

Computational Imaging using Ultrasonically-Sculpted Virtual Lenses

Hossein Baktash, Yash Belhe, Matteo Giuseppe Scopelliti, Yi Hua, *Student Member, IEEE*,
Aswin C. Sankaranarayanan, *Senior Member, IEEE*, Maysamreza Chamanzar

Abstract—Ultrasonically-sculpted waveguides provide interesting opportunities for in situ optical imaging in transparent and scattering media. The interference of ultrasonic waves can be designed to form a spatially-varying refractive index pattern in the target medium, acting as a virtual lens to guide light and relay images. The images formed by such lenses are subject to a large amount of spatially-varying blur, which significantly reduces their contrast. To alleviate this issue, the images can be computationally deblurred post experiment to restore the image. First, we demonstrate a brute force deconvolution technique to deblur the relayed images. While effective, this method proves to be computationally intensive due to the spatially-varying blur kernel. The reconfigurability of ultrasonically sculpted waveguides can be leveraged to address this issue by measuring line integrals of the image at multiple angles to form the Radon image of the target, which can be deblurred very efficiently with a simple linear model. We validate this method using simulated and experimental results. The tantalizing notion of using the reconfigurable ultrasonically-sculpted waveguides as part of the imaging systems, demonstrated in this paper, opens new opportunities for hybrid physical-computational optical systems.

Index Terms—Computational photography, Ultrasonically sculpted virtual waveguides, Spatially-varying blur, Deblurring

1 INTRODUCTION

PROPAGATION of light in a medium is governed by variations in its refractive index. Many optical elements, including gradient-index (GRIN) lenses, use this principle for confining and steering light. A recent technique leverages this principle to form virtual waveguides in a medium using ultrasound for confining light as well as relaying images from deep within the medium [1]. In this method, ultrasound waves interfere in the medium to produce pressure waves that locally compress and rarefy the medium and modulate its local refractive index, thereby creating spatially-varying refractive index patterns. A single GRIN waveguide can be sculpted by forming a uniform ultrasonic pressure standing wave in the medium. This phenomenon has been used to implement tunable external optical lenses, where the focal plane can be changed by changing the intensity of ultrasound [2]. The same principle has also been used to sculpt a virtual GRIN lens in the target medium to form a virtual GRIN lens for in situ relay imaging [3]. Ultrasound lenses are reconfigurable, which opens up numerous interesting opportunities, including the ability to actively refocus light onto different planes [3], as well as generating complex optical patterns [4]. These virtual ultrasonically sculpted lenses have been demonstrated for relay imaging and fluorescent microscopy [3].

The images acquired using ultrasound lenses exhibit significant aberrations, since the existing implementations are mainly based on using cylindrical ultrasonic transducers,

where the standing pressure waves are formed inside a cylindrical cavity such that the sculpted refractive index follows a Bessel-shape radial profile, leading to aberrations for the confined optical beam. Hence, the acquired and relayed images suffer from strong artifacts in the form of spatially-varying blur. This causes severe loss of contrast and image quality, posing challenges for the wider adoption of this technique for imaging applications.

This paper formulates a computational imaging approach for restoring images captured with ultrasonically-sculpted lenses. First, we discuss a brute-force approach to computationally deblur the relayed images formed by single-point focus virtual lenses. In this technique, we characterize the spatially-varying blur kernels of the virtual lens and then deconvolve the blur from the images similar to prior work on image deblurring [5]. While this brute-force approach is effective, it is computationally intensive and not scalable. Moreover, it requires precise measurement of the blur kernel at every point within the field of view, which is cumbersome and prone to modeling errors. Next, we show that more complex virtual lenses can be employed to perform optical measurements that would enable reconstruction of the image with less computational load. Specifically, we use a multi-element ultrasonic array to form elongated point spread functions (PSFs) that can integrate the relayed images along lines across the field of view. This has the effect of optically implementing a line-integral of the image in the focus plane. Rotating the the ultrasonic mode relative to the object allows us to rotate the PSF, thereby capturing a series of images that form the Radon transform of the target image. The aberrations in the measurements affect each line-integral identically. Therefore, each line-integral is computationally restored with a simple matrix-vector multiplication, followed by an inverse Radon transform. As a consequence, this approach is computationally more

- Baktash, Hua, Sankaranarayanan and Chamanzar are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, 15213.
E-mail: hbaktash, yhua1, saswin, mchamanz@andrew.cmu.edu
- Belhe is with University of California, San Diego.
- Scopelliti is with Cutera.
- Belhe and Scopelliti contributed to this work while at Carnegie Mellon University.

tractable than the first brute-force approach.

Contributions. We present a computational-experimental imaging technique to capture and reconstruct images using reconfigurable ultrasound lenses that suffer from spatially-varying blur. The contributions of this paper are:

- *Deblurring images from a single-point focus ultrasound lens.* First, we show results of successfully improving the quality of relayed images via careful characterization of the spatially-varying blur kernels and adopting a computational technique for deblurring the images.
- *Deblurring images from an ultrasound lens that performs line integrals.* Second, we discuss a framework for forming and deblurring the images that relies on using complex patterns of ultrasound pressure waves. Using these patterns, we experimentally measure a Radon transform of the target image, which can then be used to computationally reconstruct the target image efficiently.

These contributions are validated using both simulations and real experimental results.

2 RELATED WORK

We discuss the prior art related to our deblurring technique using ultrasonically reconfigurable virtual optical lenses. This spans work on deblurring, especially for scenarios where the blur kernel is spatially-varying, as well as the use of ultrasonically sculpted optical waveguides for light delivery and relay imaging.

2.1 Spatially-varying deblurring

Developing methods for efficient restoring of spatially-varying blur have been discussed in the literature since more than two decades old. Nagy and O’Leary [5] proposed an efficient algorithm of restoring a whole image with a spatially-varying blur with compact support, and analyzed the runtime, memory requirements, and conditioning of the problem in detail. They introduced a model for interpolating the spatially-varying PSF by performing a piecewise linear approximation using a sparse set of known samples. We adopt this approach in our work as well.

Recent advancements in spatially-varying deblurring have mainly focused on scenarios arising in the context of imperfect optics, motion and defocus blur. Our work falls in the former category of correcting spatially-varying blur from imperfect optics. Kee *et al.* [6] shows professional photo lenses have non-negligible blur kernels across different focus and aperture settings, the largest of which can be approximated as a Gaussian with a standard deviation of 3 pixels. Heide *et al.* [7] shows blur removal from images captured with simple lenses, where the blur kernels are 50 pixels large and vary as a function of the light wavelength. They introduced a cross-channel prior and presented an algorithm for finding the global optimum solution of the convex deblur problem. However, our imaging system is currently designed for operation with a monochromatic light source, and so we do not benefit from the use of cross-channel information. In the future, multi-wavelength imaging with these virtual ultrasonic lenses can be implemented to achieve and leverage cross-channel information.

The research on removing motion or defocus blur has been mainly concerned with estimation of the blur kernels, using local image statistics [8], [9], [10] and global smoothness constraints [11], as those spatially-varying blurs are scene dependent and different in each image. Blind deblurring is a challenging problem, especially when the PSFs are large, as is the case in our system. The virtual ultrasonically sculpted optical lenses in our system are fairly stable and repeatable. Therefore, we can pre-calibrate the system to obtain the blur kernels and, hence, avoid solving a blind deblurring problem.

2.2 Virtual ultrasonic optical waveguides

Accessing deep regions within the target medium for light delivery and imaging has always been challenging for external optical methods. Using external optics to confine light from outside to capture and relay images of the underlying scenes suffers from geometrical and physical limitations. To address some of these challenges, endoscopic imagers have been developed to be inserted into the medium to access deep structures. In these systems, a physical GRIN lens is usually implanted into the medium to collect light and relay images from the depth to be captured by an external imager on the surface. Inserting such a physical GRIN lens involves physical disruption of the medium. In some cases, such as imaging biological tissue, this disruption is invasive and undesirable. To minimize disruption of the target medium, these implantable imagers are miniaturized to realize miniscopes or micro-endoscopes. However, it would be helpful if such imagers could be implemented without an invasive component that penetrates through the medium.

Recent work has demonstrated that ultrasound can sculpt virtual waveguides in the target medium to act as a GRIN lens, effectively collecting and relaying images of the scenes to external microscopes without implanting a physical GRIN element [3]. It has been shown that this method works both in scattering and transparent media. These results open up new possibilities for imaging structures deep into the target medium, potentially for a whole gamut of applications such as biomedical imaging, machine vision, microscopy, and photography. However, a major issue is that the relayed images using these virtual ultrasound lenses sculpted in the target medium suffer from a substantial level of blur, mainly due to the aberrations of the sculpted lens. This outstanding issue needs to be addressed before the full potential of this technique can be harnessed.

3 ULTRASOUND FOCUSING OF LIGHT

In this section, we first explain the fundamentals of ultrasound-sculpted lenses. We then introduce the various types of such reconfigurable optical lenses created by operating a cylindrical ultrasonic transducer array at different modes. Finally, we discuss their corresponding PSFs.

3.1 Ultrasonically sculpted optical lenses

As mentioned earlier, when ultrasound waves are launched through the medium, the local refractive index changes. When a cylindrical array of piezoelectric transducers is

used to launch ultrasound along the radial direction, ultrasound waves interfere inside the cylindrical cavity to form a coherent pattern of standing waves. In this standing wave pattern, the density of the medium is increased in the locations with positive pressure (relative to the resting pressure state), whereas the medium becomes more rarefied in the negative pressure regions. According to the Lorentz-Lorenz relation [12], variations in the medium density affect its local refractive index, so the pressure waves ultimately modulate the local refractive index of the medium. Positive pressure increases the local refractive index. In contrast, negative pressure reduces the refractive index below the background level, i.e., the initial index of refraction for the medium at rest. The spatiotemporal distribution of the pressure standing wave pattern in a cylindrical cavity follows a Bessel profile. Therefore, for small pressure changes, the sculpted refractive index profile can be modeled in cylindrical coordinates as

$$n(r, \phi, t) = n_0 + \Delta n J_m(k_r r) \cos(m\phi) \sin(\Omega t), \quad (1)$$

where n_0 is the refractive index of the medium at rest, Δn is the refractive index contrast generated by ultrasound, J_m is the m^{th} -order of Bessel function of the first kind, k_r is the radial wavenumber, m is the azimuthal mode number, and Ω is the angular frequency of the ultrasound.

Therefore, when ultrasound is off (medium at rest), the refractive index of the medium is constant and homogeneous (Fig. 1(a)). When ultrasound elements are turned on, pressure waves modulate the medium creating a refractive index contrast, consequently sculpting a coherent pattern of refractive index in the medium, as shown schematically in Fig. 1(b). This spatial pattern of refractive index inside the medium acts as a waveguide that guides photons. Due to the gradual change of refractive index along the radial direction, the wavefront of the propagating light is modulated such that the guided light is gradually focused. Therefore, the sculpted refractive index pattern acts as a gradient-index (GRIN) lens that can relay the image of a target from the object plane to the image plane, analogously to a physical GRIN lens (Fig. 1(c)). This arrangement has been used to form tunable acoustic gradient (TAG) index of refraction lenses [2].

3.2 Complex ultrasonically sculpted patterns

It has been demonstrated that when different piezoelectric elements of a cylindrical transducer array are driven at specific combinations of amplitudes and phases, more complex patterns of refractive index can be sculpted into the medium, thus confining light in complex patterns [4]. When all the piezoelectric elements of the cylindrical array are driven at the same amplitude, frequency, and phase, the waves will interfere inside the cavity to generate a radially-symmetric standing pressure wave along the radial direction. During the positive half-cycle of the standing wave, the pressure waves interfere constructively at the center of the transducer, generating a peak of high pressure, surrounded by a ring of negative pressure (Fig. 2(a)). The spatial distribution of such a pressure pattern and the associated refractive index profile (Fig. 2(b)) can be described as a zeroth-order ($m=0$) Bessel function of the first kind (J_0). The

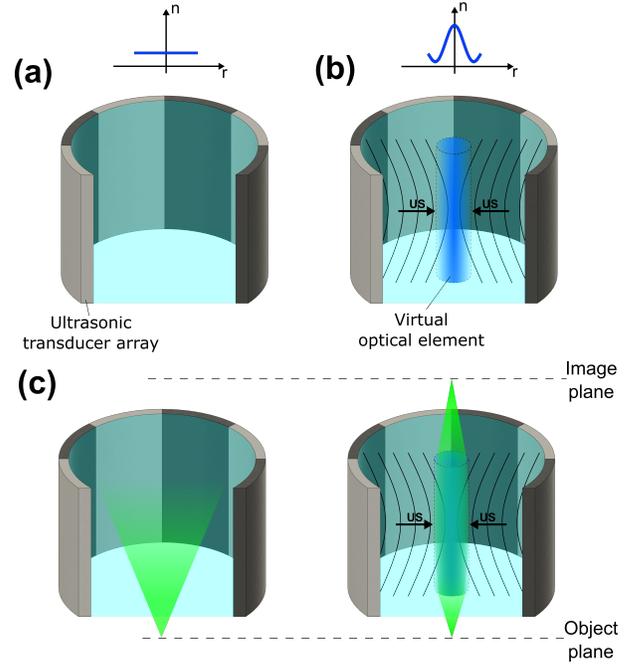


Fig. 1. An ultrasonic transducer array creates virtual optical waveguides in the medium. The medium is at rest when the ultrasonic transducer array is off (a). When ultrasound is turned on, a refractive index change is induced in the medium to form a virtual optical waveguide (b). Light from a point source is diverged and scattered when ultrasound is off (c). When ultrasound is on, the virtual waveguide acts as a lens, relaying the image from the object plane to the image plane (d).

simulated ray paths show that the resulting virtual GRIN lens confines a propagating beam of light at the center of the cavity, within the positive pressure peak. In this paper, we refer to this pattern as the zeroth-order mode, or simply **mode 0**.

By changing the relative phase and amplitude of the array elements, we can excite higher-order acoustic modes of the cylindrical cavity to generate more complex pressure patterns in the medium. For example, changing the relative phase between four different sets of pairs of adjacent elements by 180° , we can excite a quadruple mode, whose pressure distribution and refractive index profile are described by a second-order ($m=2$) Bessel function of the second kind (J_2) and are shown in Fig. 2(d-e), respectively, for the positive half-cycle of the standing wave. We can observe from the ray tracing simulation in Fig. 2(f) that light branches and converges toward the positive peak locations. We refer to this pattern as the second-order mode, or simply **mode 2**. While we restrict our experiments and analyses to these two complex interference patterns of ultrasound (i.e., mode $m=0$ and mode $m=2$) in this paper, there are additional spatial profiles that can be leveraged for the deblur analysis [4]. We have further discussed such patterns in the discussion section and in supplementary material.

We can see examples of measurements using mode 0 and 2 in Figures 3 and 5. In section 4 and 5 we show how we construct and deblur images using these measurements.

3.3 PSFs generated by different modes

Each mode of operation of our ultrasound lens produces a very specific PSF. The PSFs for a particular mode are

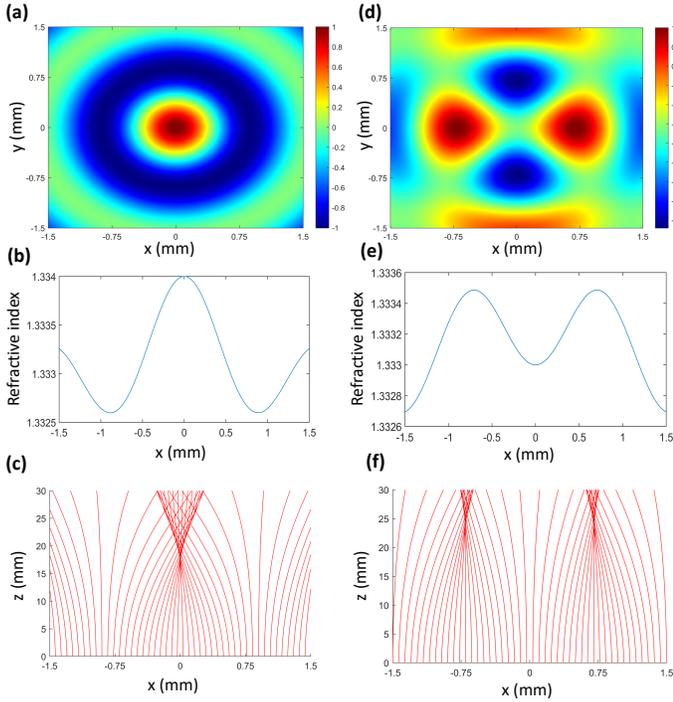


Fig. 2. Pressure profiles of two different modes and some ray tracing examples. (a) Pressure profile of mode 0, with normalized maximum intensity. (b) Horizontal cross-section of the pressure profile. (c) Ray tracing simulation of a collimated beam confined by the pressure pattern, with the assumption that the pressure profile remains invariant along the z direction. (d,e,f) same as a, b, c, but for mode 2.

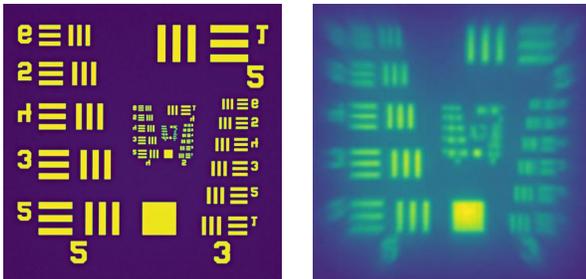


Fig. 3. A 0.5mm USAF target and its simulated mode 0 measurement.

determined by the sculpted refractive index profile for that mode. For example, the refractive index profile of the single-focus ultrasound lens (i.e., mode 0), shown in 2(b) results in a PSF that is drastically different from the PSF of the complex profile of the mode shown in 2(e). Fig. 4 shows the PSF of single-point focus of ultrasound lens, obtained by imaging a pinhole at different locations on the object plane.

Zeroth-order mode. The refractive index profile for mode 0 is radially symmetric. Therefore, the corresponding PSFs are also radially symmetric, as seen in Fig. 4. However, these PSFs are spatially varying. This suggests that mode 0 can be used as a lens with a spatially varying blur. We render the image of a USAF target relayed by an ultrasound lens with the mode 0 refractive index profile in Fig. 3.

The PSFs preserve the details at the center of the field of

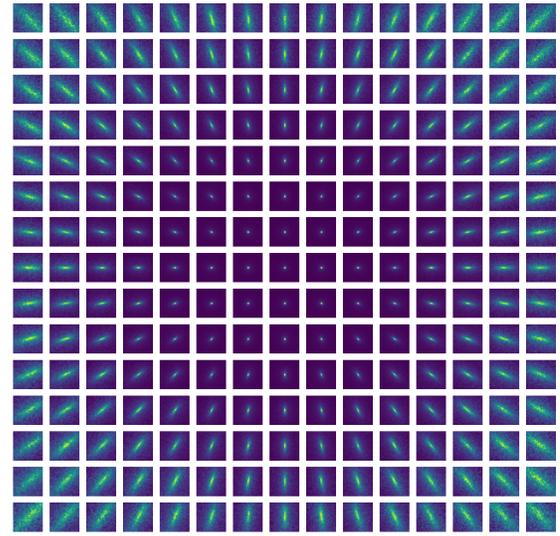


Fig. 4. PSFs within a 0.7mm wide field of view produced by mode 0 of the ultrasound lens.

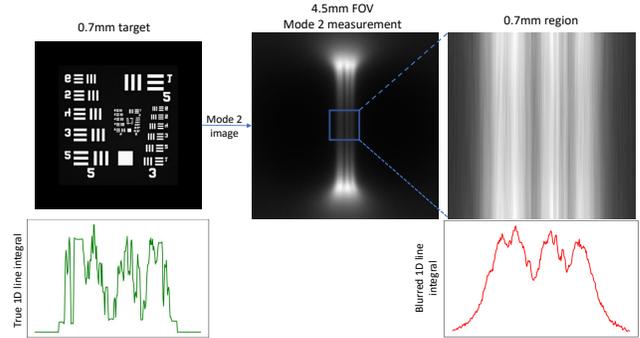


Fig. 5. Mode 2 measurement of a USAF target. The effect of positive pressure lobes working as focusing poles is shown in the 4.5mm FOV image. The region of interest for our imaging system is the central 0.7mm-wide region, which functions as a set of blurry line integrators.

view. However, the regions near the edges of the target are severely blurred and the contrast is low, making it hard to resolve the image in those regions.

Second-order mode. The refractive index profile for mode 2 is shown in Fig. 2. Some PSF samples are shown in Fig. 6(a). The red dots show the locations of the point sources. The central parts of the PSFs are magnified in Fig. 6(b). We should note that moving the point sources along the vertical direction will not change the PSF 6(c). Therefore, if we limit our field of view to the distance between the two focusing regions, we can see the transformation of the point source as a vertical line, shown in Fig. 6(c). This suggests that mode 2 can be viewed as a vertical line integrator of any target in the central region, which is confirmed in Fig. 5. This line integral, however, is not perfect and has a blur, which is spatially varying as shown in the 1D kernels in Fig. 6. We will show in section 5 that we can account for this blur and estimate the true line integral.

With these insights, we derive a spatially-varying deblurring technique for mode 0 in Section 4, and a reconstruction method based on Radon transform using mode 2

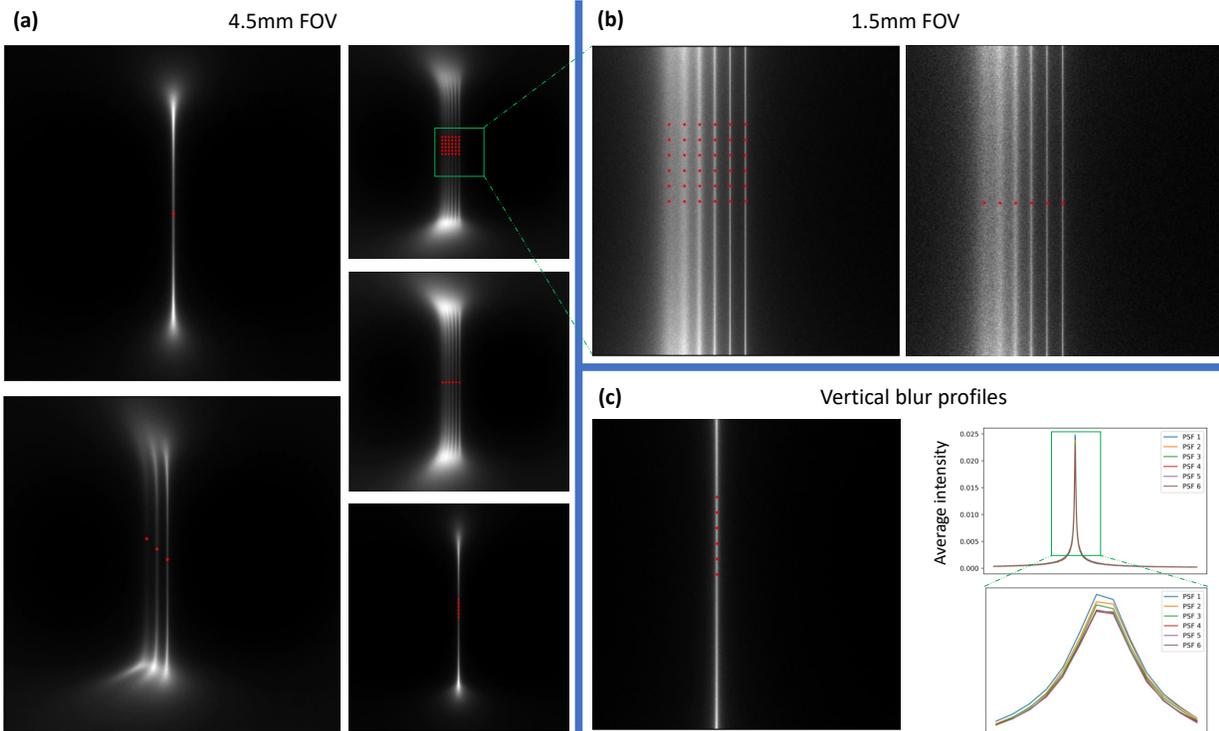


Fig. 6. PSF measurements in mode 2. Red dots represent the location of point sources. (a) PSF measurements in a large (4mm) field of view. Since the PSFs have a non-local support, we have provided multiple PSFs on single images. (b) PSFs viewed in a smaller region (1.5mm \times 1.5mm). (c) Sampled PSFs on a single column, visually confirming that blur kernels do not change along the vertical direction. Row-intensity plots confirm this.

measurements in Section 5.

4 DEBLURRING MODE 0 MEASUREMENTS

In this section, we introduce our forward model and associated multi-image deblurring technique for recovering a target being imaged by the ultrasound lens.

Consider a scene whose intensity function is given by the $n \times n$ image X . Given that we have a spatially-varying blur kernel, we denote PSF for point at location (i, j) on the input plane as an $k \times k$ matrix K_{ij} , where k is the size of the largest blur kernel. Let the resulting measurement be image Y . Since the image formation model is linear, if we vectorized Y and X to obtain \mathbf{y} and \mathbf{x} , respectively, the forward model can be written as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n} \quad (2)$$

The term \mathbf{n} denotes measurement noise that includes both signal dependant term modeling photon noise as well as a signal independent term modeling read noise. The matrix \mathbf{A} captures the operator associated with the forward model; each column of \mathbf{A} contains the vectorization of the K_{ij} , PSF for point (i, j) . Note \mathbf{A} is a large matrix, with $O(n^4)$ entries. Next, we explain how we populate entries of \mathbf{A} and how we recover the scene \mathbf{x} from this forward model.

4.1 Modeling spatially-varying blur

To fill out the entries of \mathbf{A} , ideally we should measure all PSFs by imaging pinholes at n^2 locations. This is clearly impractical for high resolution images with large n . Instead, we measure PSFs on a sparser $l \times l$ grid and use these to

estimate the entire stack of n^2 PSFs. Figure 4 shows an example of some measured mode 0 PSFs for $l = 11$. A comparison of different methods of PSF interpolation can be found in the supplementary material.

We use a piecewise linear approximation for interpolating the blur kernels. The blur kernel at each location is linearly interpolated from the four closest locations in the $l \times l$ grid. While previous research [7] opted for a piecewise constant approximation of the spatially-varying blur, we found our larger blur kernels fit the piece-wise constant approximation poorly. Please refer to the supplementary material for a detailed evaluation of the two interpolation approaches.

4.2 Deblurring

To deblur the relayed images by the single-focus ultrasound lens (i.e., mode 0), we will use the following algorithm:

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \underbrace{\|\mathbf{y}_0 - \mathbf{A}\mathbf{x}\|_2^2}_{\text{data term}} + \underbrace{\lambda \|\mathbf{x}\|_2^2}_{\text{regularization term}}$$

where \mathbf{y}_0 corresponds to measurements with mode 0. We use the $\text{TV}(\cdot)$ regularization term to denote the anisotropic total variation operator. The term λ is weight for controlling the relative importance of the two loss terms.

Implementation details. Evaluating the forward model of 2D spatially-varying blur with large PSF is expensive. Storing PSFs at all locations is introduces a high memory complexity, while interpolating PSF from sparse observations increase the time complexity as well as results in a less accurate model. Thus, constrained by the computational

complexity, we resort to small image sizes ($n = 200$) and store the blur matrix in an $n^2 \times n^2$ matrix.

We solve the deblur problem using Scipy linear solvers (GESV routine) to solve the pseudo inverse problem with ℓ_2 norm. Although this requires a high memory complexity to store $\mathcal{O}(k^2 n^2)$ non-zero entries of matrix \mathbf{A} (12.1 GB of memory).

Simulation results. For simulating the measurements captured with the ultrasound lensing system, we use a physically-accurate Monte-Carlo renderer [13] that can ray trace in the presence of spatially-varying refractive index patterns. This renderer can handle the refractive index patterns associated with both mode 0 and mode 2. To avoid rendering noise, we typically render measurements with a very large number of Monte-Carlo samples, and then add physically-accurate photon noise as well as read/thermal noise.

Setup. Our ultrasound lens simulator mimics the setup shown in 1. We place $500\ \mu\text{m} \times 500\ \mu\text{m}$ targets at the object plane and our imaging sensor at the image plane. The relayed image has $\sim 7\ \mu\text{m}/\text{pixel}$ resolution. We set the medium of transducer to water, and turn the transducer on, which creates standing pressure waves in it at a frequency of 840kHz and a maximum refractive index contrast (Δn) of 0.0021. We capture images formed at the positive peak region of the pressure profile. We also assume that measurements are integrated over a very short period of time; since the $\sin(\Omega t)$ term in 1 introduces a periodic behaviour to our pressure profile and integrating over a large period of time introduces haze to our measurements (one can imagine that at certain times the pressure profile shown in Figure 2 can even be zero, and no focusing will happen). We can assume that this happens either with a very short exposure or a infinitesimally small duty cycle light source. As we discuss in section 7, our experiments were carried out using a light source with a 25% duty cycle modulation.

Our simulated measurements are passed through a simulated camera pipeline, where we add Poisson and read noise to the measurements. Since our real world measurements are captured using a FLIR Blackfly camera (BFLY-U3-50H5M-C), we add these sources of noise according to the well capacity and read noise of this device.

A set of measurements and deblurring results for Mode 0 are shown in Fig. 7. We observe that the computationally deblurred results have higher contrast and brings out small details in the scene. We also quantify the reconstruction quality by evaluating its SNR as follows: given \mathbf{x} , the ground truth, the reconstruction $\tilde{\mathbf{x}}$ has a SNR

$$10 \log_{10} \left(\frac{\|\mathbf{x}\|_2}{\|\mathbf{x} - \tilde{\mathbf{x}}\|_2} \right). \quad (3)$$

The same measure is also used later for quantifying the amount of noise.

While this method enhances the quality of images captured by the single-focus ultrasound lens (i.e., mode 0), it is not easily scalable for two reasons. First, deconvolving 2D spatially-varying blur with non-compact kernels is computationally-intensive as solution requires repeated evaluation of the forward process. Additionally, the performance of this deblur method heavily depends on precisely

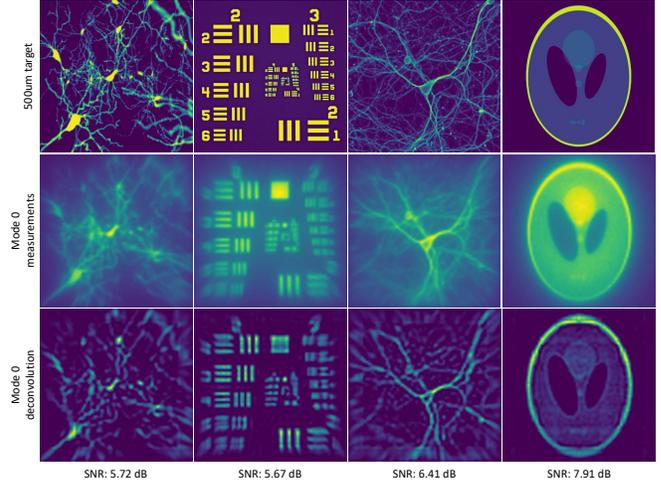


Fig. 7. Rendered measurements of several targets all with a 0.5mm size. Rendered PSF measurements are used for the deconvolution process, resulting in the bottom row images.

modeling the PSF of all points in the field of view. Since the PSF varies a lot across the field (as shown in Figure 4), even with linear interpolation, we need to image dense enough $l \times l$ grid of pinhole locations. Even for coarse grid $l = 11$, this requires over a hundred times precise translations of a pinhole and capture of images at each location. This calibration process needs to be done for each new combination of the ultrasonic transducer array and medium, and is obviously impractical for practical applications. This motivates us to leverage the reconfigurability of ultrasound lenses to deblur images more effectively.

5 RECONSTRUCTING IMAGES FROM MODE 2 MEASUREMENTS

The ultrasound lens at mode 2 exhibits a complex pressure profile with a quadruple pattern as shown in Fig. 2(d-e). At the positive peak pressure of ultrasound, only two lobes of the quadruple will be in phase and at positive peak, forming a refractive index profile that has two adjacent peaks that are separated by a weaker high refractive index region 2e. When a collimated beam of light is confined through this complex lens, the output light structure will be in the form of a line with two bright corners (Fig. 6). When using this complex PSF to image the target, all the pixels along the elongated PSF will be integrated to generate a single intensity data point. This property of mode 2 can be leveraged to form a Radon transform image of the target. In this section, we introduce our forward model for mode 2 and the associated reconstruction technique to restore the target image.

As shown in Fig. 6, imaging the scene \mathbf{x} with mode 2 produces vertical line integrals of the scene. The central horizontal slice of the measurement, $\mathbf{I}_{\theta=0}$, is given by

$$\mathbf{I}_{\theta=0} = \mathbf{A}_{1D} \mathbf{R}_{\theta=0} \mathbf{x}, \quad (4)$$

where R_{θ} denotes integration over a line, passing through the center of the ultrasonic transducer array and at angle $\theta = 0$ to the y -axis; A_{1D} denotes a one-dimensional spatially-varying blur, to account for the changes in PSF as the point

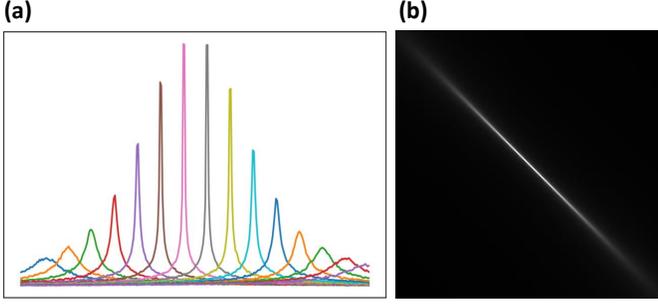


Fig. 8. (a) Several blur kernel samples across a 0.7mm region. (b) 300×300 Blur matrix across the 0.7mm field of view. Every column is a blur kernel measured at the same location using the renderer.

moves orthogonal to the direction of the line integral, as shown in Fig. 6(b).

When we rotate the pressure profile by angle θ_i , the direction of line integral changes. But we can still slice the measurement in the direction orthogonal to the line integral, L_{θ_i} , to obtain

$$l_{\theta_i} = \mathbf{A}_{1D} \mathbf{R}_{\theta_i} \mathbf{x}. \quad (5)$$

Note that R_{θ} reduces the two-dimensional scene to a one-dimensional signal by performing a line integral. R_{θ} and l_{θ} are vectors, and we can stack multiple measurements taken at $\Theta = \theta_1, \dots, \theta_m$ into a large system,

$$[l_{\theta_1} \ \dots \ l_{\theta_m}] = \mathbf{A}_{1D} [\mathbf{R}_{\theta_1} \mathbf{x} \ \dots \ \mathbf{R}_{\theta_m} \mathbf{x}], \quad (6)$$

and arrive at our forward model for mode 2,

$$\mathbf{L}_{\Theta} = \mathbf{A}_{1D} \mathbf{S}_{\Theta}, \quad (7)$$

where \mathbf{L}_{Θ} collects the measurement slices and \mathbf{S}_{Θ} is the Radon transform of \mathbf{x} evaluated at angles Θ . For simplicity, we will drop the subscript Θ from here on. In simulations, we always use 180 equally distributed rotation angles from 0 to π and in experiments we use 150 different angles.

We leverage the configurability of the ultrasonic transducer array to produce measurements that allow this formulation of the forward model at mode 2, as it allows us to use simple tools for fast reconstruction, such as one-dimensional deblurring and inverse Radon transform.

5.1 Modeling spatially-varying blur

As mentioned before, blur kernels do not change by moving the point source along the line of integration (vertically in Fig. 6). So the blur matrix \mathbf{A}_{1D} can be sampled just by using point sources on a row. In fact, since \mathbf{A}_{1D} is a 1D blur with dimensions $n \times n$, applying the forward model is very efficient. This also enabled us to sample blur kernels in experiments and use experimental measurements to build matrix \mathbf{A}_{1D} . We should note that we are limited to sparse sampling of measurements in experiments.

We rendered images of point sources in the central row of target across all columns. Some samples of the 1D blurs at different locations as well as the resulting blur matrix \mathbf{A}_{1D} are shown in Fig. 8.

5.2 Deblurring

To reconstruct the target image from the different measurements, we use filtered back projection after deblurring the Radon transform image S from the collection of mode 2 measurements \mathbf{L} :

$$\mathbf{S}^* = \arg \min_S \|\mathbf{L} - \mathbf{A}_{1D} \mathbf{S}\|_2^2 + \underbrace{\lambda \|\mathbf{S}\|_2^2}_{\text{L2 norm per column}} \quad (8)$$

After obtaining the Radon image \mathbf{S}^* , we can use filtered back projection to reconstruct image \mathbf{x} .

We implement the forward model by building the linear operator \mathbf{A}_{1D} and using linear solvers to solve the pseudo inverse problem with ℓ_2 norm. To reconstruct image \mathbf{x} , we use *iRadon* implementation with "hann" filter provided in NumPy library. The matrix \mathbf{A}_{1D} is $n \times n$, and \mathbf{L} is $n \times |\Theta|$. So unlike mode 0, solving the linear problem in equation 8 is efficient and fast and is not memory-intensive. Moreover, rendering blur kernels at various locations only requires $\mathcal{O}(n)$ rendered images, with not so many required ray samples. The latter is due to the fact that for a given point source (e.g., in Fig. 6), one can average multiple rows to obtain the 1D blur kernel, hence reducing the effect of rendering noise.

Simulation results. We use the same setup described in simulations section of 4.2 to simulate mode 2 measurements and image reconstruction. In the case of mode 2, reconstruction of the target image from the measurements involves two steps of (i) deconvolving the blur kernels from the Radon transform images and (ii) computing the inverse Radon transform. In these simulations, we assumed $\Delta n = 0.0047$. A set of mode 2 reconstructions from the simulated measurements are shown in Fig. 9 for different targets. In this figure, we show the target image, the mode 2 Radon transform measurements, the deconvolved measurements and the reconstructed images after the inverse Radon transformations. In each case, we show the results for inverse Radon transform applied to the raw measurements as well as the deblurred measurements to demonstrate the effect of deblurring before the inverse Radon transform computation.

6 ANALYSING SIMULATION RESULTS

Even though we used different experimental and computation methods for mode 0 and mode 2 images, the results can be compared from the viewpoint of different performance metrics.

Performance under varying noise levels. We change noise levels in our simulated measurements by adding different levels of read noise (from a Gaussian distribution) and photon shot noise (from a Poisson distribution). We should note that each reconstructed image is obtained from 180 different mode 2 measurements at different rotation angles. On the other hand, mode 0 images are reconstructed from a single image. To make a fair comparison between mode 0 and mode 2 reconstructed images, we should take into account the overall exposure time in each case. We scale mode 2 measurement intensities by a factor of $\frac{1}{180}$, so that the total exposure time of all mode 2 measurements equals

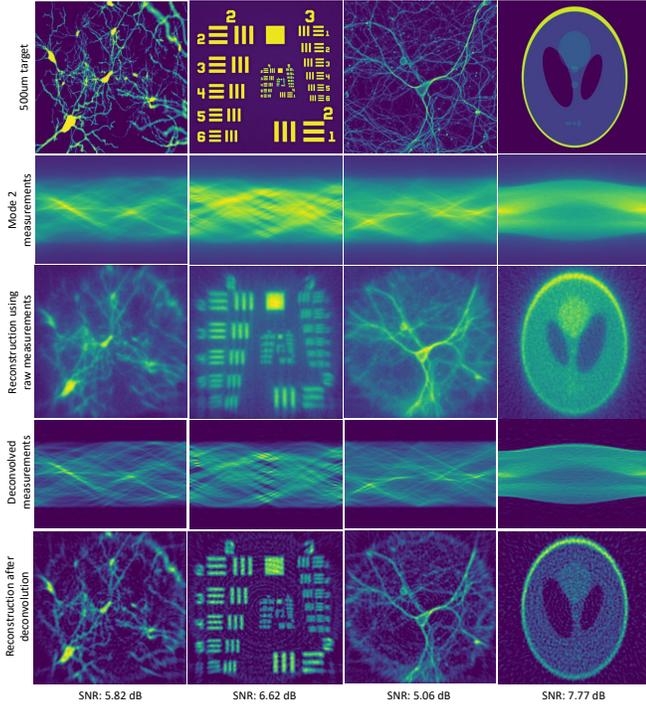


Fig. 9. Rendered measurements of different targets all with a 0.5 mm size each captured at 180 different rotation angles. Each column of the measurements is deconvolved and used in inverse Radon transform. Reconstruction SNR values of are reported below each reconstructed image. SNR of raw reconstructions is less than 1 dB for all images.

the exposure time of a single mode 0 measurement. Note that this only effects the photon noise. With this, we obtain measurements at different read noise levels.

Fig. 10 shows the reconstruction SNR as a function of different noise levels for mode 0 and mode 2. From these plots, we observe that mode 2 has a better performance compared to mode 0 in terms of the reconstruction SNR at different noise levels. The better performance of mode 2 under different noise levels can be attributed to the fact that the line integral of interest in a single mode 2 measurement is repeated along all the rows in that measurement as shown in Fig. 5 and 6. Therefore, taking an average of signals over multiple rows would result in noise reduction, since all of the rows are essentially representing the same line integral of interest.

Dependence on field of view. As we can observe in Fig. 4, mode 0 PSFs in the central regions of the field of view are sharper and more compact and the blur increases in regions further away from the center. This shows that more information is preserved in the central high pressure regions of mode 0, which is confirmed via simulations in Fig. 12. The same argument holds for mode 2 measurements of the blur kernels. As we can see in the kernel matrix image in Fig. 8 and sample PSF measurements in Fig. 6, in the central columns we have less blur. Hence, reconstruction of the image using mode 2 measurements is most accurate over an effective field of view and yields better reconstruction SNR in the central regions. Fig. 12 shows this effect for a target of size 500 μm with a central region of size 300 μm .

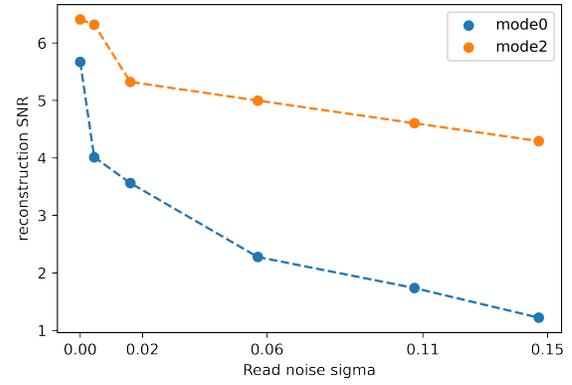


Fig. 10. Reconstruction SNR of the USAF target images obtained using mode 0 and mode 2 at different noise levels.

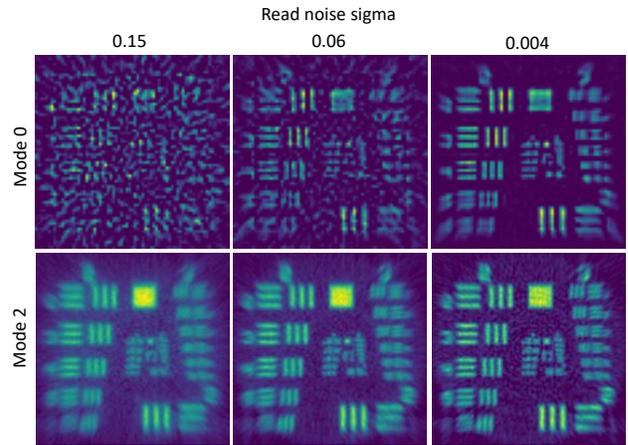


Fig. 11. Some of the deblurred targets at different noise settings. Mode 0 suffers from speckle-like artifacts at low SNR levels. Mode 2 does not get affected as much, due to the row-averaging procedure in every measurement.

Performance under different signal priors. Fig. 13 shows some samples of reconstructed images with and without TV regularization. When solving with small regularizing constant λ , we see speckle-like patterns in both cases. However, the patterns in mode 2 speckle patterns are almost radially structured due to using deconvolution on the acquired mode 2 measurements before using inverse Radon transform. In mode 0, when using very large λ , structured artifacts appear. Also in mode 2, circular artifacts appear in image reconstruction, likely because we end up with a Radon transform image that is strongly directed towards the direction of the nullspace of Radon transform.

7 EXPERIMENTAL RESULTS

7.1 Hardware setup

For the experimental demonstration, we have designed and implemented the experimental setup schematically shown in Fig. 14. The real image of the setup is shown in Fig. ???. To generate the virtual relay lens, we used a custom-designed 8-element piezoelectric transducer array made of PZT-5A, arranged around a 30mm long cylindrical geometry with

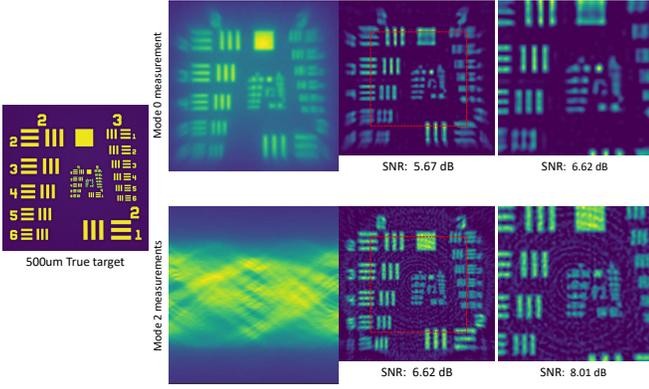


Fig. 12. Ground truth and measurements for the $500 \mu\text{m} \times 500 \mu\text{m}$ target. We see that both techniques have a better reconstruction SNR in central regions.

an inner diameter of 38mm (APC International, Ltd.), placed in a glass tank filled with deionized water. In this paper, the transducer elements was driven by a 35 V signal at 982 kHz for mode 0 and 52 V at 984.5 kHz for mode 2.

A diffused target, consisting of a negative transparency mask (from Front Range Photomask) overlaid on a speckle reducer (LSR-3005, Optotune, Inc.), was located at the base of the transducer array. The entire target assembly was placed on top of a motorized rotation stage (PRM1Z8-Ø1, Thorlabs, Inc.) to accommodate rotations required for mode 2 measurements. The target was illuminated by a red laser ($\lambda_{\text{light}} = 641\text{nm}$) from the bottom, after going through the rotation stage and being diffused and de-speckled by the speckle reducer. The images relayed by the virtual lens were captured in the transmission mode, using a top imaging system composed of a zoom lens (VZM 600i, Edmund Optics, Inc.) and a monochrome CCD camera (BFLY-U3-50H5M-C, FLIR Systems, Inc.). A clear glass optical window (WG11050-A, Thorlabs, Inc) was immersed in water and used as the interface between the external imaging system and the medium.

7.2 Calibration

The transducer was operated using a sinusoidal waveform over a frequency range of 980-985 kHz. Also the waveform was used in the burst mode with an On cycle of 1 ms and 4 ms Off cycle. The light source (red laser) was also synchronized with the transducer, both in terms of the burst period and also within each sinusoidal cycle. The laser was modulated using a square pulse train with a duty cycle of 25%.

7.3 Test targets

We used transparency masks with custom patterns as our test targets. We also used real biological tissue structures as the targets. Microscope image of the targets are shown in Fig. 15.

7.4 Imaging results

7.4.1 Mode 0 results

For mode 0 measurements, we used rendered PSFs to deconvolve from the measurements to sharpen the images.

A set of results are provided in Fig. 15. The measurements are translated to the center of the ultrasound lens focus and downsampled to match the PSF resolutions of $100\text{px} \times 100\text{px}$. Note that deconvolving larger sets of sampled PSFs for higher resolution measurements is not practical due to memory and time complexity of the deconvolution analysis for mode 0. Moreover, using a sparser set of sampled PSFs results in square-shaped artifacts similar to what is shown in Fig. 13. These artifacts also show up more often in non-central regions of the target, likely due to larger blur kernels, resulting in higher noise in measurements.

To make interpolations of PSFs robust, we need to measure PSFs at many grid points within the field of view, which would render the experiments impractical, especially for imaging large targets.

7.4.2 Mode 2 results

For Mode 2 measurements, we were able to capture experimental samples of blur kernels, due to simplicity of the model and alignment. However, for constructing the 1D blur kernel matrix, it is not practical to sample the whole matrix, unlike simulations. So we sampled the blur kernel at 40 different locations across the field of view. To prevent artifacts in image reconstruction similar to the artifacts in Fig. 13 and to account for measurement noise, we used parametric interpolation for our 1D blur kernels. The following Gaussian curve was used as our model:

$$K_{x_0}(x) = a \exp\left(-\frac{(x - x_0)^2}{2\sigma}\right) + c$$

Using the fitted curves at all desired locations yields the blur matrix shown in Fig. 16.

Our mode 2 experimental measurements are obtained by rotating the target at 150 equally spaced rotation angles from 0 to 180 degrees. In contrast to mode 0, these mode 2 measurements are not down-sampled and the deconvolution and reconstruction processes only take a few seconds even for the higher resolution measurements. The results are provided in Fig. 17. We can see that unlike mode 0 results, our mode 2 reconstructions do not suffer from artifacts and appear to have better quality.

8 DISCUSSION

In this paper, we demonstrated how single-point and multi-point focus ultrasonically sculpted optical lenses can capture, relay and process optical images of the target. In particular, we showed experimental and computational results for imaging using two different ultrasound interference patterns, generated by an array of ultrasonic transducers in a cylindrical geometry. The first pattern is a radially-symmetric profile that forms a virtual lens that can focus light onto a single focal region, similar to a traditional optical lens. However, because of the imperfect spatial refractive index profile, imaging using such a lens results in aberrations and loss of contrast. We showed that the blur can be deconvolved from the image if the blur kernels are measured across the field of view. The second pattern is a quadruple that is formed when the ultrasonic transducers are operated at a certain sequence of phase and amplitudes. The measurements using this pattern are then used to form

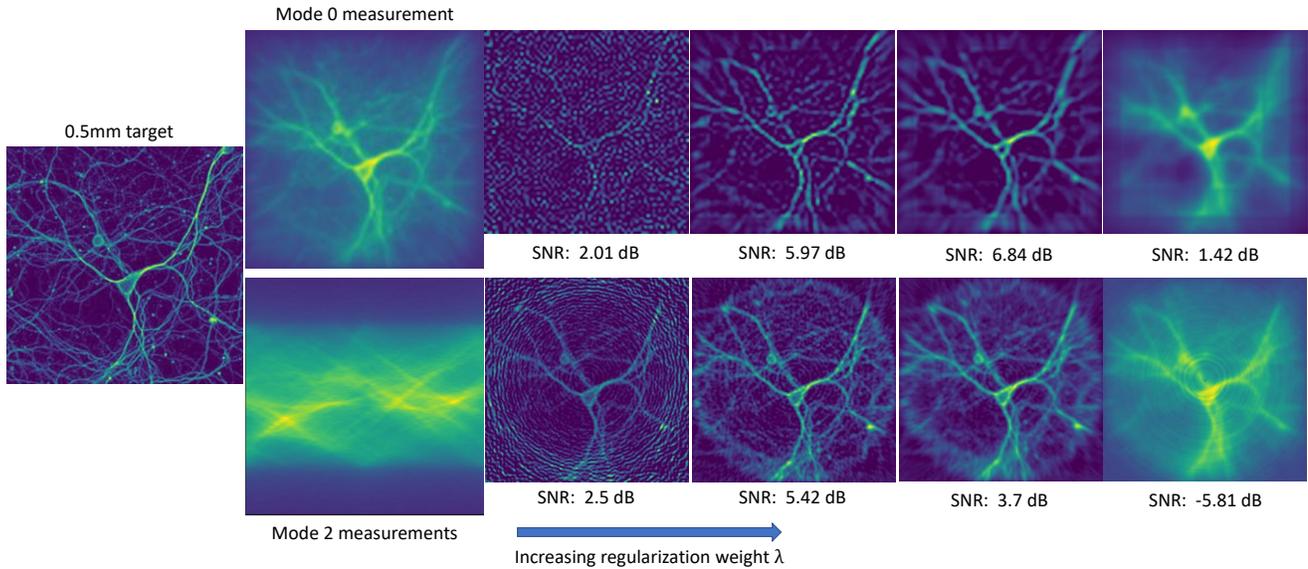


Fig. 13. Reconstruction with different on signal prior weights.

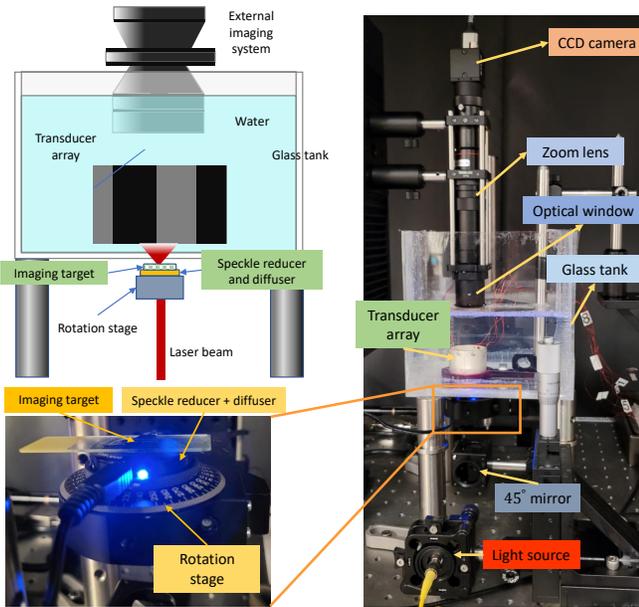


Fig. 14. Schematic and images of our experimental setup.

the Radon image of targets, which can be used to reconstruct images. This technique is scalable, and the computational load required to deblur and reconstruct images is not as intensive as the computational load of deblurring images obtained using mode 0.

8.1 Limitations

Deblurring and image acquisition. Image reconstruction using mode 0 is straightforward. However, it requires sampling PSFs at many locations and, therefore, it is computationally expensive. Instead, mode 2 gives us a robust and computationally efficient technique, but it requires collecting measurements at many different rotation angles. While the exposure time can be reduced, the main limitation for

measurement speed to cover all the different rotation angles is the relative rotation of the pattern with respect to the object. In our setup (8-element transducer array) this was performed using a motorized mechanical rotation stage. However, relying on moving elements limits the speed of experiments. In the future, ultrasonic transducer arrays with a larger number of elements (e.g., 300) can be fabricated, enabling to rotate the pattern electronically by switching between different elements of the array. As a result, we estimate that all the measurements could be performed in about 100 ms or less.

Geometric limitations. The cylindrical array of ultrasonic transducers is a limiting geometry, especially for in-situ imaging scenarios, where the target medium is supposed to be within the transducer cavity. In such scenarios, a surface array of ultrasonic transducers can be used to launch traveling ultrasound waves into the target medium and form the pressure pattern of interest. In such scenarios one can pulse the light source with a low duty cycle, similar to what we have done in this paper, to only see the effect of traveling wave, when it is at specific locations in the medium at a specific moment in time. This way, the travelling wave can be assumed to be spatially stationary at that specific moment in time.

Optical scattering. It has been shown that ultrasound can still effectively shape the trajectory of light in optically thick scattering media [3], [4]. Further studies in [14] also show in details that ultrasonic waveguides increase light throughput in scattering media. Rendering scenes in the presence of scattering is also possible [13]. Also, a blind deconvolution can be applied to deal with imprecise calibration as well as changes in the properties of the scattering tissue, based on the commonly used techniques in blind deconvolution literature. One can potentially use the available PSFs, either from renderings or measurements in a homogeneous artificial medium with similar properties as tissue to initialize an iterative optimization process.

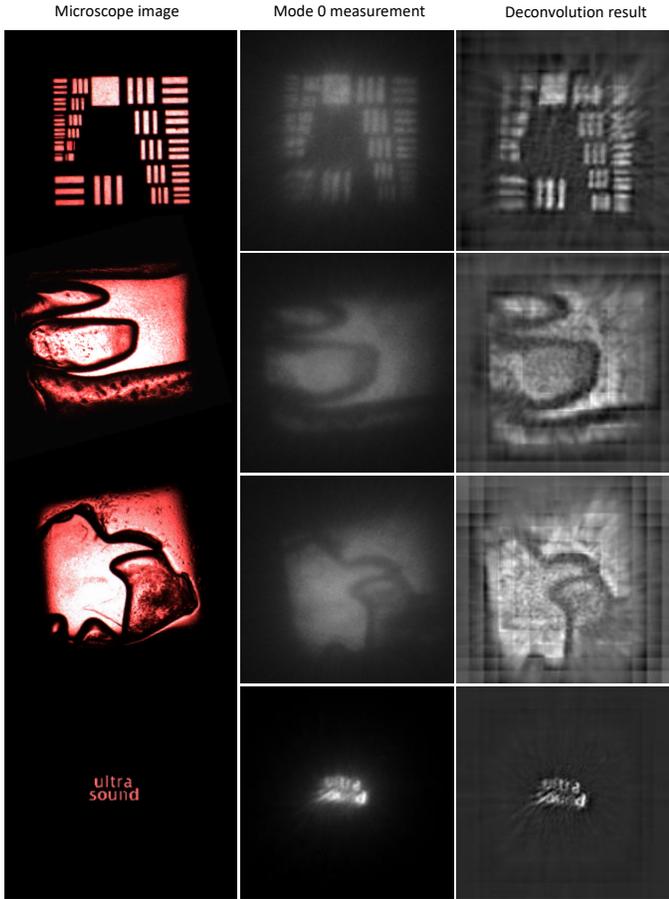


Fig. 15. Mode 0 experimental results. Target sizes vary between 0.5mm and 0.7mm. A USAF target and a multilayer 300 μ m thick mouse brain slice were used as targets. The latter were selected to have low and high contrast regions.

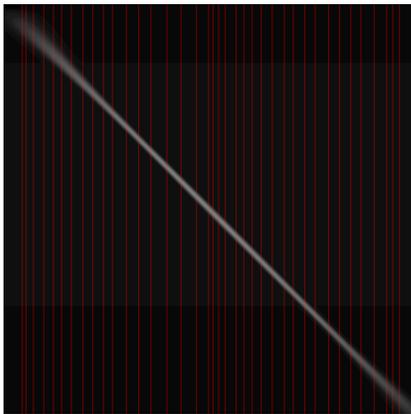


Fig. 16. Mode 2 parametric interpolated blur matrix on a 1mm \times 1mm field of view. Red lines show the sampling locations of the blur kernels.

8.2 Future directions

Other pressure patterns. While in this paper we only used mode 0 and mode 2, the tantalizing notion of using complex ultrasonic interference patterns to perform optical measurements that go hand in hand with computations can be extended to other spatial patterns. In other words, it is quite conceivable to design ultrasonic phased arrays to form the

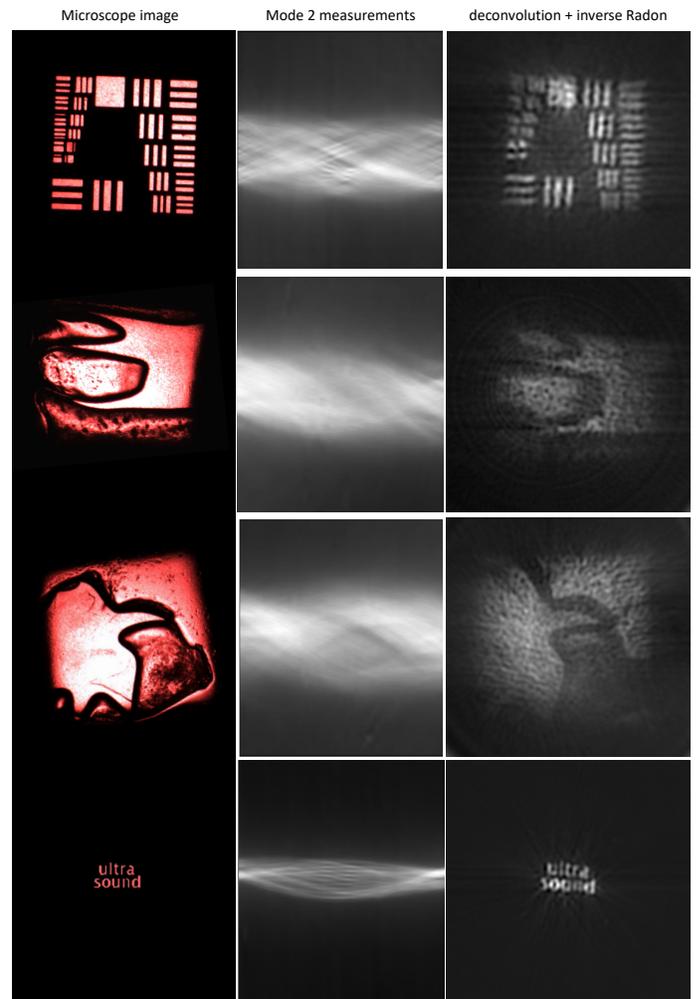


Fig. 17. Mode 2 experimental results with target sizes varying between 0.5mm and 0.7mm.

desired patterns that can perform more complex transformations of the image. In this paper, we used a small array of ultrasonic transducers (8 elements) arranged in a cylindrical geometry to produce the PSFs corresponding to mode 0 and mode 2. We also explored mode 1 pattern (first-order Bessel function). Some sample PSFs and measurements are provided in the supplementary material. Using the same array, we can generate higher order mode patterns, when changing the phase, amplitude or frequency of the array elements. An immediate example is shown in Figure 18. These time-varying patterns can be obtained by operating the transducer elements at a frequency range of 971-973 KHz with the same phase sequence and amplitudes used for mode 2.

Volumetric scanning. A volume of the target medium can be imaged by scanning multiple focal planes, simply by changing the ultrasound intensity. This capability has also been demonstrated in [3] and [4].

These features can also be extended by using a larger number of array elements in the same cylindrical geometry or by changing the geometry of the transducer arrays, for example, by arranging them in a non-cylindrical geometry. This rich design space enables engineering the system to

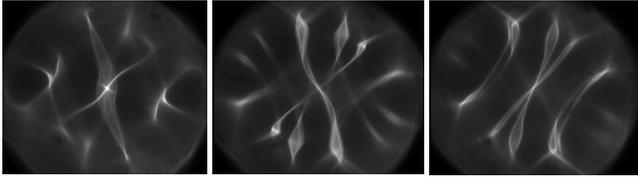


Fig. 18. Single point source measurements (at the center of FOV) acquired by operating mode 2 over the new frequency range of [971,973] KHz. Field of view is $2\text{mm} \times 2\text{mm}$.

produce the desired PSFs to perform more complex experimental processing during optical measurements and produce rich priors for computations.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation (NSF) under Grants 1935849 and 1730147. This work used the Bridges-2 system, which is supported by NSF award number OAC-1928147 at the Pittsburgh Supercomputing Center.

REFERENCES

- [1] M. N. Cherkashin, C. Brenner, G. Schmitz, and M. R. Hofmann, "Transversally travelling ultrasound for light guiding deep into scattering media," *Communications Physics*, vol. 3, no. 1, pp. 1–11, 2020.
- [2] A. Mermillod-Blondin, E. McLeod, and C. B. Arnold, "High-speed varifocal imaging with a tunable acoustic gradient index of refraction lens," *Optics letters*, vol. 33, no. 18, pp. 2146–2148, 2008.
- [3] M. G. Scopelliti and M. Chamanzar, "Ultrasonically sculpted virtual relay lens for in situ microimaging," *Light: Science & Applications*, vol. 8, no. 1, pp. 1–15, 2019.
- [4] Y. Karimi, M. G. Scopelliti, N. Do, M.-R. Alam, and M. Chamanzar, "In situ 3d reconfigurable ultrasonically sculpted optical beam paths," *Optics express*, vol. 27, no. 5, pp. 7249–7265, 2019.
- [5] J. G. Nagy and D. P. O'Leary, "Restoring images degraded by spatially variant blur," *SIMA Journal on Scientific Computing*, vol. 19, no. 4, pp. 1063–1082, 1998.
- [6] E. Kee, S. Paris, S. Chen, and J. Wang, "Modeling and removing spatially-varying optical blur," in *IEEE International Conference on Computational Photography (ICCP)*, 2011, pp. 1–8.
- [7] F. Heide, M. Rouf, M. B. Hullin, B. Labitzke, W. Heidrich, and A. Kolb, "High-quality computational imaging through simple lenses," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 5, pp. 1–14, 2013.
- [8] A. Levin, "Blind motion deblurring using image statistics," in *Advances in Neural Information Processing Systems*, 2007, pp. 841–848.
- [9] A. Chakrabarti, T. Zickler, and W. T. Freeman, "Analyzing spatially-varying blur," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2010, pp. 2512–2519.
- [10] X. Zhu, S. Cohen, S. Schiller, and P. Milanfar, "Estimating spatially varying defocus blur from a single image," *IEEE Transactions on image processing*, vol. 22, no. 12, pp. 4879–4891, 2013.
- [11] F. Couzinie-Devy, J. Sun, K. Alahari, and J. Ponce, "Learning to estimate and remove non-uniform image blur," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 1075–1082.
- [12] M. Born and E. Wolf, *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013.
- [13] A. Pediredla, Y. K. Chalmiani, M. G. Scopelliti, M. Chamanzar, S. Narasimhan, and I. Gkioulekas, "Path tracing estimators for refractive radiative transfer," *ACM Transactions on Graphics (TOG)*, vol. 39, no. 6, pp. 1–15, 2020.
- [14] A. Pediredla, M. G. Scopelliti, S. Narasimhan, M. Chamanzar, and I. Gkioulekas, "Optimized virtual optical waveguides enhance light throughput in scattering media," 2021.



Hossein Baktash is a PhD student in ECE department at Carnegie Mellon university. He received his B.Sc with a major in Computer Engineering and a minor in Mathematics from Sharif University of Technology, Tehran, Iran in 2020. His current research is in computational imaging with a focus on acousto-optics.



Yash Belhe is a PhD student in the CSE department at the University of California, San Diego. He received his B.Tech in Electronics and Communication Engineering from the Indian Institute of Technology, Guwahati in 2019, and his MS in Electrical and Computer Engineering from Carnegie Mellon University in 2020. His research interests are in computer graphics, specifically differentiable rendering and also in computational photography.



Matteo Giuseppe Scopelliti is a Sr. Biomedical Engineer. He received his PhD in Electrical and Computer Engineering from Carnegie Mellon University, PA, USA, in 2021, and his MS in Biomedical Engineering from Polytechnic University of Turin, Italy, in 2015. His research interests are in acousto-optics, light-tissue interaction, and biomedical imaging.



Yi Hua is a Ph.D. student in ECE Department at Carnegie Mellon University. She received her B.Sc in computer science in 2015 from Rice University, Houston, TX, USA, and M.S. in Computer Vision from Carnegie Mellon University, PA, USA in 2016. Her research interests are computational imaging and computer vision.



Aswin C. Sankaranarayanan is a professor in the ECE department at CMU, where he is the PI of the Image Science Lab. His research interests are broadly in compressive sensing, computational photography, signal processing and machine vision. Aswin's doctoral research was in the University of Maryland where his dissertation won the distinguished dissertation award from the ECE department in 2009. Aswin is the recipient of best paper awards at CVPR 2019 and ICCP 2021, the CIT Dean's Early Career Fellowship, the Spira Teaching award, the NSF CAREER award, the Eta Kappa Nu (CMU Chapter) Excellence in Teaching award, and the Herschel Rich Invention award from Rice University.



Maysamreza Chamanzar is the William D. and Nancy W. Strecker Associate Professor of Electrical and Computer Engineering at Carnegie Mellon University, where he is running an interdisciplinary research program at the interface of photonics and neural engineering to design next generation multimodal neural interfaces. Maysam is the recipient of the NSF CAREER Development Award, Spira Excellence in Teaching Award, HKN Teaching Award and the Dr. William D. and Nancy W. Strecker Chair Professorship. He was also the recipient of the Sigma-Xi Best Ph.D. Thesis Award in 2013 and the IEEE BRAIN Best Paper Award in 2019.