Loss Functions for Image Restoration with Neural Networks

Hang Zhao, Orazio Gallo, Iuri Frosio, and Jan Kautz

NVIDIA Research and MIT Media Lab

1 Instructions

Please view this document in full screen (e.g., for Acrobat Adobe Reader CTRL+L on Windows and COMMAND+L on Mac). To compare the images resulting from different networks click the corresponding hyperlinks in the caption.

For each image the reader can either just scroll through the images to see the changes or jump from one image to the other by clicking the figure numbers at the bottom. For each image we provide a couple of details that are worth observing. It is important to see in full-screen mode so that the images are aligned.

We show results for:

- JPEG deblocking,
- joint denoising + demosaicking,
- different training schedules, and
- super-resolution.

JPEG deblocking

Fig. 1. JPEG deblocking – ℓ_1
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 fails at removing both the artifacts in the sky and the halo introduced by the JPEG compression at the edge of the building. ℓ_1 does better than ℓ_2, but only Mix succeeds at completely removing the halo and attenuating the artifacts in the sky.

ℓ_1 1, Mix 2, ℓ_2 3, ground truth 4, JPEG 5.

Fig. 2. JPEG deblocking – Mix
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 fails at removing both the artifacts in the sky and the halo introduced by the JPEG compression at the edge of the building. ℓ_1 does better than ℓ_2, but only Mix succeeds at completely removing the halo and attenuating the artifacts in the sky.

ℓ_1 1, Mix 2, ℓ_2 3, ground truth 4, JPEG 5.

Fig. 3. JPEG deblocking – ℓ_2
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 fails at removing both the artifacts in the sky and the halo introduced by the JPEG compression at the edge of the building. ℓ_1 does better than ℓ_2, but only Mix succeeds at completely removing the halo and attenuating the artifacts in the sky.

ℓ_1 1, Mix 2, ℓ_2 3, ground truth 4, JPEG 5.

Fig. 4. JPEG deblocking – Ground truth
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 fails at removing both the artifacts in the sky and the halo introduced by the JPEG compression at the edge of the building. ℓ_1 does better than ℓ_2, but only Mix succeeds at completely removing the halo and attenuating the artifacts in the sky.

ℓ_1 1, Mix 2, ℓ_2 3, ground truth 4, JPEG 5.

Fig. 5. JPEG deblocking – JPEG
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 fails at removing both the artifacts in the sky and the halo introduced by the JPEG compression at the edge of the building. ℓ_1 does better than ℓ_2, but only Mix succeeds at completely removing the halo and attenuating the artifacts in the sky.

ℓ_1 1, Mix 2, ℓ_2 3, ground truth 4, JPEG 5.

Fig. 6. JPEG deblocking – ℓ_1
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that Mix produces a result that is sharper than the one produced by ℓ_1, while better attenuating the halos.

ℓ_1 6, Mix 7, ℓ_2 8, ground truth 9, JPEG 10.

Fig. 7. JPEG deblocking – Mix
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that Mix produces a result that is sharper than the one produced by ℓ_1, while better attenuating the halos.

ℓ_1 6, Mix 7, ℓ_2 8, ground truth 9, JPEG 10.

Fig. 8. JPEG deblocking – ℓ_2
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that Mix produces a result that is sharper than the one produced by ℓ_1, while better attenuating the halos.

ℓ_1 6, Mix 7, ℓ_2 8, ground truth 9, JPEG 10.

Fig. 9. JPEG deblocking – Ground truth
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that Mix produces a result that is sharper than the one produced by ℓ_1, while better attenuating the halos.

ℓ_1 6, Mix 7, ℓ_2 8, ground truth 9, JPEG 10.

Fig. 10. JPEG deblocking – JPEG
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that Mix produces a result that is sharper than the one produced by ℓ_1, while better attenuating the halos.

ℓ_1 6, Mix 7, ℓ_2 8, ground truth 9, JPEG 10.

Fig. 11. JPEG deblocking – ℓ_1
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Mix outperforms both ℓ_1 and ℓ_2 at removing the artifacts (left patch) and the halos (right patch).

ℓ_1 11, Mix 12, ℓ_2 13, ground truth 14, JPEG 15.

Fig. 12. JPEG deblocking – Mix
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Mix outperforms both ℓ_1 and ℓ_2 at removing the artifacts (left patch) and the halos (right patch).

ℓ_1 11, Mix 12, ℓ_2 13, ground truth 14, JPEG 15.

Fig. 13. JPEG deblocking – ℓ_2
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Mix outperforms both ℓ_1 and ℓ_2 at removing the artifacts (left patch) and the halos (right patch).

ℓ_1 11, Mix 12, ℓ_2 13, ground truth 14, JPEG 15.

Fig. 14. JPEG deblocking – Ground truth
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Mix outperforms both ℓ_1 and ℓ_2 at removing the artifacts (left patch) and the halos (right patch).

ℓ_1 11, Mix 12, ℓ_2 13, ground truth 14, JPEG 15.

Fig. 15. JPEG deblocking – JPEG
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Mix outperforms both ℓ_1 and ℓ_2 at removing the artifacts (left patch) and the halos (right patch).

ℓ_1 11, Mix 12, ℓ_2 13, ground truth 14, JPEG 15.

Joint denoising + demosaicking

Fig. 16. Joint Denoising+Demosaicking – BM3D
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Observe the loss of details for BM3D compared with Mix in the patches on the leaves. Also note the zippering artifacts due to BM3D+demosaicking on the flower, and how they are solved by Mix.

BM3D 16, ground truth 17, ℓ_1 18, Mix 19, ℓ_2 20, MS-SSIM 21, SSIM_5 22, SSIM_9 23, noisy 24.

Fig. 17. Joint Denoising+Demosaicking – Ground Truth
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Observe the loss of details for BM3D compared with Mix in the patches on the leaves. Also note the zippering artifacts due to BM3D+demosaicking on the flower, and how they are solved by Mix.

BM3D 16, ground truth 17, ℓ_1 18, Mix 19, ℓ_2 20, MS-SSIM 21, SSIM_5 22, SSIM_9 23, noisy 24.

Fig. 18. Joint Denoising+Demosaicking – ℓ_1
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Observe the loss of details for BM3D compared with Mix in the patches on the leaves. Also note the zippering artifacts due to BM3D+demosaicking on the flower, and how they are solved by Mix.

BM3D 16, ground truth 17, ℓ_1 18, Mix 19, ℓ_2 20, MS-SSIM 21, SSIM$_5$ 22, SSIM$_9$ 23, noisy 24.

Fig. 19. Joint Denoising+Demosaicking – Mix
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Observe the loss of details for BM3D compared with Mix in the patches on the leaves. Also note the zippering artifacts due to BM3D+demosaicking on the flower, and how they are solved by Mix.

BM3D 16, ground truth 17, ℓ_1 18, Mix 19, ℓ_2 20, MS-SSIM 21, SSIM$_5$ 22, SSIM$_9$ 23, noisy 24.

Fig. 20. Joint Denoising+Demosaicking – ℓ_2
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Observe the loss of details for BM3D compared with Mix in the patches on the leaves. Also note the zippering artifacts due to BM3D+demosaicking on the flower, and how they are solved by Mix.

BM3D 16, ground truth 17, ℓ_1 18, Mix 19, ℓ_2 20, MS-SSIM 21, SSIM_5 22, SSIM_9 23, noisy 24.

Fig. 21. Joint Denoising+Demosaicking – MS-SSIM
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Observe the loss of details for BM3D compared with Mix in the patches on the leaves. Also note the zippering artifacts due to BM3D+demosaicking on the flower, and how they are solved by Mix.

BM3D 16, ground truth 17, ℓ_1 18, Mix 19, ℓ_2 20, MS-SSIM 21, SSIM_5 22, SSIM_9 23, noisy 24.

Fig. 22. Joint Denoising+Demosaicking – SSIM$_5$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Observe the loss of details for BM3D compared with Mix in the patches on the leaves. Also note the zippering artifacts due to BM3D+demosaicking on the flower, and how they are solved by Mix.

BM3D 16, ground truth 17, ℓ_1 18, Mix 19, ℓ_2 20, MS-SSIM 21, SSIM$_5$ 22, SSIM$_9$ 23, noisy 24.

Fig. 23. Joint Denoising+Demosaicking – SSIM$_9$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Observe the loss of details for BM3D compared with Mix in the patches on the leaves. Also note the zippering artifacts due to BM3D+demosaicking on the flower, and how they are solved by Mix.

BM3D 16, ground truth 17, ℓ_1 18, Mix 19, ℓ_2 20, MS-SSIM 21, SSIM$_5$ 22, SSIM$_9$ 23, noisy 24.

Fig. 24. Joint Denoising+Demosaicking – Noisy
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Observe the loss of details for BM3D compared with Mix in the patches on the leaves. Also note the zippering artifacts due to BM3D+demosaicking on the flower, and how they are solved by Mix.

BM3D 16, ground truth 17, ℓ_1 18, Mix 19, ℓ_2 20, MS-SSIM 21, SSIM_5 22, SSIM_9 23, noisy 24.

Fig. 25. Joint Denoising+Demosaicking – BM3D
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 produces splotchy artifacts in the sky patch. MS-SSIM removes them, but changes the color of the sky. Mix (ℓ_1 combined with MS-SSIM) achieves the desired result. Also, focus on the patch from the top of the building to notice how SSIM_5 and SSIM_9 produce an increasingly large "halo" of noise around the edge, which is removed by MS-SSIM (See main paper for the explanation).

BM3D 25, ground truth 26, ℓ_1 27, Mix 28, ℓ_2 29, MS-SSIM 30, SSIM_5 31, SSIM_9 32, noisy 33.

Fig. 26. Joint Denoising+Demosaicking – Ground Truth
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 produces splotchy artifacts in the sky patch. MS-SSIM removes them, but changes the color of the sky. Mix (ℓ_1 combined with MS-SSIM) achieves the desired result. Also, focus on the patch from the top of the building to notice how SSIM_5 and SSIM_9 produce an increasingly large "halo" of noise around the edge, which is removed by MS-SSIM (See main paper for the explanation).

BM3D 25, ground truth 26, ℓ_1 27, Mix 28, ℓ_2 29, MS-SSIM 30, SSIM_5 31, SSIM_9 32, noisy 33.

Fig. 27. Joint Denoising+Demosaicking – ℓ_1
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 produces splotchy artifacts in the sky patch. MS-SSIM removes them, but changes the color of the sky. Mix (ℓ_1 combined with MS-SSIM) achieves the desired result. Also, focus on the patch from the top of the building to notice how SSIM$_5$ and SSIM$_9$ produce an increasingly large "halo" of noise around the edge, which is removed by MS-SSIM (See main paper for the explanation).

BM3D 25, ground truth 26, ℓ_1 27, Mix 28, ℓ_2 29, MS-SSIM 30, SSIM$_5$ 31, SSIM$_9$ 32, noisy 33.

Fig. 28. Joint Denoising+Demosaicking – **Mix**
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 produces splotchy artifacts in the sky patch. MS-SSIM removes them, but changes the color of the sky. Mix (ℓ_1 combined with MS-SSIM) achieves the desired result. Also, focus on the patch from the top of the building to notice how SSIM$_5$ and SSIM$_9$ produce an increasingly large "halo" of noise around the edge, which is removed by MS-SSIM (See main paper for the explanation).

BM3D 25, ground truth 26, ℓ_1 27, Mix 28, ℓ_2 29, MS-SSIM 30, SSIM$_5$ 31, SSIM$_9$ 32, noisy 33.

Fig. 29. Joint Denoising+Demosaicking – ℓ_2
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 produces splotchy artifacts in the sky patch. MS-SSIM removes them, but changes the color of the sky. Mix (ℓ_1 combined with MS-SSIM) achieves the desired result. Also, focus on the patch from the top of the building to notice how $SSIM_5$ and $SSIM_9$ produce an increasingly large "halo" of noise around the edge, which is removed by MS-SSIM (See main paper for the explanation).

BM3D 25, ground truth 26, ℓ_1 27, Mix 28, ℓ_2 29, MS-SSIM 30, $SSIM_5$ 31, $SSIM_9$ 32, noisy 33.

Fig. 30. Joint Denoising+Demosaicking – MS-SSIM
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 produces splotchy artifacts in the sky patch. MS-SSIM removes them, but changes the color of the sky. Mix (ℓ_1 combined with MS-SSIM) achieves the desired result. Also, focus on the patch from the top of the building to notice how SSIM_5 and SSIM_9 produce an increasingly large "halo" of noise around the edge, which is removed by MS-SSIM (See main paper for the explanation).

BM3D 25, ground truth 26, ℓ_1 27, Mix 28, ℓ_2 29, MS-SSIM 30, SSIM_5 31, SSIM_9 32, noisy 33.

Fig. 31. Joint Denoising+Demosaicking – $\mathbf{SSIM_5}$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 produces splotchy artifacts in the sky patch. MS-SSIM removes them, but changes the color of the sky. Mix (ℓ_1 combined with MS-SSIM) achieves the desired result. Also, focus on the patch from the top of the building to notice how $SSIM_5$ and $SSIM_9$ produce an increasingly large "halo" of noise around the edge, which is removed by MS-SSIM (See main paper for the explanation).

BM3D 25, ground truth 26, ℓ_1 27, Mix 28, ℓ_2 29, MS-SSIM 30, $SSIM_5$ 31, $SSIM_9$ 32, noisy 33.

Fig. 32. Joint Denoising+Demosaicking – SSIM$_9$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 produces splotchy artifacts in the sky patch. MS-SSIM removes them, but changes the color of the sky. Mix (ℓ_1 combined with MS-SSIM) achieves the desired result. Also, focus on the patch from the top of the building to notice how SSIM$_5$ and SSIM$_9$ produce an increasingly large "halo" of noise around the edge, which is removed by MS-SSIM (See main paper for the explanation).

BM3D 25, ground truth 26, ℓ_1 27, Mix 28, ℓ_2 29, MS-SSIM 30, SSIM$_5$ 31, SSIM$_9$ 32, noisy 33.

Fig. 33. Joint Denoising+Demosaicking – Noisy
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note that ℓ_2 produces splotchy artifacts in the sky patch. MS-SSIM removes them, but changes the color of the sky. Mix (ℓ_1 combined with MS-SSIM) achieves the desired result. Also, focus on the patch from the top of the building to notice how SSIM_5 and SSIM_9 produce an increasingly large "halo" of noise around the edge, which is removed by MS-SSIM (See main paper for the explanation).

BM3D 25, ground truth 26, ℓ_1 27, Mix 28, ℓ_2 29, MS-SSIM 30, SSIM_5 31, SSIM_9 32, noisy 33.

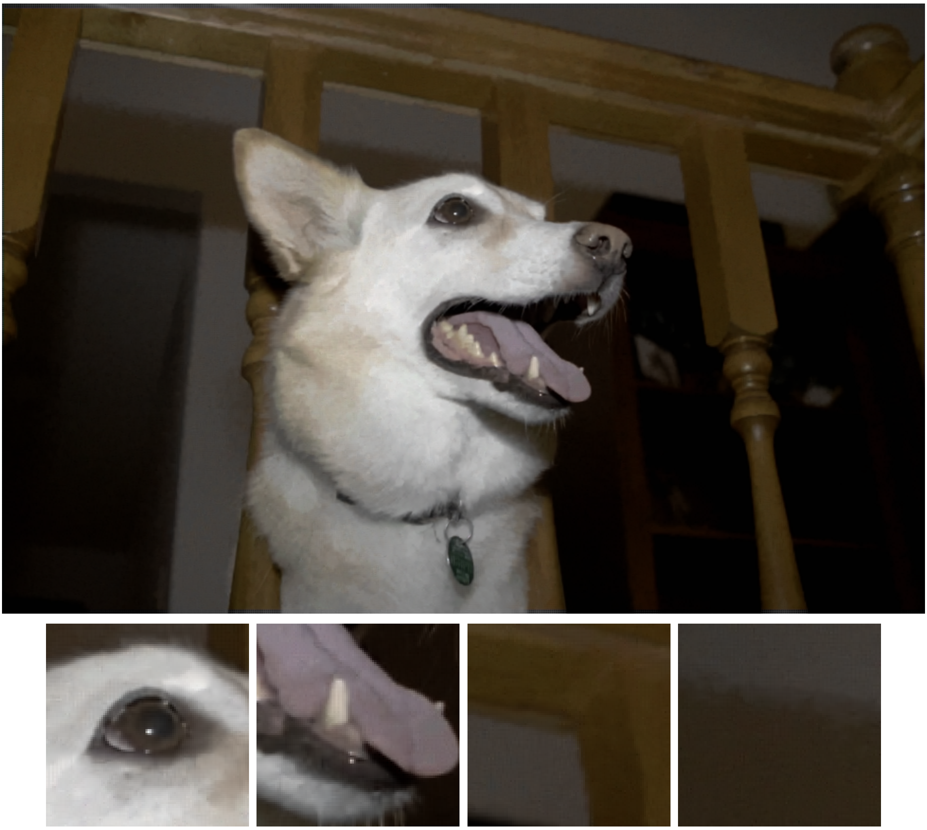

Fig. 34. Joint Denoising+Demosaicking – BM3D
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Please compare the sharpness of BM3D and Mix on the eye and mouth of the dog. Also, compare the details on the tongue between ℓ_1 and Mix, particularly the crease on the tongue.

BM3D 34, clean 35, ℓ_1 36, Mix 37, ℓ_2 38, MS-SSIM 39, $SSIM_5$ 40, $SSIM_9$ 41, noisy 42.

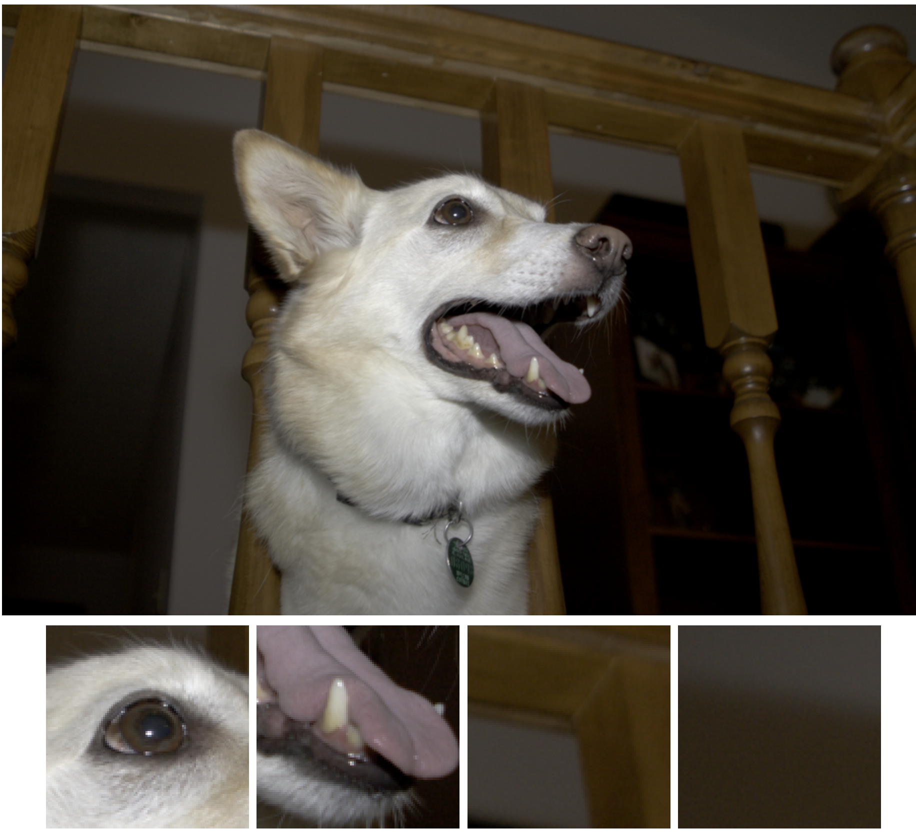

Fig. 35. Joint Denoising+Demosaicking – Ground Truth
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Please compare the sharpness of BM3D and Mix on the eye and mouth of the dog. Also, compare the details on the tongue between ℓ_1 and Mix, particularly the crease on the tongue.

BM3D 34, clean 35, ℓ_1 36, Mix 37, ℓ_2 38, MS-SSIM 39, SSIM$_5$ 40, SSIM$_9$ 41, noisy 42.

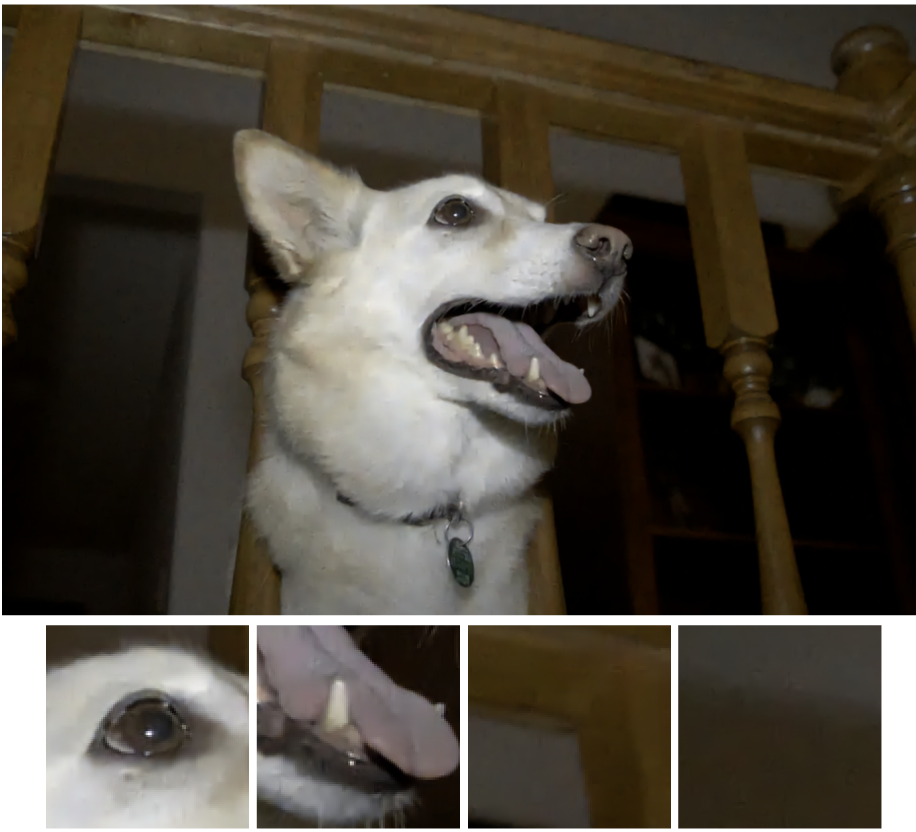

Fig. 36. Joint Denoising+Demosaicking – ℓ_1
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Please compare the sharpness of BM3D and Mix on the eye and mouth of the dog. Also, compare the details on the tongue between ℓ_1 and Mix, particularly the crease on the tongue.

BM3D 34, clean 35, ℓ_1 36, Mix 37, ℓ_2 38, MS-SSIM 39, SSIM$_5$ 40, SSIM$_9$ 41, noisy 42.

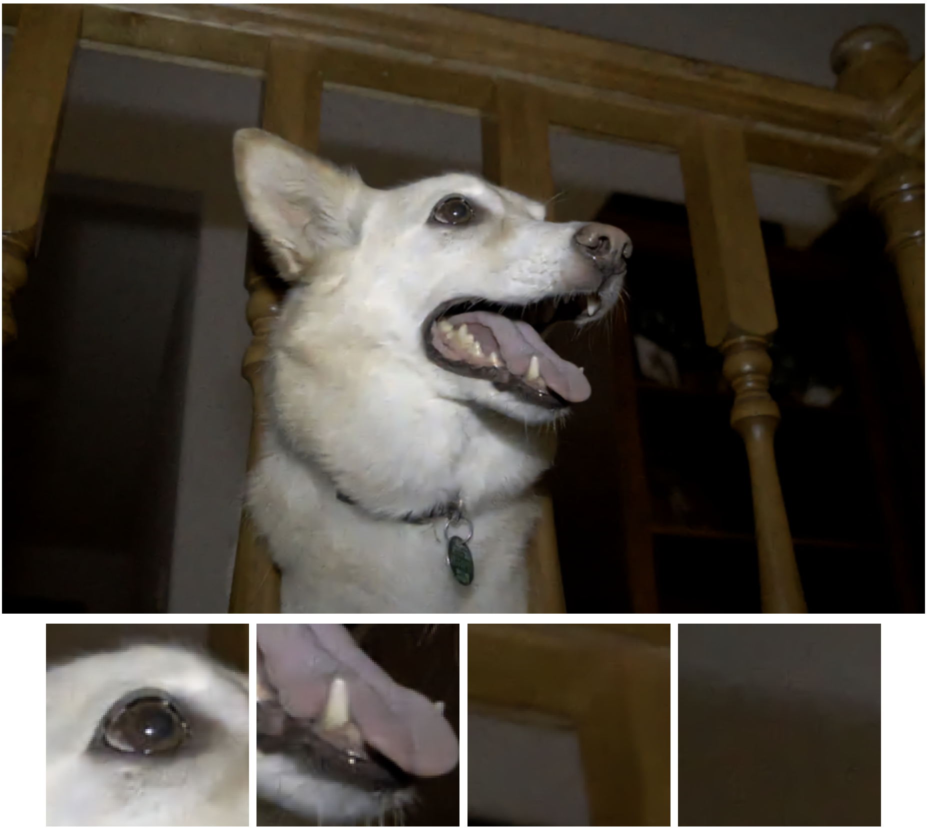

Fig. 37. Joint Denoising+Demosaicking – Mix
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Please compare the sharpness of BM3D and Mix on the eye and mouth of the dog. Also, compare the details on the tongue between ℓ_1 and Mix, particularly the crease on the tongue.

BM3D 34, clean 35, ℓ_1 36, Mix 37, ℓ_2 38, MS-SSIM 39, SSIM$_5$ 40, SSIM$_9$ 41, noisy 42.

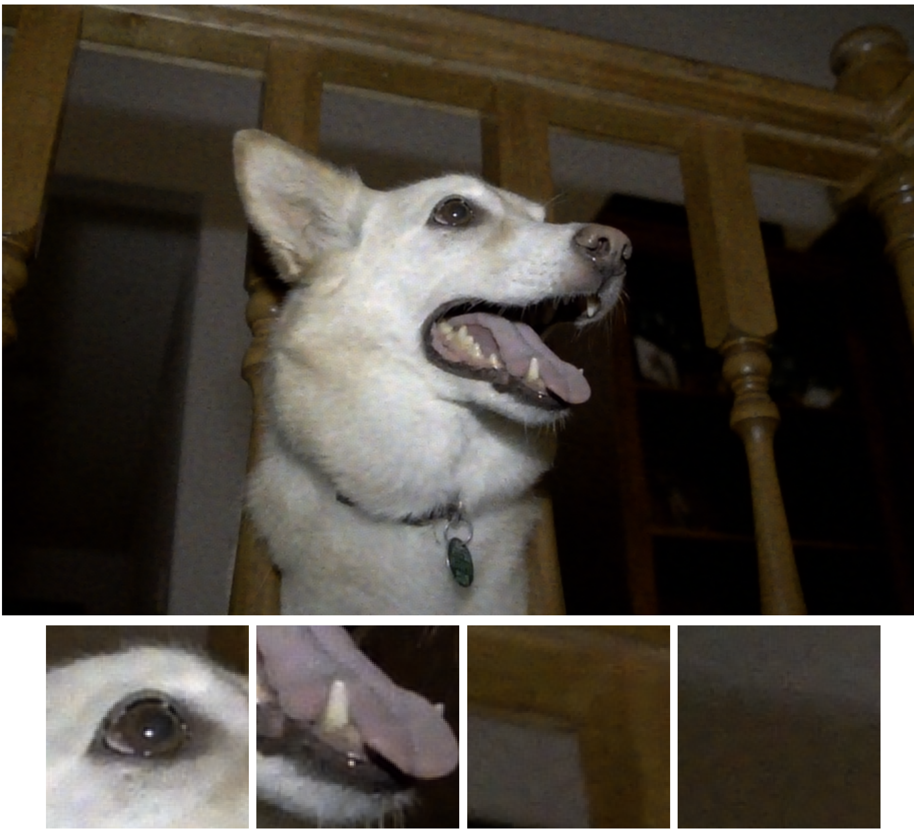

Fig. 38. Joint Denoising+Demosaicking – ℓ_2
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Please compare the sharpness of BM3D and Mix on the eye and mouth of the dog. Also, compare the details on the tongue between ℓ_1 and Mix, particularly the crease on the tongue.

BM3D 34, clean 35, ℓ_1 36, Mix 37, ℓ_2 38, MS-SSIM 39, SSIM$_5$ 40, SSIM$_9$ 41, noisy 42.

Fig. 39. Joint Denoising+Demosaicking – MS-SSIM
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Please compare the sharpness of BM3D and Mix on the eye and mouth of the dog. Also, compare the details on the tongue between ℓ_1 and Mix, particularly the crease on the tongue.

BM3D 34, clean 35, ℓ_1 36, Mix 37, ℓ_2 38, MS-SSIM 39, SSIM$_5$ 40, SSIM$_9$ 41, noisy 42.

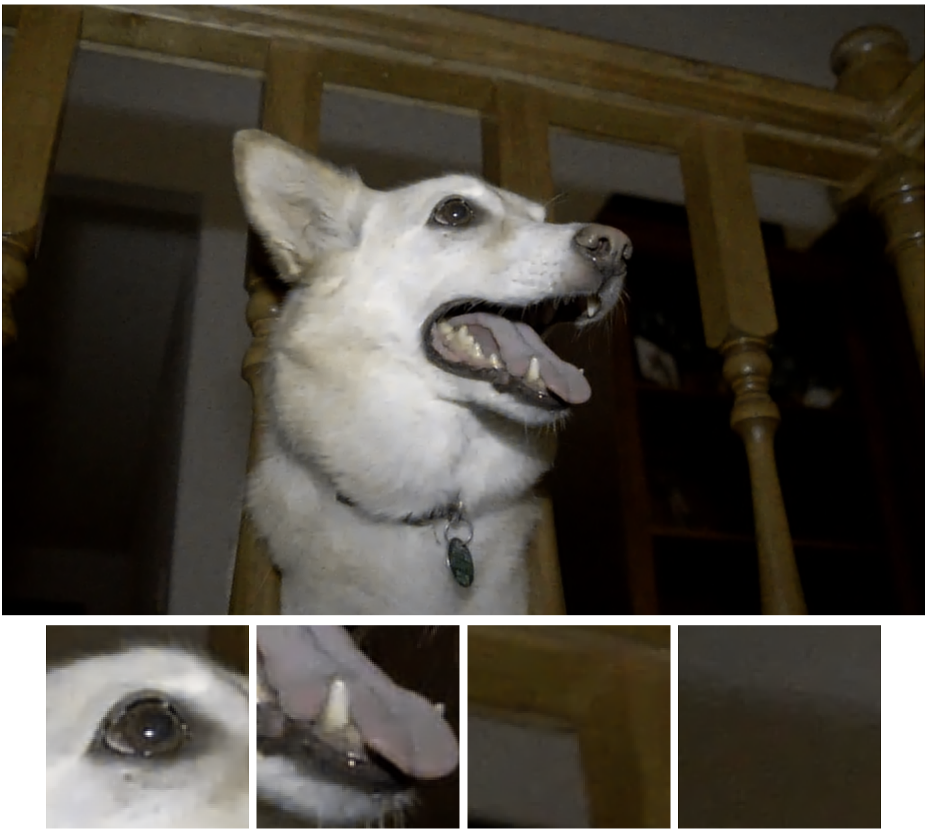

Fig. 40. Joint Denoising+Demosaicking – SSIM$_5$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Please compare the sharpness of BM3D and Mix on the eye and mouth of the dog. Also, compare the details on the tongue between ℓ_1 and Mix, particularly the crease on the tongue.

BM3D 34, clean 35, ℓ_1 36, Mix 37, ℓ_2 38, MS-SSIM 39, SSIM$_5$ 40, SSIM$_9$ 41, noisy 42.

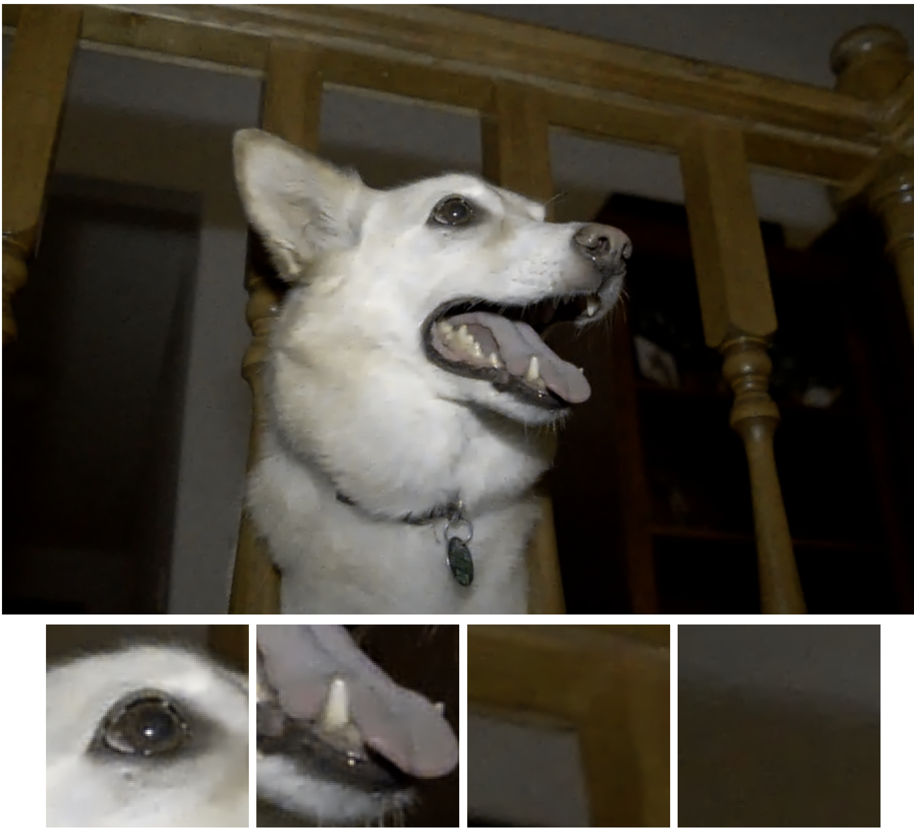

Fig. 41. Joint Denoising+Demosaicking – SSIM$_9$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Please compare the sharpness of BM3D and Mix on the eye and mouth of the dog. Also, compare the details on the tongue between ℓ_1 and Mix, particularly the crease on the tongue.

BM3D 34, clean 35, ℓ_1 36, Mix 37, ℓ_2 38, MS-SSIM 39, SSIM$_5$ 40, SSIM$_9$ 41, noisy 42.

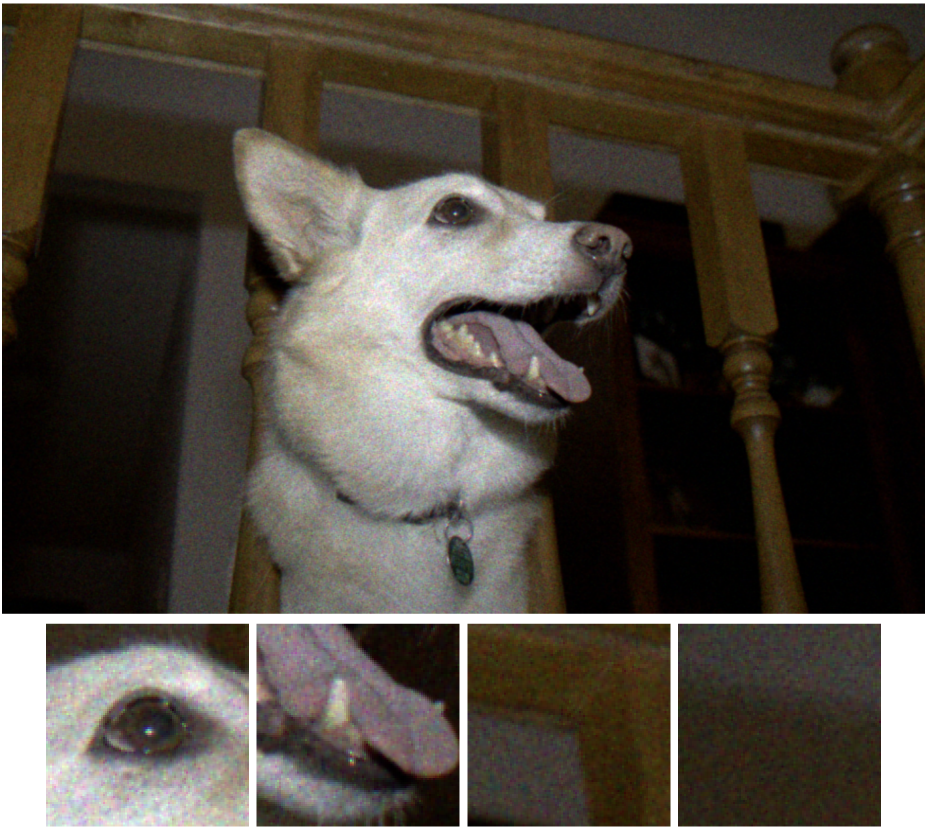

Fig. 42. Joint Denoising+Demosaicking – Noisy
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Please compare the sharpness of BM3D and Mix on the eye and mouth of the dog. Also, compare the details on the tongue between ℓ_1 and Mix, particularly the crease on the tongue.

BM3D 34, clean 35, ℓ_1 36, Mix 37, ℓ_2 38, MS-SSIM 39, SSIM$_5$ 40, SSIM$_9$ 41, noisy 42.

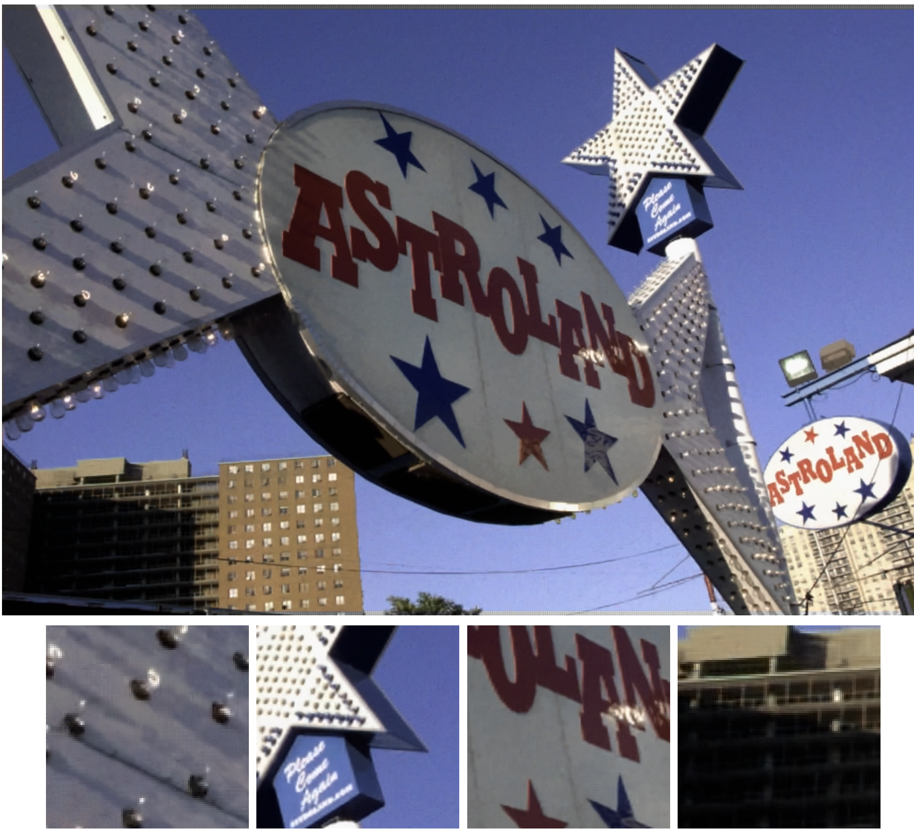

Fig. 43. Joint Denoising+Demosaicking – BM3D
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note how Mix preserves the subtle structure on the power line better than ℓ_1. Also, switch between SSIM_5, SSIM_9, and MS-SSIM to appreciate the noise halos around the star (see also paper).

BM3D 43, ground truth 44, ℓ_1 45, Mix 46, ℓ_2 47, MS-SSIM 48, SSIM_5 49, SSIM_9 50, noisy 51.

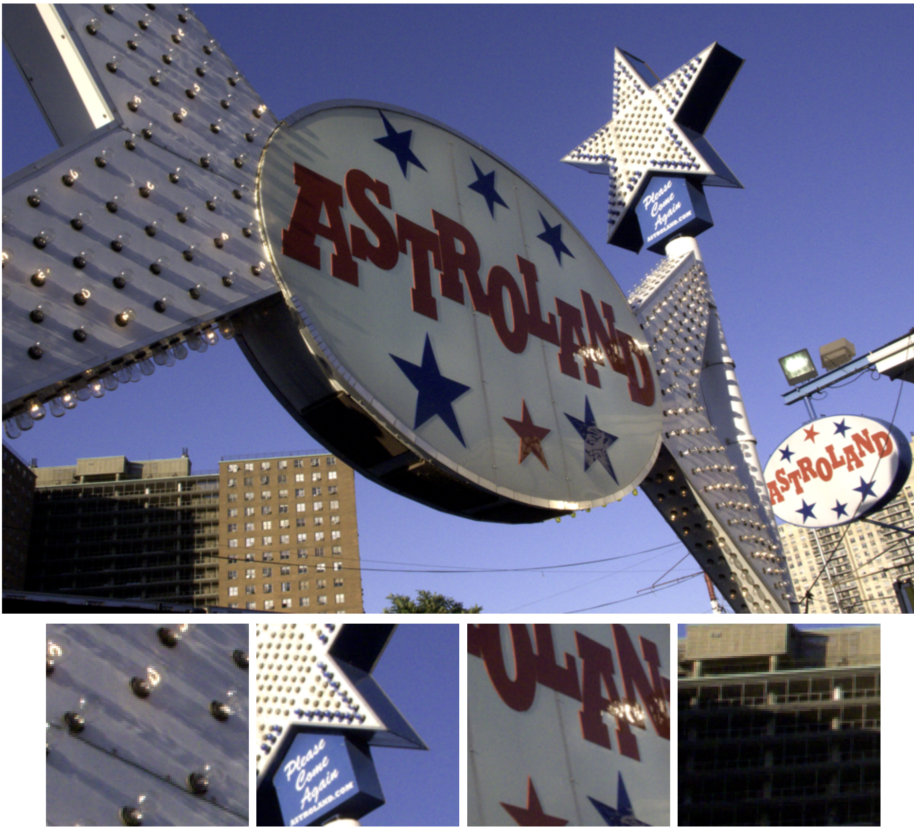

Fig. 44. Joint Denoising+Demosaicking – Ground Truth
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note how Mix preserves the subtle structure on the power line better than ℓ_1. Also, switch between SSIM_5, SSIM_9, and MS-SSIM to appreciate the noise halos around the star (see also paper).

BM3D 43, ground truth 44, ℓ_1 45, Mix 46, ℓ_2 47, MS-SSIM 48, SSIM_5 49, SSIM_9 50, noisy 51.

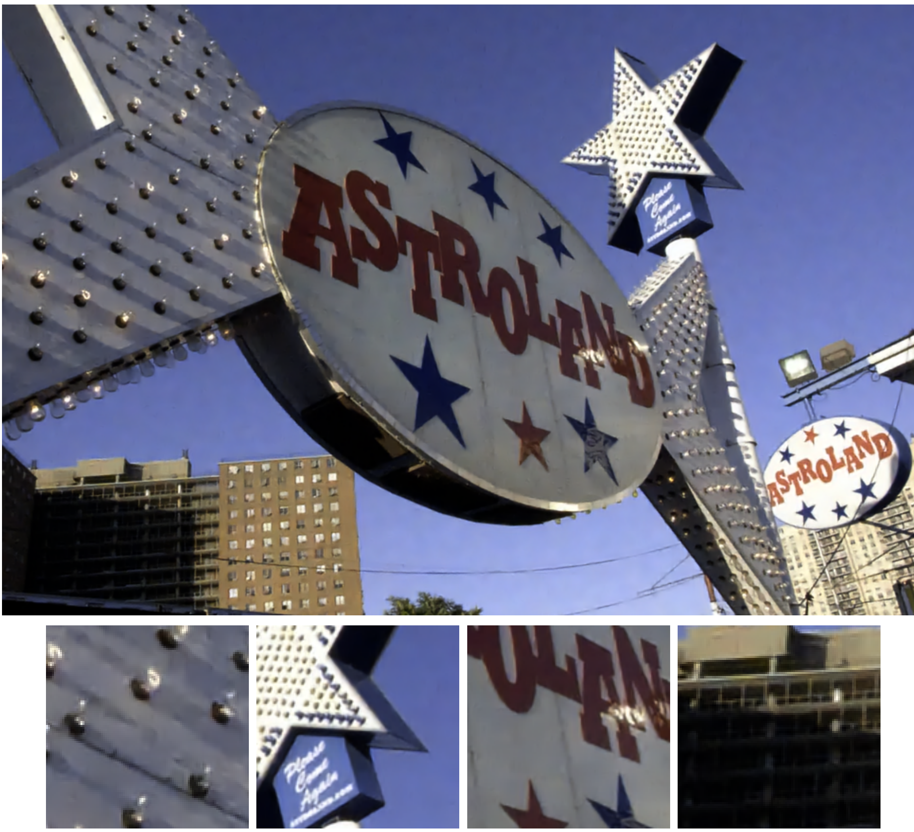

Fig. 45. Joint Denoising+Demosaicking – ℓ_1
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note how Mix preserves the subtle structure on the power line better than ℓ_1. Also, switch between SSIM_5, SSIM_9, and MS-SSIM to appreciate the noise halos around the star (see also paper).

BM3D 43, ground truth 44, ℓ_1 45, Mix 46, ℓ_2 47, MS-SSIM 48, SSIM_5 49, SSIM_9 50, noisy 51.

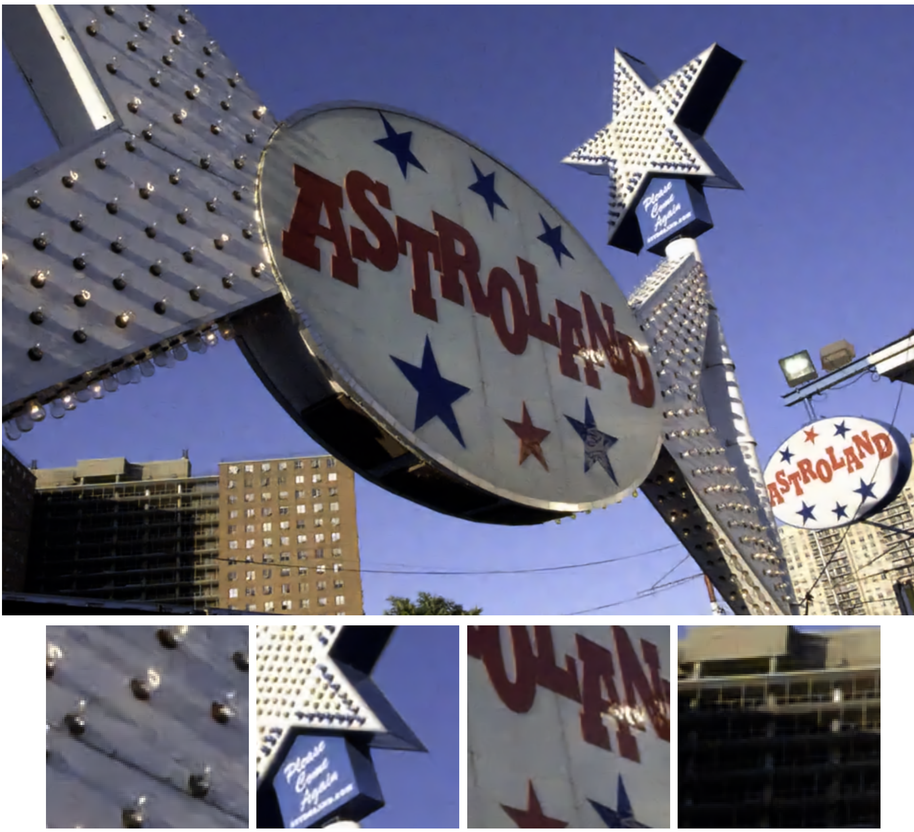

Fig. 46. Joint Denoising+Demosaicking – Mix
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note how Mix preserves the subtle structure on the power line better than ℓ_1. Also, switch between SSIM_5, SSIM_9, and MS-SSIM to appreciate the noise halos around the star (see also paper).

BM3D 43, ground truth 44, ℓ_1 45, Mix 46, ℓ_2 47, MS-SSIM 48, SSIM_5 49, SSIM_9 50, noisy 51.

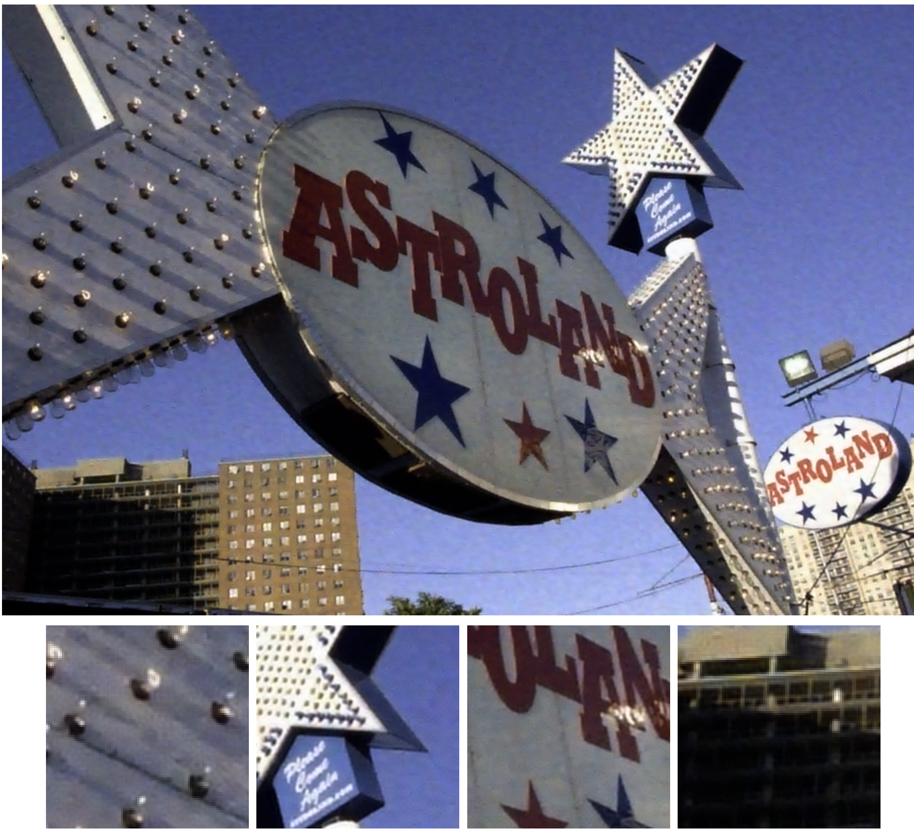

Fig. 47. Joint Denoising+Demosaicking – ℓ_2
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note how Mix preserves the subtle structure on the power line better than ℓ_1. Also, switch between SSIM_5, SSIM_9, and MS-SSIM to appreciate the noise halos around the star (see also paper).

BM3D 43, ground truth 44, ℓ_1 45, Mix 46, ℓ_2 47, MS-SSIM 48, SSIM_5 49, SSIM_9 50, noisy 51.

Fig. 48. Joint Denoising+Demosaicking – MS-SSIM
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note how Mix preserves the subtle structure on the power line better than ℓ_1. Also, switch between SSIM_5, SSIM_9, and MS-SSIM to appreciate the noise halos around the star (see also paper).

BM3D 43, ground truth 44, ℓ_1 45, Mix 46, ℓ_2 47, MS-SSIM 48, SSIM_5 49, SSIM_9 50, noisy 51.

Fig. 49. Joint Denoising+Demosaicking – SSIM$_5$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note how Mix preserves the subtle structure on the power line better than ℓ_1. Also, switch between SSIM$_5$, SSIM$_9$, and MS-SSIM to appreciate the noise halos around the star (see also paper).

BM3D 43, ground truth 44, ℓ_1 45, Mix 46, ℓ_2 47, MS-SSIM 48, SSIM$_5$ 49, SSIM$_9$ 50, noisy 51.

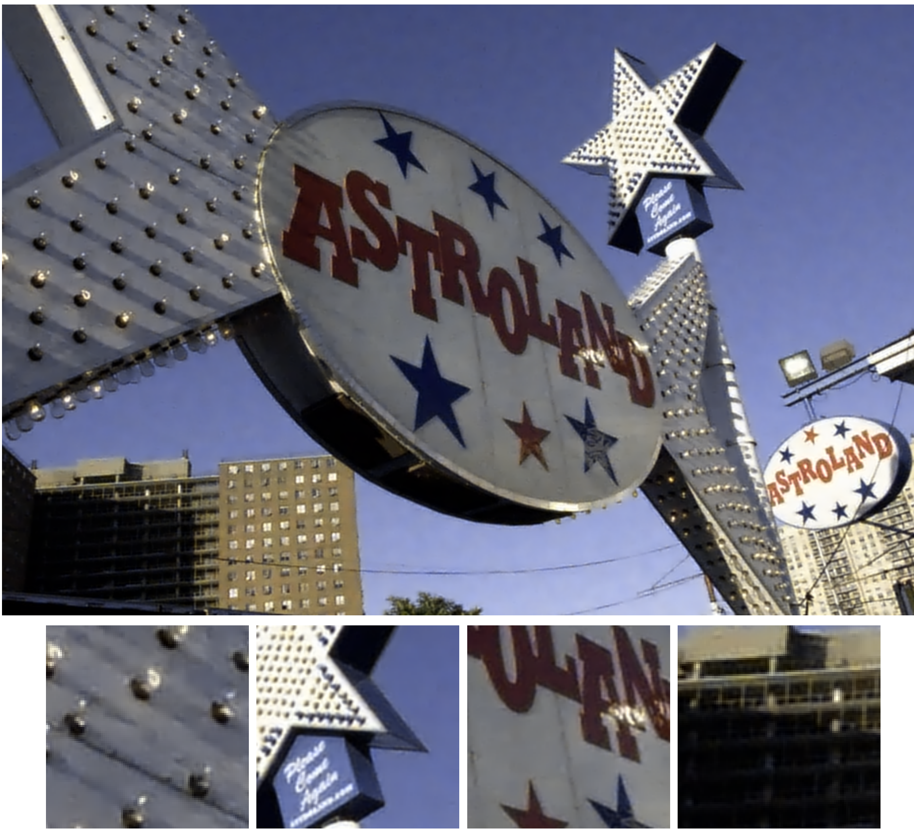

Fig. 50. Joint Denoising+Demosaicking – SSIM$_9$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note how Mix preserves the subtle structure on the power line better than ℓ_1. Also, switch between SSIM$_5$, SSIM$_9$, and MS-SSIM to appreciate the noise halos around the star (see also paper).

BM3D 43, ground truth 44, ℓ_1 45, Mix 46, ℓ_2 47, MS-SSIM 48, SSIM$_5$ 49, SSIM$_9$ 50, noisy 51.

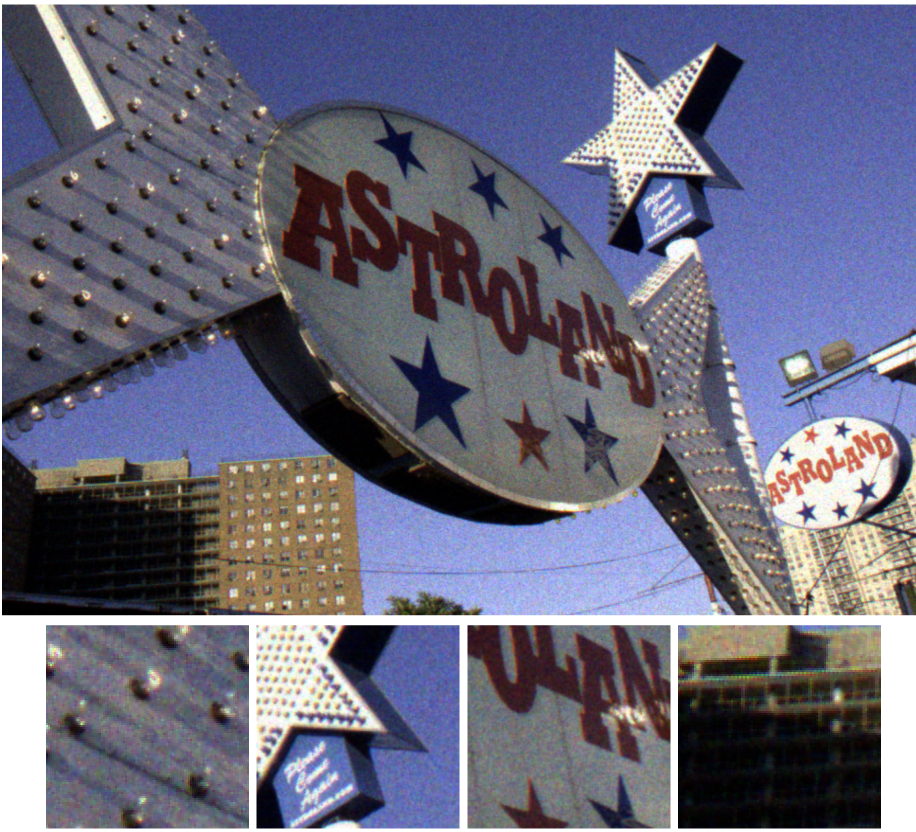

Fig. 51. Joint Denoising+Demosaicking – Noisy
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note how Mix preserves the subtle structure on the power line better than ℓ_1. Also, switch between SSIM_5, SSIM_9, and MS-SSIM to appreciate the noise halos around the star (see also paper).

BM3D 43, ground truth 44, ℓ_1 45, Mix 46, ℓ_2 47, MS-SSIM 48, SSIM_5 49, SSIM_9 50, noisy 51.

Comparisons of different training schedules

Fig. 52. Comparison of different training schedules on denoising+demosaicking $- \ell_1$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Here we compare the output of the networks trained alternating ℓ_1 and ℓ_2. In the paper we show that a network trained with ℓ_2 gets stuck in a local minimum. When ℓ_2 is used to train a network that was pre-trained with ℓ_1, the ℓ_2 loss decreases but the result is not as good as with the other losses in flat regions, such as the sky.

ℓ_1 52, Mix 53, ℓ_2 54, ℓ_1+ℓ_2 55, ℓ_2+ℓ_1 56, ground truth 57, noisy 58.

Fig. 53. Comparison of different training schedules on denoising+demosaicking – Mix
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Here we compare the output of the networks trained alternating ℓ_1 and ℓ_2. In the paper we show that a network trained with ℓ_2 gets stuck in a local minimum. When ℓ_2 is used to train a network that was pre-trained with ℓ_1, the ℓ_2 loss decreases but the result is not as good as with the other losses in flat regions, such as the sky.

ℓ_1 52, Mix 53, ℓ_2 54, ℓ_1+ℓ_2 55, ℓ_2+ℓ_1 56, ground truth 57, noisy 58.

Fig. 54. Comparison of different training schedules on denoising+demosaicking $-\ \ell_2$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Here we compare the output of the networks trained alternating ℓ_1 and ℓ_2. In the paper we show that a network trained with ℓ_2 gets stuck in a local minimum. When ℓ_2 is used to train a network that was pre-trained with ℓ_1, the ℓ_2 loss decreases but the result is not as good as with the other losses in flat regions, such as the sky.

ℓ_1 52, Mix 53, ℓ_2 54, ℓ_1+ℓ_2 55, ℓ_2+ℓ_1 56, ground truth 57, noisy 58.

Fig. 55. Comparison of different training schedules on denoising+demosaicking – $\ell_1 + \ell_2$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Here we compare the output of the networks trained alternating ℓ_1 and ℓ_2. In the paper we show that a network trained with ℓ_2 gets stuck in a local minimum. When ℓ_2 is used to train a network that was pre-trained with ℓ_1, the ℓ_2 loss decreases but the result is not as good as with the other losses in flat regions, such as the sky.

ℓ_1 52, Mix 53, ℓ_2 54, ℓ_1+ℓ_2 55, ℓ_2+ℓ_1 56, ground truth 57, noisy 58.

Fig. 56. Comparison of different training schedules on denoising+demosaicking $- \ell_2 + \ell_1$
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Here we compare the output of the networks trained alternating ℓ_1 and ℓ_2. In the paper we show that a network trained with ℓ_2 gets stuck in a local minimum. When ℓ_2 is used to train a network that was pre-trained with ℓ_1, the ℓ_2 loss decreases but the result is not as good as with the other losses in flat regions, such as the sky.

ℓ_1 52, Mix 53, ℓ_2 54, $\ell_1+\ell_2$ 55, $\ell_2+\ell_1$ 56, ground truth 57, noisy 58.

Fig. 57. Comparison of different training schedules on denoising+demosaicking – Ground Truth
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Here we compare the output of the networks trained alternating ℓ_1 and ℓ_2. In the paper we show that a network trained with ℓ_2 gets stuck in a local minimum. When ℓ_2 is used to train a network that was pre-trained with ℓ_1, the ℓ_2 loss decreases but the result is not as good as with the other losses in flat regions, such as the sky.

ℓ_1 52, Mix 53, ℓ_2 54, ℓ_1+ℓ_2 55, ℓ_2+ℓ_1 56, ground truth 57, noisy 58.

Fig. 58. Comparison of different training schedules on denoising+demosaicking – Noisy
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Here we compare the output of the networks trained alternating ℓ_1 and ℓ_2. In the paper we show that a network trained with ℓ_2 gets stuck in a local minimum. When ℓ_2 is used to train a network that was pre-trained with ℓ_1, the ℓ_2 loss decreases but the result is not as good as with the other losses in flat regions, such as the sky.

ℓ_1 52, Mix 53, ℓ_2 54, ℓ_1+ℓ_2 55, ℓ_2+ℓ_1 56, ground truth 57, noisy 58.

Super-resolution

Fig. 59. Super-resolution – Ground truth
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note the artifacts on the wing for ℓ_2.

Ground truth 59, LR interpolated 60, ℓ_1 61, Mix 62, ℓ_2 63.

Fig. 60. Super-resolution – Low Resolution interpolated
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note the artifacts on the wing for ℓ_2.

Ground truth 59, LR interpolated 60, ℓ_1 61, Mix 62, ℓ_2 63.

Fig. 61. Super-resolution – ℓ_1
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note the artifacts on the wing for ℓ_2.

Ground truth 59, LR interpolated 60, ℓ_1 61, Mix 62, ℓ_2 63.

Fig. 62. Super-resolution – Mix
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note the artifacts on the wing for ℓ_2.

Ground truth 59, LR interpolated 60, ℓ_1 61, Mix 62, ℓ_2 63.

Fig. 63. Super-resolution – ℓ_2
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Note the artifacts on the wing for ℓ_2.

Ground truth 59, LR interpolated 60, ℓ_1 61, Mix 62, ℓ_2 63.

Fig. 64. Super-resolution – Ground truth
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Both ℓ_1 and Mix produce an image that is sharper than ℓ_2, see for instance the necklace.

Ground truth 64, LR interpolated 65, ℓ_1 66, Mix 67, ℓ_2 68.

Fig. 65. Super-resolution – Low Resolution interpolated
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Both ℓ_1 and Mix produce an image that is sharper than ℓ_2, see for instance the necklace.

Ground truth 64, LR interpolated 65, ℓ_1 66, Mix 67, ℓ_2 68.

Fig. 66. Super-resolution – ℓ_1
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Both ℓ_1 and Mix produce an image that is sharper than ℓ_2, see for instance the necklace.

Ground truth 64, LR interpolated 65, ℓ_1 66, Mix 67, ℓ_2 68.

Fig. 67. Super-resolution – Mix
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Both ℓ_1 and Mix produce an image that is sharper than ℓ_2, see for instance the necklace.

Ground truth 64, LR interpolated 65, ℓ_1 66, Mix 67, ℓ_2 68.

Fig. 68. Super-resolution – ℓ_2
Please switch between the output of the different networks clicking on the links below (Figures numbers).
Both ℓ_1 and Mix produce an image that is sharper than ℓ_2, see for instance the necklace.

Ground truth 64, LR interpolated 65, ℓ_1 66, Mix 67, ℓ_2 68.