Pupil-aware Holography Supplementary Information

PRANEETH CHAKRAVARTHULA, Princeton University, USA
SEUNG-HWAN BAEK, Princeton University, USA
FLORIAN SCHIFFERS, Reality Labs Research, Meta, USA
ETHAN TSENG, Princeton University, USA
GRACE KUO, ANDREW MAIMONE, NATHAN MATSUDA, OLIVER COSSAIRT, and DOUGLAS LANMAN, Reality Labs Research, Meta, USA
FELIX HEIDE, Princeton University, USA

CCS Concepts: • **Computing methodologies** → **Computational photography**.

Additional Key Words and Phrases: computational optics

ACM Reference Format:
Praneeth Chakravarthula, Seung-Hwan Baek, Florian Schiffers, Ethan Tseng, Grace Kuo, Andrew Maimone, Nathan Matsuda, Oliver Cossairt,, Douglas Lanman, and Felix Heide. 2022. Pupil-aware Holography Supplementary Information. *ACM Trans. Graph.* 41, 6, Article 212 (December 2022), 13 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 ADDITIONAL DETAILS ON THE METHOD

In this section, we provide additional details on the proposed method for generating pupil-aware holograms.

1.1 Pupil-aware Image Formation Forward Model

We aim to present uniform image quality, irrespective of the sampling by eye across the exit-pupil or the eyebox. Before we consider pupil sampling in the proposed method, we ensure that light from every image pixel is spread out uniformly and reaches the eyebox as parallel as possible, in a directed wavefront. Specifically, if our SLM with phase Φ_{SLM} and a coplanar thin lens of focal length f is illuminated by a plane wave U_{in}, we want the resulting field U_{out} to also be a plane wave reaching the eyebox, that is

$$U_{\text{out}} = U_{\text{in}} e^{j\Phi_{\text{SLM}}} e^{\left(j\frac{k}{2f}(x^2+y^2)\right)}, \tag{1}$$

where $k = 2\pi/\lambda$ is the wave number. Since we require the input and output waves to be planar, the total phase is zero, that is

$$\Phi_{\text{SLM}} + \frac{k}{2f}(x^2 + y^2) = 0. \tag{2}$$

Therefore, the SLM phase

$$\Phi_{\text{SLM}} = -\frac{k}{2f}(x^2 + y^2) \tag{3}$$

is always a constant and, in fact, equal to the conjugate phase of the co-planar lens. However, the condition of a *constant SLM phase makes holographic image generation impossible*. We lift this restriction by placing the target

0730-0301/2022/12-ART212
https://doi.org/10.1145/nnnnnnn.nnnnnnn

image at a distance f away from the SLM, where each target pixel acts as a diverging point source spanning an extent not more than the maximum diffraction angle of the SLM. Consider a target image pixel that is *on the optical axis* of the display. While the target pixel emits a diverging beam, the co-planar lens throws it to infinity, effectively turning the diverging beam into plane waves across the eyebox. This parallel beam will then be focused onto the retina by the eye's lens. As a result, the pixel is visible at any position within the eyebox. For target image pixels that are *off-centered from the optical axis*, we apply an offset to the quadratic phase function displayed on the SLM in both dimensions

$$\Phi_{\text{SLM}} = -\frac{k}{2f}\left((x - t_x)^2 + (y - t_y)^2\right), \tag{4}$$

thereby turning the effective phase of output wave from Eq. (1) into

$$U_{\text{out}} = U_{\text{in}} e^{j\Phi_{\text{SLM}}} e^{j\frac{k}{2f}(x^2+y^2)} \tag{5}$$

$$= U_{\text{in}} e^{-j\frac{k}{2f}\left((x-t_x)^2+(y-t_y)^2\right)} e^{j\frac{k}{2f}(x^2+y^2)} \tag{6}$$

$$= U_{\text{in}} e^{-j\frac{k}{2f}\left((x-t_x)^2+(y-t_y)^2+(x^2+y^2)\right)} \tag{7}$$

$$= U_{\text{in}} e^{j\frac{k}{2f}(2t_x x - t_x^2)+(2t_y y - t_y)^2} \tag{8}$$

where t_x and t_y are the translation offsets in x and y directions respectively. Note that this only results in a tilt in the output wave while leaving the planar wavefront itself unaffected.

1.2 Trade-off between eyebox-size and field-of-view for the pupil-invariant configuration

For conventional holography the trade-off between field-of-view and eyebox-size is commonly referred to as the limited etendue problem. The etendue is well studied for both near- and far-field configurations. For our proposed setup with a collocated lens, this trade-off still exists, but it is less obvious. The main purpose of the collocated lens is to ray-bias the scattered light of the SLM towards the eye-box. The eyebox is located at the focal-plane of the collocated lens. Given a SLM with a fixed pixel-pitch, the focal length will influence the extent of the eyebox and the field-of-view. From Fig. 1 and 2 we can see that a longer focal-length will correspond to larger eyebox at the cost of a smaller FoV. The configuration with a collocated lens can be kept relatively compact which is needed for any practical near-eye display. However, but it allows us to choose an eyebox/FoV-pair that fits the design needs when given a display with a fixed etendue.

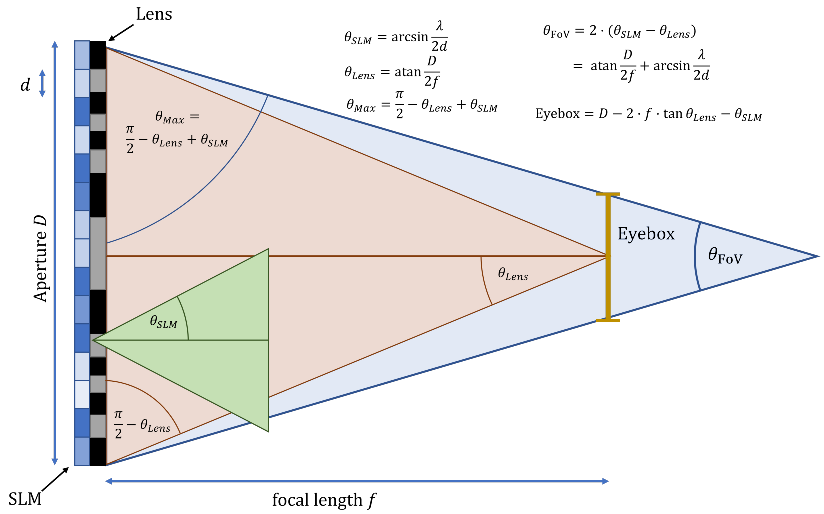

Fig. 1. The geometry of our ray-biasing non-relay display. The scattering angle of the SLM is limited by its pixel size. Together with the refraction angle of the collocated lens we can compute a new "effective maximal focal length" which is longer than the original focal length. However, if still place the eye-box at the original focal-plane, we can see that the eye-box is now actually larger, however on the cost o a smaller FoV.

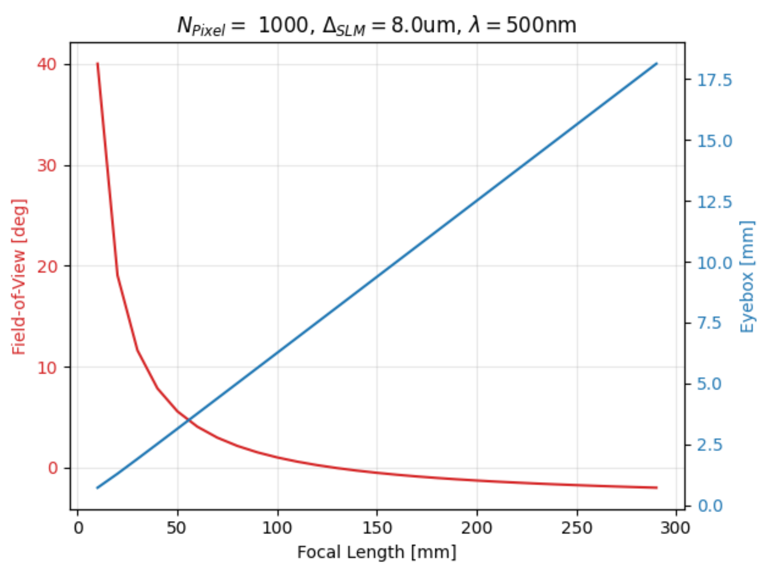

Fig. 2. This plots shows the trade-off between the Field-of-View and the eyebox for a fixed number of SLM pixels of 1000 and a pixel pitch of 8 um for a wavelength of 500nm.

2 ADDITIONAL EXPERIMENTAL SETUP DETAILS

Next, we provide more details on the hardware setups, especially the relayed-SLM configurations.

2.1 Ray-biasing Non-relay Display

We describe our prototype display implementation in detail in the main document. We show in Fig. 3 how our display configuration relates to the eyebox and the images as viewed by the eye.

2.2 Near-field Holograms on a Relayed-SLM Display

Our relayed-SLM prototype holographic display schematic is shown in Figure 4. We build this with a HOLOEYE PLUTO spatial light modulator with a resolution of 1920×1080 and a pixel pitch of $8\mu m$. This SLM is illuminated by a collimated and linearly polarized beam from a single optical fiber emitting at a wavelength of $636nm$ and controlled using a ThorLabs KLD101 Kinesis K-Cube laser diode driver. We use $150mm$ and $75mm$ achromatic lenses to demagnify the SLM and relay it to a virtual SLM plane creating an eyebox of about $6mm$. On the intermediate Fourier plane between these lenses, we place an iris to filter higher diffraction orders. The wavefield produced by the virtual SLM is then sampled by an iris and measured by a Point Grey FLIR machine vision camera with a focusing lens.

2.3 Far-field Holograms on a Relayed-SLM Display

Our relayed-SLM prototype holographic display schematic is shown in Figure 5. We build this with a HOLOEYE PLUTO spatial light modulator with a resolution of 1920×1080 and a pixel pitch of $8\mu m$. This SLM is illuminated by a collimated and linearly polarized beam from a single optical fiber emitting at a wavelength of $636nm$ and controlled using a ThorLabs KLD101 Kinesis K-Cube laser diode driver. We use $150mm$ and $75mm$ achromatic lenses to demagnify the SLM and relay it to a virtual SLM plane creating an eyebox of about $6mm$. On the intermediate Fourier plane between these lenses, we place a DC block to filter the unmodulated light from the SLM. This is typical to all far-field Fourier holograms. The wavefield produced by the virtual SLM is then sampled by an iris and measured by a Point Grey FLIR machine vision camera with a focusing lens. This configuration is similar to the near-field holographic configuration except for replacing the iris with a DC block filter. However, instead of an imaging lens we're not using a 2f-setup to measure directly the fourier-transform of the eyebox. Note that we use a tilted plane wave coming out of the SLM since a central plane-wave would be blocked by the DC-block in the Fourier plane. The tilted plane wave leads to an off-centered point focused on the detector.

2.4 SLM Phase Calibration

We calibrate the SLM to produce a linear mapping between the pixel voltages and the displayed gray level of the hologram. Our calibrated SLM look up table (LUT) for the voltages and the corresponding full range gray scale phase value can be seen in Figure 6. As can be observed, the mapping between the voltages and the gray levels is linear within the 8 bit range and hence we use 8 bit quantization of our optimized phase holograms as well as the near- and far-field holograms for our experimental prototype displays.

2.5 RGB color capture of holograms

In order to capture RGB colors we optimize a SLM-pattern for each independently. The optimized patterns are displayed and captured in time multiplex fashion. During the color capture the same hardware based phase-look-up-table was used for each wavelength. Since the look-up table depends on the wavelength, we need to adjust the phase-range on the software side. We do so by an affine scaling of our grayscale values which we pre-calibrate and is fixed during capture.

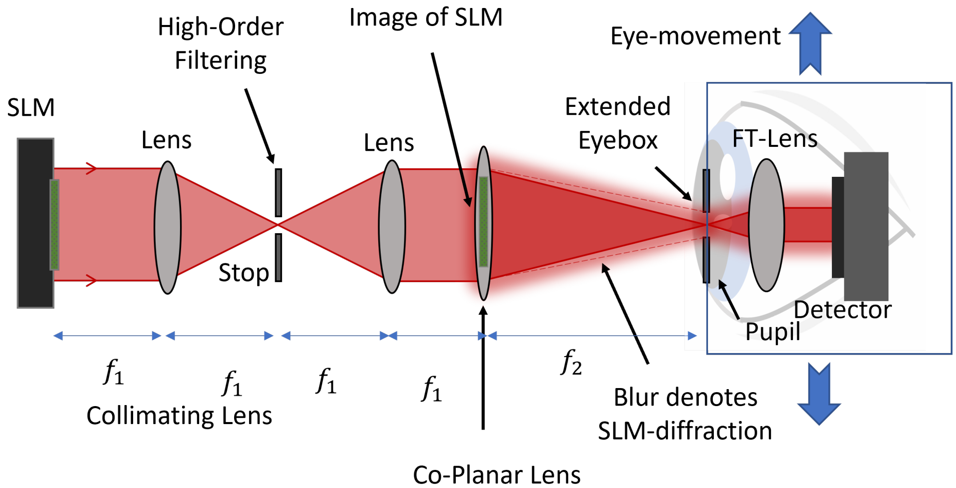

Fig. 3. The proposed hardware setup. The DPAC-encoded phase-only field creates a complex-field at the conjugate field due to high-pass filter in Fourier plane. Now, a lens is placed at the image-plane of the SLM which modulates the SLM-field with a quadratic phase function and biases the field towards the eye-box. Note that due to the diffraction of the SLM, the eye-box has actually larger support than just the focal-spot which we indicated by a blur. At the eye-box a pupil samples the field which is then Fourier-transformed on to detector plane to emulate the imaging process of the eye. In our setup we can move the eye laterally to demonstrate pupil-invariant behavior. The eye consists of iris, fourier-lens and the detector and we denote that the whole module can be moved by the blue arrows.

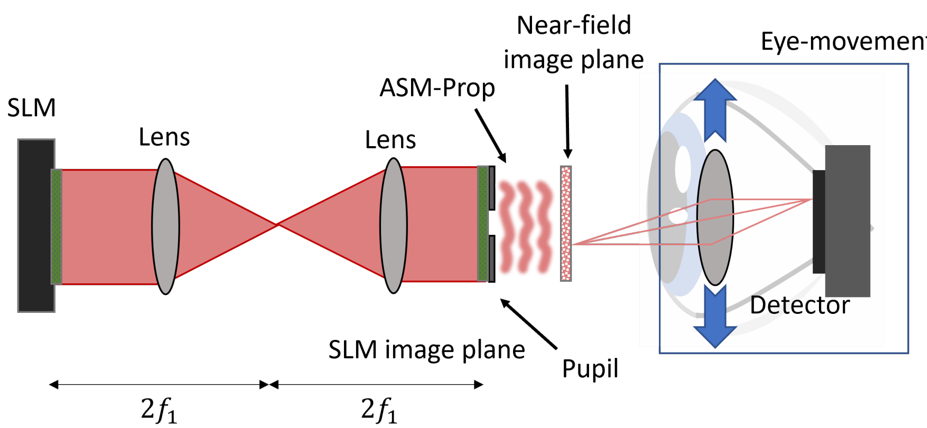

Fig. 4. Sketch of a typical relay-based hardware setup in near-field configuration. The iris stop is used to block the high frequencies created due to higher diffraction orders of the SLM. To show that the near field configuration is not robust against pupil-shifts we shift the whole eye module. Our eye module consists of iris, fourier-lens and the detector. We denote that the whole module can be moved laterally by the blue arrows in the sketch.

The the different colors have different output power which lead - if unaccounted - to a wrong color balance. We use a powermeter to adjust the exposure time for each color such that a plane wavefront creates the same measured intensity over the detector. No further pre-processing is done.

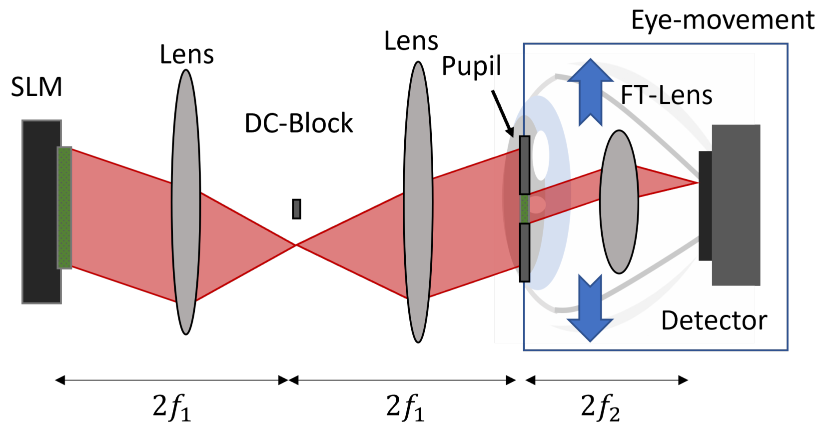

Fig. 5. Sketch of a typical relay-based hardware setup in Far-Field configuration. The only difference to the near-field configuration is that the iris is replaced with a DC-block. Note that the DC-block is conjugate to the detector plane. This is the reason we see the DC-block in our experimental images. To show that the near field configuration is not robust against pupil-shifts we shift the whole eye module. Our eye module consists of iris, fourier-lens and the detector. We denote that the whole module can be moved laterally by the blue arrows in the sketch.

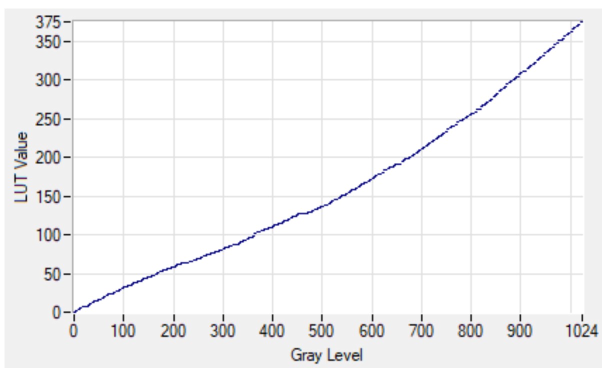

Fig. 6. We calibrate our SLM with a spatially-invariant look-up-table that maps from gray-level (10 bit == 1028 levels) to the corresponding phase-value. Note that the LUT-curve slightly deviates from a linear curve and LUT-curve calibration necessary for achieve the best image quality. We therefore use 8 bits (256 gray levels) where the phase modulation is linear for displaying and evaluating the holograms generated by various methods.

3 ADDITIONAL SYNTHETIC RESULTS

We show additional results on the ablation study of our pupil-invariant hologram phase optimization in Fig. 7. A detailed description can be found in the main document.

For our ablation study we compare three configurations, relayed SLM displaying near-field holograms, relayed-SLM displaying far-field holograms and our pupil-invariant display holograms. For all the target images, we see similar behaviour as reported in the main manuscript.

4 ADDITIONAL HARDWARE RESULTS

We show additional results captured on our prototype holographic displays.

Figure 8 and Figure 9 compare the quality of reconstructed images for various pupil sampling positions for all the display configurations. Figure 10 on the other hand extends the teaser image in the main document and shows the image reconstructions for various pupil sampling positions on our protoype displays.

Figures 11 and 12 show a sequence of color results for both SLM-patterns computed with the conventional (non-pupil aware) and our proposed pupil aware algorithm where the eye-pupil was moved from the center of the eyebox to its edge. As expected, both algorithms provide good image quality for pupil locations close to the center of the eye-box. However, compared to our proposed pupil-aware algorithm, the non-pupil aware hologram formation decreases much quicker quality much quicker since the eye-box created by the SLM has its energy concentrated mostly in the center. On the contrary, our pupil-aware algorithm distributes more of the beams energy around the whole eye-box, which allows some images (albeit noisy) to be formed at region where the conventional algorithm acts as a pure high-pass filter canceling all of the low-frequency terms.

5 TARGET PHASE COMPARISONS

In Fig. 13 we show the optimized and encoded SLM-phase patterns for six different target images. We further include zoomed insets of the amplitude and phase to highlight the spatial features of the phase patterns. The optimized phase-patterns are correlated with the spatial features in the amplitude image. The phase is mostly smooth which leads to less-noisy reconstructions on the detector plane compared to phase-patterns that show a fully randomized phase-distribution profile.

Note, that the optimized phase-patterns still exhibit regular high-frequency components. The checkerboard like features stem from the interleaved Double Phase Amplitude Coding of the optimized complex SLM wave field as described in the main manuscript. However, it can be noticed that there are further high frequencies visible in the hologram which result in a more even distribution of light within the eyebox as well as high frequency noise that is pushed into the higher diffraction orders. We filter these high-frequency components optically via a physical aperture placed in the Fourier domain.

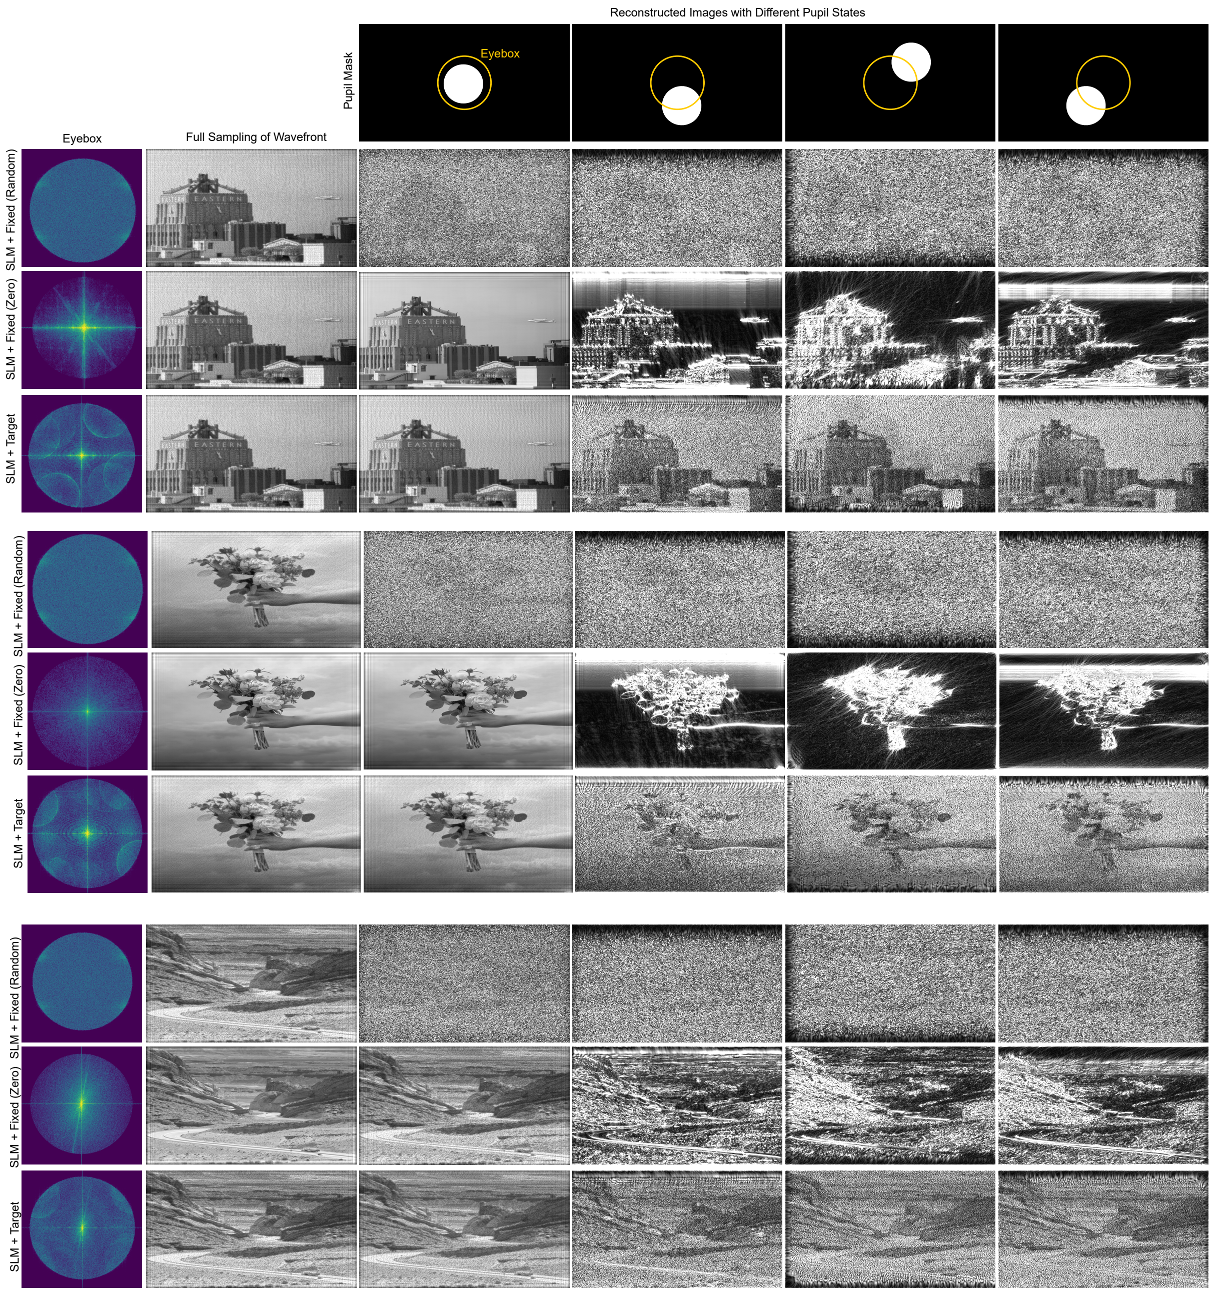

Fig. 7. Additional Ablation Experiments (Simulation). The proposed joint optimization of the target phase and SLM phase improves pupil-invariance as compared to optimizing only SLM phase with fixed target phase: constant (center row) and random (top row) phase. All ablation experiments shown here have been achieved with the same optimization method proposed in this work, either with or without optimizing target phase along with SLM phase. The first column shows the energy distribution of the arriving wavefront at the eyebox and the second column shows the image reconstructions when the full wavefront is sampled. The remaining four columns show the reconstruction results at four different pupil positions indicated as in the first row.

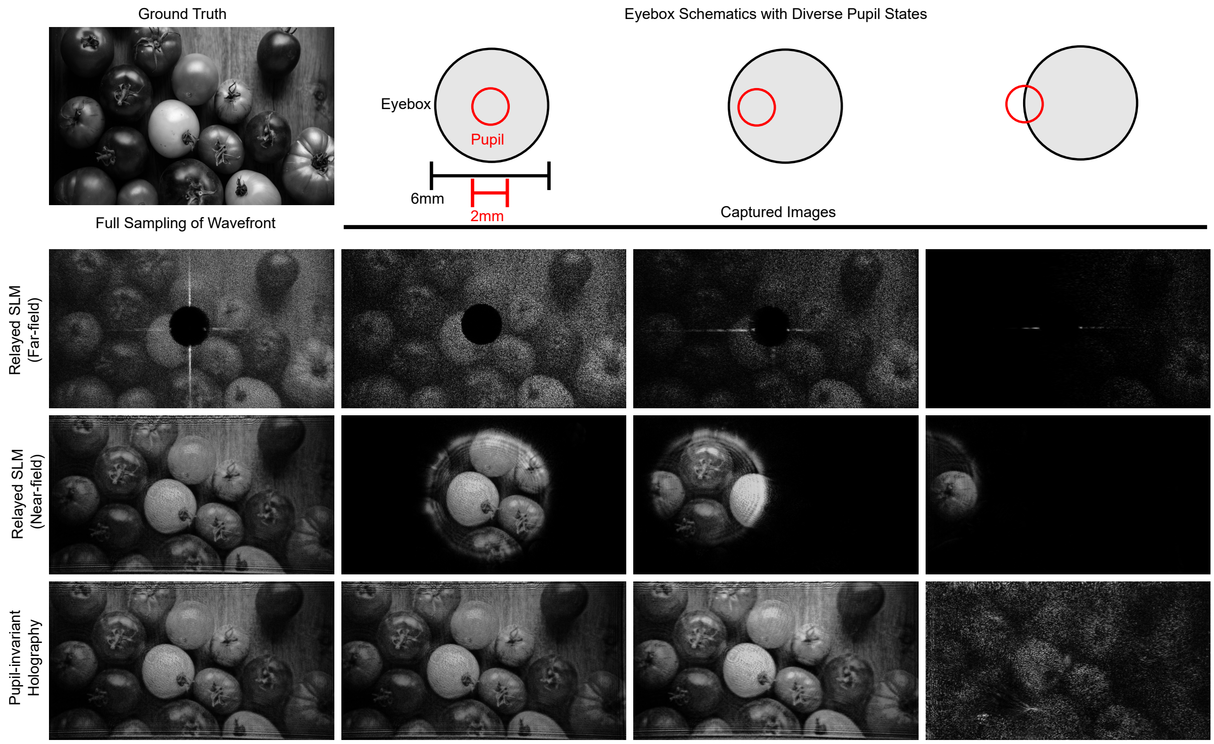

Fig. 8. Experimental Evaluations. We experimental validated the effectiveness of our pupil-invariant holography by capturing the reconstructed holograms at three different pupil states shown in the top row. While the far-field/near-field relayed SLM provides accurate reconstructions when the wavefront is fully sampled, they suffer from speckle and cropping artifacts for partial pupil sampling. In contrast, our pupil-invariant holography clearly shows better consistency for diverse pupil states. Note that our method even maintains the structure of the berry with the partial edge sampling shown on the most right column.

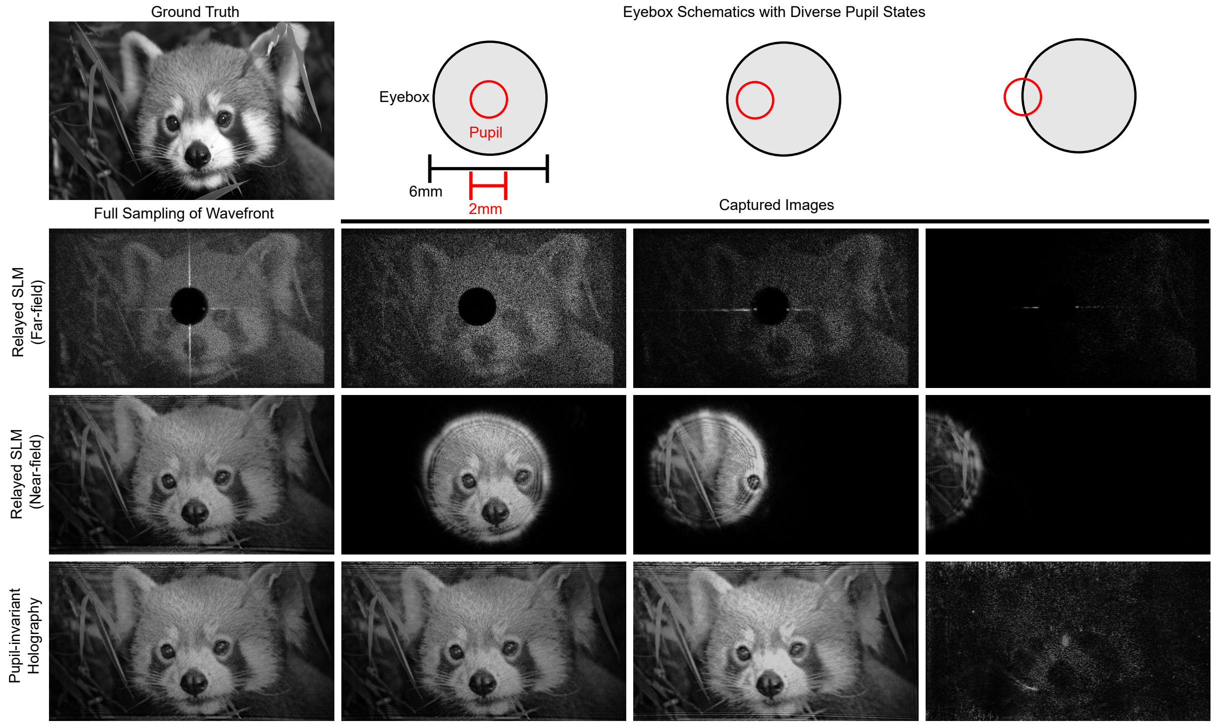

Fig. 9. Experimental Evaluations. We experimental validated the effectiveness of our pupil-invariant holography by capturing the reconstructed holograms at three different pupil states shown in the top row. While the far-field/near-field relayed SLM provides accurate reconstructions when the wavefront is fully sampled, they suffer from speckle and cropping artifacts for partial pupil sampling. In contrast, our pupil-invariant holography clearly shows better consistency for diverse pupil states. Note that our method even maintains the structure of the panda with the partial edge sampling shown on the most right column.

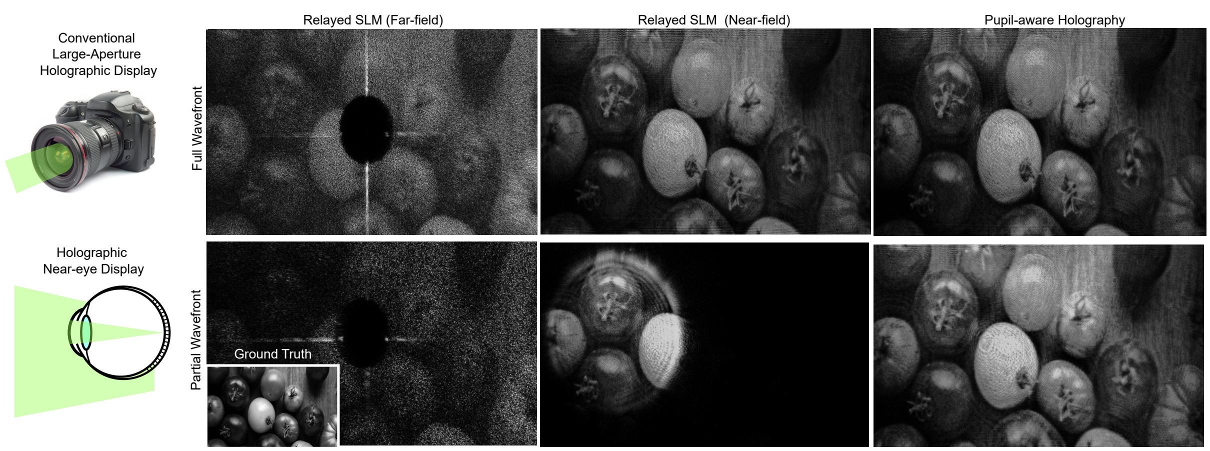

Fig. 10. We show the extended results of full and partial wavefront sampling for the teaser figure in the main paper.

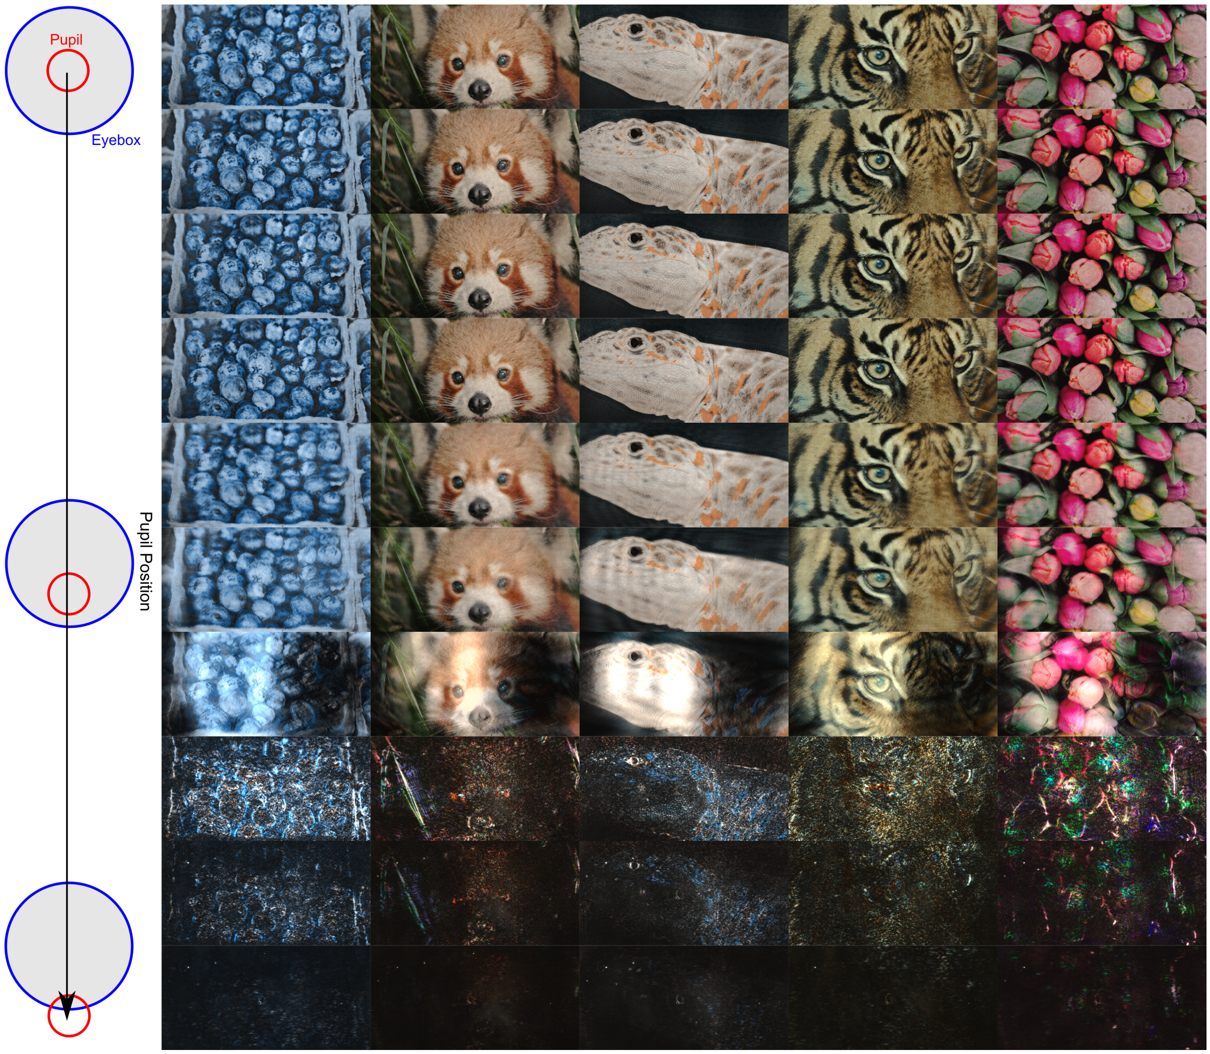

Fig. 11. Additional results for conventional reconstruction algorithm (i.e. non-pupil aware) where the pupil was moved from the center towards the edge of the eyebox.

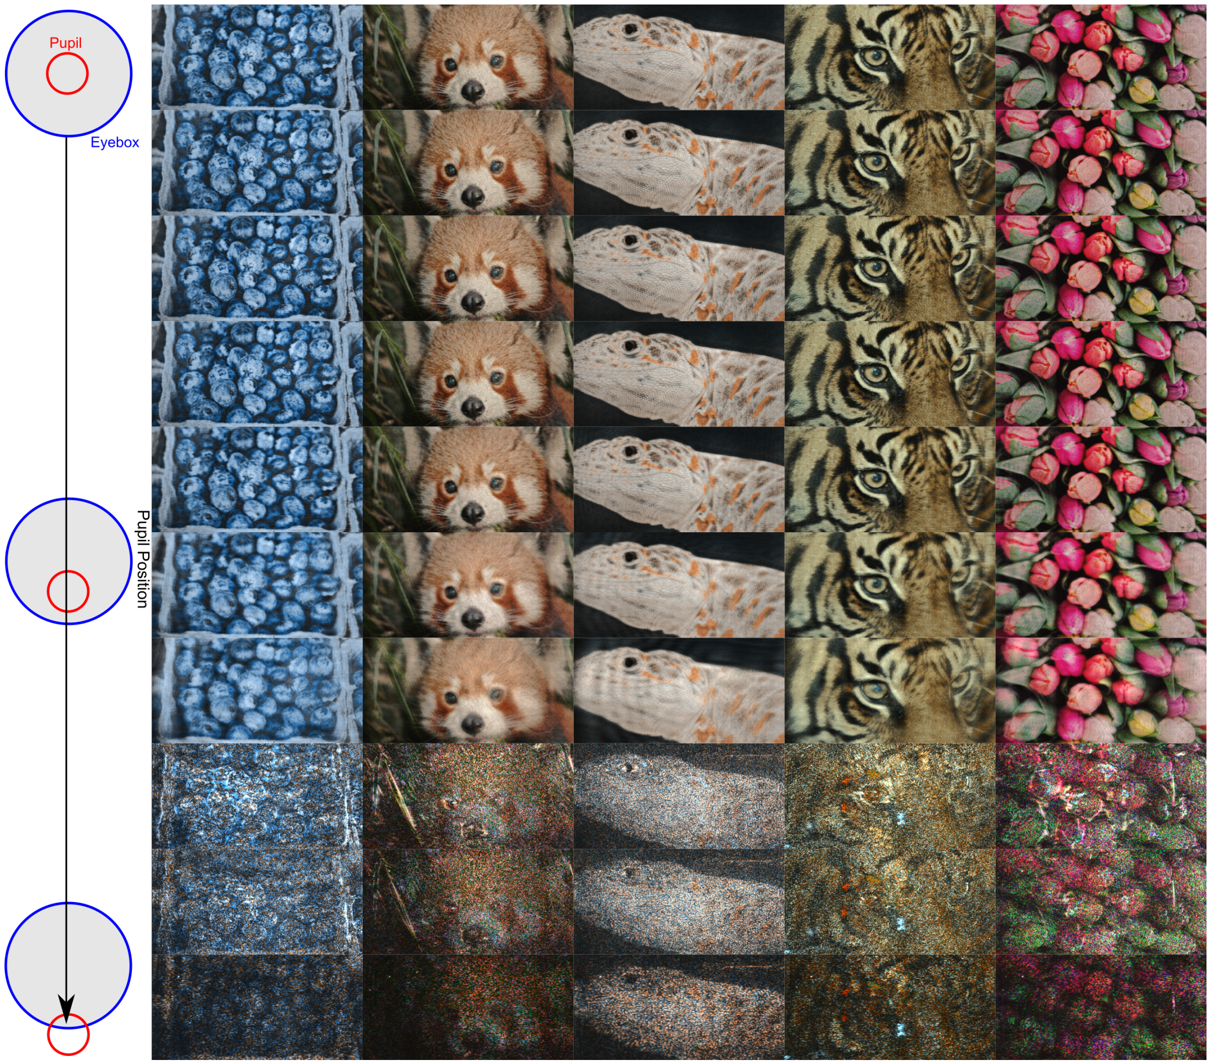

Fig. 12. Additional results for the proposed pupil-aware reconstruction algorithm where the pupil was moved from the center towards the edge of the eyebox.

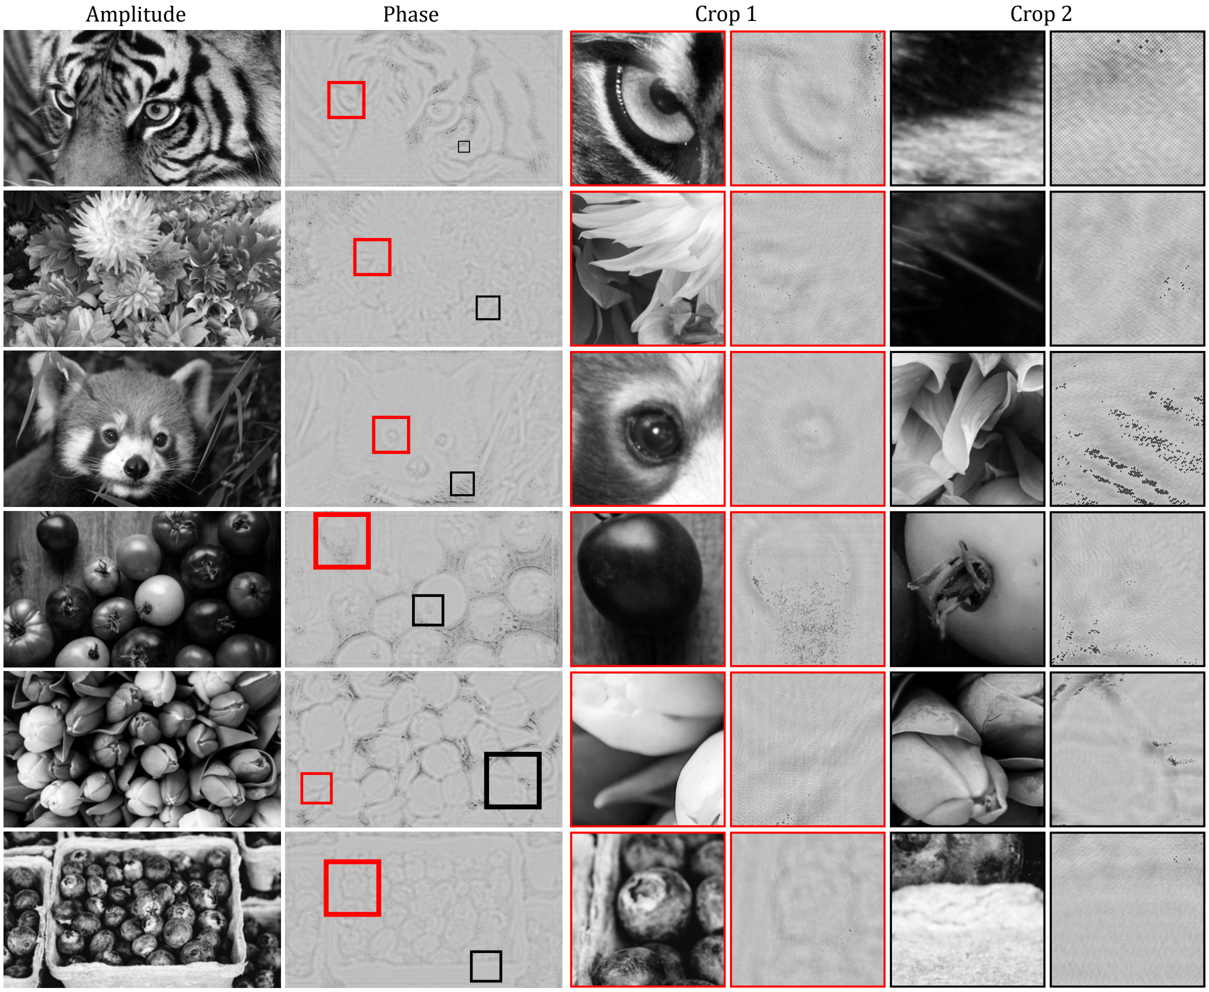

Fig. 13. Here we show the optimized SLM-phase patterns corresponding to the respective amplitude images shown on the left. We include two zoomed crops for each complex target image. Our pupil-aware optimization scheme tends to converge to SLM-phase-patterns that share many spatial features with the target amplitude. Note that the target-phase is implicitly modulated via a checkerboard pattern since we incorporate Double-Phase-Amplitude-Coding (DPAC) in our optimization scheme. We have chosen different crop-sizes, hence sometimes the checkerboard patter is less visible than at other times. This creates high-frequency noise which we filter away via a physical filter in the Fourier plane. Ignoring the high-frequency components from the checkerboard and outliers (like the black dots visible e.g. with the Panda), our optimized SLM-phases are smooth leading to less noisy reconstructions.