Supplementary material: A generic and flexible regularization framework for NeRFs

Thibaud Ehret Roger Marí Gabriele Facciolo
Université Paris-Saclay, CNRS, ENS Paris-Saclay, Centre Borelli, 91190, Gif-sur-Yvette, France
thibaud.ehret@ens-paris-saclay.fr

1. Additional visual results

We present additional visual results comparing RegNeRF [2] to the proposed regularization framework (with both depth regularization and normals regularization) in Fig. 1.

2. Parameters study for surface estimation

In Fig. 2, we show the impact of the two parameters controlling the loss, namely the regularization weight λ_{curv} and clipping value κ_{curv}, on a specific example of the DTU dataset [1] using Gaussian curvature. A very small weight and a strong clipping ($\lambda_{curv} = 0.0001$ and $\kappa_{curv} = 1$) leads to barely no regularization as expected. On the contrary, a large weight with little clipping ($\lambda_{curv} = 0.001$ and $\kappa_{curv} = 10$) learns a sort of envelope of the surface. The most interesting results, however, are obtained with intermediate parameters. For example, $\lambda_{curv} = 0.0005$ and $\kappa_{curv} = 5$ produces a surface that is both strongly regularized while keeping most of the details. Another interesting result is that of $\lambda_{curv} = 0.001$ and $\kappa_{curv} = 5$, which leads to a Cubist-style surface, with little details and exaggerated edges. Such models might be interesting to learn a surface that can be represented by a mesh with few triangles or for deriving a piece-wise developable surface [3]. Note also how the regularization has a tendency to fill small gaps (for example the mouth of the bunny is completely regularized with parameters $\lambda_{curv} = 0.0001$ and $\kappa_{curv} = 10$). This shows the importance of clipping to preserve small details and the limits of such type of regularization. This study was done using VolSDF [4].

References

[1] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014. 1

[2] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. RegNeRF: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2

[3] Silvia Sellán, Noam Aigerman, and Alec Jacobson. Developability of heightfields via rank minimization. *ACM Trans. Graph.*, 39(4):109, 2020. 1

[4] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34, 2021. 1

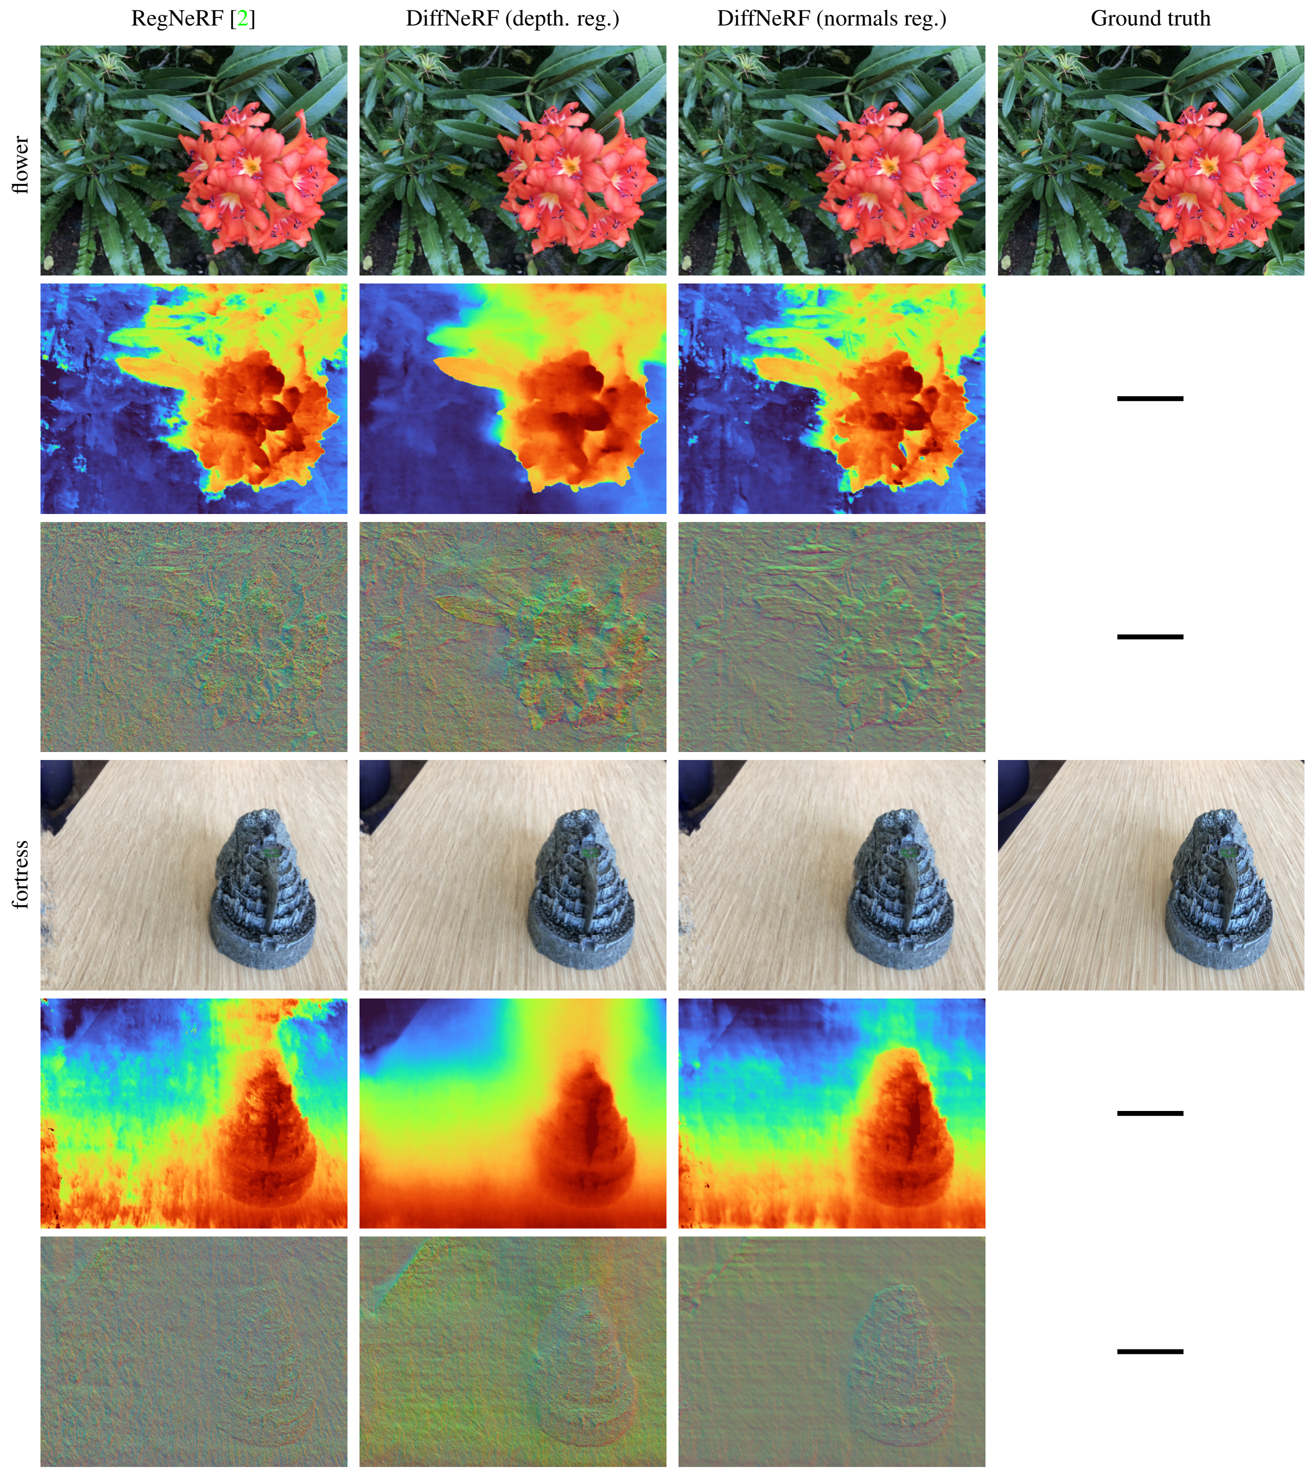

Figure 1. Visual examples of novel view synthesis for the *flower* (top) and *fortress* (bottom) sequences of the LLFF dataset after training with three views.

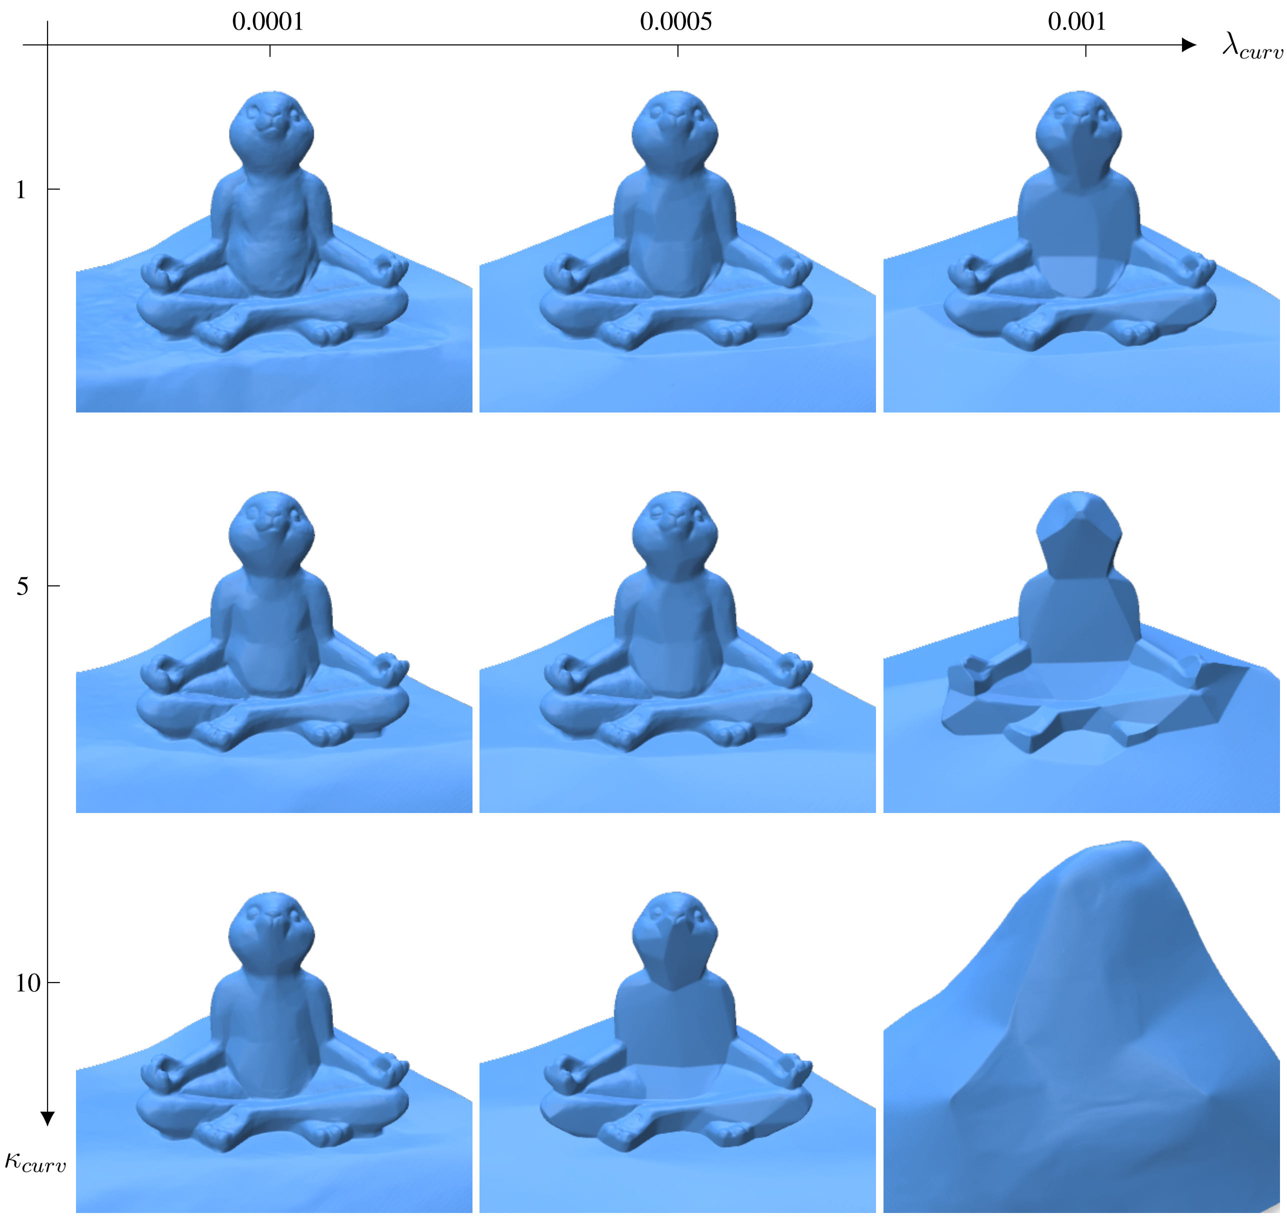

Figure 2. Study of the impact of the two parameters controlling the strength of the regularization, i.e. the regularization weight λ_{curv} and the clipping value κ_{curv}, on the scene 110 of the DTU dataset. While intermediate parameters, such as $\lambda_{curv} = 0.0005$ and $\kappa_{curv} = 5$, provide a good balance between detail preservation and surface smoothness, strong regularization can produce surprising surfaces such as the Cubist-like representation given by $\lambda_{curv} = 0.001$ and $\kappa_{curv} = 5$. These results were computed using the Gaussian curvature. Results best seen zoomed.