Paper ID: 7163
Neural Mesh Flow: 3D Manifold Mesh Generation via Diffeomorphic Flows

We thank the reviewers for their thoughtful feedback. We are delighted that (**R1**, **R3**, **R4**) voted clear accept. We appreciate their finding the paper *"a pleasure to read"* (**R1**), with experiments comprising *"very convincing visual outputs, as well as good quantitative scores"* (**R4**). They find our formulation *"novel"* (**R3**), *"interesting"* (**R1**) and giving *"new perspectives in 3D mesh reconstruction"* (**R4**).

(**R2**) **Comparison to Occupancy Networks:** We believe **R2** missed comparisons to OccNet [16] included in the paper. NMF performs better on 6 out of 7 metrics as shown in Tab 3 and Fig 8, which we discuss in L265-279. Note that inference time of NMF is only 0.18 sec compared to 6.65 sec of OccNet – a speed up of 37x with superior results.

(**R2**) **"Zero level set of an implicit function via marching cubes $\cdots$ trivial to retrieve a manifold mesh"** We respectfully disagree. As discussed in L88-91, marching cubes involves "*rasterization of iso-surface* [which] *is a purely local operation*, [thus] *often leads to ambiguities*". In contrast, as stated in L271-273, NMF is designed to yield a manifold mesh. The fact that NMF outperforms OccNet in 3 out of 4 manifoldness metrics further validates our point.

(**R2**) **Relevance of simulations and physically based renderings:** Such illustrations are important to demonstrate that manifoldness is useful in practice, as also noted by **R1**, **R3** and **R4**. Both the visualizations in Fig 6-8 and the supplementary video convincingly demonstrate the advantages of NMF over SotA.

(**R2**) **Novelty:** This is the first work to explicitly define and guarantee manifoldness in mesh generation, allowing use in simulations without any post-processing while maintaining a good runtime performance. We note that all of **R1**, **R3**, **R4** rate the paper as novel and giving "*new perspectives in 3D mesh reconstruction*" (**R4**).

(**R2**) **ShapeFlow and OccFlow:** We thank **R2** for pointing us to concurrent work of ShapeFlow [Jiang et al. (2020)], which appeared on arXiv AFTER the NeurIPS deadline. Our contrast stated in L42-44 applies to all of ShapeFlow, OccFlow [15] and PointFlow [14] – they require category specific priors for reconstruction, while NMF does not. This makes NMF more generally applicable and trainable for a diverse set of shapes.

(**R1**) **Point-shape feature and its impact on self-intersection:** It is true that 3D point trajectories can intersect when the dynamics f_θ operates on a manifold in higher dimension containing augmented states [9]. However, NMF dynamics operate in 3D state-space and as such it does not cause self-intersections. Therefore, with any combination of point and shape features, the manifoldness will remain intact. Due to finite numerical precision, some self-intersections can occur, but vanish as $tolerance \to 0$ (see Table 9 in supplementary).

(**R3**, **R4**) **Ablations:** Instance normalization is crucial to the performance of NMF. We include its ablations in Fig. 5 and Fig 11 (supplementary). The integration time was empirically determined to be large enough for flow to work, but not too large to cause overfitting. We will add more discussion on these in the final version.

(**R3**) **Smoothing:** Please note that NMF meshes are rendered directly, without Laplacian smoothing. Other methods render poorly due to non-manifoldness (Fig 2 in supplementary), so only those methods are smoothed for visualizations.

(**R4**) **Chamfer-L2:** It is possible to get better Chamfer-L2 scores despite inaccurate topology, thus, we believe they should be complemented with manifoldness metrics to better judge mesh quality.

Future work: We thank **R1** for the intriguing idea of exploring [Sitzmann et al. (2020)] for further improvement. (**R2**) While we focus on manifold shapes here, our future work will incorporate texture generation along with reconstruction.

(**R4**) **Higher genus and templates** We agree that using templates more representative of target shape may allow handling higher genus, perhaps at the cost of generalizability. Approaches that learn such templates as part of their pipeline like MeshRCNN [2] can be explored to further generalize NMF to more object categories.

Other feedbacks We thank **R1** for highlighting common motivations of NMF and [Deng et al. (2020)] with respect to physics simulation. We will add a discussion on it as well as one on limitations such as higher genus shapes in the final version (**R4**). Any formatting issues and typos will also be addressed in the final version.

References

Deng, B., Genova, K., Yazdani, S., Bouaziz, S., Hinton, G., and Tagliasacchi, A. (2020). Cvxnet: Learnable convex decomposition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 31–44.

Jiang, C., Huang, J., Tagliasacchi, A., Guibas, L., et al. (2020). ShapeFlow: Learnable deformations among 3D shapes. *arXiv preprint arXiv:2006.07982*.

Sitzmann, V., Martel, J. N. P., Bergman, A. W., Lindell, D. B., and Wetzstein, G. (2020). Implicit neural representations with periodic activation functions. *arXiv preprint arXiv:2006.09661*.