

Super Resolved Imaging with Adaptive Optics

Robin Swanson^{1,2} Esther Y. H. Lin¹ Masen Lamb^{3,4} Suresh Sivanandam^{1,2} Kiriakos N. Kutulakos¹

¹ University of Toronto ² Dunlap Institute for Astronomy & Astrophysics

³ International Gemini Observatory ⁴ University of Victoria

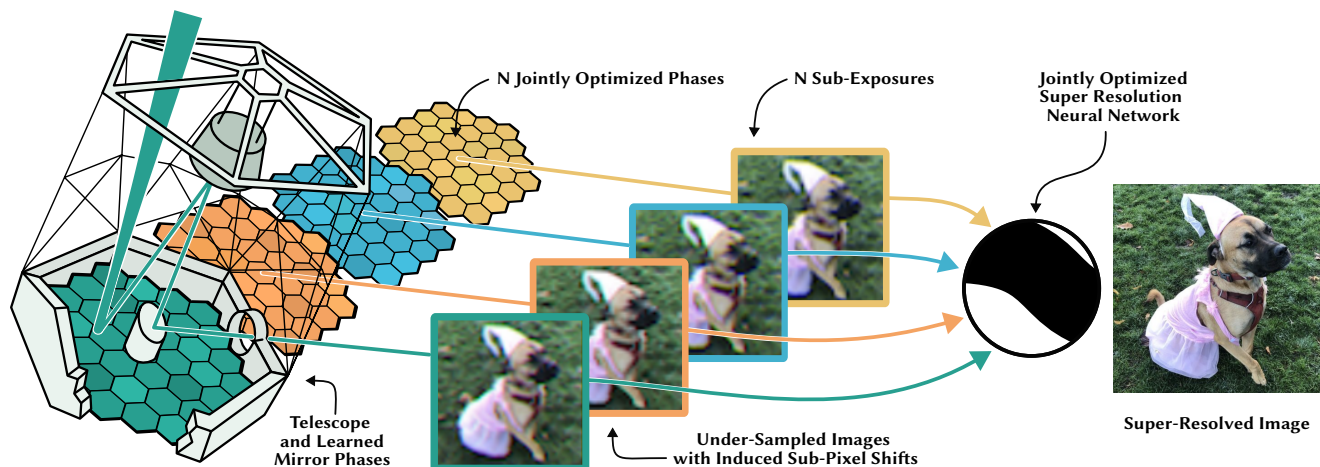


Figure 1. *Super Resolved Imaging with Adaptive Optics* is a computational imaging method to enable multi-image super-resolution in scientific telescopes. Without any additional hardware, existing adaptive mirrors are used to inject diversity into each sub-exposure. When jointly reconstructed, we are able to recover higher spatial resolution from an otherwise undersampled signal. Each exposure’s phase profiles is jointly optimized along with the image reconstruction algorithm—resulting in greatly improved performance.

Abstract

Astronomical telescopes suffer from a tradeoff between field of view (FoV) and image resolution: increasing the FoV leads to an optical field that is under-sampled by the science camera. This work presents a novel computational imaging approach to overcome this tradeoff by leveraging the existing adaptive optics (AO) systems in modern ground-based telescopes. Our key idea is to use the AO system’s deformable mirror to apply a series of learned, precisely controlled distortions to the optical wavefront, producing a sequence of images that exhibit distinct, high-frequency, sub-pixel shifts. These images can then be jointly upsampled to yield the final super-resolved image. Crucially, we show this can be done while simultaneously maintaining the core AO operation—correcting for the unknown and rapidly changing wavefront distortions caused by Earth’s atmosphere. To achieve this, we incorporate end-to-end optimization of both the induced mirror distortions and the

upsampling algorithm, such that telescope-specific optics and temporal statistics of atmospheric wavefront distortions are accounted for. Our experimental results with a hardware prototype, as well as simulations, demonstrate significant SNR improvements of up to 12 dB over non-AO super-resolution baselines, using only existing telescope optics and no hardware modifications. Moreover, by using a precise bench-top replica of a complete telescope and AO system, we show that our methodology can be readily transferred to an operational telescope.

1. Introduction

Super-resolution (SR) techniques have been integral to astronomical imaging due to inherent telescope design constraints [23]. Telescopes can have many science instruments

Corresponding author: robin@cs.toronto.edu
Project website: www.cs.toronto.edu/~robin/aosr

and in certain cases optical resolution is not their priority [67]. A large field of view (FoV), for example, necessitates undersampling to capture a wider view of the night sky. Another common case in photon-starved conditions is to increase signal-to-noise ratio by focusing the light onto fewer, larger, pixels, thereby incurring less read noise [43]. In these, and many other, cases although the telescope is capable of providing a diffraction-limited signal, the science instrument can be undersampled. Thus, the use of SR methods in astronomical imaging has a long history in both satellite and terrestrial telescopes [34].

Unlike telescopes stationed in orbit, ground telescopes face an additional challenge: the Earth’s atmosphere [5]. Our atmosphere is a source of dynamic optical aberrations which introduce additional difficulties regardless of a telescope’s location, size, or design [58]. As the atmosphere evolves, small changes in index of refraction occur across the sky. This causes incoming light to distort as it transitions from outer space and through our telescope, blurring the images. Therefore, traditional multi-image SR methods which combine multiple sharp images with small geometric distortions—whether from telescope movements [23] or Earth’s rotation [4] are not well suited for these conditions.

We present Super Resolved Imaging with Adaptive Optics, an end-to-end SR method for terrestrial telescopes that leverages existing adaptive optics (AO) systems—designed to correct for atmospheric distortions—to simultaneously correct for aberrations and achieve SR. The crux of our approach is to take advantage of the existing AO system in ground telescopes, which applies a phase mask in the telescope’s Fourier domain to counteract atmospheric aberrations. We use the AO system by optimizing and injecting additional phase displacements that produce sub-pixel shifts in the image plane. After capturing several images, these shifts serve as inputs to an end-to-end trained, multi-image SR (MISR) method. These additional phases are optimized to minimally alter the telescope’s point spread function (PSF), thus subtly modifying each exposure to best facilitate MISR. The phase displacements can be co-optimized with deep learning-based (or any differentiable) MISR to enhance overall performance. This optimization is performed entirely offline, requiring no modifications to the telescope or use of on-sky time.

In summary, our contributions include:

- A novel MISR method for telescopes that leverages existing hardware, requires no modifications, and does not impact the science payload.
- Demonstration that combining AO with MISR outperforms either technique alone, achieving performance greater than the sum of its parts.

- Showing that phase displacements and image reconstruction can be co-optimized end-to-end offline using simple telescope telemetry as a proxy for real systems, enabling unobtrusive deployment where telescope time is extremely costly [31].
- A real-hardware setup demonstrating the feasibility of integrating our method into existing telescope systems.

2. Principles of Adaptive Optics

The Earth’s atmosphere fundamentally limits the resolving power of any telescope by introducing optical aberrations—one of the reasons why space-based telescopes such as the Hubble Space Telescope have outperformed many larger ground-based observatories. This results in the final telescope PSF being several factors worse than the diffraction limit—the best achievable resolution. As such, AO systems are now widely used in modern ground-based optical and near-infrared telescopes to reverse these effects.

2.1. Components of an AO system

AO systems correct atmospheric turbulence by introducing a phase equal in magnitude but opposite in sign to that of the atmosphere. A prototypical AO setup (Fig. 2) consists of a deformable mirror (DM), wavefront sensor (WFS), and control system. The DM, placed on the Fourier plane of the telescope, takes on the opposite phase of the atmosphere, resulting in corrected wavefronts of light which are then passed on to the science camera and back to the WFS. Because the atmosphere evolves over time, the WFS must continuously measure the residual, uncorrected wavefront, and send updated commands to the DM. This loop, implemented via the control system, typically operates on the order of 1-2 kHz for modern AO systems.

The most common types of WFSs are based on Shack-Hartmann [50] and pyramid [51] designs. To operate they require a “guide-star”—a distant light source which can be approximated as a point. Usually a bright star near the science object is used, but laser beacons can also be used to create an artificial star image in the upper atmosphere. The guide star samples a similar turbulence path as the science image without reducing its signal-to-noise ratio (SNR). By comparing the measured signal to a reference, we can calculate how the atmosphere has affected the wavefront across the telescope pupil. Pre-calibrated responses are then used to update the DM shape and negate the distortions [58].

The DM can be constructed in several ways. Many of today’s telescopes use a continuous design with a single flexible reflective membrane supported by a grid of discrete actuators [20]. These actuators can be individually pushed or pulled to collectively generate complex phase patterns. In contrast, many upcoming AO systems use grids of individual mirrors, enabling much larger DMs and corrections for

even higher spatial frequencies [26].

Once converged, the effects of the AO loop, shown in Fig. 2, are visible and have enabled many recent breakthroughs in modern astronomy. Direct imaging of exoplanets [42] and the first direct evidence of a super-massive black hole within our galaxy [25] would not have been possible without the increased resolution and contrast provided by AO. For additional details on how AO has been recently used in both astronomy and microscopy see [28].




Figure 2. Starlight is aberrated by Earth’s atmosphere \mathcal{A} , blurring the telescope PSF. An AO system (a) measures the atmosphere with a wavefront sensor and applies a corrective phase, \mathcal{D} , on its deformable mirror to cancel it out. When AO is off, resolution is directly limited by the seeing conditions (i.e., turbulence) at the observatory for both short and long exposures (b). When AO is enabled the effects of the atmosphere are removed. If sampled at or above the Nyquist rate, the resolution is limited only by the mirror diameter and its PSF forms a diffraction-limited Airy disk (c). Our method enables under-sampled instruments with AO to recover spatial resolution.

2.2. The AO Image Formation Model

Suppose we are observing a distant, constant object (e.g., a natural guide star) over an exposure time T using a telescope without AO. The recorded image $I(x, y)$ in pixel coordinates (x, y) is modeled as the convolution of the object $O(x, y)$ with an atmosphere-degraded PSF \mathcal{P}_{TEL} . Here, \mathcal{P}_{TEL} corresponds to a diffraction-limited PSF $\mathcal{P}_{\text{IDEAL}}$ (i.e., an Airy Disk [59]) degraded by the time-varying turbulence of the Earth’s atmosphere. In Fourier space, this turbulence can be represented by a phase distribution $\mathcal{A}(u, v, t)$ with unit amplitude where (u, v) are its spatial frequency coordinates. By the convolution theorem, this can be described in Fourier space with a point-wise multiplication:

$$\mathcal{F}\{\mathcal{P}_{\text{TEL}}\}(u, v, t) = e^{j\mathcal{A}(u, v, t)} \mathcal{F}\{\mathcal{P}_{\text{IDEAL}}\}(u, v, t), \quad (1)$$

where $\mathcal{F}\{\cdot\}$ denotes the Fourier Transform operator. While the AO is in operation, a DM continuously applies an additional phase distortion $\mathcal{D}(u, v, t)$ in the Fourier plane of the telescope to counteract the effects of the atmosphere. The resulting corrected PSF, $\mathcal{P}_{\text{CORR}}$, can be described in the Fourier domain as

$$\mathcal{F}\{\mathcal{P}_{\text{CORR}}\}(u, v, t) = e^{j\mathcal{D}(u, v, t)} e^{j\mathcal{A}(u, v, t)} \mathcal{F}\{\mathcal{P}_{\text{IDEAL}}\}(u, v, t). \quad (2)$$

In an ideal scenario, the atmospheric phase \mathcal{A} and the deformable mirror phase \mathcal{D} would perfectly cancel at every point and time; i.e., $\mathcal{P}_{\text{CORR}} \approx \mathcal{P}_{\text{IDEAL}}$.

In practice the DM is limited by fitting errors and the finite range of its actuators which restricts the range of spatial frequencies it can correct. Additionally, the effects of non-circular mirror geometries [29] (e.g., hexagonal apertures) and mirror support structures cause the DM to deviate from ideal performance. Consequently, the final effective PSF of the telescope never truly reaches $\mathcal{P}_{\text{IDEAL}}$. We model these effects on \mathcal{P}_{EFF} in the Fourier domain as

$$\mathcal{F}\{\mathcal{P}_{\text{EFF}}\}(u, v, t) = e^{j\delta(u, v, t)} \mathcal{F}\{\mathcal{P}_{\text{CORR}}\}(u, v, t), \quad (3)$$

where $\delta(u, v, t)$ encapsulates the hardware and optical constraints described above.

Note that these contributions are predominantly static and can be approximated as $\delta(u, v, t) \approx \delta(u, v)$. For a well-performing, diffraction-limited AO system, variations will be short because the AO re-converges quickly. Therefore, over long exposures, they average out over time. Finally, we consider incoherent image formation, as astronomical sources emit light independently with random phases. Our image formation becomes

$$I(x, y) = \int_0^T |O(x, y)|^2 \otimes |\mathcal{P}_{\text{EFF}}(x, y, t)|^2 dt + \eta \quad (4)$$

where \otimes denotes the convolution operator, T is the total exposure time, and η is cumulative noise.

3. Previous Work

Image Super-Resolution (SR). SR techniques aim to reconstruct a high-resolution (HR) image from one or more observed low-resolution (LR) images [65]. However, the problem is ill-posed in nature, and many plausible solutions exist [41]. SR methods are distinguished between single-image (SISR) and multi-image (MISR) super-resolution depending on the number of images used. Deep-learning SR has emerged as the state-of-the-art for both SISR and MISR tasks [2, 41, 41, 56, 64], thanks to the growth of convolutional networks and large datasets [18, 35]. However, deep SR networks thrive on natural images but struggle with out-of-distribution data due to limited training data [2]. In

such cases, incorporating priors based on sensor, object, and acquisition conditions can improve performance [2]. Astronomical imaging is a regime with unique conditions, including extreme low-light, non-static aberrations, and image features unlike those typically seen in terrestrial images. Hence, additional modifications to existing deep-learning SR techniques are required.

SR in Astronomical Imaging: *Drizzle*—a SR technique commonly used in astronomical imaging and astrophotography—aligns and linearly blends multiple deliberately dithered frames to produce a super-resolved image [23]. The technique has been widely adopted, including by the Hubble Space Telescope, and remains in use today [27]. However, the success of drizzling and its contemporaries hinges on highly accurate image registration and precise knowledge of the offsets [3, 8], a challenging task when the atmosphere adds time-varying aberrations and distortions [24]. In contrast, our method leverages the telescope’s built-in AO system to inject optimized phase-domain “jitters” and does not use any type of image registration. SR techniques have also been leveraged to improve the resolution and quality of WFS measurements [46]. Our method, building on our previous work [54], is agnostic to the underlying WFS type and can remain applicable as WFS technology advances.

SR in Computational Imaging: Camera system design often incorporates super-resolution. One early example jointly optimized both the optics and SISR algorithms for simple, refractive lenses [52]. Burst photography from handheld cameras—where small distortions are induced by hand shake—are extremely successful at recovering high-resolution information [7, 62] even in low-light conditions [38]. Deep learning methods have also been explored in the scientific imaging community for both microscopy [21] and satellite imaging [45]. MISR has also long played a role in classic camera development; many digital cameras are able to programmatically shift the imaging sensor between subsequent images which can then be combined into one higher resolution image (e.g., Sony’s Pixel Shift Multi Shooting [15]). This work draws inspiration from previous works by incorporating end-to-end optimization of telescope mirror distortions with SR algorithms.

Computational Wavefront Sensing and Control: CNNs have been used to improve WFS reconstruction [19, 47], as well as to predict atmospheric conditions to reduce servo-lag in the system [33, 53]. Similar techniques have shown to be effective in microscopy AO systems as well [17, 66]. More general-purpose methods have been proposed such as [60] which used coded wavefront sensors and spatial light modulators (SLMs) to achieve much higher order spatial frequency correction than is capable from typical AO sys-

tems in astronomy or microscopy. However, SLMs, unlike DMs, may not be suitable for correcting rapidly changing atmospheric conditions due to their slower speeds [10].

Optical Phase and PSF Optimization: End-to-end optimization of optical systems and their image reconstruction algorithms have been used in thin lens design [57], holography [9, 49], and active illumination [12]. These methods typically focus on optimization of critical optical elements of the system, requiring either direct access to the optical system [63] or robust, differentiable simulations during optimization. Due to the complexity and operating costs of modern telescopes, modification to the underlying system or real-time access for network training would be infeasible. Instead, we focus on optimization of optical elements available to us in such a way that would not affect its operation, allowing us to side step these limitations.

4. Super Resolving Adaptive Optics

This work augments existing MISR methods by learning the optimal way to enhance each individual image captured over a single exposure. Using the DM already present in a telescope’s AO system, we inject optimized phase profiles, thereby enabling multi-image super-resolution.

In this section, we present an end-to-end approach to determine optimal phase profiles while accounting for AO hardware constraints and atmospheric effects on final images. We discuss our approach’s assumptions (Sec. 4.1), the framework for optimizing phase profiles (Sec. 4.2), compatibility with various image reconstruction methods for combining sub-exposures (Sec. 4.3), and the training details for obtaining the optimal phase profiles (Sec. 4.3).

4.1. The “Small Phase Shift” Assumption

To include the image formation model in our learning infrastructure we must make some reasonable assumptions about the process. Importantly, we assume that sub-pixel shifts negligibly impact the AO loop or science image. In our case this is reasonable due to the fact that we consider only under-sampled systems. If our phase shifts keep \mathcal{P}_{TEL} within the range of its under-sampled area, the aggregated effect over the entire exposure will be negligible¹. The shifts define the DM “flat” (a standard feature of AO real-time controllers), and can be updated easily between exposures. Therefore, the AO system will run entirely as usual—unaware of the phase offsets applied by our method.

4.2. Phase Shifting and Optimization

Because our DM is designed to create highly complex shapes, we are not restricted to planar movements like traditional methods [23]. We describe a given DM shape \mathcal{S}

¹If necessary, we can also impose constraints on the magnitude of our learned phase shifts to guarantee science performance

using either actuator positions or, as detailed below, a more general modal basis.

Specifically, one approach would be to use a combination of Zernike polynomial modes, \mathcal{Z} . The Zernike modes are well suited for this problem because of their ability to represent rich phases (including planar shifts) with minimal coefficients and are well-suited for circular apertures, such as those of telescopes [32]. They can also provide some insight into the nature of a given phase due to the relationship between Zernike modes and classical optics aberrations (e.g., focus, coma, and astigmatism [13]). Besides planar shifts (tip and tilt Zernike modes), there is no trivial or well known choice of Zernike modes for super-resolution available. We therefore optimize their values in conjunction with the MISR algorithm of choice to estimate the best possible choice of modes (i.e., an end-to-end optimization).

We can now update our image formation model from a single exposure I_i of time T to include N exposures of time T/N with individually optimizable phases $\mathcal{S}_{i=1\dots N}$

$$I_i = |O|^2 \otimes |\mathcal{P}_i \otimes e^{j\mathcal{S}_i}|^2 + \eta. \quad (5)$$

This same model can be equally applied to a simple shifting or Zernike model by replacing \mathcal{S} with \mathcal{Z} and choosing an appropriate number of Zernike coefficients.

Now, let $f_{\uparrow}(x | \theta)$ be a learned MISR function with parameters θ , and $f_{\downarrow}(x)$ be a downsampling method which reduces an image's spatial dimensions by a fixed factor. Our end-to-end optical optimization problem is then,

$$\min_{\{\theta, \mathcal{S}\}} \left\| f_{\uparrow} \left(f_{\downarrow} \left(|O|^2 \otimes |\mathcal{P}_{1\dots N} \otimes e^{j\mathcal{S}_{1\dots N}}|^2 + \eta \mid \theta \right) \right) - O \right\|_1, \quad (6)$$

where O is an image randomly sampled from a set of high-resolution images, \mathcal{P}_i is randomly chosen from a set of pre-computed AO telescope PSFs, \mathcal{S}_i are our trainable phase profiles, and $\|\cdot\|_1$ denotes the ℓ_1 norm. As described earlier, \mathcal{S} can be represented in many ways. To explore its effects, we compare the most common methods in our experiments. These include:

- **No Learned Phase:** The base case for CNN methods. The same number of exposures are given to the CNN during training—and still affected by \mathcal{P}_i —but without optimizing any per-exposure \mathcal{S}_i .
- **Tip/Tilt:** A common MISR technique where the image is translated by a sub-pixel shift [22]. The *Classic* case uses a sequence of exposures whose translational shifts form a regular sub-pixel raster scan. The *Optimized* variant initializes with the same shifts as *Classic* and then optimizes the values of those shifts.
- **Zernike Modes:** Each phase is represented using the Zernike basis and its coefficients are optimized.

- **Non-Modal:** Instead of a basis function, the DM shape is directly optimized. This requires precise knowledge of its influence functions and calibration with respect to the WFS [6]. While simple in simulation, this proves to be difficult to use in practice.

4.3. Super-Resolution Reconstruction Methods

Although our method could work equally well with any super-resolution reconstruction method, for the purposes of this paper we focus on only two:

First, we consider a linear-reconstruction method similar to Drizzle [23], also known as the shift-and-add method [22]. This may be of interest in science cases where such methods are already widely used, albeit for systems without adaptive optics and with un-optimized phases.

Second, we consider a state of the art deep learning-based SR algorithm. Specifically, we based our model on the widely known EDSR [39] single-image SR CNN. Our network was modified to incorporate the forward image formation model (Eq. (4)) as well as to accept any number of input channels, representing the phase-shifted exposures. The EDSR model was chosen for its speed, widespread adoption, and its avoidance of adversarial training (i.e., it does not rely on a GAN [16]) which can complicate training and introduce priors that may not be appropriate for scientific purposes².

4.4. Phase and Network Training Details

In addition to a training set of high-resolution images, our phase and upsampling parameters require a collection of system PSFs under the influence of the AO system and the expected turbulence. These can be generated by simulations of the system or collected from telescope telemetry.

At training time, for a model expecting N exposures, we provide the network a random batch of high-resolution images as well as a random sample of N system PSFs. The low-resolution images are generated for each PSF, as detailed in Eq. (5), using a trainable phase profile for each sub-exposure. We now have low-resolution images affected by the system PSF in addition to our optimized phases. We also have the ground-truth high-resolution image, which can be used to compute the loss function in Eq. (6).

Results shown in this work were trained on a combination of the DIV2K [1] and Microsoft ImagePairs [30] datasets. Both CNN and linear reconstructors were implemented and trained with `pyTorch` [48] following the learning procedure described in the EDSR paper [39]. Unless stated elsewhere (e.g., Sec. 5.1), the following exper-

²Our outlined procedure can be easily adapted to any SR method including diffusion [37] or transformer [40] models, or traditional convex optimization methods [11].

iments are all performed with $N = 4$ exposures. While our method can be used with any number of exposures, this value was chosen to easily compare to the classic “tip/tilt” methods when solving for a $2\times$ upscaling factor. Restricting the number of exposures to four highlights the performance differences between linear and deep learning-based reconstruction methods, as discussed in Secs. 5 to 7.

5. Simulated Results

We first evaluate our method entirely in simulation. We train SR parameters and phase profiles solely on synthetic PSFs generated by the discretized image-formation model in Eq. (5). After training, we run a new AO simulation, setting the optimized phase profiles as the deformable mirror’s “flat” position. This imposes the optimized phase on the system, replicating our convolution in the simulated hardware. The PSF generated from this new simulation is then used as the PSF for evaluating test images.

5.1. AO Simulation Environment

To create a realistic AO simulation we use the OOMAO [14] toolbox. We implemented a realistic AO simulation with multiple layers of atmosphere to generate realistic PSFs for training. Our simulated AO system was set up to match our experimental hardware (detailed in Sec. 6) with an 8 m diameter telescope and 397-actuator deformable mirror operating at 1000 Hz. Each PSF was generated by performing a full AO simulation with an evolving atmosphere over a 2-second exposure. The simulated \mathcal{P}_{TEL} was therefore imaged on a simulated science camera under the effects of AO correction and realistic sensor noise. To cover a wide range of conditions, the wind direction, wind speed, and guide star magnitude were randomized for each simulation.

Effect of Joint Optimization: We investigate how reconstruction results are affected by jointly optimizing phase profiles and the CNN reconstruction method. As shown in Fig. 4, the jointly-optimized approach yields better resolution enhancement than a naïve approach that trains the CNN without phase optimization, both qualitatively and quantitatively.

Effect of Super-Resolution Factor: We compare different reconstruction methods and phase representations for the case where $N = 4$ exposures. We compare our different phase representations against a baseline CNN case where the input to the neural network are four fixed unshifted PSFs with no optimizable phase and simple bilinear interpolation. Experimental results are shown in Sec. 5.1. Sample phase profiles and output can be found in Fig. 3.

Effect of Number of Exposures: We consider the effect of number of exposures on each method, freezing the super-resolution factor. The results are shown in Table 1.

		Scale Factor		
		$2\times$	$4\times$	$8\times$
CNN	No Learned Phase	33.72	27.98	24.04
	Tip/Tilt (Classic)	39.56	33.46	25.47
	Tip/Tilt (Optimized)	39.67	33.88	26.32
	Zernike Modes	41.51	34.09	27.74
	Non-Modal	43.13	34.24	27.79
Linear	Tip/Tilt (Classic)	28.56	26.04	23.32
	Tip/Tilt (Optimized)	29.11	26.12	23.26
	Zernike Modes	30.84	26.17	23.36
	Non-Modal	31.08	26.67	23.33

		Number of Exposures			
		2	4	6	8
CNN	Tip/Tilt (Classic)	30.77	33.89	36.03	36.31
	Zernike Modes	31.13	34.11	36.29	37.44
	Non-Modal	31.17	34.24	36.24	37.48
Linear	Tip/Tilt	26.11	26.12	26.12	26.11
	Zernike Modes	26.14	26.17	26.16	26.17
	Non-Modal	26.66	26.67	26.71	26.69

Table 1. *Effect of number of exposures and super-resolution factor on SR methods.* Values presented in PSNR (dB) \uparrow . In general the more exposures, and the more expressive the phase basis, the better the results. Details on each method can be found in Sec. 4.2.

6. Experimental Results

To demonstrate our method’s ability to work in realistic settings, we also implemented the experimental setup described in Sec. 5 on an optical table. Our setup, shown in Fig. 5, consists of a 97-actuator ALPAO deformable mirror and a custom-built 25×25 sub-aperture Shack-Hartmann wavefront sensor. Our guide star was created with an 850 nm near-infrared laser whose pupil and beam size were matched to the DM’s size (13.5 mm). Realistic atmospheric turbulence was induced by a rotating Lexitek³ phase screen driven by a stepping motor. The sensor’s native pixel pitch was $5.86 \mu\text{m}$ and the diffraction-limited spot was $22.95 \mu\text{m}$ ($f = 400$ mm, pupil size 13.5 mm, $\lambda = 635$ nm). Note that even though our experimental hardware was not undersampled by design, we used $12\times$ on-sensor binning to capture individual images. This resulted in an effective pixel pitch of $70.3 \mu\text{m}$.

The AO loop was controlled with the open source pyRTC software⁴ which allows full control of the DM and science camera. This allowed us to define the optimized phase shifts to be the reference mirror “flat”, from which the residual wavefront was measured by the AO WFS. Each image was then captured as usual, changing the DM’s “flat” between exposures as dictated by the optimized phases, and then combining these images to restore our final SR image.

³Lexitek Inc. Phase Plates – <https://lexitek.com>

⁴pyRTC on Github: <https://github.com/jacotay7/pyRTC>

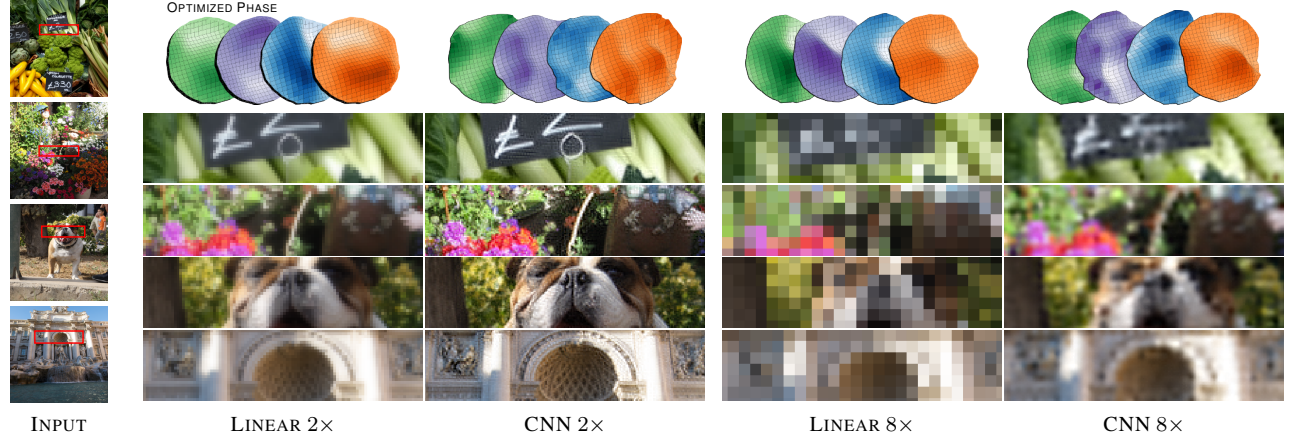


Figure 3. *Simulated SR results.* Comparisons of reconstructions with super resolution factors of $2\times$ and $8\times$ (inset regions marked in red). Images were generated with $N = 4$ exposures and phases were optimized for each method and SR factor. For a full view of the results, please refer to the supplemental material [55]




Figure 4. *Improved SR with jointly-optimized phase and CNN on simulated USAF resolution target.* End-to-end training of both phase and CNN yields superior resolution improvements compared to a naïve approach which does not optimize phases. The input to both CNNs is the inset region of the USAF target, marked in red. Qualitatively, by jointly optimizing the phase and the CNN, we can resolve smaller features and, quantitatively, achieve lower RMSE (displayed above each image).

6.1. Results

A green LED (centered at 530 nm) was used to project photographic slides on the optical bench detailed above. To demonstrate our method’s resolving power, we imaged a 1” US Air Force resolution chart (Thorlabs Positive 1951 USAF Test Target Groups 2-7, $\emptyset 1''$). Results are shown for both simulation (Fig. 4) and experiment (Fig. 6). We also show our method’s performance on more general scenes using a variety of real-world photos in Fig. 6. Because our AO system was designed specifically for sensing in infrared wavelengths, the quality of our images (captured at 635 nm) is limited by chromatic aberration and vignetting. As such, results in this paper are cropped and masked for clarity; raw images are provided in the supplemental materials [55].

7. Discussion

MTF Analysis: The modulation transfer function (MTF) is a measure of the achievable contrast for an optical system




Figure 5. *Optical setup.* We present a top-down view of the optical bench where our experiments were conducted, labeling both the wavefront sensor path and the science (image) paths.

over a range of spatial frequencies [44]. An ideal lens will linearly decrease from an MTF of 1 down to 0 as the spa-




Figure 6. *Experimental results.* We reconstruct $8\times$ upsampled images of the USAF resolution target and three natural images (Grapes, Bookshelf, Trees) from low resolution images (LR Input) captured with the optical setup depicted in Fig. 5. We present results using both linear and CNN method, each using their own optimized phase profiles. Both result in higher resolutions than the original input, as seen quantitatively on the USAF chart.

tial frequency increases and, in general, the closer a measured MTF behaves to this the better its resolving power. The MTF plots in Fig. 7 correspond to the USAF simulation of Fig. 4 and the real experiment of Fig. 6 computed using a knife-edge method [36]. Both show an increase of MTF across all spatial resolutions and thus demonstrate our method’s superiority (i.e., higher preserved contrast).




Figure 7. *Comparison of simulated and experimental MTFs.*

Phase Analysis: One advantage of representing phase profiles using Zernike modes is the ability to infer exactly which optical aberrations are most important to our reconstruction algorithms. In Fig. 8, we plot the total absolute power of the first nine basis modes identified by our method for both the CNN and the linear reconstructor, using a fixed number of exposures while varying the magnification factor. The optimized phases reveal insights. First, total modal power increases with the magnification order. Intuitively, this makes sense as more movement within a pixel is needed to sample higher resolutions. Second, CNNs apply roughly the same power across all modes for a given magnification while the linear method has more variance. Perhaps this is because linear solvers cannot leverage higher-frequency spatial information in the image as well as CNNs can.




Figure 8. *Comparison of modal powers used to induce SR.* We show modal power in the Zernike basis to explore the optimized phases (corresponding to the phases shown in Fig. 3). The linear method uses low-order, symmetric modes while the CNN spreads across all modes. Comparing $8\times$ vs. $2\times$ magnifications, they take opposite approaches: the CNN prefers low modes at higher resolutions, while the linear uses lower modes at low resolutions.

Limitations & Future Work: Currently, our work is limited to telescopes with a single-conjugate adaptive optics (SCAO) system, which treats the atmosphere as a single column. Many of today’s observatories instead use multiple WFS and DMs. This leads to more accurate corrections across a wider FoV and can even be tuned to different atmospheric layers. Such systems do not preclude the use of our method, since our method can be extended to multiple DMs and potentially enhanced by their extra degrees of freedom.

Other directions for future work include real-time viewing of video from the science camera by cycling through each phase multiple times to create individually super-resolved frames. Once extended to multiple DMs, additional multiplexing could also be achieved by varying each DM at different rates.

Most importantly, now that we have demonstrated our methodology with an optical bench prototype, our current priority is to show its effectiveness on sky with a science class telescope and AO system.

Conclusion: This work shows that science images can encode sufficient MISR information using the AO systems already present in state-of-the-art telescope observatories. We achieve resolution gains without affecting the scientific data in both simulations and optical-bench experiments. With this method, we aim to unlock new scientific capabilities at existing observatories without altering the optical system.

Acknowledgements

RS acknowledges NVIDIA for GPU hardware support. EYHL acknowledges the support of NSERC under CGSD. SS acknowledges the support of the Canada Foundation for Innovation and the Ontario Research Fund. KNK acknowledges the support of NSERC under the RGPIN and RTI programs.

References

- [1] Eiríkur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 126–135, 2017. 5
- [2] Saeed Anwar, Salman Khan, and Nick Barnes. A deep journey into super-resolution: A survey. *ACM Computing Surveys (CSUR)*, 53(3):1–34, 2020. 3, 4
- [3] Roberto J. Avila, Warren J. Hack, and Stsci Team. Astrodrizzle: Aligning images from multiple instruments. In *American Astronomical Society*, 2012. 4
- [4] Thilo Bauer. Super-resolution imaging: The use case of optical astronomy. In *Proceedings of the IADIS International Conference Computer Graphics, Visualization, Computer Vision and Image Processing, Rome*, pages 49–59, 2011. 2
- [5] Jacques M Beckers. Adaptive optics for astronomy-principles, performance, and applications. In: *Annual Review of Astronomy and Astrophysics. Vol. 31 (A94-I2726 02-90)*, p. 13-62., 31:13–62, 1993. 2
- [6] Anthony Berdeu, Michel Tallon, Éric Thiébaud, and Maud Langlois. Analytical modelling of adaptive optics systems: Role of the influence function. *Astronomy & Astrophysics*, 674:A112, 2023. 5
- [7] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Deep burst super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9209–9218, 2021. 4
- [8] David Borncamp, J. M. Anderson, Norman A. Grogin, and Warren J. Hack. Acs/wfc revised geometric distortion for drizzlepac. In *Instrument Science Report ACS/WFC*, 2015. 4
- [9] Praneeth Chakravarthula, Ethan Tseng, Tarun Srivastava, Henry Fuchs, and Felix Heide. Learned hardware-in-the-loop phase retrieval for holographic near-eye displays. *ACM Transactions on Graphics (TOG)*, 39(6):1–18, 2020. 4
- [10] Dorian Chan, Srinivas Narasimhan, and Matthew O’Toole. Holocurtains: Programming light curtains via binary holography. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 4
- [11] Stanley H Chan, Xiran Wang, and Omar A Elgendy. Plug-and-play admm for image restoration: Fixed-point convergence and applications. *IEEE Transactions on Computational Imaging*, 3(1):84–98, 2016. 5
- [12] Wenzheng Chen, Parsa Mirdehghan, Sanja Fidler, and Kiriakos N Kutulakos. Auto-tuning structured light by optical stochastic gradient descent. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5970–5980, 2020. 4
- [13] Xu Cheng, Arthur Bradley, Sowmya Ravikumar, and Larry N Thibos. Visual impact of zernike and seidel forms of monochromatic aberrations. *Optometry and Vision Science*, 87(5):300–312, 2010. 5
- [14] Rodolphe Conan and C Correia. Object-oriented matlab adaptive optics toolbox. In *Adaptive Optics Systems IV*, page 91486C. International Society for Optics and Photonics, 2014. 6
- [15] Sony Corporation. Pixel Shift Multi Shooting | SONY — support.d-imaging.sony.co.jp. <https://support.d-imaging.sony.co.jp/support/ilc/psms/ilce7rm3/en/index.html>, 2025. [Accessed 30-03-2025]. 4
- [16] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE Signal Processing Magazine*, 35(1):53–65, 2018. 5
- [17] Benjamin P Cumming and Min Gu. Direct determination of aberration functions in microscopy by an artificial neural network. *Optics Express*, 28(10):14511–14521, 2020. 4
- [18] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *Computer Vision – ECCV 2014*, pages 184–199, Cham, 2014. Springer International Publishing. 3
- [19] Theodore B DuBose, Dennis F Gardner, and Abbie T Watnik. Intensity-enhanced deep network wavefront reconstruction in shack–hartmann sensors. *Optics Letters*, 45(7):1699–1702, 2020. 4
- [20] Mark A Ealey and John F Washeba. Continuous facesheet low voltage deformable mirrors. *Optical Engineering*, 29(10):1191–1198, 1990. 2
- [21] Linjing Fang, Fred Monroe, Sammy Weiser Novak, Lindsey Kirk, Cara R Schiavon, Seungyoon B Yu, Tong Zhang, Melissa Wu, Kyle Kastner, Alaa Abdel Latif, et al. Deep learning-based point-scanning super-resolution imaging. *Nature Methods*, 18(4):406–416, 2021. 4
- [22] Sina Farsiu, Dirk Robinson, Michael Elad, and Peyman Milanfar. Robust shift and add approach to superresolution. In *Applications of Digital Image Processing XXVI*, pages 121–130. SPIE, 2003. 5
- [23] AS Fruchter and RN Hook. Drizzle: A method for the linear reconstruction of undersampled images. *Publications of the Astronomical Society of the Pacific*, 114(792):144, 2002. 1, 2, 4, 5
- [24] Andrew S. Fruchter. A new method for band-limited imaging with undersampled detectors. *Publications of the Astronomical Society of the Pacific*, 123(902):497, 2011. 4
- [25] Andrea M Ghez, BL Klein, M Morris, and EE Becklin. High proper-motion stars in the vicinity of sagittarius a*: Evidence for a supermassive black hole at the center of our galaxy. *The Astrophysical Journal*, 509(2):678, 1998. 3
- [26] Roberto Gilmozzi and Jason Spyromilio. The european extremely large telescope (e-elt). *The Messenger*, 127(11):3, 2007. 3
- [27] S Gonzaga, W Hack, A Fruchter, and J Mack. The drizzlepac handbook. *The DrizzlePac Handbook*, 2012. 4
- [28] Karen M Hampson, Raphaël Turcotte, Donald T Miller, Kazuhiro Kurokawa, Jared R Males, Na Ji, and Martin J Booth. Adaptive optics for high-resolution imaging. *Nature Reviews Methods Primers*, 1(1):1–26, 2021. 3
- [29] James E Harvey and Christ Ftaclas. Diffraction effects of telescope secondary mirror spiders on various image-quality criteria. *Applied Optics*, 34(28):6337–6349, 1995. 3
- [30] Hamid Reza Vaezi Joze, Ilya Zharkov, Karlton Powell, Carl Ringle, Luming Liang, Andy Roulston, Moshe Lutz,

- and Vivek Pradeep. Imagepairs: Realistic super resolution dataset via beam splitter camera rig. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 518–519, 2020. 5
- [31] SR Kulkarni. Instruments on large optical telescopes—a case study. *arXiv preprint arXiv:1606.06674*, 2016. 2
- [32] Vasudevan Lakshminarayanan and Andre Fleck. Zernike polynomials: a guide. *Journal of Modern Optics*, 58(7):545–561, 2011. 5
- [33] Rico Landman, Sebastiaan Y Haffert, Vikram M Radhakrishnan, and Christoph U Keller. Self-optimizing adaptive optics control with reinforcement learning. In *Adaptive Optics Systems VII*, page 1144849. International Society for Optics and Photonics, 2020. 4
- [34] Tod R Lauer. Combining undersampled dithered images. *Publications of the Astronomical Society of the Pacific*, 111(756):227, 1999. 2
- [35] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network, 2017. 3
- [36] Hang Li, Changxiang Yan, and Jianbing Shao. Measurement of the modulation transfer function of infrared imaging system by modified slant edge method. *Journal of the Optical Society of Korea*, 20(3):381–388, 2016. 8
- [37] Haoying Li, Yifan Yang, Meng Chang, Shiqi Chen, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing*, 479:47–59, 2022. 5
- [38] Orly Liba, Kiran Murthy, Yun-Ta Tsai, Tim Brooks, Tianfan Xue, Nikhil Karnad, Qirui He, Jonathan T Barron, Dillon Sharlet, Ryan Geiss, et al. Handheld mobile photography in very low light. *ACM Transactions on Graphics (TOG)*, 38(6):1–16, 2019. 4
- [39] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 136–144, 2017. 5
- [40] Zhisheng Lu, Juncheng Li, Hong Liu, Chaoyan Huang, Linlin Zhang, and Tiejong Zeng. Transformer for single image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 457–466, 2022. 5
- [41] Bahattin Can Maral. Single image super-resolution methods: A survey. *arXiv preprint arXiv:2202.11763*, 2022. 3
- [42] Christian Marois, Bruce Macintosh, Travis Barman, B Zuckerman, Inseok Song, Jennifer Patience, David Lafrenière, and René Doyon. Direct imaging of multiple planets orbiting the star hr 8799. *science*, 322(5906):1348–1352, 2008. 3
- [43] Y Mellier, Abdurrouf Abdurrouf, JA Acevedo Barroso, A Achúcarro, J Adamek, R Adam, GE Addison, N Aghanim, M Agüena, V Ajani, et al. Euclid. i. overview of the euclid mission. *Astronomy & Astrophysics*, 2024. 2
- [44] HH Nasse. How to read mtf curves. *Carl Zeiss, Camera Lens Division*, page 33, 2008. 7
- [45] Ngoc Long Nguyen, Jérémy Anger, Axel Davy, Pablo Arias, and Gabriele Facciolo. Self-supervised multi-image super-resolution for push-frame satellite images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1121–1131, 2021. 4
- [46] Sylvain Oberti, Carlos Correia, Thierry Fusco, Benoit Neichel, and Pierre Guiraud. Super-resolution wavefront reconstruction. *Astronomy & Astrophysics*, 667:A48, 2022. 4
- [47] Scott W Paine and James R Fienup. Machine learning for improved image-based wavefront sensing. *Optics Letters*, 43(6):1235–1238, 2018. 4
- [48] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 5
- [49] Yifan Peng, Suyeon Choi, Nitish Padmanaban, and Gordon Wetzstein. Neural holography with camera-in-the-loop training. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020. 4
- [50] Ben C Platt and Roland Shack. History and principles of shack-hartmann wavefront sensing, 2001. 2
- [51] Roberto Ragazzoni, Emiliano Diolaiti, and Elise Vernet. A pyramid wavefront sensor with no dynamic modulation. *Optics Communications*, 208(1-3):51–60, 2002. 2
- [52] Vincent Sitzmann, Steven Diamond, Yifan Peng, Xiong Dun, Stephen Boyd, Wolfgang Heidrich, Felix Heide, and Gordon Wetzstein. End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018. 4
- [53] Robin Swanson, Masen Lamb, Carlos M Correia, Suresh Sivanandam, and Kiriakos Kutulakos. Closed loop predictive control of adaptive optics systems with convolutional neural networks. *Monthly Notices of the Royal Astronomical Society*, 503(2):2944–2954, 2021. 4
- [54] Robin Swanson, Suresh Sivanandam, and Kiriakos N Kutulakos. Adaptive optics mediated sub-pixel super resolution. In *Adaptive Optics Systems VIII*, pages 2594–2600. SPIE, 2022. 4
- [55] Robin Swanson, Esther Y. H. Lin, Masen Lamb, Suresh Sivanandam, and Kiriakos N. Kutulakos. Super resolved imaging with adaptive optics: Supplemental document. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. 7
- [56] Jing Tian and Kai-Kuang Ma. A survey on super-resolution imaging. *Signal, Image and Video Processing*, 5(3):329–342, 2011. 3
- [57] Ethan Tseng, Shane Colburn, James Whitehead, Luocheng Huang, Seung-Hwan Baek, Arka Majumdar, and Felix Heide. Neural nano-optics for high-quality thin lens imaging. *Nature Communications*, 12(1):1–7, 2021. 4
- [58] Robert K Tyson. *Adaptive ptics engineering handbook*. Marcel Dekker New York, 2000. 2
- [59] Robert K Tyson and Benjamin West Frazier. *Principles of adaptive optics*. CRC press, 2022. 3

- [60] Congli Wang, Qiang Fu, Xiong Dun, and Wolfgang Heidrich. Megapixel adaptive optics: towards correcting large-scale distortions in computational cameras. *ACM Transactions on Graphics (TOG)*, 37(4):1–12, 2018. [4](#)
- [61] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726*, 2020. [2](#)
- [62] Bartłomiej Wronski, Ignacio Garcia-Dorado, Manfred Ernst, Damien Kelly, Michael Krainin, Chia-Kai Liang, Marc Levoy, and Peyman Milanfar. Handheld multi-frame super-resolution. *ACM Transactions on Graphics (TOG)*, 38(4): 1–18, 2019. [4](#)
- [63] Haitang Yang, Esther Y. H. Lin, Kiriakos N. Kutulakos, and George V. Eleftheriades. Sub-wavelength passive single-shot computational super-oscillatory imaging. *Optica*, 9(12): 1444–1447, 2022. [4](#)
- [64] Wenming Yang, Xuechen Zhang, Yapeng Tian, Wei Wang, Jing-Hao Xue, and Qingmin Liao. Deep learning for single image super-resolution: A brief review. *IEEE Transactions on Multimedia*, 21(12):3106–3121, 2019. [3](#)
- [65] Linwei Yue, Huanfeng Shen, Jie Li, Qiangqiang Yuan, Hongyan Zhang, and Liangpei Zhang. Image super-resolution: The techniques, applications, and future. *Signal Processing*, 128:389–408, 2016. [3](#)
- [66] Biwei Zhang, Jiazhu Zhu, Ke Si, and Wei Gong. Deep learning assisted zonal adaptive aberration correction. *Frontiers in Physics*, page 634, 2021. [4](#)
- [67] Hao Zhu, Mantang Guo, Hongdong Li, Qing Wang, and Antonio Robles-Kelly. Revisiting spatio-angular trade-off in light field cameras and extended applications in super-resolution. *IEEE Transactions on Visualization and Computer Graphics*, 27(6):3019–3033, 2019. [2](#)