

Active Exploration for Neural Global Illumination of Variable Scenes

ANONYMOUS AUTHOR(S)

SUBMISSION ID: 300

ACM Reference Format:

Anonymous Author(s). 2022. Active Exploration for Neural Global Illumination of Variable Scenes. *ACM Trans. Graph.* 1, 1 (February 2022), 4 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 SELECTED VIEWS

We show the chosen views that we use in our quantitative evaluation in Figure 1.

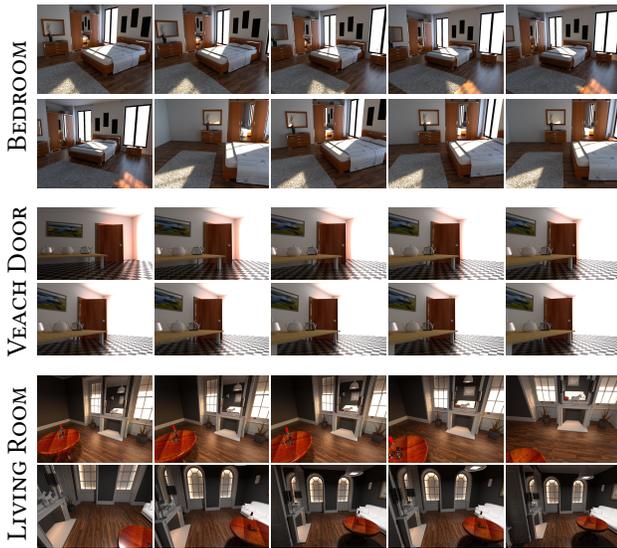


Fig. 1. The chosen views, which correspond to scene configurations with complex illumination effects, that were used for the ablations.

2 COMPARISON TO CNSR

We show additional results for same quality and same time comparisons against CNSR [Granskog et al. 2020] in Figure 6.

3 COMPARISON TO ANF

In Figure 2 we display a sample of ground truth images used during the finetuning of the ANF [Işık et al. 2021] pretrained model on our scenes. Please observe how the complex caustic effects that the models fails to reproduce, even after finetuning, exist in the ground truths. The amount of noise in the ground truths is equivalent to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.
0730-0301/2022/2-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

that in ours but our model is able to both learn these effects and average out the noise in world space during training.



Fig. 2. Sample ground truth images used for finetuning ANF on our scenes.

4 COMPARISON TO GT

In Figure 7 we provide difference images for the comparison of our method to ground truth, using the MAPE metric, to help with visual inspection.

5 NETWORK ARCHITECTURE

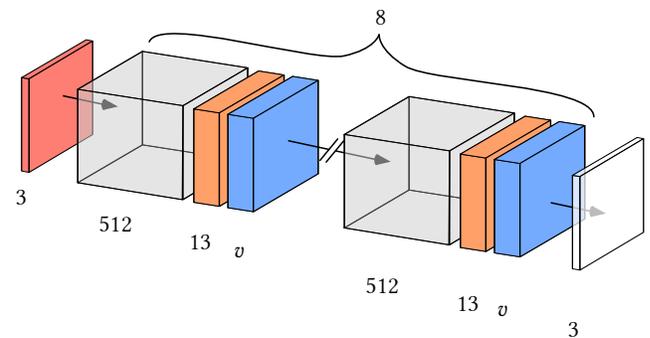


Fig. 3. The architecture of our generator. The positional buffer is shown in red, all the G-Buffers in orange, the explicit scene representation vector v in blue and the output in white.

The architecture of our generator is the Pixel Generator proposed by Granskog et al. [2020] with a preconditioning on position (Fig. 3). The Pixel Generator is an MLP (we use leaky ReLU activations)

with skip connections on every layer. We map the position buffer (red) from 3 to 512 channels and then we concatenate all the G-buffers (orange) and explicit vector v (blue) at each layer. The total hidden layers are 8 with 512 hidden features. In Figure 4 we show that using a smaller network can provide acceptable results and lower inference speed, resulting in 13 FPS in our prototype Python implementation.

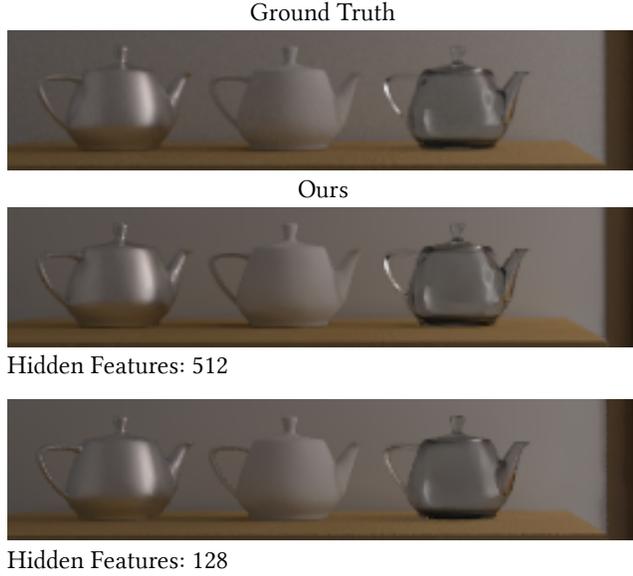


Fig. 4. Using 128 hidden features results in acceptable results and higher frame rates, but lower quality compared to using 512.

6 MCMC STATES LIFESPAN

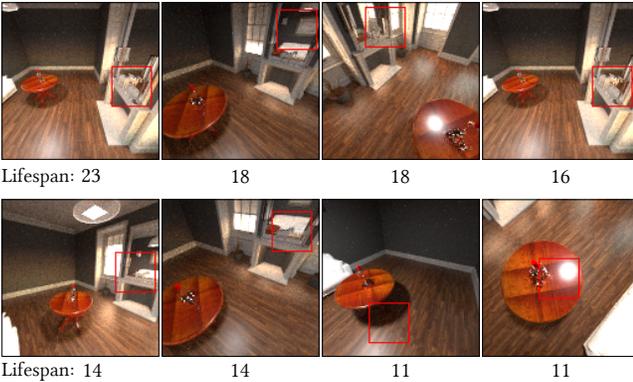


Fig. 5. The 8 longest lifespan MCMC states when training on the LIVING ROOM scene.

In order to evaluate what type of effects our Active Exploration focuses on, we visualize the MCMC states with the longest lifespan (consecutive times being the current state) for the LIVING ROOM

scene in Figure 5. We observe that our Active Exploration spends more time on effects that require more training to be represented such as reflections, glossy highlights and shadows. During training only the red patch would be rendered and used for training, here we render the whole image for visualization purposes.

7 SAMPLE REUSE DERIVATION

Given the two options to either reuse or generate a new sample with respective likelihood l_{exist} and l_{new} , a simple Bernoulli distribution that respect the likelihood ratio has a probability p of reusing defined by:

$$p = \frac{l_{new}}{l_{exist} + l_{new}}$$

This Bernoulli distribution can further be skewed as to favor the reuse case by dividing the likelihood of the reuse case l_{exist} by α :

$$p = \frac{l_{new}}{\frac{l_{exist}}{\alpha} + l_{new}}$$

For instance setting alpha to 99 skews the probability distribution so that for equal likelihood $p = \frac{99}{100}$

We then assume that the losses $Loss_{new}$ and $Loss_{exist}$ represent the negative log-likelihood of the network output with respect to a probability distribution parameterized by the ground truth, which for the L2 loss case would be a Normal distribution centered around the ground truth value and for the L1 loss is a Laplace distribution also centered around the ground truth value. We thus have:

$$\sigma(Loss_{exist} - Loss_{new} + \beta) = \frac{e^{Loss_{exist} - Loss_{new} + \beta}}{1 + e^{Loss_{exist} - Loss_{new} + \beta}}$$

$$\sigma(Loss_{exist} - Loss_{new} + \beta) = \frac{e^{-Loss_{new}}}{e^{-Loss_{exist}} e^{-\beta} + e^{-Loss_{new}}}$$

$$\sigma(Loss_{exist} - Loss_{new} + \beta) = \frac{l_{new}}{\frac{l_{exist}}{e^{\beta}} + l_{new}}$$

Which inspired our reuse strategy.

REFERENCES

- Jonathan Granskog, Fabrice Rousselle, Marios Papas, and Jan Novák. 2020. Compositional neural scene representations for shading inference. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 135–1.
- Mustafa İşık, Krishna Mullia, Matthew Fisher, Jonathan Eisenmann, and Michaël Gharbi. 2021. Interactive Monte Carlo Denoising using Affinity of Neural Features. *ACM Transactions on Graphics (TOG)* 40, 4, Article 37 (2021), 13 pages. <https://doi.org/10.1145/3450626.3459793>

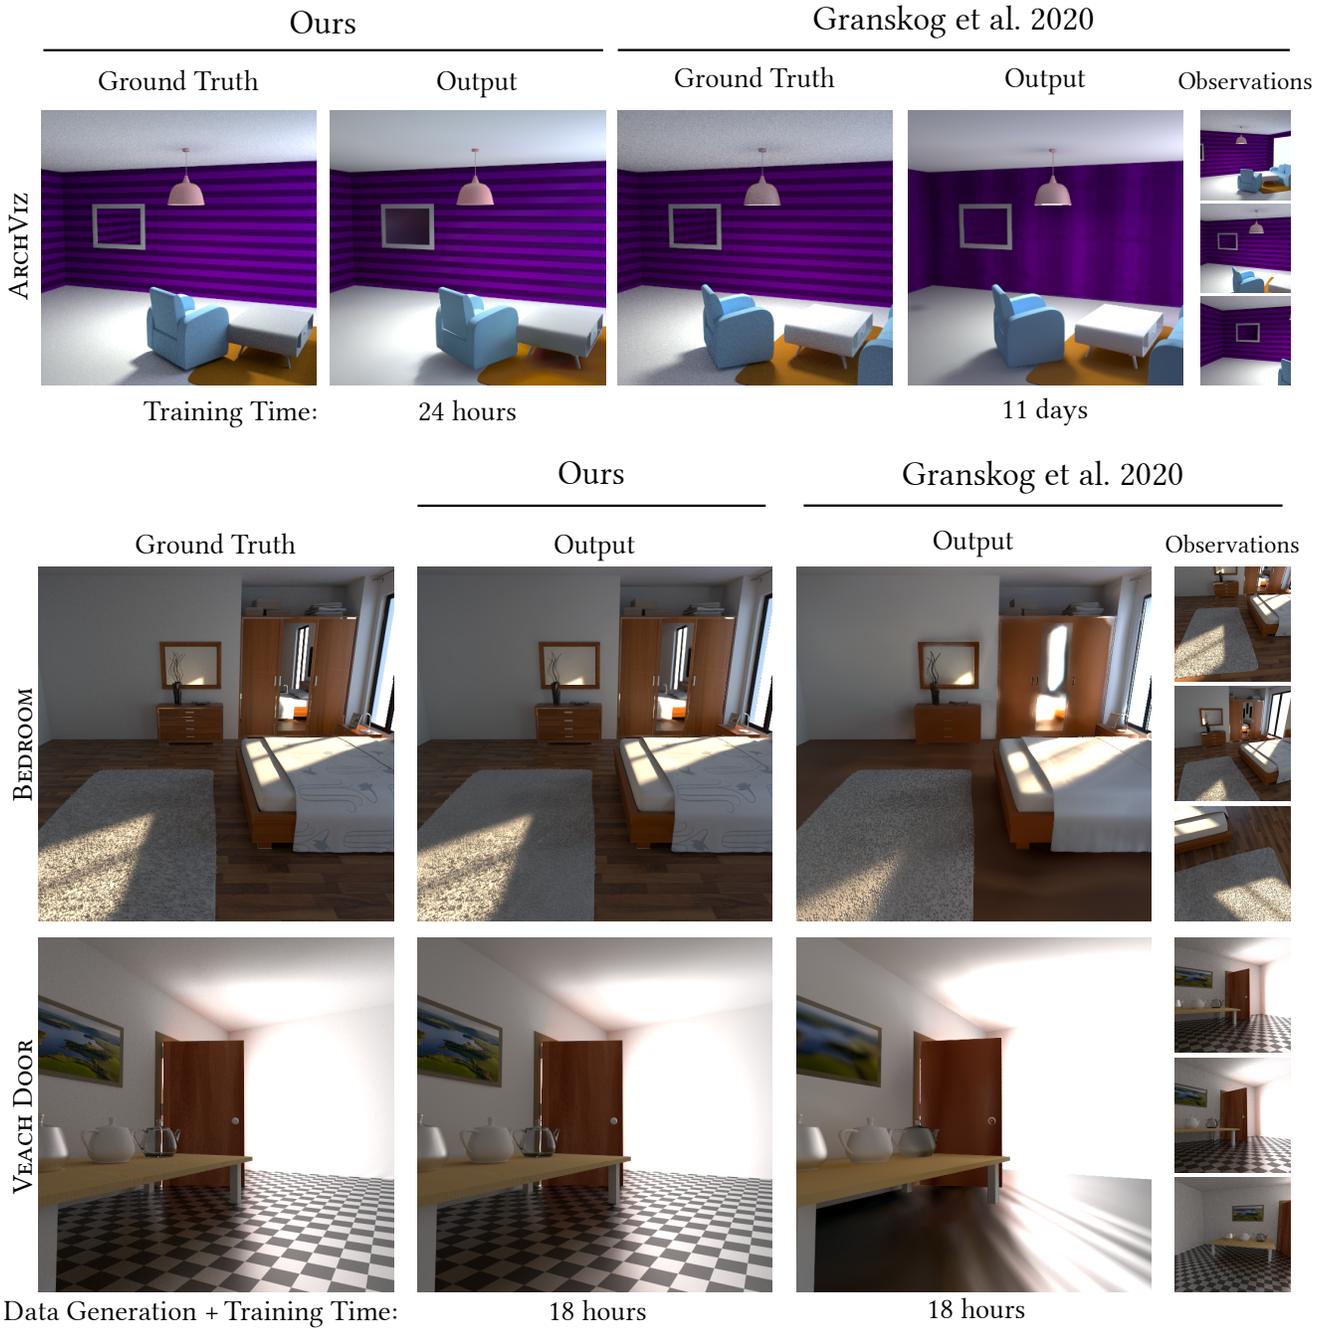
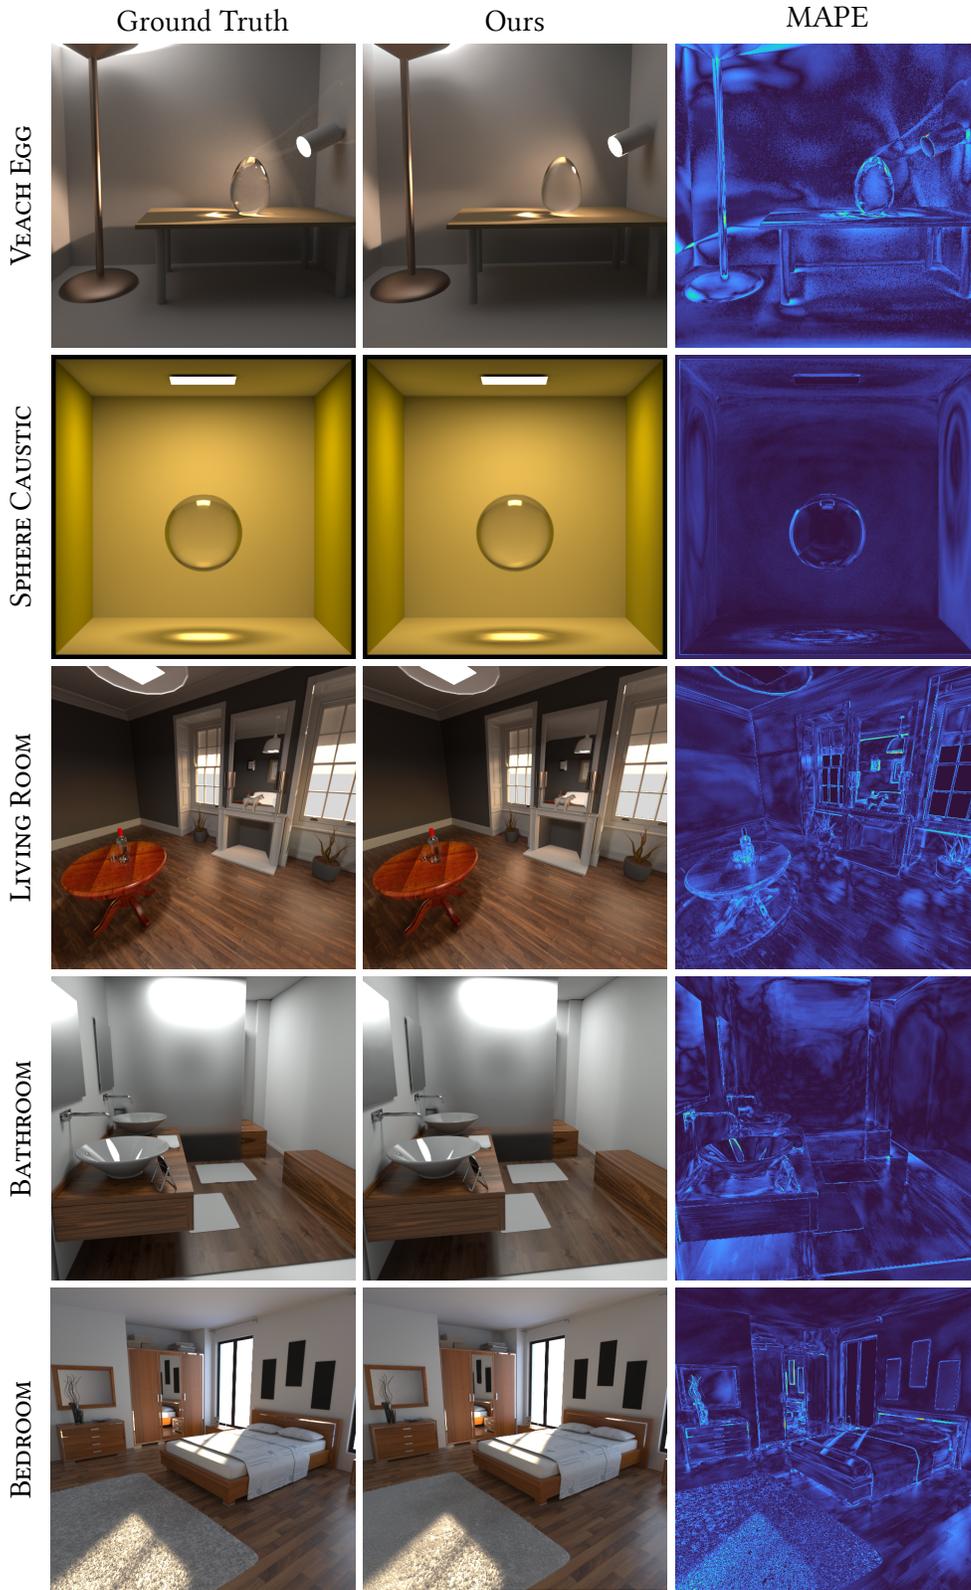


Fig. 6. Additional same quality and same time results with CNSR [Granskog et al. 2020].

343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399



400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456

Fig. 7. Comparison of our method to ground truth with additional difference images, using the MAPE metric, for visual inspection.