Deep Random Projector: Accelerated Deep Image Prior (THU-AM-162)

Taihui Li[1], Hengkang Wang[1], Zhong Zhuang[2], Ju Sun[1]

[1]Computer Science and Engineering, University of Minnesota
[2]Electrical and Computer Engineering, University of Minnesota

University of Minnesota
Driven to Discover®

JUNE 18-22, 2023
CVPR
VANCOUVER, CANADA

Deep Image Prior

Deep image prior (DIP):

$$\mathbf{x} \approx G_\theta(\mathbf{z})$$

Objective function:

$$\min_\theta \ \ell(\mathbf{y},\ f \circ G_\theta(\mathbf{z})) + \lambda R \circ G_\theta(\mathbf{z})$$

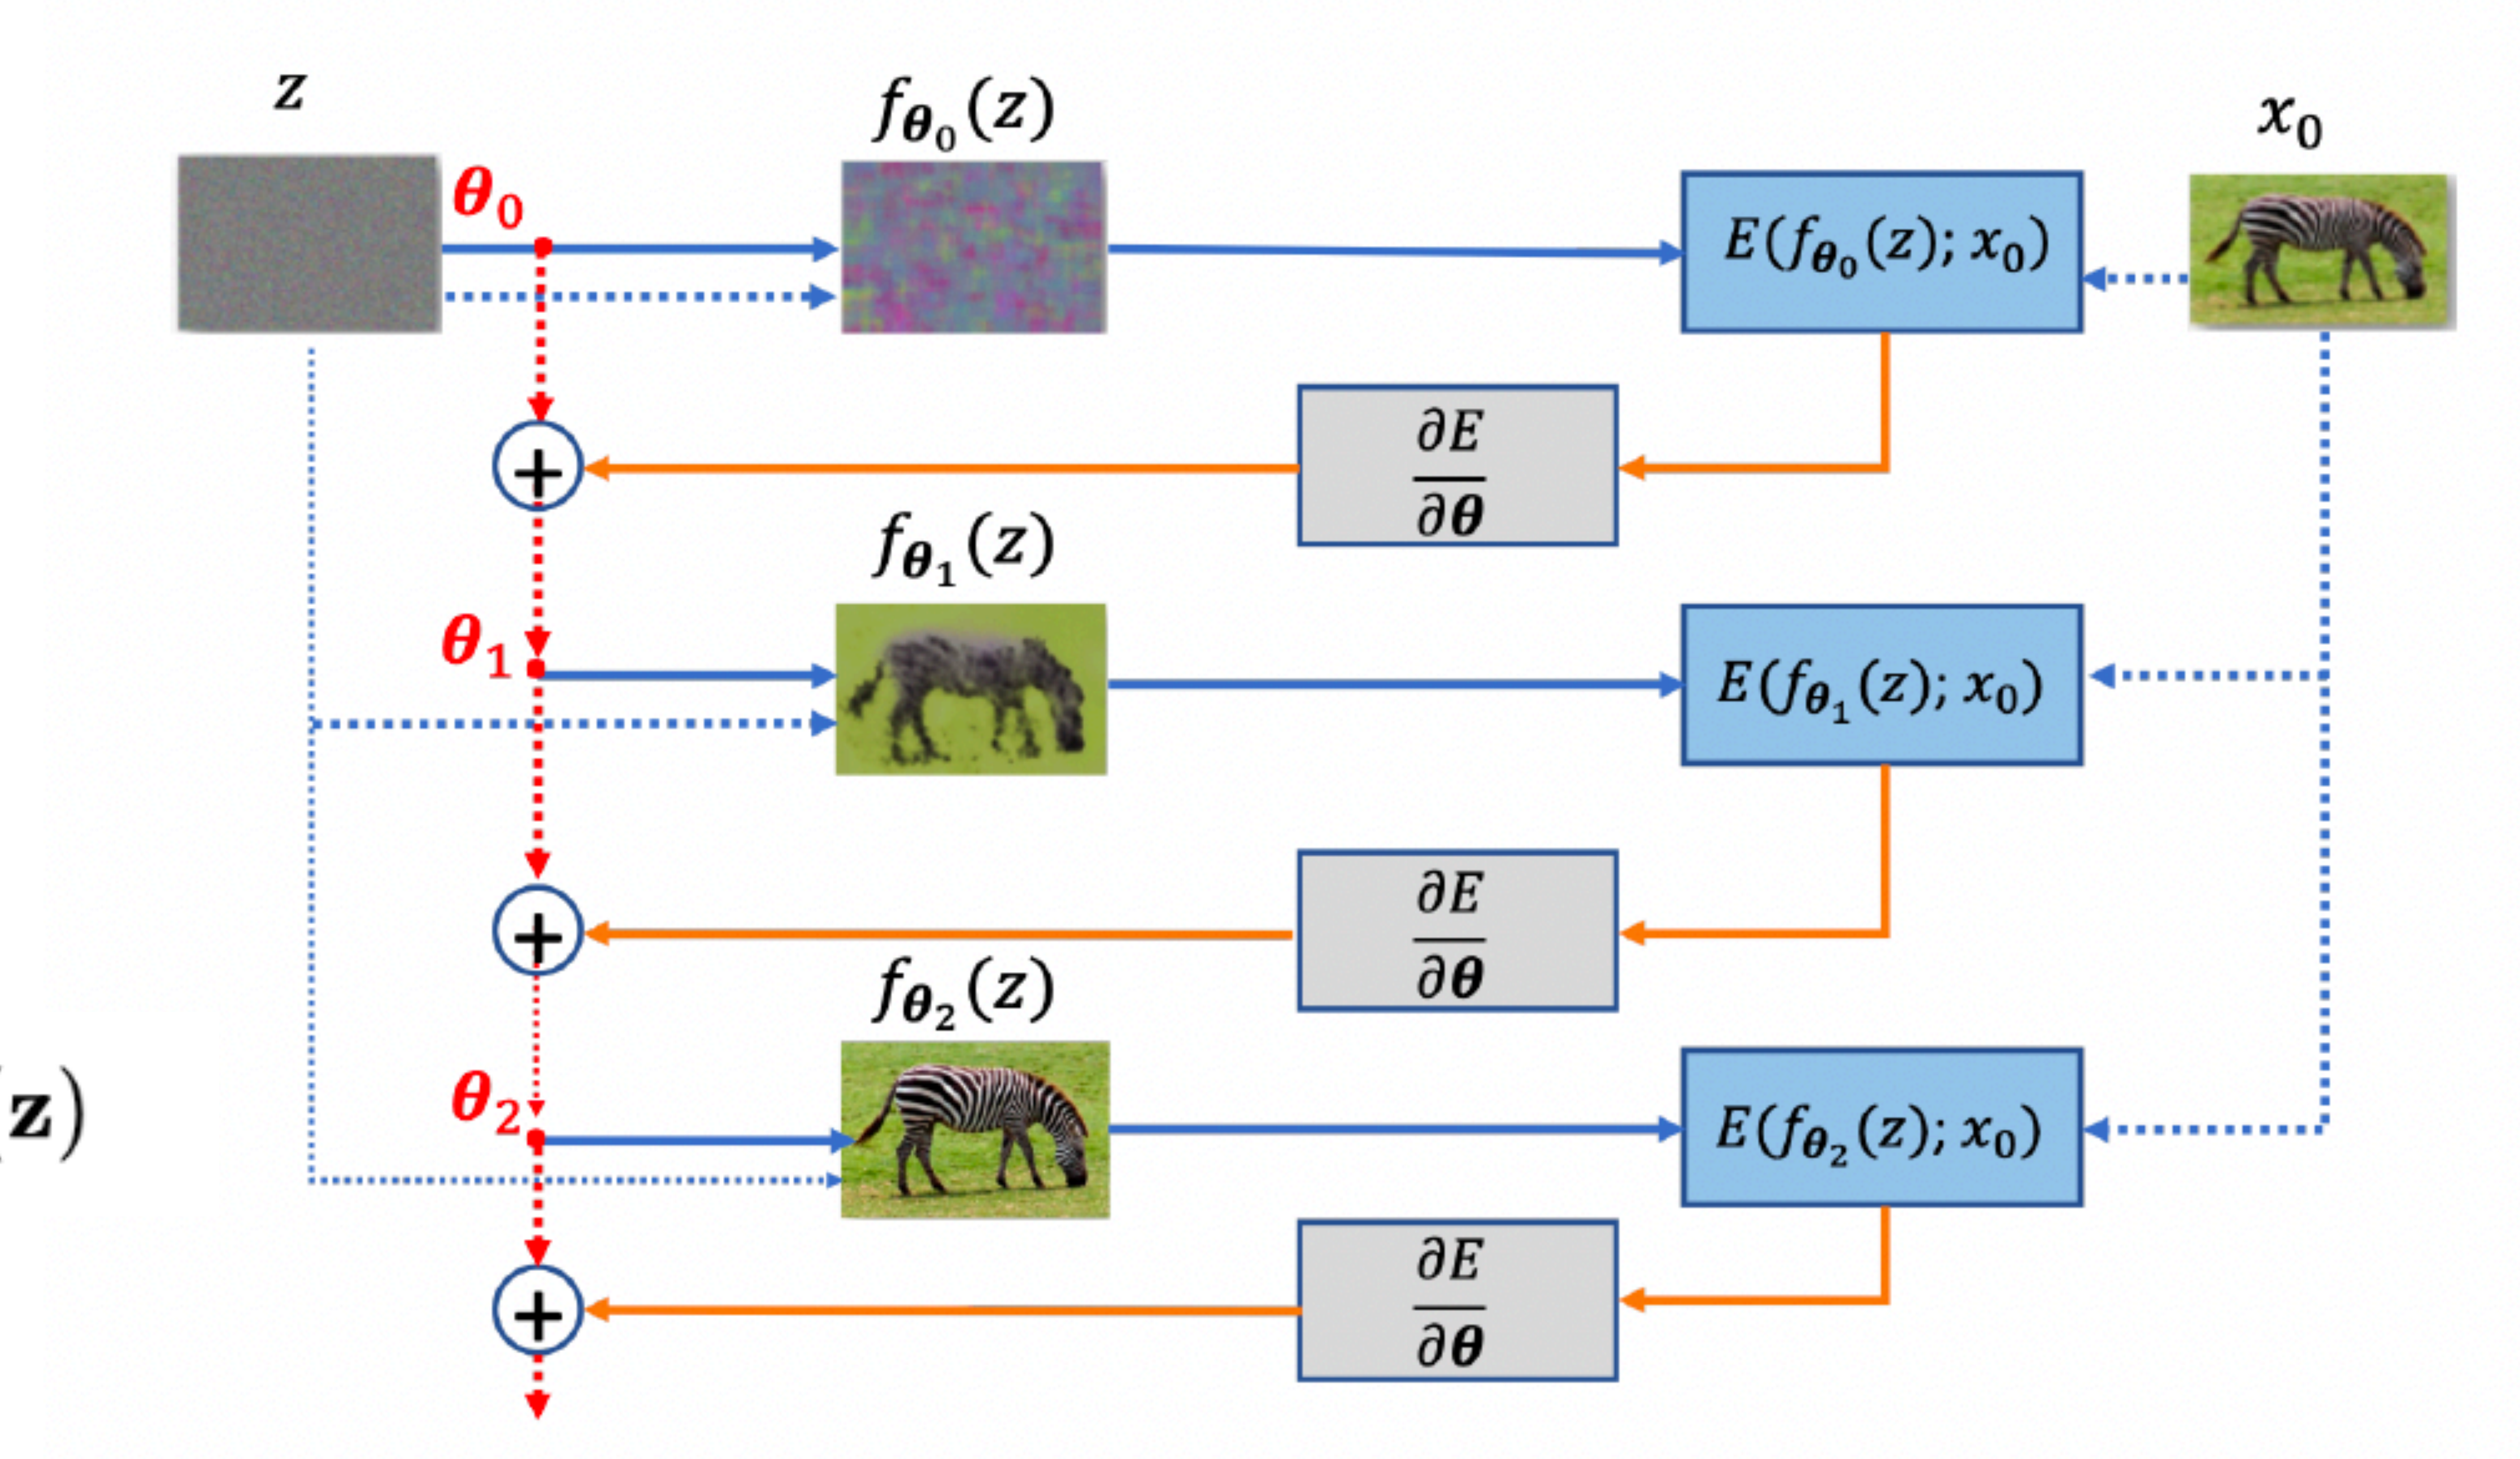

Credit to [1].

Deep Image Prior Is Good At…

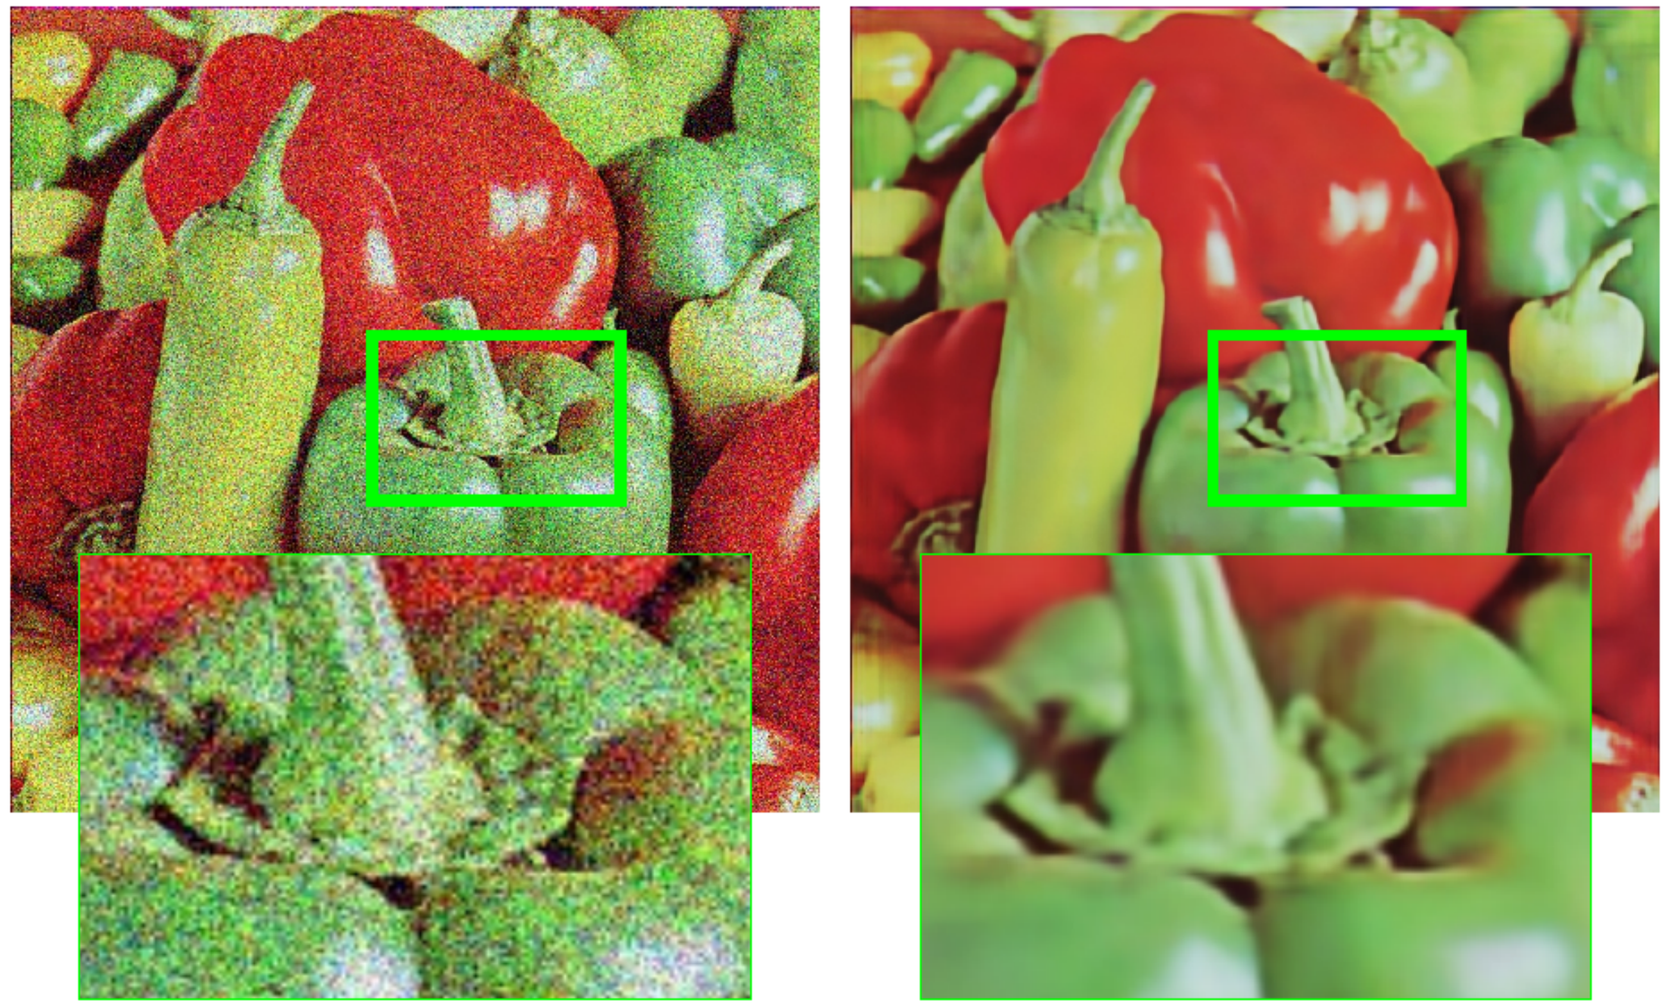

Image denoising

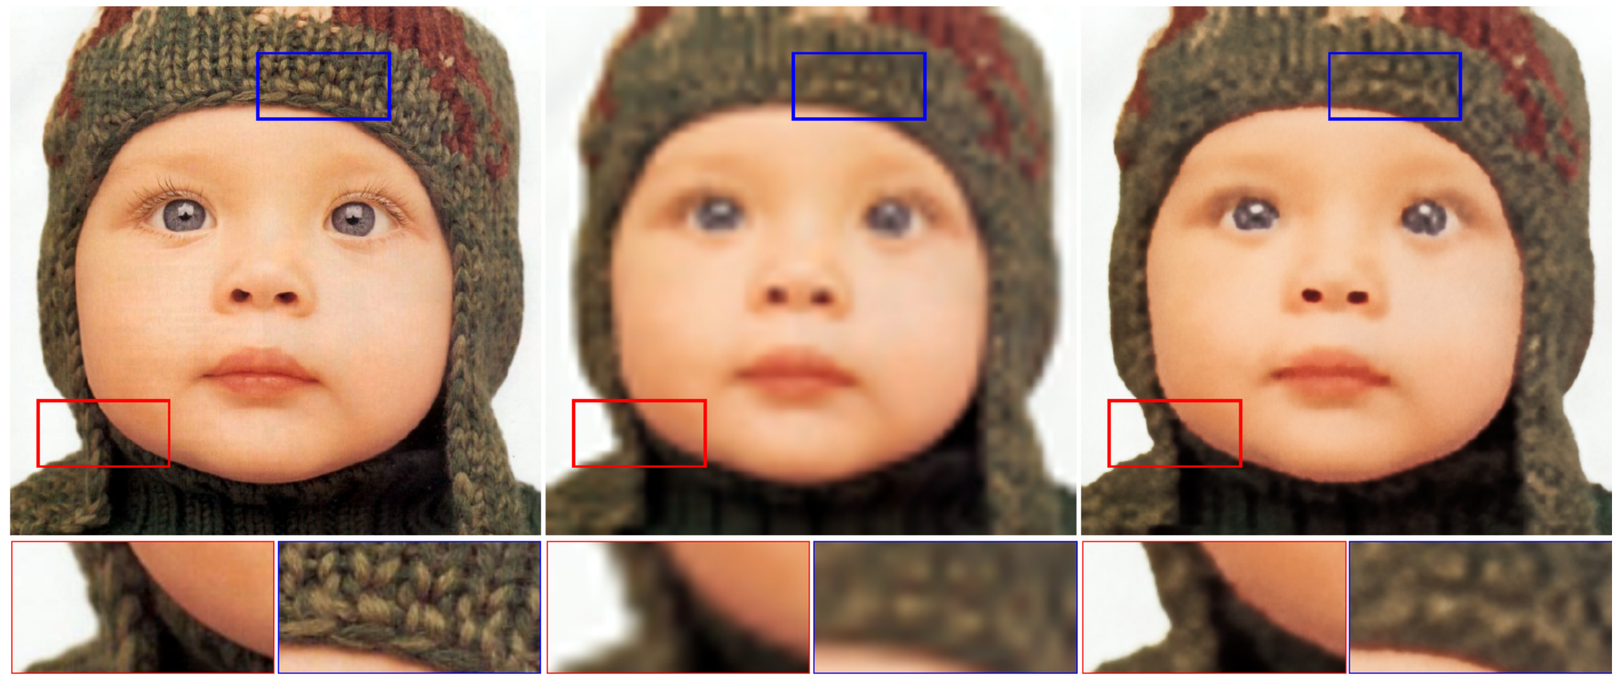

Image super-resolution

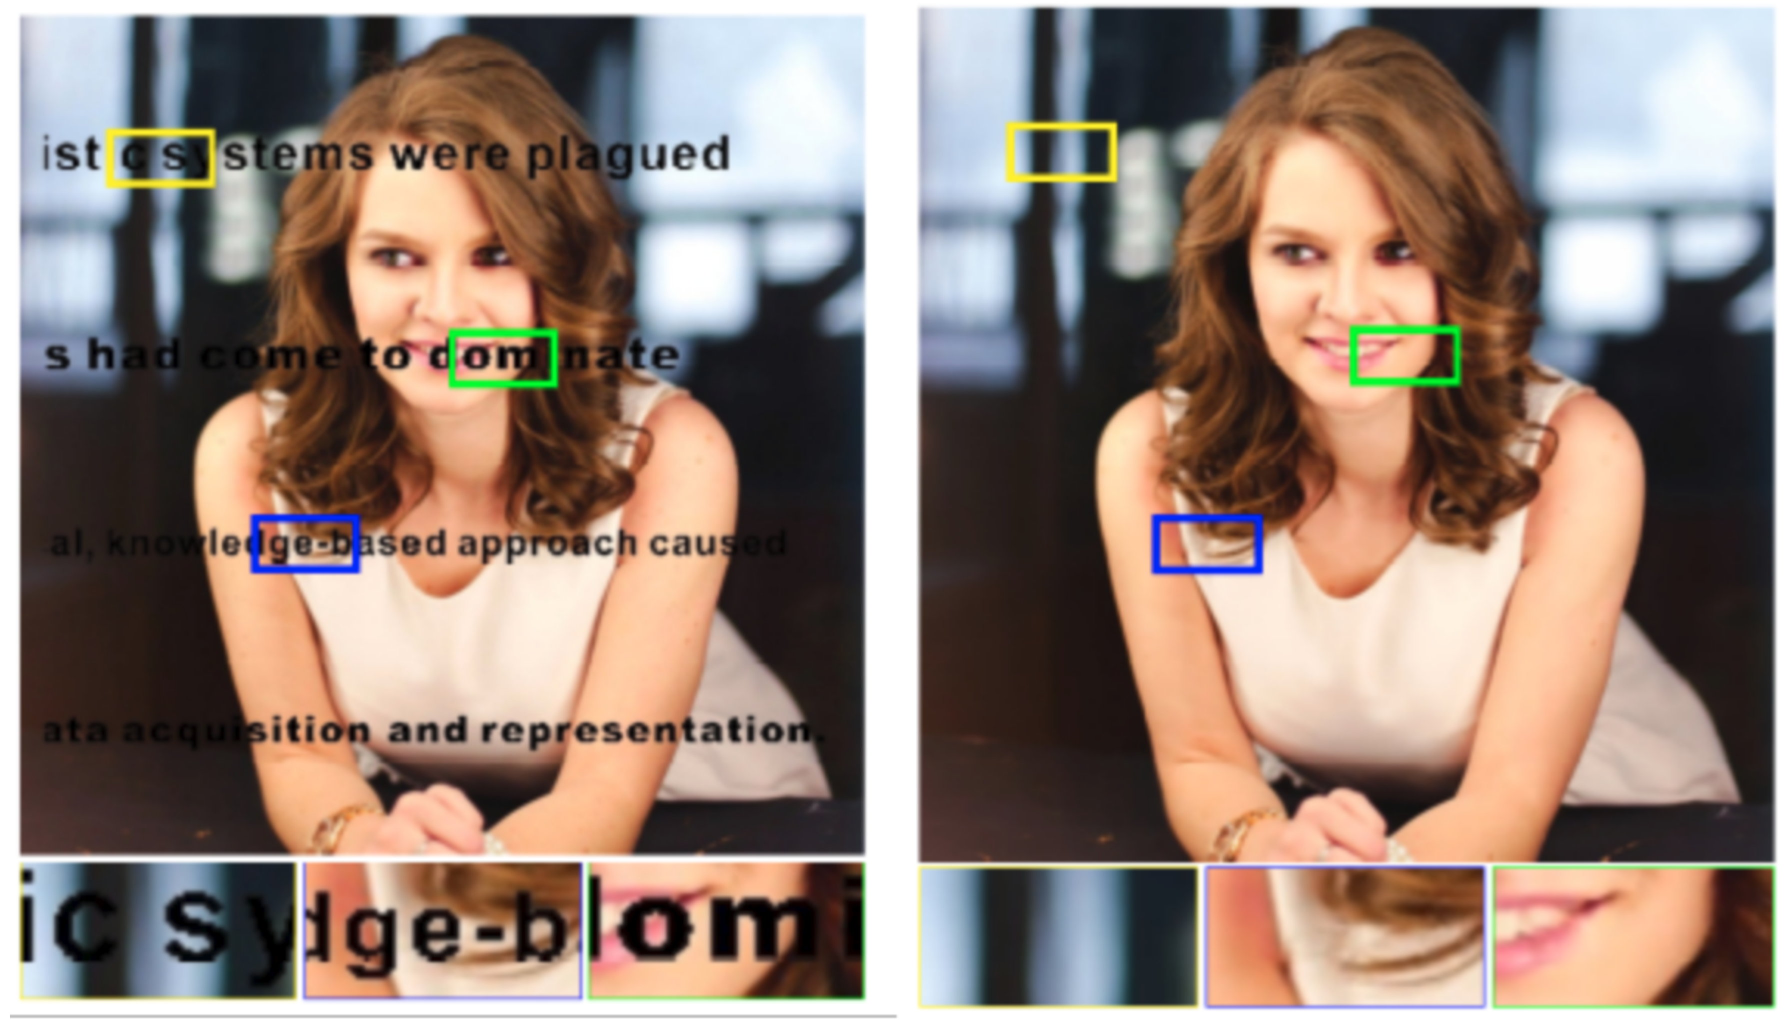

Image inpainting

All figures are credited to [1].

Deep Image Prior Suffers From…

- Overfitting to noise
- Slowness in optimization

Key(s) to Accelerating Optimization

Optimizing the input seed, not the network → improve the convergence

$$\min_{\theta} \ell(\mathbf{y}, f \circ G_\theta(\mathbf{z})) \longrightarrow \min_{\boldsymbol{z}} \ell(\boldsymbol{y}, f \circ G_\theta(\boldsymbol{z}))$$

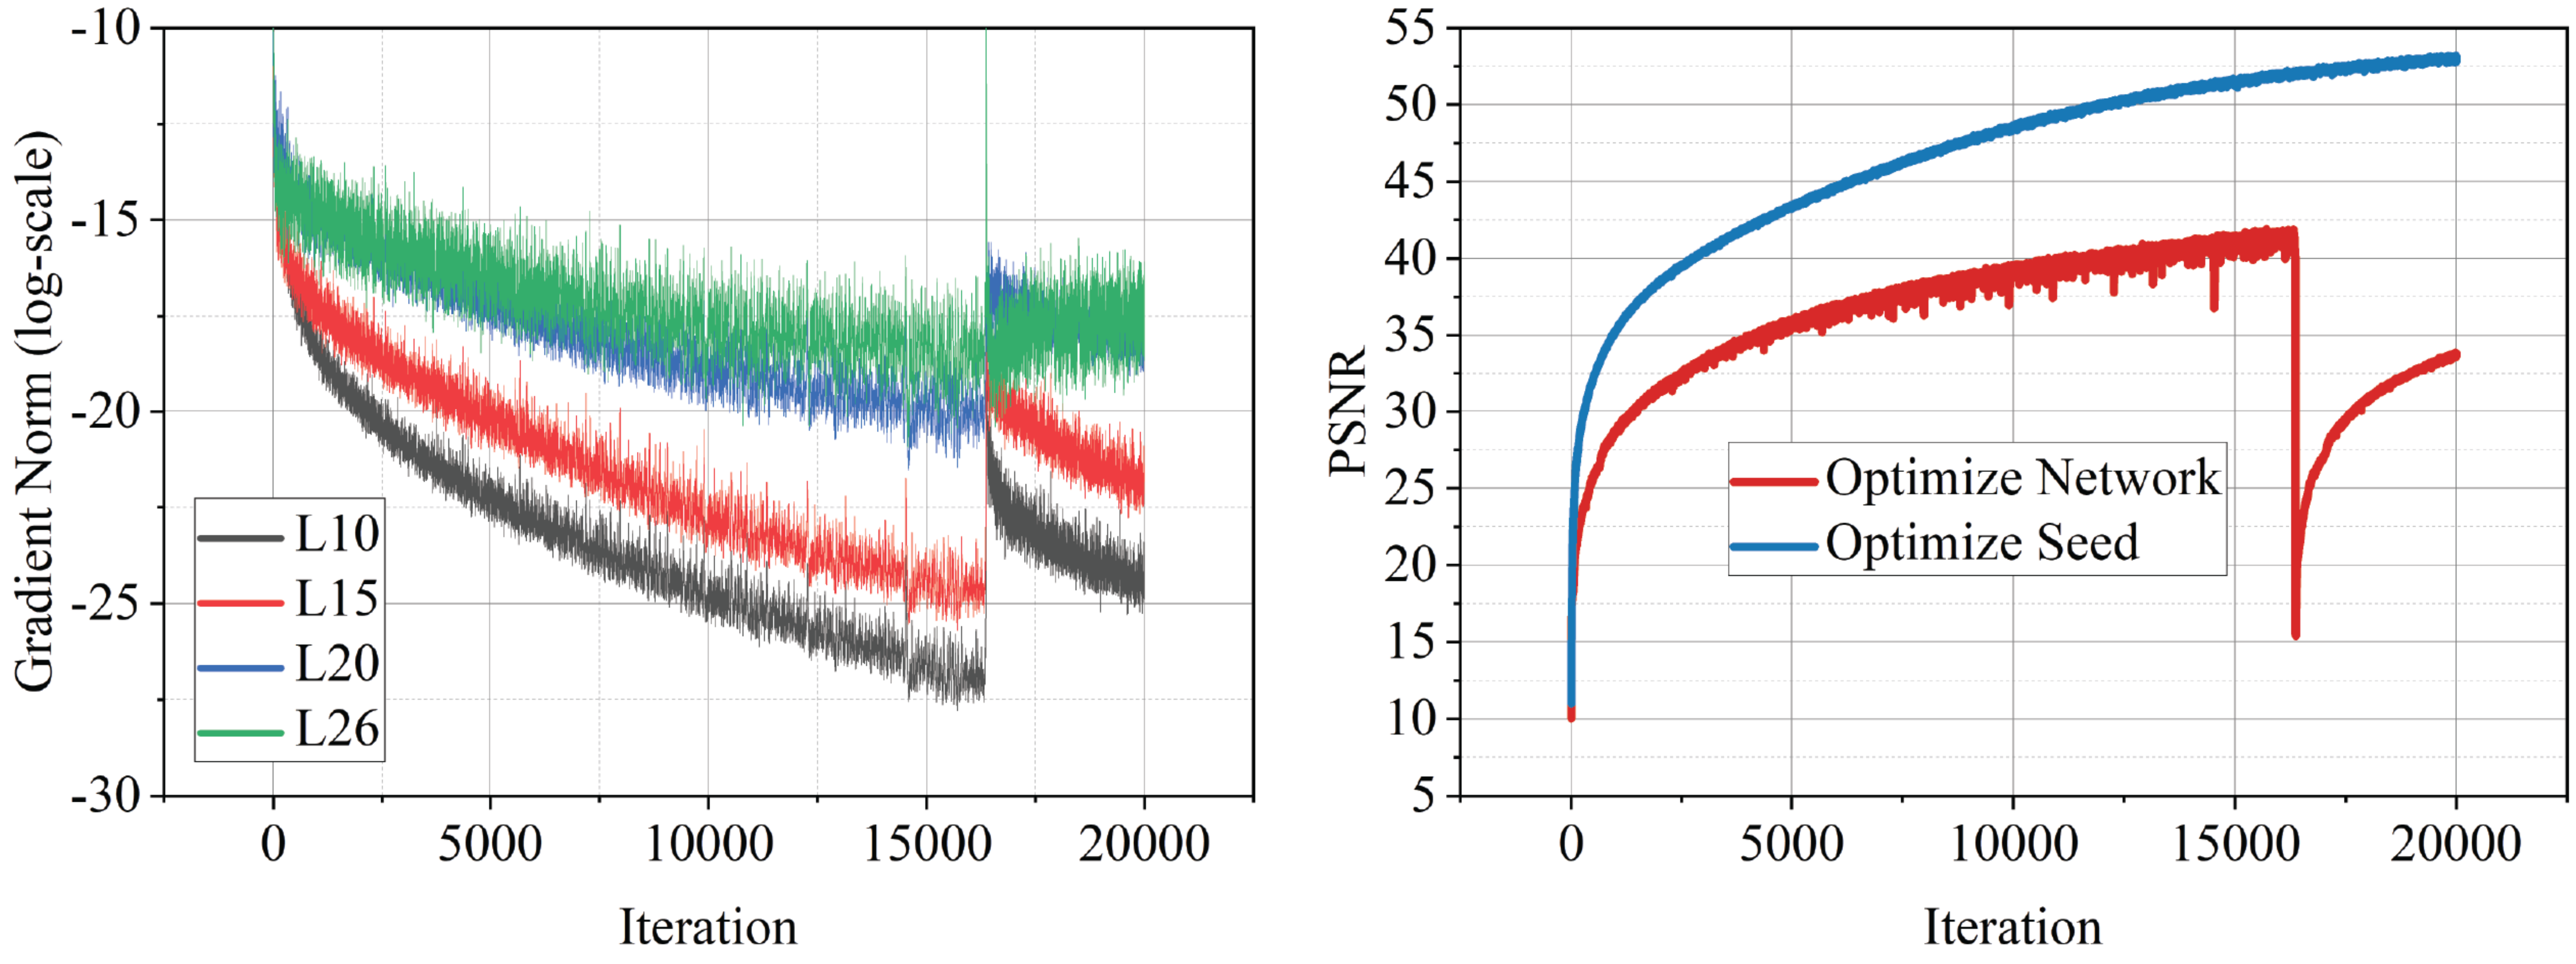

Key(s) to Accelerating Optimization

Cutting down the network depth:

1) improve per-iteration cost

2) speed up the convergence

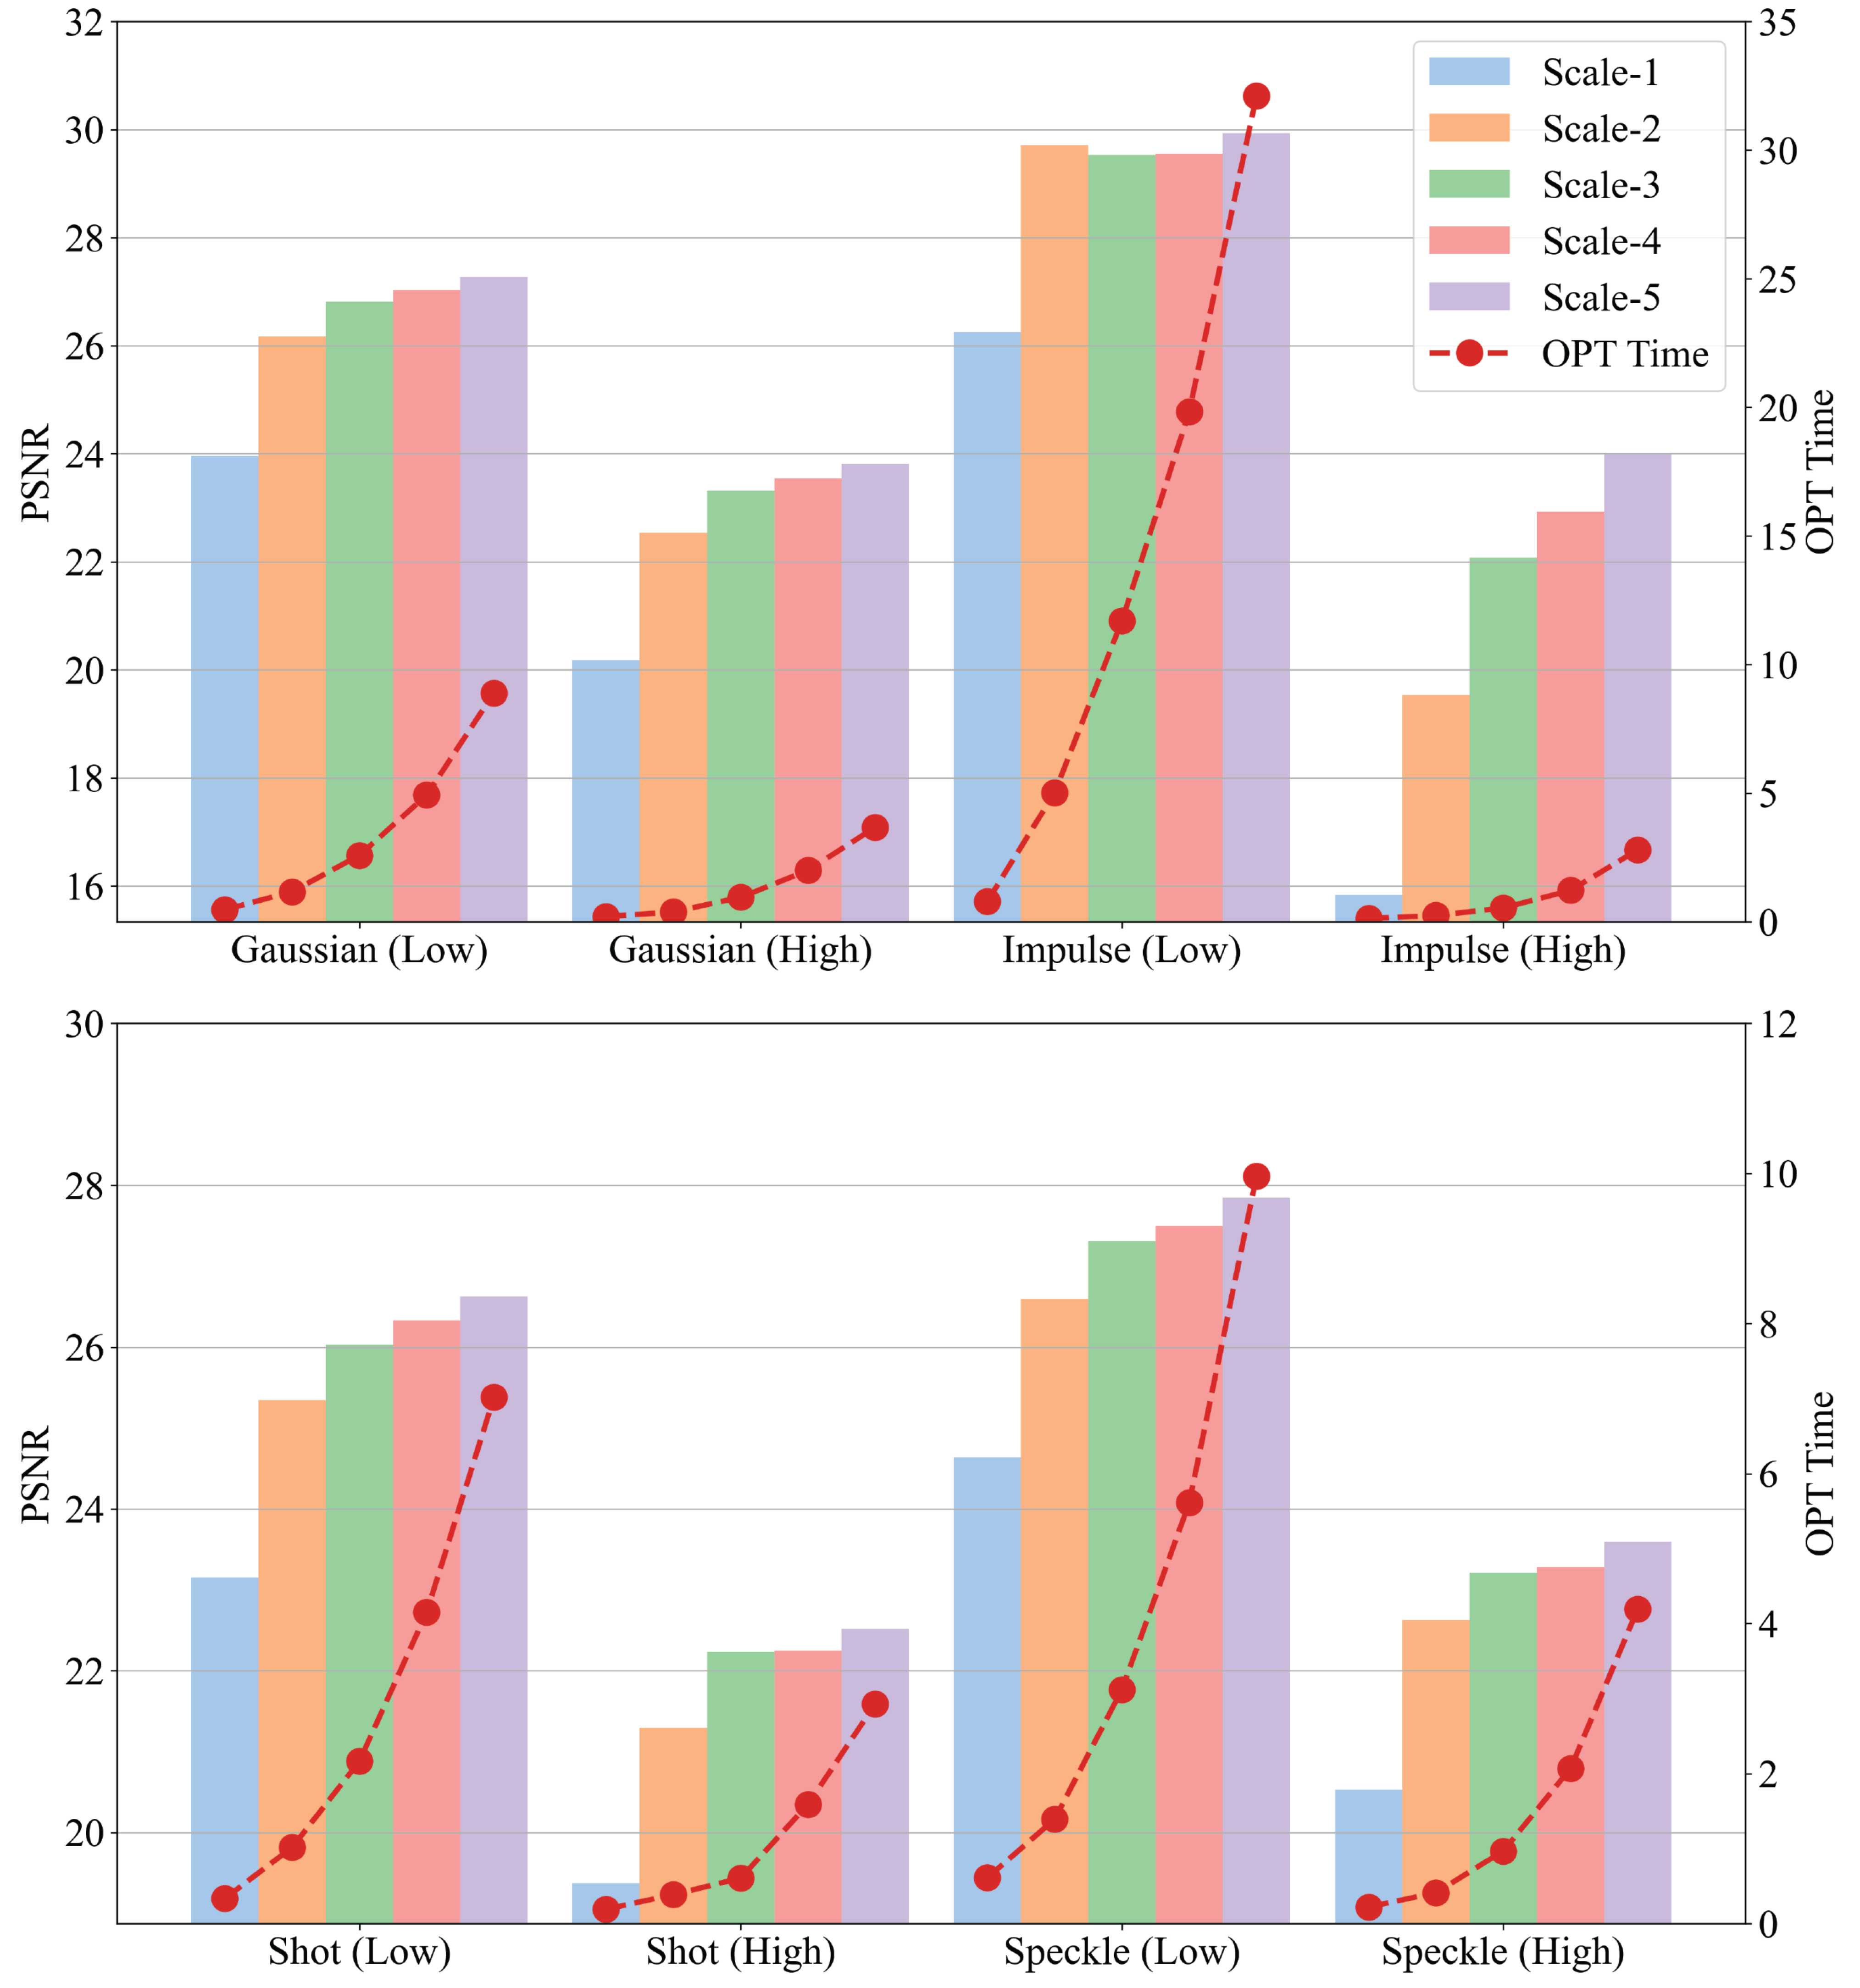

Key(s) to Retaining Restoration Quality

Adopting an explicit regularizer:

$$\min_{\boldsymbol{z},\boldsymbol{\theta}_{BN}} \|\boldsymbol{y} - f \circ G_{\boldsymbol{\theta}}(\boldsymbol{z})\| + \lambda \mathrm{TV}(G_{\boldsymbol{\theta}}(\boldsymbol{z})),$$

Experiments: Implementation

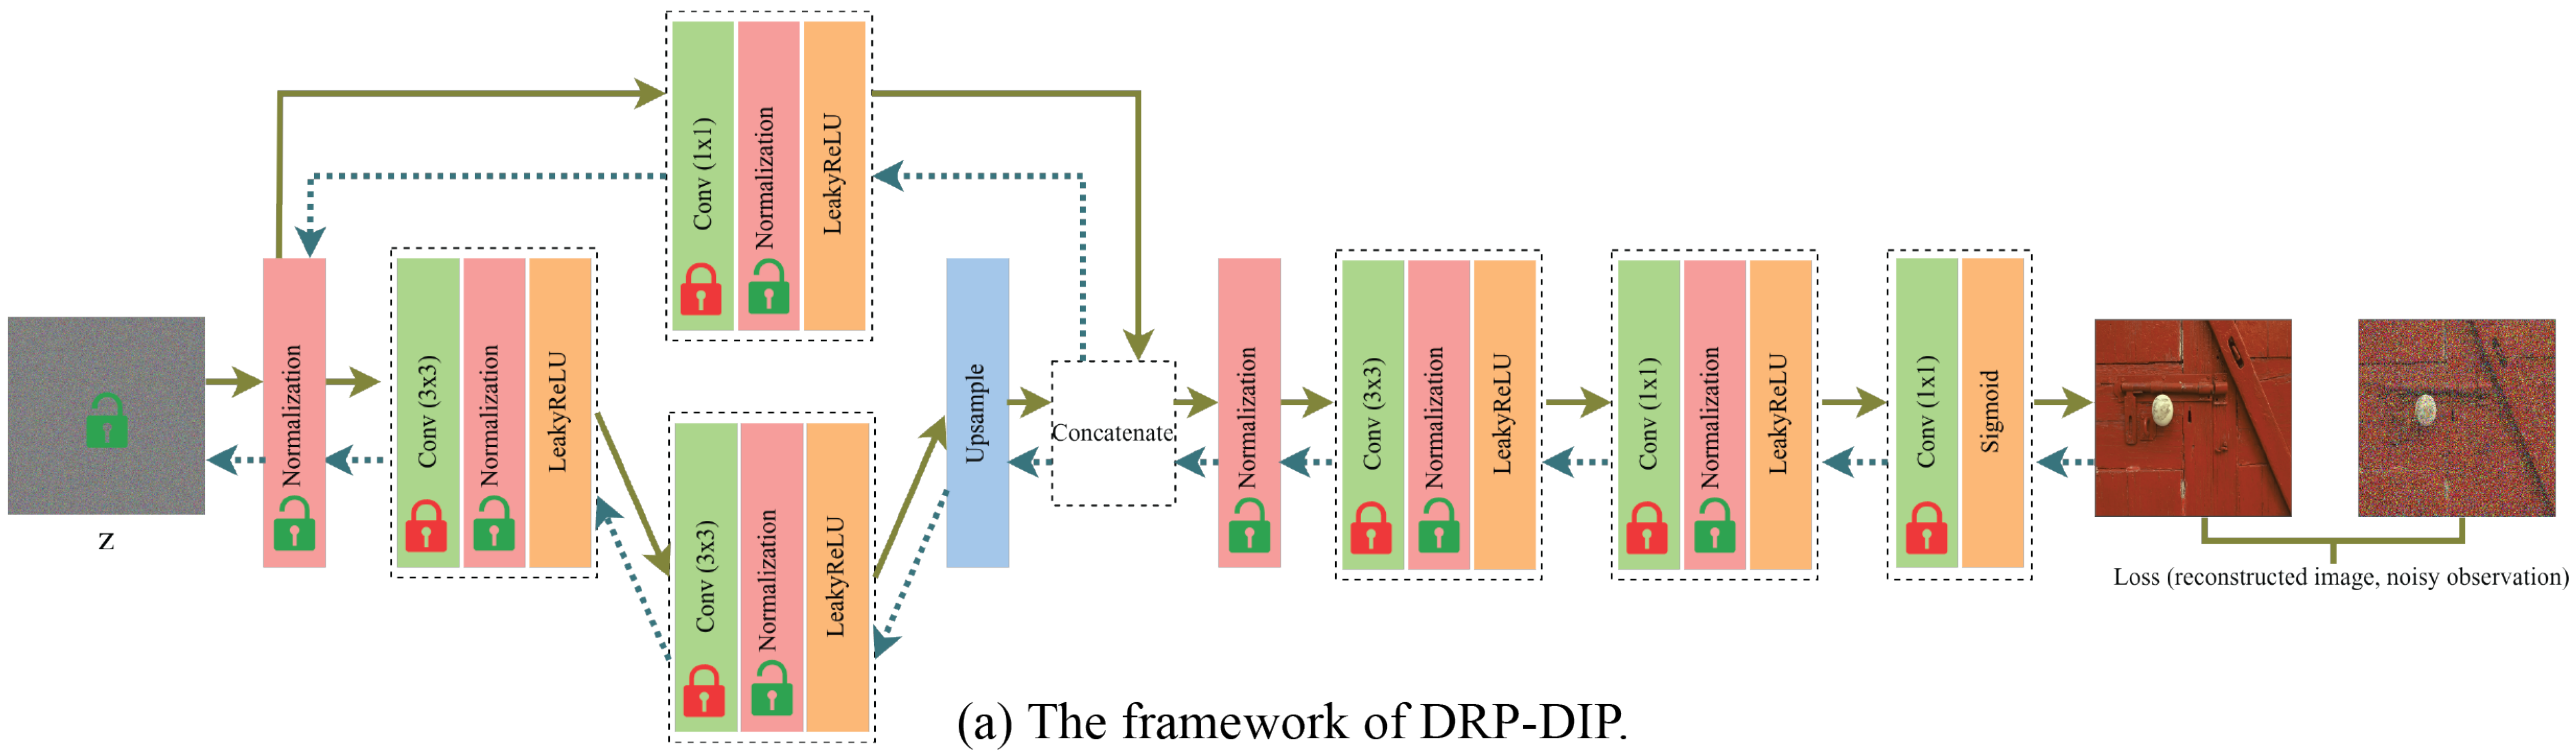

(a) The framework of DRP-DIP.

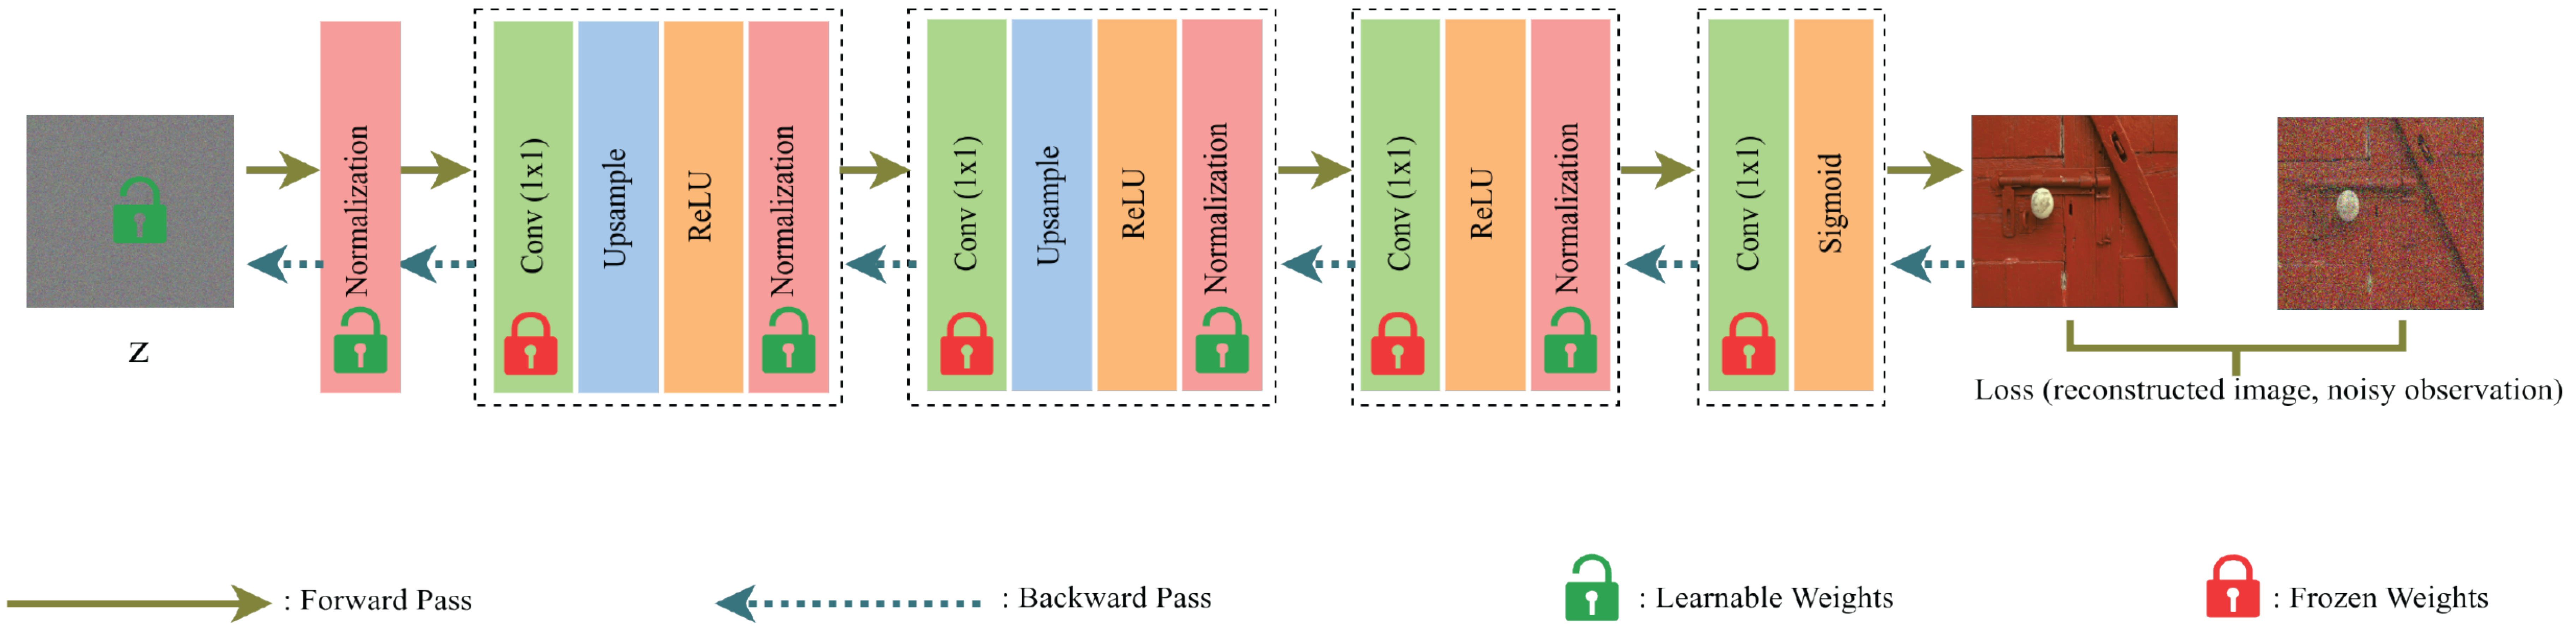

(b) The framework of DRP-DD.

Experiments: Image Denoising

Problem setting:

- given a noisy image $\mathbf{y} = \mathbf{x} + n$
- restore $\mathbf{x}$

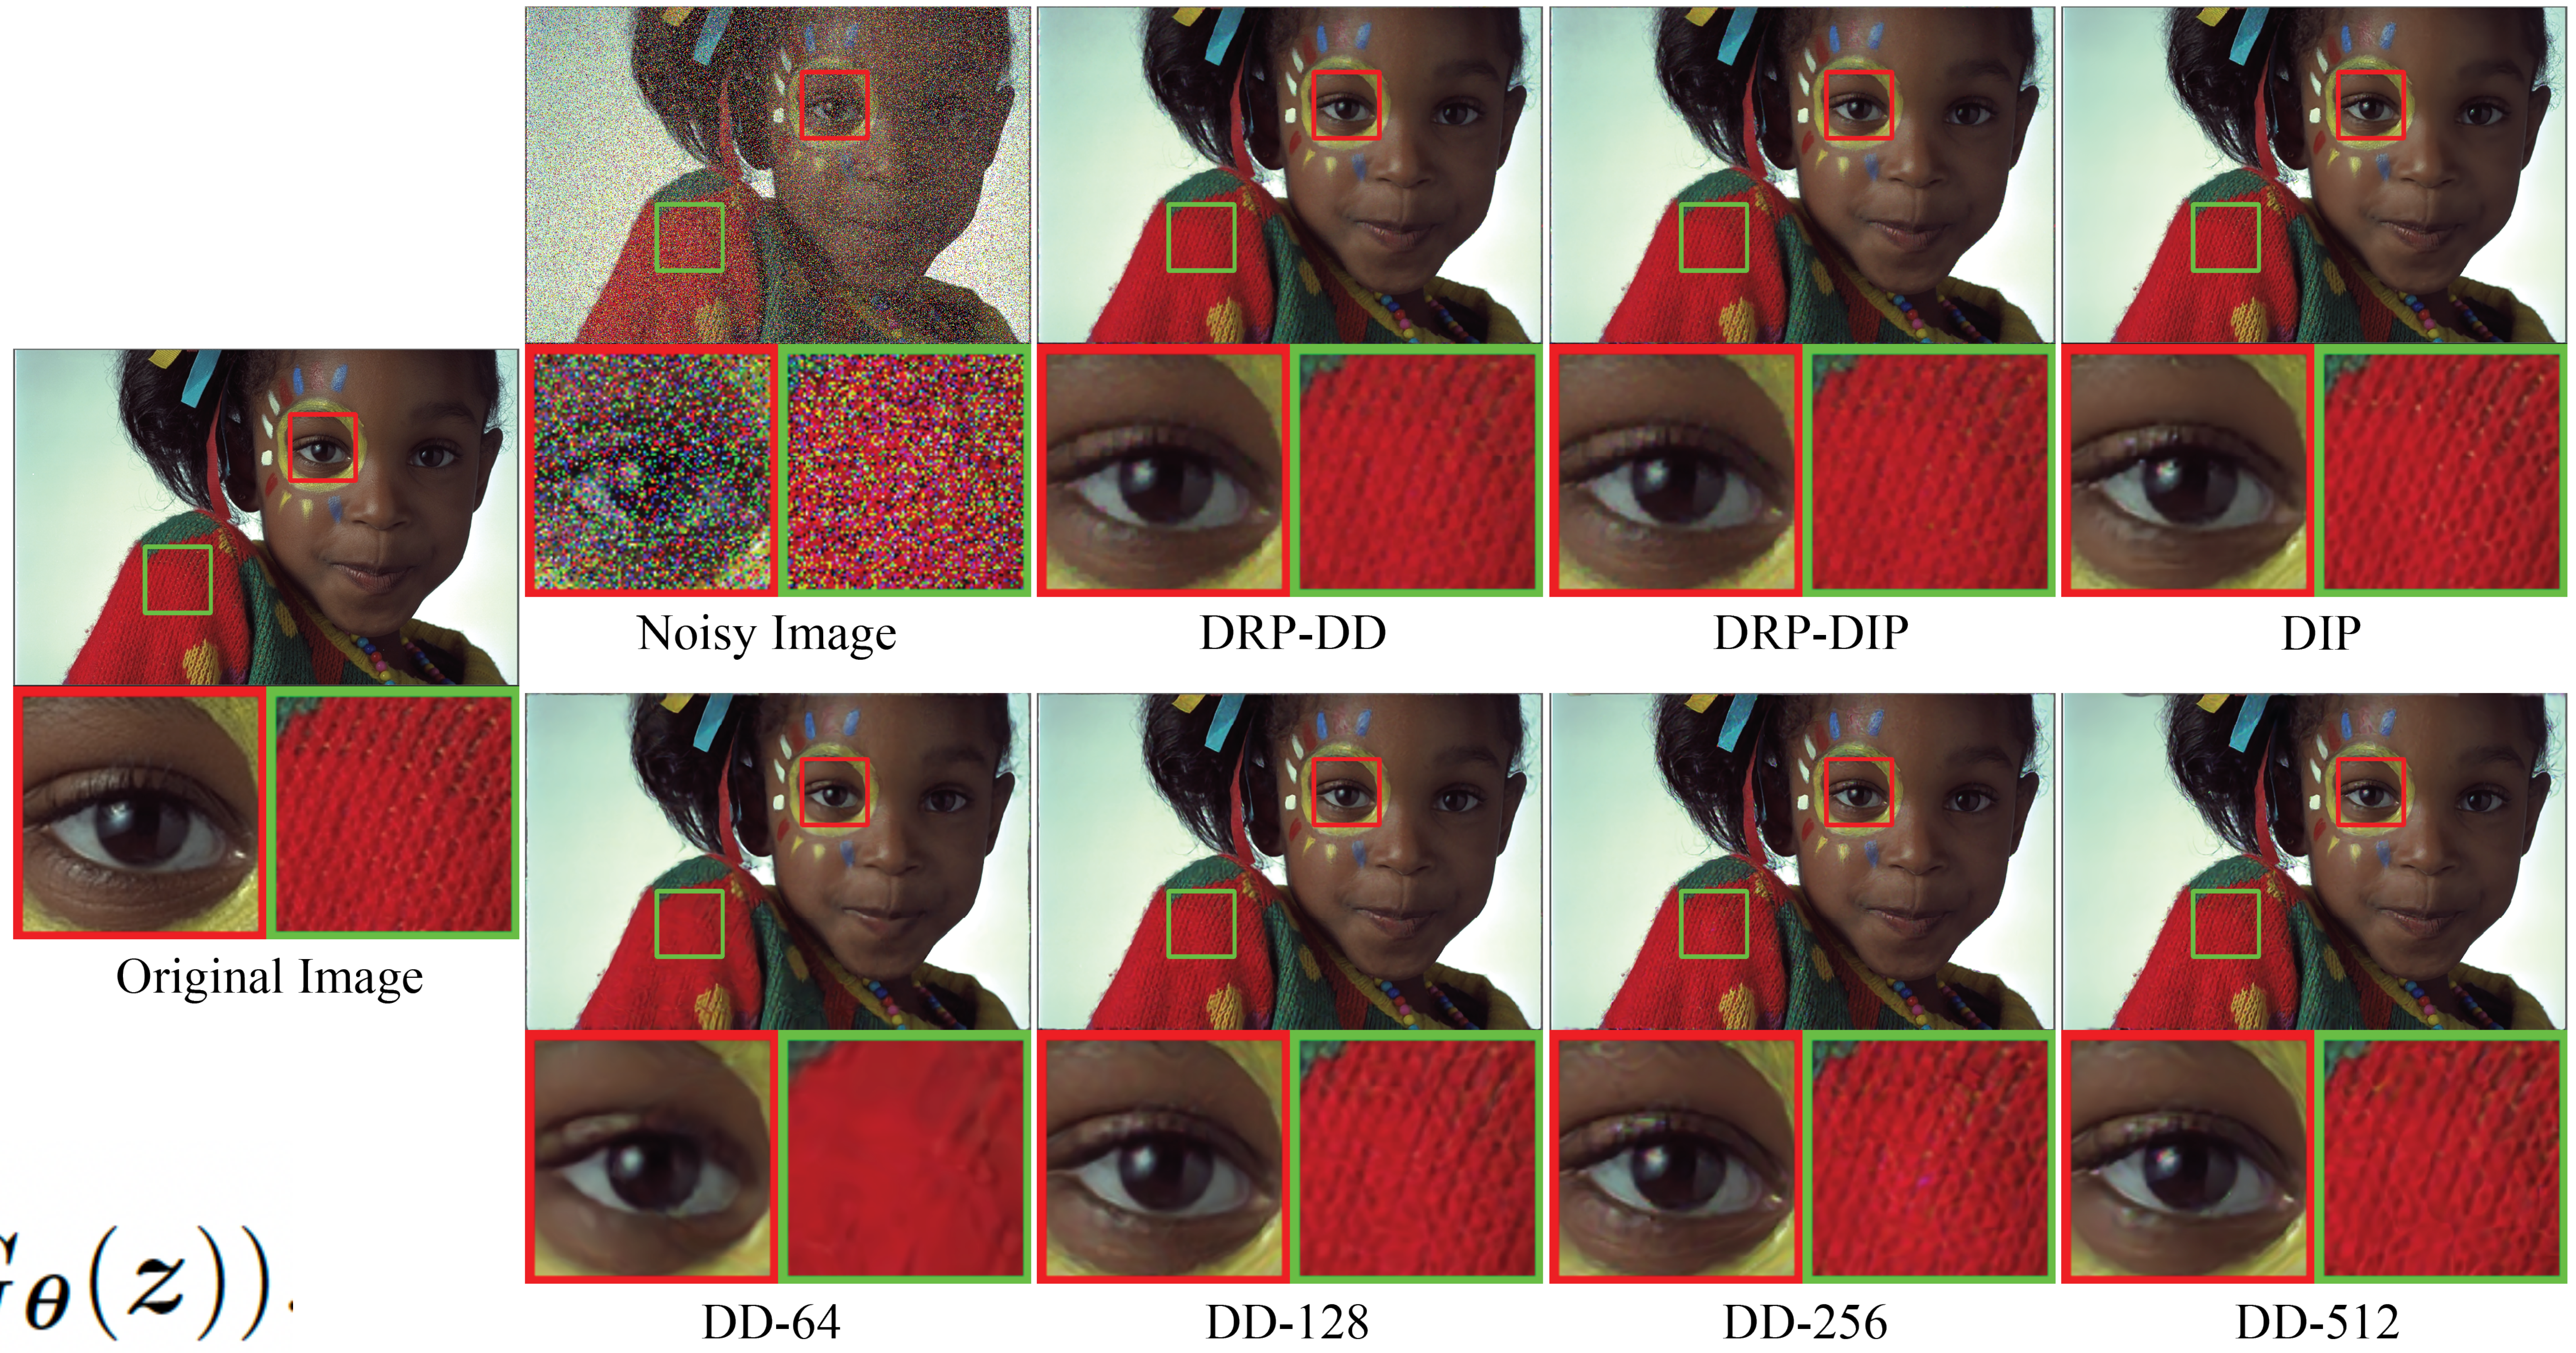

Objective function:

$$\min_{\boldsymbol{z},\boldsymbol{\theta}_{BN}} \|\boldsymbol{y} - G_{\boldsymbol{\theta}}(\boldsymbol{z})\| + \lambda \mathrm{TV}(G_{\boldsymbol{\theta}}(\boldsymbol{z})).$$

Experiments: Image Denoising

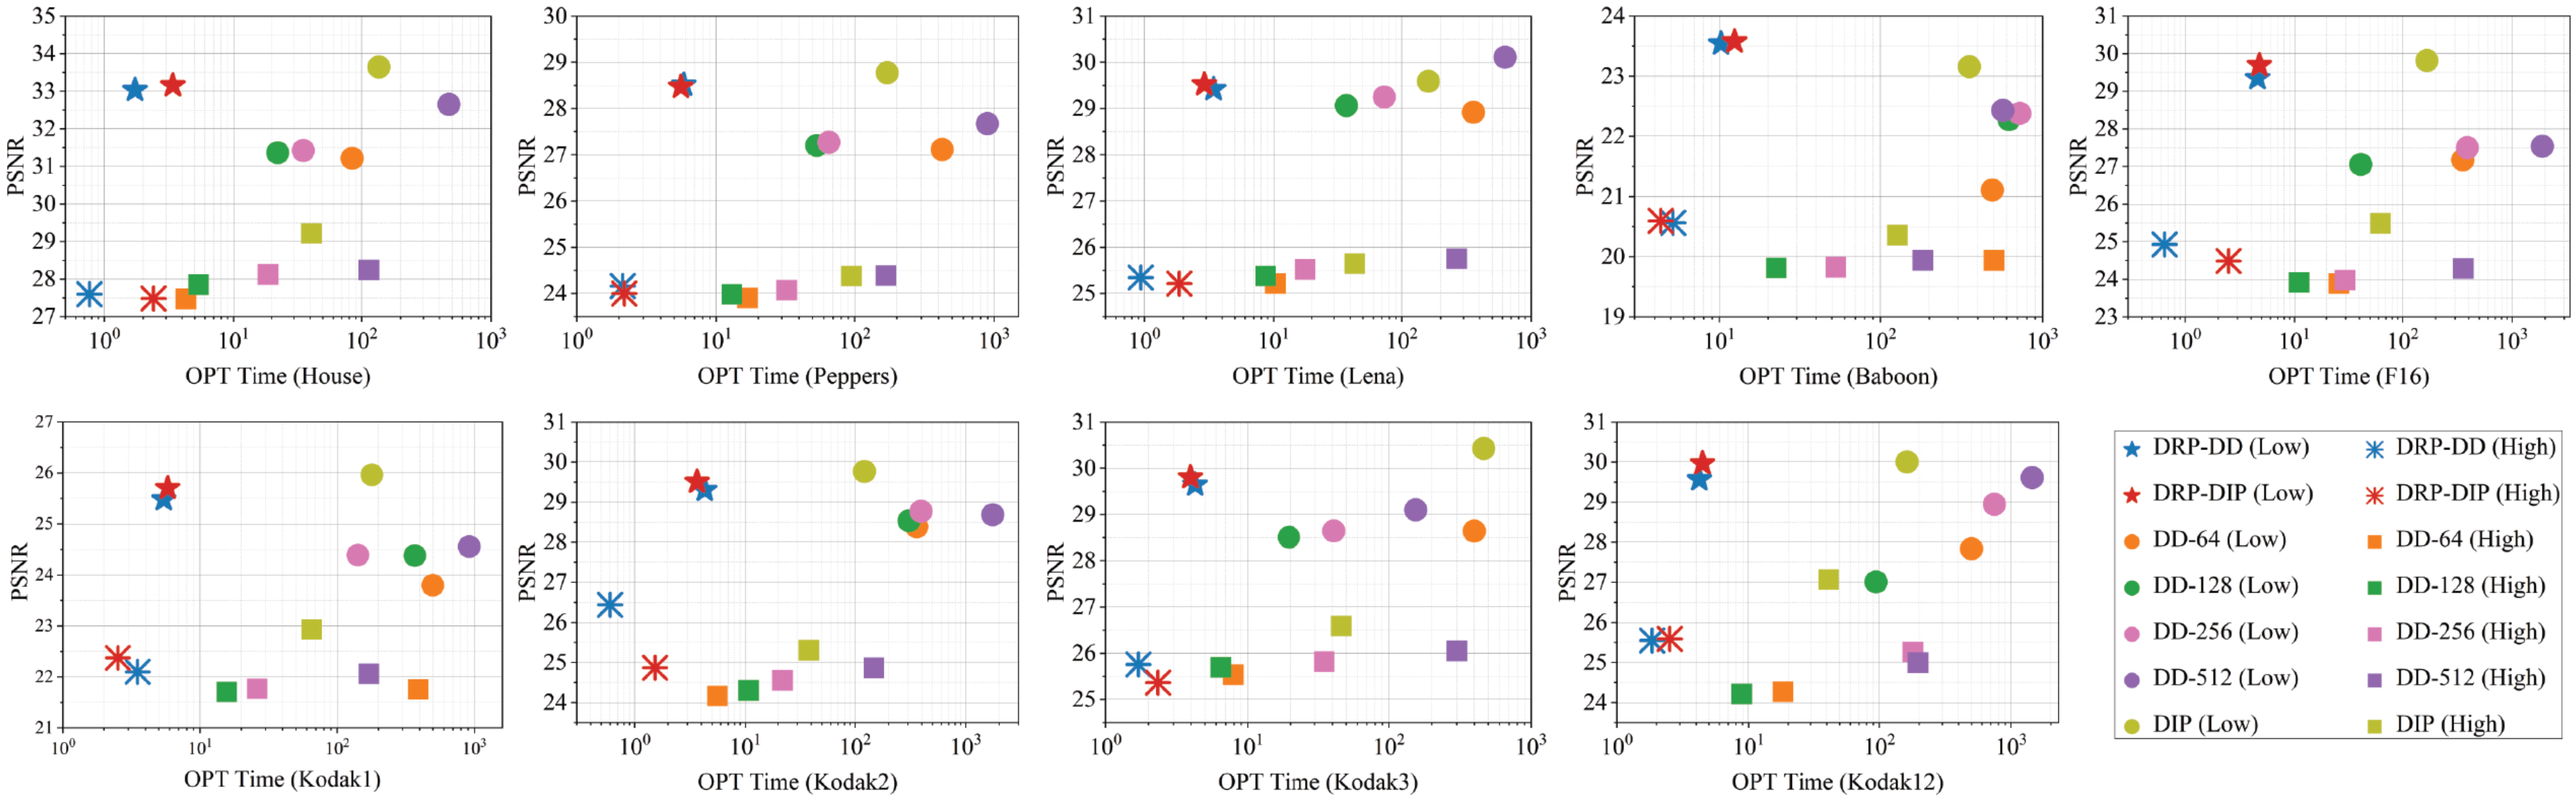

Experiments: Image Inpainting

Problem setting:

- given a noisy image $\mathbf{y} = \mathbf{x} \odot \mathbf{m}$ and the corresponding mask $\mathbf{m}$
- restore $\mathbf{x}$ by only using $\mathbf{y}$ and $\mathbf{m}$

Objective function:

$$\min_{\boldsymbol{z},\boldsymbol{\theta}_{BN}} \|\boldsymbol{y} - G_{\boldsymbol{\theta}}(\boldsymbol{z}) \odot \boldsymbol{m}\| + \lambda \mathrm{TV}(G_{\boldsymbol{\theta}}(\boldsymbol{z}))$$

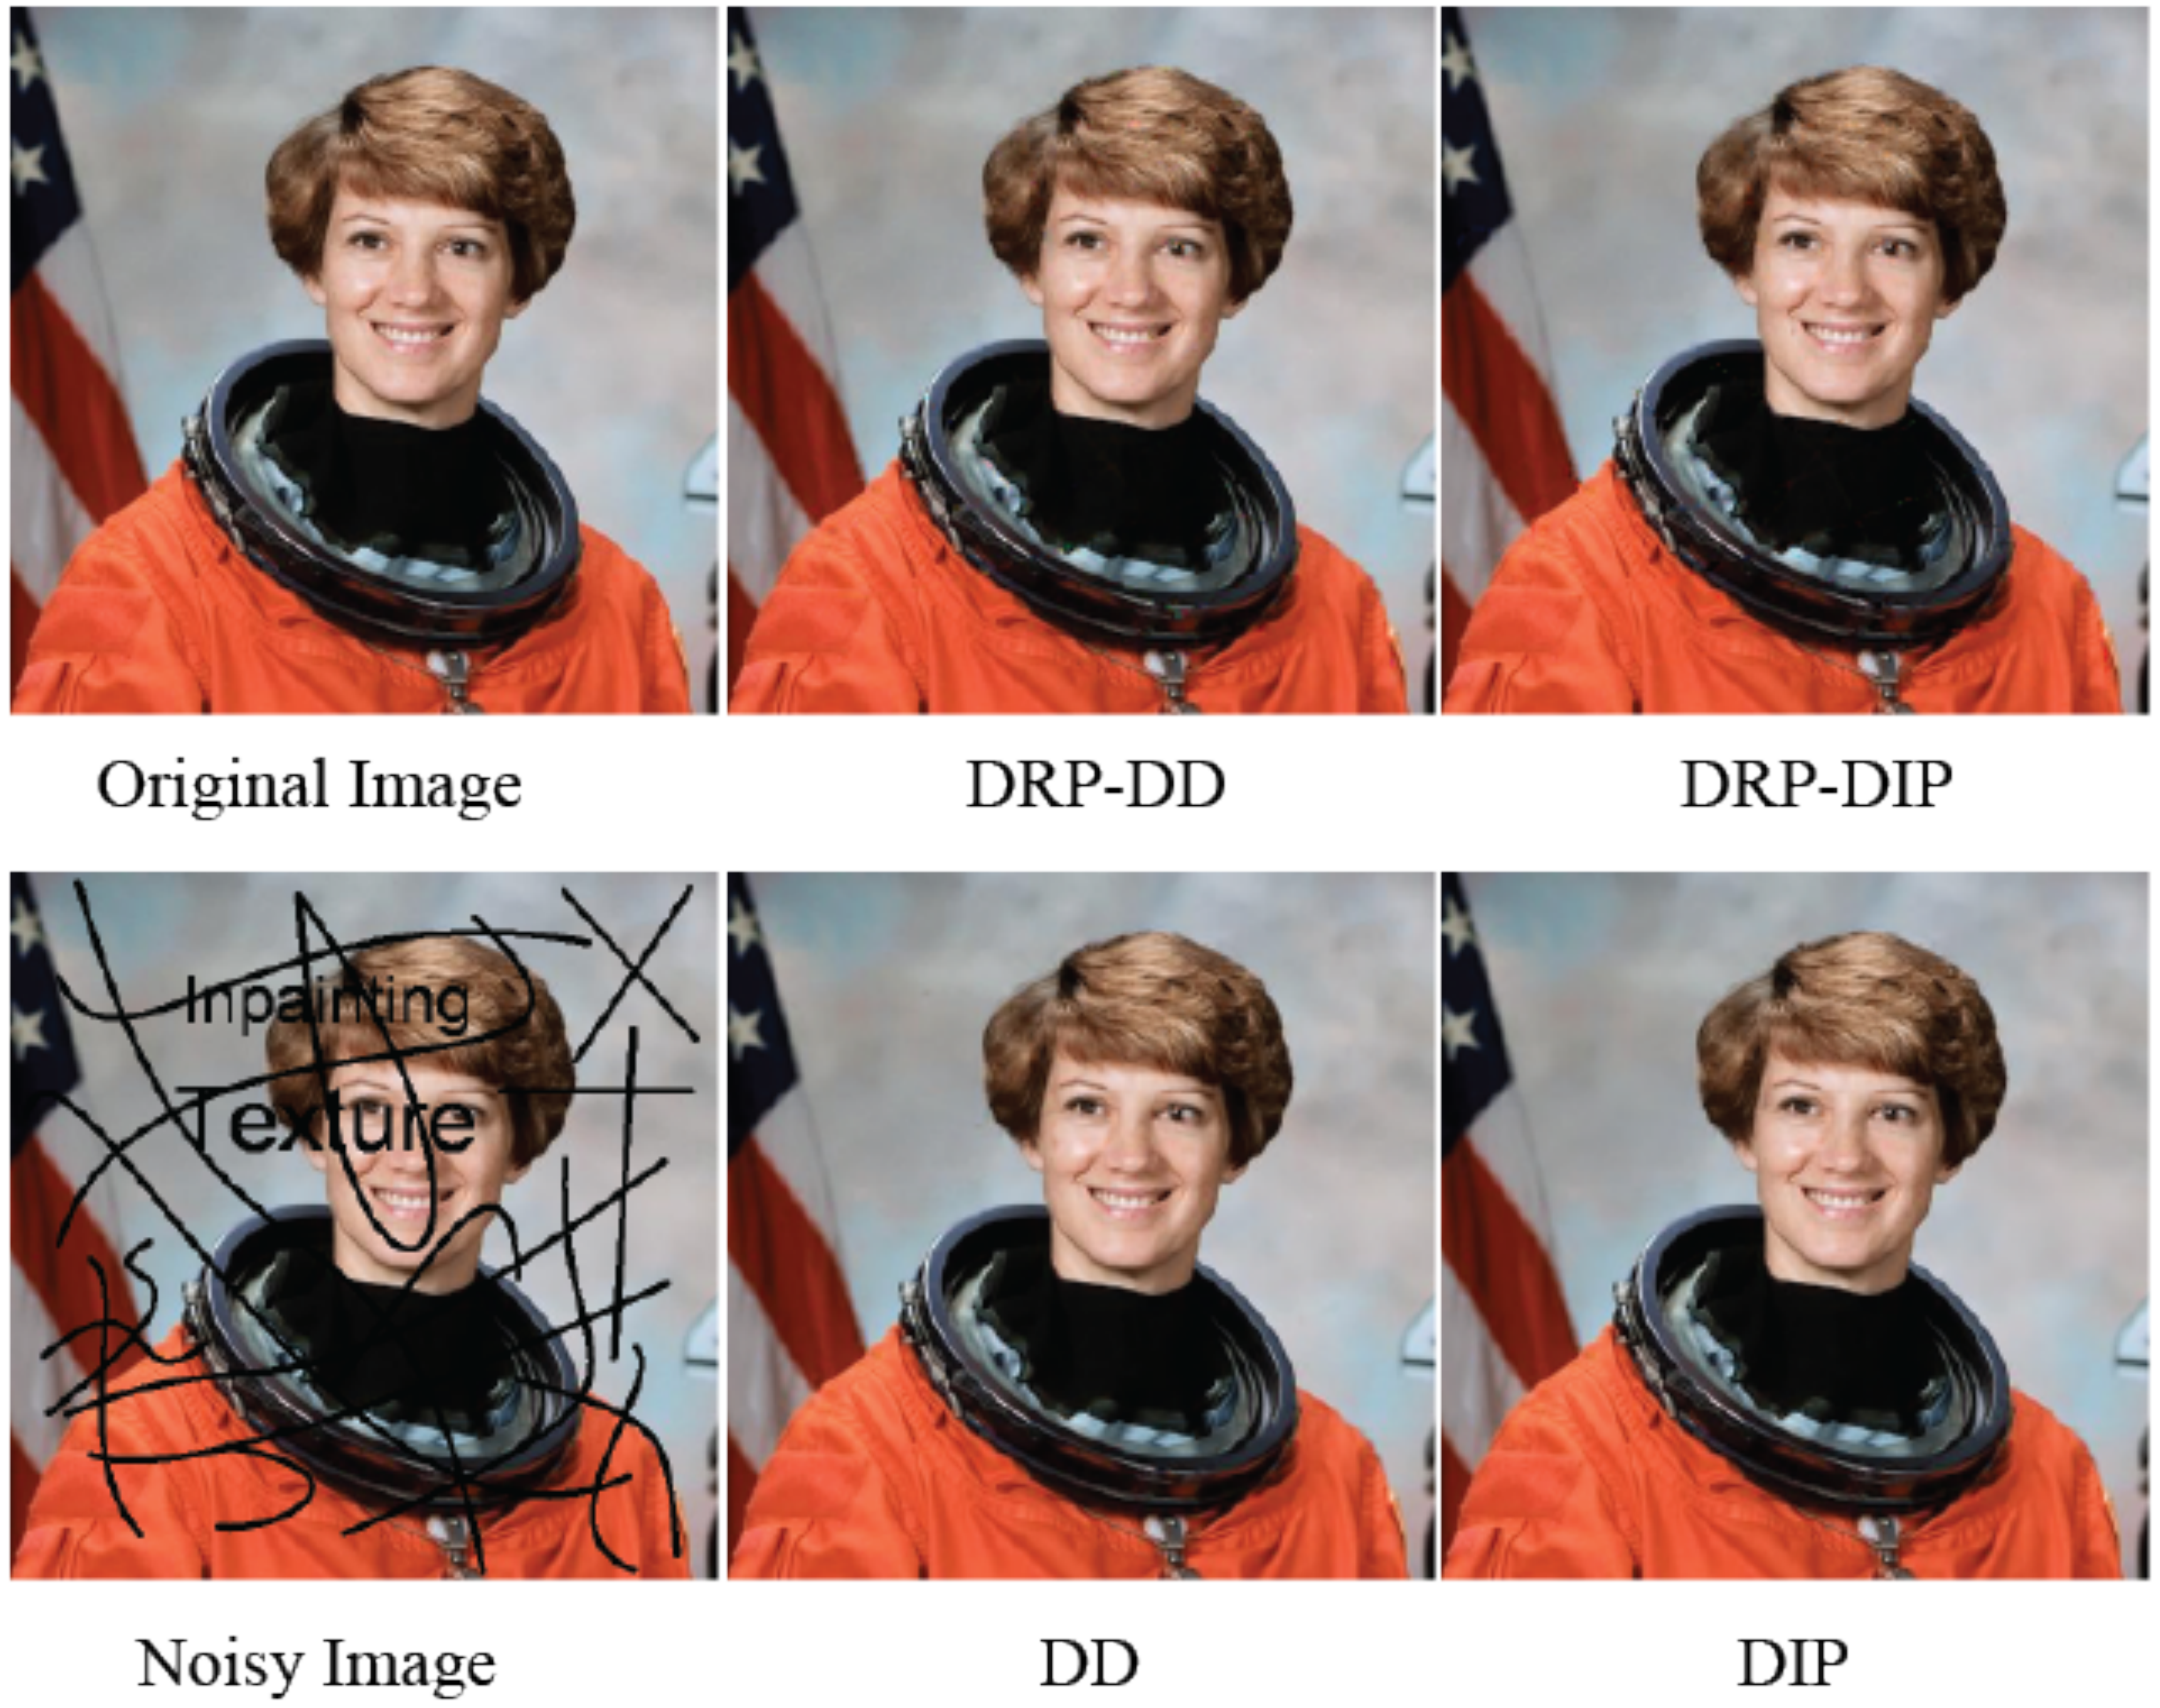

Experiments: Image Inpainting

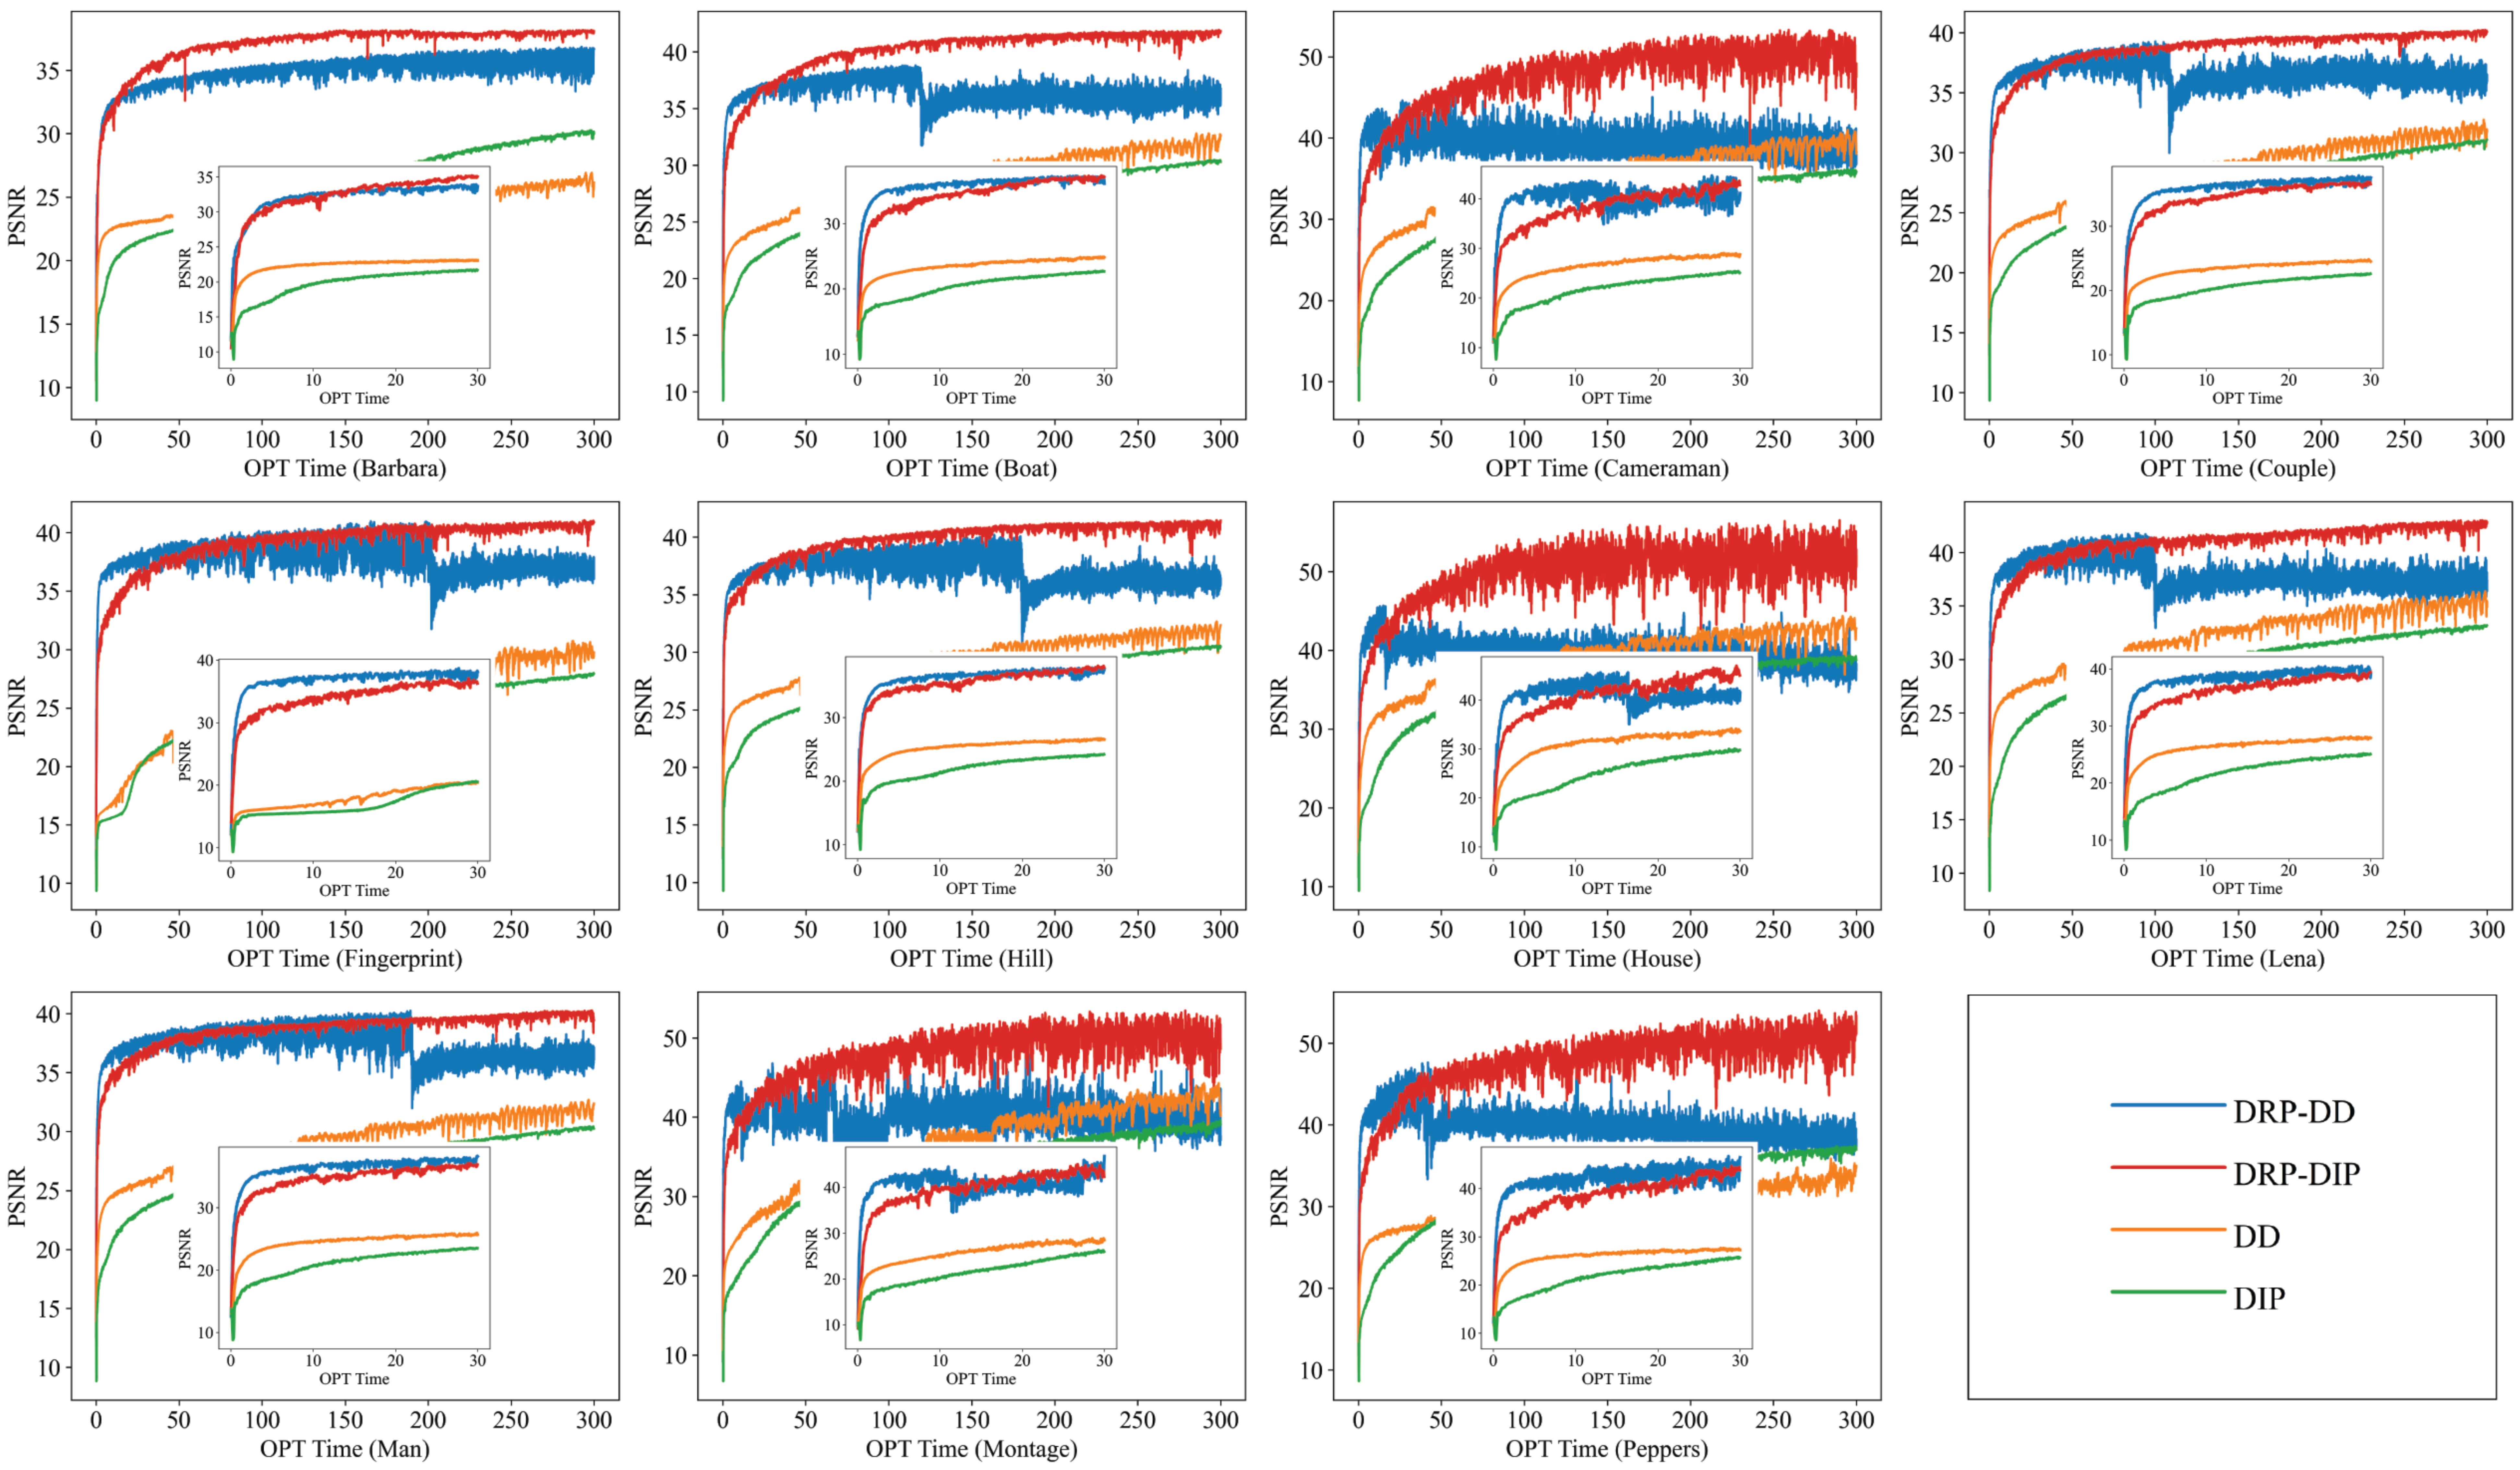

Conclusion

- We investigate the slow optimization issue of DIP
- We propose a simple-yet-effective method named Deep Random Projector to speed up DIP
- We evaluate our method on common image restoration tasks and verify its effectiveness

References

[1] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky, “Deep image prior,” in 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. 2018, pp. 9446–9454, Computer Vision Foundation / IEEE Computer Society.

Thank you!