Neural James-Stein Combiner for Unbiased and Biased Renderings (Supplementary Report)

JEONGMIN GU, Gwangju Institute of Science and Technology, South Korea
JOSE A. IGLESIAS-GUITIAN, Universidade da Coruña - CITIC, Spain
BOCHANG MOON, Gwangju Institute of Science and Technology, South Korea

CCS Concepts: • **Computing methodologies → Ray tracing**.

Additional Key Words and Phrases: James-Stein estimator, James-Stein combiner, Monte Carlo rendering, learning-based denoising

ACM Reference Format:
Jeongmin Gu, Jose A. Iglesias-Guitian, and Bochang Moon. 2022. Neural James-Stein Combiner for Unbiased and Biased Renderings (Supplementary Report). *ACM Trans. Graph.* 41, 6, Article 262 (December 2022), 4 pages. https://doi.org/10.1145/3550454.3555496

1 ADDITIONAL EVALUATION

1.1 Our combination results using classical alternatives

We employ a deep neural network to estimate the unknown variance of an unbiased input and the per-pixel blending factor (see Eqs. 6 and 12 in the main paper). A simple alternative is to apply a Gaussian filter to the sample variance of the unbiased input colors and set the alpha-blending factor to a fixed value (e.g., 0.5) so that the James-Stein combiner can be performed without relying on the neural network. Specifically, we assign two fixed bandwidths to the Gaussian filter (i.e., $\sigma = 1, 3$). We also test a more straightforward approach that directly employs the input sample variance without additional smoothing.

In Fig. 1, we visualize the variances estimated by the simple alternatives and our neural network, respectively, and show their combination results via our localized JS combiner. As shown in the figure, the JS combiner produces lower errors than its unbiased inputs when using any of the simple estimation approaches above. Nonetheless, the simple alternatives tend to leave high-frequency noise in their results for the Glass-of-water scene where firefly noise easily shows up. Note that the estimated per-pixel variances play a key role for the JS combiner since they directly control its shrinkage factor, which should be determined to minimize the errors in the final result. A deep neural network allows us to infer per-pixel optimal parameters by a learning process (e.g., supervised learning), generating more robust combination results than the tested alternatives.

Authors' addresses: Jeongmin Gu, Gwangju Institute of Science and Technology, Gwangju, South Korea, jeong755@gm.gist.ac.kr; Jose A. Iglesias-Guitian, Universidade da Coruña - CITIC, A Coruña, Spain, j.iglesias.guitian@udc.es; Bochang Moon, Gwangju Institute of Science and Technology, Gwangju, South Korea, moonbochang@gmail.com.

1.2 Equal-sample comparisons

In Figs. 2 and 3, we compare our method with existing post-denoisers, PD [Firmino et al. 2022], DC [Back et al. 2020], and ED [Zheng et al. 2021], given equal-sample counts. Note that the main paper includes equal-time comparisons (see Figs. 10 and 12 in the paper). As shown in the figures, the existing methods sometimes fail to improve their input, e.g., PD for the Veach-Ajar and Staircase, DC for the Veach-Ajar, and ED with NFOR for the Dragon and Curly-hair scenes. Our method, however, outperforms the previous techniques while improving our input denoisers consistently.

REFERENCES

Jonghee Back, Binh-Son Hua, Toshiya Hachisuka, and Bochang Moon. 2020. Deep combiner for independent and correlated pixel estimates. *ACM Trans. Graph.* 39, 6 (2020), 12 pages.

Steve Bako, Thijs Vogels, Brian McWilliams, Mark Meyer, Jan Novák, Alex Harvill, Pradeep Sen, Tony Derose, and Fabrice Rousselle. 2017. Kernel-predicting convolutional networks for denoising Monte Carlo renderings. *ACM Trans. Graph.* 36, 4 (2017), 14 pages.

Arthur Firmino, Jeppe Revall Frisvad, and Henrik Wann Jensen. 2022. Progressive Denoising of Monte Carlo Rendered Images. *Computer Graphics Forum* (2022). https://doi.org/10.1111/cgf.14454

Jiaqi Yu, Yongwei Nie, Chengjiang Long, Wenju Xu, Qing Zhang, and Guiqing Li. 2021. Monte Carlo Denoising via Auxiliary Feature Guided Self-Attention. *ACM Trans. Graph.* 40, 6 (2021), 13 pages.

Shaokun Zheng, Fengshi Zheng, Kun Xu, and Ling-Qi Yan. 2021. Ensemble denoising for Monte Carlo renderings. *ACM Trans. Graph.* 40, 6 (2021), 17 pages.

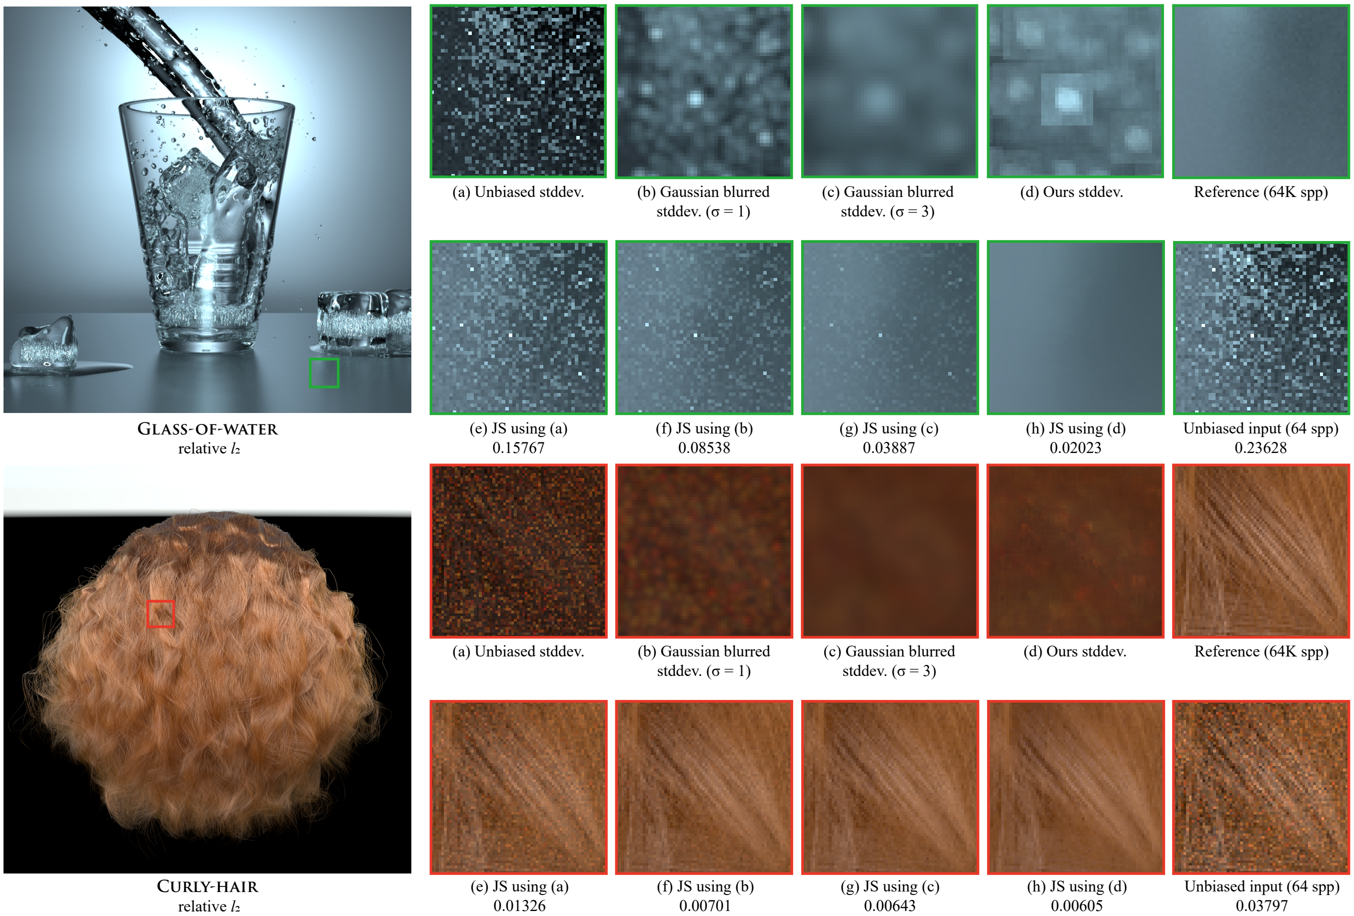

Fig. 1. James-Stein combination results with and without the use of a deep neural network. We show the sample variances of the unbiased input colors (a), their filtering results by Gaussian filtering ((b) and (c)), and our estimated variance using a neural network (d). We take the square root of the estimated variances, i.e., estimated standard deviations, to show the values more clearly. When taking the sample variance without additional smoothing, the JS combination results (e) exhibit noise propagated from the unbiased input. We can mitigate the noise using the classical Gaussian filter, but their results (f) and (g) still suffer from residual noise, especially for the GLASS-OF-WATER scene that exhibits fireflies. Our current approach (h), which exploits a deep neural network, enables the JS combiner to produce more accurate results without such noticeable noise.

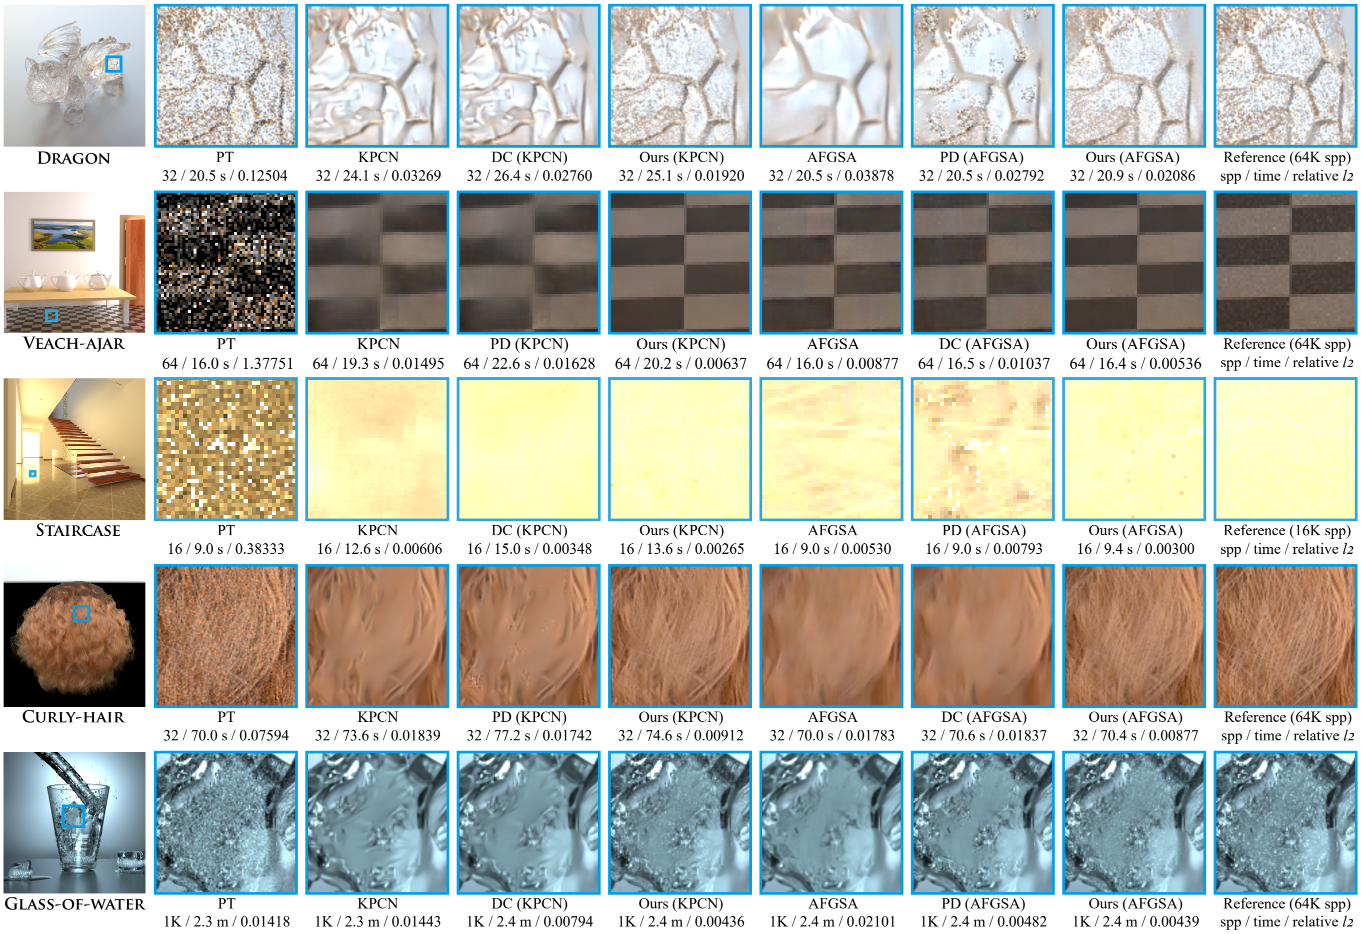

Fig. 2. Equal-sample comparisons between our technique and the post-denoisers, DC and PD. We test the two recent learning-based denoisers, KPCN [Bako et al. 2017] and AFGSA [Yu et al. 2021], as the input denoisers of the post-denoisers and our technique. While DC and PD sometimes produce higher errors than the input denoisers, e.g., DC for the VEACH-AJAR and PD for the STAIRCASE and VEACH-AJAR scenes, our technique consistently improves the input methods.

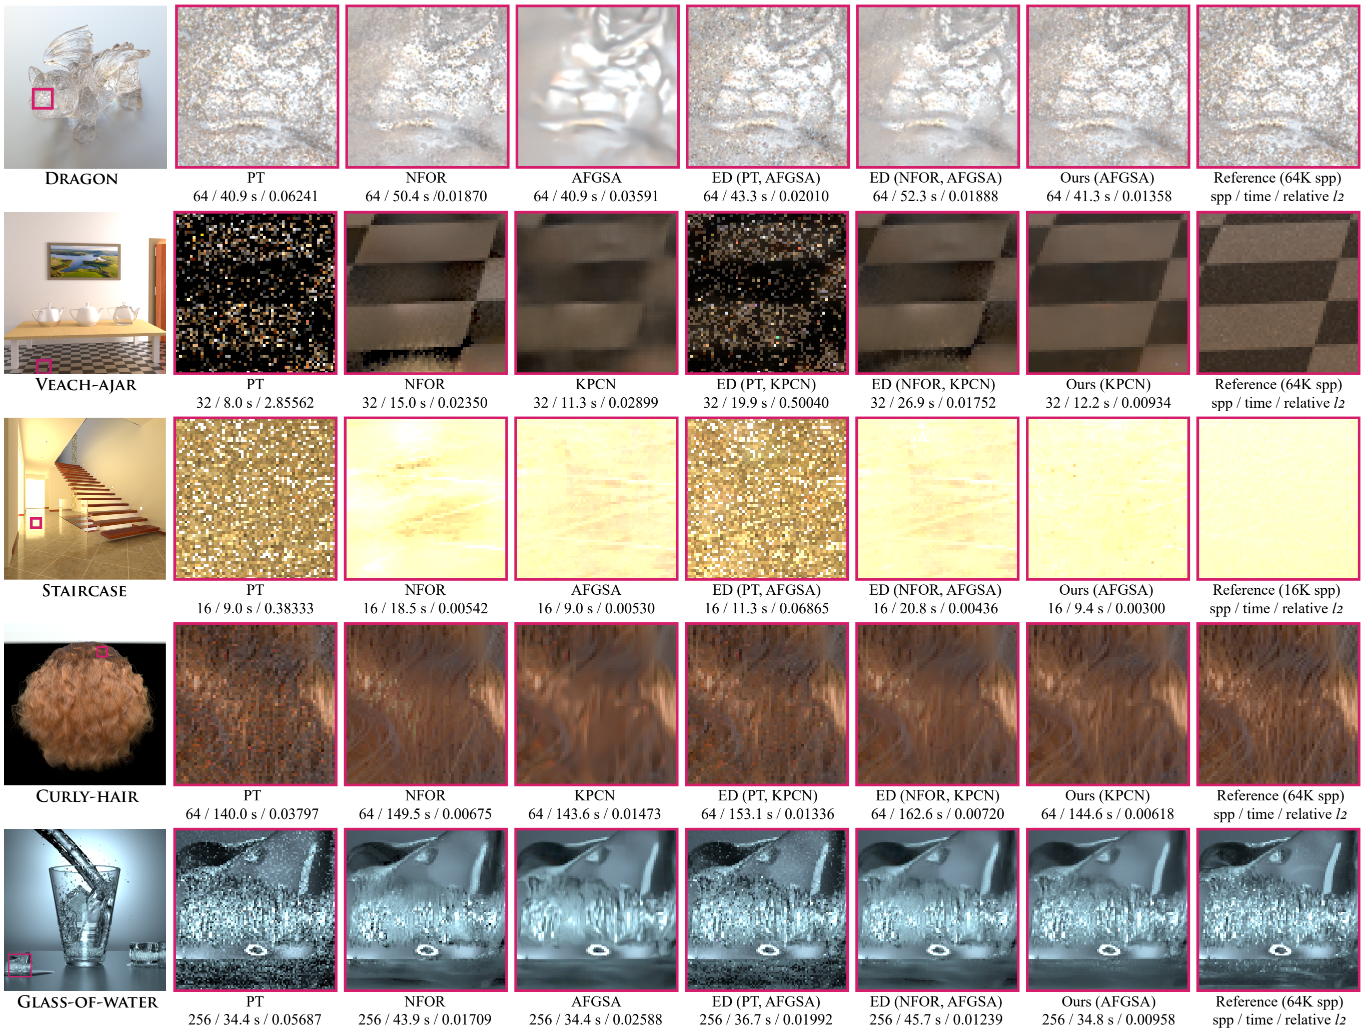

Fig. 3. Equal-sample comparisons between our technique and ED with two input configurations. As the first input configuration for ED, we use PT and a learning-based denoiser (KPCN and AFGSA), which is the same setting as ours. We also exploit a consistent denoiser (NFOR) and a learning-based denoiser for their input. It is noticeable that ED can become robust when it takes only reasonable estimates, like NFOR and KPCN or NFOR and AFGSA. Nevertheless, it sometimes fails to improve its inputs (see the results for the DRAGON and CURLY-HAIR scenes). On the other hand, our technique shows a consistent error reduction for the tested learning-based methods.