Role of Spatial Coherence in Single-Shot Lensless Image Reconstruction

Ganesh M. Balasubramaniam[1], Xiao-Liu Chu[2] and Matthew R. Foreman[1,2]*
[1] *School of Electrical and Electronic Engineering, Nanyang Technological University, 50 Nanyang Avenue, Singapore 639798*
[2] *Institute for Digital Molecular Analytics and Science, 59 Nanyang Drive, Singapore 636921*
(Dated: July 20, 2026)

Lensless imaging is a technique that recovers object information computationally from diffraction patterns recorded without imaging optics, making its performance strongly dependent on the forward model used during image reconstruction. A common source of reconstruction error is the assumption of fully coherent propagation, even when the illumination exhibits partial spatial coherence in practice. Here, the role of spatial coherence is examined for assorted object classes using a coherence-aware forward model based on generalised van Cittert-Zernike Schell propagation. Simulated measurements are generated over a controlled range of effective source spatial coherence lengths and reconstructed using either a partially coherent forward model or a conventional coherent propagation model. Our results show that decreasing spatial coherence progressively degrades reconstructions under the coherent assumption, whereas incorporating partial coherence can preserve object structure and improve image quality, particularly for dense high-spatial-frequency features. Reconstructions from experimental measurements further confirm the practical need to use coherence-aware inversion under partially spatially coherent illumination. These findings establish spatial coherence as a defining component of the inverse problem in lensless imaging and provide a route to optimization of partially coherent lensless imaging systems.

Lensless imaging provides an alternative to conventional microscopy by computationally recovering object information from diffraction patterns recorded directly on an image sensor. By eliminating imaging optics, lensless systems combine a compact form factor with a large space-bandwidth product, making them attractive for wide-field microscopy, high-throughput screening, and portable imaging platforms [1–3]. In such systems, the detector does not record an image of the specimen in the conventional sense. Instead, a diffraction-encoded intensity pattern is measured, from which the object must be estimated. In model-based reconstruction algorithms, a forward model maps a candidate object to a detector intensity under the assumed illumination and propagation conditions. Consequently, reconstruction quality is determined not only by the numerical formulation of the inverse problem, but also by the physical accuracy of the forward model.

A key component of a forward model is the spatial coherence of the illumination. Coherent propagation models are widely used because they yield a simple wave-optical description of diffraction and integrate readily into inverse reconstruction algorithms. Many practical lensless imaging systems, however, operate with extended sources, such as light-emitting diodes, for which the object plane illumination is only partially spatially coherent [4, 5]. Partial coherence is not a minor perturbation to an optical system. It is known to influence resolution, contrast transfer, and reconstruction performance across tomographic microscopy, high-numerical-aperture focusing, and quantitative phase retrieval [4–6]. Fundamentally, the transverse extent of the source determines the mutual coherence of the incident field, which in turn modifies interference visibility, diffraction contrast, and the spatial structure of the intensity at the detector, thus affecting lensless image formation [7, 8]. Since single-shot reconstruction relies entirely on the assumed forward model, these coherence-dependent changes in the diffraction pattern are incorrectly attributed to the object, erasing fine features or introducing artefacts. Whether coherent propagation suffices for a given object, or whether coherence-aware inversion is required, therefore remains an open question whose resolution necessitates isolating coherence-induced error from that associated with sampling, regularisation, and numerical optimisation.

In this work, coherence-induced model mismatch in lensless imaging is examined directly using a forward model based on the generalised van Cittert-Zernike Schell propagator, which provides an efficient description of the field diffracted by an object under partially spatially coherent illumination [9]. Lensless measurements are simulated across a controlled range of source radii for both amplitude and phase objects, and reconstructions obtained with a coherence-aware inverse model are compared with those found using a conventional coherent propagator. Our analysis establishes how decreasing spatial coherence alters the sensor measurement, quantifies the resulting degradation when coherence is neglected during inversion, and demonstrates that explicitly incorporating the illumination coherence can restore reconstruction quality across distinct object classes. We further support our findings through an experimental demonstration.

We consider a lensless forward imaging model defined for a thin transmissive object with complex transmit-

* matthew.foreman@ntu.edu.sg

tance

$$t(\mathbf{r}_o) = a(\mathbf{r}_o)\exp[i\phi(\mathbf{r}_o)]\,,\ 0 \le a(\mathbf{r}_o) \le 1,\ 0 \le \phi(\mathbf{r}_o) \le 2\pi, \tag{1}$$

where $\mathbf{r}_o = (x_o, y_o)$ is the transverse coordinate in the object plane, $a(\mathbf{r}_o)$ is the amplitude transmittance, and $\phi(\mathbf{r}_o)$ is the object-induced phase delay. The source plane, object plane, and detector plane are separated by distances z_{so} and z_{od}, as shown in Fig. 1. When the object is illuminated by a coherent spherical wave, the detector-plane intensity is given by

$$I_{\mathrm{coh}}(\mathbf{r}_d; z_{od}) = \left| \mathcal{P}_{z_{od}} \left\{ t(\mathbf{r}_o) \exp\left[\frac{ik}{2z_{so}} |\mathbf{r}_o|^2 \right] \right\} (\mathbf{r}_d) \right|^2, \tag{2}$$

where $\mathbf{r}_d = (x_d, y_d)$ denotes the transverse detector coordinate, $k = 2\pi/\lambda$ is the wavenumber, and $\mathcal{P}_{z_{od}}\{\cdot\}$ represents free-space propagation from the object plane to the detector plane, which we implement using the angular spectrum method.

Partial spatial coherence of the illumination source can additionally be incorporated using the generalised van Cittert-Zernike Schell formalism [9]. The normalised coherence function is a function of the separation between two points in the object plane, $\Delta\mathbf{r}_o = (\Delta x_o, \Delta y_o)$ and is given by

$$\mu_F(\boldsymbol{\xi}_o) = \frac{\mathcal{F}_{\mathbf{r}_s \to \boldsymbol{\xi}_o}\{S(\mathbf{r}_s)\}}{\left|\mathcal{F}_{\mathbf{r}_s \to \boldsymbol{\xi}_o}\{S(\mathbf{r}_s)\}\right|_{\max}}, \qquad \boldsymbol{\xi}_o \triangleq \frac{\Delta\mathbf{r}_o}{\lambda z_{so}}, \tag{3}$$

where $S(\mathbf{r}_s)$ is the intensity distribution at the pinhole source (see Fig. 1) and $\mathcal{F}_{\mathbf{r}_s \to \boldsymbol{\xi}_o}\{\cdot\}$ denotes the two-dimensional Fourier transform from the source-plane coordinate $\mathbf{r}_s = (x_s, y_s)$ to the spatial-frequency coordinate $\boldsymbol{\xi}_o$. In this case, the detector-plane intensity (neglecting an unimportant constant scaling factor) is

$$I_{\mathrm{pc}}(\mathbf{r}_d; z_{od}) = \mathcal{F}_{\Delta\mathbf{r}_o \to \mathbf{r}_d}\left[\mathcal{F}^{-1}_{\mathbf{r}_d \to \Delta\mathbf{r}_o} \left[I_{\mathrm{coh}}(\mathbf{r}_d; z_{od}) \right] \mu_F(\boldsymbol{\xi}_o) \right]. \tag{4}$$

The source distribution is taken here to be uniform, i.e., $S(\mathbf{r}_s) = \mathrm{circ}(|\mathbf{r}_s|/r_s)$, and the pinhole radius r_s is varied to control the spatial coherence of the illumination. In the fully coherent point-source limit, $\mu_F(\boldsymbol{\xi}_o) = 1$, such that (4) reduces to (2).

To examine coherence-dependent image formation, we simulated ensembles of amplitude-only objects ($\phi(\mathbf{r}_o) = 0$) which vary in two structural properties, namely the spatial frequency content of $a(\mathbf{r}_o)$ and the real-space feature density (see Supplementary Material). Specifically, these span high-frequency high-density (HFHD), high-frequency low-density (HFLD), low-frequency high-density (LFHD), and low-frequency low-density (LFLD) object classes. Representative examples from each class are shown in Fig. 1 (first column). The second column depicts the corresponding normalised log Fourier spectrum,

$$\widetilde{S}_O(f_x, f_y) = \mathcal{N}\left[\log\left\{1 + |\mathcal{F}\{a(\mathbf{r}_o)\}|\right\}\right], \tag{5}$$

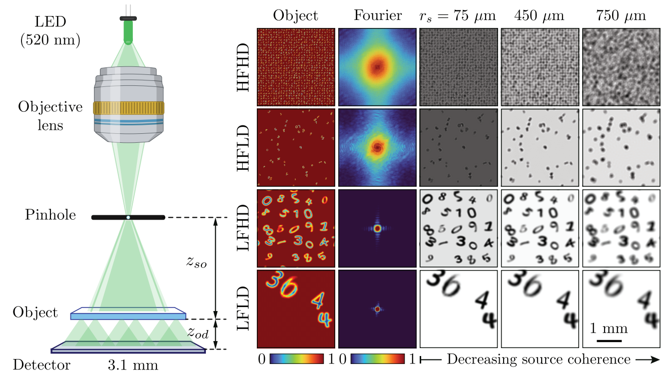

FIG. 1. **Coherence dependent lensless image formation.** (left) Lensless imaging geometry, where source spatial coherence is controlled by the pinhole diameter. (right) Columns, from left to right, show the object, corresponding normalised log Fourier spectrum $\widetilde{S}_O(f_x, f_y)$, and simulated detector intensities for $r_s = 75\ \mu$m, 450 μm, and 750 μm respectively for different object classes (top to bottom): high-frequency high-density (HFHD), high-frequency low-density (HFLD), low-frequency high-density (LFHD), and low-frequency low-density (LFLD).

where f_x and f_y are the horizontal and vertical spatial frequencies, and $\mathcal{N}[\cdot]$ (or equivalently $\widetilde{[\cdot]}$) linearly maps the input function to $[0, 1]$. For a fixed object and propagation geometry, increasing the source radius modifies the detector-plane intensity through the coherence factor in (3). The resulting recorded lensless measurement is seen to depend directly on the source's spatial coherence, as shown in the final three columns of Fig. 1. Corresponding results for phase-objects (whereby $a(\mathbf{r}_o) = 1$, and $a(\mathbf{r}_o) \to \phi(\mathbf{r}_o)$ in (5)) are also presented in the Supplementary Material.

The effect of coherence-induced model mismatch during inversion was subsequently assessed by reconstructing each object from the simulated detector-plane intensity assuming either a coherence-aware operator, $\mathcal{M}_{\mathrm{pc}}$, or a fully coherent operator $\mathcal{M}_{\mathrm{coh}}$, which evaluate (4) and (2) respectively. For a measured detector intensity I_{meas}, the reconstructed object associated with either operator is obtained from

$$\widehat{t}_\nu = \arg\min_{t \in \mathcal{C}} \mathcal{J}_\nu(t), \qquad \nu \in \{\mathrm{pc}, \mathrm{coh}\}, \tag{6}$$

where $\mathcal{C}$ denotes the object constraint set. The reconstruction objective is defined as

$$\begin{aligned} \mathcal{J}_\nu(t) = {} & w_{\mathrm{SSIM}} \left[1 - \mathrm{SSIM}\left(\widetilde{\mathcal{M}_\nu\{t\}}, \widetilde{I}_{\mathrm{meas}} \right) \right] \\ & + w_{\mathrm{MSE}} \left\| \widetilde{\mathcal{M}_\nu\{t\}} - \widetilde{I}_{\mathrm{meas}} \right\|_2^2 + \lambda_{\mathrm{TV}} \mathcal{R}_{\mathrm{TV}}(t) + \lambda_{\mathrm{bg}} \mathcal{R}_{\mathrm{bg}}(t), \end{aligned} \tag{7}$$

where $\mathcal{R}_{\mathrm{TV}}(t) = \sum_{\mathbf{r}} [|\partial_x t(\mathbf{r})|^2 + |\partial_y t(\mathbf{r})|^2 + \epsilon_{\mathrm{TV}}^2]^{1/2}$ is an isotropic total variation penalty that suppresses spurious local oscillations while preserving sharp object transitions, and $\mathcal{R}_{\mathrm{bg}}(t) = \|B_t(\mathbf{r}) \odot [a(\mathbf{r}) - (\mathcal{G}_\sigma * a)(\mathbf{r})]\|_2^2$ suppresses high-spatial-frequency background fluctuations

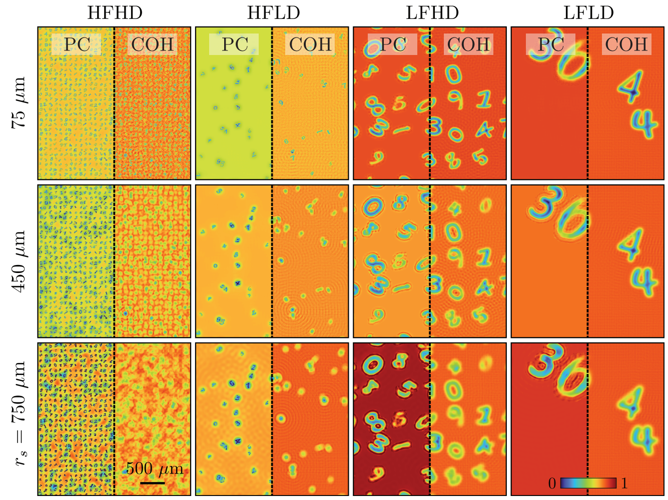

FIG. 2. **Representative reconstructions across object classes and source coherence.** Object reconstructions for HFHD, HFLD, LFHD, and LFLD object classes (left to right) and source radii (top to bottom). In each image, the left half shows the reconstruction obtained using the partially coherent (PC) forward model, while the right half shows the reconstruction obtained using the coherent (COH) forward model.

without strongly penalising object boundaries. The map $B_t(\mathbf{r})$ is a self-estimated background weight computed from the current object estimate. Pixels with weak object contrast and low local gradient are assigned high weights, whereas pixels containing object structure or sharp boundaries are assigned low weights. The Gaussian kernel $\mathcal{G}_\sigma$ defines the local smooth background estimate, SSIM is the structural similarity, $*$ denotes convolution, and $\odot$ denotes pointwise multiplication. The parameter $\epsilon_{\mathrm{TV}} = 10^{-8}$ ensures differentiability of the total variation term, while $\lambda_{\mathrm{TV}} = 10^{-4}$ and $\lambda_{\mathrm{bg}} = 10^{-5}$ set the strengths of the total variation and background regularisers respectively. $w_{\mathrm{SSIM}} = 0.3$ and $w_{\mathrm{MSE}} = 0.7$ dictate the relative weights assigned to the structural-similarity and mean-squared-error terms. The same objective, optimisation schedule, and initialisation are used for both $\mathcal{M}_{\mathrm{pc}}$ and $\mathcal{M}_{\mathrm{coh}}$, ensuring that any difference in reconstruction quality arises solely from the forward model used.

Reconstruction was performed from a single detector intensity recorded at $z_{od} = 2$ mm, under the constraint $0 \leq a(\mathbf{r}_o) \leq 1$. For all simulated measurements, we assumed a wavelength of 520 nm, $z_{so} = 3$ cm, whilst r_s was varied from 75 μm to 750 μm. This range of pinhole radii spans source spatial coherences from high to low (corresponding to small to large source radii), following the scaling used in the van Cittert-Zernike Schell formalism [9]. Twenty-five distinct objects were simulated for each individual object class and source radius (1000 amplitude objects in total). Representative amplitude reconstructions are shown in Fig. 2. Reconstructions of phase-objects, obtained using two detector planes, are further reported in the Supplementary Material. The coherence-aware model preserves the dominant object structure across the considered source radii. In contrast, reconstructions obtained under the fully coherent assumption become progressively less accurate as the source radius increases. The degradation is especially evident for dense objects, where the coherent model produces loss of fine structure, reduced feature contrast, and reconstruction artefacts that are absent or substantially suppressed with the partially coherent model. The consequences of coherence mismatch are thus seen to emerge directly in the recovered amplitude structure.

To quantitatively compare reconstruction quality we considered a number of distinct metrics Q_ν. First, the global reconstruction fidelity was quantified using the Pearson correlation coefficient (PCC) [10]. Next, to assess whether reconstruction preserved the spectral distribution of the reference object, the high-frequency retention error (HFRE) and low-frequency retention error (LFRE) were computed as the logarithmic mismatch between the reconstructed and reference spectral energy fractions above and below a fixed normalised radial-frequency cutoff of 0.25 cycles/pixel. Finally, boundary recovery was quantified using the edge F_1 score, E_{F1}, found from boundary precision and recall with a small spatial tolerance [11, 12]. For each object class and source radius, reconstruction advantage was computed as the difference between quality metrics for the partially coherent and coherent forward models. Specifically, for metrics for which larger values indicate better reconstruction (E_{F1} and PCC), we set $\Delta Q = Q_{\mathrm{pc}} - Q_{\mathrm{coh}}$, whereas for HFRE and LFRE, which are error measures, we set $\Delta Q = Q_{\mathrm{coh}} - Q_{\mathrm{pc}}$. Positive values therefore always indicate that the partially coherent model provides improved image reconstruction. The resulting average reconstruction improvement scores (and corresponding standard deviation) for each object class and source radius are depicted in Fig. 3.

The results of Fig. 3 demonstrate a strong dependence on the spatial dimensionality and frequency content of the object class. HFHD objects exhibit the largest and most consistent positive improvement from use of $\mathcal{M}_{\mathrm{pc}}$. Their dense, high-frequency structure generates many closely spaced boundaries and a broad spatial-frequency spectrum, making the recorded diffraction pattern highly sensitive to partial coherence. As the source radius increases, $\mathcal{M}_{\mathrm{coh}}$ increasingly fails to represent this partial coherence-induced redistribution of spatial frequency content. The positive ΔPCC shows improved full-image reconstruction performance, while the positive ΔHFRE and ΔLFRE show that the partially coherent model better preserves the spectral distribution of the reference object. Positive ΔE_{F1} further indicates more accurate boundary localisation. HFLD objects show a weaker, but still clear, improvement when using $\mathcal{M}_{\mathrm{pc}}$. These objects contain fine features, but their lower spatial density reduces the number of independent high-frequency structures in the field. Consequently, ΔPCC, ΔHFRE, and ΔE_{F1} remain positive, demonstrating improved reconstruction.

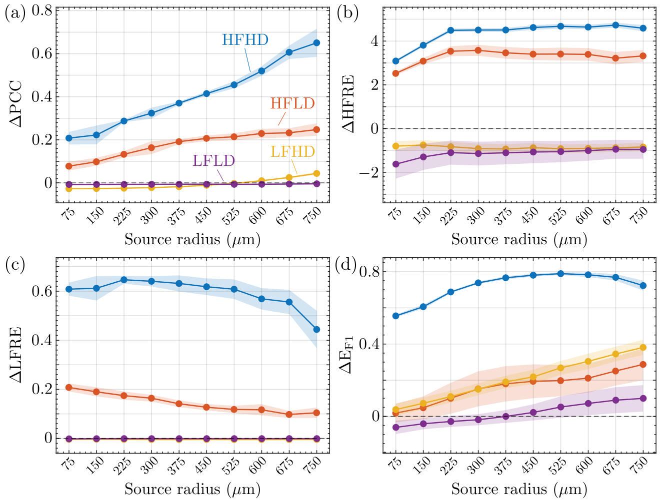

FIG. 3. **Class-resolved reconstruction advantage under varying source coherence.** Differences ΔQ are plotted as a function of source radius for the four object classes. Panels show differences in (a) PCC, (b) HFRE, (c) LFRE, and (d) E_{F1}. Solid markers denote the class-averaged value (across 25 distinct objects) at each source radius, and shaded bands denote standard deviation across objects within that class.

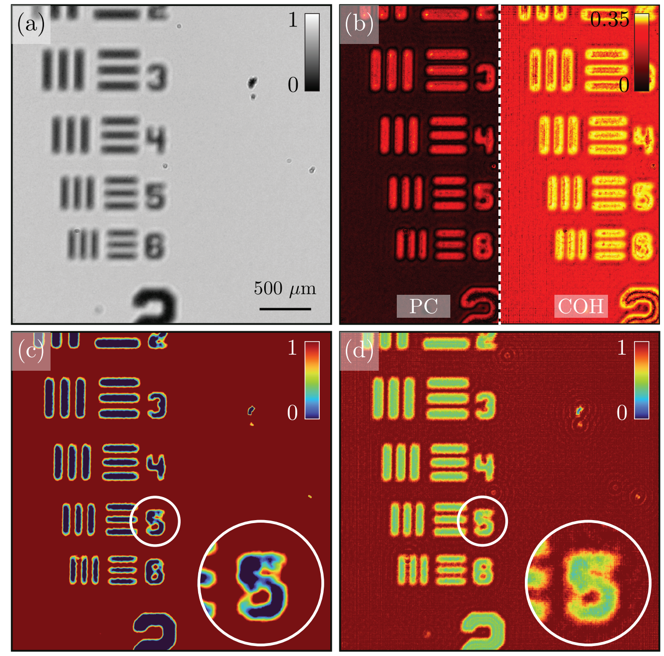

FIG. 4. **Experimental verification using a USAF resolution target.** (a) Measured lensless detector intensity. (b) Absolute error map between the measured and predicted detector intensities, for partially (left) and fully (right) coherent models, respectively. Reconstructed target obtained using a (c) partially coherent and (d) fully coherent forward model. Insets showing magnified views of the indicated feature.

LFHD amplitude objects show limited ΔPCC, because their dominant structure lies at low spatial frequency and is therefore less affected by partial coherence. $\mathcal{M}_{\rm coh}$ can capture much of the broad object envelope, so the difference between the partially coherent and coherent reconstructions remains small in global intensity and spectrum quality. Notably, the positive ΔE_{F1} however shows that the $\mathcal{M}_{\rm pc}$ still improves boundary localisation. The negative ΔHFRE arises since LFHD objects contain relatively little energy above the radial-frequency cutoff. In such a regime, small amounts of additional reconstructed edge sharpness or residual texture can increase the high-frequency spectral mismatch, even when $\mathcal{M}_{\rm pc}$ improves boundary localisation. LFLD amplitude objects show the weakest improvement among all objects, as they and their corresponding reconstructions are dominated by broad slowly varying features. Full-image metrics are thus less sensitive to local boundary differences. The small or negative spatial and spectral differences, together with the weaker E_{F1} improvement, indicate that partial-coherence modelling produces only limited boundary-level gains for this class.

Phase-object reconstructions show a related but more nuanced dependence on spatial coherence, as detailed in the Supplementary Material. Coherence-aware inversion improves boundary localisation across all phase-object classes, as indicated by consistently positive ΔE_{F1}. Global and spectral gains are however more strongly class dependent. The largest spectral improvements occur for high-frequency phase objects, whereas low-frequency phase objects can show weak or negative HFRE and LFRE changes despite improved boundary localisation.

Experimental verification was performed using a lensless measurement of a USAF resolution target under partially coherent illumination (Fig. 4(a)). The target was illuminated by a commercial green light-emitting diode coupled through a 500 μm pinhole (Thorlabs P500HMB). The source-to-object and object-to-detector separations were z_{so} = 110 mm and z_{od} = 2 mm, respectively. The diffraction pattern was recorded using a Raspberry Pi Camera Module 3 equipped with a Sony-megapixel array with a 7.4 mm sensor diagonal and a nominal pixel pitch of 1.55 μm $\times$ 1.55 μm. A central 2000 $\times$ 2000-pixel region of the recorded measurement was used for computational reconstruction. Figure 4 shows the reconstruction results obtained using both $\mathcal{M}_{\rm pc}$ and $\mathcal{M}_{\rm coh}$ operators. In the absence of a ground-truth object transmittance, the reconstruction quality was verified using the agreement between the measured and predicted detector intensities (Fig. 4(b)), which was quantified using PCC and Fourier correlation (FC). The partially coherent model yielded PCC = 0.9871 and FC = 0.9906, compared with PCC = 0.9212 and FC = 0.9396 for the coherent model. The experimental results provide validation of the forward-model analysis and show that incorporating the source coherence improves agreement with the observed lensless intensity pattern.

In conclusion, our results establish spatial coherence as a consequential component of the lensless imaging inverse problem. Incorporating partial spatial coherence into the forward model improves reconstruction quality and yields closer agreement with measured detector intensities. The improvement is greatest when the target contains fine, densely distributed features, because such objects place a greater demand on accurate spatial-frequency transfer and boundary localisation. Our findings indicate that coherence should be treated as an explicit design and modelling parameter in lensless imaging, rather than as an incidental property of the illumination. Recent information-theoretic approaches to object-dependent lensless system design provide a natural framework for extending this analysis [13]. Future work could use such criteria to optimise the source coherence for specified objects and imaging tasks, enabling object-adaptive lensless system design.

ACKNOWLEDGMENTS

Ministry of Education, Singapore (RG137/24, EDUN C-33-18-279-V12); Agency for Science, Technology and Research, Singapore (H25-MRG3483).

[1] W. Bishara, T.-W. Su, A. F. Coskun, and A. Ozcan, Optics Express **18**, 11181 (2010).
[2] A. Greenbaum, W. Luo, T.-W. Su, *et al.*, Science Translational Medicine **6**, 267ra175 (2014).
[3] G. M. Balasubramaniam, X.-L. Chu, R. V. Nair, and M. R. Foreman, arXiv preprint arXiv:2604.22270 (2026).
[4] S. O. Isikman, W. Bishara, and A. Ozcan, Applied Optics **50**, H253 (2011).
[5] Y. Chen, X. Wu, L. Lu, *et al.*, Optics Letters **47**, 6061 (2022).
[6] M. R. Foreman and P. Török, Journal of the Optical Society of America A **26**, 2470 (2009).
[7] Y. Bian, W. Wang, A. Hussain, *et al.*, Optics Communications **434**, 136 (2019).
[8] Y. Cai and F. Wang, Optics Letters **32**, 205 (2007).
[9] M. Montoya, M. J. Lopera Acosta, Y. Nie, and D. Blinder, Optics Express **33**, 46313 (2025).
[10] K. Pearson, Proceedings of the Royal Society of London **58**, 240 (1895).
[11] J. Canny, IEEE Transactions on Pattern Analysis and Machine Intelligence **PAMI-8**, 679 (1986).
[12] D. R. Martin, C. C. Fowlkes, and J. Malik, IEEE Transactions on Pattern Analysis and Machine Intelligence **26**, 530 (2004).
[13] L. A. Kabuli, H. Pinkard, E. Markley, *et al.*, Optica **13**, 227 (2026).