

Fast particle characterization using digital holography and neural networks

B. SCHNEIDER,^{1,2,*} J. DAMBRE,³ AND P. BIENSTMAN^{1,2}

¹Photonics Research Group (INTEC), Ghent University–IMEC, Sint-Pietersnieuwstraat 41, B-9000 Ghent, Belgium

²Center for Nano- and Biophotonics (NB-Photonics), Ghent University, Sint-Pietersnieuwstraat 41, 9000 Ghent, Belgium

³Ghent University, Department of Electronics and Information Systems, Sint-Pietersnieuwstraat 41, 9000 Ghent, Belgium

*Corresponding author: bendix.schneider@intec.ugent.be

Received 30 September 2015; revised 26 November 2015; accepted 29 November 2015; posted 30 November 2015 (Doc. ID 251175); published 23 December 2015

We propose using a neural network approach in conjunction with digital holographic microscopy in order to rapidly determine relevant parameters such as the core and shell diameter of coated, non-absorbing spheres. We do so without requiring a time-consuming reconstruction of the cell image. In contrast to previous approaches, we are able to obtain a continuous value for parameters such as size, as opposed to binning into a discrete number of categories. Also, we are able to separately determine both core and shell diameter. For simulated particle sizes ranging between 7 and 20 μm , we obtain accuracies of $(4.4 \pm 0.2)\%$ and $(0.74 \pm 0.01)\%$ for the core and shell diameter, respectively. © 2015 Optical Society of America

OCIS codes: (090.1995) Digital holography; (200.4260) Neural networks; (290.5850) Scattering, particles; (100.3190) Inverse problems; (100.2960) Image analysis.

<http://dx.doi.org/10.1364/AO.55.000133>

1. INTRODUCTION

Characterizing spherically shaped particles is of great importance when studying their related scattering properties. Sprays, aerosols, and colloidal suspensions, e.g., paint, reflect light depending on their particle size, density, and homogeneity. Coatings can be used to modify particle properties, such as color, adsorption, and flow properties, in a designed way, leading to numerous industrial, military, and biomedical applications [1–4]. For the latter, stability and composition of the functionalized surfaces is of the utmost importance. These examples show that a measurement of coated layer thickness is an essential feature during process verification. In this paper we focus on a more biologically inspired application of particle characterization: low refractive index particles in water solution (relative $\Delta_n = 2\text{--}4\%$) with sizes ranging between 7 and 20 μm . These values are chosen similar to what is found for white blood cells (WBCs) [5] and thus form a simplified model for real-life WBCs encountered in flow cytometers. The results achieved in this paper serve as a proof-of-concept for characterization of more realistic cell images such as circulating tumor cells that typically have a larger diameter than healthy cells [6].

Common optical techniques that are used in particle sizing are based on light scattering, velocimetry, and microscopy. In this paper we use digital holography for particle sizing. It allows for single-shot characterization and can be extended to particle tracking when recording a suite of images. Extracting particle

[7] and distribution characteristics [8] from holograms, bright field images [9], or Fraunhofer diffraction patterns [10] has already been studied in the past and is generally solved by applying different numerical algorithms involving inversion, nonlinear pattern matching, or performing image analysis decomposition. Since integrals of special functions or an extensive use of FFTs occur in most of these algorithms, they all suffer from a tremendous increase in computational cost if the image size increases or several particles are studied in parallel. Therefore severe limitations exist on real-time particle characterization for those image processing algorithms in high-throughput applications.

In this work we consider the use of artificial neural networks (ANNs) in order to cope with the requirement of fast particle characterization. Neural networks are a powerful tool, well-known in the field of machine learning. They have been successfully applied to particle shape classification in the past. What is different here is that we do not feed a (reconstructed) image to an ANN but rather directly use the holographic interference pattern as input. This avoids the considerable computational cost related to the reconstruction of the image from the interference pattern. Also, ANNs that are presented with the raw microscopy image as input features only perform matrix multiplications and standard function evaluations (tanh) and can easily be optimized for parallel computation and accelerated Graphics Processing Unit arrays. In fact, the main contribution

to the computational complexity stems from the matrix multiplication of the input vector with the input-to-hidden weights and scales as $\mathcal{O}(2N_{\text{pix}}N_{\text{HLU}})$. For a 10 GFlops/s processing unit it takes less than 1 μs to perform the prediction step of a neural network composed of $N_{\text{HLU}} = 10$ hidden units and about $N_{\text{pix}} = 1000$ inputs (pixel values). Indeed, we measured a processing time of 1 μs per prediction on a 2.8 GHz Intel Xeon X5560 CPU. We therefore observe a significant increase in image processing speed. The only drawback is the time necessary for training the ANN upon calibration, which can take seconds up to minutes depending on the network architecture, size, and training algorithm. However, this needs to happen only once and can be done offline. Only single isolated particle holograms are considered at the current stage of our work. An extension to imaging and characterizing multiple particles together is possible but not easily implemented. Although this might be a severe limitation in some fields of particle holography, it is not the case for flow experiments in microfluidic channels if the particles are well aligned in the center. For instance the experimental work on lab-on-a-chip, label-free cell classification [11] enables high-throughput operation since hundreds or thousands of microfluidic channels are used in parallel. In a stable flow each cell is isolated from the others as it passes across an illuminating pinhole. Additionally the microfluidic confinement leads to an improved control and a reduced uncertainty for the optical depth. Our research was highly motivated by these experiments because its high-speed, on-line image classification requirements are without doubt. Indeed we proved recently [12] that our method successfully applies to the classification of experimentally recorded inline holograms of white blood cells in a microfluidic channel. In that sense the method still allows for parallel computation of many spatially isolated particles but excludes currently the possibility of characterizing and tracking multiple particles that are close to each other.

A recent study reports on the direct application of support vector machines (SVMs) to digital holographic images [13] yielding a 1000-fold gain in computational speed as compared to nonlinear fitting techniques. Parameters relevant for particle tracking, such as refractive index, depth, and radius, are recovered. However, a disadvantage of that approach lies in the fact that SVMs are built from a discrete set of target vectors that span the underlying sample space (dictionary) and to which collected data is matched subsequently. If a fine resolution is needed the dictionary size increases substantially and requires the creation of an unrealistically big, expensive training set. This important problem is addressed and solved by the use of ANNs. Additionally they have the possibility to output continuous variables, which is not straightforward with SVMs.

To summarize, the main contribution of this paper is a method that can not only obtain continuous-valued estimates for particle size (as opposed to binning in discrete sets like in [13]), but also separately determine both core and shell radius in coated particles. Moreover, the method is computationally efficient, because it skips a time-consuming image reconstruction step and has an ANN directly operating on the holographic interference pattern.

The rest of this paper is structured as follows: Section 2 describes how digital inline holography in conjunction with

artificial neural networks can be successfully applied to extract shell and core diameter of coated spheres directly from the hologram record. In Section 3 the outcomes of the suggested method under different imaging conditions, such as change in optical depth z , numerical aperture (NA), or pixel pitch (Δ_{pix}), are investigated. Concluding remarks are presented in the final Section 4 of the paper.

2. PARTICLE CHARACTERIZATION USING DIGITAL HOLOGRAPHY AND ARTIFICIAL NEURAL NETWORKS

Digital inline holography (DIH) has widely been applied to various studies of microscopic objects, e.g., for flows around obstructions and tracking in microfluidics [14], microsphere imaging [15], and the characterization of optical components [16]. As a holographic approach, DIH relies on the fact that the superposition of a wave scattered off the object and a reference wave leads to a unique interference pattern on a recording medium. Inline holography uses a robust common path configuration in which the reference and the object wave are completely aligned (see Fig. 1).

Making the assumptions of an object wave, E_O , with magnitude much smaller than the reference, E_R , and constant background intensity that can be subtracted from the recorded hologram, one derives the well-known expression for the contrast image [14]:

$$\tilde{I} \propto |E_O + E_R|^2 - |E_R|^2 \approx E_O E_R^* + E_O^* E_R. \quad (1)$$

It contains two terms, which are conjugate to each other and are known as the image and twin image contribution. There exist a variety of integral transforms that allow the numerical reconstruction of the hologram, of which the inverse Fresnel transform is one of the most popular choices.

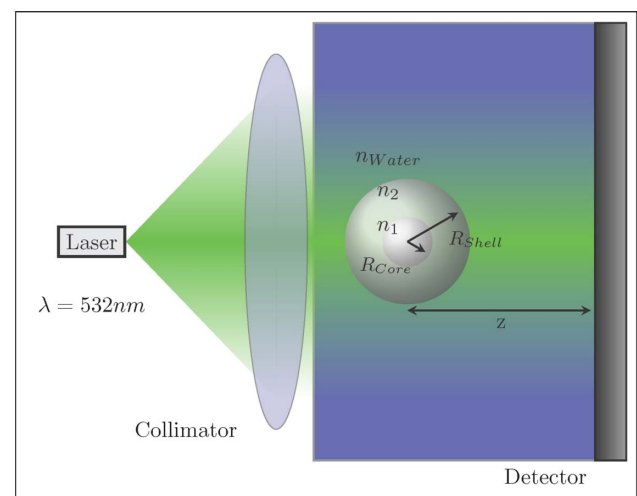


Fig. 1. Simulation setup: A green light laser beam is collimated by a lens before impinging on a coated particle suspended in water. The scattered part of the incident plane wave (described by Mie theory) interferes with its direct part on an image sensor, spaced at a distance z from the particle. The digitized version of the fringe patterns form the inline hologram.

In traditional holography, the twin term, if not suppressed by an appropriate filter, creates a defocused virtual image that spatially overlaps with the real image. This led to the numerous off-axis holography systems. In DIH microscopy, however, the twin image spreads out over the whole reconstructed image.

ANNs are machine learning techniques that are used to learn an arbitrary input-to-output mapping. Feed-forward neural networks are ANNs that are devoid of feedback. They are typically arranged in multiple layers: an input layer, one or several hidden layers, and an output layer. Each layer is solely connected to the following one and is composed of a varying number of neurons. These neurons are biologically inspired and function in a similar, although more abstract, way. Every single neuron receives weighted inputs from the neurons located in the preceding layer. The inputs are summed up (neuron activation) and a nonlinear (tanh) transfer function is applied. This forms the neuron output response that connects to the subsequent layer. At the output layer the network response is compared to the target signal. Any deviation results in an error signal, which is used to adjust the network weights so as to minimize the error metric, e.g., the sum of squared errors. Presenting a large number of input–output pairs to the network eventually results in a supervised training effect: the neural network “learns” the mapping by tuning its weights accordingly. The trained network is then able to generalize for unseen input.

As is the case with all supervised machine learning techniques, ANNs require a dedicated training set that grows with the number of unknown weights inside the net. In order to avoid a time-consuming training procedure, we limit ourselves to rather small-size, feed-forward neural nets solely composed of a single hidden layer. The prediction accuracy depends on the number of hidden units. Therefore a heuristic search for a good number of hidden units as a trade-off between accuracy and model complexity is necessary. In a similar way, the number of input units constitutes a compromise between model complexity and resolution capability of the device. We will denote a particular network architecture as (*number of input units, number of hidden units, number of output units*), e.g., (1024, 10, 1). For each parameter characterizing the particle, a distinct network is trained using a stochastic gradient descent backpropagation algorithm. The learning rate decays following a search-then-converge schedule [17]. L1 regularization of the weight norms prevents the neural network from overfitting [18]. We preferred L1 over the more common L2 regularization method because of its tendency to drive many network parameters (weights) to zero, leading to a sparse input vector representation. This is advantageous for later hardware implementations, in which zero or near-to-zero weights can be omitted and thus simplify the architecture. The neural net is trained with 10 different initial weight distributions so as to eliminate cases where the training algorithm is trapped in a local, non-optimal minimum.

The training, validation, and test sets consist of random partitions of a catalogue of diffraction patterns. We used rigorous Mie-scattering theory [19,20] to calculate the diffraction holograms of concentric spheres at various depths and with different radii under plane wave laser illumination (see Fig. 1).

Because of the translation and rotation invariance of the diffraction patterns in the detection plane, we assume the particle

to be located at the origin of the detector coordinate system and only record the radial dependence. In a practical setup this can be done by locating and shifting the first moment of the diffraction spot back to the origin. Figure 2 shows that in this case the hologram is a 1024×1 pixel line image recorded by a sensor (pixel pitch $\Delta_{\text{pix}} = 1.5 \mu\text{m}$). This allows for a considerable speed-up of the sensor frame rate in real-time applications. We stress the fact that the argument of symmetry only applies to the single particle studies, which means in low density solutions or microfluidic channel flows, e.g., the cell flow cytometry application mentioned above. In these cases the particles are well-separated upon detection. If this condition cannot be met in a particular application, our method needs to be adapted in order to take into account the overlap of diffraction holograms of nearby particles. Although we believe that this is in principle possible, the necessary changes to the proposed method are complex and beyond the scope of this work.

3. SIMULATION RESULTS

In order to exploit the full learning capacity of the neural network a large enough training data set is mandatory. Therefore the Mie scattering patterns of 20,000 transparent concentric spheres of different sizes in water were simulated and used as inputs to the network. More specifically the input to the i -th unit, x_i , is determined as

$$x_i = \text{const} \cdot \int_0^{T_{\text{exposure}}} \int_{(i+0.5)\Delta-0.5ff\Delta}^{(i+0.5)\Delta+0.5ff\Delta} S_z dx dt, \quad (2)$$

$$i = 0 \dots N_{\text{pix}} - 1,$$

with $N_{\text{pix}} = 1024$ being the number of pixels, $ff = 0.9$ the pixel fill factor, Δ the pixel pitch, and S_z the projection of the Poynting vector on the detector surface normal. The magnitude of S_z corresponds to the hologram intensity value as indicated by Eq. (1). In our simulations we also retained the weaker




Fig. 2. Selection of four different hologram records. The underlying radial symmetry allows one to select only a one-dimensional line scan as input vector to the subsequent neural network (intersecting plane). The dimension of the input vector is determined by the number of pixels in one line of the sensor (shortened in the image for better visibility).

contribution proportional to $|E_o|^2$. The neural network then maps the input vector x onto hidden layer state vector z according to Eq. (3a). A prediction value \hat{y} at the output layer neuron results from a linear combination of all the hidden layer states in Eq. (3b). All the trainable network weights and biases are contained inside the weight matrices $W_{\text{Inpto Hid}}$, $W_{\text{Hidto Out}}$, and the bias vectors bias_{Hid} , bias_{Out} , respectively. As a result of the training the neural network, the weights are incrementally adjusted so as to bring the predicted value \hat{y} closer to its corresponding target y :

$$z = \tanh(W_{\text{Inpto Hid}}^T x + \text{bias}_{\text{Hid}}), \quad (3a)$$

$$\hat{y} = W_{\text{Hidto Out}}^T z + \text{bias}_{\text{Out}}. \quad (3b)$$

All solutions to the Mie scattering problem assumed an x -polarized plane wave of wavelength $\lambda = 532$ nm incident on the layered particle with core refractive index (RI) of 1.39 and shell RI of 1.37. The RI of the surrounding water was taken as 1.34. Particle diameters were chosen according to the probability density functions in Eqs. (4a) and (4b) with $a = 7$ μm , $b = 20$ μm , and $c = 4$ μm . The initial joint distribution of core and shell diameters is shown in Fig. 3. Two distinct depth values, $z_1 = 100$ μm and $z_2 = 250$ μm , of the particles along the optical axis were investigated. We emphasize that at those optical depths neither the Fraunhofer nor the Fresnel approximation holds. We chose fixed values for the optical depth as it agrees with the fact that the depth is well controlled in microfluidic flow channels, which triggered our interest in this research topic (see [11,12]). Nonetheless our method is not restricted to any specific choice of the optical depth, which is often unknown. The artificial neural network can in fact be trained to continuously predict the optical depth value z as well. The reason that we did not include this prediction parameter yet is based on the fact that a much more comprehensive data set is required in order to train this augmented network properly and that it is a very time-consuming process:

$$p_{\text{Core}}(x_C) = (b - a)^{-1} 1_{[a,b]}, \quad (4a)$$

$$p_{\text{Shell}}(x_S | x_C) = (x_C - 1 - c)^{-1} 1_{[c, x_C - 1]}, \quad c < a - 1. \quad (4b)$$

We measure the overall prediction accuracy as the achieved normalized mean square error norm (NRMSE) across the test data:




Fig. 3. Artificial data set distribution of coated spheres in the model parameter space ($d_{\text{Core}}, d_{\text{Shell}}$). A total of 20,000 samples is divided into training, validation, and test set with splitting ratios 70%, 15%, and 15%, respectively.

$$\text{NRMSE} = \frac{\sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2}}{y_{\text{max}} - y_{\text{min}}}. \quad (5)$$

In order to obtain the best possible results, we optimize the number of units in the hidden layer. On the one hand, more hidden units lower the model bias and lead to a better fit of the training data. On the other hand, the model complexity and computational cost grow as more and more network weights need to be adjusted. This is shown in Fig. 4, in which the number of hidden units is swept for the core diameter prediction task. For fewer than 30 hidden units we observed a substantial prediction error, i.e., the model is not accurate enough. Using a lot more than 30 hidden layer units, however, does not give rise to improved performances. In contrast, the computation time increases and the model is easily prone to overfitting the training data. As a consequence of these findings we reported optimal results at a hidden unit number of 30, whereas the more time-consuming sweeps of detector configurations are accomplished with only 10 hidden layer units.

From an exploratory, but non-exhaustive, search in parameter space the following optimal NRMSE performances were derived: $(4.4 \pm 0.2)\%$ and $(0.74 \pm 0.01)\%$ for the core and shell diameter prediction, respectively. Table 1 summarizes the corresponding parameter settings. The error bounds were calculated as the two-sided, 95% confidence intervals obtained from fivefold cross-validation. For each fold the data set split into training, validation, and test set is repeated randomly.

Next we are interested in the effect of lowering the model complexity by reducing the number of input units. There are two approaches to achieve this. The first one is to prune the network iteratively, choosing the least significant weight for removal and retraining the network each time. The second simply reduces the input space to contain fewer pixels. We used the latter since it allowed us to study the practically more relevant input configurations of regularly spaced pixels. For our purpose

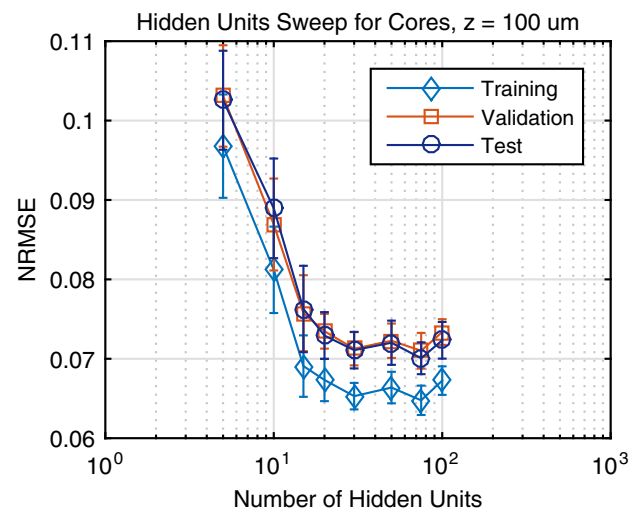


Fig. 4. Normalized mean square error as a function of the number of hidden layer units for core prediction at $z = 100$ μm . Hyperparameters were selected from an extensive scan of the (1024,10,1)-network: learning rate 0.01, annealing time 500, annealing constant 300, L1-regularizer 10^{-4} .

Table 1. Simulation Parameters at Best Performance

Parameter	Symbol	Value
Optics		
Optical depth	z	250 μm
Numerical aperture	NA	1.18
Core RI	n_{Core}	1.39
Shell RI	n_{Shell}	1.37
Solution RI	n_{Water}	1.34
Pixel number/input units	N_{pix}	305
Pixel pitch	Δ_{pix}	1.5 μm
Pixel fill factor	ff	0.9
Machine learning		
Hidden layer units	N_{HLU}	30
Learning rate	lr	0.01
L1-regularization	L1	10^{-4}
Annealing time	T	500
Annealing constant	C	30

we were decimating the recorded diffraction patterns by powers of two, which results in fewer pixels. One could equally replace the current imager (1024 pixels, 1.34 NA) by an imager with fewer pixels but the same numerical aperture and pixel size (e.g., 512 pixels, 1.34 NA). This decrease in the number of input units is an important step on the route toward real-time processing, in which fast on-line retraining algorithms adjust the network weights over time. That way it would be possible to refocus quickly on the particle under study while it is diffusing.

Because of the weak index contrast in our simulations it is more challenging to extract the core diameter precisely. This is confirmed in Figs. 5 and 6 in which the NRMSE for the core diameter prediction is always bigger than the one corresponding to the shell diameter. Another trend confirmed in those figures is the better prediction accuracy at larger optical depth ($z = 250 \mu\text{m}$).

By inspecting the upper part of Fig. 5 it becomes clear that the shell diameter is strongly affected by the downsampling at short optical depths and its error grows approximately linearly for small downsampling factors. At long optical depths the increase in the shell prediction error is more gradual. In both cases the error stays below 5% even for diffraction patterns of only 64 pixels. In the bottom part of Fig. 5 the same decimation is repeated for the core diameter prediction. The NRMSE error is increasing sublinearly with better results at longer optical depth. Only 128 pixels (downsampling integer factor of eight) are necessary to achieve 10% relative accuracy.

As in conventional microscopy, the resolution in digital inline holographic microscopy is inversely proportional to the numerical aperture of the imager. Since we assume a diffraction pattern that is rotationally invariant in the plane of detection, the subtended half-angle entirely contains the imager. On the one hand, the size of the imager or the range of the pixels used is the limiting factor for the NA. On the other hand, a restricted number of pixels to be read out means higher processing speed, and there is a trade-off between both. We are thus interested in the dependence of the prediction accuracy on the NA. These results are shown in Fig. 6, in which the NA is varied by one-sided cropping of the digital image. As a consequence thereof, the number of pixels N_{pix} used as input units depends on the

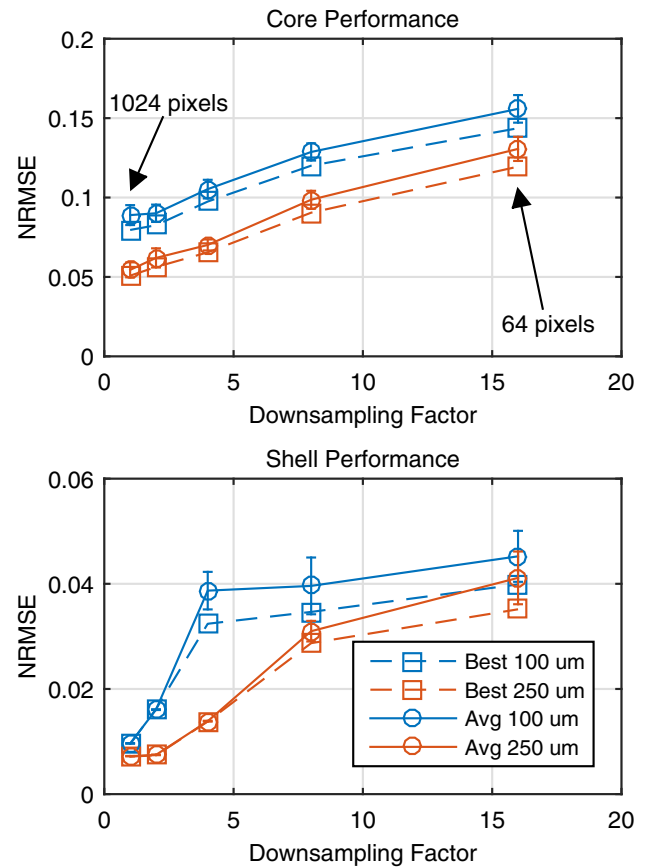


Fig. 5. Normalized root mean square error as a function of the downsampling factor. Top: results for the shell diameter prediction. Bottom: results for the core diameter prediction. The best and the average performance for 10 randomly initialized networks are shown at two optical depths (100 and 250 μm).

optical depth z and can be determined by the following relation [Eq. (6)], in which n stands for the ambient RI, i.e., water (see also 1):

$$N_{\text{pix}} = \frac{z}{\Delta_{\text{pix}}} \frac{\frac{\text{NA}}{n}}{\sqrt{1 - \left(\frac{\text{NA}}{n}\right)^2}}. \quad (6)$$

Analyzing Fig. 6, one concludes that the shell diameter prediction is nearly unaffected by the change in NA and shows constant performance throughout a large range of aperture sizes. Only at very small NA (0.14) and short optical depths (100 μm) the prediction accuracy drops. That is a strong indication for the fact that the network infers the shell size correctly from the forward scatter at small angles, similar to what is done in standard scattering theory to estimate the particle size. In contrast to this, the core diameter prediction requires a large enough NA close to 1 in order to work most accurately. This is expected since a good resolution is necessary to resolve very thin coatings too. Interestingly the optimal value of the core prediction accuracy is found at numerical apertures of 1.0 to 1.2, slightly below the maximal achievable NA. This suggests that the neural network is overfitting for large angles. Hence, it is helpful to reduce the range of captured pixels to about 100 and 300 at optical depths of 100 μm and 250 μm , respectively.

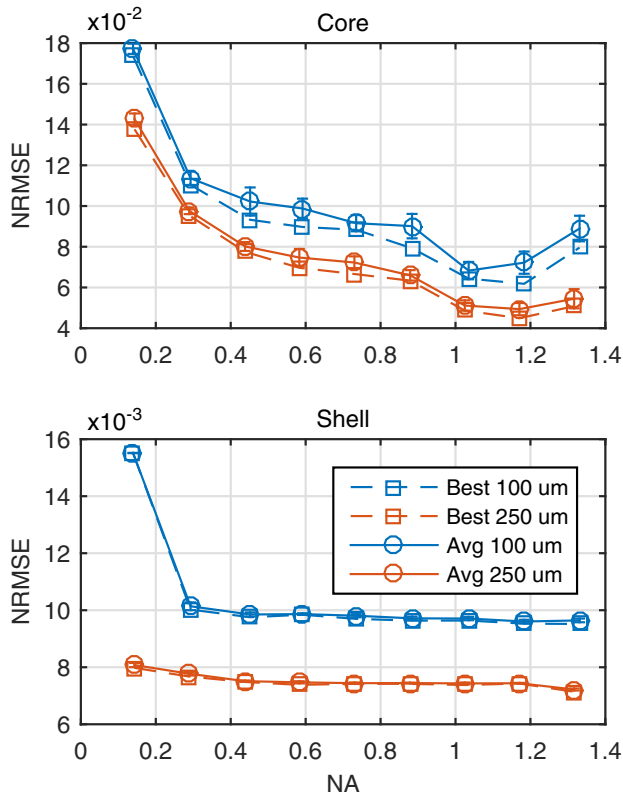


Fig. 6. Normalized root mean square error as a function of the numerical aperture of the imaging system. Top: results for the core diameter prediction. Bottom: results for the shell diameter prediction. The best and the average performance for 10 randomly initialized networks are shown at two optical depths (100 and 250 μm).

This measure decreases the model complexity, decreases the readout time, and avoids the overfitting.

Finally we take a glance at the detailed error map in order to see which points in parameter space are the most difficult to predict. This information is obscured by the use of a global error measure. Therefore we study the optimal network (82-10-1) determined from Fig. 6 at an optical depth of 100 μm .

Figure 7 clearly shows that the core diameter prediction accuracy is worsening for tiny cores, especially for increasing shell diameters. In these situations the core contribution to the scattering amplitude is particularly small. Shell diameters are predicted very accurately within the entire configuration space. The worst possible relative errors of about 4% are those obtained for small-sized spheres. We can explain the larger error for the core diameter prediction, especially in regions where $D_{\text{Core}} \ll D_{\text{Shell}}$, by considering the scattering of light by particles with a small index contrast as a first-order scattering process under the Born approximation. The small index contrast leads to a small magnitude of the scattering potential that can be treated as a perturbation to the free particle problem. As a consequence one can show that the differential cross section σ_{diff} scales as

$$\sigma_{\text{diff}} \propto \Delta n V^2 k_0^4, \quad (7)$$

in the limit of small angles. That is at a fixed incident wavelength (wave vector k_0) and index contrast Δn the small angle

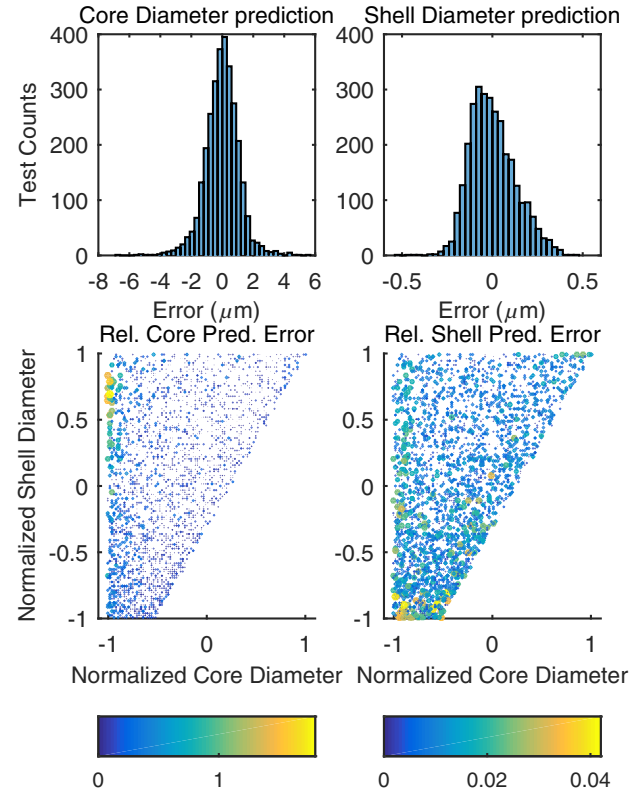


Fig. 7. Prediction errors for the test data and a (82-10-1)-network ($\text{NA} = 10$, $N_{\text{pix}} = 1024$). The focal length is 100 μm . Top row: histograms for the absolute error distribution for both core and shell diameter prediction. The error histogram takes the shape of a Gaussian distribution for the core diameter prediction and a shifted beta distribution for the shell diameter prediction. Bottom row: the relative prediction error for both core and shell diameter is displayed in parameter space with bigger markers indicate bigger error magnitudes.

approximation of the differential cross section scales with the square of the particle volume V^2 , or as D^6 in terms of the particle diameter D . Practically this implies that the overall particle size can be inferred from the hologram by inspecting the scattered intensity pattern at small angles. Indeed the relative prediction error of the shell diameter is very small for all the particles under consideration. However, Eq. (7) also implies that the scattered signal originating from the core is very weak in comparison to the one that stems from the shell in the limit of small core diameter and large shell diameter, $\frac{D_{\text{Core}}}{D_{\text{Shell}}} \ll 1$. It scales as $(\frac{D_{\text{Core}}}{D_{\text{Shell}}})^6$. Such a weak signature naturally falls beyond the detection limit in any practical arrangement or is shadowed by rounding and averaging in simulations. Therefore we conclude that the nature of the larger core prediction errors shown in Fig. 7 is not due to our specific prediction method but a fundamental problem intrinsic to weak scattering processes.

4. CONCLUSION

We demonstrated numerically that characteristic parameters of a particle can be reliably retrieved by direct investigation of its holographic interference pattern with the help of a single-layered feed-forward neural network. We implemented a

simple model for light scattering off WBCs by studying the digital inline holograms of concentric, transparent spheres in the Mie regime. In this way important cell parameters such as overall cell size and nucleus size can be predicted. Those cell parameters are significant for classification of different groups of WBCs. Our best simulation results for sphere diameters varying between 7 and 20 μm achieve NRMSE accuracy of $(4.4 \pm 0.2)\%$ and $(0.74 \pm 0.01)\%$ for the core and shell diameter, respectively. Future efforts are necessary in building more realistic cell models, including multi-lobed, non-concentric nuclei and granules.

The neural network boosts real-time application because of its intrinsic parallelism and easy-to-implement matrix operations. No time-consuming image reconstruction is needed and all the training experience of the network is stored in its connection weights. No time-consuming look-up procedure in a dictionary-based solution is necessary. Further investigation is necessary in order to include more realistic noise models, particle shapes (e.g., aspherical core and/or shell), and a larger configuration space including multiple particle imaging events.

Funding. Belgian Science Policy Office (BELSPO); European Research Council (ERC) (NARESCO); Interuniversity attraction pole (IAP) (Photonics@be).

Acknowledgment. We would like to thank the interuniversity attraction pole (IAP) Photonics@be of the Belgian Science Policy Office and the ERC NaResCo starting grant for supporting this work.

REFERENCES

1. R. A. Jain, "The manufacturing techniques of various drug loaded biodegradable poly(lactide-co-glycolide) (PLGA) devices," *Orthop. Polym. Biomater.* **21**, 2475–2490 (2000).
2. M. Lattuada and T. A. Hatton, "Synthesis, properties and applications of Janus nanoparticles," *Nano Today* **6**, 286–308 (2011).
3. A. Walther and A. H. E. Müller, "Janus particles: Synthesis, self-assembly, physical properties, and applications," *Chem. Rev.* **113**, 5194–5261 (2013).
4. X. Han, Y. Liu, and Y. Yin, "Colorimetric stress memory sensor based on disassembly of gold nanoparticle chains," *Nano Lett.* **14**, 2466–2470 (2014).
5. V. P. Maltsev, A. G. Hoekstra, and M. A. Yurkin, "Optics of white blood cells: optical models, simulations, and experiments," in *Advanced Optical Flow Cytometry*, V. V. Tuchin, ed. (Wiley-VCH, 2011), pp. 63–93.
6. H. W. Hou, M. E. Warkiani, B. L. Khoo, Z. R. Li, R. A. Soo, D. S.-W. Tan, W.-T. Lim, J. Han, A. A. S. Bhagat, and C. T. Lim, "Isolation and retrieval of circulating tumor cells using centrifugal forces," *Sci. Rep.* **3**, 1259 (2013).
7. M. Seifi, L. Denis, and C. Fournier, "Fast and accurate 3D object recognition directly from digital holograms," *J. Opt. Soc. Am. A* **30**, 2216–2224 (2013).
8. A. Ishimaru, S. Kitamura, R. J. Marks, L. Tsang, C. M. Lam, and D. C. Park, "Particle-size distribution determination using optical sensing and neural networks," *Opt. Lett.* **15**, 1221–1223 (1990).
9. F. Strubbe, S. Vandewiele, C. Schreuer, F. Beunis, O. Drobchak, T. Brans, and K. Neyts, "Characterizing and tracking individual colloidal particles using Fourier-Bessel image decomposition," *Opt. Express* **22**, 24635–24645 (2014).
10. M. S. Marshall and R. E. Benner, "Particle characterization using a holographic ring-wedge detector and an optical neural network," *Proc. SPIE* **1772**, 251–257 (1993).
11. D. Vercauteren, A. Dusa, R. Stahl, G. Vanmeerbeeck, K. de Wijs, C. Liu, D. Prodanov, P. Peumans, and L. Lagae, "Three-part differential of unlabeled leukocytes with a compact lens-free imaging flow cytometer," *Lab Chip* **15**, 1123–1132 (2015).
12. B. Schneider, G. Vanmeerbeeck, R. Stahl, L. Lagae, and P. Bienstman, "Using neural networks for high-speed blood cell classification in a holographic-microscopy flow-cytometry system," *Proc. SPIE* **9328**, 93281F (2015).
13. A. Yevick, M. Hannel, and D. G. Grier, "Machine-learning approach to holographic particle characterization," *Opt. Express* **22**, 26884–26890 (2014).
14. J. Garcia-Sucerquia, W. Xu, S. K. Jericho, P. Klages, M. H. Jericho, and H. J. Kreuzer, "Digital in-line holographic microscopy," *Appl. Opt.* **45**, 836–850 (2006).
15. W. Xu, M. H. Jericho, I. A. Meinertzhagen, and H. J. Kreuzer, "Digital in-line holography of microspheres," *Appl. Opt.* **41**, 5367–5375 (2002).
16. V. Kebbel, H.-J. Hartmann, and W. P. Jüptner, "Application of digital holographic microscopy for inspection of micro-optical components," *Proc. SPIE* **4398**, 189–198 (2001).
17. C. Darken, J. Chang, and J. Moody, "Learning rate schedules for faster stochastic gradient search," in *Neural Networks for Signal Processing (IEEE, 1992)*, Vol. **2**.
18. J. F. T. Hastie and R. Tibshirani, *Neural Networks*, 2nd ed. (Springer-Verlag, 2009), Chap. 11, p. 398.
19. M. Born and E. Wolf, *Optics of Metals*, 6th (corrected) ed. (Pergamon, 1980), Chap. XIII, pp. 611–664.
20. J. Leinonen, "Python code for calculating Mie scattering from single- and dual-layered spheres," <http://code.google.com/p/pymiecoated/>.