3D Localization of Microparticles from Holographic Images using Neural Networks

Ayush Paliwal
MPI for Dynamics and Self-Organization
Göttingen, Germany
ayush.paliwal@ds.mpg.de

Oliver Schlenczek
MPI for Dynamics and Self-Organization
Göttingen, Germany

Birte Thiede
MPI for Dynamics and Self-Organization
Göttingen, Germany

Gholamhossein Bagheri
MPI for Dynamics and Self-Organization
Göttingen, Germany

Alexander Ecker
Institute of Computer Science and Campus Institute Data Science, University of Göttingen
MPI for Dynamics and Self-Organization
Göttingen, Germany
ecker@cs.uni-goettingen.de

Abstract

Digital in-line holography is a versatile and reliable imaging technique for characterizing the size and spatial distribution of particles in flows. It maps 3D particle snapshots into a 2D image – a hologram. The analysis bottleneck of in-line holographic imaging is the immense cost of computation and manual effort that must be invested. Here we propose a learning-based approach to extract the size and 3D spatial distribution of objects from holograms. We outperform the standard propagation based method in terms of detection rate, precision and processing time on a small sensor aperture and at $1/4$ of the original resolution. Our method performs with an F1 score of 0.91 in suspensions with more than 60 particles$/\mathrm{cm}^3$, which is an 18 percent performance boost in comparison to propagation-based software commonly used in practice. At the same time, our method is also several orders of magnitude faster, eliminating the computational bottleneck.

1 Introduction

Holography is an imaging technique that captures 3D information of an object on a 2D image. In everyday photography, only the intensity variations are recorded, while the phase, which encodes the 3D character of the scene, is lost [Ghatak, 2008]. In holography the phase information can be reconstructed using the principles of diffraction of coherent light.

Holography has been used for many years in atmospheric sciences to study particle flows [Thompson, 1974], which is crucial for understanding the Earth's climate and hydrological cycle [Bodenschatz et al., 2010] and the effects of anthropogenic changes in aerosols on clouds [Toll et al., 2019, Rosenfeld et al., 2019]. The unique feature of holography is that it provides a non-intrusive, instantaneous measurement of particle size distribution in a large and representative sample volume, which is difficult to achieve with other imaging techniques. However, the analysis of hologram images remains a tedious task that is fraught with significant difficulties. The analysis of a hologram takes about

Machine Learning and the Physical Sciences Workshop, NeurIPS 2023.

five orders of magnitude longer than its acquisition, so the method cannot be efficiently applied to a large-scale problem.

Our goal in this paper is to work towards making digital in-line holography accurate, affordable and real-time efficient. In particular, we focus on settings with a large sample depth and small sensor size (Fig. 1). We developed a learning based method to extract the coordinates and sizes of all the particles in a hologram. On synthetic holograms with added camera read noise and quantization, we outperform the classical propagation approach[Fugal et al., 2009] on a small sensor/camera aperture and at $1/4$th of the original resolution. In particular, we show that we can retrieve the 3d position and size of particles beyond the diffraction-based resolution limit of classical reconstruction-based methods.

2 Methods for holography reconstruction

Reconstruction based on FFT-propagation The basic setup for in-line holography is illustrated in Fig. 1. Based on scalar diffraction theory, the hologram at distance z_0 from camera can be propagated to other distances by convolving it with a Huygens-Fresnel kernel h_{z-z_0},

$$I(x, y, z) = (h_{z-z_0} * I)(x, y, z_o), \quad (x, y, z) \in \mathbb{R}^3. \tag{1}$$

For sampling reasons [Fugal et al., 2009], the above is done in Fourier space using the Huygens-Fresnel filtering kernel [Goodman, 2005].

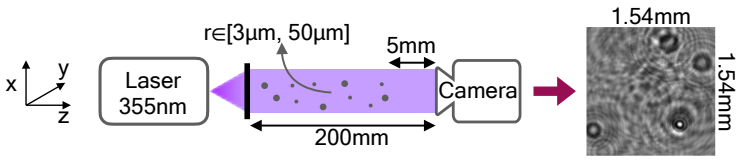

Figure 1: Digital in-line holography setup with large sample depth and small sensor area as we use it in this paper.

The operation is repeated for each z required to cover the entire depth of the volume. Typically, in cloud holography, one desires $100\mu\text{m}$ depth resolution, where one has to then reconstruct thousands of z planes, thereby leading to a computational bottleneck.

After reconstruction, further image processing is needed. One common approach is to do intensity thresholding and other filtering methods described in Fugal et al. [2009], Bagheri et al. [2023] to extract object parameters, followed by particle–artefact classification using a small amount of human-supervised training data, denoted as 'Std. method' in (Fig. 4). Alternatively, one can apply criteria for particle shape [Schreck et al., 2022], but particles close to the each other or to the edges of the detector may be missed as they appear distorted in shape. At best, this method takes 30 minutes$/25$M pixels to output particle parameters.

There are methods which take the inverse problem approach [Soulez et al., 2007, Seifi et al., 2012, Fournier et al., 2011]. Here, convex optimization is performed directly on the parameter space (x_j, y_j, z_j, r_j). Overall, above methods are accurate, but not suitable for processing voluminous data because they are slow.

ML based approaches. [Shimobaba et al., 2019] showed that such problem can easily be mapped to the deep learning framework directly on the particle parameter space. Later, efforts by [Shao et al., 2020, Chen et al., 2021, Zhang et al., 2022] have been made to improve on detection and precision. These approaches have been evaluated only on holograms with small sample depths $\sim 2\text{cm}$, which have relatively localized interference patterns that are straightforward to detect. Whether they generalize to larger sample depth on the order of 20cm is unclear. [Schreck et al., 2022] describe a hybrid approach using classical reconstruction followed by a neural network for particle detection. Their method has a lot of potential, but does not avoid the computational bottleneck of reconstruction.

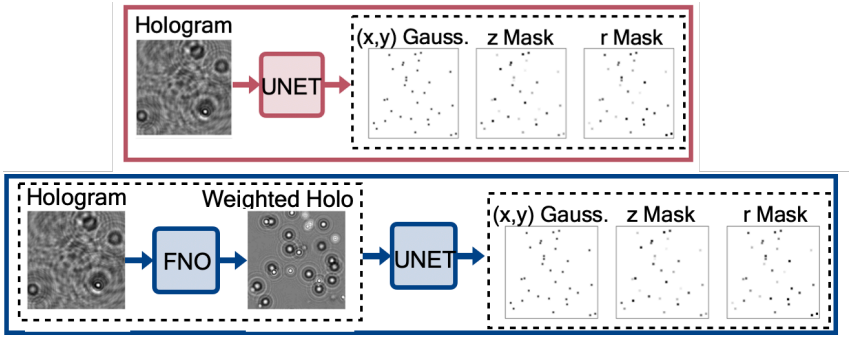

Figure 2: Illustration of our two approaches. The output is the same in both approaches. **B**: We use a pre-processing step to localise fringes and enhance contrast using an FNO-based neural network.

3 Our approach

The two main challenges our approach addresses are (1) avoiding the computational cost of explicit 3d reconstruction or iterative optimization and (2) extracting particle locations out of highly non-local interference patterns with strong contrast variations across particles as they arise in holograms from setups with large sample depth (like Fig. 1).

To address the first challenge, we circumvent explicitly reconstructing the 3d volume and treat the problem as an (x, y) localization problem combined with regression of the z coordinate and size of each particle. For this, we employ the U-Net architecture [Ronneberger et al., 2015] combined with the "objects as points" approach to object detection [Zhou et al., 2019]. The U-Net takes as input a hologram and outputs three feature maps (Fig. 2A): (1) a heatmap that predicts the probability of a particle being at any (x, y) location; (2) the z coordinate of the particle from the sensor; (3) the radius r of the particle. From the first heatmap, we extract the particle locations by straightforward detection of local maxima. The latter two feature maps then indicate z and r of the particle. They are defined only at particle locations.

The second challenge is the nature of holograms when imaging a large sample depth: particles far away from the sensor give rise to highly non-local interference patterns and variation in particle size results in highly variable contrast of the particles' fringe patterns. We found these properties to result in less salient particles often being missed. To address this challenge, we devised a learned pre-processing step that redistributes the contrast and makes the fringe pattern of all particles more localized by predicting what the hologram would look like if all particles in the 3d volume had the same size and were imaged with a much lower (1/8) wavelength (see Fig. 2B "Weighted Holo"). For this pre-processing step, we use another image-to-image network, which is based on the Fourier Neural Operator (FNO) architecture [Kovachki et al., 2021, Li et al., 2020]. We found the FNO architecture to be more suitable for this intermediate step than a regular U-Net, presumably due to its operating in Fourier space, which gives rise to global receptive fields in every layer. After this pre-processing step, the resulting modified (non-physical) hologram, along with the original hologram, is fed into the U-Net described above to detect particles and regress their size and radius (Fig. 2B).

4 Experiments

Datasets. We used the Fraunhofer diffraction formula [Goodman, 2005] to calculate the expected photon count at the sensor, added shot noise by sampling from a Poisson distribution and added Gaussian noise to simulate the sensor read noise. The LASER wavelength was $355\,\mathrm{nm}$. Synthetic holograms had 512×512 pixel resolution with $3\,\mu\mathrm{m}$ pixel pitch. For training (99%) and validation (1%) we downsampled these holograms to 128×128.

We created two datasets. The first consisted of 50,000 holograms that contained two particles of size $6\,\mu\mathrm{m}$ and $100\,\mu\mathrm{m}$, respectively, placed at $200\,\mathrm{mm}$ from the camera at a random (x_j, y_j) location. We refer to this dataset as the *toy dataset*, as it presents a minimal example of hard cases that demonstrate the limitation of diffraction-limited propagation methods and the success of neural networks (Fig. 3).

The second, our main dataset, consisted of 900,000 holograms. In this dataset the particles' coordinates and size were drawn uniformly ($z \in [5\,\mathrm{mm}, 200\,\mathrm{mm}]$; $r \in [6\,\mu\mathrm{m}, 100\,\mu\mathrm{m}]$). The number of particles in a hologram were drawn from $[1, 2, 4, 8, 16, 32]$ for 150,000 holograms each.

Our test dataset consisted of 2,000 holograms. Here, location and size were again drawn uniformly, but all particle counts between 1 and 32 were possible, drawn according to the following distribution, which roughly resembles the distribution at training time.

$$p(N) = \begin{cases} \frac{1}{4} & 1 \leq N < 5 \\ 2^{-\frac{N}{4}} & 5 \leq N \leq 32 \end{cases}, \quad \text{for } N \text{ particles.} \tag{2}$$

Architecture details. We use the standard U-Net architecture [Ronneberger et al., 2015], with the exception that we use six blocks in both encoder and decoder, and a sigmoid activation just before the output. For the z and r feature maps, the output is scaled between $[0, 200]$ and $[0, 50]$, respectively.

In the FNO network, we concatenate the input hologram with a positional encoding (consisting of two channels coding the x and y coordinate, respectively) and lift the result to 512 channels using 1×1 convolutions. The result is passed through five FNO blocks [Kovachki et al., 2021, Li et al., 2020]. Each FNO block has 128 Fourier filters and GeLU activation function, and operates only on the 64×64 lowest spatial frequency modes. Finally, the 512 channels are projected back to one channel, which predicts the weighted hologram. This weighted hologram along with the original hologram is then fed as a two-channel input to a U-Net with the same architecture and outputs as above. Batch normalization is used after every convolution-activation pair for both the models.

Training. We use the Adam optimizer with a learning rate of 5×10^{-4}, which is reduced by a factor of 0.9 with patience of 10 epochs. We use the following loss functions: a weighted combination of mean of squared errors (weight$= 0.9999$) with total variation (TV) loss (weight $= 0.0001$) for the heatmap, Huber loss ($\delta = 0.001$) for the z and r regression [Shao et al., 2020] and an MSE loss for the weighted holograms output by the FNO network. The overall loss being optimized was the sum of all individual loss functions.

Evaluation. For evaluation, we first determine all correctly detected, missed and falsely detected particles as follows: Any detection in a 3×3 pixel area and within $\pm 8\mathrm{mm}$ in z of the ground truth was counted as a match. We report the performance in terms of F1 score, recall and mean absolute error (MAE) in z and r.

5 Results

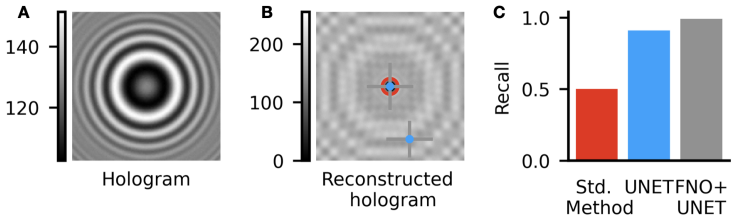

Figure 3: **A:** Two particles (3 µm and 50 µm radius), placed at $z = 200\,\mathrm{mm}$. **B:** Propagation-based reconstruction at $z = 200\,\mathrm{mm}$ only detects $50\mu\mathrm{m}$ particle (red circle), but neural nets detect both (blue dot, grey cross). **C:** Accuracy of the three approaches.

We start with an illustrative toy example of just two particles (Fig. 3) that highlights the limitations of the propagation-based approach and the challenges associated with detecting small, distant particles. While the propagation-based reconstruction does not detect the small particle (Fig. 3B), the neural networks reconstruct such cases almost perfectly (Fig. 3B, C).

The superior performance of our neural network based approach also holds in the more general case (Fig. 4). We compared it to the standard method described by [Fugal et al., 2009] as well as an oracle based on reconstruction followed by thresholding. The oracle was given all z slices that

contained a particle and determined only the (x, y) coordinate based on intensity thresholding of the reconstruction (it represents an upper bound to what can be achieved with thresholding). The neural networks outperformed the classical methods with respect to both recall and precision (Fig. 4A). The improvements were evident across all distances[1] from the sensor (Fig. 4D), particle sizes (Fig. 4E) and particle densities (Fig. 4F) tested. Regressing the z location worked with reasonable accuracy, but with somewhat higher error than the standard method (Fig. 4B). In terms of size estimation, our approach outperformed the standard method considerably (Fig. 4C). The 10-run average processing time to extract particle parameters is 0.1 seconds$/25$M pixels which is 18000 times faster than the standard method.

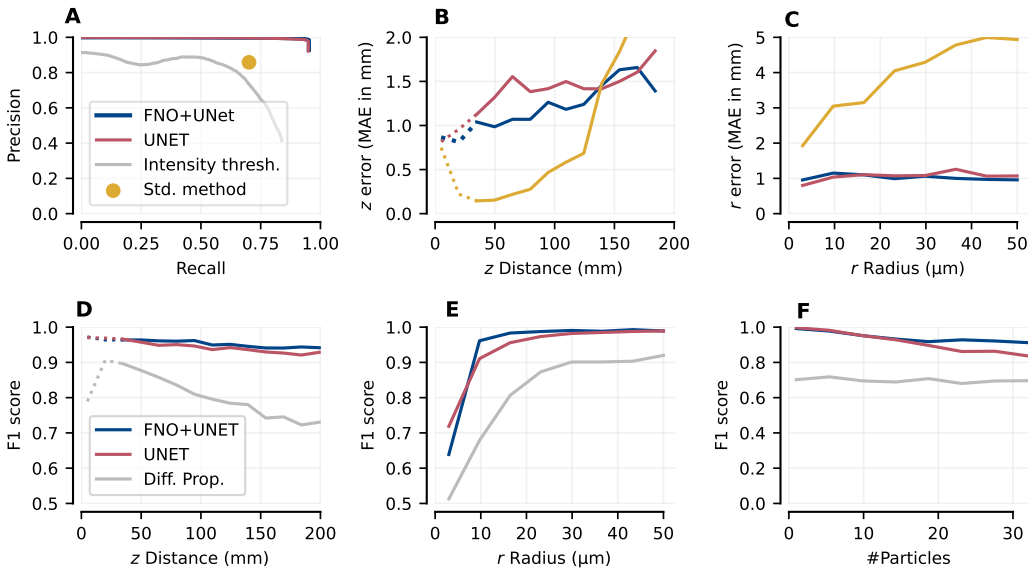

Figure 4: **A:** Precision over recall for the test dataset. **B,C:** Comparison for Mean Average Error (MAE) in z and r with respect to z (B) and radius r (C) of particles. **D–F:** Accuracy (F1 score) with respect to z (D), r (E) and number of particles (F).

6 Discussion

Our results show clear improvements in detecting particles in large volumes with limited sensor size while at the same time improving inference speed. Estimating the z localization could still be improved. It is noteworthy that our neural networks are currently working on $1/4$ of the original resolution used for propagation-based reconstruction. Analysis with half- and full- resolution holographic images, physical models to include the near-field as well as scaling the approach to higher-resolution sensors often used in practice (>25 megapixels) is work in progress.

References

A. Ghatak. *Optics*. Tata McGraw-Hill, 2008. ISBN 9780071336659. URL `https://books.google.de/books?id=psTXtwEACAAJ`.

B. J. Thompson. Holographic particle sizing techniques. *Journal of Physics E: Scientific Instruments*, 7:781, 1974.

[1]The far-field approximation sets a bound that $z \ggg \frac{r^2}{4\lambda} \sim 28\text{mm}$, for our setup. For near-field (dotted lines in Fig. 4) understanding, one shouldn't trust above results.

E. Bodenschatz, S. P. Malinowski, R. A. Shaw, and F. Stratmann. Can we understand clouds without turbulence? *Science*, 327(5968):970–971, 2010. doi: 10.1126/science.1185138. URL https://www.science.org/doi/abs/10.1126/science.1185138.

V. Toll, M. Christensen, and J. et al. Quaas. Weak average liquid-cloud-water response to anthropogenic aerosols. *Nature*, 572:51–55, 2019. doi: 10.1038/s41586-019-1423-9.

Daniel Rosenfeld, Yannian Zhu, Minghuai Wang, Youtong Zheng, Tom Goren, and Shaocai Yu. Aerosol-driven droplet concentrations dominate coverage and water of oceanic low-level clouds. *Science*, 363(6427):eaav0566, 2019. doi: 10.1126/science.aav0566. URL https://www.science.org/doi/abs/10.1126/science.aav0566.

Jacob P Fugal, Timothy J Schulz, and Raymond A Shaw. Practical methods for automated reconstruction and characterization of particles in digital in-line holograms. 20(7):075501, 2009. ISSN 0957-0233, 1361-6501. doi: 10.1088/0957-0233/20/7/075501. URL https://iopscience.iop.org/article/10.1088/0957-0233/20/7/075501.

J.W. Goodman. *Introduction to Fourier Optics*. McGraw-Hill physical and quantum electronics series. W. H. Freeman, 2005. ISBN 9780974707723. URL https://books.google.de/books?id=ow5xs_Rtt9AC.

Gholamhossein Bagheri, Oliver Schlenczek, Laura Turco, Birte Thiede, Katja Stieger, Jana M. Kosub, Sigrid Clauberg, Mira L. Pöhlker, Christopher Pöhlker, Jan Moláček, Simone Scheithauer, and Eberhard Bodenschatz. Size, concentration, and origin of human exhaled particles and their dependence on human factors with implications on infection transmission. *Journal of Aerosol Science*, 168:106102, 2023. ISSN 0021-8502. doi: https://doi.org/10.1016/j.jaerosci.2022.106102. URL https://www.sciencedirect.com/science/article/pii/S0021850222001380.

J. S. Schreck, G. Gantos, M. Hayman, A. Bansemer, and D. J. Gagne. Neural network processing of holographic images. *Atmos. Meas. Tech.*, 15:5793–5819, 2022. doi: https://doi.org/10.5194/amt-15-5793-2022,2022.

Ferréol Soulez, Loïc Denis, Corinne Fournier, Éric Thiébaut, and Charles Goepfert. Inverse-problem approach for particle digital holography: accurate location based on local optimization. *J. Opt. Soc. Am. A*, 24(4):1164–1171, Apr 2007. doi: 10.1364/JOSAA.24.001164. URL https://opg.optica.org/josaa/abstract.cfm?URI=josaa-24-4-1164.

Mozhdeh Seifi, Corinne Fournier, Loic Denis, Delphine Chareyron, and Jean-Louis Marié. Three-dimensional reconstruction of particle holograms: a fast and accurate multiscale approach. 29 (9):1808, 2012. ISSN 1084-7529, 1520-8532. doi: 10.1364/JOSAA.29.001808. URL https://opg.optica.org/abstract.cfm?URI=josaa-29-9-1808.

Corinne Fournier, Loic Denis, Eric Thiebaut, Thierry Fournel, Mozhdeh Seifi, and Laboratoire Hubert Curien. Inverse problems approaches for digital hologram reconstruction. 2011.

Tomoyoshi Shimobaba, Takayuki Takahashi, Yota Yamamoto, Yutaka Endo, Atsushi Shiraki, Takashi Nishitsuji, Naoto Hoshikawa, Takashi Kakue, and Tomoyosh Ito. Digital holographic particle volume reconstruction using a deep neural network. *Appl. Opt.*, 58(8):1900–1906, Mar 2019. doi: 10.1364/AO.58.001900. URL https://opg.optica.org/ao/abstract.cfm?URI=ao-58-8-1900.

Siyao Shao, Kevin Mallery, S. Santosh Kumar, and Jiarong Hong. Machine learning holography for 3d particle field imaging. *Opt. Express*, 28(3):2987–2999, Feb 2020. doi: 10.1364/OE.379480. URL https://opg.optica.org/oe/abstract.cfm?URI=oe-28-3-2987.

Ni Chen, Congli Wang, and Wolfgang Heidrich. Holographic 3d particle imaging with model-based deep network. *IEEE Transactions on Computational Imaging*, 7:288–296, 2021. doi: 10.1109/TCI.2021.3063870.

Yunping Zhang, Yanmin Zhu, and Edmund Y. Lam. Holographic 3d particle reconstruction using a one-stage network. *Appl. Opt.*, 61(5):B111–B120, Feb 2022. doi: 10.1364/AO.444856. URL https://opg.optica.org/ao/abstract.cfm?URI=ao-61-5-B111.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597, 2015. URL `http://arxiv.org/abs/1505.04597`.

Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019.

Nikola B. Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces. *CoRR*, abs/2108.08481, 2021.

Zongyi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *CoRR*, abs/2010.08895, 2020. URL `https://arxiv.org/abs/2010.08895`.