

# Neural Implicit Surface Reconstruction from Noisy Camera Observations (Student Abstract)

Sarthak Gupta<sup>1,2</sup>, Patrik Huber<sup>2</sup>

<sup>1</sup> Indian Institute of Technology Roorkee, Roorkee, Uttarakhand, India - 247667

<sup>2</sup> University of York, Deramore Lane, Heslington, York, YO10 5GH, United Kingdom  
mrsarthakgupta@gmail.com, patrik.huber@york.ac.uk

## Abstract

Representing 3D objects and scenes with neural radiance fields has become very popular over the last years. Recently, surface-based representations have been proposed, that allow to reconstruct 3D objects from simple photographs. However, most current techniques require an accurate camera calibration, i.e. camera parameters corresponding to each image, which is often a difficult task to do in real-life situations. To this end, we propose a method for learning 3D surfaces from noisy camera parameters. We show that we can learn camera parameters together with learning the surface representation, and demonstrate good quality 3D surface reconstruction even with noisy camera observations.

## Introduction

The idea of representing 3D objects and scenes with neural networks, instead of traditional mesh-like representations, has gained significant traction recently. In (Mildenhall et al. 2020), an approach called *NeRF* is proposed, Neural Radiance Fields, where a neural network is used together with a volumetric representation to learn the representation of a scene from a collection of calibrated multi-view 2D images. However, a volumetric representation is not the best representation in many cases; for example, many objects like faces are better represented using surfaces. This was tackled in a follow-up work by (Wang et al. 2021a), called *NeuS*, who proposed to use neural implicit surfaces together with volume rendering for multi-view reconstruction. This year, (Wang et al. 2021b) addressed another shortcoming of classical NeRF’s, by extending the NeRF method to work on image data when camera calibration data is not present. This is achieved by jointly estimating the scene representation and optimising for the camera parameters, and the authors have shown promising results for frontal-facing scenes.

In this work, we propose to marry the benefits of each of these approaches: We propose a method to learn a neural implicit surface based representation of objects from noisy camera observations. We show that the classical NeuS method fails to learn an object completely if camera parameters are not precise, whereas our approach succeeds. Furthermore, what is often taken as ground-truth camera parameters

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: (*Left*): Reconstruction using (Wang et al. 2021a) with ground truth camera parameters. (*Centre*): Reconstruction using (Wang et al. 2021a) with noisy camera parameters, where the approach completely fails. (*Right*): Reconstruction using the proposed approach with noisy camera parameters. For each of the two objects, the lower images represent the actual image of the object while the upper image is the rendered image of the reconstructed surface.

come themselves from a multi-view reconstruction software like COLMAP, and thus likely with estimation errors. By making camera parameters learnable, the approach gains the ability to correct for these errors, and to potentially achieve better quality than methods that take camera parameters as given truth.

## Methodology

We draw inspiration from (Wang et al. 2021b), where the authors make the camera parameters of a NeRF learnable to allow them to converge to values that result in desirable 3D scene reconstructions. While (Wang et al. 2021b) works with unknown camera parameters for only forward-facing input images with rotational and translational perturbations of up to  $\pm 20^\circ$ , our approach works successfully on images from  $360^\circ$  view angles. The latter results in a problem that is much harder to solve and is prone to local optima that do

not produce good results when trying to learn the camera parameters from scratch.

The implicit surface based network of (Wang et al. 2021a), *NeuS*, consists of two MLPs to encode a signed distance function (SDF) and colour, respectively. Based on that, we add two modules that make both the extrinsic and intrinsic camera parameters learnable (see the Supplementary Material<sup>1</sup> heading *Network* for the full network architecture). The extrinsic parameters are expressed as a  $4 \times 4$  camera-to-world space transformation matrix with  $T_{wc} = [R|t]$  where  $R \in SO(3)$  and  $t \in \mathbb{R}^3$  denote the camera rotation and translation, respectively. The camera intrinsics include the focal length and the principal point, using a pinhole camera model. We take the sensor’s centre as the principal point, and we assume that the same camera takes all input images. Thus, estimating the camera intrinsics simplifies to finding the focal length.

The network, including the camera parameters, is trained on images of a scene using a weighted linear combination of Eikonal-loss (Gropp et al. 2020), colour-loss, and mask-loss masking out the background. The loss is then back-propagated through the network, and back to the camera estimation modules. We further use separate optimisers for the learnable extrinsics and intrinsics, along with their own learning rate schedulers, to help the values to converge efficiently.

## Experiments & Results

Case	Reconstruction loss	PSNR
bmvs_man-baseline-gt	0.032	36.1
bmvs_man-baseline-noisy	0.948	9.3
bmvs_man-learnable-gt	<b>0.028</b>	<b>37.5</b>
bmvs_man-learnable-noisy	0.038	33.8
bmvs_bear-baseline-gt	0.102	23.9
bmvs_bear-baseline-noisy	0.690	10.9
bmvs_bear-learnable-gt	0.115	23.4
bmvs_bear-learnable-noisy	<b>0.098</b>	<b>26.2</b>

Table 1: We show results for two objects and four cases. Our approach with *learnable* camera parameters performs as well or better than the *baseline* (vanilla NeuS), whereas vanilla NeuS completely fails if camera parameters are noisy.

We consider four cases to evaluate the performance of our method: We compare vanilla NeuS (*baseline*) to our approach with *learnable* camera parameters, both with ground-truth (*gt*) and with noisy camera parameters (*noisy*). We evaluate our approach quantitatively and qualitatively on two objects from the BlendedMVS dataset (Yao et al. 2020), *bmvs\_bear* and *bmvs\_man*.

For the extrinsics parameters, we sample noise from a normal distribution with a mean of zero and a standard deviation of 0.1. While for the intrinsics parameters, we add noise with

<sup>1</sup>Our supplementary material can be found on <https://arxiv.org/abs/2210.01548>.

a mean of zero and standard deviation equal to 0.2 times the mean value of the learnable parameters to make the noise of the same order as the values in the intrinsics matrix. Under these noise levels, we run the training process for 300,000 iterations while using the same hierarchical sampling as in (Wang et al. 2021a).

We find that the baseline NeuS completely fails, as soon as camera parameters are not accurate. In the noisy camera parameter setting, our method produces reconstructions that are visually indistinguishable from that of baseline NeuS and produces a final reconstruction loss and Peak Signal-to-Noise Ratio (PSNR) similar to the baseline with an equal number of iterations. Table 1 summarises the results, and Figure 1 shows example renderings for two objects.

## Conclusions & Future Work

While comparing our work to the current 3D surface reconstruction state-of-the-art in a noisy parameter setting, we have found that the existing method fails as soon as camera parameters are not accurate. Our method outperforms the existing method by being able to learn an equal or more accurate object representation, even in the presence of significant noise in the camera parameters. Our work broadens the use cases of neural implicit surface based object reconstruction, by removing the need for accurate camera calibration information, and increasing the robustness to errors.

However, there is still much scope for work in this field, as reconstructing the surface from completely unknown camera parameters is still an open problem for  $360^\circ$  view angles. To tackle this, we would like to investigate an approach akin to (Jeong et al. 2021; Lin et al. 2021), who learn a volumetric NeRF from unknown camera parameters. Furthermore, we plan to apply our method to a multi-view reconstruction benchmark, where 3D shape accuracy is evaluated.

## References

- Gropp, A.; Yariv, L.; Haim, N.; Atzmon, M.; and Lipman, Y. 2020. Implicit Geometric Regularization for Learning Shapes. In *ICML*.
- Jeong, Y.; Ahn, S.; Choy, C. B.; Anandkumar, A.; Cho, M.; and Park, J. 2021. Self-Calibrating Neural Radiance Fields. In *ICCV*. IEEE.
- Lin, C.; Ma, W.; Torralba, A.; and Lucey, S. 2021. BARF: Bundle-Adjusting Neural Radiance Fields. In *ICCV*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV (1)*, LNCS. Springer.
- Wang, P.; Liu, L.; Liu, Y.; Theobalt, C.; Komura, T.; and Wang, W. 2021a. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. In *NeurIPS*.
- Wang, Z.; Wu, S.; Xie, W.; Chen, M.; and Prisacariu, V. A. 2021b. NeRF--: Neural Radiance Fields Without Known Camera Parameters. *CoRR*, abs/2102.07064.
- Yao, Y.; Luo, Z.; Li, S.; Zhang, J.; Ren, Y.; Zhou, L.; Fang, T.; and Quan, L. 2020. BlendedMVS: A Large-Scale Dataset for Generalized Multi-View Stereo Networks. In *CVPR*.