



Soft Matter

**Machine learning enables precise holographic
characterization of colloidal materials in real time**

Journal:	<i>Soft Matter</i>
Manuscript ID	SM-ART-09-2022-001283.R3
Article Type:	Paper
Date Submitted by the Author:	16-Mar-2023
Complete List of Authors:	Altman, Lauren; NYU, Department of Physics Grier, David; NYU, Department of Physics

SCHOLARONE™
Manuscripts

Cite this: DOI: 10.1039/xxxxxxxxxx

Machine learning enables precise holographic characterization of colloidal materials in real time

Lauren E. Altman^a, and David G. Grier^{*,a}Received Date
Accepted Date

DOI: 10.1039/xxxxxxxxxx

www.rsc.org/journalname

Holographic particle characterization uses in-line holographic video microscopy to track and characterize individual colloidal particles dispersed in their native fluid media. Applications range from fundamental research in statistical physics to product development in biopharmaceuticals and medical diagnostic testing. The information encoded in a hologram can be extracted by fitting to a generative model based on the Lorenz-Mie theory of light scattering. Treating hologram analysis as a high-dimensional inverse problem has been exceptionally successful, with conventional optimization algorithms yielding nanometer precision for a typical particle's position and part-per-thousand precision for its size and index of refraction. Machine learning previously has been used to automate holographic particle characterization by detecting features of interest in multi-particle holograms and estimating the particles' positions and properties for subsequent refinement. This study presents an updated end-to-end neural-network solution called CATCH (Characterizing and Tracking Colloids Holographically) whose predictions are fast, precise, and accurate enough for many real-world high-throughput applications and can reliably bootstrap conventional optimization algorithms for the most demanding applications. The ability of CATCH to learn a representation of Lorenz-Mie theory that fits within a diminutive 200 kB hints at the possibility of developing a greatly simplified formulation of light scattering by small objects.

1 Introduction

Machine learning algorithms are revolutionizing measurement science by decoupling quantitative analysis of experimental data from the mathematical representation of the underlying theory^{1,2}. The abstract representation of a measurement principle that is encoded in a well-designed and well-trained machine-learning system can rival the precision and accuracy attained by fitting to an analytic theory and typically yields results substantially faster. Gains in speed and robustness have been particularly impressive for measurement techniques based on video streams^{3–6}, which typically involve distilling small quantities of valuable information from large volumes of noisy data. Previous studies have demonstrated that machine-learning algorithms dovetail well with holographic video microscopy⁷, identifying features of interest within experimentally recorded holograms^{4,8,9} and extracting individual particles' positions and characteristics from the information encoded in those features^{4,9,10}.

Using holography to count, track and characterize colloidal particles provides unprecedented insights into the composition

and microscopic dynamics of colloidal dispersions^{11–13}, with applications ranging from fundamental research in statistical physics^{14,15} to formulation and manufacture of biopharmaceuticals^{16–19} and medical testing^{20,21}. Hologram analysis is a challenging inverse problem^{7,22} both because recorded intensity patterns necessarily omit half of the information about the light's amplitude and phase profiles and also because the underlying Lorenz-Mie theory of light scattering is notoriously complicated^{23–25}. Extracting quantitative information from holograms is an unusual application for machine learning in two respects: (1) it involves regression of continuously varying properties from experimental data and (2) the machine-learning system can be trained with synthetic data generated from an exact theory^{4,10,26}. The trained system therefore embodies a simplified representation of the underlying theory over a specified parameter domain that can be computed rapidly enough to be useful for real-world applications.

Previous machine-learning implementations of holographic particle characterization surpassed conventional algorithms²⁷ for detecting features associated with particles in complicated multi-particle holograms^{4,8,9}. They fared less well, however, at reliably extracting information from those features^{4,9,10}, typically resolving particle radius and refractive index with 5% accuracy⁹, compared with the part-per-thousand resolution obtained with iter-

^a Department of Physics and Center for Soft Matter Research, New York University, New York, NY 10003, USA

* Corresponding author




Fig. 1 Schematic representation of Lorenz-Mie microscopy using CATCH machine-learning analysis. (a) Collimated laser light illuminates a colloidal sample. Light scattered by a particle interferes with the rest of the illumination in the focal plane of a microscope, which magnifies and relays the interference pattern to a video camera. (b) A typical recorded hologram of micrometer-diameter silica and polystyrene colloidal spheres. Superimposed boxes denote features corresponding to individual particles. (c) The CATCH machine-learning pipeline consists of two modules. The Localizer, implemented with YOLOv5, finds in-plane coordinates, x_p and y_p , for detected features and generates bounding boxes such as the examples in (b). Feeding these features into the Estimator yields predictions for \mathbf{r}_p , a_p and n_p .

ative optimization¹⁴. Even so, the precision afforded by such machine-learning implementations is competitive with standard particle-resolved sizing techniques such as electron microscopy and is good enough to bootstrap iterative optimization for especially demanding applications. Most importantly, machine-learning analysis can be applied to novel systems without requiring *a priori* knowledge of their composition.

Guided by an analysis of the information content encoded in colloidal particles' holograms, we designed and implemented a deep neural network called CATCH that rapidly performs fully-automated analyses of in-line holographic microscopy images to detect, localize and characterize individual colloidal particles⁴. Here, we introduce enhancements to the CATCH architecture that improve the precision and accuracy of parameter estimation substantially enough to rival iterative optimization algorithms across best-case parameter ranges. The availability of a fast end-to-end solution for colloidal characterization creates opportunities for high-throughput applications in areas such as medical diagnostics^{20,28,29} and industrial process control¹⁸. The ability of CATCH to encapsulate the complexities of Lorenz-Mie theory in a small memory footprint furthermore hints at the existence of a simplified representation of light-scattering theory that would benefit areas as diverse as astrophysics and industrial materials characterization.

1.1 Lorenz-Mie Microscopy

Figure 1(a) schematically represents an in-line holographic microscope that is suitable for characterizing and tracking colloidal particles^{11,30}. A sample containing colloidal particles is illuminated by a collimated laser beam whose electric field may be modeled as a plane wave of frequency ω and vacuum wavelength λ propagating along the \hat{z} axis,

$$\mathbf{E}_0(\mathbf{r}, t) = E_0 e^{ikz} e^{-i\omega t} \hat{x}. \quad (1)$$

Here, E_0 is the field's amplitude and $k = 2\pi n_m / \lambda$ is the wavenumber of light in a medium of refractive index n_m . The beam is

assumed to be linearly polarized along \hat{x} . Our implementation uses a fiber-coupled diode laser (Coherent, Cube) operating at $\lambda = 447 \text{ nm}$. The 10 mW beam is collimated at a diameter of 3 mm, which more than fills the input pupil of the microscope's objective lens (Nikon Plan Apo, 100 \times , numerical aperture 1.4, oil immersion). The objective lens relays images through a 200 mm tube lens to a gray-scale camera (FLIR, Flea3 USB 3.0) with a 1280 pixel \times 1024 pixel sensor, yielding a system magnification of 48 nm/pixel and a dynamic range of 8 bits/pixel.

A colloidal particle located at \mathbf{r}_p scatters a small proportion of the illumination to position \mathbf{r} in the focal plane of the microscope,

$$\mathbf{E}_s(\mathbf{r}, t) = E_0 e^{-ikz_p} \mathbf{f}_s(k(\mathbf{r} - \mathbf{r}_p)) e^{-i\omega t}. \quad (2)$$

The scattered wave's relative amplitude, phase and polarization are described by the Lorenz-Mie scattering function, $\mathbf{f}_s(k\mathbf{r})$, which generally depends on the particle's size, shape, orientation and composition^{23–25}. For simplicity, we model the particle as an isotropic homogeneous sphere, so that $\mathbf{f}_s(k\mathbf{r})$ depends only on the particle's radius, a_p , and refractive index, n_p .

The incident and scattered waves interfere in the microscope's focal plane. The resulting interference pattern is magnified by the microscope and is relayed to the camera³¹, which records its intensity. Each snapshot in the camera's video stream constitutes a hologram of the particles in the observation volume. The image in Fig. 1(b) is a typical experimentally recorded hologram of colloidal silica and polystyrene spheres.

The distinguishing feature of Lorenz-Mie microscopy is the method used to extract information from recorded holograms. Rather than attempting to reconstruct the three-dimensional light field that created the hologram, Lorenz-Mie microscopy instead treats the analysis as an inverse problem, modeling the intensity pattern recorded in the plane $z = 0$ as¹¹

$$I(\mathbf{r}) = E_0^2 \left| \hat{x} + e^{-ikz_p} \mathbf{f}_s(k(\mathbf{r} - \mathbf{r}_p)) \right|^2 + I_0, \quad (3)$$

where I_0 is the calibrated dark count of the camera. Fitting

Eq. (3) to a measured hologram yields estimates for the three-dimensional position, \mathbf{r}_p , radius, a_p , and refractive index, n_p , for each particle in the field of view.

2 Algorithms for hologram analysis

2.1 Feature Detection and Localization

Analyzing a hologram such as the example in Fig. 1(b) begins with detecting features of interest in the recorded image. This is a challenging image analysis problem because the number of features typically is not known *a priori*, each feature can cover a large area with alternating bright and dark fringes, and neighboring particles' fringes can interfere with each other. Circular Hough transforms^{27,32,33}, voting algorithms³² and symmetry-based transforms^{27,34} leverage a feature's radial symmetry to coalesce its concentric rings into a simple peak that can be detected with standard particle-tracking algorithms³⁵. Image noise and interference artifacts can violate the assumptions underlying these algorithms, leading to poor localization and an undesirable rate of false-positive and false-negative detections⁴.

2.2 Pixel Selection

Having detected and localized a feature, the analytical pipeline selects pixels for further analysis. Limiting the selection to too small a range discards information from the diffraction pattern's outer fringes. Selecting too large a range reduces the sample's signal-to-noise ratio and, worse, can introduce interference from neighboring spheres. A suitable range can be estimated by counting diffraction fringes²⁷. Additional efficiency can be gained by sampling a subset of the pixels within that range³⁶.

2.3 Parameter Estimation

Information is extracted from the selected pixels by fitting their intensity values with the generative model in Eq. (3). Such fits typically involve iterative nonlinear refinement of the adjustable parameters whose convergence to an optimal solution is never certain³⁷. Successful optimization relies on good starting estimates for the adjustable parameters and typically yields values with part-per-thousand precision¹⁴.

Pioneering implementations of holographic particle characterization relied on manual annotation of features in holograms and *a priori* knowledge of particle properties to initialize fits to generative models¹¹. Automated initialization might use wavefront curvature to estimate axial position^{11,38,39} and fringe spacings to estimate particle size^{11,38,40,41}. These methods typically work well over a limited range of parameters. Monte Carlo methods can cover a wider range by initializing fits from multiple starting points and selecting the best solution overall⁴². This approach achieves robust convergence, but at a considerably higher computational cost.

2.4 Effective parameter ranges

The Lorenz-Mie theory for light scattering by homogeneous spheres is the simplest and most effective model for analyzing holograms of colloidal particles. Real-time implementations re-

turn tracking and characterization data as fast as the camera records holograms. Applying this analysis to holograms of aspherical and inhomogeneous particles yields values for particle position, size and refractive index that reflect the properties of an effective sphere enclosing the particle⁴³. Effective-sphere properties can be related to an inhomogeneous particles' true properties through effective-medium theory^{16,43-47}. When applied to colloidal dimers, for example, effective-sphere analysis can yield the asymmetric particle's three-dimensional orientation in addition to its three dimensional position²⁹. Practical implementations of Lorenz-Mie analysis, however, are limited by instrumental and computational constraints.

2.4.1 Axial Position

The scale and nature of a recorded holographic feature depends on how far the scattered light propagates before it reaches the imaging plane and the phase of the reference beam at that plane. As the particle approaches the imaging plane, the separation between diffraction fringes becomes smaller than the camera's pixel size, and information about the particle's properties are lost. The spatial resolution of our reference instrument sets the lower bound for axial tracking at roughly $z_p \geq 5\mu\text{m}$. Conversely, as the particle moves away from the focal plane, its scattering pattern spreads over increasingly many pixels to the detriment of the signal-to-noise ratio. This sets an upper limit on axial tracking in our microscope to $z_p \leq 50\mu\text{m}$ ³¹.

2.4.2 Particle Size

Both commercial and academic implementations of holographic particle characterization work with particles ranging in diameter from 500 nm to $10\mu\text{m}$ ¹¹. The lower limit is set by the low signal-to-noise ratio for in-line holograms created by weak scatterers. Switching to off-axis or dark-field holography improves the signal-to-noise ratio for weak scatterers and extends the lower size limit down to 50 nm⁴⁸. The upper limit is set by the tendency of large particles to scatter light strongly enough to saturate the camera. Mitigating this effect by moving to lower magnification, including lensless implementations, provides tracking and sizing information for particles as large as $50\mu\text{m}$ using simplified generative models⁴⁹.

2.4.3 Refractive Index

Hologram fitting has been demonstrated for dielectric spheres with refractive indexes ranging from bubbles with $n_p = 1.0$ ¹⁸ up to titanium dioxide with $n_p = 2.8$ ¹¹. The particles can have refractive indexes either higher or lower than that of the surrounding medium⁴⁷. Successful tracking and characterization require only that the refractive index of the particle differ from that of the medium by $\Delta n = n_p - n_m = \pm 0.002$ ¹⁸.

2.4.4 Morphology

The choice of scattering function, $\mathbf{f}_s(k(\mathbf{r} - \mathbf{r}_p))$, in Eq. (3) establishes what kinds of particles can be analyzed. The present study focuses on the Lorenz-Mie scattering function for homogeneous spheres. Other choices include scattering functions for core-shell particles and layered spheres, for ellipsoids, and for spherocylinders⁵⁰⁻⁵². Elementary scattering functions can be combined to




Fig. 2 Schematic overview of the Estimator network. A cropped hologram with dimensions $w \times w$ is scaled to a standard size of 201 pixels \times 201 pixels before being fed in to a cascade of convolutional layers that distill it into a 400-wide vector. These values together with the scale factor are analyzed by a fully-connected layer to produce a 20-wide vector that serves as an optimal representation of the information in the original hologram. This representation is then parsed by separate fully-connected layers into estimates for the particle's axial position, z_p , radius, a_p , and refractive index, n_p .

treat more highly structured particles such as dimers and clusters of spheres^{53,54}. Increasing the complexity of the model increases the demands on the analytical pipeline to find optimal solutions for each of the model's adjustable parameters.

3 CATCH

Figure 1(c) presents the CATCH machine-learning system that performs all of the analytical operations identified in Sec. 2 with an integrated pipeline⁴. Machine-learning algorithms have been adopted for holographic particle characterization to expand the effective range, improve robustness against false positive and false negative detections, and reduce processing time compared with conventional image-analysis techniques^{4,8,10,55}. The original implementations used distinct types of trainable algorithms for feature detection^{4,8}, pixel selection, and parameter regression^{4,10}. CATCH, by contrast, uses the widely-adopted YOLO family of object-detection networks^{56–59} to detect and localize features^{4,55} and then feeds the selected pixels directly into a custom Estimator network that extracts optimal values for the adjustable parameters.

The original implementation of CATCH⁴ uses YOLOv3, which is based on the open-source darknet library⁵⁸. The complementary Estimator network is written in TensorFlow⁶⁰. Having both darknet and TensorFlow as separate requirements complicates installation, maintenance and customization. The two systems, furthermore, require distinct training protocols.

The version of CATCH developed for this study (CATCHv2) uses YOLOv5, which is built with the PyTorch⁶¹ machine-learning framework. The updated Estimator also is defined in PyTorch, thereby fully integrating the two stages, simplifying installation and maintenance, and facilitating training.

3.1 Continuous Scaling

CATCHv2 features a set of critical innovations that dramatically improve its performance. The most important of these involves how features identified by the Localizer are transferred into the Estimator. Features vary in size depending on the nature and

position of the particle. As shown in Fig. 2, each feature must be scaled from its true dimensions, $w \times w$, to a standard size of 201 pixels \times 201 pixels before it can be processed by the Estimator. The original implementation of CATCH either cropped a given feature to this size or else scaled it by an integer factor before cropping, depending on the ideal size determined by the Localizer. CATCHv2 instead continuously scales the block of pixels to the required size with bilinear interpolation.

Continuous scaling allows features to be precisely cropped for analysis, which improves the signal-to-noise ratio. It also increases the system's reliance on the Localizer to estimate feature extents accurately. CATCHv2 achieves this by adopting the “small” variant of YOLOv5, which estimates feature extent with single-pixel precision. The original CATCH implementation, by contrast, achieved ten-pixel precision using the “tiny” variant of YOLOv3⁴. Improving the estimate for feature extent also encodes more information about the particle's size and axial position in the scale factor. CATCHv2 leverages this to improve the precision of its overall parameter estimation.

3.2 Architecture of the Estimator

CATCH's Estimator uses four convolutional layers to distill the 40401 8-bit values that constitute a scaled feature into a set of 400 single-precision floating-point values. This set together with the scale factor is then processed by a fully-connected neural-network layer into a 20-element vector that optimally represents the particle's properties in an abstract vector space that is parameterized by z_p , a_p and n_p . The vector of values computed from the feature is parsed by three specialized fully-connected layers into each of these parameters.

The Estimator's convolutional layers distill information from a feature using sets of 3 pixel \times 3 pixel masks. Intermediate results are combined with three stages of two-fold max-pooling and one stage of four-fold max-pooling. The fully-connected layers use rectified linear unit (ReLU) activation, which has been shown to facilitate rapid training in regression networks⁶². Final results are scaled into physical units by a linear fully-connected layer.

The vector space of optimal representations may be viewed as an idealized model of the Lorenz-Mie scattering theory in the relevant range of parameter values. If the network can be appropriately trained, the 20 values that span the space could encode 10^{96} distinct values of z_p , a_p and n_p , which would be more than enough to estimate parameters with the part-per-thousand precision provided by conventional optimization algorithms. The entire Estimator has 34983 trainable parameters that can be stored in 200 kB. Whether such a small network can achieve the potential suggested by the naive interpretation of the optimal representation depends on the success of the training protocol.

3.3 Training

Like its predecessors^{4,8,10}, CATCHv2 streamlines training by using a generative model such as Eq. (3) to produce synthetic data with established ground-truth parameters, rather than relying on manually annotated experimental data. Synthetic training data not only eliminates annotation errors but also can cover the parameter space more comprehensively than would be feasible with experimentally-derived data.

3.3.1 Training Data

The training set consists of 10^5 synthetic holograms for training and an additional 10^4 for validation. Within the same field of view as the reference microscope, each hologram is computed as the superposition of up to 6 particles' fields. This number corresponds to experimental concentrations up to 10^6 particles/mL. The simulated particles' properties are drawn at random from the range $a_p \in [200\text{ nm}, 5\text{ }\mu\text{m}]$ and $n_p \in [1.38, 2.500]$. Particles are placed at random in the field of view and are located in the range $z_p \in [2.5\text{ }\mu\text{m}, 29\text{ }\mu\text{m}]$ with the caveat that collisions between particles are not allowed. To further mimic experimental data, each calculated hologram is scaled to a mean intensity value of 100 and cast to 8 bits/pixel. These images are degraded with 5% additive Gaussian noise, which is consistent with the median-absolute-deviation noise estimate for holograms recorded by the reference holographic microscope, including the example in Fig. 1 (b). Incorporating noise into the training holograms and allowing for overlapping features helps to prevent overtraining and improves the network's performance with experimental holograms.

3.3.2 Training Protocol

Both the Localizer and the Estimator are trained by backpropagation. The Localizer uses stochastic gradient descent (SGD) for its optimizer, and the Estimator is trained using root-mean-square propagation (RMSprop)⁶³. The original implementation of CATCH minimized the L2 loss for both the Localizer and the Estimator, which is equivalent to minimizing squared errors in the estimates for the features' centroids and extents, and also for the particles' properties and axial positions. While effective for training YOLO, L2 loss overemphasizes outliers due to particularly problematic holograms. Such bad outcomes are inherent in the optimization problem because the Lorenz-Mie theory admits near-degeneracies in which distinct sets of parameters produce nearly identical holograms³⁷. CATCHv2 deemphasizes degeneracies by minimizing the smooth-L1 loss, which interpolates

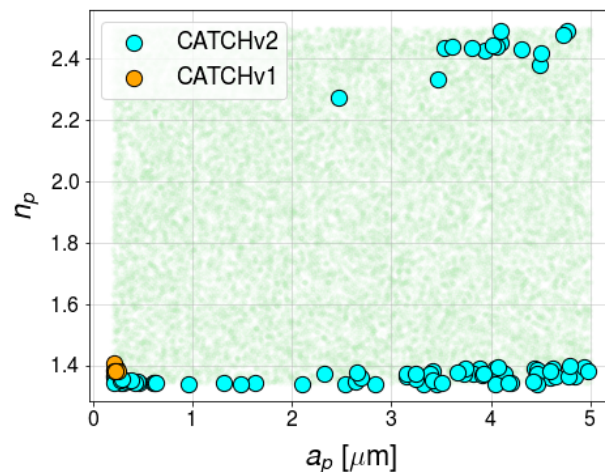


Fig. 3 False negative detections from the CATCHv1 and CATCHv2 Localizers as a function of a_p and n_p . Green points represent the 25000 sets of properties tested. The 6 features not detected by CATCHv1 (orange) are all small and weakly scattering. The 86 particles missed by CATCHv2 (teal) are either weakly scattering because they are nearly index matched to the medium, or else scatter light especially strongly because they are large and have high refractive indexes.

between the mean-square error for small errors and the mean-absolute error for large errors⁶⁴.

The training protocol for CATCHv2 achieves convergence for all of its outputs without overtraining any of them by incorporating early-stopping callbacks that monitor the loss metric of each output for validation data. Once the network converges on a solution for one of its outputs, the callback freezes the values of all of the network parameters that contribute to that output, thereby preventing overfitting. Training continues for the remaining network parameters until the second output converges, and then the third. This training protocol requires no user input and accounts naturally for differences in learning speed for each of the three output parameters of our Estimator model.

Using this protocol, the Estimator was trained for 7882 epochs with a batch size of 64. The Localizer was trained using a batch size of 32 for 3163 epochs. Training was performed on a desktop workstation outfitted with an NVIDIA Titan RTX graphical processing unit (GPU) for hardware acceleration. The two modules were trained sequentially using an average of 80% of the GPU's processors for a total of seven days.

The trained networks can be adapted through transfer learning⁶⁵ to work with microscopes with different wavelength and magnification and media with different refractive indexes. Typically, this only requires retraining the fully-connected layers in the final stage of the Estimator, which can be completed in two hours once a suitable set of training data has been computed.

4 Performance

4.1 Validation with Synthetic Data

We first evaluate the performance of CATCH on a set of synthetic images similar to those used in training. The Localizer is evaluated on a set of 25000 full-frame images of size $1280\text{ pixels} \times 1024\text{ pixels}$, each containing exactly one holographic

feature with randomized properties. The images are degraded with 5 % Gaussian noise. This test establishes the system's performance under ideal conditions without the added complication of overlapping features.

4.1.1 Detection Accuracy

Detection accuracy is assessed with two metrics: the rates of false positive detections and false negative detections. Of these two kinds of errors, false negatives pose a greater challenge because they correspond to a loss of information. False positives generally can be identified and eliminated at later stages of analysis. The false positive detections from both Localizer versions are plotted in Fig. 3. CATCHv1 (orange) boasted an impressive false-negative rate of 0.02 %, missing only 6 out of 25000 holographic features. By comparison, CATCHv2 (teal) performs with a false negative rate of 0.3 % on the same data set, missing a total of 86 holographic features. Neither CATCHv1 nor CATCHv2 has false-positive detections for these test images. In both cases, losses are limited to either the most weakly scattering particles or the largest and most strongly scattering particles. The overall detection efficiency of 99.7 % achieved by CATCHv2 greatly improves upon the 60 % rate previously reported for conventional algorithms over the same range of parameters⁴ and matches that of CATCHv1 over most of the parameter range. The slight loss of detection efficiency is compensated by a very substantial gain in localization accuracy, which is critical for accurate parameter estimation.

4.1.2 Localization Accuracy

We evaluate localization accuracy using the true positive detections from the previous analysis. Of those 24914 detections, 14997 were situated such that their bounding box was not cut off by the edge of the field of view. We compute the radial distance, Δr , of those features' predicted centroids from the ground truth. The original implementation of CATCH has a mean in-plane localization error of $\Delta r = 2.7 \text{ pixels} = 130 \text{ nm}$ for this data set. The updated Localizer achieves a mean localization error of $\langle \Delta r \rangle = 0.63 \text{ pixel} = 30 \text{ nm}$. Its performance across the range of parameters is summarized in Fig. 4 and Table 1.

Localization errors contribute to errors in parameter estimation if the Estimator is trained to expect perfectly centered features. This source of error can be mitigated by training the Estimator with synthetic holograms that are randomly offset to reflect the Localizer's performance. Improving the Localizer reduces these offsets and therefore reduces the complexity of the data-analysis problem that the Estimator is required to solve. This, in turn, improves the Estimator's performance.

Alternative deep-learning particle trackers such as DeepTrack⁶⁶ and LodeSTAR⁶⁷ offer substantially better in-plane localization accuracy than YOLOv5. YOLO, however, provides the reproducibly accurate bounding boxes that the CATCHv2 Estimator requires for successful particle characterization⁵⁹.

4.1.3 Parameter Estimation

We evaluate the Estimator using a separate data set consisting of 10^4 cropped holograms of spheres with randomly selected properties in the range $z_p \in [2.5 \mu\text{m}, 29 \mu\text{m}]$, $a_p \in [0.2 \mu\text{m}, 5.0 \mu\text{m}]$, and




Fig. 4 In-plane localization error, Δr , for CATCHv2 averaged over axial position, z_p . The maximum error is smaller than 2 pixel over the entire parameter space, with the largest errors occurring for the largest, most weakly-refracting particles.

Table 1 Median and maximum errors in $r = (x_p, y_p)$, z_p , a_p and n_p predicted by CATCHv1 and CATCHv2.

		CATCHv1	CATCHv2
Δr [pixel]	Median	2.7	0.63
	Max	59	7
Δz_p [pixel]	Median	15.0	7.2
	Max	659	167
Δa_p [μm]	Median	0.09	0.04
	Max	2.62	1.01
Δn_p [ppt]	Median	45.5	26.6
	Max	781	886

$n_p \in [1.38, 2.5]$. Consistent with training conditions, a feature's ideal extent is set to twice the radius of the twentieth interference node. We introduce 5 % Gaussian random offsets into the feature's extent to simulate errors by the Localizer, and then add 5 % Gaussian noise to the feature's calculated intensity.

Figure 5 and Table 1 illustrate the extent to which the modified architecture improves the network's performance. Errors in parameter estimation are presented in Fig. 5 as the mean absolute differences between CATCH estimates and the ground truth averaged over axial positions, z_p . These deviations are dominated by systematic errors that vary across the parameter domain and thus reflect the accuracy of the Estimator. Accuracy in axial localization improves by better than a factor of two across the entire range of parameters, with a median error of $\Delta z_p = 0.35 \mu\text{m}$. Accuracy in particle sizing also improves by a factor of two, with a median error of just $\Delta a_p = 40 \text{ nm}$. Errors in refractive index are reduced to a median value of $\Delta n_p = 0.027$, which is more than sufficient to differentiate particles by their composition^{11,68}. The updated network also resolves property-dependent variations in the error that are most evident when comparing Fig. 5(b) with Fig. 5(e). Relative errors, plotted in Fig. 5(g), (h) and (i), are smaller than 10 % over the entire parameter domain and are smaller than 3 % for all but the smallest and most weakly scattering particles.

Most of the improvements in CATCHv2's accuracy relative to

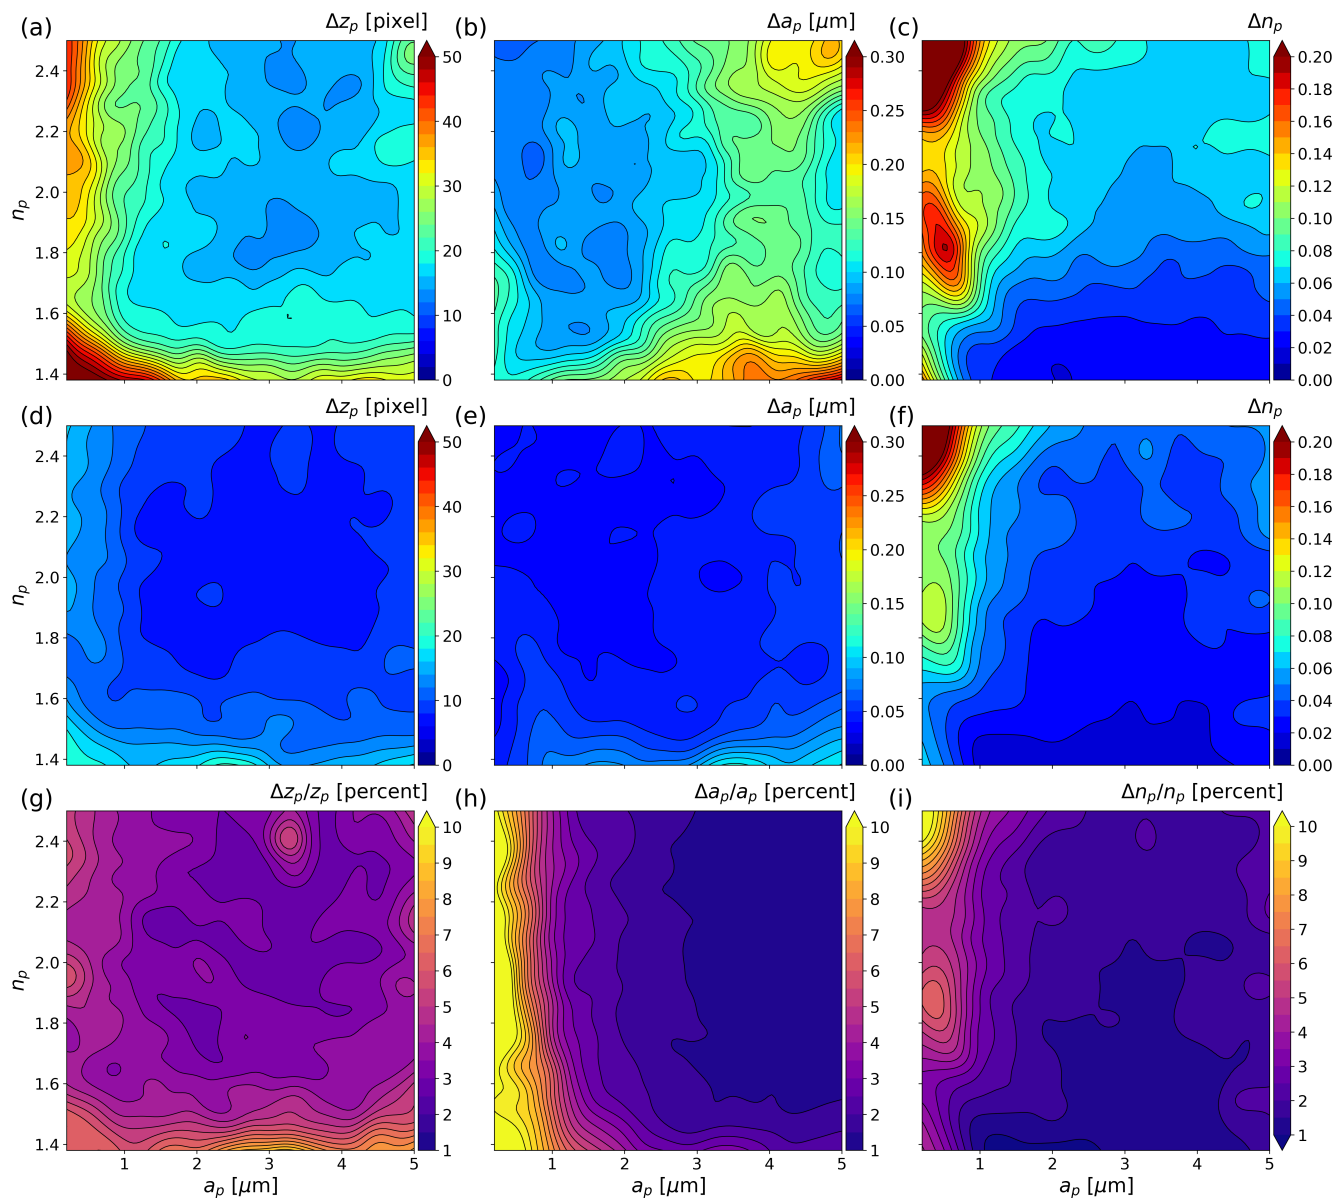


Fig. 5 Performance of the Estimator module on synthetic data for (a-c) CATCHv1 and (d-i) CATCHv2. Results are presented as a function of particle radius, a_p , and refractive index, n_p , and are averaged over axial position, z_p . (a) and (d) Absolute error in axial position, Δz_p . (b) and (e) Absolute error in particle radius, a_p . (c) and (f) Absolute error in refractive index, n_p . (g), (h) and (i) recast the CATCHv2 errors from (d), (e) and (f), respectively, as percentages of the ground-truth values. Color bars have consistent scales to aid with comparison.

CATCHv1 can be ascribed to incorporating continuous scaling into the network's optimal representation for a particle's properties. This innovation's effectiveness hinges on coordinated improvements in localization and feature-extent estimation afforded by adopting a larger and more capable Localizer network. Training with the robust smooth-L1 loss metric speeds convergence and also contributes secondarily to improvements in accuracy.

4.2 Experimental Validation

Training and validation with synthetic data does not guarantee a network's performance with experimental data. Confounding factors such as correlated noise, artifacts from normalization, and instrumental imperfections may cause experimental data to differ enough from ideal synthetic data that the trained model cannot make accurate predictions.

We illustrate the performance of CATCH on experimental data by measuring the sedimentation of a colloidal sphere between parallel walls^{4,69}. A 3 μm diameter silica bead (Bangs Laboratories, catalog number SS05N) is dispersed in a 30 μL aliquot of water that is confined between a glass #1.5 cover slip and a glass microscope slide. Silica being twice as dense as water, the bead tends to settle to the bottom of the chamber. Using a holographic optical trap⁷⁰, we lift the bead to the top of the chamber, release it and record its subsequent trajectory⁷¹. Examples of experimentally recorded holograms of this particle are presented in Fig. 6(a) when the particle is at the top of the chamber, and Fig. 6(b) when the particle is at the bottom of the chamber. We analyze the resulting holographic video with CATCHv1, CATCHv2, and by fitting to the generative model in Eq. (3) using a conventional least-squares fitter. These analyses also yield estimates for the particle's radius and refractive index.

Figure 6(c) and (e) compare results obtained with the original implementation of CATCH with results obtained by conventional fitting. Figure 6(d) and (f) present complementary results for CATCHv2. In both cases, we treat the nonlinear least-squares fit as the ground truth for the comparison. The sedimenting particle's axial position, plotted as (red) points in Fig. 6(c) and (d), follows the sigmoidal trajectory expected for confined sedimentation in a horizontal slit pore⁷². Predictions for $z_p(t)$ by CATCHv1 generally follow this trend, but with substantial random and systematic errors. CATCHv2, by contrast, tracks the particle's motion in excellent quantitative agreement with the ground truth. The updated machine-learning system improves mean errors in axial tracking by nearly a factor of ten, from $\Delta z_p = 4.2 \mu\text{m}$ to $\Delta z_p = 0.46 \mu\text{m}$, which is consistent with expectations based on the numerical validation data in Fig. 5.

The particle's radius and refractive index, plotted in Fig 6(e) and (f), form a tight cluster when reported by the conventional fitter (red). All of these values represent properties of a single particle that should not change as the particle moves through the sample cell. Results from CATCHv1 in Fig. 6(e) generally cluster in the correct region of parameter space, albeit with a systematic error of $\Delta a_p = 1 \mu\text{m}$ and a standard deviation of 500 nm. CATCHv2, by contrast, accurately estimates the radius and refractive index of the particle. The $\Delta a_p = 12 \text{ nm}$ precision and

accuracy for particle size is consistent with expectations from the numerical study in Fig. 5, and would suffice for the differential measurements required for holographic molecular binding assays^{20,21,28,73,74}. CATCHv2 therefore could provide a computationally cost-effective basis for label-free bead-based medical diagnostic testing.

Whereas the single-particle study in Fig. 6 is useful for illustrating the performance of CATCHv2 for an individual colloidal sphere moving in three dimensions, Fig 7 illustrates its performance for heterogeneous dispersions of colloidal particles. The sample for this demonstration is composed of equal concentrations of silica spheres (Thermo Fisher, catalog no. 8150) and polystyrene spheres (Bangs Laboratories, catalog no. NT16N), each with a nominal radius of $a_p = 0.75 \mu\text{m}$, dispersed in water. The hologram in Fig. 7(a) captures four of those particles as they move through the $61 \mu\text{m} \times 49 \mu\text{m}$ field of view in a pressure-driven flow. Superimposed bounding boxes are identified automatically by the Localizer stage of CATCHv2. The scatter plot in Fig. 7(b) presents characterization results from the Estimator stage for 1133 particles that flowed through the observation volume in 5 min, together with refined estimates for those particles' characteristics that were obtained by fitting to Lorenz-Mie theory. Each data point represents the radius and refractive index of one particle. Machine-learning estimates are colored by the relative density of observations, $\rho(a_p, n_p)$. The two populations of particles are clearly differentiated by refractive index, even though their size distributions overlap.

	CATCHv2	Lorenz-Mie	Error	Expected
	a_p [μm]	a_p [μm]	Δa_p [nm]	Δa_p [nm]
PS	0.770 ± 0.053	0.729 ± 0.008	42 ± 54	57 ± 74
SiO ₂	0.953 ± 0.070	0.779 ± 0.035	173 ± 79	65 ± 78
	n_p	n_p	Δn_p [ppt]	Δn_p [ppt]
PS	1.631 ± 0.040	1.598 ± 0.006	33 ± 40	42 ± 44
SiO ₂	1.455 ± 0.008	1.449 ± 0.006	6 ± 11	26 ± 40

Table 2 Particle-characterization performance of CATCHv2 for the two-component dispersion presented in Fig. 7. Population-averaged values for the radius and refractive index for the two types of particles are compared with refined estimates obtained by fitting the same set of holograms to Eq. (3). Differences between estimates and refined values are compared with the expected performance from Fig. 5.

Table 2 reports the average radii and refractive indexes estimated by CATCHv2 for the two types of particles in the sample. These population-average values are compared with the averages obtained by fitting the same holographic features to the Lorenz-Mie model. The ranges of values are computed as standard deviations of the single-particle results and therefore combine estimation errors with intrinsic particle-to-particle variations in the two populations. We interpret differences between machine-learning estimates and refined values as errors in the machine-learning estimates, Δa_p and Δn_p . Table 2 compares these discrepancies with expectations for the performance of CATCHv2 based on the numerical validation results presented in Fig. 5. In this case, the range of expected errors is computed as the mean absolute deviation of the validation results averaged over axial position, z_p . This is a reasonable treatment for a flow experiment in which particles pass through the imaging volume at different heights.

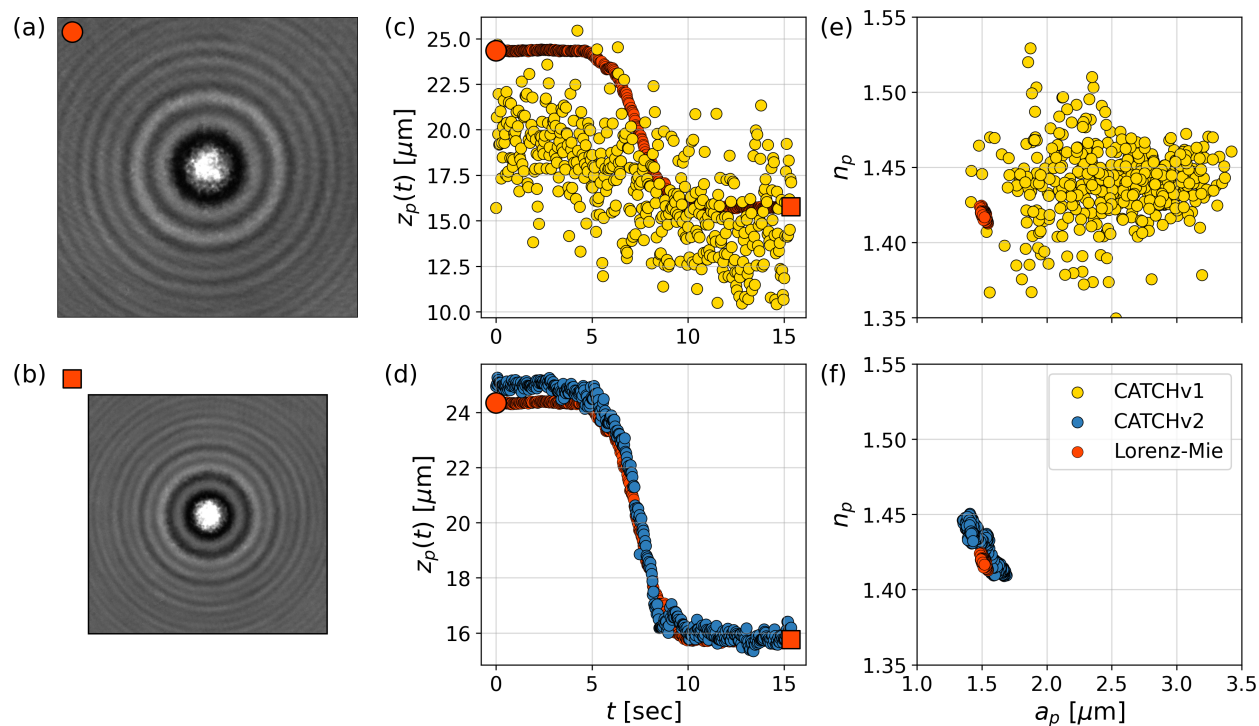


Fig. 6 (a) Experimentally recorded hologram of a colloidal silica sphere at the upper wall of a rectangular channel that is filled with water. (b) Hologram of the same sphere after it sediments to the lower wall of the channel. Both holograms are cropped to include 20 diffraction fringes. This sphere's trajectory is used to assess the performance of CATCHv1 and CATCHv2 for tracking (c, d) and characterization (e, f). Machine-learning estimates are compared with ground truth values obtained from fitting to the generative model (red). Values for the axial position, z_p , obtained from the holograms in (a) and (b) are plotted with large (red) symbols in (c) and (d). (c) Axial trajectory, $z_p(t)$, compared with predictions of CATCHv1 (orange) and (d) CATCHv2 (blue). (e) Values for the particle radius, a_p , and refractive index, n_p , estimated at each time step in the trajectory by CATCHv1 (orange) and (f) CATCHv2 (blue). Ground truth values for a_p and n_p are estimated by conventional optimization.

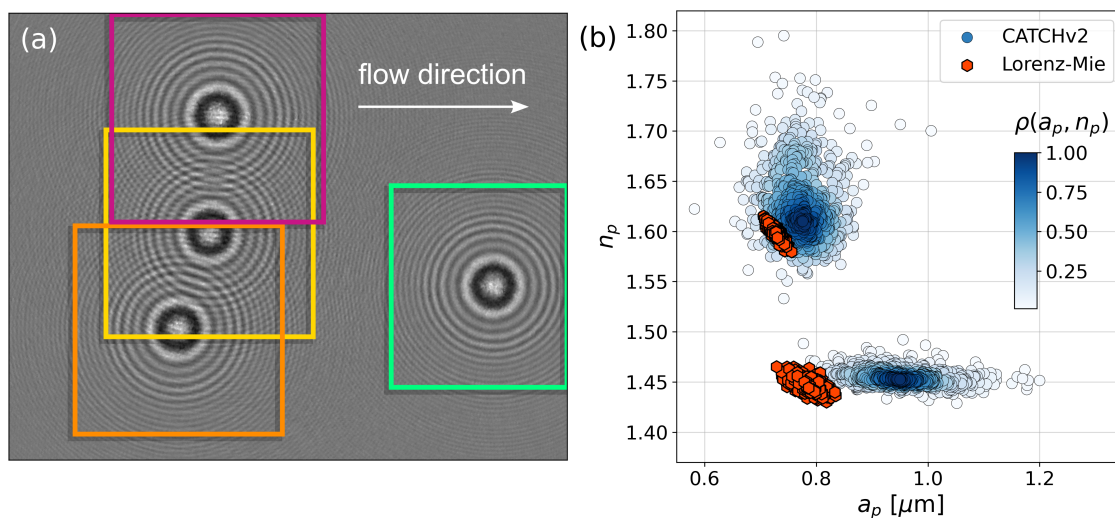


Fig. 7 Detection and characterization of particles in a heterogeneous sample. (a) Typical hologram of colloidal spheres flowing through the microscope's field of view overlaid with bounding boxes automatically detected by CATCHv2. (b) Characterization estimates provided by CATCHv2 (circles) for 1133 particles together with refined fits (hexagons) that were initialized with those estimates. CATCH results are colored by the relative density of observations, $\rho(a_p, n_p)$.

Systematic discrepancies between machine-learning estimates for the particle radii and refined values are not surprising because errors in machine-learning estimates generally are not normally distributed. Such discrepancies are likely to be exacerbated by defects such as aberrations in experimentally recorded holograms^{31,50} that are not accounted for in the generative model used to synthesize training data. Even so, errors in the polystyrene spheres' radii fall within the expected range of $\Delta a_p = \pm 60$ nm, as do the sizing errors for the $1.5 \mu\text{m}$ -radius silica sphere reported in Fig. 6. By contrast, the 170 nm systematic offset for the more weakly scattering $0.8 \mu\text{m}$ -radius silica spheres in Fig. 7 is nearly three times larger than expected. In assessing these performance, it should be noted that the only particle-resolved sizing technique that consistently surpasses the precision and accuracy of CATCHv2 in this range is the full Lorenz-Mie implementation.

Machine-learning estimates for the two populations' refractive indexes agree well with the refined values. Results for the refractive index, in particular, are sufficiently precise to distinguish the two types of particles by composition¹⁰, which represents a very substantial improvement over CATCHv1.

Because CATCHv2 performs independent estimates for size and refractive index, errors in these estimates tend not to be correlated⁴. This contrasts with conventional optimization techniques whose results more directly reflect the structure of the error surface for the theory. Correlations are less prominent, for example, in the machine-learning estimates plotted in Fig. 7(b) than in the iteratively optimized results.

Predictions by CATCHv2 are sufficiently accurate and precise that refinement may not be necessary for many applications, including classifying impurity particles for quality control in biopharmaceuticals^{17,18}, semiconductor processing⁷⁵ and environmental monitoring⁷⁶. Reliable detection of particles across a wide range of parameters ensures accurate measurement of particle concentrations^{19,76}, which is valuable across many industries. Combining this with accurate sizing and differentiation by material composition affords a particularly detailed view into the composition of heterogeneous dispersions. All such applications will benefit from the processing speed afforded by an end-to-end machine-learning implementation.

4.3 Speed

Both CATCHv1 and CATCHv2 make efficient use of hardware acceleration on CUDA-capable graphics cards. CATCHv1 processes a single $1280 \text{ pixel} \times 1024 \text{ pixel}$ hologram in 21 ms on an NVIDIA Titan Xp GPU, with 20 ms required for the Localizer and 0.9 ms for the Estimator⁴. These times are reduced by 5% when the same code is run on an NVIDIA Titan RTX GPU.

Moving from the C-language implementation of CATCHv1 to the pure python implementation of CATCHv2 incurs a performance penalty, all the more so because of the increased size and complexity of the Localizer module. Running on the Titan RTX platform, the updated Localizer requires 24 ms to process one frame, and the Estimator requires an additional 1.3 ms to analyze each feature. This is still fast enough to process frames in real time at 30 frames/s.

The dramatic improvement in prediction accuracy gained with CATCHv2 translates into particularly substantial performance gains for those applications that no longer require optimization by conventional algorithms. Even when further refinement is required, the improved initial estimates provided by CATCHv2 increase the likelihood of successful convergence and reduce the time to convergence.

Performance differences between C-based and python-based implementations should decrease as development efforts continue to improve the processing speed of python programs. CATCHv2 also would benefit from optimizations such as parameter pruning and quantization, neither of which have been applied to the demonstration implementation presented here.

5 Discussion

The implementation of the CATCH machine-learning system presented in this study solves a central problem in soft-matter science: characterizing and tracking individual colloidal particles in their native media in real time. When analyzing data from the reference holographic microscopy instrument, CATCH provides three-dimensional tracking data with $\Delta r = 50$ nm accuracy in-plane and $\Delta z_p = 350$ nm along the axial direction in a $100 \mu\text{m} \times 100 \mu\text{m} \times 30 \mu\text{m}$ observation volume. CATCH simultaneously measures a micrometer-scale particle's radius with a median accuracy of $\Delta a_p = 40$ nm for particles ranging in radius from $a_p = 200$ nm to $a_p = 5 \mu\text{m}$. In large regions of parameter space, CATCH achieves precision and accuracy which rivals that of the conventional algorithms.

Holographic characterization offers the substantial advantage relative to other particle-characterization technologies of measuring a recorded particle's refractive index, thereby providing information about its composition. CATCH estimates the refractive index with an accuracy of $\Delta n_p = 0.026$ over the range from near-index-matching, $n_p = 1.38$, to very strong scattering, $n_p = 2.5$.

These estimates for the system's accuracy are consistent with the illustrative example of a colloidal sphere sedimenting through water. The training protocol therefore appropriately accounts for instrumental imperfections that might otherwise degrade prediction accuracy.

The ability of CATCH to estimate particles' positions and characteristics with an accuracy of one or two percent is sufficient for many of the applications that already have been identified for holographic particle characterization, including particle characterization in biopharmaceuticals^{18,19}, agglomerate detection in semiconductor polishing slurries⁷⁵ and process control for materials synthesis⁷³. In more specialized applications where per-thousand accuracy is desirable, predictions from CATCH can be used to initialize parameter refinement using the generative model from Eq. (3). Both the speed and reliability of iterative optimization are improved by the high quality of the starting estimates provided by CATCH³⁷.

CATCH can be readily adapted to work with new instruments and can be trained automatically to cover different parameter ranges. For example, CATCH can be trained to accommodate particles with refractive indexes smaller than that of the medium. Training over a smaller parameter range can improve the accu-

racy of CATCH's predictions to the point that machine-learning estimates rival the precision and accuracy of state-of-the-art optimization while retaining their substantial speed advantage.

Having demonstrated that a machine-learning system can provide precise end-to-end holographic analysis of colloidal spheres in real time, we can speculate on possible generalizations of our implementation. CATCH currently treats the refractive index of the medium as a fixed parameter, for example. More generally, n_m can be allowed to vary at the cost of increased training complexity, and indeed could be obtained as an output of the Estimator. Such a generalized model would be useful for analyzing most dispersions of micrometer-scale colloids without any *a priori* knowledge about their composition and without requiring any retraining. The Localizer can be trained to differentiate holograms into categories such as "spherical", "rod-like" and "irregular", "large" and "small", "high-index" and "low-index". Such classifications could be used to transfer holograms to specialized variants of the Estimator for detailed analysis. The value of such elaborations hinges on the rapidly increasing variety of applications for holographic particle characterization.

As the simplest and presumably smallest machine-learning implementation of holographic particle characterization, the present implementation CATCH can be incorporated readily into commercial instrumentation. CATCH is small enough, for example, to be realized on a field-programmable gate array (FPGA) suitable for board-level integration.

The diminutive 200 kB memory footprint of the CATCH model also hints at opportunities for recasting Lorenz-Mie theory itself. The standard formulation of light scattering by small particles is technically challenging to compute. It is possible that the condensed representation learned by CATCH can guide the development of a greatly simplified analytic formulation⁷⁷, which would be broadly useful. CATCH therefore can play a role in the emerging paradigm shift toward machine-driven discovery of fundamental principles.

The full open-source implementation of CATCHv2 is available at <https://github.com/laltman2/CATCH/>. The open-source `pylorenzmie` package for Lorenz-Mie analysis is available at <https://github.com/davidgrier/pylorenzmie/>

Conflicts of Interest

David Grier is a founder of Spheryx, Inc., which manufactures instruments for holographic particle characterization.

Acknowledgements

This work was primarily supported by the SBIR program of the National Institutes of Health under Award Number R44TR001590. Additional support was provided by the National Science Foundation under Award Number DMR-2104837. The Titan Xp and Titan RTX GPUs used for this work were provided by GPU Grants from NVIDIA. The custom holographic characterization instrument was constructed with support from the MRI program of the NSF under Award Number DMR-0922680.

Notes and references

- 1 M. Wittwer and M. Seita, *npj Comput. Mater.*, 2022, **8**, 1–9.

- 2 G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto and L. Zdeborová, *Rev. Mod. Phys.*, 2019, **91**, 045002.
- 3 A. W. Long, J. Zhang, S. Granick and A. L. Ferguson, *Soft Matter*, 2015, **11**, 8141–8153.
- 4 L. E. Altman and D. G. Grier, *J. Phys. Chem. B*, 2020, **124**, 1602–1610.
- 5 E. N. Minor, S. D. Howard, A. A. S. Green, M. A. Glaser, C. S. Park and N. A. Clark, *Soft Matter*, 2020, **16**, 1751–1759.
- 6 W. F. Reinhart, A. W. Long, M. P. Howard, A. L. Ferguson and A. Z. Panagiotopoulos, *Soft Matter*, 2017, **13**, 4733–4745.
- 7 C. Martin, L. E. Altman, S. Rawat, A. Wang, D. G. Grier and V. N. Manoharan, *Nat. Rev. Methods Primers*, 2022, **2**, 1–17.
- 8 M. D. Hannel, A. Abdulali, M. O'Brien and D. G. Grier, *Opt. Express*, 2018, **26**, 15221–15231.
- 9 B. Midtvedt, E. Olsèn, F. Eklund, F. Höök, C. B. Adiels, G. Volpe and D. Midtvedt, *ACS nano*, 2021, **15**, 2240–2250.
- 10 A. Yevick, M. Hannel and D. G. Grier, *Opt. Express*, 2014, **22**, 26884–26890.
- 11 S.-H. Lee, Y. Roichman, G.-R. Yi, S.-H. Kim, S.-M. Yang, A. Van Blaaderen, P. Van Oostrum and D. G. Grier, *Opt. Express*, 2007, **15**, 18275–18282.
- 12 P. Memmolo, L. Miccio, M. Paturzo, G. Di Caprio, G. Coppola, P. A. Netti and P. Ferraro, *Adv. Opt. Photonics*, 2015, **7**, 713–755.
- 13 Z. Wang, L. Miccio, S. Coppola, V. Bianco, P. Memmolo, V. Tkachenko, V. Ferraro, E. Di Maio, P. L. Maffettone and P. Ferraro, *Light Adv. Manuf.*, 2022, **3**, 151–176.
- 14 B. J. Krishnatreya, A. Colen-Landy, P. Hasebe, B. A. Bell, J. R. Jones, A. Sunda-Meya and D. G. Grier, *Am. J. Phys.*, 2014, **82**, 23–31.
- 15 J. Sheng, E. Malkiel and J. Katz, *Exp. Fluids*, 2008, **45**, 1023–1035.
- 16 C. Wang, X. Zhong, D. B. Ruffner, A. Stutt, L. A. Philips, M. D. Ward and D. G. Grier, *J. Pharm. Sci.*, 2016, **105**, 1074–1085.
- 17 P. N. Kasimbeg, F. C. Cheong, D. B. Ruffner, J. M. Blusewicz and L. A. Philips, *J. Pharm. Sci.*, 2019, **108**, 155–161.
- 18 A. Winters, F. C. Cheong, M. A. Odete, J. Lumer, D. B. Ruffner, K. I. Mishra, D. G. Grier and L. A. Philips, *J. Pharm. Sci.*, 2020, **109**, 2405–2412.
- 19 H. Rahn, M. Oeztuerk, N. Hentze, F. Junge and M. Hollmann, *J. Pharm. Sci.*, 2022, in press.
- 20 K. Snyder, R. Quddus, A. D. Hollingsworth, K. Kirshenbaum and D. G. Grier, *Soft Matter*, 2020, **16**, 10180–10186.
- 21 L. E. Altman and D. G. Grier, *Biomed. Opt. Express*, 2020, **11**, 5225–5236.
- 22 M. Bertero, P. Boccacci and C. De Mol, *Introduction to Inverse Problems in Imaging*, CRC press, 2021.
- 23 C. F. Bohren and D. R. Huffman, *Absorption and Scattering of Light by Small Particles*, Wiley Interscience, New York, 1983.
- 24 M. I. Mishchenko, L. D. Travis and A. A. Lacis, *Scattering, Absorption and Emission of Light by Small Particles*, Cambridge University Press, Cambridge, 2001.

- 25 G. Gouesbet and G. Gréhan, *Generalized Lorenz-Mie Theories*, Springer-Verlag, Berlin, 2011.
- 26 S. Shao, K. Mallery, S. S. Kumar and J. Hong, *Opt. Express*, 2020, **28**, 2987–2999.
- 27 B. J. Krishnatreya and D. G. Grier, *Opt. Express*, 2014, **22**, 12773–12778.
- 28 Y. Zagzag, M. F. Soddu, A. D. Hollingsworth and D. G. Grier, *Sci. Rep.*, 2020, **10**, 1932.
- 29 L. E. Altman, R. Quddus, F. C. Cheong and D. G. Grier, *Soft Matter*, 2021, **17**, 2695–2703.
- 30 J. Sheng, E. Malkiel and J. Katz, *Appl. Opt.*, 2006, **45**, 3893–3901.
- 31 B. Leahy, R. Alexander, C. Martin, S. Barkley and V. N. Manoharan, *Opt. Express*, 2020, **28**, 1061–1075.
- 32 F. C. Cheong, B. Sun, R. Dreyfus, J. Amato-Grill, K. Xiao, L. Dixon and D. G. Grier, *Opt. Express*, 2009, **17**, 13071–13079.
- 33 R. Parthasarathy, *Nat. Methods*, 2012, **9**, 724–726.
- 34 A. D. Kashkanova, A. B. Shkarin, R. G. Mahmoodabadi, M. Blessing, Y. Tuna, A. Gemeinhardt and V. Sandoghdar, *Opt. Express*, 2021, **29**, 11070–11083.
- 35 J. C. Crocker and D. G. Grier, *J. Colloid Interface Sci.*, 1996, **179**, 298–310.
- 36 T. G. Dimiduk, R. W. Perry, J. Fung and V. N. Manoharan, *Appl. Opt.*, 2014, **53**, G177–G183.
- 37 D. B. Ruffner, F. C. Cheong, J. M. Blusewicz and L. A. Philips, *Opt. Express*, 2018, **26**, 13239–13251.
- 38 L. Denis, C. Fournier, T. Fournel, C. Ducottet and D. Jeulin, *Appl. Opt.*, 2006, **45**, 944–952.
- 39 J. Öhman and M. Sjö Dahl, *Appl. Opt.*, 2016, **55**, 7503–7510.
- 40 B. J. Thompson, *J. Phys. E: Sci. Instru.*, 1974, **7**, 781–788.
- 41 S. Soontaranon, J. Widjaja and T. Asakura, *Opt. Commun.*, 2004, **240**, 253–260.
- 42 S. Barkley, T. G. Dimiduk, J. Fung, D. M. Kaz, V. N. Manoharan, R. J. McGorty, R. W. Perry and A. Wang, *Comput. Sci. Eng.*, 2020, **22**, 72–82.
- 43 F. C. Cheong, K. Xiao, D. J. Pine and D. G. Grier, *Soft Matter*, 2011, **7**, 6816–6819.
- 44 V. Markel, *J. Opt. Soc. Am. A*, 2016, **33**, 1244–1256.
- 45 C. Wang, F. C. Cheong, D. B. Ruffner, X. Zhong, M. D. Ward and D. G. Grier, *Soft Matter*, 2016, **12**, 8774–8780.
- 46 J. Fung and S. Hoang, *J. Quant. Spectr. Rad. Trans.*, 2019, **236**, 106591.
- 47 M. A. Odete, F. C. Cheong, A. Winters, J. J. Elliott, L. A. Philips and D. G. Grier, *Soft Matter*, 2020, **16**, 891–898.
- 48 A. Ray, M. A. Khalid, A. Demčenko, M. Daloglu, D. Tseng, J. Reboud, J. M. Cooper and A. Ozcan, *Nature Commun.*, 2020, **11**, 1–10.
- 49 G. Ding, J. Wang, J. Zou, C. Wang, T. Wang, J. Meng and D. Li, *IEEE Sensors J.*, 2020, **20**, 6928–6936.
- 50 R. Alexander, B. Leahy and V. N. Manoharan, *J. Appl. Phys.*, 2020, **128**, 060902.
- 51 A. Wang, W. B. Rogers and V. N. Manoharan, *Phys. Rev. Lett.*, 2017, **119**, 108004.
- 52 A. Wang, T. G. Dimiduk, J. Fung, S. Razavi, I. Kretzschmar, K. Chaudhary and V. N. Manoharan, *J. Quant. Spectr. Rad. Trans.*, 2014, **146**, 499–509.
- 53 M. I. Mishchenko, L. D. Travis and D. W. Mackowski, *J. Quant. Spectr. Rad. Trans.*, 1996, **55**, 535–575.
- 54 D. W. Mackowski and M. I. Mishchenko, *J. Opt. Soc. Am. A*, 1996, **13**, 2266–2278.
- 55 Y. Zhang, Y. Zhu and E. Y. Lam, *Appl. Opt.*, 2022, **61**, B111–B120.
- 56 J. Redmon, S. Divvala, R. Girshick and A. Farhadi, *Proc. IEEE Comp. Vision Pattern Recog.*, 2016, pp. 779–788.
- 57 J. Redmon and A. Farhadi, *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 7263–7271.
- 58 J. Redmon and A. Farhadi, *CoRR*, 2018, **abs/1804.02767**, 1–6.
- 59 G. Jocher, A. Chaurasia, A. Stoken, J. Borovec, NanoCode012, Y. Kwon, TaoXie, J. Fang, imyhxy, K. Michael, Lorna, A. V. D. Montes, J. Nadar, Laughing, tkianai, yxNONG, P. Skalski, Z. Wang, A. Hogan, C. Fati, L. Mammana, AlexWang1900, D. Patel, D. Yiwei, F. You, J. Hajek, L. Diaconu and M. T. Minh, *ultralytics/yolov5: v6.1 - TensorRT, TensorFlow Edge TPU and OpenVINO Export and Inference*, 2022, <https://doi.org/10.5281/zenodo.6222936>.
- 60 M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu and X. Zheng, *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*, 2015, <https://www.tensorflow.org/>, Software available from tensorflow.org.
- 61 A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai and S. Chintala, *Advances in Neural Information Processing Systems 32*, Curran Associates, Inc., 2019, pp. 8024–8035.
- 62 X. Glorot, A. Bordes and Y. Bengio, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, Fort Lauderdale, FL, USA, 2011, pp. 315–323.
- 63 T. Kurbiel and S. Khaleghian, *Training of Deep Neural Networks based on Distance Measures using RMSProp*, 2017.
- 64 R. Girshick, *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 1440–1448.
- 65 H.-C. Shin, H. R. Roth, M. Gao, L. Lu, Z. Xu, I. Nogues, J. Yao, D. Mollura and R. M. Summers, *IEEE Trans. Med. Imag.*, 2016, **35**, 1285–1298.
- 66 B. Midtvedt, S. Helgadottir, A. Argun, J. Pineda, D. Midtvedt and G. Volpe, *Appl. Phys. Rev.*, 2021, **8**, 011310.
- 67 B. Midtvedt, J. Pineda, F. Skärberg, E. Olsén, H. Bachimanchi, E. Wesén, E. K. Esbjörner, E. Selander, F. Höök, D. Midtvedt

- et al.*, *Nat. Commun.*, 2022, **13**, 7492.
- 68 K. Xiao and D. G. Grier, *Phys. Rev. Lett.*, 2010, **104**, 028302.
- 69 E. R. Dufresne, D. Altman and D. G. Grier, *Europhys. Lett.*, 2001, **53**, 264–270.
- 70 D. G. Grier, *Nature*, 2003, **424**, 810–816.
- 71 M. J. O'Brien and D. G. Grier, *Opt. Express*, 2019, **27**, 24866–25435.
- 72 L. E. Altman and D. G. Grier, *Phys. Rev. E*, 2022, **106**, 044605.
- 73 C. Wang, H. Shpaisman, A. D. Hollingsworth and D. G. Grier, *Soft Matter*, 2015, **11**, 1062–1066.
- 74 C. Wang, H. W. Moyses and D. G. Grier, *Appl. Phys. Lett.*, 2015, **107**, 051903.
- 75 F. C. Cheong, P. Kasimbeg, D. B. Ruffner, E. H. Hlaing, J. M. Blusewicz, L. A. Philips and D. G. Grier, *Appl. Phys. Lett.*, 2017, **111**, 153702.
- 76 L. A. Philips, D. B. Ruffner, F. C. Cheong, J. M. Blusewicz, P. Kasimbeg, B. Waisi, J. R. McCutcheon and D. G. Grier, *Water Res.*, 2017, **122**, 431–439.
- 77 M. Cranmer, A. Sanchez Gonzalez, P. Battaglia, R. Xu, K. Cranmer, D. Spergel and S. Ho, *Adv. Neural Inf. Process. Syst.*, 2020, **33**, 17429–17442.