

Light Microscopy at Maximal Precision

Matthew Bierbaum, Brian D. Leahy, Alexander A. Alemi,^{*} Itai Cohen, and James P. Sethna

Department of Physics, Cornell University, Ithaca, New York 14853, USA

(Received 8 April 2017; revised manuscript received 18 July 2017; published 13 October 2017)

Microscopy is the workhorse of the physical and life sciences, producing crisp images of everything from atoms to cells well beyond the capabilities of the human eye. However, the analysis of these images is frequently little more accurate than manual marking. Here, we revolutionize the analysis of microscopy images, extracting all the useful information theoretically contained in a complex microscope image. Using a generic, methodological approach, we extract the information by fitting experimental images with a detailed optical model of the microscope, a method we call parameter extraction from reconstructing images (PERI). As a proof of principle, we demonstrate this approach with a confocal image of colloidal spheres, improving measurements of particle positions and radii by 10–100 times over current methods and attaining the maximum possible accuracy. With this unprecedented accuracy, we measure nanometer-scale colloidal interactions in dense suspensions solely with light microscopy, a previously impossible feat. Our approach is generic and applicable to imaging methods from brightfield to electron microscopy, where we expect accuracies of 1 nm and 0.1 pm, respectively.

DOI: [10.1103/PhysRevX.7.041007](https://doi.org/10.1103/PhysRevX.7.041007)

Subject Areas: Optics, Soft Matter

I. INTRODUCTION

Microscope technology has progressed to near perfection. Crisp images speak of precisely engineered microscope components: large-aperture and nearly aberration-free lenses, high-frame-rate and low-noise cameras, and powerful and uniform light sources. Nanometer-scale details boast of super-resolution techniques thought impossible mere decades ago: photoactivated localization microscopy (PALM) [1], stochastic optical reconstruction microscopy (STORM) [2], and stimulated emission depletion (STED) [3]. The continued development of ever more powerful techniques—structured illumination [4], lattice-light sheet [5], and holographic microscopy [6,7]—reassures us that resolution will continue to improve.

However, our ability to extract *quantitative* information from microscopy images has not kept pace. In fields from electron microscopy to super-resolution localization, current methods mimic human perception with heuristic approaches, such as looking for the centers of bright spots or regions of contrast in an image [8–13]. The simplicity of these methods necessarily ignores physical complexities in the image formation. As a result, systematic errors and inefficient estimates plague these techniques [14,15].

In this paper, we present a universal method of scientific image analysis that extracts *all* the useful information

theoretically contained in a complex image. Our method, dubbed parameter extraction from reconstructing images (PERI), uses a detailed model of the physics of image formation to fit experimental images. From the fit, we then extract information about the image at the information-theoretic limit. Importantly, since PERI relies on simple comparisons between computationally reconstructed images and those generated by physical experiments, it can be broadly applied to improve precision measurements in images generated by a wide range of techniques, from electron microscopy to SQUID scanning-probe microscopy [16,17]. Our method does not require modifying the microscope or the image acquisition. As a result, any researcher with a microscope can readily apply our technique to push their data to the information-theoretic limit. Here, we illustrate this approach on confocal images of colloidal spheres, measuring each particle's position and radius to within 3 nm, a 10–100-times improvement over current methods, with especially large improvements for images of dense suspensions (see Supplemental Material (SM) [18], Sec. VI.). We use this extreme accuracy to measure colloidal interactions at the nanometer scale, measuring deviations from hard-sphere interactions in dense suspensions for the first time with light microscopy.

How precisely can an object be located in an image? The fundamental limitation in locating an object arises from statistical noise in the image formation, not directly from diffraction or optical limitations [19]. This limit is determined through the interplay of the image signal and noise, as described by the Cramér-Rao bound. Specifically, the Cramér-Rao bound states that the covariance matrix of the estimated parameters is always larger than the inverse of the Fisher information matrix of the noise distribution [20].

^{*}Work done while at Cornell; currently affiliated with Google Inc.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

For an image with Gaussian white noise of variance σ^2 , sampled at points \mathbf{x}_k , the minimum uncertainty in the parameters $\boldsymbol{\theta}$ measured from the image is

$$\text{cov}\theta_{ij} \geq \sigma^2 \left(\sum_k \frac{\partial \mathcal{I}(\mathbf{x}_k)}{\partial \theta_i} \frac{\partial \mathcal{I}(\mathbf{x}_k)}{\partial \theta_j} \right)^{-1}, \quad (1)$$

where $\mathcal{I}(\mathbf{x})$ is the image that would be measured in the absence of noise.

We can use this equation to estimate the minimum uncertainty in measuring a colloidal sphere's radius and position from a three-dimensional confocal image. For a particle of radius R blurred by diffraction over a width w , the derivatives with respect to particle radius in Eq. (1) are only nonzero on a shell at the particle's edge of approximately $4\pi R^2 w$ voxels. At the particle's edge, the intensity changes from a characteristic brightness of approximately I to 0 over a width of approximately w , and the derivatives are thus approximately of magnitude I/w . Substituting these values gives a minimum uncertainty in a particle's radius as $\sigma_R \sim \sqrt{w/4\pi R^2}/\text{SNR}$, where $\text{SNR} = I/\sigma$ is the signal-to-noise ratio. Likewise, changing the particle's position only affects the edge voxels in the direction of the particle's motion. The positional derivatives will thus be approximately of magnitude I/w only on a projected shell of approximately $\pi R^2 w$ voxels, giving the minimal uncertainty in the particle's position as $\sigma_x \sim \sqrt{w/\pi R^2}/\text{SNR}$. For a colloidal particle of diameter $1 \mu\text{m}$, imaged with a confocal microscope with voxel size of 100 nm and diffractive blur of $w \approx 200 \text{ nm}$ at a $\text{SNR} = 25$, these uncertainties correspond to $\sigma_R \approx 1.5 \text{ nm}$ and $\sigma_x \approx 3 \text{ nm}$, a fantastically high precision [21].

II. RESULTS

Actually achieving this localization without serious systematic errors requires a detailed knowledge of the image formation process. To incorporate this knowledge, we create a generative model of the microscope image based on the physics of the light interacting with the sample and with the microscope's optical train. We then fit every parameter in the model by comparing the image produced by the model to the experimental image. Our model describes the physics of image formation in the order that it occurs: (1) Fluorescent dye is distributed unevenly throughout the sample, (2) the dyed sample is illuminated unevenly by the laser, (3) the resultant image is blurred because of diffraction, and (4) the final image is noisy.

Dye distribution.—To reconstruct the image, we start with the continuous distribution of the fluorescent dye in the sample. For the image in Fig. 1, the dye is distributed everywhere except in a slab, representing the glass coverslip slide, and in a collection of spherical lacunae, representing the colloidal particles. To represent this continuous dye distribution on a pixelated grid, we draw these objects in real space using a function that is tuned to match the exact Fourier representation of a sphere (see SM [18] for an extensive discussion of this and the rest of the generative model). We call this correctly aliased representation on a pixelated grid the platonic image. While we focus on featuring only spheres in this work, PERI is flexible enough to include any parametrizable object in the generative model, such as ellipsoidal [22,23], rodlike [24], or polyhedral [25] particles.

Illumination field and background.—This distribution of dye is illuminated by a scanned laser. Because of imperfections and dirt in the optics, the illumination is not uniform but instead varies in space. For instance, our




FIG. 1. *PERI overview.* A demonstration of model information recovered from real confocal microscope images of $\langle a \rangle = 1.343(8) \mu\text{m}$ colloidal spheres at a volume fraction of $\phi = 0.130(5)$. (a) The generative model consists of a platonic image of dye distributed around perfect spheres and a coverslip (top panels), illuminated with a spatially varying intensity (middle panels), and convolved with a physical point-spread function (bottom panels). The left panels show the model components, and the right panels show the combined image. (b) These components combine to form a realistic generative model (bottom right), which, aside from noise, is visually indistinguishable from the data (top left). (c) From the fit parameters of the model, we extract information such as particle positions and radii (orange highlights) to within a few percent of a pixel—corresponding to an accuracy of 1 and 3–4 nm, respectively. The inset zooms into the center of the particle (white square), highlighting the precise localization and calculated uncertainty of the particle's position.

line-scanning confocal illumination field is highly striped, as any imperfections in the line illumination are dragged across the field of view. We describe this spatially varying illumination as a continuous field that varies throughout the image. Empirically, we find that combining a Barnes interpolant along the scan direction and Legendre polynomials in the perpendicular directions accurately describes both the rapidly varying stripes and the slowly varying changes in the illumination of our line-scan confocal. Additionally, the microscope always registers a nonzero background signal, which we include in our model. We parametrize this background similarly to the illumination field.

Point-spread function.—Diffraction prevents the illuminated dye from being imaged exactly onto the detector. Instead, each dye molecule in the sample projects a comparatively large blur, known as the point-spread function (PSF), onto the imaging camera. As a result, the image captured on the camera is a convolution of the illuminated platonic image with the PSF, and not simply the illuminated dye itself. While complicated, this PSF has been calculated exactly by many researchers for different geometries [26–33]. For microscope samples with a refractive index different from what the optical train is designed for, the PSF worsens with depth, becoming significantly broader and more aberrated. We use an adaptation of these exact PSF calculations for a line-scanning confocal as our PSF model, optimizing over parameters such as the numerical aperture of the lens and the index mismatch of the sample to the optics.

Putting these components together as shown in Fig. 1(a), our model image \mathcal{M} sampled at pixels \mathbf{x} is described by

$$\mathcal{M}(\mathbf{x}) = B(\mathbf{x}) + \int d^3\mathbf{x}' [I(\mathbf{x}') (1 - (1 - c)\Pi(\mathbf{x}'))] \times P(\mathbf{x} - \mathbf{x}'; \mathbf{x}), \quad (2)$$

where I is the illumination field, B is the background, Π is the platonic image, and P is the spatially varying PSF; we include a constant offset c to partially capture rapidly varying variations in the background, as discussed in SM [18]. The model image is highly realistic, as shown by the comparison with real data in Fig. 1(b).

Noise.—Finally, noise degrades the image recorded on the camera. We treat the noise using a Bayesian framework and look for the maximum-likelihood model, given the microscope data, complete with possible priors on parameter values. Since the noise is empirically Gaussian (see SM [18]), the most likely model is the least-squares fit of the model to the microscope image.

To find the most likely model, we least-squares fit every parameter in our generative model to find the correct particle positions, radii, illumination field, and point-spread function, as illustrated in Fig. 1(c). A typical confocal image contains a few times 10^3 particles, each with four fit parameters (x, y, z, R). In addition, there are a few hundred

global parameters to optimize, such as the illumination and PSF parameters and the z -step size of the lens along the optical axis, resulting in approximately 10^4 parameters per image—a daunting optimization problem. We begin with an initial guess for the positions using standard particle-locating techniques [34], and we simultaneously fit the particle positions and the global variables using a Levenberg-Marquardt algorithm modified for large parameter spaces [35–38]. From here, we ensure that we have correctly identified every particle in the image by automatically adding and subtracting particles based on the difference between the model and the microscope image. After finding the best-fit parameters, we can sample from the log-likelihood using standard Monte Carlo techniques [39] to estimate the errors in the image reconstruction. In practice, we frequently estimate the parameter uncertainties using the fit covariance, which provides results similar to full Monte Carlo sampling. (See SM [18] for a detailed description of the fitting method and numerical optimizations.) It is important to note that this fit is over all the pixels in the image—to get a meaningful extraction of parameters, every pixel must be described accurately. Imperfectly fit regions—due to, e.g., severely deformed particles or PSF leakage from objects outside the image—can bias the extracted positions of particles in the region and even affect the entire image reconstruction through the influence on image-scale variables. Overall, this fit procedure takes between 1 and 24 hours for a large confocal image, depending on prior knowledge of the microscope’s global parameters (see SM [18], Timing and Performance).

Using PERI to measure positions with nanometer accuracy requires rigorous checks on our method, with both generated and experimental data. We first generate images with a detailed physical model, employing an exact, spatially varying, point-spread function [26]; experimentally measured, spatially varying illumination; dense collections of particles with varying radii; and a realistic amount of noise. PERI successfully fits these generated data, converging to the global fit minimum in the extremely large-dimensional parameter space, despite a host of possible numerical complications, such as local minima in the fit space or a failure of the fit to converge. From this fit, PERI extracts both the particle positions and radii at the Cramér-Rao bound (approximately 2 nm and 1 nm, respectively). In contrast, current heuristic-based algorithms cannot measure the particle positions to better than 60 nm on realistically generated data sets. (See SM [18] for a detailed comparison of PERI to other particle location algorithms.)

Emboldened by this success, we next test PERI on real experimental data. We record movies of a suspension of 1.34- μm -diameter silica spheres suspended in a glycerol and water mixture, and we measure particle positions and radii in each three-dimensional frame of the movie with PERI. We use the extracted positions and radii to test for any systematic errors in PERI.

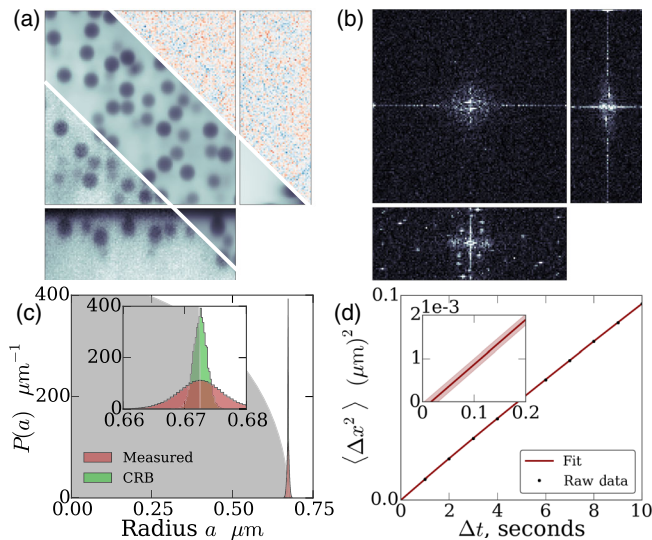


FIG. 2. *Fitting the generative model to experimental data.* (a) A representative image (lower left), its best-fit model (center), and the difference between the two (upper right), shown as cross sections in the xy , xz , and yz planes. The residuals show nearly perfect Gaussian white noise at the expected signal-to-noise ratio, demonstrating the quality of the generative model. (b) The Fourier power spectrum of the same residuals, displayed as three orthogonal slices in the $q_x q_y$, $q_x q_z$, and $q_y q_z$ planes. In addition to scanning noise, visible as the stripes along $q_x = 0$ and $q_z = 0$ as well as the isolated poles, excess power is visible at scales larger than the particles themselves but smaller than the features given by the illumination field (ILM). These residuals are associated with the incomplete description of the point-spread function. (c) PERI measures the particle radius within an uncertainty of 3–4 nm, as estimated from changes in featured particle radii with time (red histogram) measured by tracking approximately 1200 particles over 800 time points. Improving the description of the PSF would allow for radii to be featured at the CRB (green histogram), with a precision of 1 nm. (d) The experimental mean-squared displacement along the x direction $\langle \Delta x^2 \rangle$ (black dots; error bars are smaller than symbol size) from approximately 1200 particles measured over 800 time points provides an estimate of PERI’s average positional errors. Extrapolating the fitted mean-square displacement (red curve; the shaded band denotes the fit uncertainty) to $t = 0$ gives a positional error indistinguishable from 0 nm. The mean-square displacement along y and z provide the same constraints.

We first analyze the residuals of our fits to the experimental data. Figures 2(a) and 2(b) show these residuals in both real and Fourier space. If our fit to the experimental image were perfect, the residuals would be perfectly Gaussian white noise. Instead, while the overall probability distribution of the residuals is nearly Gaussian in both domains (see SM [18]), in Fourier space there are distinct wave vectors above the noise floor. Comprising roughly 10^{-5} of the power in the experimental image, the extremely small size of this remaining signal demonstrates the quality of our generative model. The deviations of our model from the experimental data occur at length scales slightly larger

than the particle diameter but smaller than typical illumination variations. These unexplained residuals most likely arise from approximations in models of line-scanning point-spread functions, excess aberrations in the microscope, and the artificially finite but large size we use in our PSF calculation to speed up optimization. Additionally, sharp peaks at high wave vectors can be seen in one slice of the Fourier-space residuals, which arise from noise in the scanning of the lens and the line illumination. If these residuals were in fact perfectly uncorrelated Gaussian noise, this would imply that we had correctly captured all meaningful information about the system aside from the shot noise in the detector. The remaining question is how much these residuals affect the parameters of interest, the particle positions and radii.

We can use the extracted particle positions and radii over time to test the accuracy of PERI. During the movies, the particles diffuse about, sampling different regions of the spatially varying illumination and point-spread function and changing the configuration of neighboring particles. However, the true particle radii remain constant in time. Measuring individual radii fluctuations over time provides a stringent model-independent measurement of errors in PERI, as the changing configuration of the particles includes all the possible sources of systematic error. Tracking these radii fluctuations over time suggests that we can measure the particle radius to within 3–4 nm [Fig. 2(c)], a fantastically high precision compared to the 672-nm particle radius and even the 125-nm pixel size. Importantly, PERI accurately measures each particle’s radius directly from the image. This individual-particle measurement allows PERI to directly measure the polydispersity in the particle radii; the sample in Figs. 2 and 3 has a very small polydispersity of $\Delta a/\bar{a} = 2\%$. A better understanding of the image formation in the microscope could increase the precision in measured radii even further, to the 1.5-nm minimal error from the Cramér-Rao bound. We can also constrain the positional errors. Since the particle positions undergo Brownian motion, their mean-square displacement grows linearly in time $\langle \Delta x^2(t) \rangle = 2Dt$ [40]. Any error in the particle positions that is uncorrelated with the particle position will manifest itself as a nonzero intercept when the fitted mean-square displacement is extrapolated to $t = 0$. By measuring the mean-square displacement along each coordinate direction and separately extrapolating to zero [panel (d)], we find that PERI’s positional errors are indistinguishable from zero and are less than 10 nm, with this constraint being limited only by statistics. Additionally, we check PERI on a data set of 2- μm diameter particles fixed in place via strong interactions—a less demanding test since immobilizing the particles also fixes most of the sources of systematic error. In these data, we find x and y errors of 1–2 nm, z errors of 3 nm, and radii errors of 0.8 nm (see SM [18]). Combined, these measurements demonstrate that we are able to measure particle positions and radii to within 3 nm.

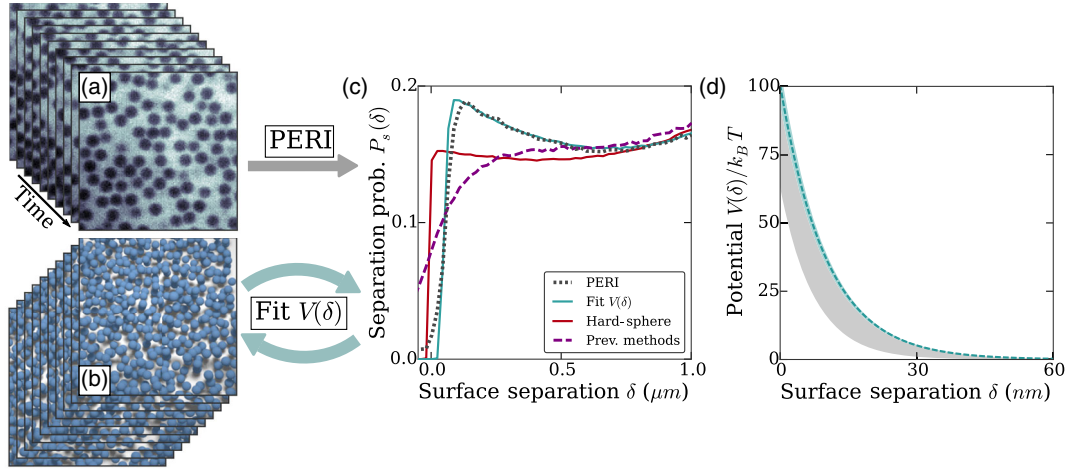


FIG. 3. *Extracting interparticle potentials.* (a) We use PERI to analyze a large ensemble of three-dimensional images of a dilute suspension of approximately 1200 Brownian particles; a small section of these images is shown in the upper left. From these data, we extract an experimental $P_s(\delta)$. (b) We use molecular dynamics to create a simulated $P_s(\delta)$, and we iteratively update the interaction potential $V(\delta)$ to find the $P_s(\delta)$ that best fits the experimental data. (c) The extracted $P_s(\delta)$ from PERI (gray dashed line) and from the fitted potential (solid cyan line) agree excellently. In contrast, the $P_s(\delta)$ from a strictly hard-sphere potential (red line) does not fit the data. The difference between these potentials depends on resolving particle separations at the nanometer level. Previous centroid-based methods [41] (purple line) produce a $P_s(\delta)$ with nonsensical features, such as significant overlaps, that cannot be fit by a reasonable interaction potential. We discuss the overlaps in PERI’s extracted $P_s(\delta)$ in SM [18], Sec. VII. (d) From the best-fit simulation, we extract the interparticle potential $V(\delta)$. The shaded bands show the uncertainty in the potential, with the teal band describing uncertainty in the fit and the gray band the uncertainty due to systematic errors (see SM [18] for further discussion).

Why is PERI able to measure particle positions and radii so accurately while heuristic methods fail? Heuristic methods produce poor measurements with large systematic errors simply because they ignore complexities of the image formation, such as the spatially varying illumination and point-spread function. In contrast, PERI includes these complexities. Fitting the entire image ensures that all the complexities are accounted for—any portion of the image formation not included in the model will manifest itself as strong residuals in the fit, declaring that the model is incomplete and suggesting what additional effect must be included. This process of model selection is described in detail in SM [18].

This extraordinary accuracy in measuring particle positions from microscopy images creates a new window into nanometer-range particle interactions in dense suspensions. When colloidal particles are suspended in an aqueous solution, the particles charge, as the polar solvent dissociates ions on the particles’ surface groups. This charge results in an electrostatic repulsion, which is in turn screened by counterions in the bulk [42,43]. The screening creates an interparticle potential that deviates from a hard-sphere potential only at nanometer separations. This potential ever so slightly biases the distribution of particle positions away from that expected for a hard-sphere suspension.

Previous efforts measured these interactions only in idealized, isolated surfaces such as between two surfaces [44] or a single colloidal particle interacting with a wall [45,46]. However, by their nature, these idealized

measurements frequently cannot include possible complications present in a real suspension, such as many-body interactions, realistic surface asperities, or increases in dissolved ion concentration from dissociated surface groups on multiple particles. Measuring the interaction potential in a dense colloidal suspension includes these and many other possible complications in the interaction.

We measure these nanometer-scale interactions by using PERI to analyze a large set of images of 1.3- μm silica spheres suspended in a water-glycerol mixture. To prevent kinetic effects from confounding our measurements, we allow the sample to fully sediment for an hour. This produces an open layer of sediment approximately 2–3 particle layers deep, shown in Fig. 3(a). We then image this suspension repeatedly over the course of several hours, extracting simulation-level detail of approximately 720 000 particle positions and radii over all the images. The particle interactions determine the structure of the suspension. We quantify this structure with the probability $P_s(\delta)$ of finding a pair of particles with surface-to-surface separation δ , accounting for radii polydispersity and sedimentation in a manner preferable to the usual pair-correlation function. To reconstruct the interparticle potential, we use the extracted particle radii and particle number from the data, and we simulate the particle dynamics using Brownian dynamics. We incorporate both gravitational settling and the interparticle potential, which we model as an exponentially decaying electrostatic repulsion. We then fit the mass density and potential by simulating, reconstructing $P_s(\delta)$

from the simulation at each set of potential parameter values, and iterating to find the best $P_s(\delta)$ that matches experiment [Fig. 3(b)].

The $P_s(\delta)$ from the best-fit simulation and from the experimental data analyzed by PERI agree excellently, at both large and small separations. At small separations, $P_s(\delta)$ rises rapidly over the first approximately $0.1 \mu\text{m}$ near contact in both the simulation data and the data extracted by PERI, as shown in Fig. 3(c). At longer distances, the probability grows because of the increased volume where particles can be located, with slight oscillations reflecting second- and third-nearest-neighbor interactions. In contrast, previous centroid-based methods [47] produce a $P_s(\delta)$ with nonsensical features, such as significant overlaps, that cannot be fit by a simulation.

We use the extracted $P_s(\delta)$ to measure nanometer-scale interactions in dense colloidal suspensions for the first time. The $P_s(\delta)$ measured by PERI is well fit by an exponentially decaying repulsive potential, as expected from electrostatic repulsion in standard colloidal theory [42] [Fig. 3(c)]. From the fit, we measure the potential's screening length as $10.1 \pm 2.5 \text{ nm}$ and the repulsion strength near contact as $100 \pm 30 \text{ kT}$, corresponding to surface potentials and screening lengths similar to those previously measured from the interaction of a single particle with a wall [45]. Our data strongly exclude hard-sphere interactions as the interparticle potential. Importantly, this resolving power between potentials results from the values of $P_s(\delta)$ near contact. Without the accurate localization provided by PERI, it is impossible to measure the potential at these separations.

III. DISCUSSION

Our technique and the ideas within it provide more than just a description of colloidal interactions. Nanometer accuracy in locating colloidal particle positions would revolutionize fields as diverse as the study of colloidal glasses and the measurement of biological forces with force-traction microscopy. With our open-source code [48], other researchers can immediately analyze existing images of these systems.

The principle of accurately reconstructing an image to extract parameters applies to a wide range of fields. While we only demonstrate our method on a single imaging system—3D line-scanning confocal images of nearly monodisperse spheres in a dyed fluid—the ideas behind our technique are generic and should apply to a wide range of imaging modalities. As a first step, extending PERI to describe image formation in brightfield microscopy by correctly describing aberrations and light scattering from the sample [49] would provide nanometer-scale precision for a simple and widespread imaging setup. Creating an accurate generative model of TEM or STM, by describing the focusing and scattering of the electron beam [50] or understanding the STM probe tip's response [51], could provide unprecedented detail of atomic positions. By

combining the logic from Eq. (1) with estimates of the error and pixel size from brightfield and TEM images, we estimate that PERI could provide localization accuracies of 1 nm for brightfield and 0.1 pm for TEM. In each case, the general approach of comparing a computationally reconstructed image to one that is collected using physical tools will enable extraction of detail at the information theoretic limit. Importantly, these improvements come by exploiting cheap computing rather than by expensive redesigns and modifications of the imaging system. Applying the idea of PERI to these imaging modalities will usher in a new era of precision measurements, for objects whose sizes range from microns to angstroms.

The data, scripts, and output files for the experiments in this manuscript are available upon request from Brian D. Leahy, bdl48@cornell.edu.

ACKNOWLEDGMENTS

We would like to acknowledge N. Lin, W. Zipfel, M. Transtrum, C. Clement, D. Koch, C. Schneider, L. Bartell, S. Whitehead, T. Beatus, and other members of the Cohen Lab for useful discussions. This work was supported in part by National Science Foundation (NSF) Grant No. DMR-1507607 (M. B., A. A., J. S., and I. C.), NSF Grant No. DMR-1120296 (B. L.), and ACS Grant No. PRF 56046-ND7 (B. L.). This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by NSF Grant No. ACI-1053575.

M. B. and B. D. L. contributed equally to this work.

APPENDIX A: MATERIALS AND METHODS

We use a Zeiss LSM 5 Live inverted confocal microscope, in conjunction with an infinity-corrected 100-times immersion oil lens (Zeiss Plan-Apochromat, 1.4 NA, immersion oil with index $n = 1.518$). The LSM 5 Live confocals operate by line scanning. Rather than rastering a single point at a time to form the image, a line-scanning confocal images an entire line at once. An image of a line is focused onto the sample, and the sample fluorescence is detected on a line CCD. Rastering this line allows images to be collected extremely rapidly; the data in the text were taken at 108 in-plane frames per second. However, the different line-scanning optics worsen the point-spread function compared to a point-scanning confocal, and they cause illumination imperfections such as dirt to be smeared out over one direction in the image. Importantly, our confocal is outfitted with a hyperfine piezo scanner, which gives precise z positioning of the lens. This precise z positioning is important for accurate reconstruction of images—with the less-precise standard positioning, our image reconstruction and results suffer considerably.

Our experimental images consist of approximately $1.3\text{-}\mu\text{m}$ silica particles (MicroPearl) suspended in a mixture of glycerol and water. The glycerol-water mixture is tuned

to match the refractive index of the particles by minimizing the sample scattering. For these particles, we find that the optimal refractive index is $n \approx 1.437$, corresponding to approximately 76% glycerol and 24% water. Since glycerol is hygroscopic, we controlled the concentration of glycerol and water by measuring the index of refraction rather than by measuring out the glycerol and water. We match the index of refraction of the spheres and the suspending fluid to within a few parts per thousand, resulting in practically zero scattering by the spheres of either the laser or fluorescent light. The glycerol has the additional advantage of creating a very viscous suspension, slowing down the Brownian motion of the particles. We add fluorescein sodium salt to dye the suspending fluid, at a concentration of 0.4 mg/mL. The fluorescein diffuses rapidly compared to the particles and is effectively uniformly distributed throughout the regions occupied by the fluid. By using a considerable amount of dye and a low laser power, we minimize photobleaching during our experiments. Fluorescein sodium salt (molar weight 376.27) consists of two sodium ions bound to a dye molecule. Thus, this dye concentration corresponds to approximately 2×10^{-3} moles/L of monovalent sodium ions and 10^{-3} moles/L of divalent fluorescein ions. To this solution, we added the 1.3- μm silica particles (MicroPearl) at a concentration of 6.8 mg particles per 1 mL of solution. These particles are placed in a 100- μm -deep sample cell; since the particles sediment, the experimental volume fraction is determined equally by settling and the sample cell height, as opposed to simply the density of particles in the original suspension. We allow the suspension to sediment for several hours to achieve equilibrium before taking any measurements. The data are collected over the course of 1–2 hours; we do not observe any change in the $P_s(\delta)$ from the earlier samples to the later ones.

APPENDIX B: INTERPARTICLE POTENTIAL FITTING

To extract the interparticle potential, we use molecular dynamics simulations to find $P_s(\delta)$ and vary the parameters to find the best-fit $P_s(\delta)$. Since we know the particles' positions and radii via PERI, we seed the simulation with the featured particle positions and radii and relax the particle positions thoroughly before sampling for $P_s(\delta)$. Using the extracted particle parameters enforces both the correct amount of particle radii polydispersity and the number density of particles. In the simulation, we use a standard DLVO potential, consisting of nonretarded van der Waals attractions and Debye-Hückel repulsion [52], augmented by gravitational settling with the coverslip represented by a hard potential. The free parameters we fit are the strength of the attraction, the strength and screening length of the repulsion, and the gravitational settling strength; physically, these correspond to the Hamaker constant, a combination of the particle zeta potential and

salt concentration, and the average particle density. Note that while one could incorporate additional effects, such as charge polydispersity or dipolar interactions, we found that the model above was sufficient to describe our data within error bounds, and we have omitted these other interaction terms.

Since the $P_s(\delta)$ is measured from the simulation as a histogram with a finite number of samples, each simulated $P_s(\delta)$ is somewhat noisy. In light of this noise, we use a Nelder-Mead algorithm to find a good initial estimate of the fit parameters. We then refine this estimate of the fit parameters. First, we fit the ensemble of simulations to an approximate model, which is locally linear in the fit parameters. We then use this linear model to estimate a new set of best-fit interaction parameters and refine our estimate of the potential; the curve plotted in Fig. 3 of the main text is the $P_s(\delta)$ generated from the linear model at the best-fit parameters. To estimate uncertainties in the fit, we repeat this process by fitting a random subset of half of the simulations to a model and finding the new set of best-fit interactions. Repeating this 1000 times provides an estimate on the best-fit parameters as the mean of these best-fit parameters, and the uncertainty as the standard deviation of those parameters. Finally, we also obtain an estimate of systematic errors due to misfeaturing of particle positions and radii by fixing each particle's radius to be its mean value throughout all the images it is measured in. Surprisingly, fixing each particle's radius to a value that does not fluctuate in time worsens both the reconstruction and the experimentally measured $P_s(\delta)$, producing about 3 times as many overlaps. This probably arises because, in some sense, PERI directly measures the particle separations from the microscope image—changing the separation of two particles slightly will considerably change the fraction of fluorescing dye separating them. Nevertheless, the fixed-radius data give an order-of-magnitude estimate of any systematics in the experimentally measured $P_s(\delta)$. In addition, we fit the interparticle interactions for several different forms of the potential: hard spheres, electrostatic repulsion only, electrostatic repulsion and van der Waals attraction (DLVO theory), and DLVO theory combined with a short-ranged hydrophilic repulsion. Also, it is of note that the $P_s(\delta)$ data as shown in Fig. 3 of the main text contain a non-negligible number of overlapping particles with $\delta < 0$. These overlaps are significantly outside of our allowed error ranges. However, they are not of significant concern. Tracking down the offending particles, we find that these overlaps were caused by a small cluster of aggregated particles tracked throughout the duration of the experiment. As these particles are likely partially non-spherical, the fitted overlaps are not surprising. We found that including or excluding this group of particles did not influence the fit parameters significantly.

Table I shows the extracted potential parameters for all the interparticle interactions. Each interaction potential is fit

TABLE I. *Fitted interaction potentials.* The fitted interaction parameters for three sets of interparticle potentials, for the positions and radii extracted from PERI (“fitted a ”) and for each particle’s radius fixed to its mean value over the duration of many frames (“fixed a ”). The interparticle potentials are as follows: pure electrostatic repulsion $U(\delta) = U_{\text{el}}e^{-\delta/\lambda_{\text{el}}}$, full DLVO theory (i.e., electrostatic repulsion and van der Waals attraction) with Hamaker constant A , and full DLVO with an additional, short-ranged hydrophilic repulsion $U_{\text{hyd}}e^{-\delta/\lambda_{\text{hyd}}}$. The uncertainties are the uncertainties in the fit only.

	U_{el} , kT	λ_{el} , nm	mg/kT , nm	A , kT	U_{hyd} , kT	λ_{hyd} , nm
Electrostatics, fitted a	100.6 ± 3.4	10.1 ± 0.06	385 ± 2
Electrostatics, fixed a	67.8 ± 2.3	7.3 ± 0.09	378 ± 4
DLVO, fitted a	103.6 ± 2.3	10.4 ± 0.05	390 ± 2	0.231 ± 0.006
DLVO, fixed a	100.9 ± 3.2	7.5 ± 0.06	374 ± 3	0.286 ± 0.010
DLVO + Hyd, fitted a	121.5 ± 0.7	18.6 ± 0.07	376 ± 2	0.496 ± 0.004	105.5 ± 0.6	4.7 ± 0.03
DLVO + Hyd, fixed a	121.5 ± 5.0	25.2 ± 1.8	350 ± 8	0.513 ± 0.02	107.3 ± 2.9	4.9 ± 0.1

two ways, by allowing the fitted particle’s radius to vary with time and by fixing each individual particle’s radius to its average value over the frames. With the exception of a pure hard-sphere potential, all of the various interaction potentials fit the data equally well. In particular, fitting the data with just an exponentially decaying electrostatic repulsion fits the data no better than including the van der Waals interaction. However, our data do not exclude a nonzero Hamaker constant, allowing Hamaker constants of up to approximately 0.2–0.5 kT. A hydrophilic repulsion is similarly not necessary to fit the data, but our data can accommodate hydrophilic repulsion of a reasonable strength and length scale. Since there are considerably more overlaps in the fixed radii $P_s(\delta)$, we use the interaction potentials from the data set with radii fitted by PERI as the best estimate of the fitting parameters, and the difference between the fits as an estimate of the systematic uncertainties from imperfect experimental data.

The fitted values are comparable to what would be expected from electrostatic repulsion in standard DLVO theory. Two charged spheres of equal radius a placed in a salt solution repel each other, with a potential energy Φ that depends on separation h as [52]

$$\Phi(h) = \Phi_0 \ln(1 - e^{-\kappa h}), \quad (\text{B1})$$

$$\kappa^{-1} = \left[\epsilon \epsilon_0 kT / \sum q^2 z_k^2 n_k^b \right]^{1/2}, \quad (\text{B2})$$

$$\Phi_0 = 2\pi \epsilon \epsilon_0 a \frac{1}{z^2} \phi_s^2. \quad (\text{B3})$$

At long distances $e^{-\kappa h} \ll 1$, this potential is equivalent to a simple exponential repulsion $\Phi(h) \approx \Phi_0 e^{-\kappa h}$. Here, the screening length κ^{-1} and the repulsion strength Φ_0 are determined by the suspension properties, through the suspension’s relative permittivity ϵ , the valence z_k , and bulk number density n_k^b of each of the dissolved ion species, and the electrostatic potential ϕ_s at each particle’s surface; q is the electron charge and ϵ_0 the permittivity of free space. For a 76:24 mixture of glycerol [$\epsilon(0) = 47$] and deionized water [$\epsilon(0) = 80$], we expect $\epsilon \approx 55$.

The only additional ions added to the suspension arise from the fluorescein sodium salt, which dissociates to two valence-1 sodium ions and one valence-2 fluorescein ion. For 0.4 mg/mL of dissolved fluorescein sodium salt dye, the screening length should be approximately 4.5 nm, slightly smaller than what we measure experimentally. The repulsion strength in Eq. (B3) depends on the effective ion valence z . For our suspension, the ion valence is between $z = 2$ for the fluorescein ions and $z = 1$ for the sodium ions. Approximating the ion valence as $z = 2$ and the zeta potential the same as for silica in pure water at $\phi_s = -40$ mV gives a repulsive barrier of $\Phi_0 \approx 200$ kT, similar to but slightly larger than what we find experimentally.

Likewise, while we cannot measure a Hamaker constant directly, our data are consistent with a Hamaker constant of approximately 0.5 kT, which is what we expect for silica particles in water. While an exact calculation of van der Waals interactions is difficult, an approximate effective Hamaker constant A can be calculated following Israelachvili [43]. For an index-matched suspension like ours, the zero-frequency component of the dielectric should be the dominant contribution to the Hamaker constant:

$$A_{\text{eff}} = \frac{3}{4} kT \left(\frac{\epsilon_1(0) - \epsilon_2(0)}{\epsilon_1(0) + \epsilon_2(0)} \right)^2, \quad (\text{B4})$$

where ϵ_1 and ϵ_2 are the zero-frequency relative permittivities of the medium and particle, respectively. For silica, $\epsilon_2 = 4$, and the suspending fluid has $\epsilon_1 = 55$, giving a Hamaker constant of $A_{\text{eff}} \approx 0.6kT$. This value is right on the border of our measurable uncertainty.

-
- [1] E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess, *Imaging Intracellular Fluorescent Proteins at Nanometer Resolution*, *Science* **313**, 1642 (2006).
- [2] M. J. Rust, M. Bates, and X. Zhuang, *Sub-Diffraction-Limit Imaging by Stochastic Optical Reconstruction Microscopy (Storm)*, *Nat. Methods* **3**, 793 (2006).

- [3] S. W. Hell and J. Wichmann, *Breaking the Diffraction Resolution Limit by Stimulated Emission: Stimulated-Emission-Depletion Fluorescence Microscopy*, *Opt. Lett.* **19**, 780 (1994).
- [4] P. Kner, B. B. Chhun, E. R. Griffis, L. Winoto, and M. G. L. Gustafsson, *Super-Resolution Video Microscopy of Live Cells by Structured Illumination*, *Nat. Methods* **6**, 339 (2009).
- [5] B.-C. Chen, W. R. Legant, K. Wang, L. Shao, D. E. Milkie, M. W. Davidson, C. Janetopoulos, X. S. Wu, J. A. Hammer, Z. Liu *et al.*, *Lattice Light-Sheet Microscopy: Imaging Molecules to Embryos at High Spatiotemporal Resolution*, *Science* **346**, 1257998 (2014).
- [6] S.-H. Lee and D. G. Grier, *Holographic Microscopy of Holographically Trapped Three-Dimensional Structures*, *Opt. Express* **15**, 1505 (2007).
- [7] J. Fung, K. E. Martin, R. W. Perry, D. M. Kaz, R. McGorty, and V. N. Manoharan, *Measuring Translational, Rotational, and Vibrational Dynamics in Colloids with Digital Holographic Microscopy*, *Opt. Express* **19**, 8051 (2011).
- [8] S. S. Rogers, T. A. Waigh, X. Zhao, and J. R. Lu, *Precise Particle Tracking Against a Complicated Background: Polynomial Fitting with Gaussian Weight*, *Phys. Biol.* **4**, 220 (2007).
- [9] R. Pathasarathy, *Rapid, Accurate Particle Tracking by Calculation of Radial Symmetry Centers*, *Nat. Methods* **9**, 724 (2012).
- [10] F. Gröll, M. Kirchgessner, R. Kaufmann, M. Hausmann, and U. Kbschull, *Accelerating Image Analysis for Localization Microscopy with FPGAs*, in *2011 International Conference on Field Programmable Logic and Applications (FPL)*, (IEEE, New York, 2011), pp. 1–5.
- [11] S. M. Anthony and S. Granick, *Image Analysis with Rapid and Accurate Two-Dimensional Gaussian Fitting*, *Langmuir* **25**, 8152 (2009).
- [12] C. S. Smith, N. Joseph, B. Rieger, and K. A. Lidke, *Fast, Single-Molecule Localization that Achieves Theoretically Minimum Uncertainty*, *Nat. Methods* **7**, 373 (2010).
- [13] S. Andersson, *Localization of a Fluorescent Source without Numerical Fitting*, *Opt. Express* **16**, 18714 (2008).
- [14] Y. Gao and M. L. Kilfoil, *Accurate Detection and Complete Tracking of Large Populations of Features in Three Dimensions*, *Opt. Express* **17**, 4685 (2009).
- [15] P. J. Lu, M. Shutman, E. Sloutskin, and A. V. Butenko, *Locating Particles Accurately in Microscope Images Requires Image-Processing Kernels to be Rotationally Symmetric*, *Opt. Express* **21**, 30755 (2013).
- [16] M. E. Huber, N. C. Koshnick, H. Bluhm, L. J. Archuleta, T. Azua, P. G. Björnsson, B. W. Gardner, S. T. Halloran, E. A. Lucero, and K. A. Moler, *Gradiometric Micro-Squid Susceptometer for Scanning Measurements of Mesoscopic Samples*, *Rev. Sci. Instrum.* **79**, 053704 (2008).
- [17] K. C. Nowack, E. M. Spanton, M. Baenninger, M. König, J. R. Kirtley, B. Kalisky, C. Ames, P. Leubner, C. Brüne, H. Buhmann *et al.*, *Imaging Currents in HgTe Quantum Wells in the Quantum Spin Hall Regime*, *Nat. Mater.* **12**, 787 (2013).
- [18] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevX.7.041007> for an extended discussion of the generative model, computational details of its implementation, examination of additional complications such as Brownian motion, and comparison of PERI with other particle location methods.
- [19] S. Ram, E. S. Ward, and R. J. Ober, *Beyond Rayleigh's Criterion: A Resolution Measure with Application to Single-Molecule Microscopy*, *Proc. Natl. Acad. Sci. U.S.A.* **103**, 4457 (2006).
- [20] C. R. Rao, *Information and Accuracy Attainable in the Estimation of Statistical Parameters*, *Bull. Calcutta Math. Soc.* **37**, 81 (1945).
- [21] During the exposure of a recording camera, 3 nm is almost always less than a free colloidal particle's Brownian motion. Nevertheless, we show in SM [18] (Sec. IV) that the particle's mean position during the camera's exposure is measurable to within 3 nm.
- [22] K. Keville, E. Franses, and J. Caruthers, *Preparation and Characterization of Monodisperse Polymer Microspheres*, *J. Colloid Interface Sci.* **144**, 103 (1991).
- [23] A. Mohraz and M. J. Solomon, *Direct Visualization of Colloidal Rod Assembly by Confocal Microscopy*, *Langmuir* **21**, 5298 (2005).
- [24] A. Kuijk, A. van Blaaderen, and A. Imhof, *Synthesis of Monodisperse, Rodlike Silica Colloids with Tunable Aspect Ratio*, *J. Am. Chem. Soc.* **133**, 2346 (2011).
- [25] S. Guttman, Z. Sapir, M. Schultz, A. V. Butenko, B. M. Ocko, M. Deutsch, and E. Sloutskin, *How Faceted Liquid Droplets Grow Tails*, *Proc. Natl. Acad. Sci. U.S.A.* **113**, 493 (2016).
- [26] S. Hell, G. Reiner, C. Cremer, and E. H. K. Stelzer, *Aberrations in Confocal Fluorescence Microscopy Induced by Mismatches in Refractive Index*, *J. Microsc.* **169**, 391 (1993).
- [27] T. D. Visser and S. H. Wiersma, *Electromagnetic Description of Image Formation in Confocal Fluorescence Microscopy*, *J. Opt. Soc. Am. A* **11**, 599 (1994).
- [28] B. Zhang, J. Zerubia, and J.-C. Olivo-Marin, *Gaussian Approximations of Fluorescence Microscope Point-Spread Functions*, *Appl. Opt.* **46**, 1819 (2007).
- [29] M. J. Nasse and J. C. Woehl, *Realistic Modeling of the Illumination Point Spread Function in Confocal Scanning Optical Microscopy*, *J. Opt. Soc. Am. A* **27**, 295 (2010).
- [30] J.-A. Conchello and J. W. Lichtman, *Theoretical Analysis of a Rotating-Disk Partially Confocal Scanning Microscope*, *Appl. Opt.* **33**, 585 (1994).
- [31] E. Dusch, T. Dorval, N. Vincent, M. Wachsmuth, and A. Genovesio, *Three-Dimensional Point Spread Function Model for Line-Scanning Confocal Microscope with High-Aperture Objective*, *J. Microsc.* **228**, 132 (2007).
- [32] R. Wolleschensky, B. Zimmermann, R. Ankerhold, and M. Kempe, *High-Speed Scanning Confocal Microscope for the Life Sciences*, in *European Conference on Biomedical Optics 2005* (International Society for Optics and Photonics, 2005), pp. 58600N–58600N, <http://dx.doi.org/10.1117/12.632957>.
- [33] E. J. Botcherby, M. J. Booth, R. Juškaitis, and T. Wilson, *Real-Time Slit Scanning Microscopy in the Meridional Plane*, *Opt. Lett.* **34**, 1504 (2009).
- [34] J. C. Crocker and D. G. Grier, *Methods of Digital Video Microscopy for Colloidal Studies*, *J. Colloid Interface Sci.* **179**, 298 (1996).

- [35] D. W. Marquardt, *An Algorithm for Least-Squares Estimation of Nonlinear Parameters*, *J. Soc. Ind. Appl. Math.* **11**, 431 (1963).
- [36] M. K. Transtrum, B. B. Machta, and J. P. Sethna, *Why Are Nonlinear Fits to Data So Challenging?*, *Phys. Rev. Lett.* **104**, 060201 (2010).
- [37] M. K. Transtrum, B. B. Machta, and J. P. Sethna, *Geometry of Nonlinear Least Squares with Applications to Sloppy Models and Optimization*, *Phys. Rev. E* **83**, 036701 (2011).
- [38] M. K. Transtrum and J. P. Sethna, *Improvements to the Levenberg-Marquardt Algorithm for Nonlinear Least-Squares Minimization*, arXiv:1201.5885.
- [39] R. M. Neal, *Slice Sampling*, *Ann. Stat.* **31**, 705 (2003).
- [40] A. Einstein, *On the Movement of Small Particles Suspended in Stationary Liquids Required by the Molecular-Kinetic Theory of Heat*, *Ann. Phys. (Berlin)* **322**, 549 (1905).
- [41] D. B. Allan, T. A. Caswell, and N. C. Keim, Trackpy v0.2 (2014), <http://soft-matter.github.io/trackpy/v0.3.2/>.
- [42] W. B. Russel, D. A. Saville, and W. R. Schowalter, *Colloidal Dispersions* (Cambridge University Press, Cambridge, England, 1989).
- [43] J. N. Israelachvili, *Intermolecular and Surface Forces*, 3rd ed. (Academic Press, New York, 2011).
- [44] J. N. Israelachvili and G. E. Adams, *Measurement of Forces Between Two Mica Surfaces in Aqueous Electrolyte Solutions in the Range 0–100 nm*, *J. Chem. Soc., Faraday Trans. 1* **74**, 975 (1978).
- [45] W. A. Ducker, T. J. Senden, and R. M. Pashley, *Direct Measurement of Colloidal Forces Using an Atomic Force Microscope*, *Nature (London)* **353**, 239 (1991).
- [46] D. C. Prieve, *Measurement of Colloidal Forces with TIRM*, *Adv. Colloid Interface Sci.* **82**, 93 (1999).
- [47] J. C. Crocker and D. G. Grier, *Methods of Digital Video Microscopy for Colloidal Studies*, *J. Colloid Interface Sci.* **179**, 298 (1996).
- [48] Source code available with documentation and tutorials at <http://www.lassp.cornell.edu/sethna/peri/index.html>.
- [49] B. Ovryn and S. H. Izen, *Imaging of Transparent Spheres through a Planar Interface Using a High-Numerical-Aperture Optical Microscope*, *J. Opt. Soc. Am. A* **17**, 1202 (2000).
- [50] L. Reimer and H. Kohl, *Transmission Electron Microscopy: Physics of Image Formation*, 5th ed. (Springer, New York, 2008), Vol. 36.
- [51] C. J. Chen, *Introduction to Scanning Tunneling Microscopy* (Oxford University Press on Demand, New York, 1993), Vol. 4.
- [52] W. B. Russel, D. A. Saville, and W. R. Schowalter, *Colloidal Dispersions* (Cambridge University Press, Cambridge, England, 1989).