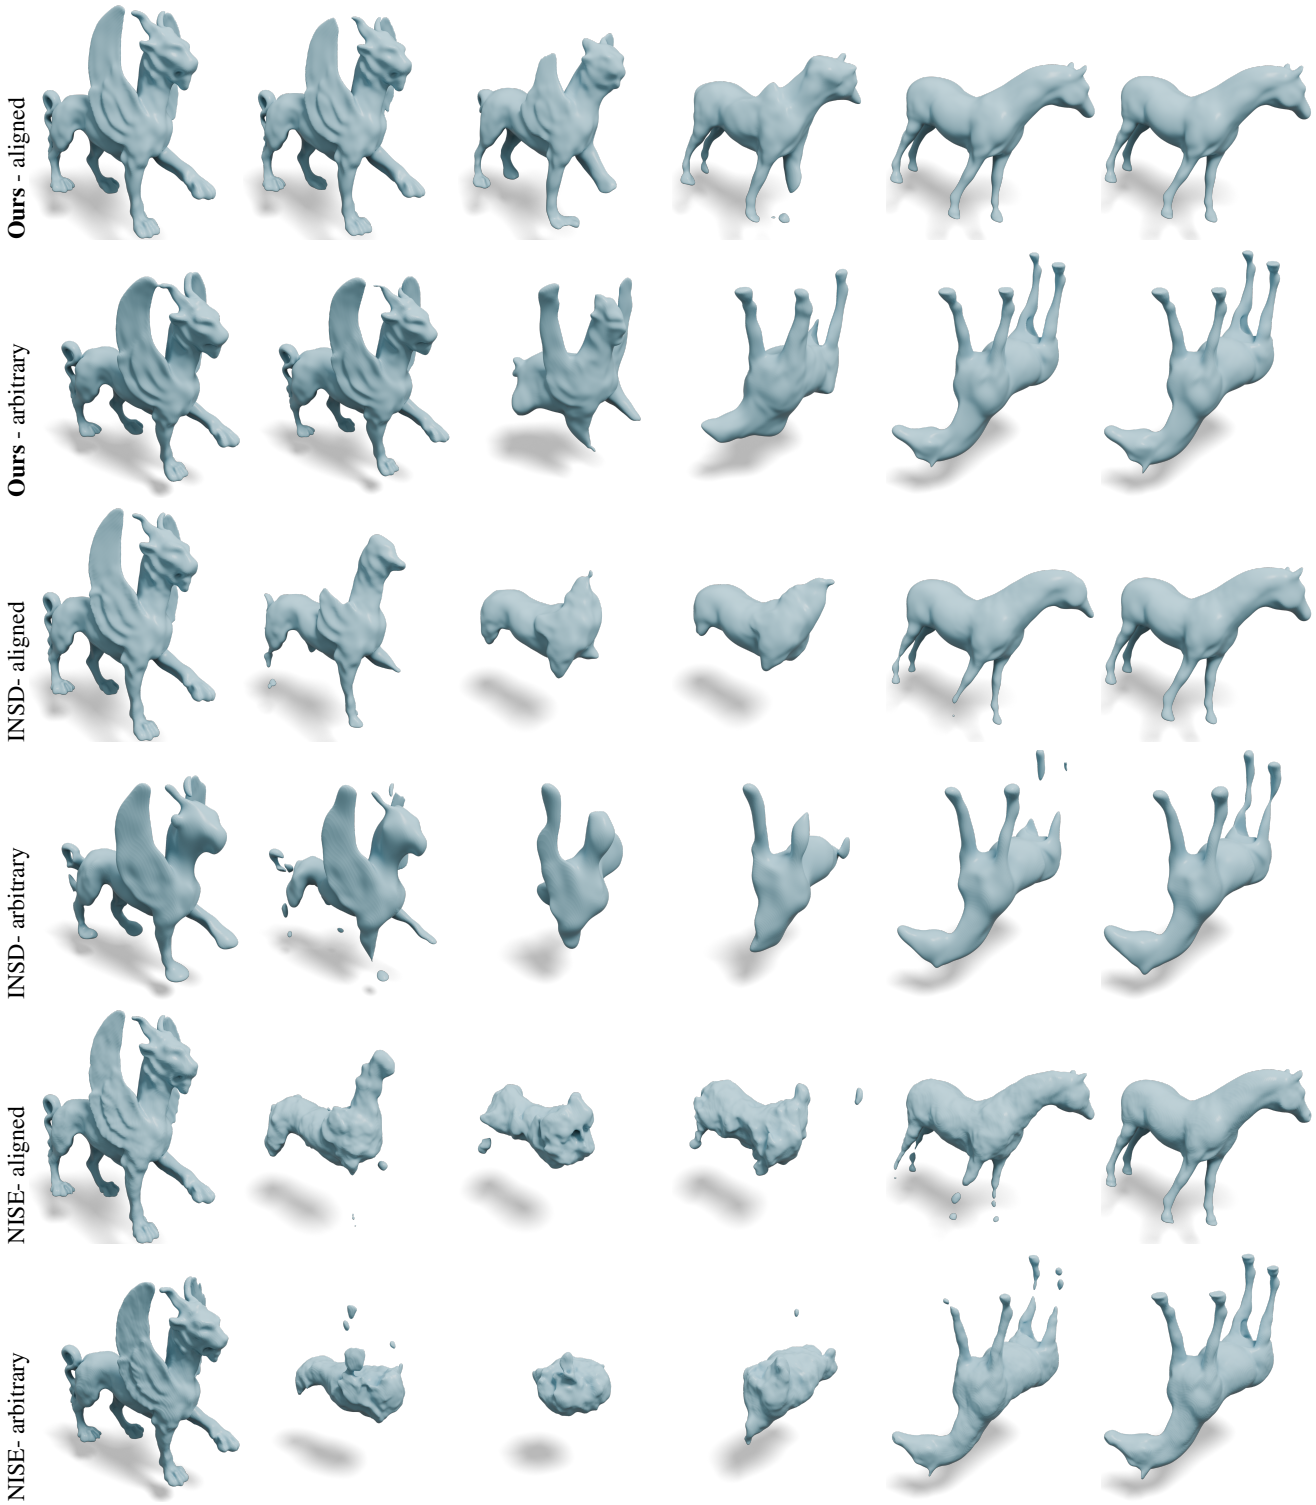

Figure 18: *Comparison on Feline-Horse in aligned and arbitrary poses.*

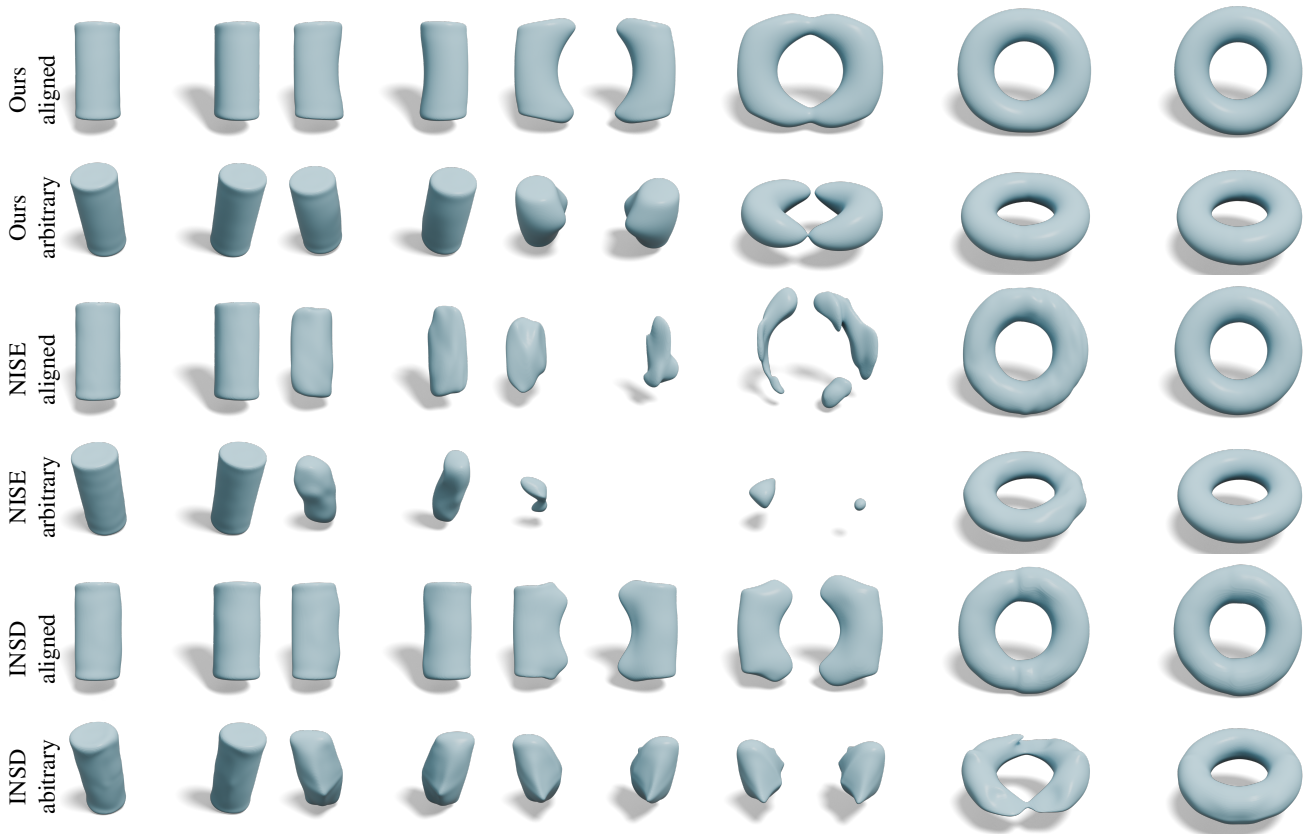

Figure 19: *Comparison on the 2 cylinders to torus experiment in aligned or abitrary poses using our method, NISE or INSD.*

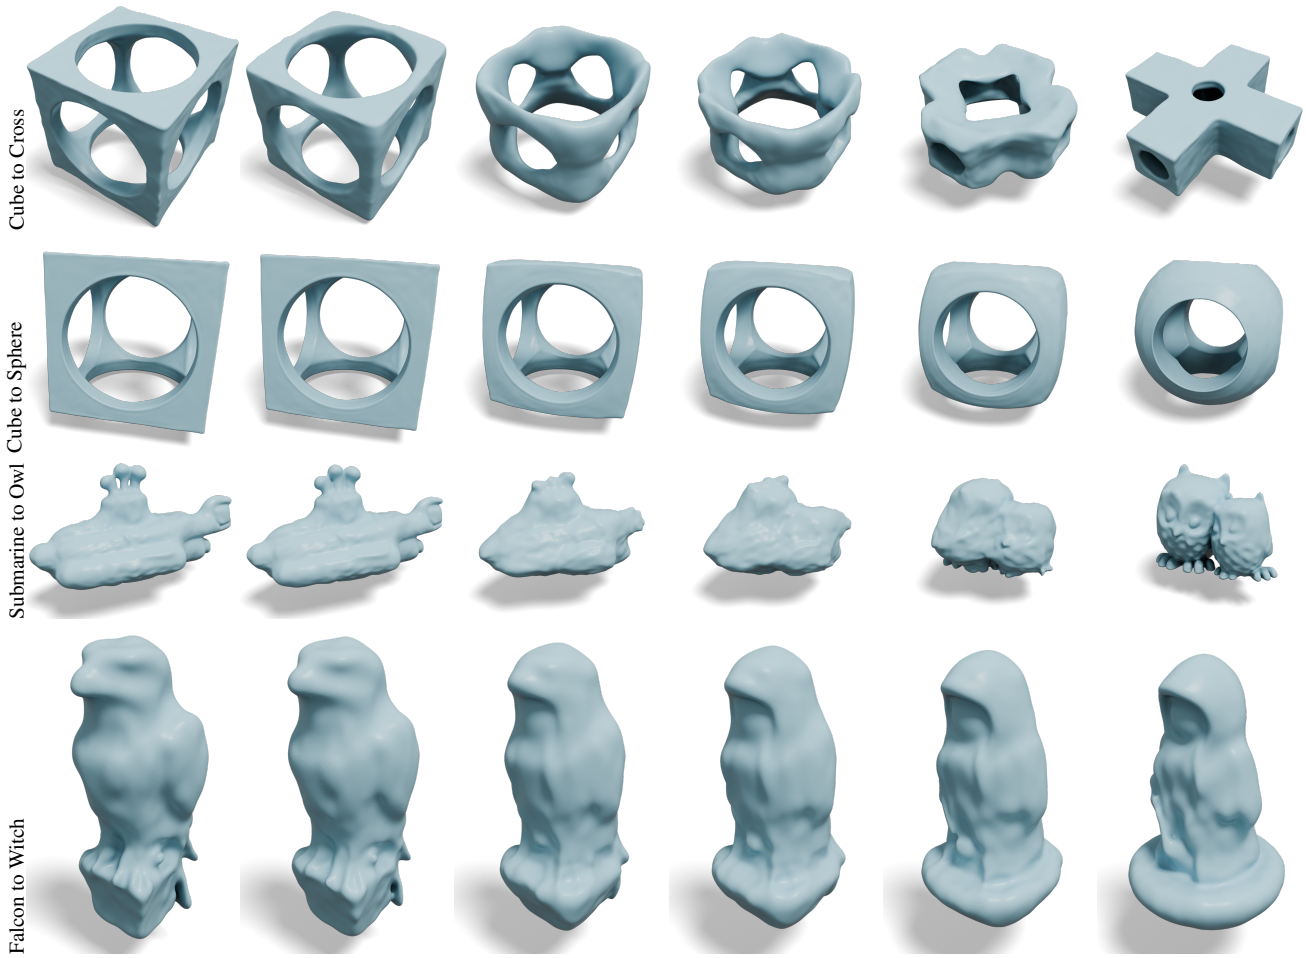

Figure 20: *Additional morphing results .*