TM-BSN: Triangular-Masked Blind-Spot Network for Real-World Self-Supervised Image Denoising

Supplementary Material

S1. Ablation on Knowledge Distillation Strategies

Table S1. Quantitative comparison of knowledge distillation approaches on the SIDD [1] and DND [10] datasets.

Methods	SIDD Validation		SIDD Benchmark		DND Benchmark	
	PSNR↑(dB)	SSIM↑	PSNR↑(dB)	SSIM↑	PSNR↑(dB)	SSIM↑
Multi Teacher [2]	37.92	0.947	38.16	0.893	39.21	0.946
Recharged Distillation [5]	38.08	0.952	38.31	0.900	39.41	0.949

To enhance both computational efficiency and denoising performance, we distill the knowledge of a pretrained TM-BSN into a lightweight student network. We compare two representative knowledge distillation strategies for distilling our model, specifically the Multi-Teacher framework and Recharged Distillation (RD). Table S1 shows the quantitative results on the SIDD validation set, SIDD benchmark, and DND benchmark. The RD approach consistently outperforms the Multi-Teacher strategy across all evaluation settings, achieving higher PSNR and SSIM scores on every dataset. Based on these observations, we adopt Recharged Distillation as our default knowledge transfer strategy.

S2. Additional Qualitative Comparison

We provide additional visual comparisons to further illustrate the effectiveness of our method. Figures S1, S2, and S3 present qualitative results on the SIDD validation dataset, while Figures S4 and S5 provide additional examples from the DND benchmark. These results offer a more comprehensive comparison across diverse scenes and noise levels.

Across all figures, our method consistently produces cleaner and more visually faithful restorations compared to existing self-supervised denoising approaches. Competing methods often leave residual noise or suffer from oversmoothing, removing essential texture details along with the noise. In contrast, our method effectively suppresses noise while preserving high-frequency structures, resulting in sharper details and clearer patterns. These observations further validate the strong detail-preserving capability of TM-BSN under real-world noise conditions.

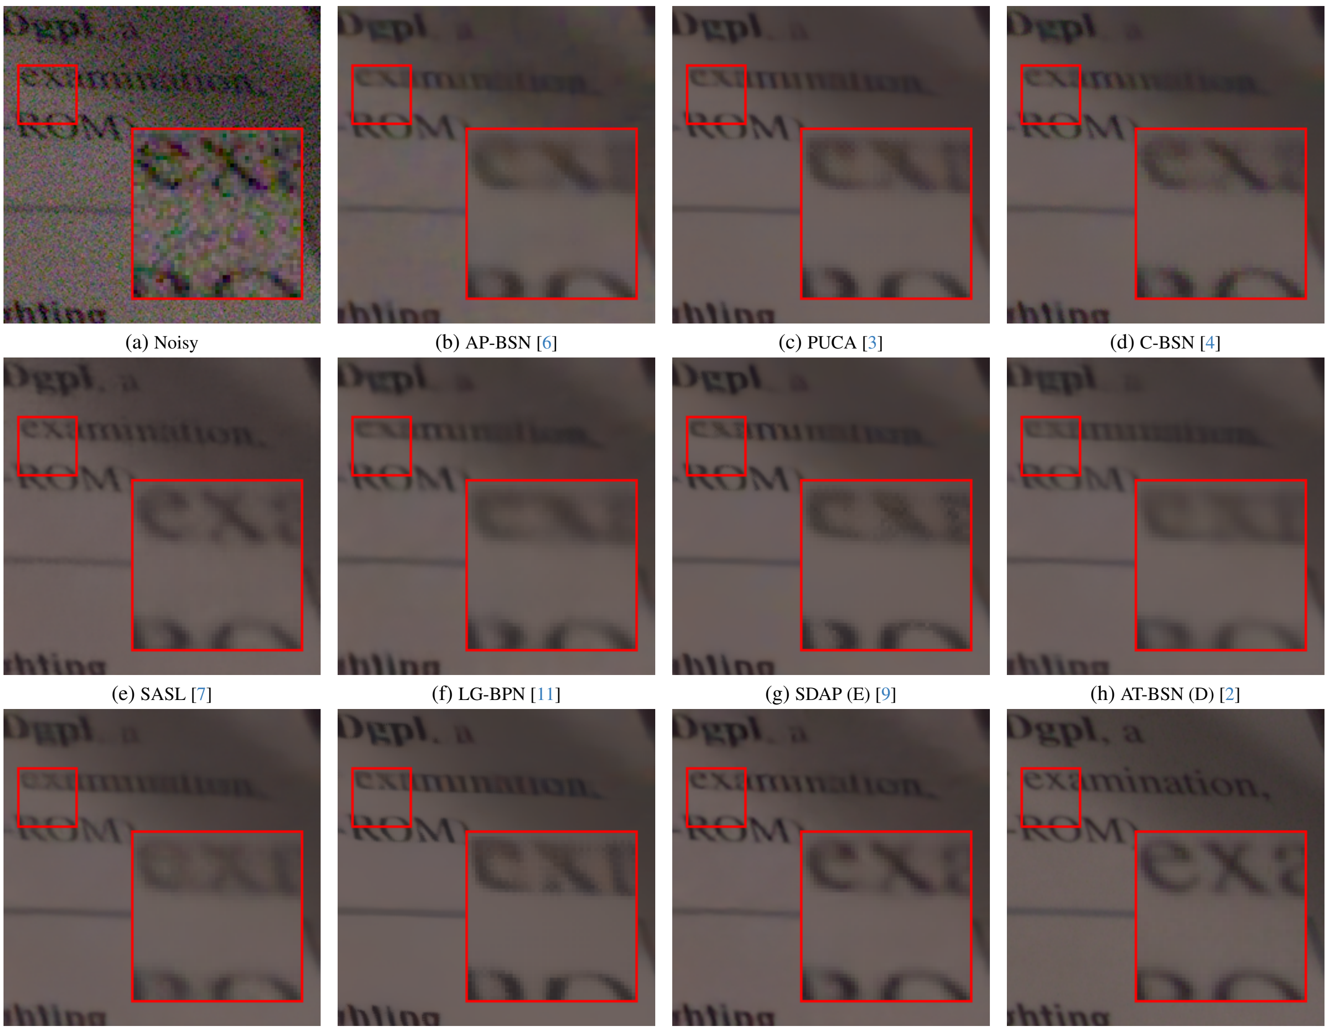

(a) Noisy (b) AP-BSN [6] (c) PUCA [3] (d) C-BSN [4]

(e) SASL [7] (f) LG-BPN [11] (g) SDAP (E) [9] (h) AT-BSN (D) [2]

(i) APR (RD) [5] (j) TBSN [8] (k) TM-BSN (D) (Ours) (l) GT

Figure S1. Additional qualitative comparison on the SIDD Validation dataset [1] (Sample A).

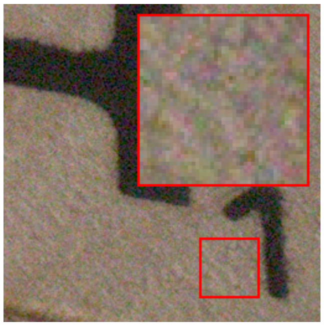
(a) Noisy

(b) AP-BSN [6]

(c) PUCA [3]
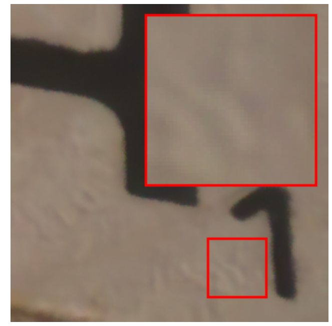
(d) C-BSN [4]
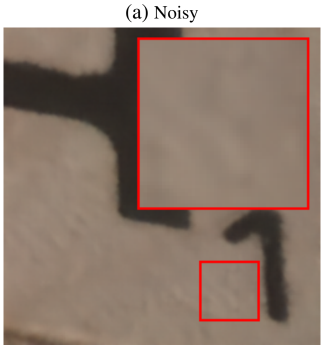
(e) SASL [7]
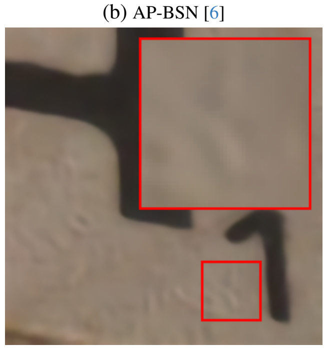
(f) LG-BPN [11]
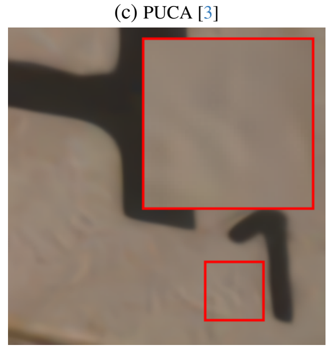
(g) SDAP (E) [9]
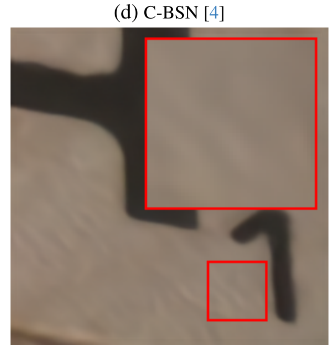
(h) AT-BSN (D) [2]
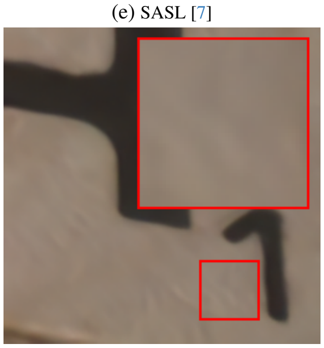
(i) APR (RD) [5]

(j) TBSN [8]
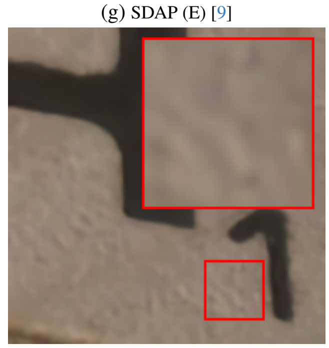
(k) TM-BSN (D) (Ours)
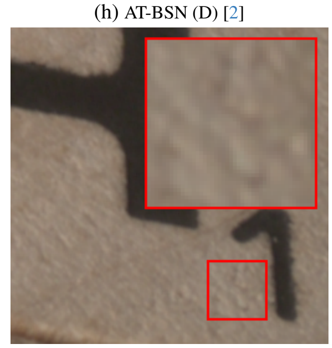
(l) GT

Figure S2. Additional qualitative comparison on the SIDD Validation dataset [1] (Sample B).

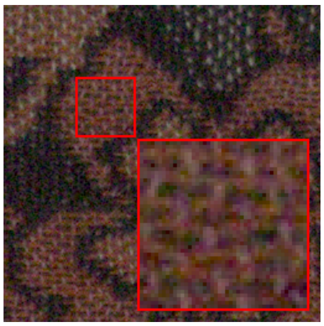

(a) Noisy

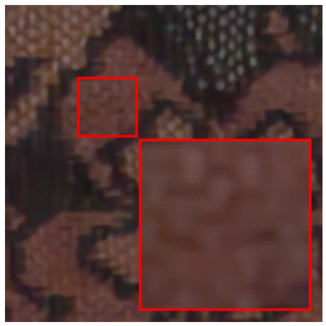

(b) AP-BSN [6]

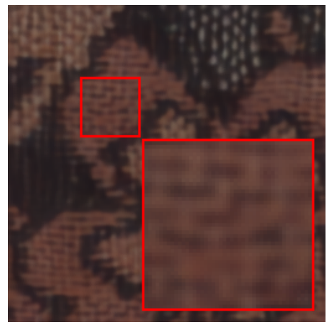

(c) PUCA [3]

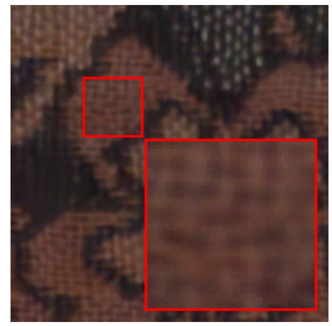

(d) C-BSN [4]

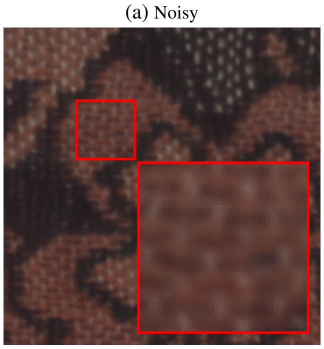

(e) SASL [7]

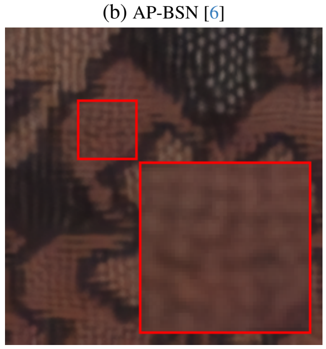

(f) LG-BPN [11]

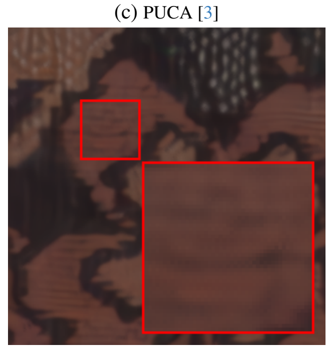

(g) SDAP (E) [9]

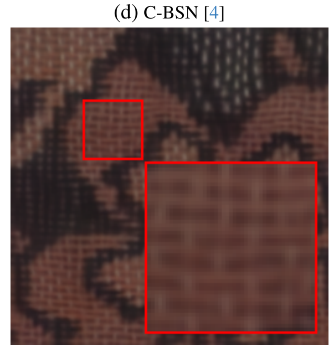

(h) AT-BSN (D) [2]

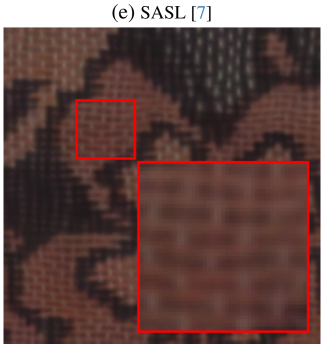

(i) APR (RD) [5]

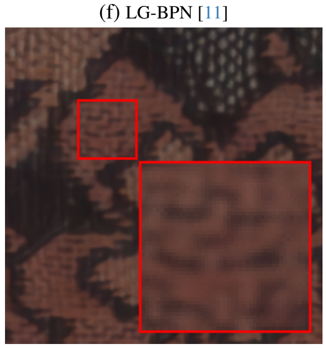

(j) TBSN [8]

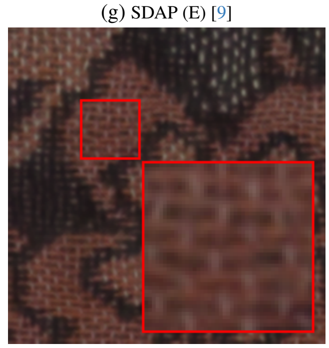

(k) TM-BSN (D) (Ours)

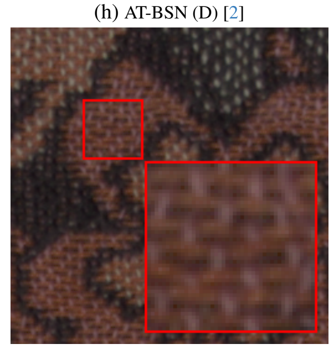

(l) GT

Figure S3. Additional qualitative comparison on the SIDD Validation dataset [1] (Sample C).

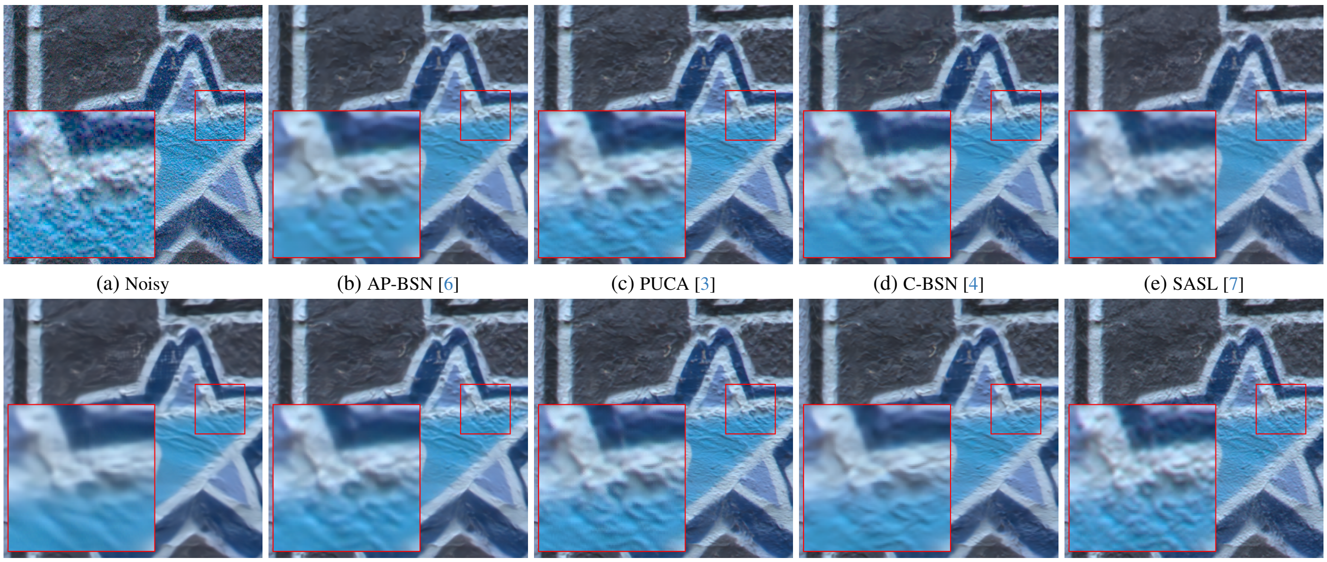

(a) Noisy (b) AP-BSN [6] (c) PUCA [3] (d) C-BSN [4] (e) SASL [7]

(f) AT-BSN (D) [2] (g) APR (RD) [5] (h) TBSN [8] (i) LG-BPN [11] (j) TM-BSN (D) (Ours)

Figure S4. Additional qualitative comparison on the DND benchmark [10] (Sample A).

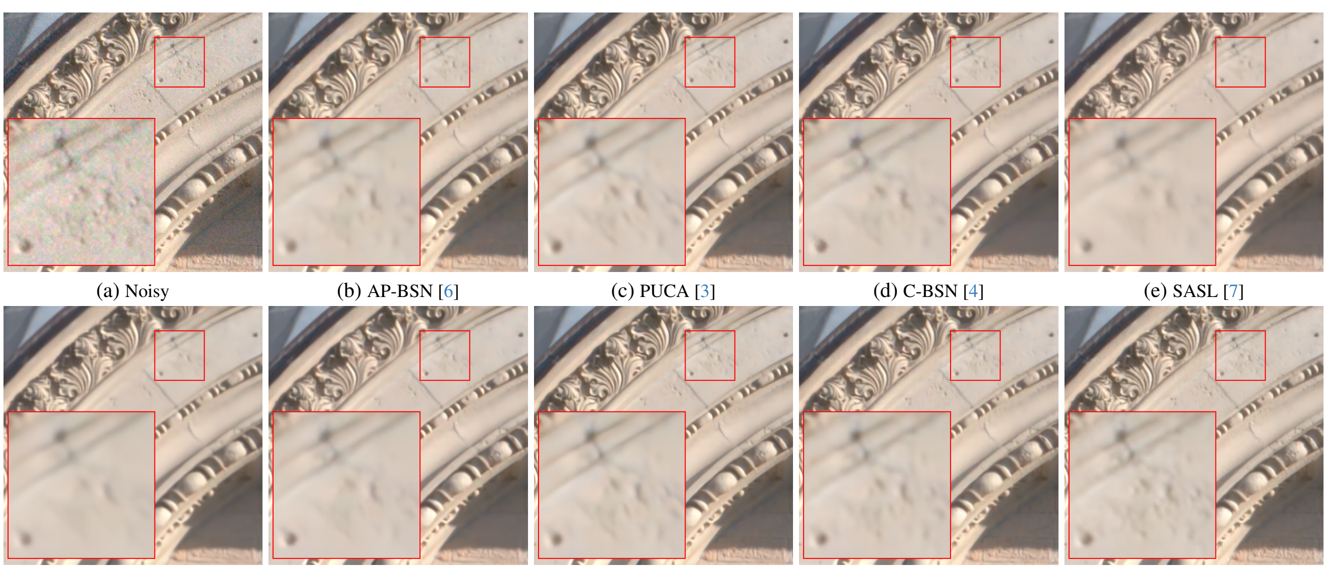

(a) Noisy (b) AP-BSN [6] (c) PUCA [3] (d) C-BSN [4] (e) SASL [7]

(f) AT-BSN (D) [2] (g) APR (RD) [5] (h) TBSN [8] (i) LG-BPN [11] (j) TM-BSN (D) (Ours)

Figure S5. Additional qualitative comparison on the DND benchmark [10] (Sample B).

References

[1] Abdelrahman Abdelhamed, Stephen Lin, and Michael S Brown. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1692–1700, 2018. 1, 2, 3, 4

[2] Shiyan Chen, Jiyuan Zhang, Zhaofei Yu, and Tiejun Huang. Exploring efficient asymmetric blind-spots for self-supervised denoising in real-world scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2814–2823, 2024. 1, 2, 3, 4, 5

[3] Hyemi Jang, Junsung Park, Dahuin Jung, Jaihyun Lew, Ho Bae, and Sungroh Yoon. Puca: patch-unshuffle and channel attention for enhanced self-supervised image denoising. *Advances in Neural Information Processing Systems*, 36:19217–19229, 2023. 2, 3, 4, 5

[4] Yeong Il Jang, Keuntek Lee, Gu Yong Park, Seyun Kim, and Nam Ik Cho. Self-supervised image denoising with downsampled invariance loss and conditional blind-spot network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12196–12205, 2023. 2, 3, 4, 5

[5] Hyunjun Kim and Nam Ik Cho. Apr-rd: Complemental two steps for self-supervised real image denoising. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4257–4265, 2025. 1, 2, 3, 4, 5

[6] Wooseok Lee, Sanghyun Son, and Kyoung Mu Lee. Ap-bsn: Self-supervised denoising for real-world images via asymmetric pd and blind-spot network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17725–17734, 2022. 2, 3, 4, 5

[7] Junyi Li, Zhilu Zhang, Xiaoyu Liu, Chaoyu Feng, Xiaotao Wang, Lei Lei, and Wangmeng Zuo. Spatially adaptive self-supervised learning for real-world image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9914–9924, 2023. 2, 3, 4, 5

[8] Junyi Li, Zhilu Zhang, and Wangmeng Zuo. Rethinking transformer-based blind-spot network for self-supervised image denoising. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4788–4796, 2025. 2, 3, 4, 5

[9] Yizhong Pan, Xiao Liu, Xiangyu Liao, Yuanzhouhan Cao, and Chao Ren. Random sub-samples generation for self-supervised real image denoising. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12150–12159, 2023. 2, 3, 4

[10] Tobias Plotz and Stefan Roth. Benchmarking denoising algorithms with real photographs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1586–1595, 2017. 1, 5

[11] Zichun Wang, Ying Fu, Ji Liu, and Yulun Zhang. Lg-bpn: Local and global blind-patch network for self-supervised real-world denoising. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18156–18165, 2023. 2, 3, 4, 5