Supplementary Materials
BlueNeg: A 35mm Negative Film Dataset for Restoring Channel-Heterogeneous Deterioration

Hanyuan Liu[1] Chengze Li[2] Minshan Xie[4] Zhenni Wang[3]
Jiawen Liang[1] Chi-Sing Leung[1,*] Tien-Tsin Wong[5]
[1] City University of Hong Kong [2] Saint Francis University [3] Huawei Technologies Ltd
[4] Centre for Perceptual and Interactive Intelligence [5] Monash University
meethyliu@gmail.com, czli@sfu.edu.hk, msxie92@gmail.com, zhenni126@126.com,
jiawliang6-c@my.cityu.edu.hk, eeleungc@cityu.edu.hk, tt.wong@monash.edu

Contents

A. Dataset Documentation

A.1. Dataset Access

Due to the large dataset size (276 GB), we use HuggingFace Hub as the hosting platform for the public release of our BlueNeg dataset. Our dataset is publicly available at https://huggingface.co/datasets/ttgroup/blueneg-release.

A.2. Dataset License

Tien-Tsin Wong holds the copyright for all released data. The BlueNeg dataset will be released under a customized license that shares the same spirit as the Creative Commons Attribution 4.0 International License (CC BY 4.0):

> All photos are owned and copyrighted by Tien-Tsin Wong. You are automatically granted with permission to use the images for academic and commerical usages, provided that the image credit "Copyrighted by Tien-Tsin Wong" is included in any forms of publication, reproduction, redistribution, or derivatives of the images.

*Corresponding author.

A.3. Maintenance Plan

The current dataset is self-contained. However, we will continue to update the dataset for at least one or two years. The planned maintenance includes updating the dataset with new images, fixing any issues with the dataset, and improving evaluation protocols.

A.4. Dataset Composition

We present the dataset composition in Table 1.

Data Source	Partition	Number	Folder	Bits per channel	Resolution (may varies)	File extension	Public
Film negatives	blue-corrupted	297	negative-16bit	16	10,128 x 6,840	dng	Yes
	blue-intact	194					
Printed photo	blue-corrupted	247	printed-16bit	16	2,994 x 1,920	tif	Yes
	blue-intact	151					
Negative preview (after negation)	blue-corrupted	297	negative-preview-8bit	8	1,322 x 892	preview.png	Yes
	blue-intact	194					
Pseudo ground truth	blue-corrupted	247	pseudogt-8bit	8	<1,322 x 892	pseudogt.png	Yes
Printed photo (testset)	blue-corrupted	30	-	16	2,994 x 1,920	tif	No
Pseudo ground truth (testset)	blue-corrupted	30	-	8	<1,322 x 892	pseudogt.png	No

Table 1. Dataset composition

A.5. Dataset Motivation

This set of images was collected to study the research problem of restoring corrupted film negatives. Due to the physical nature of film negatives, the red, green, and blue light-sensitive layers are located differently inside the film negatives. Therefore, the rates of deterioration of these three layers differ. Tien-Tsin Wong found that the blue channel is relatively vulnerable compared to the other two channels because the blue light-sensitive layer is on the outermost layer on the emulsion side. That is, the deterioration rates are heterogeneous. This characteristic is inherited from its ancestor, i.e., storing the three color negatives (R, G, B) on three separate glass plates by photographer Prokudin-Gorsky (Fig. 2 in the main paper). Since the blue channel is more vulnerable and the other two channels are relatively well-preserved, this means we can restore the blue channel by exploiting the retained information from the red and green channels to restore the color photograph. This is especially feasible with the latest AI technologies. Unfortunately, most existing photo restoration techniques are developed based on printed photographs, in which the nature of deterioration differs from that of negatives. This is why we created this dataset.

A.6. Dataset Usage

We will provide all film negatives (297 blue-corrupted and 194 blue-intact) with their preview versions in the released dataset. Additionally, we will provide the matched printed photos (247 blue-corrupted and 151 blue-intact; note that here "blue-corrupted" indicates the data partition and the blue channels do not contain deterioration in the printed photos) and the estimated pseudo ground truths (247 blue-corrupted). We will not release the printed photos and estimated pseudo ground truths of the selected testing set (30 photos), and we will host an evaluation server on HuggingFace Space and provide an evaluation SDK.

The film negatives are stored in Adobe DNG format. Users can use Adobe Photoshop, GNU Image Manipulation Program, or RawTherapee to view and manipulate the DNG files. For programmatic usage, users are suggested to use RawPy, LibRaw, or ImageMagick to access the DNG files. The printed photos are stored in TIFF format. The negative previews and estimated pseudo ground truths are stored in PNG format. We provide metadata for each photo in the `meta.json` file, which includes the photo taken date, photo taken location, film type, film roll number, and scene properties. In addition, the printed photo alignment information is stored in the `transformations.pkl` file, which is a Python pickle file containing the perspective projection matrix for each printed photo and the corresponding bounding box of the warped printed photo in the negative preview. Users need to crop the results using the bounding box to obtain the matched view with the pseudo ground truth. We will provide a Python evaluation SDK and example programs to facilitate dataset usage after the public release of our dataset.

B. Experiment Setup

Baseline Methods. We obtained the source code and pre-trained models from the official source code repositories of each mentioned photo restoration frameworks:

- DDRM [5]. We used the model checkpoint `256x256_diffusion_uncond.pt` and default hyperparameters except `DEGREDATION` was set to `inp` and we modified the source code to apply the mask on the blue channel of the input image.
- DDNM [11]. We used the model checkpoint `256x256_diffusion_uncond.pt` and default hyperparameters except `sigma_y=0.05` and `DEGREDATION` was set to `inpainting` and we modified the source code to apply the mask on the blue channel of the input image.
- PromptIR [8]. Default settings were used.
- DA-CLIP [6]. Default settings were used. The model checkpoints are `DA-CLIP-mix` and `Universal-IR-mix`. The results are resized back the the original input resolution.
- DGP [7]. We used the default `experiments/examples/run_inpainting_list.sh` and modified the source code to apply the mask on the blue channel of the input image.
- BringOld [10]. Default settings and full pipeline were used.
- NAFNet [2]. We used `NAFNet-width64.yml` with modified source code to accept 5-channel inputs. For supervised baseline, we use the `NAFNet-width64.yml` and original source codes.
- Restormer [12]. We used `Deraining_Restormer.yml` with the original source code.

Finetuning Setup As mentioned in our main paper, in order to ensure fair comparisons with supervised learning-based restoration baselines on the blue-corrupted testing split, we have finetuned the generic photo restoration baselines on our BlueNeg dataset (blue-corrupted training split). It's important to note that different photo restoration methods are based on distinct formulations and different framework designs, leading to variations in the finetuning configurations. We provide detailed configurations below:

- DDRM/DDNM: These methods are zero-shot image restoration techniques that utilize OpenAI's guided diffusion model [4] (trained on ImageNet [3]) as the generative prior. We fine-tuned the diffusion model using our blue-intact data, following standard practices like DreamBooth [9] that involve a balanced combination of original and new data.
- DGP: DGP employs BigGAN [1] (trained on ImageNet) as the generative prior for photo restoration. We applied the same fine-tuning configuration as used for DDRM/DDNM.
- BringOld: We fine-tuned the model using our data while maintaining the use of synthetic data as in the original implementation.
- PromptIR/DACLIP: Both methods are designed for generic photo restoration and support multiple degradation types. We added our dataset as an extra degradation type into the their original training data and fine-tuned the models with the expanded training sets.

Runtime Environments. The model training were conducted on 8 NVIDIA A40 GPUs with Intel Xeon(R) Gold 6248R CPU, and 1TB RAM. The evaluations were conducted on 1 NVIDIA A100 GPU with AMD EPYC 7532 CPU, and 1TB RAM.

C. Additional Visual Comparisons

In the visual comparison, to isolate the impact of blue-channel deterioration, we also replaced the R/G channels of restored images with those from the original input prior to evaluation. This ensures that artifacts newly introduced by the restoration models in the R/G channels do not confound the results, while also providing a clearer visualization of blue-channel deterioration. In this setting, elevated B-channel values typically produce a bluish cast, while reduced values yield a yellowish tint. Thus, regional deterioration can be inferred from B-channel deviations and color shifts.

C.1. Generic Photo Restoration Frameworks

Here P-GT denotes the pseudo ground truth. We show the blue channel of each photo under its RGB version.

C.2. Supervised Learning-based Baselines

We show the blue channel of each photo under its RGB version.

D. Acknowledgements

The work is jointly supported by a research grant from City University of Hong Kong (No. 9678295) and a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. UGC/FDS11/E02/23).

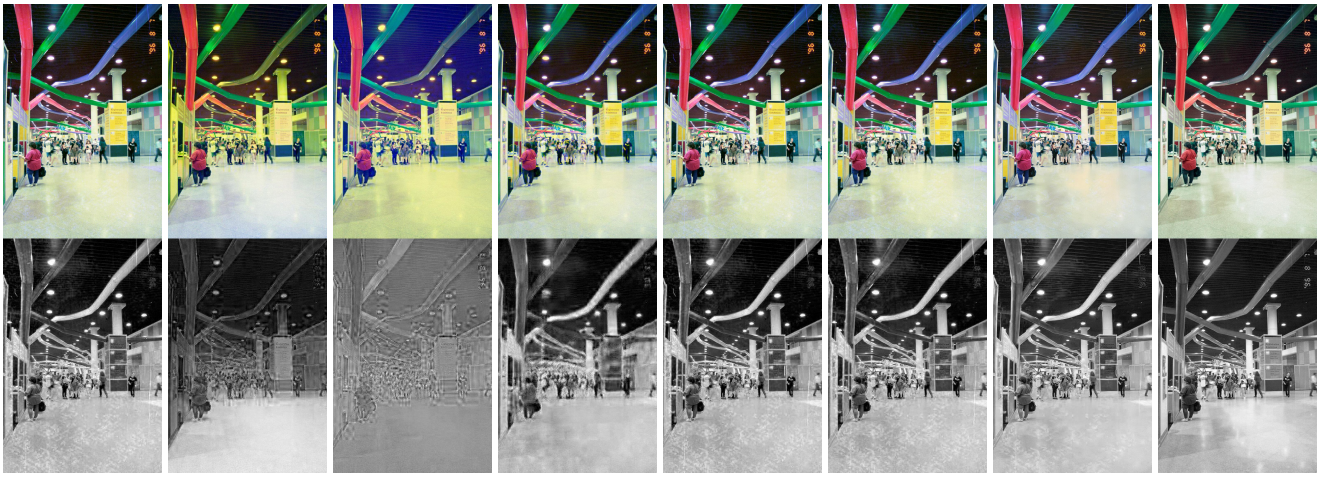

Input DDNM DDRM DGP PromptIR DA-CLIP BringOld P-GT

Figure 1. Visual results of generic photo restoration frameworks. ©Tien-Tsin Wong.

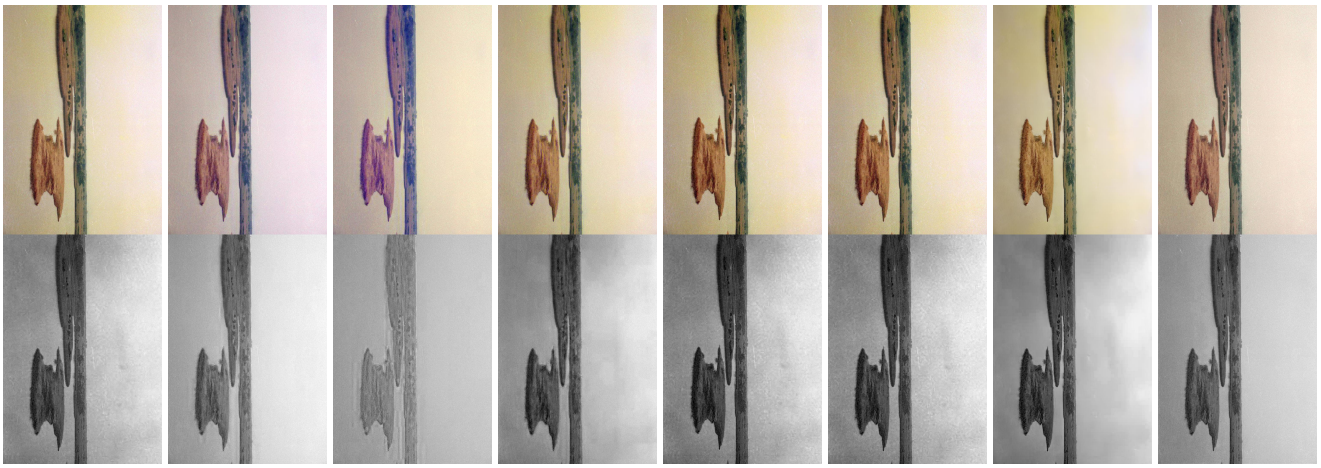

Input DDNM DDRM DGP PromptIR DA-CLIP BringOld P-GT

Figure 2. Visual results of generic photo restoration frameworks. ©Tien-Tsin Wong.

References

[1] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis, 2019. 3

[2] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *European conference on computer vision*, pages 17–33. Springer, 2022. 3

[3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 3

[4] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis, 2021. 3

[5] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *Advances in Neural Information Processing Systems*, 2022. 3

[6] Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling vision-language models for universal image restoration. *arXiv preprint arXiv:2310.01018*, 2023. 3

[7] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2021. 3

[8] Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Khan. Promptir: Prompting for all-in-one image restoration. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 3

[9] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation, 2023. 3

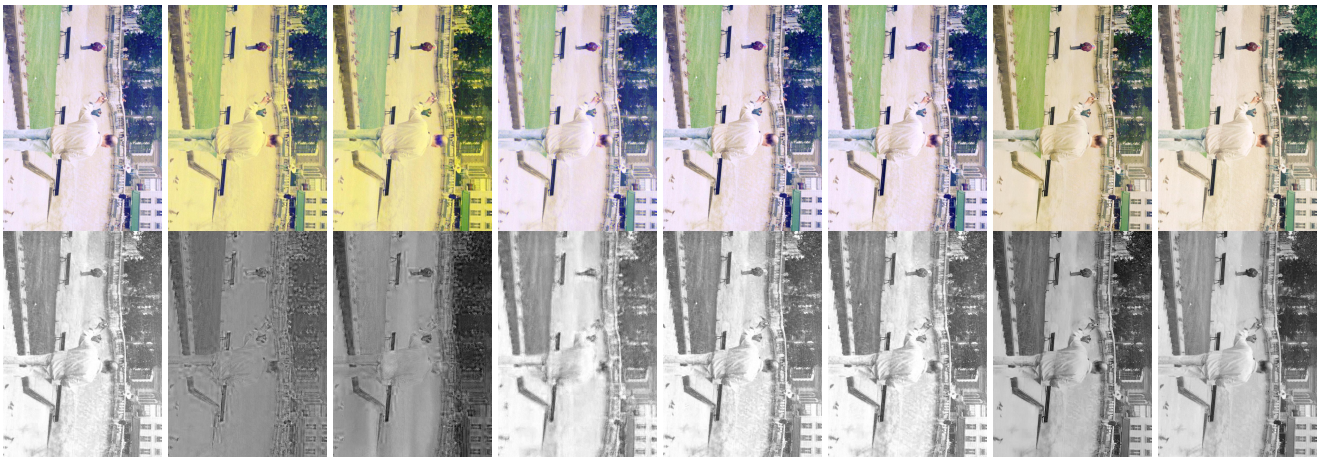

Figure 3. Visual results of generic photo restoration frameworks. ©Tien-Tsin Wong.

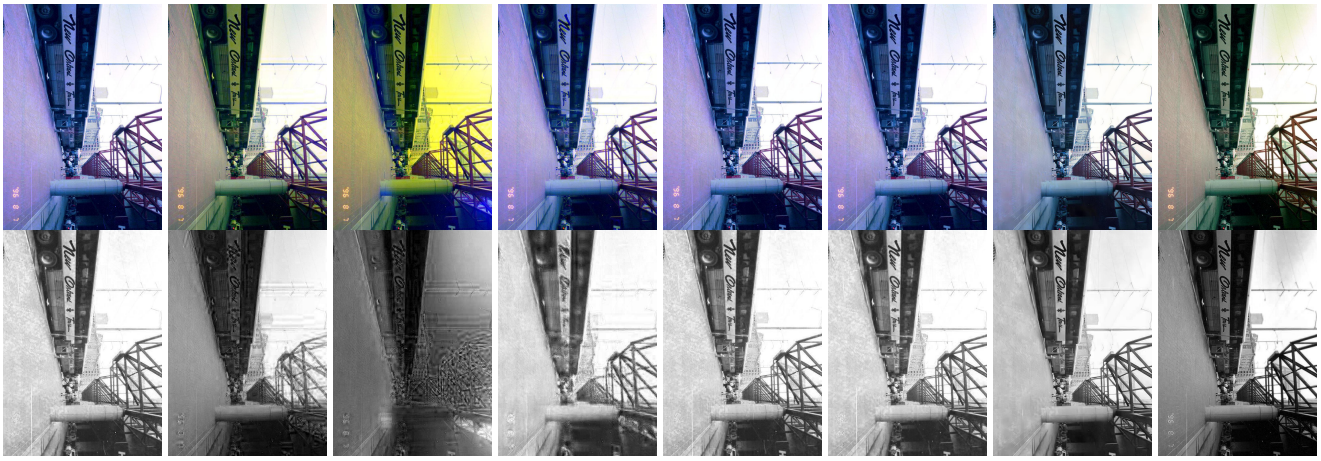

Figure 4. Visual results of generic photo restoration frameworks. ©Tien-Tsin Wong.

[10] Ziyu Wan, Bo Zhang, Dongdong Chen, Pan Zhang, Dong Chen, Jing Liao, and Fang Wen. Bringing old photos back to life. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2747–2757, 2020. 3

[11] Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. *The Eleventh International Conference on Learning Representations*, 2023. 3

[12] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5728–5739, 2022. 3

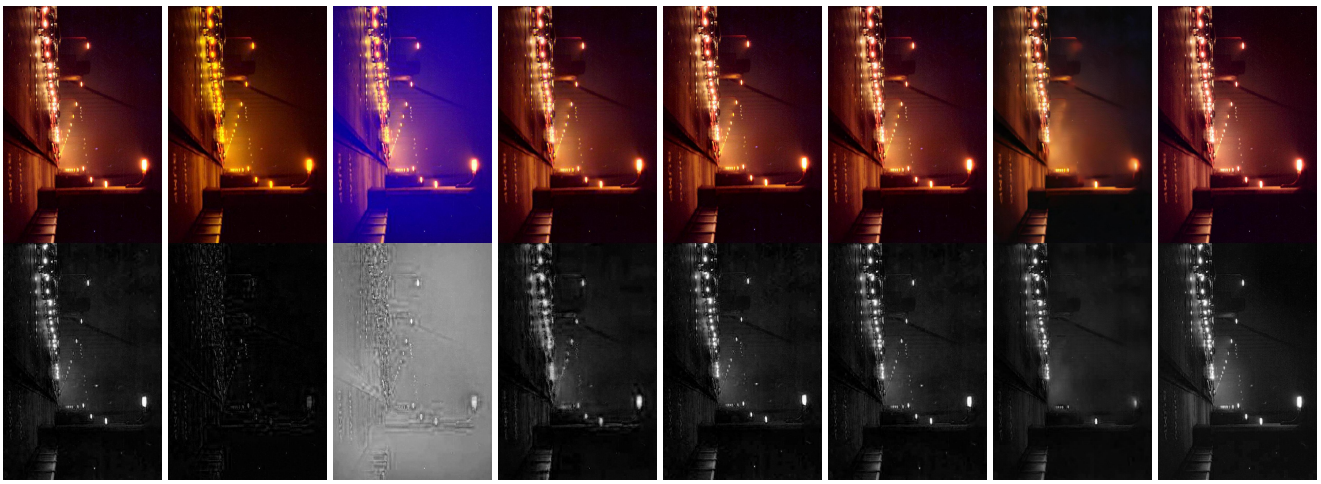

Input DDNM DDRM DGP PromptIR DA-CLIP BringOld P-GT

Figure 5. Visual results of generic photo restoration frameworks. ©Tien-Tsin Wong.

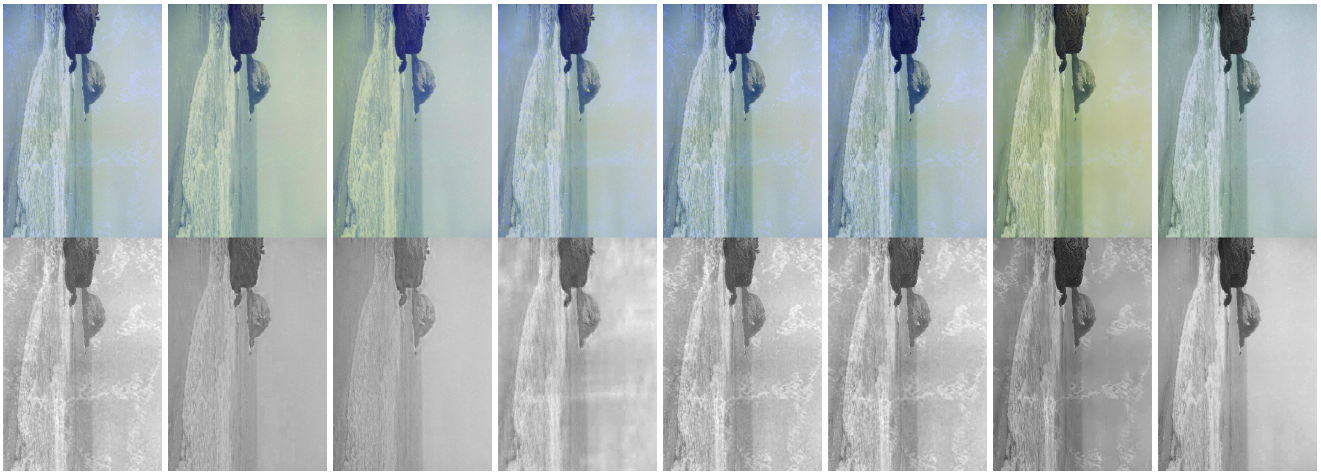

Input DDNM DDRM DGP PromptIR DA-CLIP BringOld P-GT

Figure 6. Visual results of generic photo restoration frameworks. ©Tien-Tsin Wong.

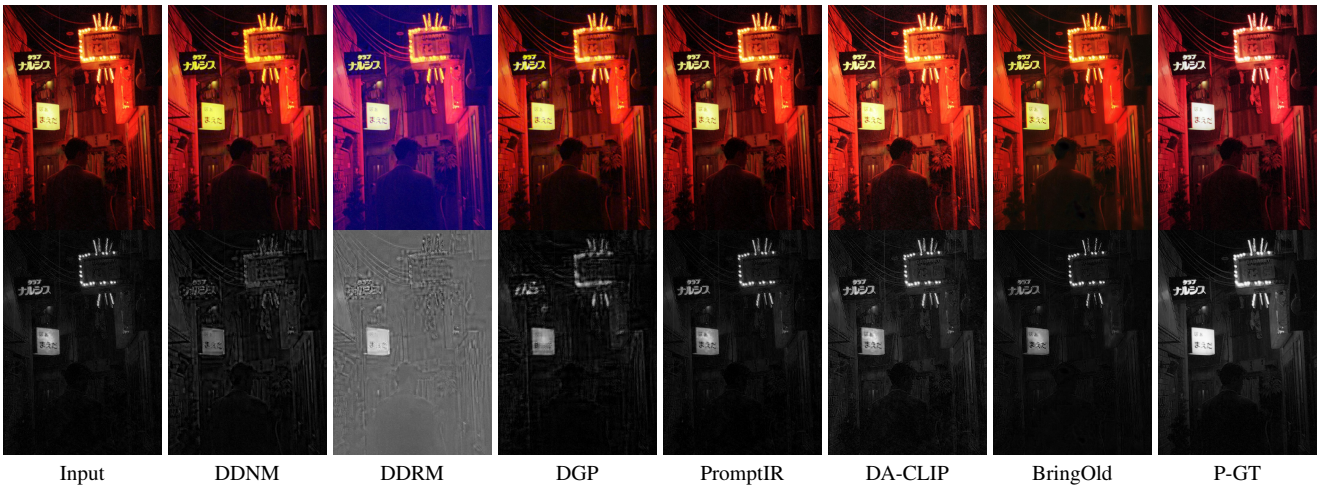

Input DDNM DDRM DGP PromptIR DA-CLIP BringOld P-GT

Figure 7. Visual results of generic photo restoration frameworks. ©Tien-Tsin Wong.

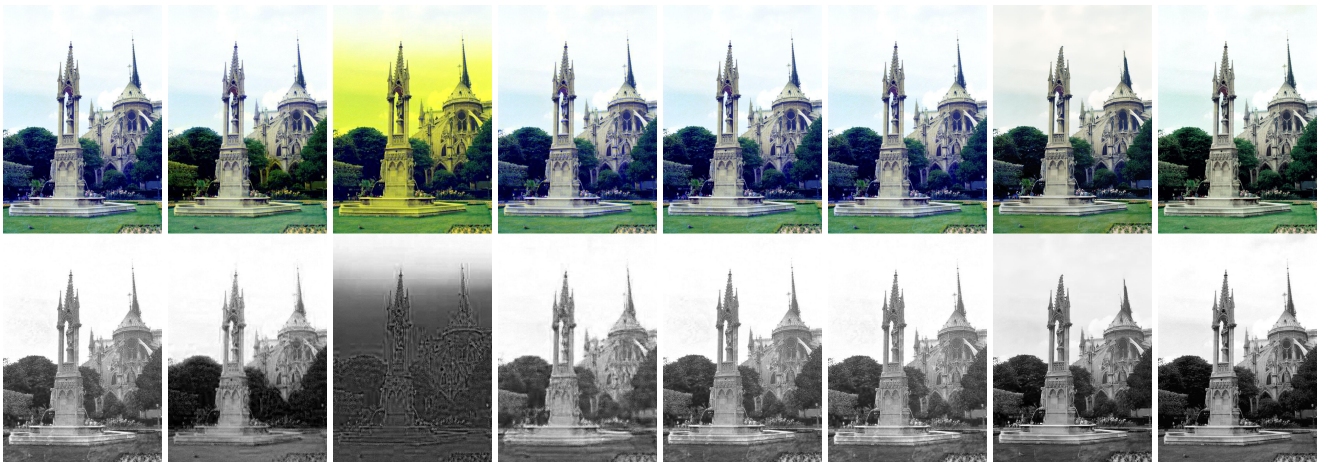

Input DDNM DDRM DGP PromptIR DA-CLIP BringOld P-GT

Figure 8. Visual results of generic photo restoration frameworks. ©Tien-Tsin Wong.

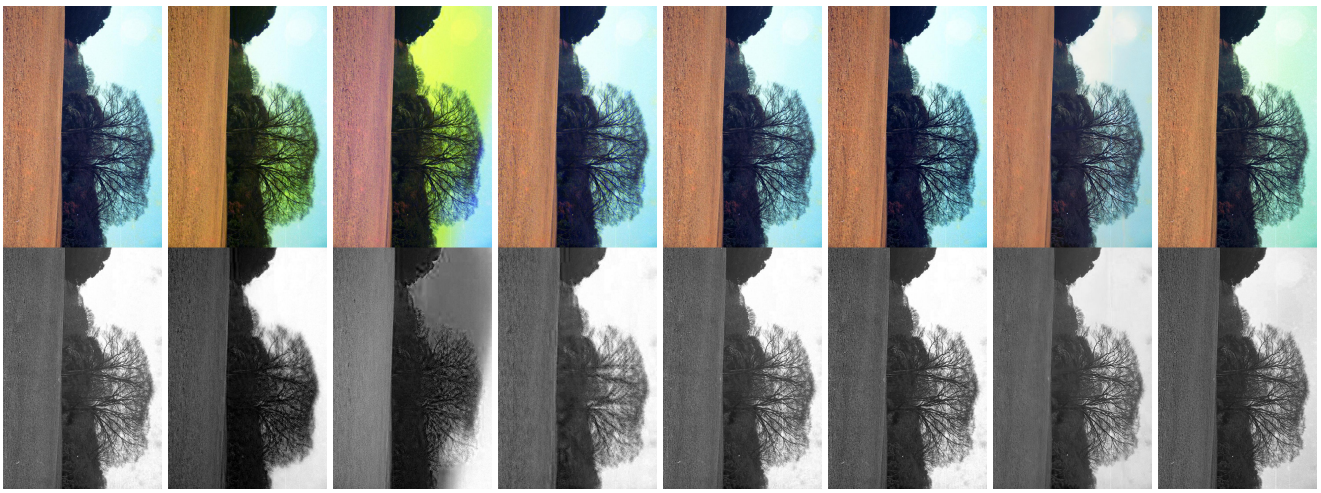

Input DDNM DDRM DGP PromptIR DA-CLIP BringOld P-GT

Figure 9. Visual results of generic photo restoration frameworks. ©Tien-Tsin Wong.

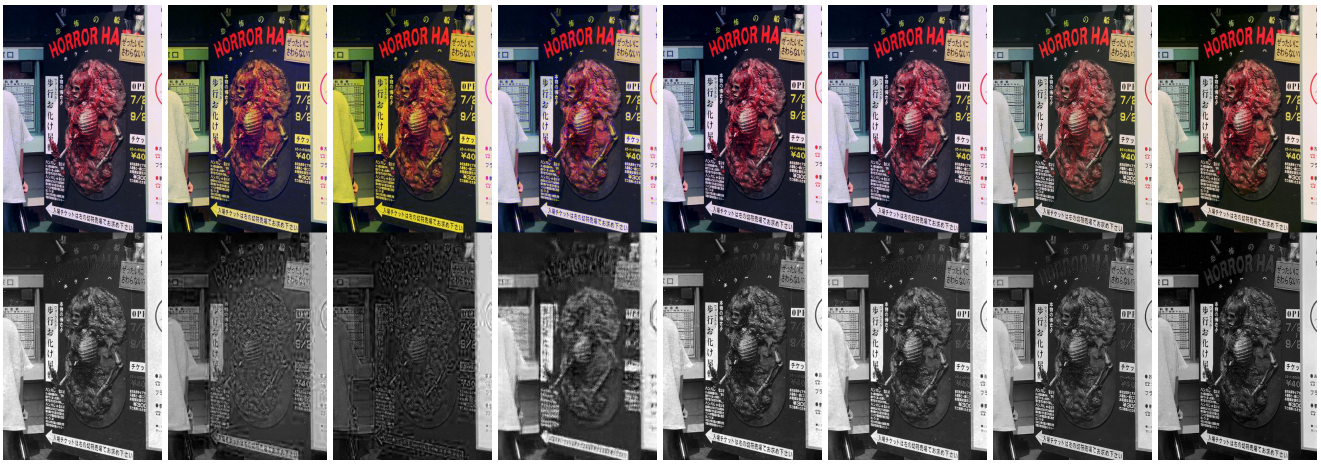

Input DDNM DDRM DGP PromptIR DA-CLIP BringOld P-GT

Figure 10. Visual results of generic photo restoration frameworks. ©Tien-Tsin Wong.

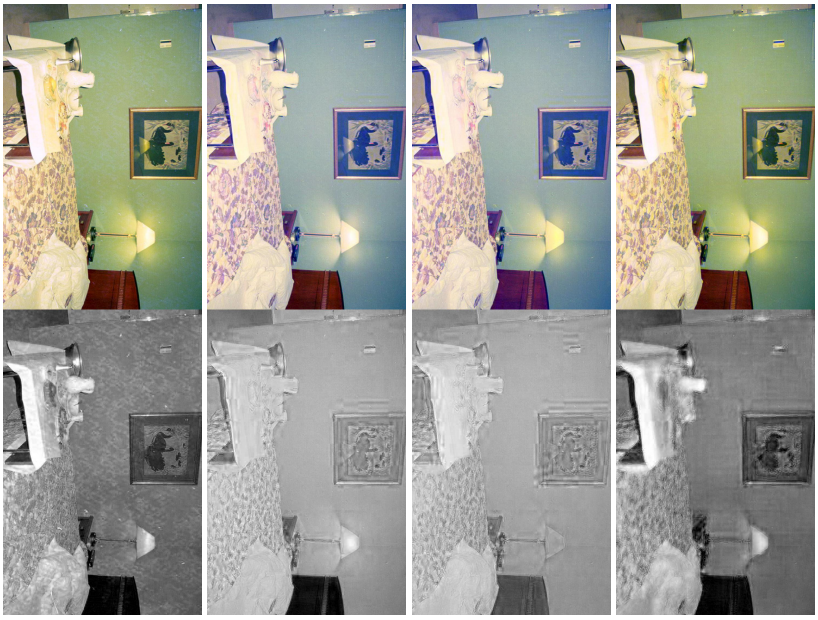

(a) Input (b) DDNM (c) DDRM (d) DGP

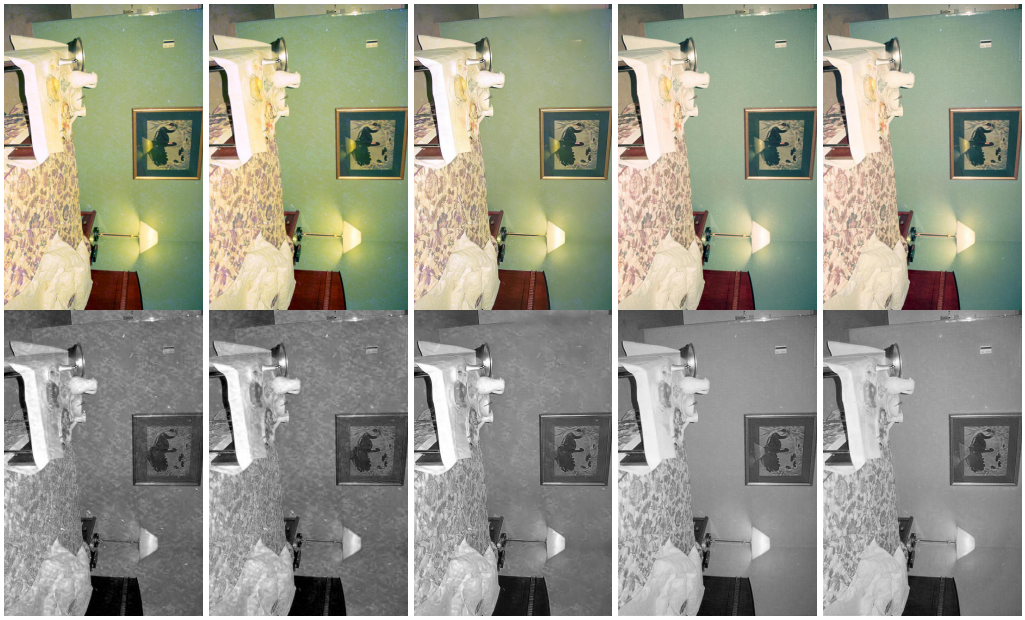

(e) PromptIR (f) DA-CLIP (g) BringOld (h) NAFNet (i) Restormer

Figure 11. Visual results of supervised learning-based baselines (NAFNet and Restormer) with generic photo restoration baselines on testing set. ©Tien-Tsin Wong.

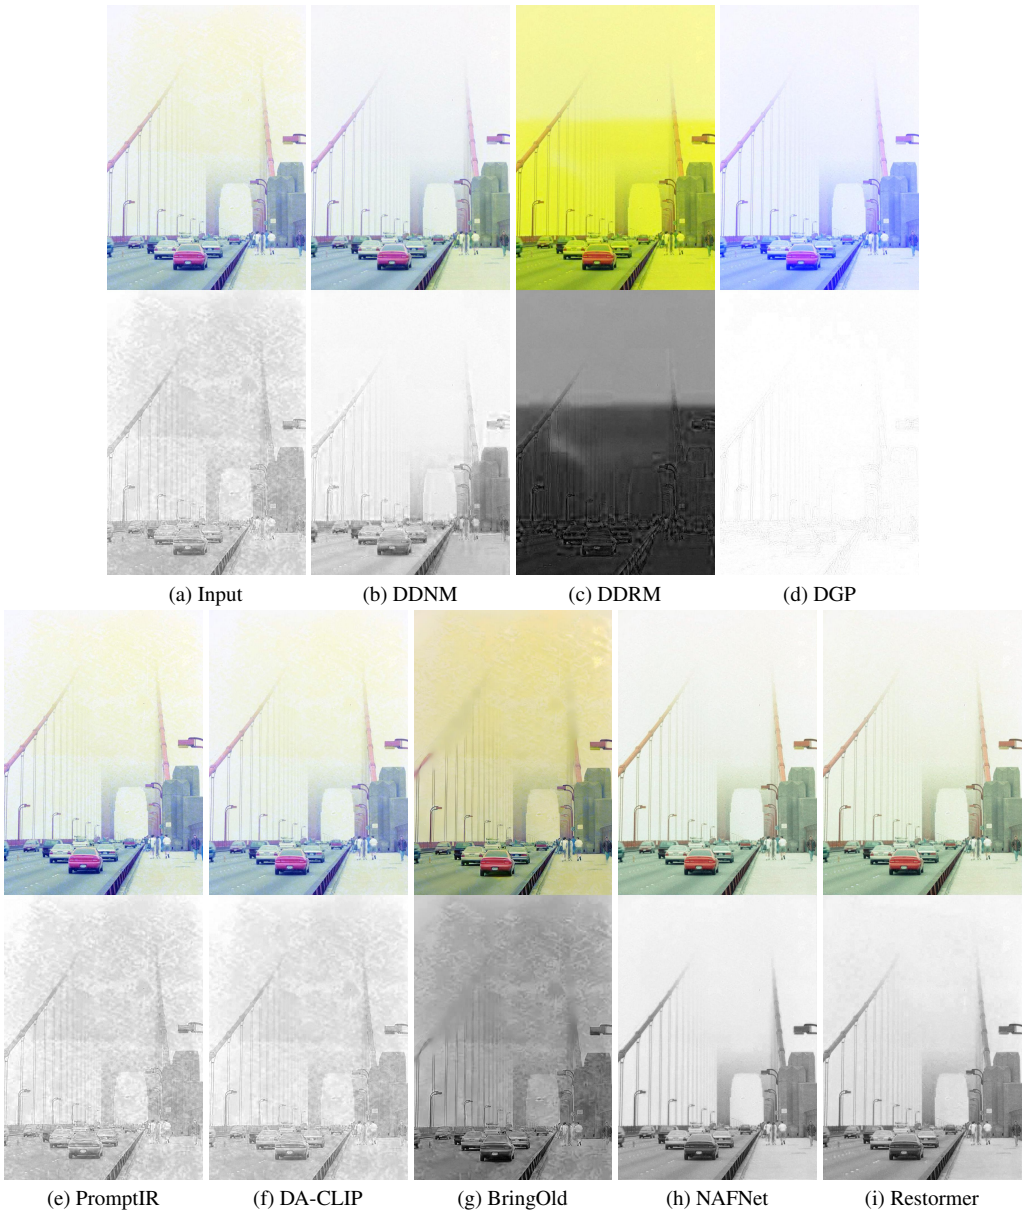

(a) Input (b) DDNM (c) DDRM (d) DGP

(e) PromptIR (f) DA-CLIP (g) BringOld (h) NAFNet (i) Restormer

Figure 12. Visual results of supervised learning-based baselines (NAFNet and Restormer) with generic photo restoration baselines on testing set. ©Tien-Tsin Wong.

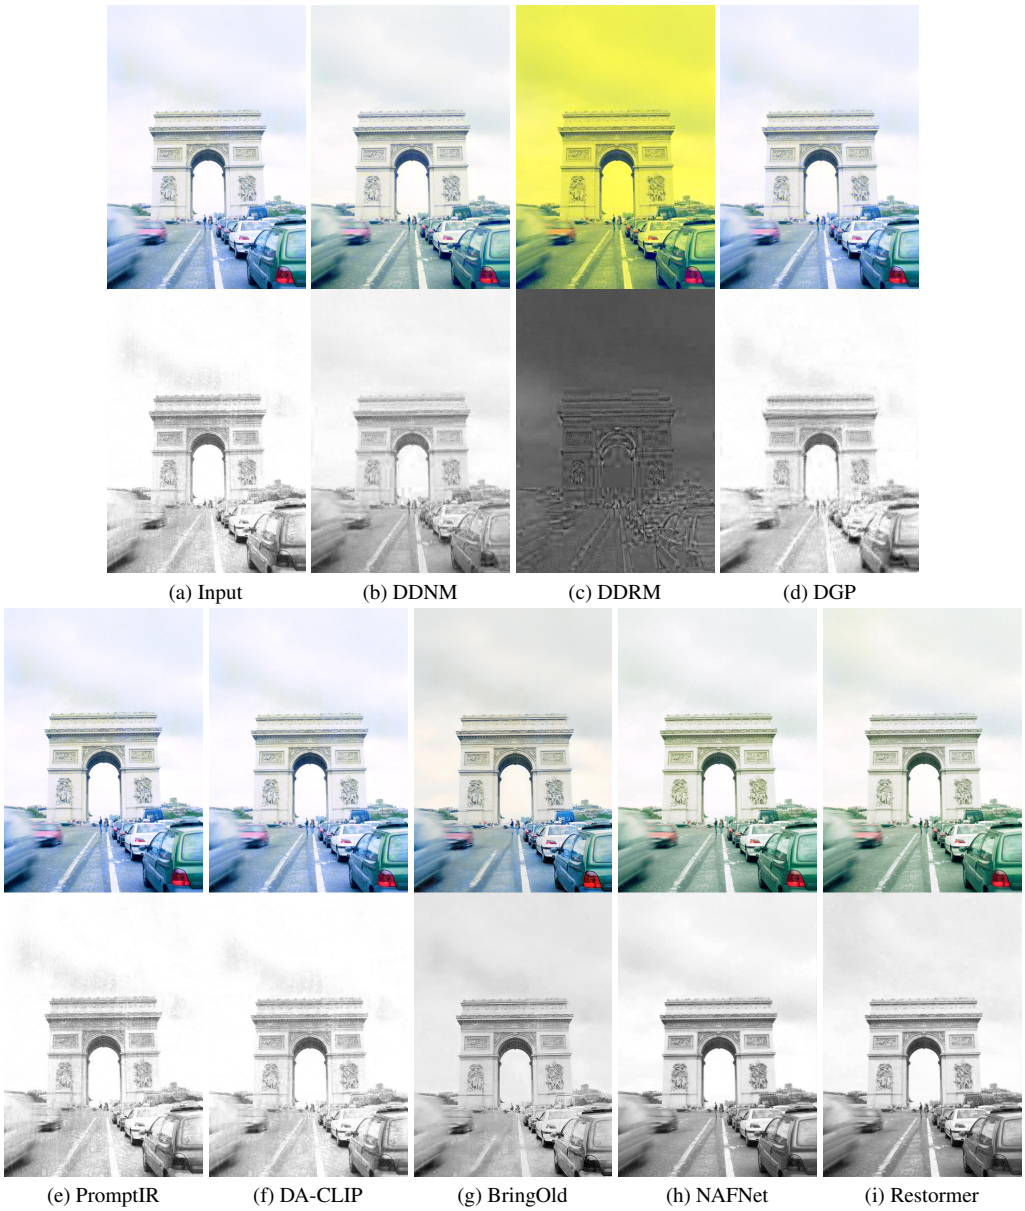

(a) Input (b) DDNM (c) DDRM (d) DGP

(e) PromptIR (f) DA-CLIP (g) BringOld (h) NAFNet (i) Restormer

Figure 13. Visual results of supervised learning-based baselines (NAFNet and Restormer) with generic photo restoration baselines on testing set. ©Tien-Tsin Wong.

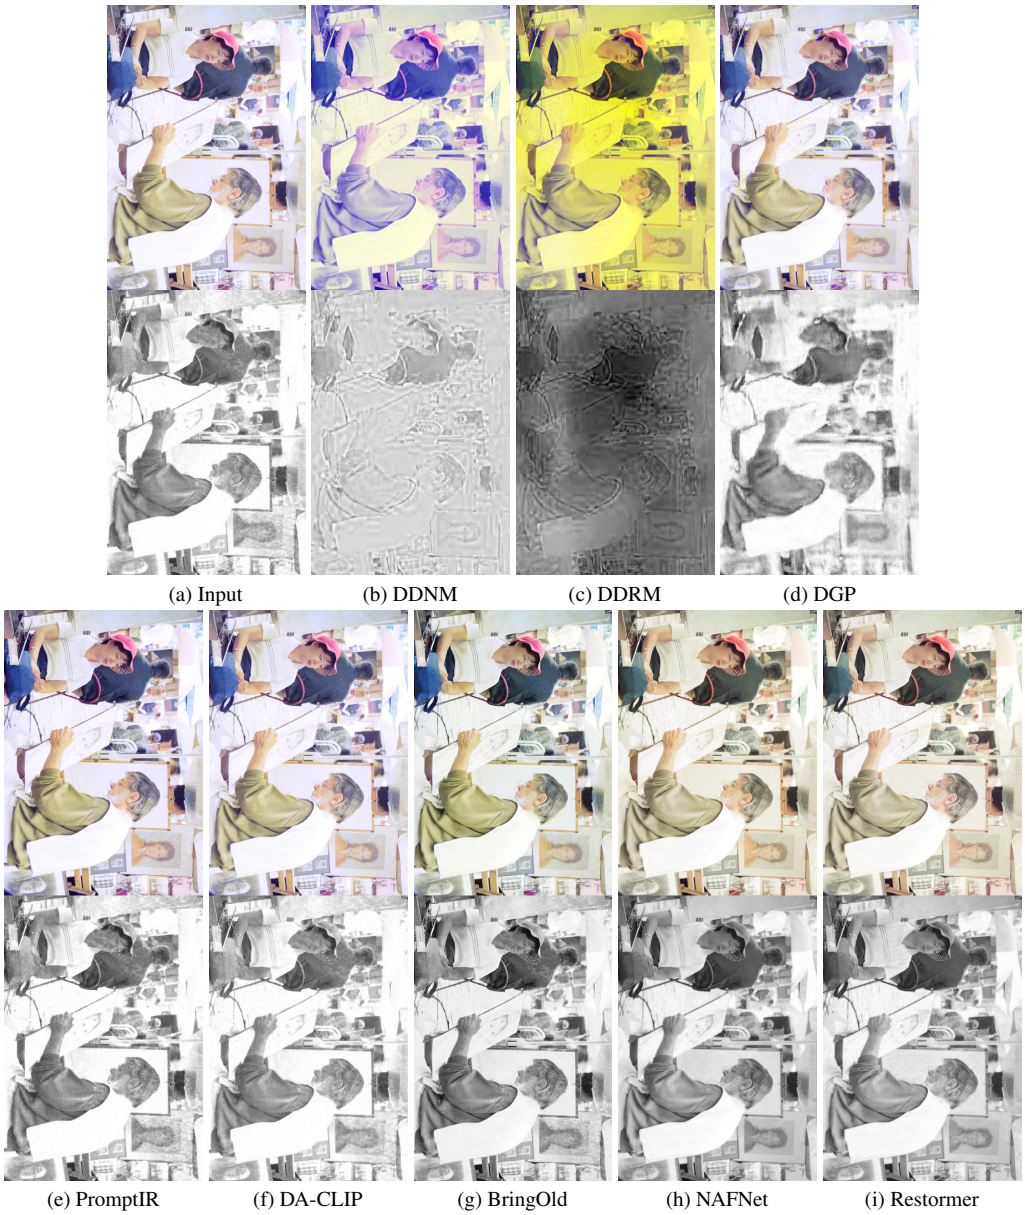

(a) Input (b) DDNM (c) DDRM (d) DGP

(e) PromptIR (f) DA-CLIP (g) BringOld (h) NAFNet (i) Restormer

Figure 14. Visual results of supervised learning-based baselines (NAFNet and Restormer) with generic photo restoration baselines on testing set. ©Tien-Tsin Wong.

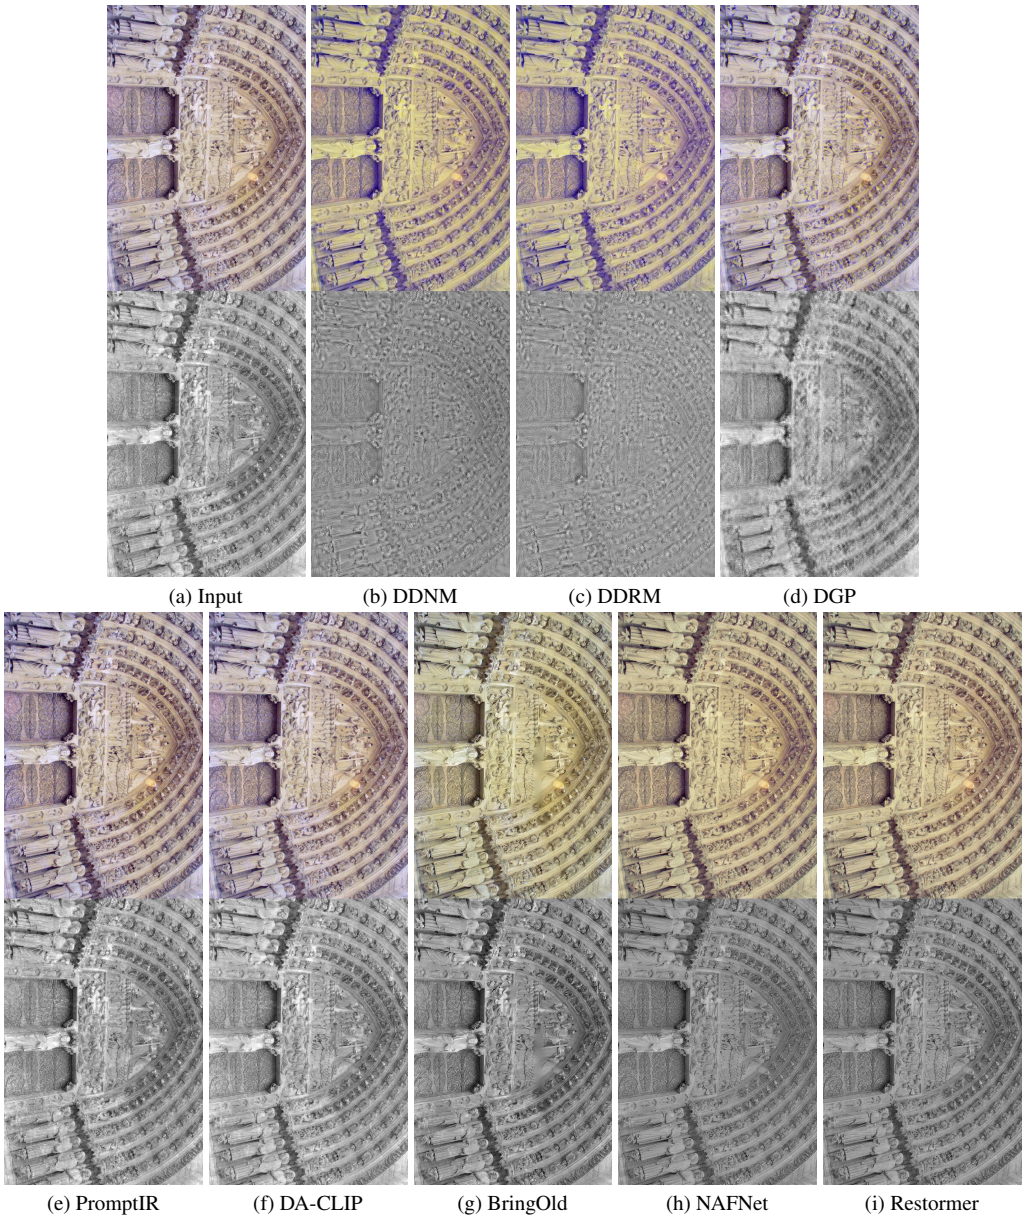

(a) Input (b) DDNM (c) DDRM (d) DGP

(e) PromptIR (f) DA-CLIP (g) BringOld (h) NAFNet (i) Restormer

Figure 15. Visual results of supervised learning-based baselines (NAFNet and Restormer) with generic photo restoration baselines on testing set. ©Tien-Tsin Wong.

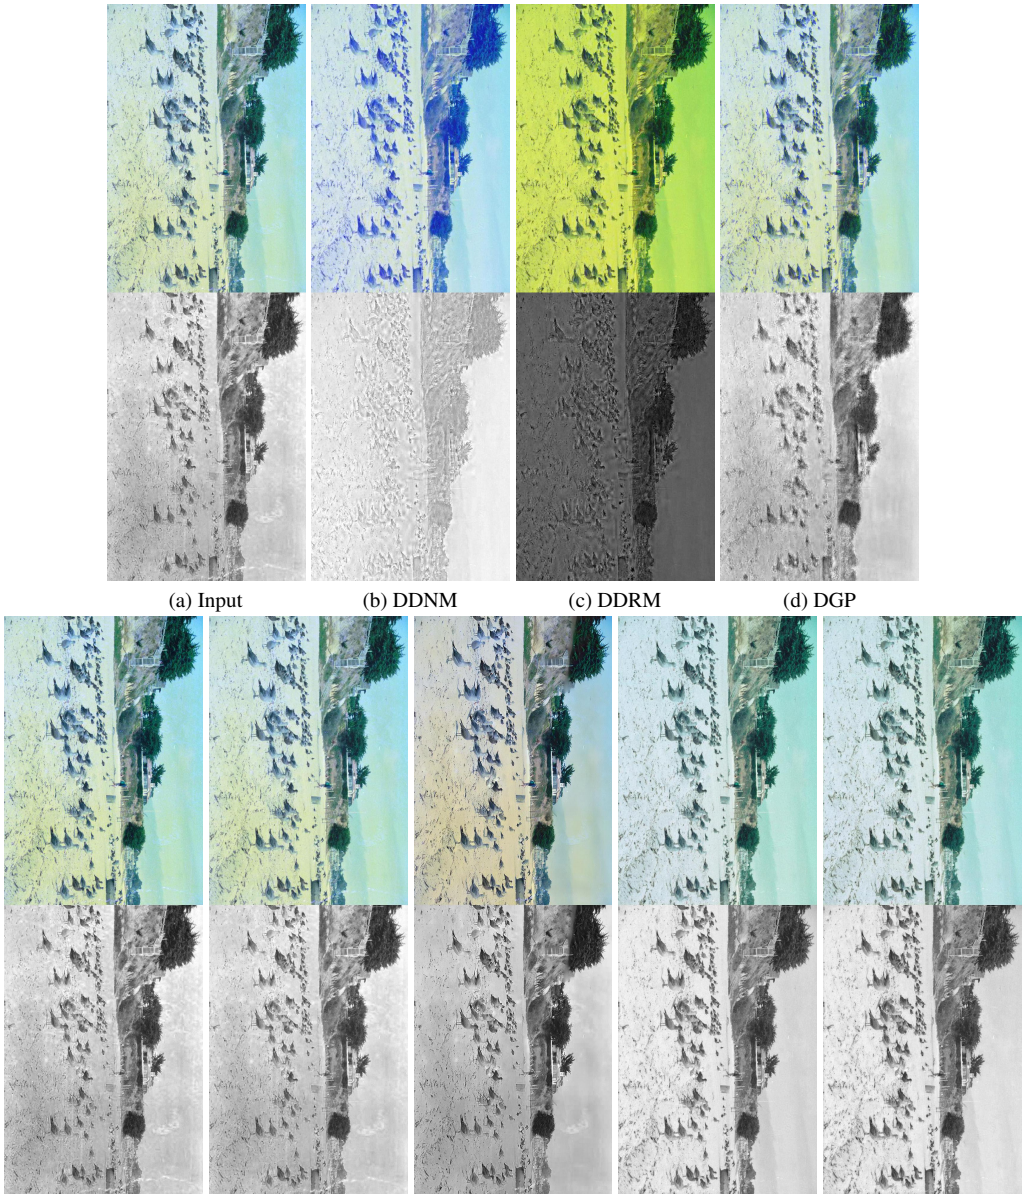

(a) Input (b) DDNM (c) DDRM (d) DGP

(e) PromptIR (f) DA-CLIP (g) BringOld (h) NAFNet (i) Restormer

Figure 16. Visual results of supervised learning-based baselines (NAFNet and Restormer) with generic photo restoration baselines on testing set. ©Tien-Tsin Wong.

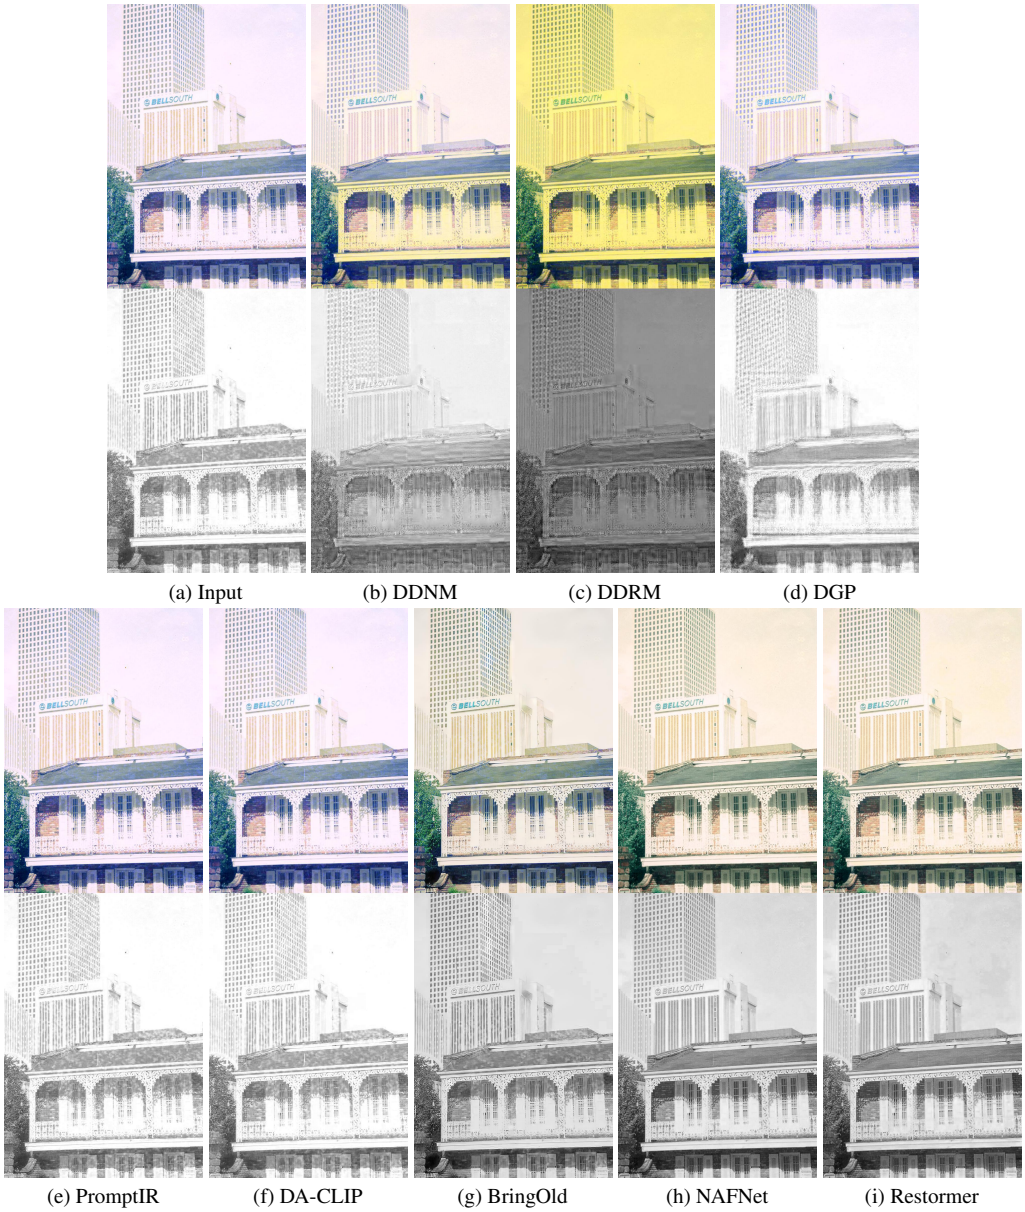

(a) Input (b) DDNM (c) DDRM (d) DGP

(e) PromptIR (f) DA-CLIP (g) BringOld (h) NAFNet (i) Restormer

Figure 17. Visual results of supervised learning-based baselines (NAFNet and Restormer) with generic photo restoration baselines on testing set. ©Tien-Tsin Wong.

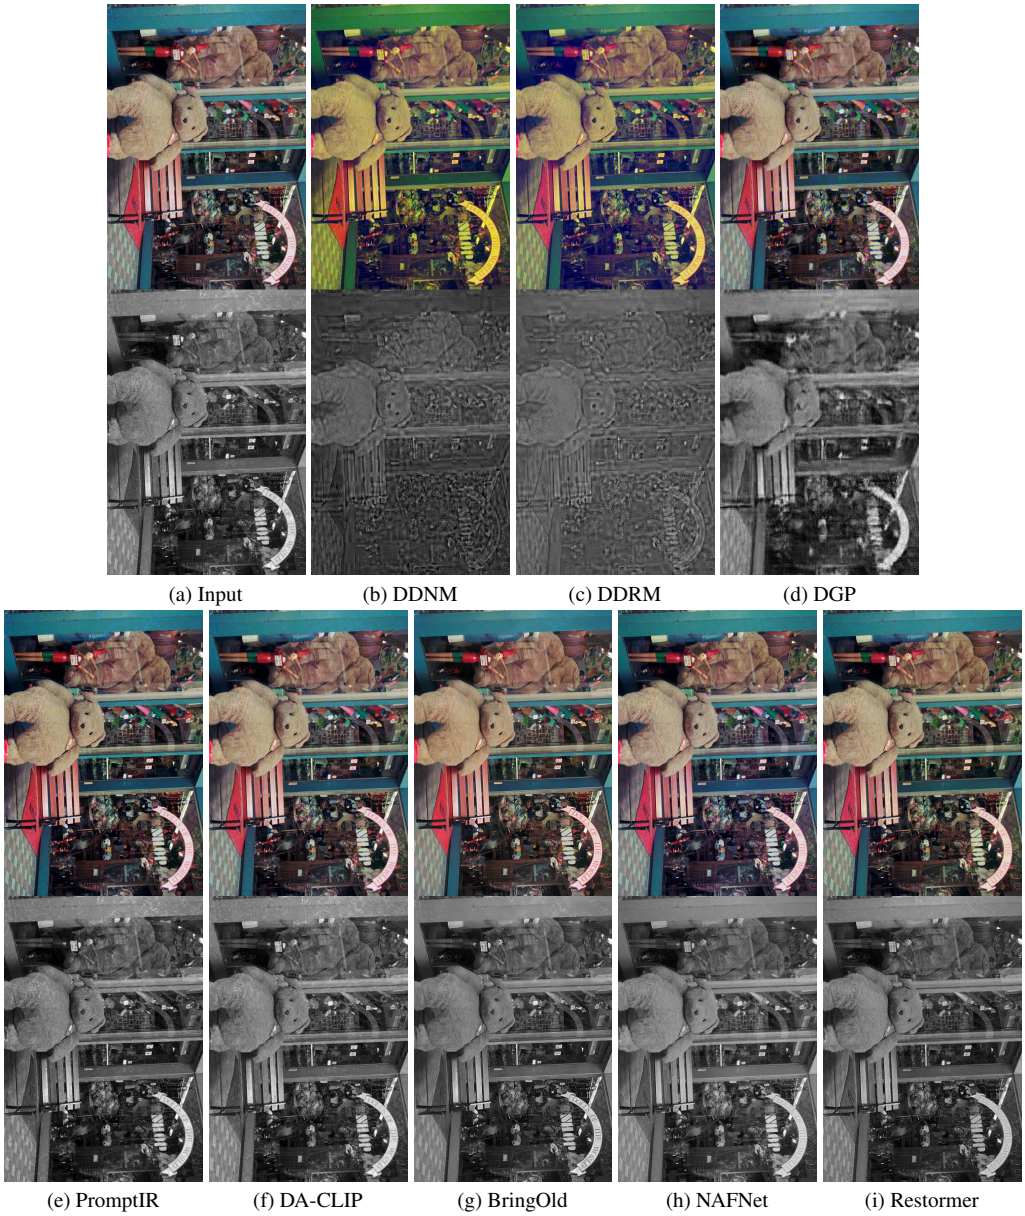

(a) Input (b) DDNM (c) DDRM (d) DGP

(e) PromptIR (f) DA-CLIP (g) BringOld (h) NAFNet (i) Restormer

Figure 18. Visual results of supervised learning-based baselines (NAFNet and Restormer) with generic photo restoration baselines on testing set. ©Tien-Tsin Wong.

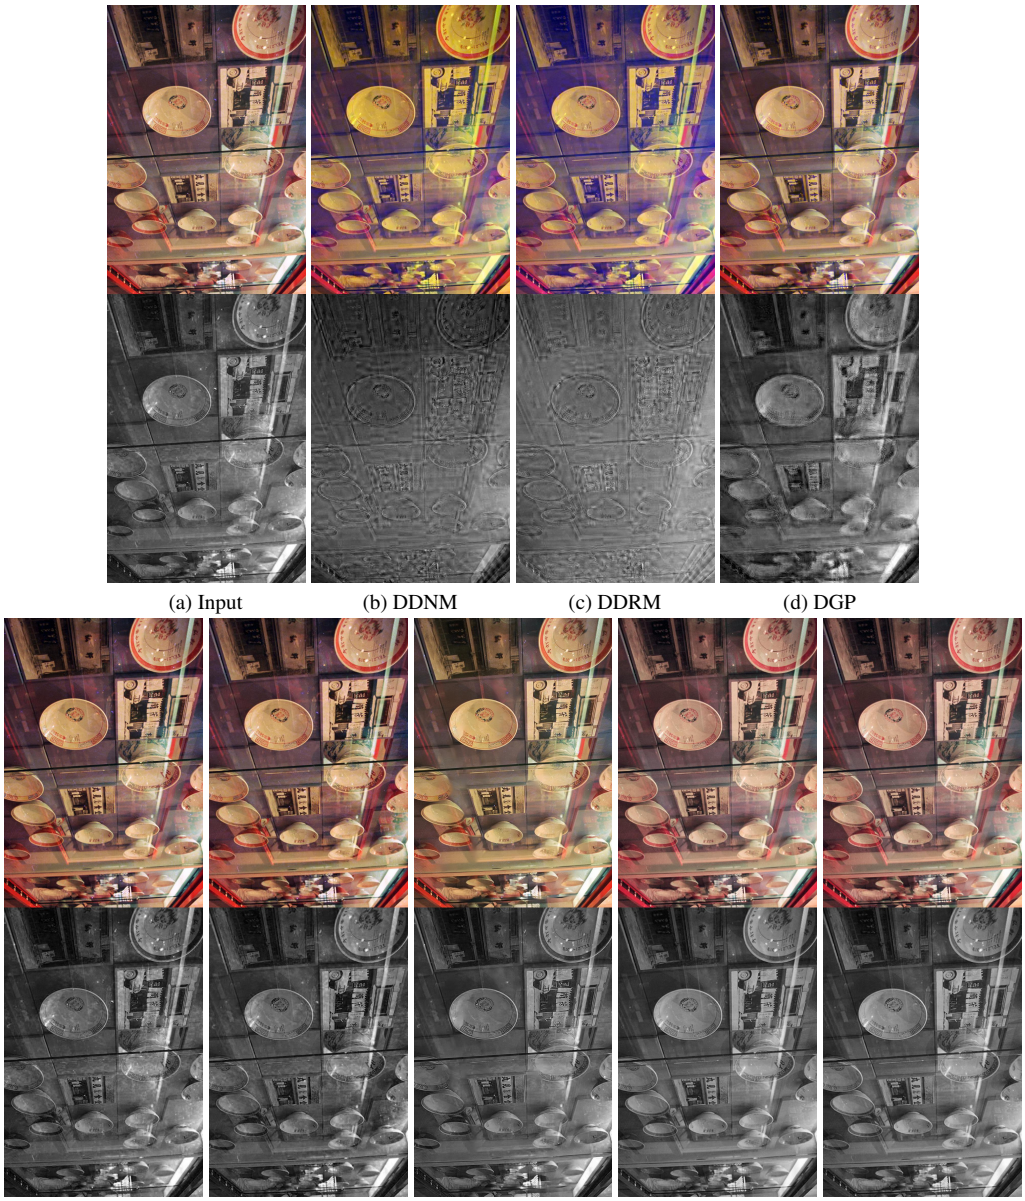

(a) Input (b) DDNM (c) DDRM (d) DGP

(e) PromptIR (f) DA-CLIP (g) BringOld (h) NAFNet (i) Restormer

Figure 19. Visual results of supervised learning-based baselines (NAFNet and Restormer) with generic photo restoration baselines on testing set. ©Tien-Tsin Wong.

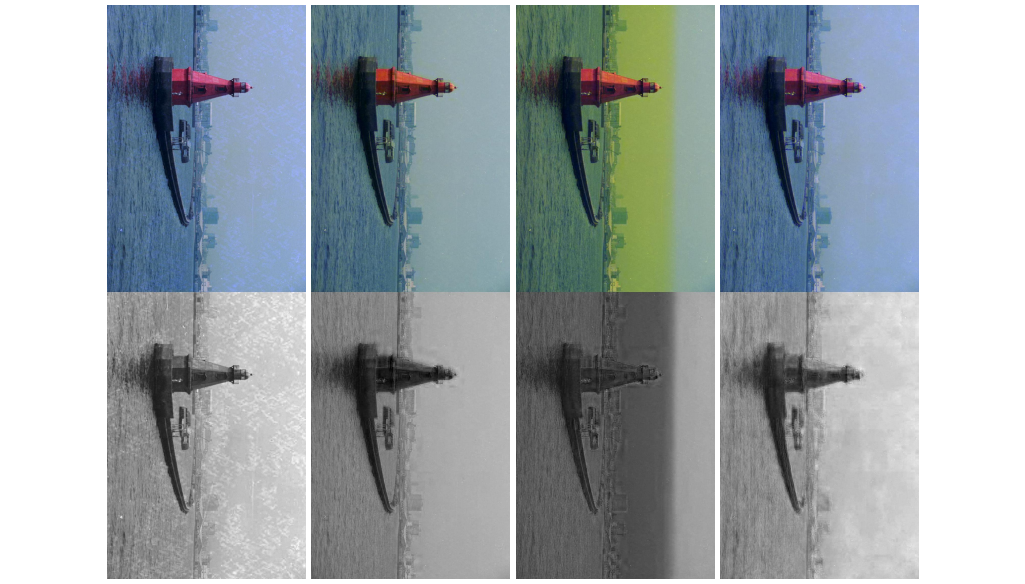

(a) Input (b) DDNM (c) DDRM (d) DGP

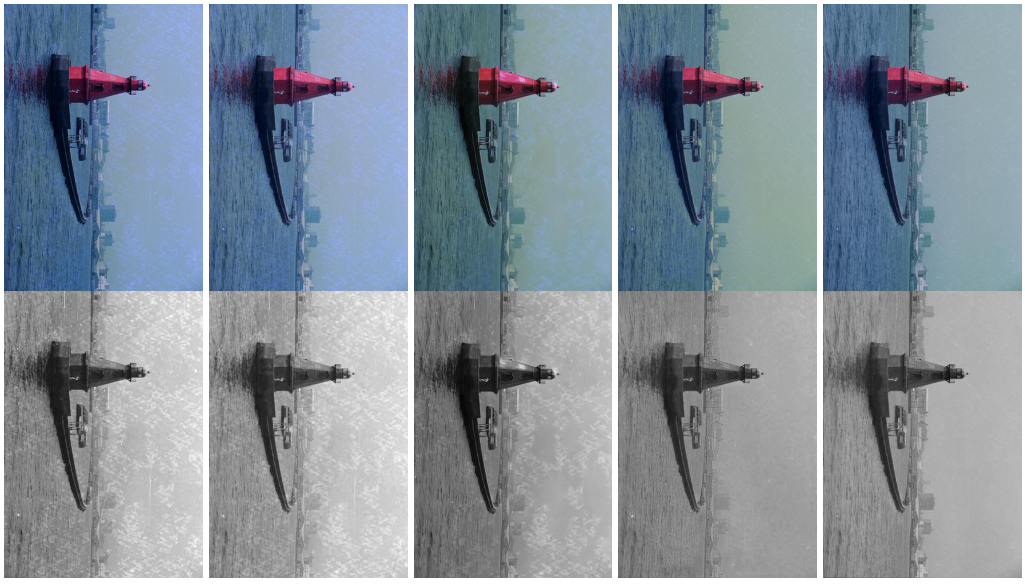

(e) PromptIR (f) DA-CLIP (g) BringOld (h) NAFNet (i) Restormer

Figure 20. Visual results of supervised learning-based baselines (NAFNet and Restormer) with generic photo restoration baselines on testing set. ©Tien-Tsin Wong.