

Ring Deconvolution Microscopy: An Exact Solution for Spatially-Varying Aberration Correction

Amit Kohli*, Anastasios N. Angelopoulos*, David McAllister,
Esther Whang, Sixian You, Kyrollos Yanny, Laura Waller

August 14, 2023

Abstract

The most ubiquitous form of computational aberration correction for microscopy is deconvolution. However, deconvolution relies on the assumption that the point spread function is the same across the entire field-of-view. It is well recognized that this assumption is often inadequate, but space-variant deblurring techniques generally require impractical amounts of calibration and computation. We present a new imaging pipeline, *ring deconvolution microscopy* (RDM), that leverages the rotational symmetry of most optical systems about their optical axis to provide simple and fast spatially-varying aberration correction. We formally derive theory and algorithms for exact image recovery and additionally propose a neural network based on Seidel coefficients as a fast alternative. We demonstrate significant visual and quantitative improvements over standard deconvolution in simulation and experimentally on images obtained using the UCLA Miniscope.

1 Introduction

Much of optical engineering is focused on reducing aberrations by adding additional corrective optical elements to an imaging system—consider a microscope objective, comprised of numerous lenses stacked in a housing. While such designs allow for high-performance imaging, they also incur added cost, weight, and complication. Even with very large and expensive lens stacks, it is difficult and, in some cases, impossible to correct all aberrations across a large area, and so aberrations are often what limit the usable field-of-view (FoV) of a system. Furthermore, some systems cannot accommodate any aberration-correction optics; for example, additional elements may not fit in miniature microscopes (e.g. Miniscope [1]) and are prohibitively expensive for large-aperture telescopes [2].

Faced with a poorly-corrected imaging system, the modern microscopist must instead turn to computational aberration correction, where the burden is shifted onto computer algorithms applied after the image is acquired. The most commonly used correction technique, image deconvolution, captures a calibration image of a small point-like source or bead, known as the point spread function (PSF), in order to characterize the aberrations. The PSF can then be used to computationally deconvolve any image taken with the system via simple and fast algorithms, to yield a deblurred result. A main limitation of this approach is that it assumes that the system’s PSF does not vary spatially (i.e. the system is linear space-invariant, LSI). Paradoxically, LSI systems are often close to ideal, which means that systems that are in most need of deconvolution benefit the least from it.

To bypass the space-invariant limitations, a large community effort has gone towards heuristic forms of spatially-varying deconvolution, wherein one measures PSFs at multiple points in the FoV and uses them to correct the image. Such heuristics include: assuming each region of an image is locally LSI [3], adaptively splitting the FoV by first quantifying the degree of space-variance [4, 5], interpolating PSFs [6, 7], decomposing the PSF into space-invariant orthogonal modes [8–14], and doing the same in Fourier space [15]. As the number of PSFs collected grows—into the hundreds of thousands—these heuristics approach exact

*These authors contributed equally to this work. Code available here

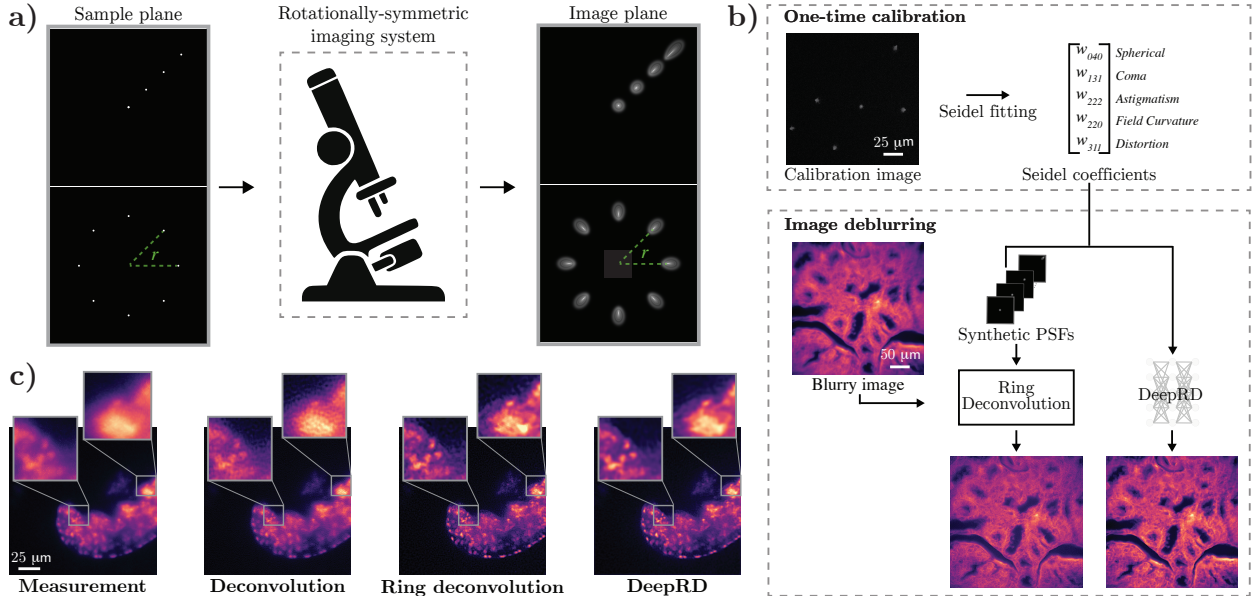


Figure 1: **Ring Deconvolution Microscopy (RDM)**. **a)** Point sources at the sample plane (left) are imaged (right) to point spread functions (PSFs) with a rotationally symmetric imaging system. The PSFs are linear revolution-invariant (LRI)—they vary with distance from the center of the field-of-view (FoV) (top row), but maintain the same shape at a fixed radius r , just revolved around the center (bottom row). **b)** The RDM pipeline. A one-time calibration procedure (top) captures a single image of randomly-placed point sources (e.g., fluorescent beads), then fits the primary Seidel coefficients (see Sec. 4.2). Next, we either use the Seidel coefficients to generate a radial line of synthetic PSFs, if using ring deconvolution, or we feed the coefficients directly into Deep Ring Deconvolution (DeepRD). After calibration, we can deblur images (bottom) using either ring deconvolution or DeepRD. **c)** Experimental deblurring of live tardigrade samples imaged with the UCLA Miniscope [1]. From left to right: measurement, standard deconvolution, ring deconvolution, and DeepRD. Ring deconvolution and DeepRD consistently outperform standard deconvolution.

recovery. But there is also a tradeoff in terms of the complexity of calibration and computation, which quickly becomes intractable. For example, in patch-wise deconvolution, the FoV is divided into patches, each of which gets deconvolved by a PSF measured at its center. The maximum accuracy is achieved when the patch size is reduced down to a single pixel, resulting in a calibration of a million PSFs and a computation time of hundreds of hours (see Fig. 3c) to deblur a 1 megapixel image.

Another emerging modality is deep deblurring [16–22], in which varying amounts of system information are incorporated into a deep neural network. On one end of the spectrum, networks that are primarily data-driven struggle with extrapolation beyond the training data, and tend to reproduce whatever biases existed therein, a particularly relevant point since many of them are trained on simulated data. On the other end, networks that incorporate physical information, such as PSFs, may have better generalization properties but suffer the same accuracy/efficiency tradeoff as the patch-based methods. For these reasons, spatially-varying deblurring has not become commonplace among practitioners. Thus, there is a need for a new method of calibrating and deblurring spatially-varying imaging systems which is effective, efficient, and robust.

Here, we explore an in-between regime where the PSFs can vary, but only radially. Such a setting occurs when the system is *rotationally symmetric*—i.e., symmetric about its optical axis. This symmetry occurs in almost all imaging systems by design, and a significant portion of optical theory is developed under this assumption. While some existing deblurring techniques have leveraged this symmetry for modeling motion blur [23, 24], for suppressing radial variance [25], and by assuming the FoV can be divided into a few (3 or 4) space-invariant radial segments [26–30], there still does not exist a general-purpose, theoretically exact approach for deblurring rotationally symmetric imaging systems.

To that end, we propose Ring Deconvolution Microscopy (RDM), a new imaging pipeline that achieves the

best of both accuracy and efficiency. We rigorously prove that RDM enables exact recovery for a rotationally symmetric imaging system, but it is also practical, both computationally and in terms of calibration. The first step in RDM is a simple, single-shot calibration scheme, wherein the system’s primary Seidel aberration coefficients are estimated from a single image of randomly-distributed point sources. These coefficients quantify the severity of spatial variance and provide the necessary system information for the second step, deblurring. Within this second step there are two proposed image deblurring algorithms. Our first and main algorithm, ring deconvolution, uses a new and exact theory for rotationally symmetric imaging to deblur the image at all points in the FoV, with only order $N^3 \log(N)$ (N is the image side length) compute time, as compared to N^4 for full spatially-varying deblurring. For even faster computation, but without guarantees on out-of-distribution accuracy, we further propose an alternate neural network-based algorithm called Deep Ring Deconvolution (DeepRD), which constrains learning with physical knowledge provided by the system’s Seidel aberration coefficients.

These algorithms outperform standard deconvolution and approach isotropic performance across the field. See Fig. 1 for a summary of the pipeline along with an example result on images of live tardigrades. An open source implementation of RDM can be found here.

2 Results

2.1 Single-shot calibration

The system is calibrated by measuring its response to a point source, i.e., its PSF. The PSFs of space-varying systems will be different at different parts of the FoV, and many PSF measurements may be required to fully characterize the system. RDM leverages the fact that rotationally symmetric imaging systems require fewer measurements for exact system characterization.

The reduction in complexity occurs because PSFs that are the same distance from the center of the FoV will all be the same shape, but rotated at different angles (see Figure 1a). We call this property of the PSFs *linear revolution-invariance* (LRI), and denote it mathematically as

$$\tilde{h}(\rho, \phi; r, \theta) = \tilde{h}(\rho, \phi - \theta; r, 0), \tag{1}$$

where $\tilde{h}(\rho, \phi; r, \theta)$ is the (spatially-varying) PSF in polar coordinates from a point source at location (r, θ) . Note that (1) does not necessarily mean that the shape of the PSF itself is rotationally symmetric. Intuitively, each point along a concentric ring in the FoV will have the same PSF shape, and the system can be exactly characterized by measuring the PSF only once for each radius r from the optical center.

In practice, directly measuring the PSF at every radius r is impractical. Instead, under LRI, there is a simple and effective method for simultaneously estimating these PSFs from a single image. The method is described in Section 4.2 and is the first step of the RDM pipeline. It works by estimating the Seidel aberration coefficients from an image of point sources randomly scattered across the field-of-view (see Fig. 1b). No motion stage is needed. The Seidel aberrations are similar to the Zernike modes [31, 32] but much simpler to estimate—Seidel polynomials are explicitly functions of the field position and thus use the same coefficients regardless of the position in the FoV [33]. Meanwhile, Zernike modes are a general-purpose spherical basis set, which is not a function of field position and whose coefficients are therefore different at each different field position. Our experiments demonstrate that the 5 primary (3rd order) coefficients (commonly called sphere, coma, astigmatism, field curvature, and distortion) suffice to characterize spatially-varying, LRI systems.

2.2 Reconstruction algorithm

After calibration, the next step is to use the estimated Seidel coefficients to algorithmically deblur the blurry measured image. LRI systems, much like LSI systems, allow for computationally efficient and exact forward and inverse models. Section 4.1 provides a description of these algorithms, which we call *ring convolution* and *ring deconvolution*, respectively. Therein is also included a theoretical proof of their exactness. We briefly describe ring deconvolution here as our method for aberration correction, along with *DeepRD*, its extension to deep learning. These methods form the second step in the RDM pipeline (Fig. 1b).

1. **Ring deconvolution.** We derive a provably exact algorithm for reconstructing the underlying sample from a blurry image given the PSF at each radius of the FoV of a rotationally symmetric imaging system. These PSFs are efficiently obtained via the single-shot calibration procedure described above. This is our main algorithm.
2. **Deep Ring Deconvolution (DeepRD).** Although ring deconvolution is significantly faster to compute than a full patch-based spatially-varying deblur technique, it may still be relatively slow (on the order of a few minutes) for very large image sizes (e.g., beyond 1024×1024) or video data. We used deep learning to perform a fast (but approximate) version of ring deconvolution called *DeepRD*. As input, it takes an blurry image and a list of the 5 primary Seidel coefficients. DeepRD is trained on a dataset of natural images synthetically blurred using ring convolution. The experimental images are completely out of the training domain, yet the algorithm generalizes effectively.

Even when the system is LSI or close to LSI, RDM still offers an improvement over standard deconvolution by instead deconvolving with a synthetic PSF generated by the Seidel coefficients (see Sec. 2.5).

2.3 Deblurring of fluorescent images from a miniature microscope

The first step of our experimental pipeline is to calibrate the imaging system, which in this case is the UCLA Miniscope [1]. We take a single image of randomly-placed fluorescent beads (obtained by smearing bead solution on a slide) and use it to fit Seidel coefficients. We obtain 0.7, 0.3, 0.6, 0.2, and 0 waves of spherical aberration, coma, astigmatism, field curvature and distortion, respectively. These numbers, while specific to our particular assembly of the Miniscope, are consistent with aberration profile of a radial GRIN lens [34], which is the objective lens used by the Miniscope. Radial GRIN lenses are dominated mostly by spherical aberration followed by astigmatism, field curvature, and a small amount of coma. The fact that the off-axis coefficients (i.e. all of the primary coefficients except for spherical) are nonzero confirms that the system is indeed spatially-varying and should benefit from ring deconvolution. For comparison, we also run the standard deconvolution calibration, in which a single fluorescent bead is isolated through repeated dilutions of bead solution on a slide and centered in the FoV using a high-resolution Prior motion stage, thus allowing us to measure only the center PSF. The PSF is then denoised before its use in deconvolution (see 4.5).

Having completed the one-time calibration for the Miniscope, we image a variety of different samples: a USAF resolution target, rabbit liver tissue, and live fluorescence-stained tardigrades. For each sample, we run ring deconvolution and DeepRD, and show a comparison with standard deconvolution (using a single PSF measured at the center of the FoV) and a baseline U-Net (see Fig. 2). We find that ring convolution is consistently the best method, a result that reflects the fact that it is theoretically exact for rotationally symmetric systems and is not biased toward any particular distribution of images. DeepRD performs similarly to ring deconvolution, but it is faster and less consistent. Both perform better than the U-Net and deconvolution in all of our examples.

It is worth noting that, though the baseline U-Net we use is inspired by the popular Content-Aware Image Restoration (CARE) model [20], we train both DeepRD and the baseline U-Net using a novel synthetic training strategy: applying ring convolutions using PSFs generated from random Seidel coefficients to natural images from Div2k [35](see Sec. 4.3). We will release this synthetic dataset upon publication. Despite being trained on simulated data, both the U-Net and DeepRD generalize surprisingly well to the images taken with the Miniscope.

Further experimental details can be found in Sec. 4.5, including (in Appendix A) a demonstration that DeepRD is interpretable: it understands the individual effects of each aberration coefficient.

USAF resolution target. Our first sample is a USAF resolution target placed at 9 separate locations in the FoV. We stitch these 9 images together, using only the region of each constituent image that contains the high resolution group, to form the measurement image shown in Fig. 2 (first row). Ring deconvolution and DeepRD—having knowledge of the field-varying aberrations via the Seidel coefficients—give the most improvement near the edges and corners of the image. Both methods allow for more consistent image contrast across the field, unlike the other methods, which degrade away from the center. The raw deconvolution produces a noisy, low contrast result in all cases because it uses a measured PSF as opposed to the RDM methods, which benefit from synthetic, noiseless PSFs. We also note that both of the learning methods (U-Net and DeepRD) offer the best denoising performance, a well-known property of neural networks [36];




Figure 2: **Deblurring of Miniscope images.** After calibrating the Miniscope with a single image of fluorescent beads (see Fig. 1b), we estimate the Seidel coefficients. We then show the measurement and results from several deblurring methods for comparison: standard deconvolution (using a single measured PSF), a U-Net trained on our spatially-varying blur dataset, and our methods: ring deconvolution and DeepRD. Deconvolution assumes space-invariance while the remaining methods are designed to handle spatially-varying aberrations. In the first row, ring deconvolution and DeepRD can resolve resolution target elements near the edges of the FoV, which are not resolved by the other methods. Zoom-ins show RDM resolves up to element 6 of group 9 (blue inset) and element 5 of group 8 (green inset). In the middle row, both RDM methods more sharply resolve membranes in the rabbit liver tissue (blue inset) and reveal a hook-shaped feature not seen using the other methods (green inset). In the bottom row, our methods are also able to resolve features on the tardigrade not well resolved by the other methods (yellow inset) and improve the contrast compared to deconvolution throughout the image (green and blue insets).

however, this comes at the cost of inconsistent performance—both models perform worse on the resolution target than on the other samples.

Rabbit liver tissue. The RDM methods perform well on the rabbit liver tissue, revealing features in the corners of the image including the outlines of membranes, which are not clear in the raw measurement or after the other methods. Such clarity is critical for downstream tasks like image segmentation.

Live tardigrades. For our final experiment, we fluorescently stain live, semi-starved tardigrades with a DNA gel stain, and capture a series of videos. We apply deblurring to each frame and display one such frame in the bottom row of Fig. 2; the full videos can be found here. As with the previous experiments, ring deconvolution and DeepRD provide increased image contrast and detail in the corners of the frames as compared to the other methods. In particular, the small, dot-like features within the tardigrade are better resolved.

2.4 Quantifying performance improvements in simulation

To quantify the performance of our methods, we conducted a series of simulations in which we have access to the unblurred ground-truth image. First, we verify the accuracy of ring convolution as a forward model. Second, we quantify the accuracy of fitting Seidel coefficients to a single calibration image, even in high-noise conditions. Finally, we demonstrate the effectiveness of ring deconvolution and DeepRD in deblurring images from a spatially-varying system.

2.4.1 Forward model: Ring convolution

Simulating the blurring operation of an imaging system is a critical part of image deblurring. It mathematically specifies the blur operator of the system, which can then be inverted in order to obtain the desired deblurring algorithm. Ring convolution produces a near exact simulation of system blur, even when the PSFs are significantly spatially-varying.

To verify, we blur each test image using spatially-varying PSFs rendered from a randomly chosen set of Seidel coefficients. For the ground truth (true blur), we superimpose the PSFs at each pixel in the image. The top left panel of Fig. 3 compares the error maps for two forward models: traditional LSI convolution and our ring convolution. Traditional convolution gives some error near the edges for the case of small aberrations, and much more error for the case of large aberrations (Fig. 3a). Since the severity of spatial variance is quantified by the norm of the off-axis Seidel coefficients, the error of standard convolution scales up approximately linearly. In contrast, ring convolution continues to produce an accurate blur, independent of the strength of the aberrations (Fig. 3b). While the ring deconvolution error is not exactly 0 (due to discretization effects, Sec. 4.1) it is negligible compared to the deconvolution error. In Fig. 3c, we compare compute times for these forward models, showing that our ring convolution method is significantly faster than the other exact method (true blur).

These results are consistent with theory, which maintains that while standard convolution is fast, it is generally inaccurate; meanwhile the true blur is exact but prohibitively slow. Ring convolution is the best of both worlds: it is exact for any rotationally symmetric system while also running quickly—on the order of seconds—even for image sizes upward of 512×512 .

2.4.2 Seidel fitting

The next set of simulations verifies our ability to accurately fit Seidel coefficients from a single noisy image of randomly scattered point sources. As detailed in Sec. 4.2, the Seidel fitting procedure involves searching for the set of 5 primary Seidel coefficients that best fit the measured PSFs at their given positions in the calibration image. We simulate this process many times by randomly generating sets of Seidel coefficients, using them to produce a calibration image with additive Gaussian noise, and then estimating those coefficients only using the noisy calibration image (see example in Fig. 3d). Section 4.2 shows that this optimization problem is nonconvex, yet we find that our estimated Seidel coefficients almost always converge correctly. This is not entirely surprising since the Seidel procedure is optimal in the maximum likelihood sense [37]. To quantify, we ran many Seidel fits over random coefficients, randomly located PSFs, and random draws of additive noise, and plotted the error between estimated and true coefficients in Fig. 3e. Due to the nonconvexity, we see that not every case produces errors which converge to 0. However, two things provide

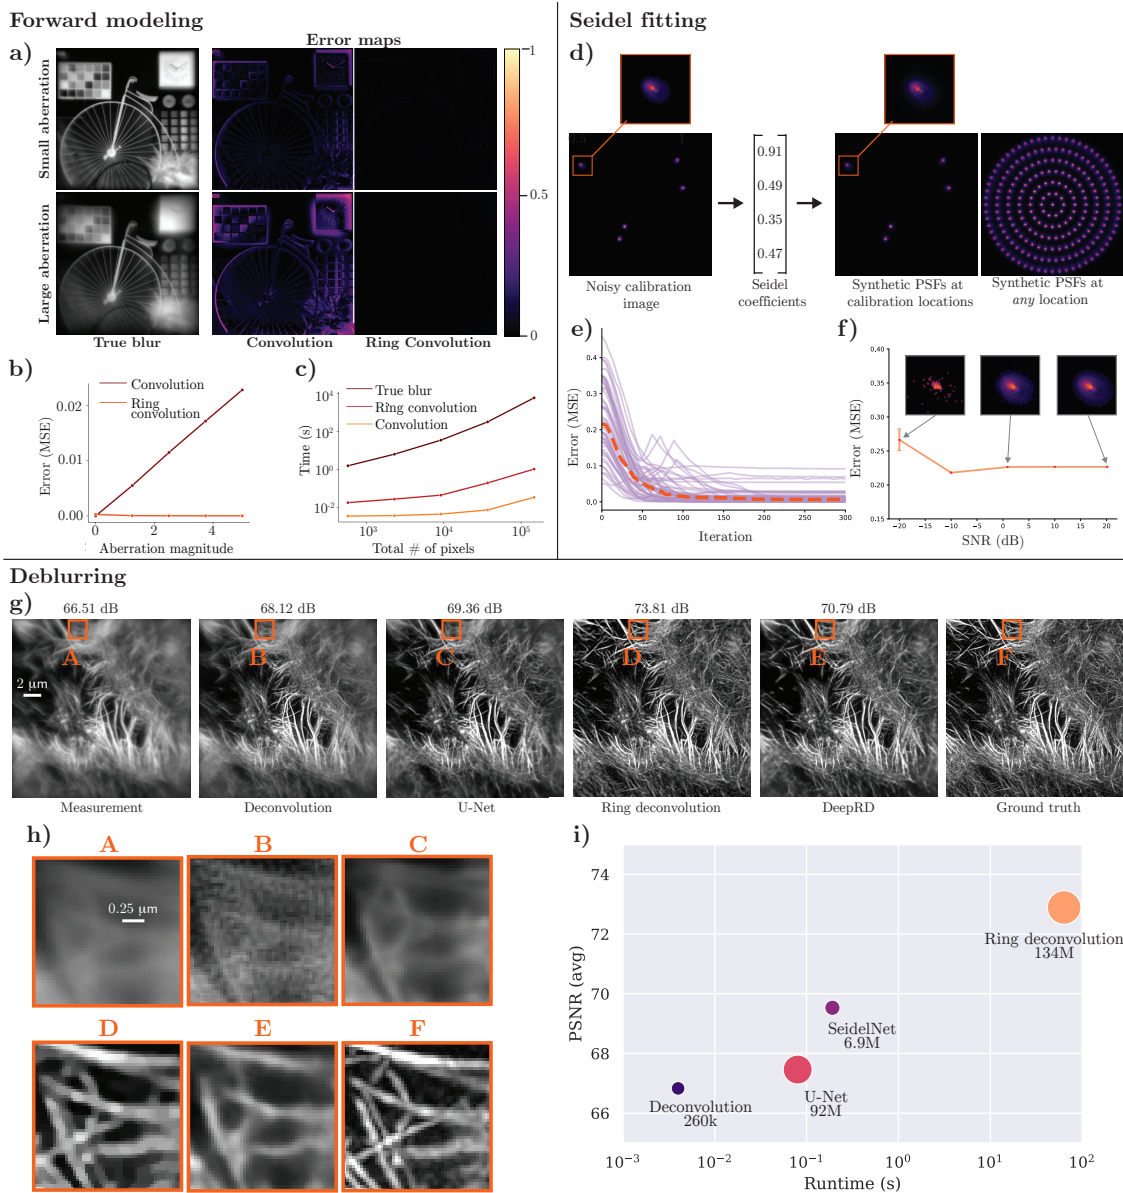


Figure 3: **Simulations to quantify RDM performance.** **a)** Error maps show the normalized absolute difference between the blur produced by deconvolution/ring deconvolution and the “true” blur (produced by manually superimposing every PSF at every pixel). When off-axis aberrations are small (top row), both forward models are accurate. When aberrations are large (bottom row), convolution becomes significantly worse, yet ring convolution remains accurate. **b)** Forward model mean squared error (MSE) as a function of off-axis aberration magnitude. **c)** Runtime of each method as a function of size of the image (in pixels); computing the true blur quickly becomes intractable, but convolution and ring convolution remain relatively fast. **d)** A noisy image of randomly placed PSFs is used to estimate the underlying Seidel coefficients, which can then be used to generate PSFs at any location. **e)** Seidel fitting error is plotted as a function of iteration in the optimization algorithm to demonstrate convergence. Each purple line is a different trial with a randomly sampled set of underlying Seidel coefficients. The red dashed line is the per-iteration median. **f)** Mean squared error of the fitted Seidel coefficients plotted as a function of SNR of the calibration image. The average coefficient error is plotted along with the variance (error bar) over 50 random trials. Some example calibration PSFs are shown. Even under severe noise, the Seidel fit is still accurate. **g)** Deblurring results on noisy images from the CARE dataset, with PSNR values above each method. **h)** Zoom-ins of an off-axis patch in each deblurred image; ring deconvolution (D) and DeepRD (E) have the highest quality. **i)** Accuracy (PSNR) vs runtime of each method, with the number of model parameters (written below each circle) determining the size of the circle.

assurance: 1) even the runs which do not converge to the global minimum produce PSFs close to the true ones, and 2) the median convergence approaches 0, meaning that a majority of optimizations will produce the optimal solution.

Finally, since the Seidel fitting procedure acts as a PSF denoiser, we tested the fit with varying amounts of noise (up to -20 dB SNR) and found that the fit only reduces performance slightly, even with severe additive noise (Fig. 3f).

2.4.3 Spatially-varying deblurring

After verifying that our forward model and Seidel fitting methods perform well, we next use them in an inverse problem to deblur images and quantify the performance. We compare our methods (Ring deconvolution and DeepRD) with traditional deconvolution and the baseline CARE-based U-Net. The learning methods are trained with images from the CARE dataset [20], after synthetically blurring them with space-varying PSFs rendered from a random set of Seidel coefficients. It was computationally infeasible to generate the training set using the true blur technique, so instead we used our ring convolution to generate the blurred images for training. This should not pose a major issue since, as the first simulation demonstrates, ring convolution is theoretically equivalent with the true blur. However, we do construct the test set using the true blur. Each model is first pretrained on natural images from the Div2k dataset [35, 38] and then fine-tuned on a small subset of the CARE dataset (84 images). After training, each method is tested on unseen images from the CARE dataset, which are blurred using the true blur method and noised with additive Gaussian noise.

The results of each method on one representative test image are shown in Fig. 3g,h, with the PSNR (in dB) listed above. Both DeepRD and ring deconvolution deblur better near the edges and corners of the image, where the PSF most deviates from the center PSF. Despite using ring convolution for the training set and the true blur for the test set, neither of the networks (U-Net and DeepRD) show signs of model mismatch.

In Fig. 3i, we plot the average accuracy (PSNR) vs runtime across the full test dataset, with the size of the circle representing the number of parameters needed (memory footprint) for each method. Ring deconvolution provides the best reconstruction, but it is also the largest and slowest of the methods tested (note that inverting the full blur would take hundreds of hours per image and is consequently not included). DeepRD performs nearly as well and has the fewest needed parameters of all the space-varying techniques, allowing it to be fast and memory efficient. The baseline U-Net and standard deconvolution PSNR is significantly worse.

2.5 Extensions and variations

Several extensions and variations of RDM, discussed below, may find use even when a system is actually LSI, when calibration data is not available, or if the system is spatially-varying, but not LRI.

Space-invariant systems

If aberrations are not space-variant, traditional deconvolution should perform as well as RDM and be more computationally efficient. However, the RDM pipeline can still provide value. When the system is space-invariant, only the spherical coefficient needs to be estimated from the calibration image via Seidel fitting. With this coefficient, we can generate a synthetic center PSF and perform deconvolution. We call this procedure *Seidel deconvolution*, and find that it essentially denoises the PSF measurement since it finds the closest synthetic PSF to the measured one. In the results shown in Fig. 4, Seidel deconvolution resolves smaller features and gives a cleaner reconstructed image as compared to traditional deconvolution.

Fitting synthetic PSFs for the purpose of deconvolution is not new; existing work has fit a variety of parametric models to the experimentally measured PSF including a 2D Gaussian distribution [39], Gaussian mixture model [40], Zernike basis [31, 32, 41–44], spherical aberration diffraction model [45], and Seidel ray model [46, 47], but to our knowledge, the fitting of Seidel polynomials in the pupil function is novel.

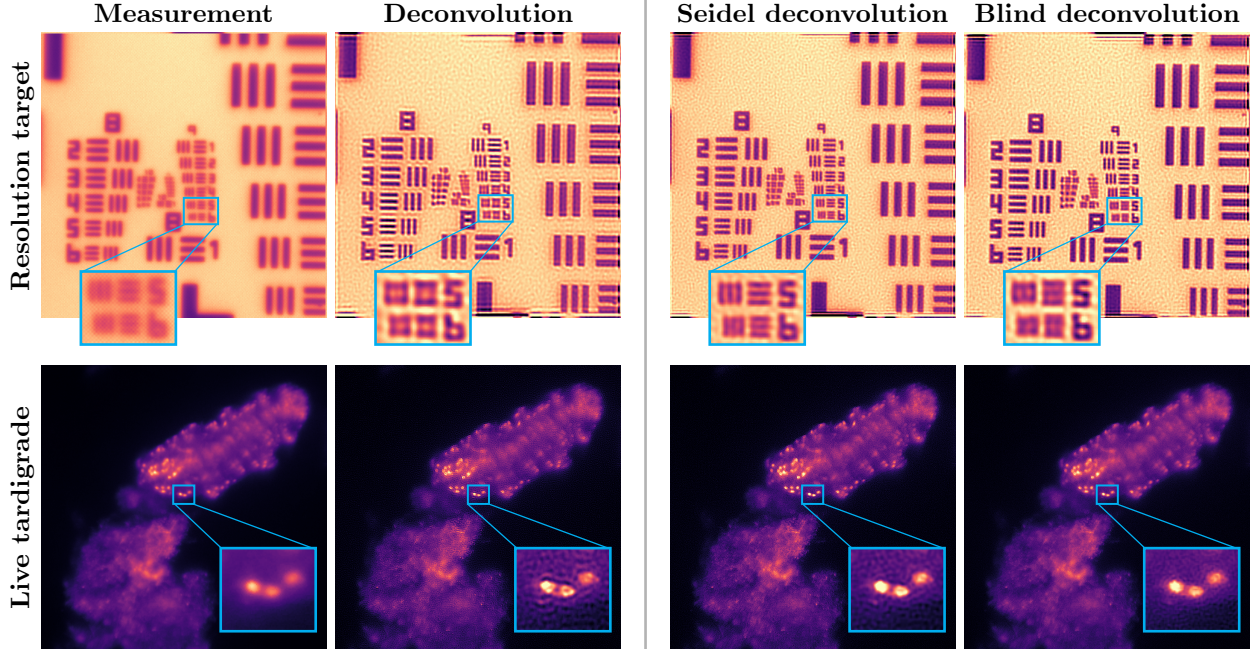


Figure 4: **RDM for shift-invariant systems.** We run our Seidel and blind deconvolution (left of gray line) on raw images (leftmost column) of tardigrades (bottom two rows) and rabbit liver tissue (top row). We also run standard deconvolution (2nd from the left) for comparison. Our methods outperform standard deconvolution by using a synthetic PSF, which prevents artifacts and loss of resolution from noise-based artifacts. In the first row, our Seidel deconvolution is able to resolve elements 5 and 6 of group 9 in the USAF resolution chart, which are not resolvable by the other methods. In the second row, both Seidel and blind deconvolution reduce ringing artifacts present in standard deconvolution. Our blind methods usually matches the non-blind Seidel deconvolution since it usually estimates the same spherical aberration coefficient; however, as in row 1, it can overestimate the coefficient resulting in an overly sharp image.

Blind deblurring

One intriguing extension of RDM is to perform deblurring without calibration data (*blind deblurring*). Such a situation may arise because measuring PSFs is difficult, inconvenient, or not available for images captured in the past. Blind deblurring involves joint estimation of the PSF and the deblurred object from a single image (see Methods 4.4).

In principal, blind deblurring can be applied to any spatially-varying system; however, it is computationally intensive, so we demonstrate here only the simple case of a shift-invariant system. Thus, only the spherical Seidel coefficient needs to be estimated. Results are shown in Fig. 4 along with non-blind deconvolution methods for comparison. The blind deconvolution result is similar to that of Seidel deconvolution, though cannot resolve the smallest features and has overly high-contrast due to noise. This can be mitigated by regularization.

Future work extending this idea to spatially-varying systems could use DeepRD, which is more computationally efficient than ring deconvolution, to iteratively search over the space of deblurring networks and choose the network with the sharpest reconstruction.

Non-LRI systems

RDM is designed for shift-varying systems that are linear revolution-invariant (LRI), but even when the system is not rotationally symmetric, it may still be useful. In such cases, RDM will not provide a complete recovery of the underlying object. Instead, it can undo the components of the blur that are revolution-invariant, which follows from the linearity of the PSF model. This may be a useful trade-off of performance and computational burden, given that full shift-varying deblurring methods are often computationally in-

tractable. It is also worth noting that the mathematical arguments made in Theorem 1 can be readily extended to other forms of symmetry, not just rotational symmetry. This could prove useful for reflection-mode microscopes, which often have a left-right symmetry due to reflection from a mirror, or even for light-sheet microscopes whose field is symmetrically aberrated due to the light sheet’s symmetry about the beam waist.

3 Discussion

We develop a new pipeline for image deblurring called Ring Deconvolution Microscopy (RDM) which encompasses both an accurate and analytically-derived deblurring technique, ring deconvolution, and a fast alternative, Deep Ring Deconvolution (DeepRD). Like standard deconvolution microscopy, our methods only require a single calibration image; however, they offer space-varying aberration correction. We back RDM with a new theory of imaging under rotational symmetry, which we call linear revolution-invariance (LRI), and an accurate implementation of the LRI forward model. For RDM calibration we also develop a procedure for fitting Seidel aberration coefficients from a single calibration image of randomly-placed point sources. Finally, we verify the accuracy of our RDM deblurring methods in both simulation and experiments using the UCLA Miniscope.

We hope that RDM will ultimately replace deconvolution microscopy as standard practice in widespread applications from biology to astronomy. The notion of rotational symmetry is deeply embedded in the theory of optics, and as such, we find it a natural and acceptable modeling assumption that can enable significant improvements in basic image processing. We believe that RDM will find most use in systems that approach optical extremes such as miniature microscopes, but may also empower optical designers to simplify hardware knowing they have the ability to better correct for aberrations digitally.

4 Methods

4.1 Ring convolution and ring deconvolution

We begin with a primer on notation. Let $g(u, v)$ describe the object’s intensity at (u, v) and $h(x, y; u, v)$ describe the space-varying PSF - the intensity at (x, y) of the PSF generated by a point source at (u, v) . We further use the notation \tilde{g} to denote the transformation of g to polar coordinates. Then the final image intensity $f(x, y)$ of a linear optical system is formed by the superposition integral [48]:

$$f(x, y) = \int \int g(u, v)h(x, y; u, v)dudv. \tag{2}$$

This equation is the system *forward model* for a linear space-varying system—it describes the image as a function of the object and PSFs at different locations in the FoV.

Traditional image deconvolution approximates the system as linear space-invariant (LSI), which means the PSF is the same at all positions in the FoV, $h(x, y; u, v) = h(x - u, y - v; 0, 0)$. This simplifies the forward model in Eq. (2) by reducing it to a convolution with a single PSF. This greatly simplifies the computation for forward and inverse problems, but at the cost of being inaccurate for space-varying aberrations.

In this paper, we incorporate radially-varying aberrations analytically into the forward model in order to provide a middle ground between purely space-invariant and completely space-varying systems. The assumption is that the system is *linear revolution invariant* (LRI) - i.e. its physical configuration is symmetric about the optical axis (see Fig. 1a). This is true of many typical optical imaging systems. Our core observation is that all LRI optical systems satisfy

$$\tilde{h}(\rho, \phi; r, \theta) = \tilde{h}(\rho, \phi - \theta; r, 0).$$

Under this assumption, the object intensity can be written as what we call a *ring convolution*, denoted by $f \triangleq g \circledast h$, where

$$f(x, y) = (g \circledast h)(x, y) = \int_0^\infty r(\tilde{g} *_{\theta} \tilde{h})(\sqrt{x^2 + y^2}, \tan^{-1}(y/x); r, 0)dr. \tag{3}$$

Algorithm 1 Ring convolution

Input: $N \times N$ pixel image g ; PSFs along one radial line $h^{(j)}$, $j = 1, \dots, K$; corresponding distances r_j , $j = 1, \dots, K$ of each PSF from the center.

Output: LRI blurred image f

- 1: $\tilde{g} \leftarrow \text{polarTransform}(g)$ ▷ polar dimensions are $M \times K$, angle by radius
 - 2: $\tilde{f} \leftarrow \text{zeros}(M \times K)$ ▷ initialize the output in polar form as an all zero matrix
 - 3: **for** $j = 1, \dots, K$ **do**
 - 4: $\tilde{h}^{(j)} \leftarrow \text{polarTransform}(h^{(j)})$
 - 5: **for** $i = 1, \dots, K$ **do**
 - 6: $\tilde{f}_{:,i} \leftarrow \tilde{f}_{:,i} + \text{iFFT}\{r_j \text{FFT}\{\tilde{g}_{:,j}\} \text{FFT}\{\tilde{h}_{:,i}^{(j)}\}\}$ ▷ compute polar output ring by ring, FFT is 1D
 - 7: $f \leftarrow \text{inversePolarTransform}(\tilde{f})$
-

This ring-wise computation, wherein points at different radii are filtered heterogeneously, is consistent with the underlying intuition in LRI: the blur varies radially. Our first main result is an efficient, FFT-based inversion of the ring convolution.

Theorem 1 (Ring Convolution Theorem). *Under LRI, where \mathcal{F}_Θ is a 1D Fourier transform over θ ,*

$$\tilde{f}(\rho, \phi) = \mathcal{F}_\Theta^{-1} \left\{ \int r \mathcal{F}_\Theta \{ \tilde{g}(r, \theta) \} \mathcal{F}_\Theta \{ \tilde{h}(\rho, \theta; r) \} dr \right\}(\phi).$$

Proof. Since the given system is LRI, (3) holds. Substituting $\rho = \sqrt{x^2 + y^2}$ and $\phi = \tan^{-1}(y/x)$, we can rewrite (3) as

$$f(\rho, \phi) = \int_0^\infty r (\tilde{g} *_{\theta} \tilde{h})(\rho, \phi; r, 0) dr.$$

Applying the Fourier Convolution Theorem to the one-dimensional convolution on the right-hand side yields

$$\tilde{f}(\rho, \phi) = \int r \mathcal{F}_\Theta^{-1} \left\{ \mathcal{F}_\Theta \{ \tilde{g}(r, \theta) \} \mathcal{F}_\Theta \{ \tilde{h}(\rho, \theta; r) \} \right\}(\phi) dr,$$

where \mathcal{F}_Θ is the one dimensional Fourier transform over θ . By Fubini's theorem, we pull the inverse Fourier Transform outside of the integral, which gives

$$\tilde{f}(\rho, \phi) = \mathcal{F}_\Theta^{-1} \left\{ \int r \mathcal{F}_\Theta \{ \tilde{g}(r, \theta) \} \mathcal{F}_\Theta \{ \tilde{h}(\rho, \theta; r) \} dr \right\}(\phi). \quad \blacksquare$$

Ring convolution has an efficient and convex formulation for computing its inverse, ring deconvolution:

$$\hat{g} = \arg \min_{\tilde{g}} \|f - \tilde{g} \circledast h\|_2^2. \quad (4)$$

This problem can be solved efficiently via an iterative least squares solver using Algorithm 1 as a substep. A Fourier interpretation of ring convolution is provided in Appendix C. While the above results are all exact, the discrete time implementations of them have small, but nonzero errors due to discretization. For example, the polar transformation in Algorithm 1 requires a small amount of interpolation.

4.2 Fitting Seidel Coefficients to PSFs

Ring (de)convolution algorithms require h , the collection of PSFs along one radial line in the FoV. Fortunately, there is a convenient and compact alternative to measuring these manually. The Seidel coefficients [33, 49] are a polynomial basis that can represent any rotationally symmetric system. We mathematically describe the form of these aberrations in Appendix B.

The estimation procedure for estimating the Seidel coefficients involves fitting them to a single, sparse image of a few randomly scattered point sources (e.g. fluorophores on a microscope slide)—such an image is

usually easier to obtain than an image of an isolated point source in the center of the FoV. The presence of off-axis PSFs in the calibration image provides information about all of the aberration coefficients. Though it may be possible to fit these coefficients from a single off-axis PSF, we find that a few, randomly placed PSFs provides a robust fit. Let r_1, \dots, r_J be the object-plane radii of the J points in the calibration image. We then find the primary Seidel coefficients $\hat{\omega}$ whose generated PSFs best match the measured PSFs. Once again, this searching procedure is succinctly stated as an optimization problem,

$$\hat{\omega} = \arg \min_{\bar{\omega}} \sum_{j=1}^J \|h^{(j)} - \mathcal{F}^{-1}\{P(\bar{\omega})^{(j)}\}\|_2^2,$$

where $P(\bar{\omega})^{(j)}$ is the pupil function with Seidel coefficients $\bar{\omega}$ from a point source at distance r_j from the center of the FoV.

It has been shown that for LRI systems, the optimal fit $\hat{\omega}$ achieves a diminishing error [33]. Furthermore, the 5 primary Seidel coefficients index the dominant aberrations present in practical imaging systems: sphere, coma, astigmatism, field curvature, and distortion. For more complex aberrations it is possible to add higher-order Seidel coefficients to the fit. In practice, we fit these 5 coefficients via the ADAM optimizer [50] and obtain reasonable local minima even though the problem is nonconvex (see Fig. 3). Armed with the estimated Seidel coefficients, we can generate synthetic PSFs at any radius.

4.3 Deep Ring Deconvolution (DeepRD)

DeepRD is designed to incorporate the Seidel coefficients into the deblurring process in a parameter-efficient and interpretable manner. To that end, we propose a neural network architecture inspired by the physical LRI image formation model. The first key design element is to use a modified Hypernetwork [51], a network which predicts the weights of another, task-specific “primary” network. In DeepRD, an MLP-based hypernetwork takes in Seidel coefficients and produces a deep deblurring network that specifically works for the given coefficients. Our second key design element is the use of ringwise convolution kernels. Specifically, the hypernetwork produces CNN kernels for each radius which the primary network applies ring-by-ring. This replicates the rotation-invariance assumption central to ring deconvolution and eases the space-invariant constraint of typical convolutional kernels. Together, this design enables a neural network that is a fraction the size of a conventional U-Net with improved performance and generalization. To our knowledge, we are the first to utilize learnable ring convolutions in the context of deep learning.

To produce a training dataset for DeepRD, we synthetically generate blurred input images from the Div2k and CARE fluorescence microscopy datasets using ring convolution with randomly sampled sets of Seidel coefficients. We must sample coefficients which adequately cover the attributes of realistic imaging systems. Each coefficient, which is in units of waves, is drawn from a uniform distribution and a noise-perturbed grid in the range 0 – 3 waves—we empirically find this range to cover the aberrations of systems ranging from perfect to highly aberrated. Note that without ring convolution, such a dataset would be prohibitively slow to produce. With that in mind, we will release both an implementation of ring convolution as well as our dataset.

DeepRD accepts a two-part input, a blurred image and its corresponding primary Seidel coefficients. The model is then supervised with the unblurred ground truth image. We find that this physically-grounded synthetic dataset generation is effective to train models that generalize to real-world evaluation.

4.4 Blind deblurring

Our version of blind deconvolution also takes advantage of the Seidel coefficients. Given just a blurry image, we start by randomly picking a value for the spherical Seidel coefficient. Then we use this value to synthetically generate a point spread function, and use it to deconvolve the blurry image. We then compute the sum of the spatial gradient of the resulting deconvolved image (this acts as a surrogate measure of image sharpness) and use its negative as a loss. We then minimize this loss (maximize the sharpness) by updating our initial guess of the spherical aberration coefficient using its gradient with respect to the loss function. Running this iteratively, we eventually converge to a final spherical aberration coefficient, generate a final synthetic PSF, and deconvolve the blurry image with this PSF to get the final result.

Note that this procedure generalizes to spatially-varying systems. We would instead jointly estimate all 5 Seidel coefficients and use ring deconvolution instead deconvolution at each step. This, however, is computationally expensive and requires generating N (the image sidelength) PSFs per iterative step. We believe it is possible to do this more efficiently with DeepRD by replacing the ring deconvolution operation at each step with DeepRD. That is, we search the space of DeepRD networks for the one that produces the sharpest reconstruction. However, this is out of scope for this project, and we leave it as future work.

4.5 Experimental details

4.5.1 Sample Preparation

A mixture of starved and nonstarved tardigrades were stained overnight with Invitrogen nucleic acid gel fluorescent stain, whose excitation and emission maxima are 502 nm and 530 nm, respectively. Individual stained tardigrades were then isolated onto a glass slide for imaging. Meanwhile, the non-fluorescent samples (USAF resolution targets, and rabbit liver tissue) were obtained imaging-ready on glass slides.

4.5.2 Imaging

For all imaging experiments, we used the UCLA Miniscope v3 with the Ximea MU9PM-MBRD 12 bit, 2.2 micron pixel sensor. Optically, the Miniscope is comprised of a gradient index objective and achromat tube lens; further details are provided in [1]. In order obtain the system PSFs, we imaged 1 μm fluorescent beads smeared on a glass slide. For deconvolution microscopy calibration, we repeatedly diluted the bead solution with isopropyl alcohol until we were able to sufficiently isolate a single bead, whereas for RDM calibration we used a single dilution and imaged a slide containing a sparse collection of beads.

4.5.3 Computation

PSF generation for the simulation experiments was done by synthetically generating pupil functions with the given Seidel coefficients [49] (see Github). The remaining computation was done using Python on a single Nvidia GeForce GPU. For standard deconvolution the measured PSF was denoised through background subtraction and pixel-wise thresholding.

All non-learning, iterative methods are solving linear least squares optimization problems (see (4)); we additionally add TV regularization to these and run them till convergence using an ADAM optimizer [50]. For each method, the hyperparameters—including learning rate and regularization strength—that provided the smallest loss and best qualitative results were used. For deconvolution we tried a variety of algorithms in addition to the iterative scheme, including Wiener filtering and Richardson-Lucy deconvolution, and used the best reconstruction, which was unsupervised wiener filtering [52].

Open source implementations of ring convolution, polar transform, Seidel fitting, and ring deconvolution can be found in the codebase. Our intent is for this codebase to function as a easy-to-use library such that any practitioner with any imaging system can utilize RDM with little-to-no overhead.

The baseline U-Net and DeepRD were both trained on ring-convolved images from the Div2k dataset. For the simulation results, both models were additionally fine-tuned on images from the CARE dataset. All models were trained till convergence of the validation loss and optimized over hyperparameters.

Acknowledgements

We would like to acknowledge Neerja Aggarwal for help and access to 3D printing. A.K. was funded by the Berkeley Fellowship for Graduate Study. A.N.A. was supported by the Berkeley Fellowship for Graduate Study and the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 1752814. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. S.Y. is supported by Moore Foundation data-driven discovery. This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-22-1-0521. This publication has been made possible in part by CZI grant DAF2021-225666 and grant DOI <https://doi.org/10.37921/192752jrgbn>

from the Chan Zuckerberg Initiative DAF, an advised fund of Silicon Valley Community Foundation (funder DOI 10.13039/100014989).

References

- [1] D. Aharoni, *UCLA miniscope*. [Online]. Available: http://www.miniscope.org/index.php/Main_Page.
- [2] H. P. Stahl, F. A. Prince, C. Smart, K. Stephens, and T. Henrichs, "Preliminary cost model for space telescopes," *UV/Optical/IR Space Telescopes: Innovative Technologies and Concepts IV*, vol. 7436, pp. 11–22, 2009.
- [3] H. Trussell and B. Hunt, "Image restoration of space variant blurs by sectioned methods," in *ICASSP '78. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 3, 1978, pp. 196–198. DOI: 10.1109/ICASSP.1978.1170472.
- [4] S.-H. Lin, T. F. Krile, and J. F. Walkup, "Piecewise isoplanatic modeling of space-variant linear systems," *J. Opt. Soc. Am. A*, vol. 4, no. 3, pp. 481–487, 1987. DOI: 10.1364/JOSAA.4.000481. [Online]. Available: <http://opg.optica.org/josaa/abstract.cfm?URI=josaa-4-3-481>.
- [5] A. W. Lohmann and D. P. Paris, "Space-variant image formation," *J. Opt. Soc. Am.*, vol. 55, no. 8, pp. 1007–1013, 1965. DOI: 10.1364/JOSA.55.001007. [Online]. Available: <http://opg.optica.org/abstract.cfm?URI=josa-55-8-1007>.
- [6] L. Denis, E. Thiébaud, and F. Soulez, "Fast model of space-variant blurring and its application to deconvolution in astronomy," in *2011 18th IEEE International Conference on Image Processing*, 2011, pp. 2817–2820. DOI: 10.1109/ICIP.2011.6116257.
- [7] J. G. Nagy and D. P. O'Leary, "Restoring images degraded by spatially variant blur," *SIAM Journal on Scientific Computing*, vol. 19, no. 4, pp. 1063–1082, 1998. DOI: 10.1137/S106482759528507X. [Online]. Available: <https://doi.org/10.1137/S106482759528507X>.
- [8] L. Denis, E. Thiébaud, F. Soulez, J.-M. Becker, and R. Mourya, "Fast approximations of shift-variant blur," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 253–278, 2015.
- [9] R. C. Flicker and F. J. Rigaut, "Anisoplanatic deconvolution of adaptive optics images," *J. Opt. Soc. Am. A*, vol. 22, no. 3, pp. 504–513, 2005. DOI: 10.1364/JOSAA.22.000504. [Online]. Available: <http://opg.optica.org/josaa/abstract.cfm?URI=josaa-22-3-504>.
- [10] H. Deng, W. Zuo, H. Zhang, and D. Zhang, "An additive convolution model for fast restoration of nonuniform blurred images," English, *International Journal of Computer Mathematics*, vol. 91, no. 11, pp. 2446–2466, 2014. DOI: 10.1080/00207160.2013.811235.
- [11] D. Miraut and J. Portilla, "Efficient shift-variant image restoration using deformable filtering (part i)," *EURASIP Journal on Advances in Signal Processing*, vol. 2012, no. 1, pp. 1–20, 2012.
- [12] T. Popkin, A. Cavallaro, and D. Hands, "Accurate and efficient method for smoothly space-variant gaussian blurring," *IEEE Transactions on image processing*, vol. 19, no. 5, pp. 1362–1370, 2010.
- [13] F. Sroubek, J. Kamenicky, and Y. M. Lu, "Decomposition of space-variant blur in image deconvolution," *IEEE Signal Processing Letters*, vol. 23, no. 3, pp. 346–350, 2016. DOI: 10.1109/LSP.2016.2519764.
- [14] T. Lauer, "Deconvolution with a spatially variant PSF," in *Astronomical Data Analysis II*, J.-L. Starck and F. D. Murtagh, Eds., International Society for Optics and Photonics, vol. 4847, SPIE, 2002, pp. 167–173. DOI: 10.1117/12.461035. [Online]. Available: <https://doi.org/10.1117/12.461035>.
- [15] S. Lee, K. Lee, S. Shin, and Y. Park, "Generalized image deconvolution by exploiting the transmission matrix of an optical imaging system," *Scientific Reports*, vol. 7, no. 1, pp. 1–10, 2017.
- [16] K. Yanny, K. Monakhova, R. W. Shuai, and L. Waller, "Deep learning for fast spatially varying deconvolution," *Optica*, vol. 9, no. 1, pp. 96–99, 2022. DOI: 10.1364/OPTICA.442438. [Online]. Available: <http://opg.optica.org/optica/abstract.cfm?URI=optica-9-1-96>.

- [17] W. Ren, J. Zhang, L. Ma, J. Pan, X. Cao, W. Zuo, W. Liu, and M.-H. Yang, “Deep non-blind deconvolution via generalized low-rank approximation,” *Advances in neural information processing systems*, vol. 31, 2018.
- [18] Q. Deng, Z. Wen, Z. Dong, J. Tang, W. Chen, X. Liu, and Q. Yang, “Spatially variant deblur and image enhancement in a single multimode fiber imaged by deep learning,” *Opt. Lett.*, vol. 47, no. 19, pp. 5040–5043, Oct. 2022. DOI: 10.1364/OL.469034. [Online]. Available: <https://opg.optica.org/ol/abstract.cfm?URI=ol-47-19-5040>.
- [19] J. Dong, S. Roth, and B. Schiele, “Deep wiener deconvolution: Wiener meets deep learning for image deblurring,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 1048–1059, 2020.
- [20] M. Weigert, U. Schmidt, T. Boothe, A. Müller, A. Dibrov, A. Jain, B. Wilhelm, D. Schmidt, C. Broadus, S. Culley, *et al.*, “Content-aware image restoration: Pushing the limits of fluorescence microscopy,” *Nature methods*, vol. 15, no. 12, pp. 1090–1097, 2018.
- [21] Y. Fang, H. Zhang, H. S. Wong, and T. Zeng, “A robust non-blind deblurring method using deep denoiser prior,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, Jun. 2022, pp. 735–744.
- [22] Y. Li, Y. Su, M. Guo, X. Han, J. Liu, H. D. Vishwasrao, X. Li, R. Christensen, T. Sengupta, M. W. Moyle, *et al.*, “Incorporating the image formation process into deep learning improves network performance,” *Nature Methods*, pp. 1–11, 2022.
- [23] C. B. Webster and S. J. Reeves, “Radial deblurring with ffts,” in *2007 IEEE International Conference on Image Processing*, vol. 1, 2007, pp. I–101–I–104. DOI: 10.1109/ICIP.2007.4378901.
- [24] G. Boracchi, A. Foi, V. Katkovnik, and K. Egiazarian, “Deblurring noisy radial-blurred images: Spatially adaptive filtering approach,” in *Image Processing: Algorithms and Systems VI*, International Society for Optics and Photonics, vol. 6812, 2008, p. 68121D.
- [25] S. Luan, S. Xie, T. Wang, X. Hao, M. Yang, and Y. Li, “A space-variant deblur method for focal-plane microwave imaging,” *Applied Sciences*, vol. 8, no. 11, p. 2166, 2018.
- [26] Y. Zhang and T. Ueda, “Deblur of radially variant blurred image for single lens system,” *IEEJ Transactions on Electrical and Electronic Engineering*, vol. 6, no. S1, S7–S16, 2011. DOI: <https://doi.org/10.1002/tee.20615>. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/tee.20615>.
- [27] Y. Zhang, I. Minema, and T. Ueda, “Restoration of radially blurred image created by spherical single lens system of cell phone camera,” *Proceedings of IEEE Sensors*, pp. 1333–1337, 2010. DOI: 10.1109/ICSENS.2010.5690562.
- [28] Y. Zhang, I. Minema, L. G. Zimin, and T. Ueda, “Analysis of radially restored images for spherical single lens cellphone camera,” *IEEE Sensors Journal*, vol. 11, no. 11, pp. 2834–2844, 2011. DOI: 10.1109/JSEN.2011.2167504.
- [29] Y. Zhang, L. G. Zimin, J. Ji, S. Ikezawa, and T. Ueda, “Real scene capturing using spherical single-element lens camera and improved restoration algorithm for radially variant blur,” *Opt. Express*, vol. 20, no. 25, pp. 27 569–27 588, 2012. DOI: 10.1364/OE.20.027569. [Online]. Available: <http://opg.optica.org/oe/abstract.cfm?URI=oe-20-25-27569>.
- [30] T. Yue, J. Suo, J. Wang, X. Cao, and Q. Dai, “Blind optical aberration correction by exploring geometric and visual priors,” Jun. 2015, pp. 1684–1692. DOI: 10.1109/CVPR.2015.7298777.
- [31] G. Zheng, X. Ou, R. Horstmeyer, and C. Yang, “Characterization of spatially varying aberrations for wide field-of-view microscopy,” *Opt. Express*, vol. 21, no. 13, pp. 15 131–15 143, Jul. 2013. DOI: 10.1364/OE.21.015131. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?URI=oe-21-13-15131>.
- [32] Y. Shao, M. Loktev, Y. Tang, F. Bociort, and H. P. Urbach, “Spatially varying aberration calibration using a pair of matched periodic pinhole array masks,” *Opt. Express*, vol. 27, no. 2, pp. 729–742, Jan. 2019. DOI: 10.1364/OE.27.000729. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?URI=oe-27-2-729>.

- [33] M. Born and E. Wolf, *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013.
- [34] D. Y. H. Wang and D. T. Moore, “Third-order aberration theory for weak gradient-index lenses,” *Appl. Opt.*, vol. 29, no. 28, pp. 4016–4025, Oct. 1990. DOI: 10.1364/AO.29.004016. [Online]. Available: <https://opg.optica.org/ao/abstract.cfm?URI=ao-29-28-4016>.
- [35] E. Agustsson and R. Timofte, “Ntire 2017 challenge on single image super-resolution: Dataset and study,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 126–135.
- [36] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, “Extracting and composing robust features with denoising autoencoders,” in *Proceedings of the 25th international conference on Machine learning*, 2008, pp. 1096–1103.
- [37] W. H. Southwell, “Wave-front analyzer using a maximum likelihood algorithm,” *J. Opt. Soc. Am.*, vol. 67, no. 3, pp. 396–399, Mar. 1977. DOI: 10.1364/JOSA.67.000396. [Online]. Available: <https://opg.optica.org/abstract.cfm?URI=josa-67-3-396>.
- [38] A. Ignatov, R. Timofte, *et al.*, “Pirm challenge on perceptual image enhancement on smartphones: Report,” in *European Conference on Computer Vision (ECCV) Workshops*, Jan. 2019.
- [39] L. Yang, X. Zhang, and J. Ren, “Adaptive wiener filtering with gaussian fitted point spread function in image restoration,” in *2011 IEEE 2nd International Conference on Software Engineering and Service Science*, 2011, pp. 208–212. DOI: 10.1109/ICSESS.2011.5982291.
- [40] D. K. Samuylov, P. Purwar, G. Székely, and G. Paul, “Modeling point spread function in fluorescence microscopy with a sparse gaussian mixture: Tradeoff between accuracy and efficiency,” *IEEE Transactions on Image Processing*, vol. 28, no. 8, pp. 3688–3702, 2019. DOI: 10.1109/TIP.2019.2898843.
- [41] E. Maalouf, B. Colicchio, and A. Dieterlen, “Fluorescence microscopy three-dimensional depth variant point spread function interpolation using zernike moments,” *J. Opt. Soc. Am. A*, vol. 28, no. 9, pp. 1864–1870, Sep. 2011. DOI: 10.1364/JOSAA.28.001864. [Online]. Available: <https://opg.optica.org/josaa/abstract.cfm?URI=josaa-28-9-1864>.
- [42] N. Becherer, H. Jödicke, G. Schlosser, J. Hesser, F. Zeilfelder, and R. Männer, “On soft clipping of Zernike moments for deblurring and enhancement of optical point spread functions,” in *Computational Imaging IV*, C. A. Bouman, E. L. Miller, and I. Pollak, Eds., International Society for Optics and Photonics, vol. 6065, SPIE, 2006, p. 60650C. DOI: 10.1117/12.642272. [Online]. Available: <https://doi.org/10.1117/12.642272>.
- [43] B. M. Hanser, M. G. Gustafsson, D. A. Agard, and J. W. Sedat, “Phase retrieval for high-numerical-aperture optical systems,” *Optics letters*, vol. 28, no. 10, pp. 801–803, 2003.
- [44] B. M. HANSER, M. G. L. GUSTAFSSON, D. A. AGARD, and J. W. SEDAT, “Phase-retrieved pupil functions in wide-field fluorescence microscopy,” *Journal of Microscopy*, vol. 216, no. 1, pp. 32–48, 2004. DOI: <https://doi.org/10.1111/j.0022-2720.2004.01393.x>. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.0022-2720.2004.01393.x>.
- [45] P. Pankajakshan, L. Blanc-Feraud, Z. Kam, and J. Zerubia, “Point-spread function retrieval for fluorescence microscopy,” in *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2009, pp. 1095–1098. DOI: 10.1109/ISBI.2009.5193247.
- [46] C. Wang, J. Chen, H. Jia, B. Shi, R. Zhu, Q. Wei, L. Yu, and M. Ge, “Parameterized modeling of spatially varying psf for lens aberration and defocus,” *J. Opt. Soc. Korea*, vol. 19, no. 2, pp. 136–143, 2015. [Online]. Available: <http://opg.optica.org/josk/abstract.cfm?URI=josk-19-2-136>.
- [47] J. Simpkins and R. L. Stevenson, “Mapping measurable quantities of point-spread function observations to seidel aberration coefficients,” in *2012 19th IEEE International Conference on Image Processing*, 2012, pp. 369–372. DOI: 10.1109/ICIP.2012.6466872.
- [48] J. W. Goodman, “Introduction to fourier optics. 3rd,” *Roberts and Company Publishers*, 2005.
- [49] D. G. Voelz, *Computational fourier optics: a MATLAB tutorial*. SPIE press Bellingham, Washington, 2011.

- [50] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv:1412.6980*, 2014.
- [51] D. Ha, A. M. Dai, and Q. V. Le, “Hypernetworks,” in *International Conference on Learning Representations*, 2017. [Online]. Available: <https://openreview.net/forum?id=rkpACe1lx>.
- [52] F. Orieux, J.-F. Giovannelli, and T. Rodet, “Bayesian estimation of regularization and point spread function parameters for wiener–hunt deconvolution,” *J. Opt. Soc. Am. A*, vol. 27, no. 7, pp. 1593–1607, Jul. 2010. DOI: 10.1364/JOSAA.27.001593. [Online]. Available: <https://opg.optica.org/josaa/abstract.cfm?URI=josaa-27-7-1593>.
- [53] U. S. N. B. of Standards, *Miscellaneous Publication - National Bureau of Standards* (Miscellaneous Publication - National Bureau of Standards). The Bureau, 1934. [Online]. Available: <https://books.google.com/books?id=Cct00yTHR3QC>.
- [54] A. V. Oppenheim and R. W. Schaffer, “Digital signal processing(book),” *Research supported by the Massachusetts Institute of Technology, Bell Telephone Laboratories, and Guggenheim Foundation. Englewood Cliffs, N. J., Prentice-Hall, Inc., 1975. 598 p*, 1975.
- [55] A. Dominguez, *A history of the convolution operation*, 2015. [Online]. Available: <https://www.embs.org/pulse/articles/history-convolution-operation/>.
- [56] N. Wiener *et al.*, *Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*. MIT press Cambridge, MA: 1964, vol. 8.
- [57] W. H. Richardson, “Bayesian-based iterative method of image restoration,” *J. Opt. Soc. Am.*, vol. 62, no. 1, pp. 55–59, 1972. DOI: 10.1364/JOSA.62.000055. [Online]. Available: <http://opg.optica.org/abstract.cfm?URI=josa-62-1-55>.
- [58] L. Lucy, “An iterative technique for the rectification of observed distributions,” *Astronomical Journal*, vol. 79, p. 745, 1974. DOI: 10.1086/111605.

A Interpretability of DeepRD

One key property of any deep learning model is its interpretability. How well we understand what the network is doing to produce its predictions corresponds to how confident we are in its performance. One such mark of interpretability when working with encoder-decoder frameworks, is whether the latent space, i.e., the space of encodings, is disentangled, meaning each coordinate of the encoding has a distinctive meaning. In the case of DeepRD, the latent space is the space of Seidel coefficients. In the ideal case we expect the network to learn about the individual effects of each coordinate (each aberration coefficient) it takes in as input; say we only increase the coefficient value for radial distortion, then we expect the network to further unwarp the predicted image correspondingly.

This property indeed emerges in practice. We observe this by deblurring images of insect cornea taken by the Miniscope using a sweep over each aberration coefficient independently. For each coefficient, we set all other coefficients to 0 and sweep the chosen coefficient from 0 to 3 waves. What we observe is that the network only corrects for the chosen aberration in increasing amounts while leaving the other characteristics of the image alone. For example, when we sweep distortion, the image remains blurry but radially warps in the opposite direction compared to if it were distorted. Results for each aberration coefficient are shown in video form here. Such results confirms our hypothesis that DeepRD is indeed learning to perform ring deconvolution given a systems Seidel coefficients, and does not rely on knowledge of the specific image distribution. This further aids our confidence in its ability to generalize.

B Seidel Coefficients

Here, we give a brief background on Seidel coefficients. This is not a complete treatment and we encourage an interested reader to refer to Voelz et al.[49] for an overview, and the standard Born et al.[33] for a fully rigorous derivation.

Consider an rotationally symmetric—and consequently LRI—imaging system. It is common to consolidate the system aberrations into a single function and apply it at the exit pupil plane of the system. This complex-valued function, known as the generalized pupil function p is composed of a binary-valued amplitude distribution (set by the shape of the pupil) and a phase distribution w , which quantifies the deviation of the pupil wavefront from the ideal spherical shape necessary for diffraction-limited imaging. Since LRI systems are generally spatially-varying (albeit only radially), w becomes a function of the distance of the source from the optical axis r . Letting (s, t) be the pupil plane coordinates, we can write p as

$$p(s, t; r) = \text{circ}\left(\frac{r}{R}\right) e^{w(s, t; r)}.$$

For rotationally symmetric systems, as in our case, it is possible to expand w as an infinite power series (see the aberration function in [33]). Usually only the third-order terms of this series are used, which yields the following 2D polynomial of the pupil plane coordinates,

$$w_\omega(s, t; r) = \omega_s(s^2 + t^2)^2 + \omega_c(s^2 + t^2)sr + \omega_a s^2 r^2 + \omega_f(s^2 + t^2)r^2 + \omega_d sr^3,$$

where $\omega = (\omega_s, \omega_c, \omega_a, \omega_f, \omega_d)$ are the five primary Seidel coefficients and w_ω only depends on the radial location due to the LRI assumption. Note that while the 5 primary Seidel coefficients are a subset of the infinitely many available coefficients, they represent the most common optical aberrations: spherical, coma, astigmatism, field curvature, and distortion. In particular, these aberrations are inherent to all spherically-shaped optics. Finally, the pupil function contains the same information as the PSFs, since they are related by a Fourier transform,

$$h(x, y; r) = \left| \mathcal{F}^{-1}\{p(-\lambda df_x, -\lambda df_y; r)\} \right|^2,$$

where λ is the wavelength of light, d is the distance from the pupil plane to the image plane, and f_x, f_y are the variables which the Fourier transform is taken over. Thus, knowledge of the 5 Seidel coefficients provides an approximation of the pupil function, which in turn, can accurately estimate PSFs for LRI systems, including the radial line of PSFs needed for computing ring convolution.

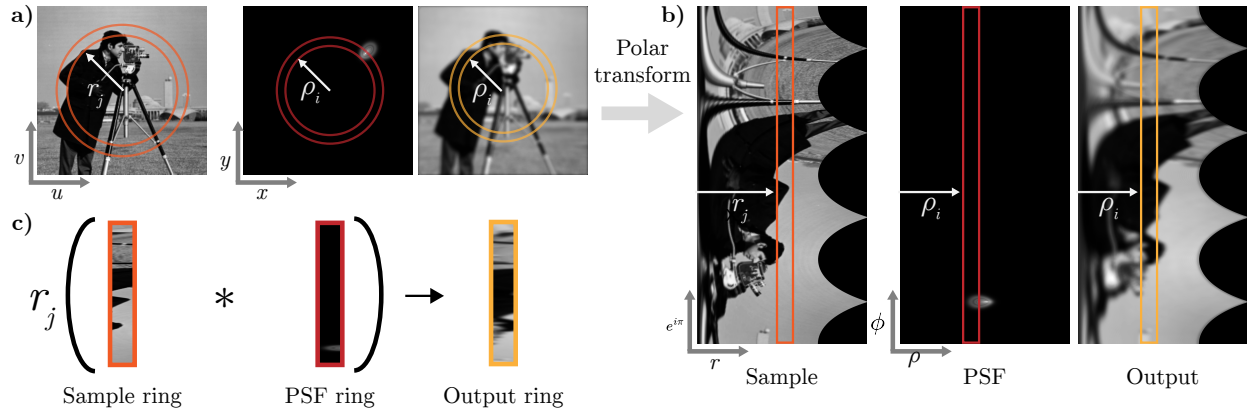


Figure 5: **Ring convolution.** A single step of the LRI convolution algorithm. Here we are solving for a single concentric ring of the output denoted by its distance ρ_i from the center. **a)** Left to right: sharp ‘sample’, the system PSF at radius r_j , and blurry output of ring convolution. **b)** Corresponding polar resamplings of the images in a), this allows us to extract rings. **c)** The output ring at ρ_i is the sum of 1D convolutions; the ρ_i ring of the PSF at position r_j is convolved with the r_j ring of the sample for every radius r_j .

C A Fourier analysis of LRI systems

Perhaps the main reason to make the linear space-invariant assumption is access to its Fourier space interpretation. Through an application of the Convolution Theorem to the space-invariant forward model, we see that an LSI system’s output spectrum is a product of the input’s spectrum with the spectrum of the impulse response, i.e., the transfer function. Thus any LSI imaging system can be thought of as a *filter* which individually scales each frequency component of the input object’s intensity distribution. In the context of imaging, the transfer function—called the Optical Transfer Function (OTF)—describes how the imaging system scales each spatial frequency in the sample. The OTF gives rise to valuable intuitions; for example, since all practical imaging systems have bandlimited transfer functions, the output image is a lowpass-filtered version of the sample whose maximum-achievable resolution is directly proportional to the bandlimit. Moreover system aberrations can be analyzed by comparing the OTF in the presence of aberrations with the ideal OTF.

In this section we will formally develop an analogous Fourier interpretation of ring convolution and explain how its features can similarly be used to characterize the performance of an imaging system. While the LRI interpretation is more complicated than its LSI counterpart, it still provides a rich and realistic view of how LRI imaging systems transmit frequencies. As an example of its utility, we will see that the LRI model allows for a more general, radially-dependent notion of system resolution. For a review of the ring convolution operation please refer to Fig. 5

To begin, we will define a slightly different notion of spatial frequency, called rotational frequency and denoted Θ , which is the quantity that gets filtered by an LRI system.

Definition 1 (Rotational Fourier Transform (RoFT)). *Let $f : \mathbb{R}^2 \rightarrow \mathbb{C}$ represent a (potentially complex-valued) image, and let \tilde{f} be its polar counterpart (as per 1). The RoFT of f is given by*

$$\tilde{F}(r, \xi) = \int \tilde{f}(r, \theta) e^{i2\pi\theta\xi} d\theta.$$

Intuitively, one can think of the values of rotational frequencies as quantifying how quickly the image can change as one travels in a ring of radius r around the center. However, this notion of oscillation speed depends on the radius—think about the spokes on the wheel of a bicycle, even though the spokes are evenly spaced for any given radius, they become further spaced for larger radii. However, under polar sampling they are scaled to have the same frequency. This scaling can be seen in Fig. 6 as the RoFT of each quantity tend to have higher frequencies at larger radii.

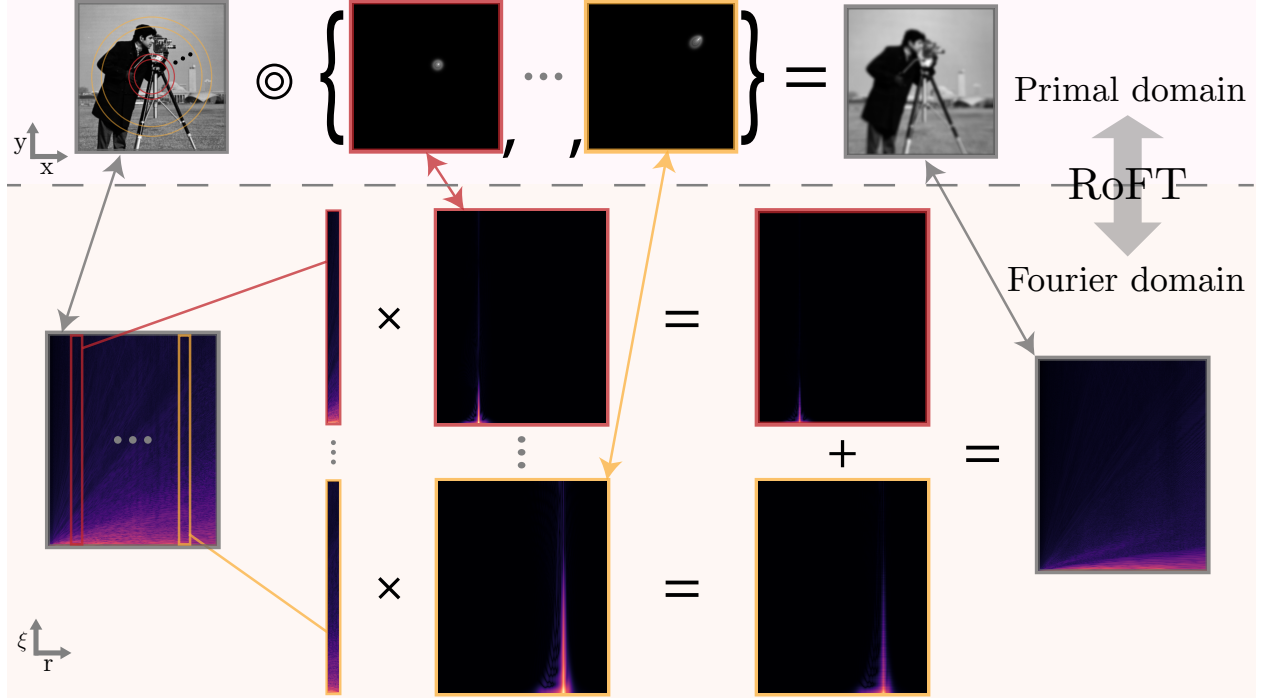


Figure 6: **LRI filtering.** The top, pink box shows LRI filtering in the spatial domain: an object (left) is LRI filtered by a radial set of PSFs (middle) to give a blurry image (right). The bottom, orange box shows the Fourier domain equivalent. Double arrows indicate Rotational Fourier Transform (RoFT) pairs. Strips at each r of the object RoFT (left) are individually multiplied by the RoFTs of the corresponding PSF at r (middle left) producing filtered contributions (middle right). Finally each of these contributions for every r is summed to form the RoFT of the blurry image (right).

Now, writing the Ring Convolution Theorem in terms of the RoFTs of each quantity (denoted with capital letter and tilde), we see that

$$\tilde{F}(\rho, \xi) = \int r \tilde{G}(r, \xi) \tilde{H}(\rho, \xi; r) dr.$$

Like its LSI counterpart, the above model can be thought of a filtering operation, but now a more complex one; the object's values r away from the center are *filtered* and *mixed* by the r^{th} PSF. To more fully understand this interpretation, consider the object's values at some radius r ; they form a ring which will be filtered by the r^{th} PSF. Specifically, the spectrum of this ring, or equivalently its RoFT at r , is point-wise multiplied by each ring in the RoFT of the r^{th} PSF, which yields a set of filtered rings indexed by ρ . The ρ^{th} filtered ring represents the contribution of the object ring at r to the image ring at ρ (see Fig. 6).

It follows that the shape of the r^{th} PSF determines both how much the object's rotational frequencies are filtered through the system and which parts of the object are mixed together to form the image. By looking at the shape of the r^{th} PSF we can determine the exact nature of this filtering and mixing. As shown in Fig. 7 the angular length of the PSF determines the height of the PSF's RoFT at ρ , which tells us which of the object rotational frequencies at r will make it to the image at ρ —this is analogous to the OTF bandwidth in the LSI case. Meanwhile, the radial extent of the PSF (i.e., how many concentric rings it covers) controls the extent of the image rings (i.e., which ρ values) are effected by the object at r . This manifests as the width of the PSF's RoFT which we call the *mix width*.

Now our notion of resolution, assuming the ability to perfectly de-mix, is radially dependent and, at radius r , is proportional to the angular arc length at $\rho = r$ of the r^{th} PSF. The LRI filtering interpretation offers a more realistic understanding of how an imaging system filters an object and opens the door for a host of new imaging techniques which optimize for key features, such as resolution, under the LRI model.

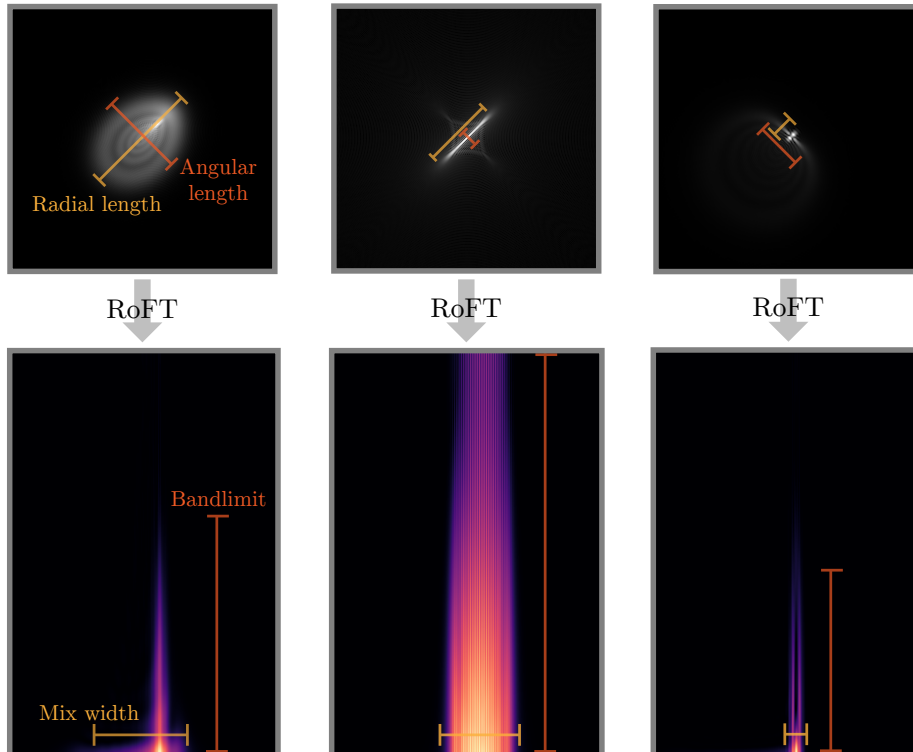


Figure 7: **LRI PSF interpretation.** Evaluating an LRI optical system amounts to considering the length of each PSF in the angular and radial directions. These lengths are inversely proportional to the bandwidth and mix width of the corresponding RoFT, respectively (see leftmost column). In the middle column we see an elongated PSF in the radial direction (i.e., astigmatism) which has a large bandwidth, but incurs more off-radius mixing due to a large mix width. The rightmost column shows an elongated PSF in the angular direction which has a small mix width, but has a relatively small bandwidth.

D Additional Related Work

The idea of linear space-invariance (LSI)—or rather its 1D predecessor time invariance—is difficult to trace historically, but the first mention of it that we were able to find was by Richard Hamming in 1934 [53], followed by the far more popular definition provided in [54]. The concept also has a close connection to the convolution integral—which had existed separately as far back as the 1800s [55]—through the famous Convolution Theorem. Much of the foundational mathematics of convolution and its frequency domain interpretation comes from the work of Laplace and Fourier; these techniques enabled the application of linear time (space)-invariant systems to communications and, of course, imaging. The practical use of LSI systems for imaging analysis arose from Bell labs during the emergence of digital signal processing for communication in the 1960s and 1970s.

Image deblurring has its roots much later; it was only in 1949 when Norbert Wiener published the Wiener Filter for linear time-invariant deconvolution [56]. While the original technique was developed in 1940 for time series data in order to track enemy planes in World War II, it was soon adapted for image deblurring. With the rise of digital computing in the following decades came an explosion of image deconvolution techniques which are commonly employed today such as Richardson-Lucy deconvolution [57, 58]. It is interesting to note that some of these algorithms have existed for nearly a century and yet there is no clear “best” method.