Relative Illumination Fields:
Learning Medium and Light Independent Underwater Scenes

Supplementary Material

1. Videos

The supplementary material contains extended video results for the experiments discussed in the main paper in Fig.4 that show consistent results from many perspectives, both for scene colors as well as for the lights.

2. Evaluation on Color Charts

We additionally assess the effectiveness of water removal by comparing to Levy et al. [10] based on relative scene color accuracy, tested on 3 different Jerlov-like tank waters mixed with food colorants and a scattering agent. For each we capture image sequences with a GoPro Hero9 with attached artificial illumination. Restoration is then evaluated for each water type using a MacBeth chart.

	Type `IB`	Type `II`	Type `3C`
Levy *et al.* [10]	**5.71**	5.83	7.76
Proposed	11.70	**3.72**	**5.36**

Table 2. Evaluation of water removal performance using a Macbeth chart based method, reporting angular RGB error in degrees[6]. For both methods, we pre-scale the restored value of the white MacBeth field to the RGB value predicted from spectrophotometer readings of the patch and then transformed into GoPro space.

We report average angular errors from color chart images in the held-out test set not used during training. Test image poses are estimated using structure-from-motion and the evaluated methods synthesize images from these poses. With both methods we predict wideband RGB albedo resp. restored images. In order to avoid the error being contaminated by global white balancing effects, we pick the white MacBeth field and compute global factors for R,G and B respectively in both restorations such that the white fields become equal. Afterwards we compute the error of Finlayson et al.[6] with respect to independently obtained spectrophotometer readings projected to GoPro space on all other fields. Note that the implementation of [10] requires 16 bit integers, where we noticed quantization effects during our fusion of input image to a common HDR space. If we reimplemented the interface to float similarly to what our method is using, results for [10] would likely slightly improve. We still believe that the trend in the data will persist that the more turbid/denser waters exibit strong light cones and scattering effects that our method tends to compensate better. Again, please note that [10] is not designed to cope with underwater light cones, but we still find it instructive to compare our method to something.

3. Real-World Demonstration

We additionally, demonstrate the method on two real-world ocean datasets, featuring significantly varying water properties and camera-light systems, see Fig. 7. Despite having no ground truth, Fig. 8 demonstrates qualitatively good results of our approach (please zoom in!), showing that the training converges reasonably without heavy artifacts although the sequences had not been captured with NeRF in mind.

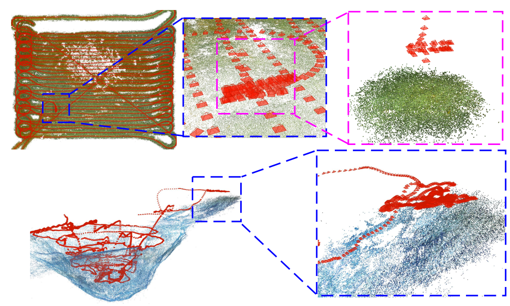

Figure 7. SfM results of two real-world ocean datasets. The first set was acquired by an AUV on a coastal site with a lawnmower-pattern trajectory. The second set is taken by an ROV at a deep-sea crater. Left: trajectory overview. Right: the selected patch for evaluation.

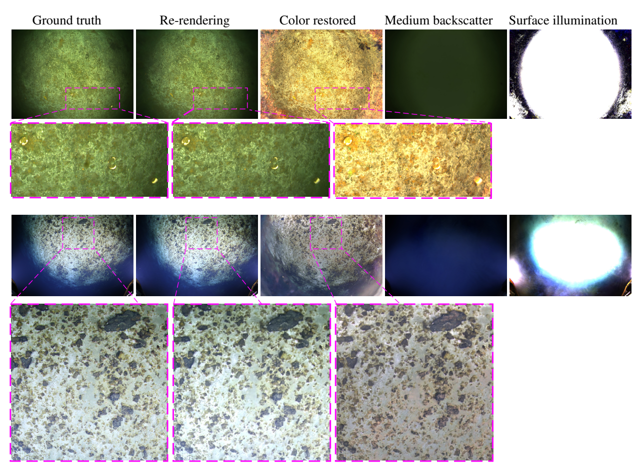

Figure 8. Results of our approach on two real-world datasets (top: greenish turbid coastal water, bottom: clearer blueish deep ocean).