

SUPPLEMENTARY MATERIAL

Multi-image blind super-resolution of 3D scenes

Abhijith Punnappurath, T. M. Nimisha, A. N. Rajagopalan, *Senior Member, IEEE*

This supplementary material carries additional results that could not be provided in the main paper due to space constraints. All references correspond to our original manuscript.

S1. EXPERIMENTS

We first perform a synthetic experiment to demonstrate our proposed method's ability to handle even inclined planes and smoothly-varying depth values. Next, we show two more real examples along with comparisons against competing techniques. The first of these two, similar to the real example in Fig. 11 of our main paper, is of a scene that is relatively close to the camera. The second, akin to the experiment in Fig. 12 of the original manuscript, involves objects at greater distances from the camera.

To test the effectiveness of our method in more general scenarios involving gradual depth changes, we set up the synthetic experiment depicted in Fig. S1. The super-resolution factor is two. Here the depth layer in the foreground is gradually varying as can be seen from the ground truth depth map in row three. The depth map estimated using optical flow correctly identifies the background layers but fails to detect the gradual depth variations as well as the depth boundaries. The edges appear smeared in the OF depth map as compared to the depth map estimated by our method. Note that we have discretized the depth range and limited the number of depth layers to a maximum of 50. Even with this quantization, we are able to capture fine variations in depth. The depth boundaries are also accurately recovered using our approach. To check the importance of our alternating minimization scheme, we ignored the depth refinement step, and the result thus produced is also shown in row two of Fig. S1. Observe that there is residual blur in the image since the depth map is incorrect. Zoomed-in regions from the first of the five LR observations, the HR image, the output obtained without refining depth, and the output of our

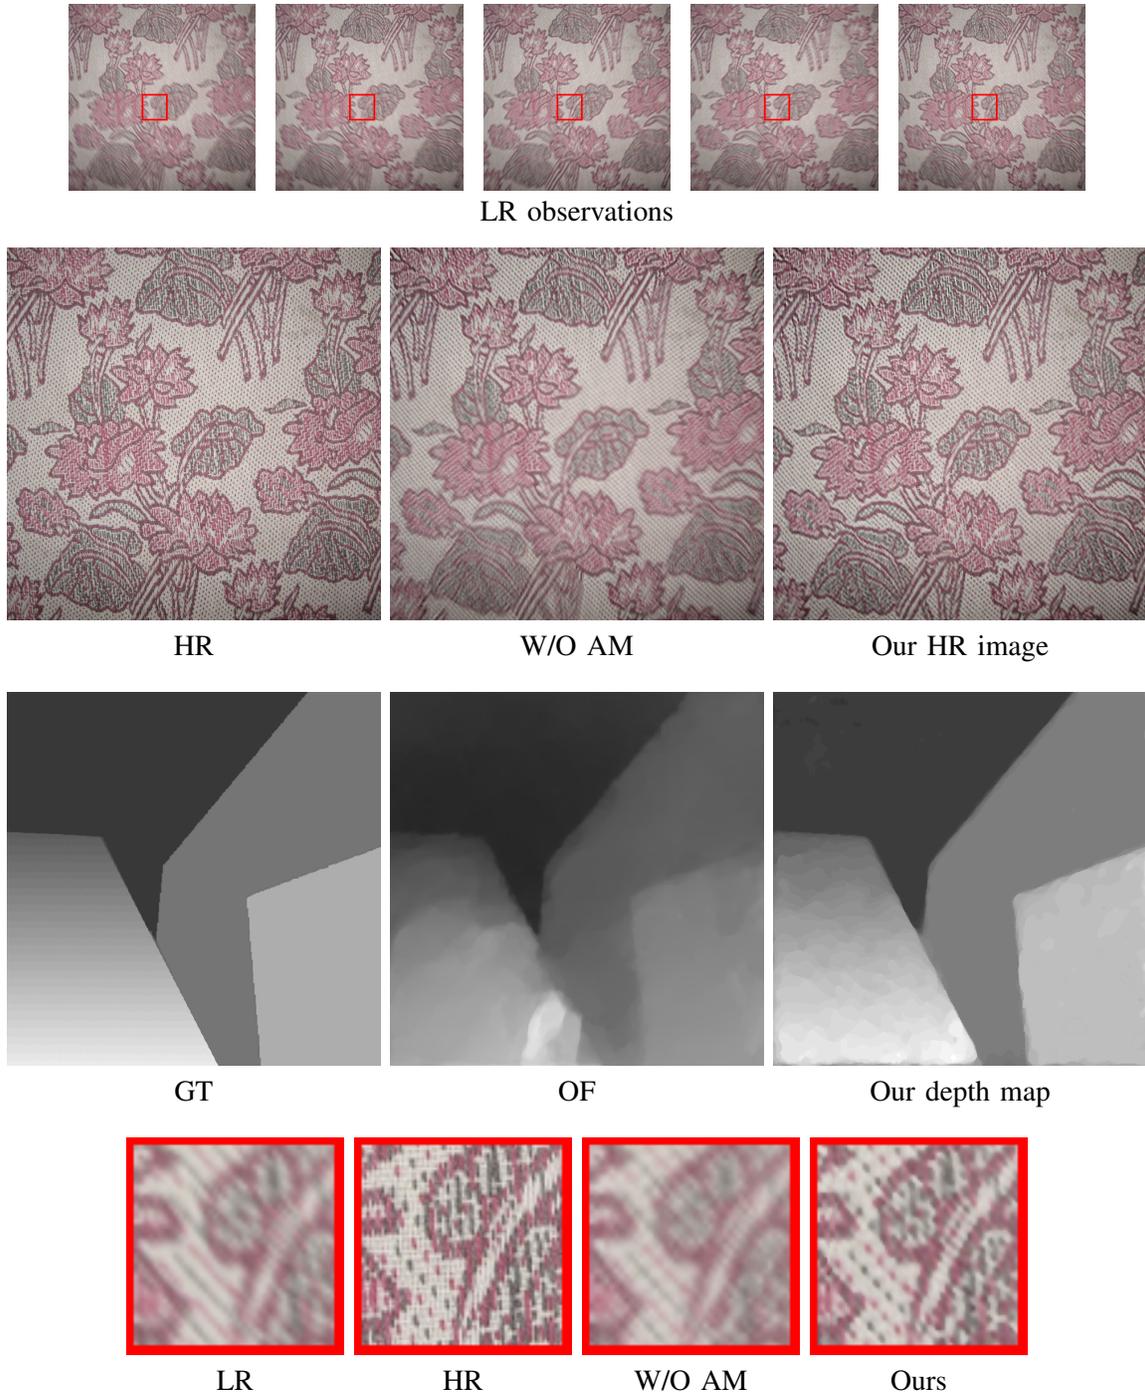


Fig. S1. A synthetic experiment demonstrating our method’s ability to handle gradual depth changes. Row one: blurred LR observations, row two: latent HR image, result obtained without AM, and result of the proposed method, row three: ground truth depth map with gradual depth variations, depth map from optical flow, and our estimated depth map, and row four: zoomed-in regions from the first of the five LR observations in row one, and the three images in row two, respectively.

proposed method are shown in row four. Our output is sharp and compares closely to the original as is apparent from the dots visible in the zoomed-in region.

Yet another real example with objects placed relatively close to the camera is shown in Fig. S2. Books and a textured wallpaper placed at three different depth layers constitute the scene. Zoomed-in regions extracted from the first input LR image and output HR images are also shown. It can be observed that our result is sharper and has fewer deblurring artifacts as compared to other methods. The ‘Nvidia’ logo and the text is clearly discernible. The dots near the stem of the flower are also nicely recovered in our SR output.

For the example in Fig. S3, the zoomed-in regions show that the Roman numerals on the dial of the clock which is in the foreground, and the text in the background are best recovered by our approach. The outputs of both [18] and [5] have deblurring artifacts, while the SR framework of [4] fails to remove the blur on the clock.

We set $\alpha = 1000$ and $\lambda = 0.01$ for all our experiments. The patches extracted from the LR images for HR PSF estimation were of size 81×81 pixels.

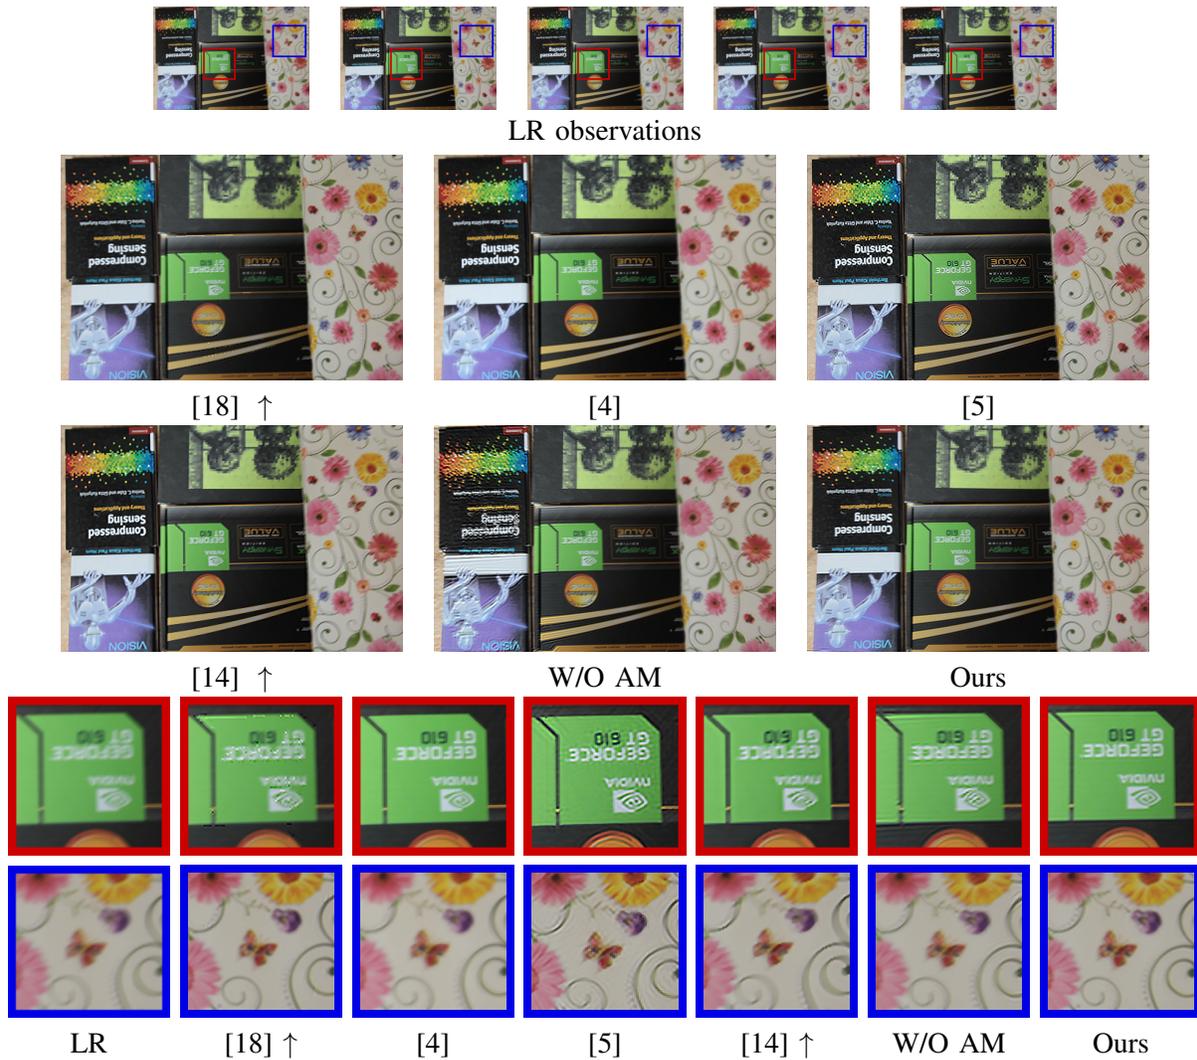


Fig. S2. Row one: blurred LR observations, row two: upsampled output of Paramanand and Rajagopalan [18], results of Sroubek et al. [4] and Ma et al. [5], row three: upsampled output of Hu et al. [14], result obtained without AM, and our HR output image, and row four: zoomed-in regions from the first LR observation in row one, [18] ↑, [4], [5] in row two, [14] ↑, without AM and our HR output in row three, respectively.

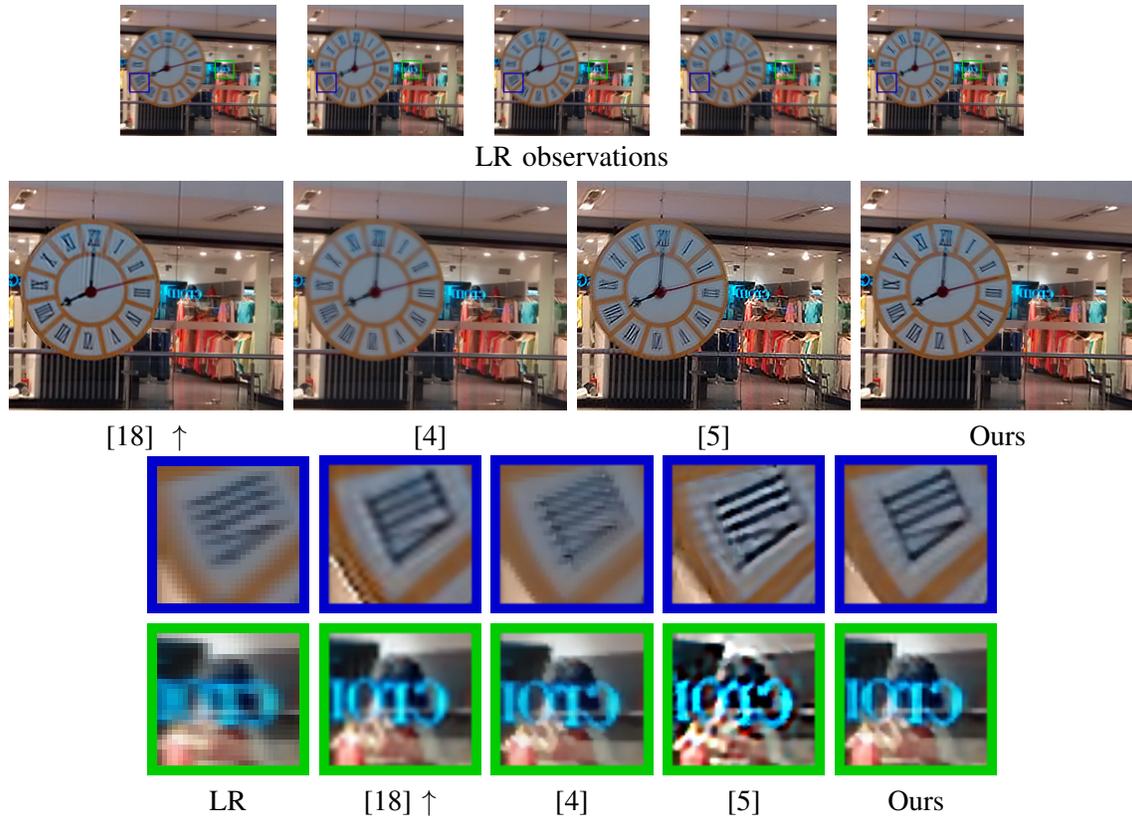


Fig. S3. Row one: blurred LR observations, row two: upsampled output of Paramanand and Rajagopalan [18], results of Sroubek et al. [4] and Ma et al. [5], and our HR output image, and row three: zoomed-in regions from the first LR observation in row one, [18] ↑, [4], [5], and our HR output in row two, respectively.