Proceedings of Machine Learning Research 1:1–4, 2022 Geometric Deep Learning in Medical Image Analysis 2022

Structure preserving implicit shape encoding via flow regularization

Sven Dummer[1] S.C.DUMMER@UTWENTE.NL

Nicola Strisciuglio[2] N.STRISCIUGLIO@UTWENTE.NL

Christoph Brune[1] C.BRUNE@UTWENTE.NL

[1] *Mathematics of Imaging & AI, Department of Applied Mathematics, University of Twente, NL*

[2] *Data Management & Biometrics, Department of Computer Science, University of Twente, NL*

Editor: Under Review for GeoMedIA 2022

Abstract

Recently, implicit neural representations (INRs) emerged as an effective method for reconstructing shapes. Several of such methods transform templates to target shapes. Current template-based methods lack proper regularization. In this work, we add a novel regularization to a deformable template approach and discuss the benefits of this regularization with a simple test case.

Keywords: latent space, shape reconstruction, implicit neural representation, neural ODE

1. Introduction

3D shape reconstruction using deep learning has important applications in many fields, for instance, in computer vision (Huang et al., 2021) or in the medical domain (Wiesner et al., 2022; Balashova et al., 2019). To perform 3D shape reconstruction, several works use implicit neural representations (INRs) to represent the shape (Park et al., 2019; Mescheder et al., 2019). The main idea behind this representation is to indicate for each point in the domain whether the point is inside or outside the shape. The points where a transition occurs from 'inside' to 'outside' correspond to points on the shape.

When using this implicit neural representation for 3D shapes, it is difficult to calculate point-correspondences between different shapes. To overcome this difficulty, methods are proposed that combine templates with INRs (Sun et al., 2022). These approaches represent a template shape as an implicit neural representation and transform this template into other shapes, allowing for point-to-point correspondences.

Similar to the fact that there are an infinite number of implicit representations to represent a specific shape, in the template-based approach there are an infinite number of template and transformation pairs that are able to reconstruct the same shape. This is an issue as not every choice of template and transformation pair reflects the shape prior of the data. While the INR approaches are able to select a specific implicit function by neural network construction (Mescheder et al., 2019) or by regularization (Sitzmann et al., 2020), not much research has been done in this direction regarding the template-based approaches.

In this work, we address this issue by extending an existing deformable template approach with a regularization strategy inspired by mathematical shape space analysis. Using a simple example, we show that adding this regularization selects a specific type of template and deformation pair. More precisely, we show that this choice results in a template that reflects the shape prior and in learned deformations that give physically plausible interpolations between the template and the data. Moreover, we discuss an additional advantage of our strategy when compared to more standard INR approaches.

2. Method

2.1. Neural diffeomorphic flow

The model we focus on is neural diffeomorphic flow (NDF) (Sun et al., 2022). It consists of a template shape and a warping module. The template shape is parameterized by an INR. More precisely, the template shape is defined via a neural network $f_\theta(\boldsymbol{x}) : \mathbb{R}^3 \to \mathbb{R}$ that maps a spatial coordinate $\boldsymbol{x}$ to a number $y = f_\theta(\boldsymbol{x})$. In this case, the template shape corresponds to the points $\boldsymbol{x}$ such that $f_\theta(\boldsymbol{x}) = c$ for some predefined constant c.

The warping module is defined via a neural ordinary differential equation of the form $\frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}t} = v_\theta(\boldsymbol{x}(t), t, \boldsymbol{z})$ with $\boldsymbol{x}(0) = \boldsymbol{x}_0$, $\boldsymbol{z}$ a latent code, and $v_\theta(\boldsymbol{x}(t), t, \boldsymbol{z})$ a neural network. The warping function is defined as $D(\boldsymbol{x}_0, \boldsymbol{z}, t) := x(t)$.

Combining the template shape representation with the warping module, we define the implicit representation $I(\boldsymbol{x}, \boldsymbol{z}, t)$ as

$$I(\boldsymbol{x}, \boldsymbol{z}, t) := f_\theta(D(\boldsymbol{x}, \boldsymbol{z}, t)) \tag{1}$$

where $I(\boldsymbol{x}, \boldsymbol{z}, 0)$ equals the implicit representation of the template shape and $I(\boldsymbol{x}, \boldsymbol{z}, 1)$ is the implicit representation of the shape corresponding to latent code $\boldsymbol{z}$. The template shape, the implicit representations $I(\boldsymbol{x}, \boldsymbol{z}, 1)$, and the latent codes of the training samples are learned via an autodecoder strategy using training shapes represented via an implicit function (Sun et al., 2022; Park et al., 2019).

2.2. Flow regularization

To let the neural networks learn a specific template and warping module pair, we introduce a novel regularization on the velocity vector field v_θ in the NDF. We regularize v_θ via:

$$L_{reg} := \mathbb{E}_{\mathcal{S}_i} \left(\int_\Omega \int_0^1 L_v(v_\theta(D(\boldsymbol{x}, \boldsymbol{z}_i, t), t, \boldsymbol{z}_i)) \mathrm{d}t \mathrm{d}\boldsymbol{x} \right) \tag{2}$$

where $\mathcal{S}_i$ is a training shape, $\boldsymbol{z}_i$ is the corresponding latent code, Ω is the region in which the shape should be present, and L_v is a regularization functional on the velocity vector field. There are many different choices of L_v. In this paper, we choose an isometric (rigid) deformation prior on the flow via the Killing energy (Solomon et al., 2011; Tao et al., 2016):

$$L_v := \frac{1}{2} \|(J_{\boldsymbol{x}} v_\theta) + (J_{\boldsymbol{x}} v_\theta)^T\|_F \tag{3}$$

where $J_{\boldsymbol{x}} v_\theta$ is the Jacobian of v_θ with respect to $\boldsymbol{x}$. As a consequence, we expect $I(\boldsymbol{x}, \boldsymbol{z}, t)$ to preserve the template's structure over time.

3. Numerical results

In the neural diffeomorphic flow model, we represent the template and training shapes as occupancy functions (Mescheder et al., 2019). As test case, we use a dataset consisting of two squares of equal size that are translated and rotated in space. On this dataset, we train the model with and without the proposed regularization term (Equation (2)) in the loss function. The result of training can be seen in Figure 1.

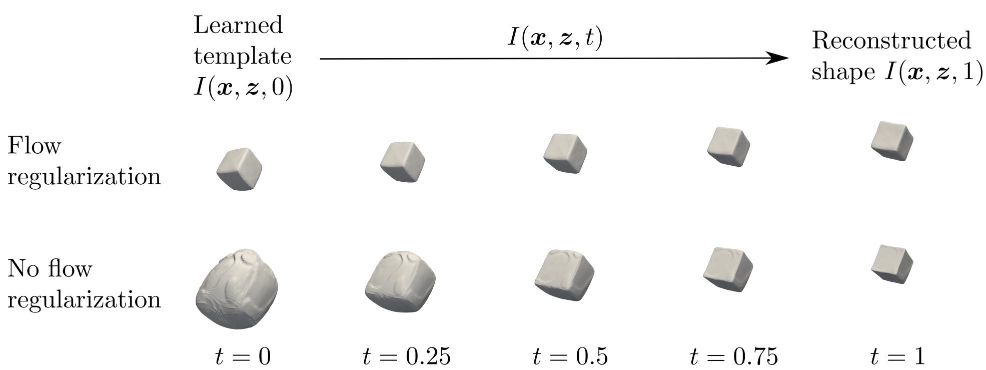

Figure 1: Decoding process of the shape $I(\boldsymbol{x}, \boldsymbol{z}, t)$ over time, when flow regularization is used (first row) or not (second row). The flow regularization enables the preservation of the shape structure prior in time.

From Figure 1, we see that in both cases the square is well reconstructed. Without regularization a template and deformation pair is learned that does not resemble the data class with its prior assumptions. However, when applying regularization, the template shape rotates to the reconstructed shape as desired. Hence, the deformable template approach allows for regularizing the template and the interpolations by regularizing the flow in the warping module. Regularizing the interpolations is much harder to do in standard INRs like DeepSDF (Park et al., 2019). In this case, one has to infer some particle velocity vector field from the changes of implicit functions and regularize these velocity vector fields (Atzmon et al., 2021). In contrast, with the deformable template approach one can immediately regularize the warping module.

4. Conclusion and future work

We introduced a shape space based regularizer to the neural diffeomorphic flow model and showed that this regularizer enables template learning to find a shape mean that preserves the main structure of the data. We showed that learning with the regularizer yields proper deformations of the learned template shape into the reconstructed training shapes. Moreover, we highlighted that regularizing shape interpolations via the flow is straightforward compared to regularizing shape interpolations in a more standard implicit neural representation approach like DeepSDF. In future work, we will apply and analyse the method in more complex scenarios. In addition, we will investigate the quality of the latent space in terms of extrapolation and latent code interpolation.

References

Matan Atzmon, David Novotny, Andrea Vedaldi, and Yaron Lipman. Augmenting implicit neural shape representations with explicit deformation fields. *arXiv preprint arXiv:2108.08931*, 2021.

Elena Balashova, Jiangping Wang, Vivek Singh, Bogdan Georgescu, Brian Teixeira, and Ankur Kapoor. 3D organ shape reconstruction from topogram images. In *Information Processing in Medical Imaging*, pages 347–359. Springer International Publishing, 2019. ISBN 978-3-030-20351-1.

Qixing Huang, Xiangru Huang, Bo Sun, Zaiwei Zhang, Junfeng Jiang, and Chandrajit Bajaj. Arapreg: An as-rigid-as possible regularization loss for learning deformable shape generators. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5815–5825, October 2021.

Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33:7462–7473, 2020.

Justin Solomon, Mirela Ben-Chen, Adrian Butscher, and Leonidas Guibas. As-killing-as-possible vector fields for planar deformation. In *Computer Graphics Forum*, volume 30, pages 1543–1552. Wiley Online Library, 2011.

Shanlin Sun, Kun Han, Deying Kong, Hao Tang, Xiangyi Yan, and Xiaohui Xie. Topology-preserving shape reconstruction and registration via neural diffeomorphic flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20845–20855, June 2022.

Michael Tao, Justin Solomon, and Adrian Butscher. Near-isometric level set tracking. In *Computer Graphics Forum*, volume 35, pages 65–77. Wiley Online Library, 2016.

David Wiesner, Julian Suk, Sven Dummer, David Svoboda, and Jelmer M Wolterink. Implicit neural representations for generative modeling of living cell shapes. *arXiv preprint arXiv:2207.06283*, 2022.