Supplementary Materials

Yuhui Quan[1], Mingqin Chen[1], Tongyao Pang[2] and Hui Ji[2]
[1]School of Computer Science and Engineering, South China University of Technology, Guangzhou 510006, China
[2]Department of Mathematics, National University of Singapore, 119076, Singapore
csyhquan@scut.edu.cn, csmingqinchen@mail.scut.edu.cn, matpt@nus.edu.sg and matjh@nus.edu.sg

1. Details on Partial Convolution

Partial convolution [1] is originally designed for inpainting image holes, which allows progressively filling the holes from the outside to the inside. Let $\boldsymbol{k} \in \mathbb{R}^Z$ be the weights of a convolution kernel and $b \in \mathbb{R}$ the corresponding bias. Let $\boldsymbol{f} \in \mathbb{R}^Z$ denote the feature values (pixels values) for the current convolution (sliding) window and $\boldsymbol{m} \in \mathbb{R}^Z$ is the corresponding binary mask. The partial convolution at every location is expressed as

$$\boldsymbol{f}' = 1[\|\boldsymbol{m}\|_1 > 0](\boldsymbol{k}^\top(\boldsymbol{f} \odot \boldsymbol{m})\frac{Z}{\|\boldsymbol{m}\|_1} + b), \tag{1}$$

where $\odot$ is the Hadamard's product. It can be seen that the output of the function only depends on the unmasked inputs. The scaling factor $\frac{Z}{\|\boldsymbol{m}\|_1}$ applies appropriate scaling to adjust for the varying amount of valid (unmasked) inputs. At the beginning, we initialize the mask $\boldsymbol{m}$ such that it excludes the dropped pixels of the input Bernoulli sampled instance as well as those of the input images (*e.g.* in removing salt-and-pepper noise). After crossing the current PConv layer, we then update the mask for the next PConv layer as follows: if the convolution was able to condition its output on at least one valid input value, then we mark that location to be valid. This can be expressed as $\boldsymbol{m}' = 1[\|\boldsymbol{m}\|_1 > 0]$, which can be easily implemented as a part of forward pass. See [1] for more details.

2. Proof of Proposition 1

Proof. Rewrite the loss function as follows:

$$\begin{aligned}\sum_{m=1}^{M}\|\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}}_m)-\bar{\boldsymbol{y}}_m\|_{\boldsymbol{b}_m}^2 &= \sum_{m=1}^{M}\|\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}}_m)-\boldsymbol{y}\|_{\boldsymbol{b}_m}^2 = \sum_{m=1}^{M}\|\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}}_m)-(\boldsymbol{x}+\boldsymbol{n})\|_{\boldsymbol{b}_m}^2\\ &= \sum_{m=1}^{M}\|\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}})-\boldsymbol{x}\|_{\boldsymbol{b}_m}^2 + \sum_{m=1}^{M}\|\boldsymbol{n}\|_{\boldsymbol{b}_m}^2 - 2\sum_{m=1}^{M}((\boldsymbol{1}-\boldsymbol{b}_m)\odot\boldsymbol{n})^\top(\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}}_m)-\boldsymbol{x})\\ &= \sum_{m=1}^{M}\|\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}})-\boldsymbol{x}\|_{\boldsymbol{b}_m}^2 + \sum_{m=1}^{M}\|\boldsymbol{n}\|_{\boldsymbol{b}_m}^2 - 2\boldsymbol{n}^\top\big(\sum_{m=1}^{M}(\boldsymbol{1}-\boldsymbol{b}_m)\odot(\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}}_m)-\boldsymbol{x})\big).\end{aligned} \tag{2}$$

Regarding the second term in (2), its expectation is

$$\mathbb{E}_{\boldsymbol{n}}\big[\sum_{m=1}^{M}\|\boldsymbol{n}\|_{\boldsymbol{b}_m}^2\big] = \mathbb{E}_{\boldsymbol{n}}\big[\sum_{m=1}^{M}\|(\boldsymbol{1}-\boldsymbol{b}_m)\odot\boldsymbol{n}\|_2^2\big] = \sum_{m=1}^{M}\|(\boldsymbol{1}-\boldsymbol{b}_m)\odot\boldsymbol{\sigma}\|_2^2 = \sum_{m=1}^{M}\|\boldsymbol{\sigma}\|_{\boldsymbol{b}_m}^2. \tag{3}$$

Regarding the last term in (2), for simplicity we define

$$\boldsymbol{r} = \sum_{m=1}^{M}(\boldsymbol{1}-\boldsymbol{b}_m)\odot(\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}}_m)-\boldsymbol{x}) = \sum_{m=1}^{M}(\boldsymbol{1}-\boldsymbol{b}_m)\odot(\mathcal{F}_{\boldsymbol{\theta}}(\boldsymbol{b}_m\odot\boldsymbol{x}+\boldsymbol{b}_m\odot\boldsymbol{n}))-\boldsymbol{x}). \tag{4}$$

Note that $\mathcal{F}_{\boldsymbol{\theta}}(\boldsymbol{b}_m \odot \boldsymbol{x} + \boldsymbol{b}_m \odot \boldsymbol{n})$ contributes to $\boldsymbol{r}(i)$ only if $\boldsymbol{b}_m(i) = 0$. But in this case, $\boldsymbol{n}(i)$ is erased by $\boldsymbol{b}_m(i)$. This means that $\boldsymbol{n}(i)$ has no contribution to $\boldsymbol{r}(i)$. Together with that $\boldsymbol{n}(i)$ is independent of $\boldsymbol{n}(j)$ for any $i \neq j$, we can conclude that $\boldsymbol{r}(i)$ is independent to $\boldsymbol{n}(i)$ for all i. Therefore, we have

$$\mathbb{E}_{\boldsymbol{n}}[\boldsymbol{n}^\top \boldsymbol{r}] = (\mathbb{E}_{\boldsymbol{n}}[\boldsymbol{n}])^\top (\mathbb{E}_{\boldsymbol{n}}[\boldsymbol{r}]) = 0. \tag{5}$$

Combining (2), (3) and (5) yields

$$\begin{aligned} \mathbb{E}_{\boldsymbol{n}}\Big[\sum_{m=1}^{M} \|\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}}_m) - \bar{\boldsymbol{y}}_m\|_{\boldsymbol{b}_m}^2\Big] &= \mathbb{E}_{\boldsymbol{n}}\Big[\sum_{m=1}^{M} \|\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}}) - \boldsymbol{x}\|_{\boldsymbol{b}_m}^2\Big] + \mathbb{E}_{\boldsymbol{n}}\Big[\sum_{m=1}^{M} \|\boldsymbol{n}\|_{\boldsymbol{b}_m}^2\Big] - 2\mathbb{E}_{\boldsymbol{n}}[\boldsymbol{n}^\top \boldsymbol{r}] \\ &= \sum_{m=1}^{M} \|\mathcal{F}_{\boldsymbol{\theta}}(\widehat{\boldsymbol{y}}) - \boldsymbol{x}\|_{\boldsymbol{b}_m}^2 + \sum_{m=1}^{M} \|\boldsymbol{\sigma}\|_{\boldsymbol{b}_m}^2 . \end{aligned} \tag{6}$$

The proof is done.

□

3. More Results on Blind Gaussian Denoising

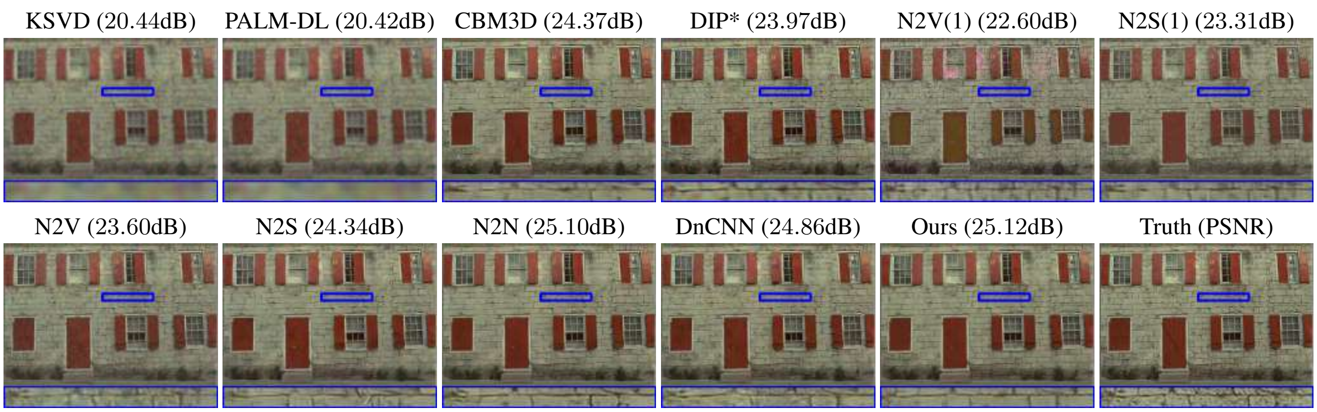

Figure 1: Visual results of blind AWGN denoising on image 'Kodim01' of Set9 with noise level $\sigma = 75$.

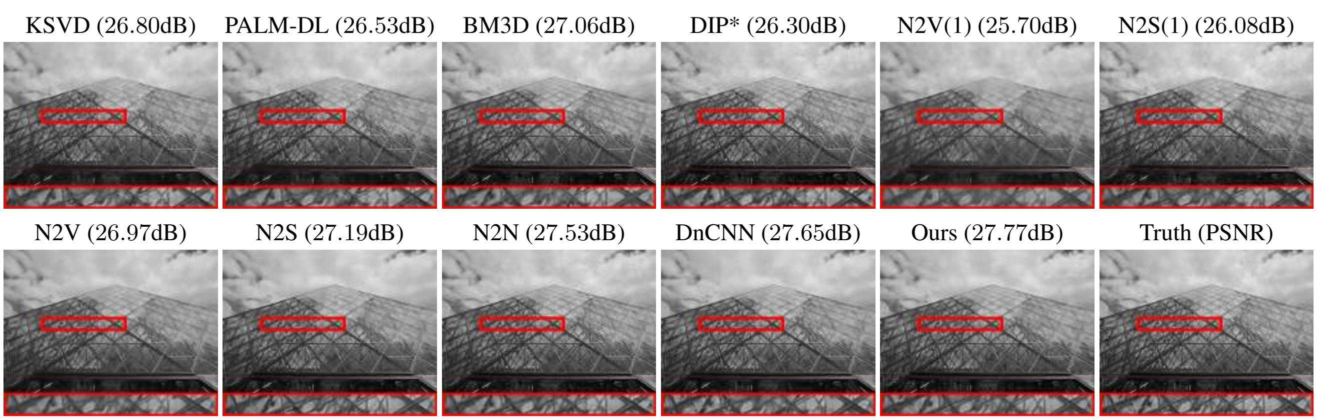

Figure 2: Visual results of blind AWGN denoising on image '223061' on BSD68 with noise level $\sigma = 25$.

4. More Results on Removal of Real-World Image Noise

Due to space limitation, the quantitative results of N2V(1) and N2S(1) are not listed in Table 2 in our main paper. The following are their results. (a) N2V(1): PSNR=34.14dB, SSIM=0.95; (b) N2S(1): PSNR=34.69dB, SSIM=0.97. Also, regarding the visual comparison in Fig. 3 in our main paper, the results of some methods are not presented. For completeness, we show the results of all compared methods in Fig. 3. See also Fig. 4 for one more example on visual comparison.

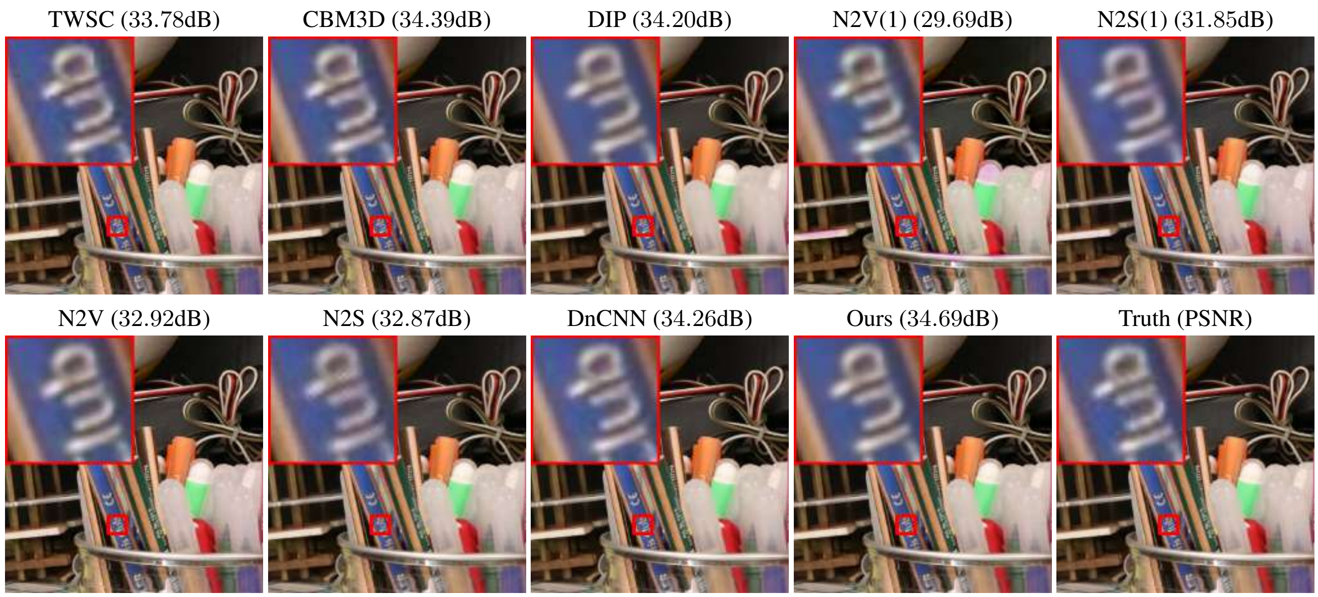

Figure 3: Denoising results on a real-world noisy image by different methods.

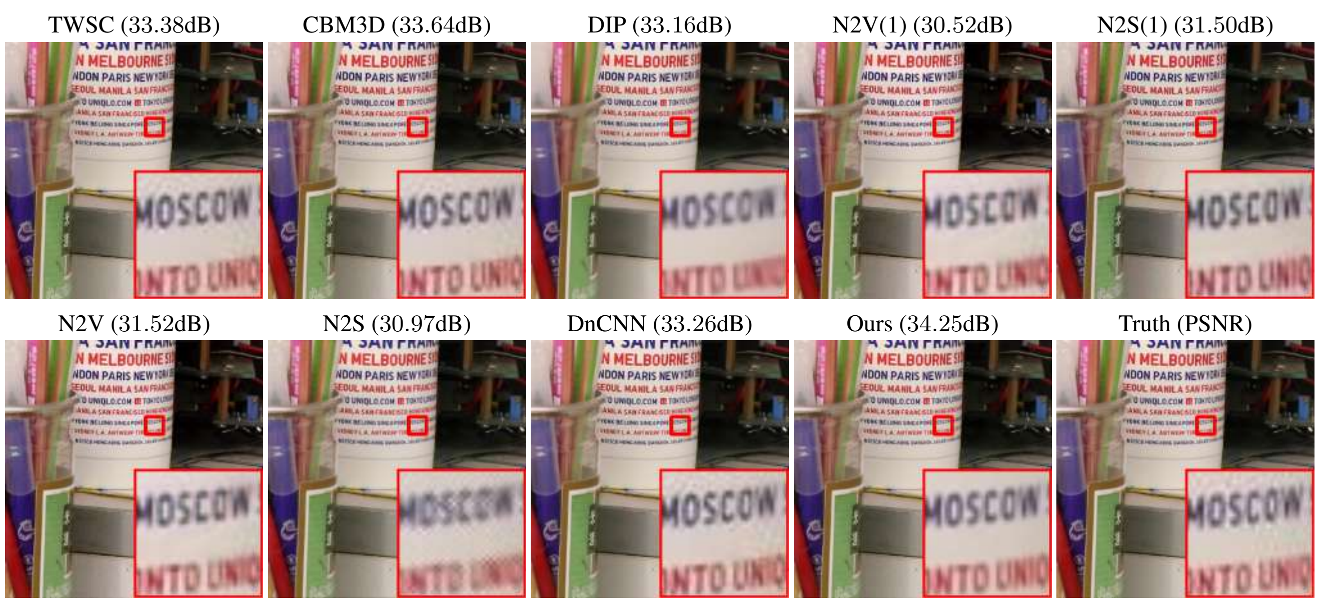

Figure 4: Denoising results on a real-world noisy image by different methods.

References

[1] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proc. ECCV*, pages 85–100, 2018. 1