

ERRATA

EFFICIENT 4D NEURAL PATH GUIDING

In Sec. 6.2 we compared the performance of decomposed feature grids to two alternatives for conditioning the network on an additional dimension (Fig. 4), i.e., the Fourier positional encoding (frequency encoding) and a 4D hashed feature grid. It turns out that these alternatives could perform better using a more reasonable configuration. We explored different choices of these hyperparameters and used those well-behaved ones for a more fair comparison.

The updated experiment (Fig. 4) with refined parameters and related discussions, as well as specific parameters we used are included in the revised author version of the paper in this page. We also show the main results below. As the result shows, the more complex feature-grid-based generally achieves slightly better reconstruction quality than the simpler frequency encoding. However, in simpler cases where modeling the extra correlation brings less benefits (e.g., the `Liquids` scene), the more lightweight approaches like the frequency encoding might be a more favorable choice.

RelMSE	3D+Aux	4D+Hash	Ours(4D)	Ours(Progressive)
Living-Room	0.052 (97s)	0.045 (105s)	0.040 (103s)	0.037 (97s)
Tower	0.051 (112s)	0.045 (122s)	0.041 (117s)	0.035 (112s)
Liquids	0.205 (78s)	0.207 (85s)	0.194 (86s)	0.203 (85s)

TABLE 1. The relative mean squared error (RelMSE) of each alternative on a subset of example scenes. Experiment conducted on an RTX 4080S.