






End-to-end metasurface inverse design for single-shot multi-channel imaging

ZIN LIN,^{1,*} RAPHAËL PESTOURIE,¹  CHARLES ROQUES-CARMES,²  ZHAOYI LI,³ FEDERICO CAPASSO,³ MARIN SOLJAČIĆ,^{2,4} AND STEVEN G. JOHNSON¹ 

¹Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02138, USA

²Research Lab of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02138, USA

³John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA

⁴Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02138, USA

*zinlin@mit.edu

Abstract: We introduce end-to-end inverse design for multi-channel imaging, in which a nanophotonic frontend is optimized in conjunction with an image-processing backend to extract depth, spectral and polarization channels from a single monochrome image. Unlike diffractive optics, we show that subwavelength-scale “metasurface” designs can easily distinguish similar wavelength and polarization inputs. The proposed technique integrates a single-layer metasurface frontend with an efficient Tikhonov reconstruction backend, without any additional optics except a grayscale sensor. Our method yields multi-channel imaging by spontaneous demultiplexing: the metaoptics front-end separates different channels into distinct spatial domains whose locations on the sensor are optimally discovered by the inverse-design algorithm. We present large-area metasurface designs, compatible with standard lithography, for multi-spectral imaging, depth-spectral imaging, and “all-in-one” spectro-polarimetric-depth imaging with robust reconstruction performance ($\lesssim 10\%$ error with 1% detector noise). In contrast to neural networks, our framework is physically interpretable and does not require large training sets. It can be used to reconstruct arbitrary three-dimensional scenes with full multi-wavelength spectra and polarization textures.

© 2022 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

Metasurfaces—large-area subwavelength-patterned structures for free-space optics—have been heralded as a revolutionary platform for realizing complex functionalities and compact form factors inaccessible to conventional refractive or diffractive optics [1–4]. Meanwhile, an emerging “end-to-end” paradigm in computational imaging, in which an optical frontend is optimized in conjunction with a computational-imaging [5] backend, has received increasing attention due to successful applications in diffractive optics [6,7]. More recently, the end-to-end paradigm has been introduced to full-wave vectorial nanophotonic and metasurface frontends [8–10], demonstrating an enhanced capability for physical data acquisition and manipulation. So far, these early endeavors have been limited to two-dimensional (2D) RGB imaging or classification problems. In this paper, we present end-to-end metaoptics inverse design for single-shot *multi-channel* imaging beyond 2D RGB information: reconstruction of several depth, spectral and polarization channels *simultaneously* from a single monochrome image (Section 2.). As a key result, we show that, even though demultiplexing is not a designated/prescribed objective, inverse design automatically leads to *spatial demultiplexing* of the multiple channels into *spontaneous domains*—distinct regions in the detected image for different channels, whose locations are not pre-designated but are optimally discovered during the course of optimization (Sections 3 and 4). In contrast to data-driven approaches such as neural networks [7], our framework is physically interpretable, does not overfit despite a small generic training set, and is fully validated

against ground truths vastly different from those of the training set. Specifically, we present metasurface designs for 16-color imagers with 5–12% reconstruction error (under 1% image noise), a 4-color/4-depth imager with 5% error, and a 2-color/2-depth/4-polarization imager with 2% error (Section 3.). All the presented designs take into account fabrication constraints and are compatible with large-scale metasurface lithography [11]. In practice, our method only requires a single calibration step (via measurement or calculation of the point spread function) and is amenable to arbitrary material platforms and differentiable reconstruction algorithms. These results highlight the power of full-wave optics design with subwavelength components, whereas scalar diffractive optics could struggle to distinguish different wavelengths and polarizations due to limited dispersion and polarization sensitivity [4], and previous work could only do so by imposing strong spatio-spectral priors (e.g., limiting reconstruction to images that change slowly with wavelength) [7].

A major aspiration of metasurface technology has been to realize aberration-free focusing via an ultra-thin interface, directly replacing traditional bulky lenses [2]. While there has been significant progress towards this goal [3], most metalenses suffer from fundamental delay-bandwidth limits on wave focusing [12]. Although nanophotonic inverse design has introduced several innovations to metaoptics architectures [13–18], further disruptive improvements await the advent of mature three-dimensional (3D) nanofabrication [19–21]. In contrast, recent studies in end-to-end inverse design [8–10] have unveiled “computationally aware” nanostructures that bear little semblance to a lens and offer capabilities beyond optics-only or computation-only designs. On the other hand, several computational techniques have been developed for retrieving depth, spectral and polarization information from a scene [22–31]. Such techniques operate by combining multiple bulky refractive, diffractive and/or absorptive elements, often involve time-domain multiplexing (for example, scanning a scene to accumulate different shots), and, in most cases, enable the reconstruction of a single additional dimension (e.g. depth, color, or polarization). A universal framework is still lacking, by which a *single-piece* nanophotonic structure can be optimally designed to extract *any and all* channels *simultaneously* from a *single filter-free* monochrome exposure, exploiting the full Maxwell physics. Our proposed end-to-end framework enables inverse design of an ultra-thin single-layer metasurface in conjunction with a simple Tikhonov-regularized reconstruction algorithm. In particular, the Tikhonov regularization is agnostic to the nature of the information channels under consideration and is thus capable of extracting any and all channels (whether they be depth, spectral, polarization, or any combination thereof); alternative backends such as compressed sensing, which increase the channel capacity at the tradeoff of stronger priors, are discussed in Sec. 4.

2. Theory

2.1. Image formation model

In conventional imaging, the optical frontend is usually modeled by an elementary phase-shift function $e^{i2\pi h(x,y)/\lambda}$, where λ is the free-space wavelength and $h(x,y)$ the surface profile of the diffractive optical element [32]. In nanophotonics and metaoptics, involving sub-wavelength scatterers, more detailed electromagnetic simulations are required, which must take into account richer wave effects such as multiple scattering [1,15–17,33,34]. In this work, we use a Chebyshev-interpolated surrogate model (\mathbf{T}) under a locally periodic approximation (LPA) to efficiently simulate the transmitted electric field through a large-area metasurface [17]. Specifically, a metasurface is defined by a vector \mathbf{g} characterizing the geometry of meta-atoms (such as width, height and orientation of nanopillars) while the surrogate model maps each parameter g in a periodic unit cell to complex transmission coefficients. The transmitted electric field is then given by $\mathbf{E}_{\text{transmitted}} = \mathbf{T}(\mathbf{g}) \cdot \mathbf{E}_{\text{incident}}$.

In general, any ground-truth object \mathbf{u} can be numerically discretized into a tensor of five dimensions including three-dimensional real space as well as color and polarization dimensions.

For convenience, we denote \mathbf{u} as a set of 2D (x, y) intensity arrays: $\mathbf{u} \equiv \{u_{z,\lambda,p}\}$, where each 2D array u is indexed by depth (z), wavelength (λ) and polarization (p) channels (see Fig. 1(a,b)). Such a “multi-channel” representation is naturally made for a multi-channel image-formation model, in which a single 2D monochrome image v is formed by the sum of convolutions of the object channels with the corresponding point spread functions (PSFs) also indexed by (z, λ, p) :

$$\begin{aligned} v &= \sum_{z,\lambda,p} \text{PSF}_{z,\lambda,p} \circledast u_{z,\lambda,p} + \eta, \\ z &\in \{z_1, z_2, \dots, z_{n_z}\}, \\ \lambda &\in \{\lambda_1, \lambda_2, \dots, \lambda_{n_\lambda}\}, \\ p &\in \{p_x, p_y, p_{xy}^R, p_{xy}^I\}, \end{aligned} \quad (1)$$

where η is a generic noise term (typically modeled by zero-mean Gaussian white noise with standard deviation σ : $\eta \sim \mathcal{N}(0, \sigma^2)$ [6]). Note that in our model, we assume shift-invariant PSFs (valid in the paraxial regime) and only consider object intensities under incoherent illumination [32]. The PSFs are computed from the surrogate model followed by near-to-far-field propagation, given a specific metasurface geometry \mathbf{g} . While there is no limit to the number of depths (n_z) or wavelength channels (n_λ), four polarization channels are sufficient to reconstruct the full Stokes vector [35] ($n_p \leq 4$). Those components can be understood as follows: x -polarized intensity channel (p_x), y -polarized intensity channel (p_y), the real part of the correlation between x and y polarizations (p_{xy}^R) and the imaginary part (p_{xy}^I).

2.2. Inverse scattering and end-to-end optimization

Given the multi-channel image formation model, the corresponding reconstruction problem (also called inverse scattering problem) is posed as:

$$\min_{\{\mu_{z,\lambda,p}\}} \left\| v - \sum_{z,\lambda,p} \text{PSF}_{z,\lambda,p} \circledast \mu_{z,\lambda,p} \right\|^2 + R(\{\mu_{z,\lambda,p}\}). \quad (2)$$

The reconstructed object, denoted by $\hat{\mathbf{u}} = \{\hat{u}_{z,\lambda,p}\}$, is the solution that minimizes the problem (2). Here, a regularization term $R(\cdot)$ is usually needed to make the inverse problem well-posed and well-conditioned as well as to impose any prior information such as sparsity or smoothness. A simplest choice (with minimal prior information) is the so-called Tikhonov regularization or L_2 norm [36] where $R(\cdot) = \alpha \|\cdot\|^2$, leading to:

$$\hat{\mathbf{u}} = \left(\mathbf{G}^T \mathbf{G} + \alpha \mathbf{I} \right)^{-1} \mathbf{G}^T \mathbf{v}, \quad (3)$$

where, for convenience, the convolutions have been recast into a matrix notation,

$$\mathbf{G} = \begin{bmatrix} \text{PSF}_{z_1, \lambda_1, p_x} \circledast & \dots & \text{PSF}_{z, \lambda, p} \circledast & \dots \end{bmatrix} \quad (4)$$

Typically, the matrix \mathbf{G} is large and dense, easily reaching over $10^5 \times 10^5$ in dimension. We use matrix-free FFT-based convolutions [32,37] in both forward and inverse scattering models to efficiently compute the action of \mathbf{G} or \mathbf{G}^T on arbitrary vectors without storing \mathbf{G} explicitly. In particular, in Eq. (3), $\hat{\mathbf{u}}$ can be obtained by the iterative conjugate-gradient method [38], within ~ 100 iterations, instead of directly computing a matrix inverse. Our end-to-end inverse design considers the entire pipeline (see Fig. 1(c)) and can be formulated as minimizing the average

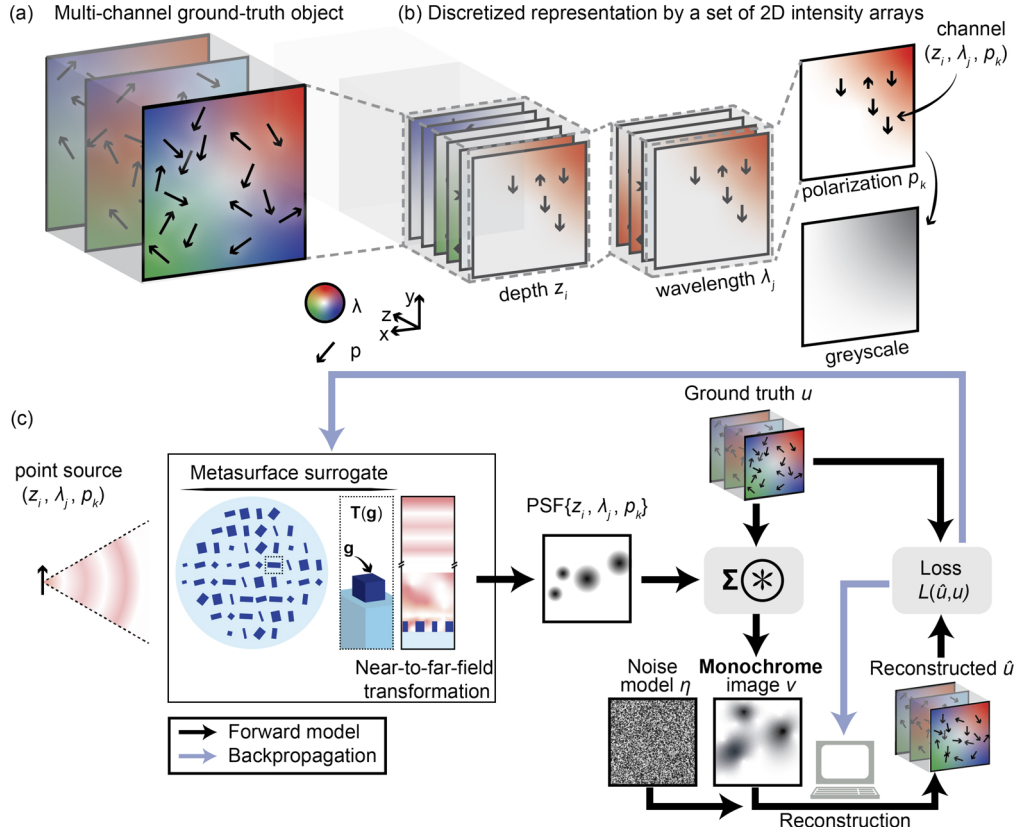


Fig. 1. (a,b) A multi-channel ground truth object consists of depth, spectral and polarization channels and can be represented by a set of two-dimensional (2D) intensity arrays indexed by (z, λ, p) : a three-dimensional (3D) object can be naturally sectioned into a collection of 2D “depth” slices; each depth slice can be further decomposed into different “color” slices; each color slice is, in turn, decomposed into different “polarization” slices. (c) End-to-end inverse design: a metasurface frontend is optimized in conjunction with a computational reconstruction backend to minimize the reconstruction error evaluated at the end of the full pipeline.

reconstruction error:

$$\begin{aligned}
 \min_{\mathbf{g}, \alpha} L(\hat{\mathbf{u}}, \mathbf{u}) &\triangleq \langle \|\mathbf{u} - \hat{\mathbf{u}}\|^2 \rangle_{\mathbf{u}, \eta} \\
 \hat{\mathbf{u}} &= (\mathbf{G}^T \mathbf{G} + \alpha \mathbf{I})^{-1} \mathbf{G}^T \mathbf{v} \\
 \mathbf{v} &= \sum_{z, \lambda, p} \text{PSF}_{z, \lambda, p} \otimes u_{z, \lambda, p} + \eta \\
 \text{PSF} &= |\text{FF}(\mathbf{T}(\mathbf{g}) \cdot \mathbf{E}_{\text{incident}})|^2.
 \end{aligned} \tag{5}$$

Here, $\langle \cdot \rangle_{\mathbf{u}, \eta}$ denotes averaging over training data as well as image noise; the training dataset consists of a few randomly-generated ground truths (e.g., random patterns drawn from a uniform distribution). FF denotes the near-to-far-field propagation to the detector plane—a convolution of the transmitted electric fields with the free-space Green’s function [32]. In our end-to-end framework, the gradients are back-propagated through the entire pipeline all the way to the metasurface parameters, and are efficiently handled by an in-house implementation of the adjoint

method [38] (see Appendix A)—an innovation unique to our framework as opposed to solely relying on popular automatic differentiation libraries [39] which perform poorly for differentiating through iterative algorithms such as the conjugate-gradient method. The reconstruction accuracy of an optimized design is validated over vastly different ground truths (distinct from training objects).

3. Results

We now show how our framework can be utilized to inverse-design metaoptics with multi-channel reconstruction capability (depth, spectral, and polarization). We denote the dimensions of a ground truth object as $n_{\text{ch}} \times m \times m$ —a set of n_{ch} arrays each with $m \times m$ pixels (note that $n_{\text{ch}} = n_z n_\lambda n_p$), while the monochrome image is a *single* 2D array of $n \times n$ pixels. In this work, we choose $n^2 \geq n_{\text{ch}} m^2$, that is, there are at least as many image pixels as the *total* size of the object—an over-determined inverse problem, suitable for Tikhonov regularization which harbors minimal assumptions about the nature of the object. This limit on the channel capacity in our designs can be overcome—that is, one can solve *under*-determined inverse problems ($n^2 < n_{\text{ch}} m^2$)—if one incorporates stronger priors on the object, such as sparsity [40], into the end-to-end framework (Sec. 4.).

First, we design a 16-color metasurface imager made up of 600 nm-tall TiO_2 pillars on silica—a design platform compatible with large-area lithographic fabrication as recently demonstrated in millimeter-scale achromatic metasurfaces [11]. We initialize our metasurface to a uniform array of pillars with the same width (the initial value of the width is chosen to be between 60 nm and 405 nm). A Chebyshev-interpolated surrogate model maps the width of each pillar inside a unit cell (465 nm period) to transmission coefficients at 16 different wavelengths across the visible spectrum (450 – 660 nm). The optimized metasurface (Fig. 2(a,b)) is used to image a superimposed picture of letters, each emitting a different wavelength (Fig. 2(d), left inset). Note that the letters in the picture cannot be distinguished by naked eyes or by an ordinary imaging lens. In contrast, the *single-shot monochrome* image shows *spatial demultiplexing* of the wavelength channels, separating out the letters (Fig. 2(c)). We emphasize that the example of the lettered ground truth is chosen for an easy visualization of the demultiplexing effect while we train the metasurface with completely random patterns (see Section 2.2). Interestingly, the imager does not solely rely on the demultiplexing effect; for example, there is channel replication, e.g. λ_2 (460 nm) channel, and a small degree of hybridization, e.g. between λ_1 (450 nm) and λ_2 (460 nm) channels. While the human eye is not equipped to recover all the information encoded in hybridization and redundancy, these apparent “imperfections” do not necessarily represent information loss. In particular, the apparent mixing between different channels does not preclude information recovery since the channels can be readily reconstructed by computation (as long as the corresponding PSFs are distinctly non-degenerate), leading to a reconstruction error of 12.5% under 1% Gaussian image noise (Fig. 2(e)). We note that a signal-to-noise ratio of ~ 100 (1% noise) can be readily achieved by modern electronic sensors [6]. Furthermore, the apparent residual fine-grain noise in the reconstruction of non-random objects can be easily removed by simple de-noising routines.

In this example, we considered a simple geometry (a square pillar) suited for photo-lithographic mass production; utilizing a more complex geometry, such as a holey pillar, allowing for more degrees of freedom to manipulate incident wavefronts, leads to even better performance (5% error with 1% noise, see Appendix B, Fig. 5). Our methods are also amenable to inverse-design techniques allowing for freeform geometries, involving domain decomposition methods with larger unit cells and full topology optimization [13,34]. We emphasize that our framework does not seek a “heavily-processed imitation” of the ground truth; it looks for a faithful reconstruction which is *stable* under moderate noise, and should be applicable for imaging *any* object, including random ones (see Appendix B). If desired, additional processing may be used, such as convolutional

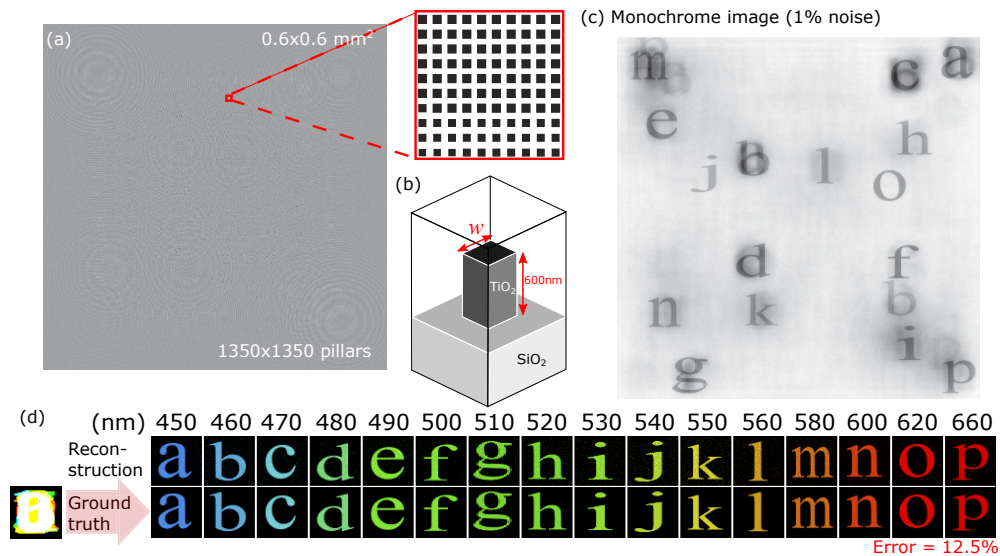


Fig. 2. (a) Design of a metasurface multi-spectral imager that can reconstruct 16 color channels from 450 nm (blue) to 660 nm (red). Inset: zoom-in of the design. (b) Each unit cell has a period of 465 nm, consisting of a square nanopillar. The pillar has a height of 600 nm and a width of $60 \text{ nm} \leq w \leq 405 \text{ nm}$. (c) Monochrome image of the synthetic object shown in (d). The wavelengths are spatially demultiplexed onto distinct domains on the single-shot monochrome image, captured 1 mm away from the metasurface by a CCD array of 400×400 pixels (each pixel has an area of $1.4 \times 1.4 \mu\text{m}^2$). (d) Reconstruction of a synthetic ground truth—a multi-spectral picture of letters ‘a’ to ‘p’, situated 2 cm away from the metasurface and each letter emitting a different wavelength. Note that the letters in the ground truth cannot be distinguished by the naked eye (inset on the left of the ground truth row). Computationally, the ground truth is represented by a set of 16 intensity arrays, each of which is a 50×50 -pixel image of a letter (with $25 \mu\text{m}$ resolution). The ground truth and the reconstruction are color-coded for a visual interpretation of the wavelengths. The reconstruction error is 12.5% under 1% image noise.

neural networks, which can be trained to “interpolate” a particular distribution of objects, enhance the reconstruction of *that* class of objects, perform image segmentation, or classification on the reconstructed objects.

Apart from spectral imagers, our framework is powerful in that it is straightforward to extract *any and all* kinds of channels. For example, we design a depth-spectral imager (Fig. 3) that can reconstruct 4 depth channels \times 4 wavelength channels. Additionally, as a proof of concept, we also design an “all-in-one” imager (Fig. 4) that can reconstruct 2 depth channels \times 2 wavelength channels \times 4 polarization channels. In that case, spontaneous spatial demultiplexing discovered via inverse design is observed for channels that are a combination of a given depth and polarization (i.e. channels sharing the same depth but having different polarizations are also demultiplexed, and vice-versa). On the other hand, a greater degree of hybridization is seen to arise in between the depth channels. This originates from the limited geometric control of the *local* metasurface design we have chosen, which cannot provide sufficiently strong *spatial* dispersion to fully separate the depth channels. In future works, we will engineer larger unit cells, higher diffraction orders, and cascaded metamaterials to induce strongly *non-local*, spatially-dispersive effects [15,19].

In our optimization problems, we found that end-to-end inverse design converges to an optimal design in about 300 to 500 iterations. An optimal design achieves accurate reconstruction in

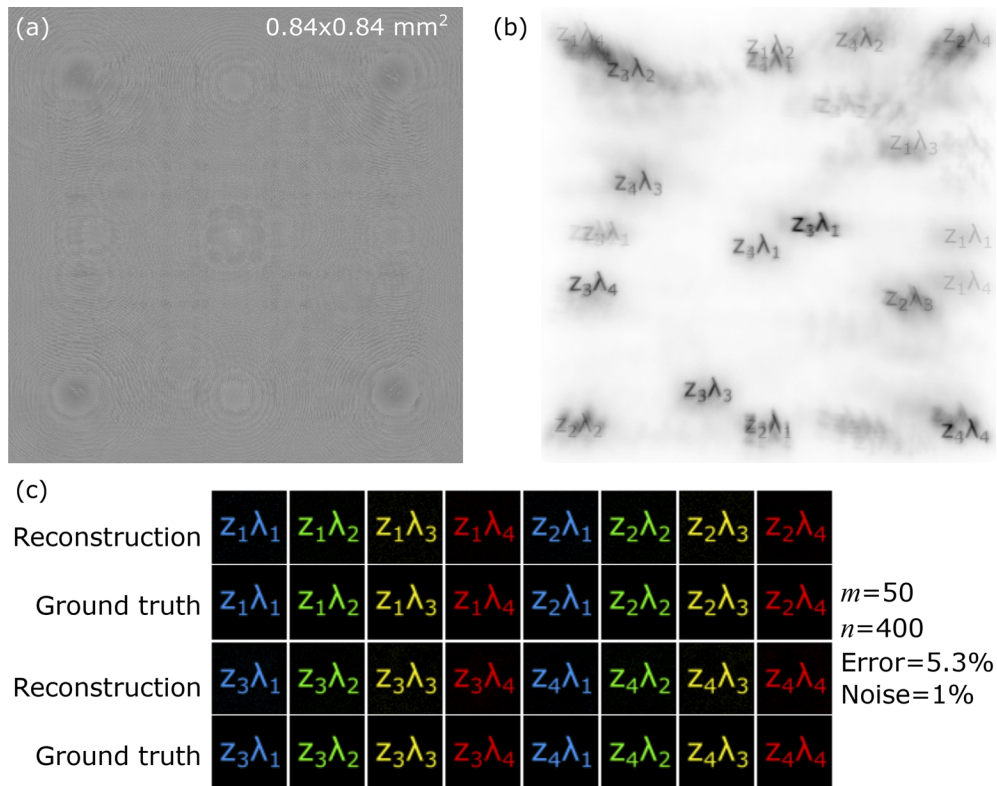


Fig. 3. Depth-spectral imager. (a) Metasurface depth-spectral imager design. (b) Monochrome image of a synthetic multi-dimensional object consisting of 4×4 color channels. The channels in the test object are artificially synthesized as $m \times m$ -pixel images of the channel indices ($m = 50$). The monochrome image has $n \times n$ pixels ($n = 400$) and is corrupted by 1% noise, leading to (c) a reconstruction error of 5.3%. Note that $z_i \in \{2, 4, 6, 8\}$ cm, $\lambda_j \in \{470, 520, 582, 660\}$ nm.

about 100 steps of the conjugate gradient (CG) descent. This can be even further accelerated by preconditioning approaches such as an incomplete LU factorization of the optimal PSF matrix. The bulk of computation in each step is made up of fast fourier transforms (FFT) which can be accelerated by optimized software and hardware.

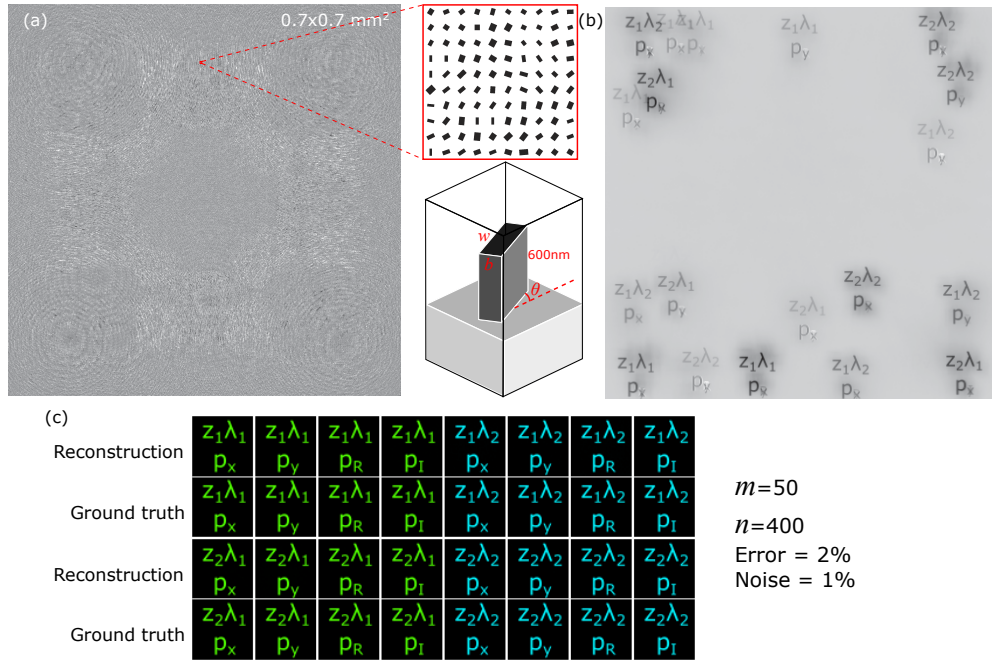


Fig. 4. Spectro-polarimetric-depth imager. (a) The metasurface consists of TiO₂ nanopillars, each characterized by width (w), breadth (b) and orientation angle (θ), where $60 \text{ nm} \leq w, b \leq 299 \text{ nm}$. (b) Monochrome image of a synthetic multi-dimensional object consisting of 2 depths \times 2 colors \times 4 polarization channels. The channels in the test object are artificially synthesized as pictures of the channel indices with $m \times m$ pixels ($m = 50$). The monochrome image has $n \times n$ pixels ($n = 400$) and is corrupted by 1% noise, leading to (c) a reconstruction error of 2%. Note that $z_i \in \{1.7, 3.4\} \text{ cm}$, $\lambda_j \in \{532, 488\} \text{ nm}$.

4. Discussion and outlook

The central result of this work is the realization of metaimaging based on *spontaneous* demultiplexing of multi-channel information into distinct spatial domains, whose locations appear irregular but are optimally determined by end-to-end inverse design. This is in contrast to the situation where such domains would be dictated by a user, as would be the case for conventional optics-only designs such as color splitters [12], or even hybrid systems [29]. The end-to-end automated discovery naturally leads to an optimal demultiplexing scheme with minimal hybridization between channels. In contrast, a human-designated scheme, such as a regular lattice of focal spots, may not be permitted by the available degrees of freedom in the *single-layer* metasurface, resulting in noisy crosstalk, sub-optimal PSFs, and poor resolution; relatedly, we note that traditional inverse design has to resort to three-dimensional *volumetric* structures to split three colors to three designated spots [21]. Intuitively, the demultiplexing effect is enabled by the Tikhonov reconstruction backend, which does not attempt to learn from the specific training data, but “judiciously” nudges the optical frontend to separate the incoming channels in order to reduce the reconstruction error. Therefore, our result is physically interpretable and data-efficient. It only requires a small training set, for example, as few as 30 training objects drawn from a uniform distribution in the case of the 16-color imager (Fig. 2). At the same time, the final optimized designs achieve robust reconstruction performance with consistent accuracy over vastly different sets of ground truths (whether they are pictures like letters or patterns like random dots, see Appendix B). This is in contrast to recent works [7] using phase masks and neural networks,

which required large training sets ($\gtrsim 10^4$ objects) and yielded similarly shaped PSFs at different wavelengths and depths (making reconstruction difficult without strong spatio-spectra priors, i.e. images resembling the training set). In general, the limited spectral dispersion in diffractive phase masks (or coded apertures) make it challenging for them to resolve PSFs at nearby wavelengths, whereas subwavelength-metasurface (nanophotonic) optics can exploit the full Maxwell physics to respond strongly to wavelength and polarization changes.

One limitation of our current multi-channel imagers is the channel capacity, and in particular that the transverse dimensions of the object must be smaller than those of the detector divided by the number of channels; therefore, the device is not suitable for reconstructing the entire natural field of view corresponding to the size of the detector. In practice, a narrower operational field of view may be realized by an appropriate aperture, a directed flash, or by selective illumination (a common technique in microscopy) [41,42]. The field of view can be enlarged by designing larger-area metasurfaces or by taking into account out-of-field-of-view light in the image-formation model. On the other hand, the inverse problem is under-determined if we choose the same transverse dimensions for the object and the detector. Such a problem requires additional priors on the object, and a regularization scheme like Tikhonov may not be sufficient. One powerful prior in image processing is sparsity, and a theoretically rigorous technique for reconstructing sparse objects is called compressed sensing [40]. In another manuscript under preparation, we will present a fully end-to-end inverse design framework with a compressed-sensing backend. Ultimately, future backends may be realized by new architectures that combine classical algorithms (such as Tikhonov and CS, which are theoretically rigorous and physically interpretable) and deep neural networks (which are best suited for learning deep data priors). Moreover, the performance of ultra-compact nanophotonic devices (such as depth and spectral sensitivities) can be further enhanced if we transcend the limitations imposed by LPA and expand the available degrees of freedom to encompass the full Maxwell physics. In future work, we hope to explore end-to-end inverse design using more sophisticated domain decomposition methods [34] and scattering-matrix formulations [43], or by cascading non-local metamaterials and 3D photonic crystals with local metasurfaces.

Appendix A. Adjoint gradient

The reconstruction $\hat{\mathbf{u}}$ under $\eta = 0$, $\alpha > 0$ is obtained by iteratively solving the equation:

$$\left(\mathbf{G}^T\mathbf{G} + \alpha\mathbf{I}\right)\hat{\mathbf{u}} = \mathbf{G}^T\mathbf{G}\mathbf{u}, \quad (6)$$

using the conjugate-gradient method without forming an explicit matrix. Note that the transpose of a convolution kernel is a convolution with the mirror image of the original kernel. Therefore, \mathbf{G}^T is simply convoluting with mirrored PSFs and then vertically stacking the results (in order to have output of the same size and shape as \mathbf{u} .)

Given a function $f(\hat{\mathbf{u}}(\mathbf{g}))$, we outline how to use the adjoint method [38] to find $\frac{\partial f}{\partial \mathbf{g}}$.

$$\frac{\partial f}{\partial \mathbf{g}} = \frac{\partial f}{\partial \hat{\mathbf{u}}} \cdot \frac{\partial \hat{\mathbf{u}}}{\partial \mathbf{g}} \quad (7)$$

$$= \Lambda \cdot \frac{\partial \mathbf{G}^T\mathbf{G}}{\partial \mathbf{g}} (\mathbf{u} - \hat{\mathbf{u}}) \quad (8)$$

where the adjoint variable Λ is given by

$$\left(\mathbf{G}^T\mathbf{G} + \alpha\mathbf{I}\right)\Lambda = \frac{\partial f}{\partial \hat{\mathbf{u}}} \quad (9)$$

We may find Λ using the same iterative solver that we used to find $\hat{\mathbf{u}}$.

The trickier issue is to find $\frac{\partial \mathbf{G}^T \mathbf{G}}{\partial \mathbf{g}}$. If one carries out the algebra faithfully, one may find that the inner product sandwiching the tricky derivative boils down to a cross-correlation between Λ and $\mathbf{G}(\mathbf{u} - \hat{\mathbf{u}})$. However, we can exploit the `autograd` automatic-differentiation (AD) package [39] in Python to compute this derivative effortlessly as follows (pseudo-code):

```
def innerdv(x, a, b):  
    def aGTGb(x):  
        ... # compute and return aGTGb given design parameters x  
        ... # by "autograd-able" convolutions  
    g = autograd.grad (aGTGb)  
    return g(x)
```

The desired product $\Lambda \cdot \frac{\partial \mathbf{G}^T \mathbf{G}}{\partial \mathbf{g}}(\mathbf{u} - \hat{\mathbf{u}})$ is then simply given by `innerdv (g, Lambda, u-uhat)`.

Appendix B. Metasurface with holey pillars

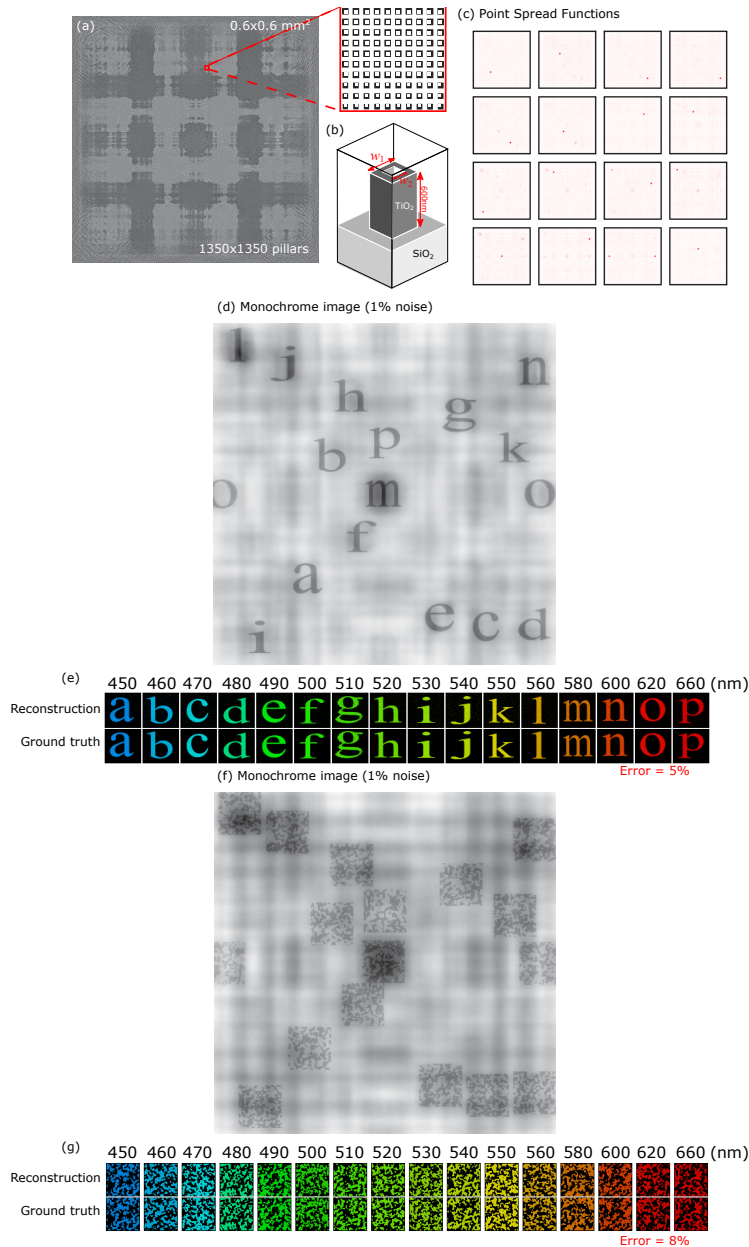


Fig. 5. (a-c) A 16-color imager with “holey pillars” and the PSFs. The metasurface has an average transmission efficiency of $> 55\%$. It can accurately reconstruct colored letters (d,e) as well as a random ground truth (f,g).

Funding. MIT-IBM Watson AI Laboratory (Challenge 2415); Multidisciplinary University Research Initiative (FA9550-21-1-0312); Army Research Office (W911NF-18-2-0048).

Disclosures. The authors declare no conflicts of interest.

Data availability. The code used in this paper is publicly available at [44]. The data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

References

1. N. Yu, P. Genevet, M. A. Kats, F. Aieta, J.-P. Tetienne, F. Capasso, and Z. Gaburro, "Light propagation with phase discontinuities: generalized laws of reflection and refraction," *Science* **334**(6054), 333–337 (2011).
2. M. Khorasaninejad, W. T. Chen, A. Y. Zhu, J. Oh, R. C. Devlin, C. Roques-Carmes, I. Mishra, and F. Capasso, "Visible wavelength planar metalenses based on titanium dioxide," *IEEE J. Sel. Top. Quantum Electron.* **23**(3), 43–58 (2017).
3. W. T. Chen, A. Y. Zhu, V. Sanjeev, M. Khorasaninejad, Z. Shi, E. Lee, and F. Capasso, "A broadband achromatic metalens for focusing and imaging in the visible," *Nat. Nanotechnol.* **13**(3), 220–226 (2018).
4. J. Engelberg and U. Levy, "The advantages of metalenses over diffractive lenses," *Nat. Commun.* **11**(1), 1991 (2020).
5. J. N. Mait, G. W. Euliss, and R. A. Athale, "Computational imaging," *Adv. Opt. Photonics* **10**(2), 409–483 (2018).
6. V. Sitzmann, S. Diamond, Y. Peng, X. Dun, S. Boyd, W. Heidrich, F. Heide, and G. Wetzstein, "End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging," *ACM Trans. Graph.* **37**(4), 1–13 (2018).
7. S.-H. Baek, H. Ikoma, D. S. Jeon, Y. Li, W. Heidrich, G. Wetzstein, and M. H. Kim, "End-to-end hyperspectral-depth imaging with learned diffractive optics," arXiv preprint arXiv:2009.00463 (2020).
8. Z. Lin, C. Roques-Carmes, R. Pestourie, M. Soljačić, A. Majumdar, and S. G. Johnson, "End-to-end nanophotonic inverse design for imaging and polarimetry," *Nanophotonics* **10**(3), 1177–1187 (2021).
9. C. M. V. Burgos, T. Yang, Y. Zhu, and A. N. Vamivakas, "Design framework for metasurface optics-based convolutional neural networks," *Appl. Opt.* **60**(15), 4356–4365 (2021).
10. E. Tseng, S. Colburn, J. Whitehead, L. Huang, S.-H. Baek, A. Majumdar, and F. Heide, "Neural nano-optics for high-quality thin lens imaging," arXiv preprint arXiv:2102.11579 (2021).
11. Z. Li, R. Pestourie, J.-S. Park, Y.-W. Huang, S. G. Johnson, and F. Capasso, "Inverse design enables large-scale high-performance meta-optics reshaping virtual reality," arXiv preprint arXiv:2104.09702 (2021).
12. F. Presutti and F. Monticone, "Focusing on bandwidth: achromatic metalens limits," *Optica* **7**(6), 624–631 (2020).
13. S. Molesky, Z. Lin, A. Y. Piggott, W. Jin, J. Vucković, and A. W. Rodriguez, "Inverse design in nanophotonics," *Nat. Photonics* **12**(11), 659–670 (2018).
14. D. Sell, J. Yang, S. Doshay, R. Yang, and J. A. Fan, "Large-angle, multifunctional metagratings based on freeform multimode geometries," *Nano Lett.* **17**(6), 3752–3757 (2017).
15. Z. Lin, B. Groever, F. Capasso, A. W. Rodriguez, and M. Lončar, "Topology-optimized multilayered metaoptics," *Phys. Rev. Appl.* **9**(4), 044030 (2018).
16. Z. Lin, V. Liu, R. Pestourie, and S. G. Johnson, "Topology optimization of freeform large-area metasurfaces," *Opt. Express* **27**(11), 15765–15775 (2019).
17. R. Pestourie, C. Pérez-Arancibia, Z. Lin, W. Shin, F. Capasso, and S. G. Johnson, "Inverse design of large-area metasurfaces," *Opt. Express* **26**(26), 33732–33747 (2018).
18. Z. Shi, A. Y. Zhu, Z. Li, Y.-W. Huang, W. T. Chen, C.-W. Qiu, and F. Capasso, "Continuous angle-tunable birefringence with freeform metasurfaces for arbitrary polarization conversion," *Sci. Adv.* **6**(23), eaba3367 (2020).
19. Z. Lin, C. Roques-Carmes, R. E. Christiansen, M. Soljačić, and S. G. Johnson, "Computational inverse design for ultra-compact single-piece metalenses free of chromatic and angular aberration," *Appl. Phys. Lett.* **118**(4), 041104 (2021).
20. C. Roques-Carmes, Z. Lin, R. E. Christiansen, Y. Salamin, S. E. Kooi, J. D. Joannopoulos, S. G. Johnson, and M. Soljačić, "Towards 3d-printed inverse-designed metaoptics," arXiv preprint arXiv:2105.11326 (2021).
21. P. Camayd-Muñoz, C. Ballew, G. Roberts, and A. Faraon, "Multifunctional volumetric meta-optics for color and polarization image sensors," *Optica* **7**(4), 280–283 (2020).
22. A. P. Pentland, "A new sense for depth of field," *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-9**(4), 523–531 (1987).
23. Q. Guo, Z. Shi, Y.-W. Huang, E. Alexander, C.-W. Qiu, F. Capasso, and T. Zickler, "Compact single-shot metalens depth sensors inspired by eyes of jumping spiders," *Proc. Natl. Acad. Sci.* **116**(46), 22959–22965 (2019).
24. A. Levin, R. Fergus, F. Durand, and W. T. Freeman, "Image and depth from a conventional camera with a coded aperture," *ACM Trans. Graph.* **26**(3), 70–es (2007).
25. A. Greengard, Y. Y. Schechner, and R. Piestun, "Depth from diffracted rotation," *Opt. Lett.* **31**(2), 181–183 (2006).
26. K. Monakhova, K. Yanny, N. Aggarwal, and L. Waller, "Spectral diffusercam: Lensless snapshot hyperspectral imaging with a spectral filter array," *Optica* **7**(10), 1298–1307 (2020).
27. S. K. Sahoo, D. Tang, and C. Dang, "Single-shot multispectral imaging with a monochromatic camera," *Optica* **4**(10), 1209–1213 (2017).
28. Z. Yang, T. Albrow-Owen, H. Cui, J. Alexander-Webber, F. Gu, X. Wang, T.-C. Wu, M. Zhuge, C. Williams, P. Wang, A. V. Zayats, W. C. Dai, S. Hofmann, M. Overend, L. Tong, Q. Yang, Z. Sun, and T. Hasan, "Single-nanowire spectrometers," *Science* **365**(6457), 1017–1020 (2019).
29. N. A. Rubin, G. D'Aversa, P. Chevalier, Z. Shi, W. T. Chen, and F. Capasso, "Matrix Fourier optics enables a compact full-stokes polarization camera," *Science* **365**(6448), 1 (2019).
30. A. Wagadarikar, R. John, R. Willett, and D. Brady, "Single disperser design for coded aperture snapshot spectral imaging," *Appl. Opt.* **47**(10), B44–B51 (2008).

31. W. Feng, H. Rueda, C. Fu, G. R. Arce, W. He, and Q. Chen, "3d compressive spectral integral imaging," *Opt. Express* **24**(22), 24859–24871 (2016).
32. J. W. Goodman, *Introduction to Fourier Optics* (Roberts and Company, 2005).
33. M. Khorasaninejad, W. T. Chen, R. C. Devlin, J. Oh, A. Y. Zhu, and F. Capasso, "Metalenses at visible wavelengths: Diffraction-limited focusing and subwavelength resolution imaging," *Science* **352**(6290), 1190–1194 (2016).
34. Z. Lin and S. G. Johnson, "Overlapping domains for topology optimization of large-area metasurfaces," *Opt. Express* **27**(22), 32445–32453 (2019).
35. J. N. Damask, *Polarization optics in telecommunications*, vol. 101 (Springer Science & Business Media, 2004).
36. A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation* (SIAM, 2005).
37. M. Frigo and S. G. Johnson, "The design and implementation of FFTW3," *Proc. IEEE* **93**(2), 216–231 (2005).
38. G. Strang, *Computational Science and Engineering*, vol. 791 (Wellesley-Cambridge, 2007).
39. D. D. Maclaurin and M. Johnson, "Autograd: Efficiently computes derivatives of numpy code," (2015).
40. D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory* **52**(4), 1289–1306 (2006).
41. B. E. Saleh and M. C. Teich, *Fundamentals of photonics* (John Wiley & Sons, 2019).
42. J. Mertz, *Introduction to optical microscopy* (Cambridge University, 2019).
43. M. Benzaouia, J. D. Joannopoulos, S. G. Johnson, and A. Karalis, "Quasi-normal mode theory enforcing fundamental constraints for truncated expansions," arXiv preprint arXiv:2105.01749 (2021).
44. Z. Lin, R. Pestourie, C. Roques-Carmes, Z. Li, F. Capasso, M. Soljačić, and S. G. Johnson, "Code for "End-to-end metasurface inverse design for single-shot multi-channel imaging"," Github (2022), https://github.com/zlin-opt/end2end_multich_tikhonov.git.