

MIT Open Access Articles

Lensless computational imaging through deep learning

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Sinha, Ayan et al. "Lensless Computational Imaging through Deep Learning." *Optica* 4, 9 (September 2017): 1117 © 2017 Optical Society of America

As Published: <http://dx.doi.org/10.1364/OPTICA.4.001117>

Publisher: Optical Society of America

Persistent URL: <http://hdl.handle.net/1721.1/118935>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Lensless computational imaging through deep learning

AYAN SINHA^{1*}, JUSTIN LEE², SHUAI LI¹, AND GEORGE BARBASTATHIS^{1,3}

¹Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139

²Institute for Medical Engineering Science, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139

³Singapore-MIT Alliance for Research and Technology (SMART) Centre, One Create Way, Singapore 117543, Singapore

*Corresponding author: sinhayan@gmail.com

Compiled June 27, 2017

Deep learning has been proven to yield reliably generalizable answers to numerous classification and decision tasks. Here, we demonstrate for the first time, to our knowledge, that deep neural networks (DNNs) can be trained to solve inverse problems in computational imaging. We experimentally demonstrate a lens-less imaging system where a DNN was trained to recover a phase object given a raw intensity image recorded some distance away. © 2016 Optical Society of America

OCIS codes: (100.3190) Inverse problems; (100.4996) Pattern recognition, neural networks; (100.5070) Phase retrieval; (110.1758) Computational imaging.

<http://dx.doi.org/10.1364/optica.XX.XXXXXX>

1. INTRODUCTION

Neural network training can be thought of as generic function approximation: given a training set (*i.e.*, examples of matched input and output data obtained from a hitherto-unknown model), a neural network attempts to generate a computational architecture that accurately maps all inputs in a test set (distinct from the training set) to their corresponding outputs. In this paper, we propose that deep neural networks may “learn” to approximate solutions to inverse problems in computational imaging.

A general computational imaging system consists of a physical part and computational part. In the physical part, light propagates through one or more objects of interest as well as optical elements such as lenses, prisms, etc. finally producing a raw intensity image on a digital camera. The raw intensity image is then computationally processed to yield object attributes, *e.g.* a spatial map of light attenuation and/or phase delay through the object—what we traditionally call an “intensity image” or “quantitative phase image,” respectively. The computational part of the system is then said to have produced a solution to the inverse problem.

The study of inverse problems is traced back at least a century to Tikhonov [1] and Wiener [2]. A good introductory book with rigorous but not overwhelming discussion of the underlying mathematical concepts, especially regularization, is [3]. During the past decade, the field experienced a renaissance due to the almost simultaneous maturation of two related mathematics disciplines: convex optimization and harmonic analysis, especially sparse representations. A light technical introduction to these fascinating developments can be found in [4] and

a more detailed exposition in [14].

Neural networks have their own history of legendary ups-and-downs [5] culminating with an even more recent renaissance. This was driven by empirical findings that deep multi-layer architectures, dubbed as “deep neural networks” (DNNs), could generalize better than had been previously thought possible. Vast improvements in the available computational power were certainly helpful; most effective, however, were revivals of older concepts combined with new insights on these concepts’ function and realization. These have included: architectures, such as convolutional connectivity [6–9] for regularization and pruning; nonlinearities, such the now widespread use of non-differentiable piecewise linear units [10] as opposed to the older sigmoidal functions that were differentiable but also prone to stagnation [11]; and algorithms, such as more efficient backprop [12, 13]. Within the last four-five years, neural networks have exhibited spectacular success at solving “hard” computational problems: playing complex games like Atari [23] and Go [24]; object generation [15]; object detection [25]; and image restoration: colorization [26], deblurring [27–29], and in-painting [30].

The hypothesis that we set out to test in this paper is whether a neural network can be trained to recover object estimates from raw intensity images (*i.e.* solve the inverse problem). This is a rather general question and may take several flavors, depending on the nature of the object, the physical design of the imaging system, etc. We chose to test our hypothesis in a very specific “heavy” computational imaging scenario: a lensless optical setup where diffraction patterns of pure phase objects under coherent illumination were captured as “raw




Fig. 1. DNN training. Rows (a) and (b) denote the networks trained on Faces-LFW and ImageNet dataset, respectively. (i) randomly selected example drawn from the database; (ii) calibrated phase image of the drawn sample; (iii) diffraction pattern generated on the CMOS by the same sample; (iv) DNN output before training (*i.e.* with randomly initialized weights); (v) DNN output after training.

images”.

Our experimental arrangement, described in more detail in Section 2, falls in-between two categories of imaging systems that could be traditionally called “digital holographic imaging,” [34] and “transport-of-intensity imaging” [18, 21]. It is neither, because it violates the necessary assumptions of sparse objects leaving most of the incoming light unscattered to serve as reference beam for the digital hologram; and of sparse object gradients that avoid singularities in the transport-of-intensity equation. Hence, either technique would be expected to require significant fine-tuning of regularization parameters to yield satisfactory results.

The idea of using neural networks to clean up images isn’t exactly new. For example, Hopfield’s associative memory network [31] was capable of retrieving entire faces from partially obscured inputs, and was implemented in an all-optical architecture [32] when computers weren’t nearly as powerful as they are now. Recently, Horisaki et al. [33] used support-vector machines, a form of bi-layer neural network with nonlinear discriminant functions, also to recover face images when the obscuration is caused by scattering media. Here, we extend upon the aforementioned efforts and train a deep neural network to recover images of objects given “raw image” measurements of the modulus of their diffraction patterns.

Our results demonstrate that DNNs are capable of “learning” the inverse mapping between raw intensity image and object *directly* from experimental data. Our results also suggest that the neural network “learns” the underlying governing equations of the system, including its forward operator and possible deviations from underlying idealizations and assumptions. This lack of need for a prior model is notable because it removes the difficulty of correctly specifying the forward operator; many optimization approaches are sensitive to errors due to inaccurate or incomplete forward models.

Neural network approaches often come under criticism because the quality of training depends on the quality of the examples given to the network during the training phase. For instance, if the inputs used to train a network are not diverse enough, then the DNN will learn priors of the input images instead of generalized rules for “cleaning up images.” This was the case in [33], where an SVM trained using images of faces could adequately reconstruct faces, but when given the task of reconstructing images of natural objects such as a pair of scissors, the trained SVM still returned an output that resembled a human face.

For our specific problem, an ideal training set would encompass




Fig. 2. Experimental arrangement. SF: spatial filter; CL: collimating lens; M: mirror; POL: linear polarizer; BS: beam splitter; SLM: spatial light modulator.

all possible “phase objects.” Unfortunately, “phase objects,” generally speaking, constitute a rather large class of objects and it would be unrealistic to attempt to train a network sampling from across all possible objects from this large class. Instead, we synthesize phase objects in the form of natural images derived from the ImageNet [36] database because it is readily available and widely used in the study of various machine learning problems. For comparison, we also trained a separate network using a narrower class of (facial) images from the Faces-LFW [35] database.

As expected, our network did well when presented with unknown phase objects in the form of faces or natural images that it had been trained to. Notably, the network also performed well when presented with objects outside of its “training class” – the DNN trained using images of faces was able to reconstruct images of natural objects, and the DNN trained using images of natural objects was able to reconstruct images of faces. Additionally, both DNNs were able to reconstruct completely distinct images including: handwritten digits, characters from different languages (Arabic, Mandarin, English), and images from a disjoint natural image dataset.

Both trained networks yielded accurate results even when the object-to-sensor distance(s) in the training set slightly differed from that of the testing set, suggesting that the network is not merely pattern-matching but instead has actually “learned” a generalizable model approximating the underlying system.

The details of our experiment, including the physical system and the computational training and testing results, are described in Section 2. The neural network itself is analyzed in Section 3, and concluding thoughts are in Section 4.




Fig. 3. Detailed schematic of our DNN architecture, indicating the number of layers, nodes in each layer, etc.

2. EXPERIMENT

Our experimental arrangement is as shown in Figure 2. Light from a He-Ne laser source (Thorlabs, HNL210L, 632.8nm) first transmits through a spatial filter, which consists of a microscope objective (Newport, M-60X, 0.85NA) and a pinhole aperture ($D = 5\mu\text{m}$), to remove spatial noise. After being collimated by the lens ($f = 150\text{mm}$), the light is reflected by a mirror and then passes through a linear polarizer to set the appropriate polarization. After that, the light is split by a beam splitter. A spatial light modulator (Holoeye, LC-R 720, reflective) is placed normally incident to the transmitted light and acts as a pixel-wise phase object. The SLM-modulated light is then reflected by the beam splitter and passes through a linear polarization analyzer, before being collected by a CMOS camera (Basler, A504k). Images recorded are then processed on an Intel i7 CPU, with neural network computations performed on a GTX1080 graphics card (NVIDIA).

According to its user manual, the LC-R 720 SLM can realize (approximate) pure-phase modulation if we modulate the light polarization properly. Specifically, for He-Ne laser light, if we set the polarization of the incident beam at 45° linearly polarized with respect to the vertical direction and also set the linear polarization analyzer to be oriented at 340° with respect to the vertical direction, then the amplitude modulation of the SLM will become almost independent of the assigned (8-bit gray-level) input. In this arrangement, the phase modulation of the SLM follows a monotonic, almost-linear relationship with the assigned pixel value (with maximum phase depth: $\sim 1\pi$). We experimentally evaluated the correspondence between 8-bit grayscale input images projected onto the SLM and phase values in the range $[0, -\pi]$ (see supplement). In this paper, we approximate our SLM as a pure-phase object and computationally recover the phase using a neural network.

The CMOS detector was placed after a free-space propagation distance d , which ranged from $\sim 37.5 - 97.5\text{cm}$ to record diffraction patterns. Our experiment consists of two phases: training and testing. During the training phase, we modulate the phase SLM according to samples randomly selected from the Faces-LFW or ImageNet database. We resize, and pad selected images before displaying them on our SLM. Two examples of inputs, as they are sent to the SLM, and their corresponding raw intensity images (diffraction patterns) as captured on the CMOS are shown in Figure 1. Our training set consisted of 10,000 such faces/images - diffraction pattern pairs. The raw intensity images from all these training examples were used to train the weights in our DNN. We used a Zaber A-LST1000D stage with repeatability $2.5\mu\text{m}$ to translate the camera in order to analyze the robustness of the learnt network to perturbations (See: Network Analysis).

Our DNN uses a convolutional residual neural network (ResNet) architecture. In a convolutional neural network (CNN), inputs are passed from nodes of each layer to the next, with adjacent layers connected by convolution. Convolutional ResNets extend CNNs by adding short term memory to each layer of the network. The intuition behind ResNets is that one should only increase the depth of the neural network if he/she stands to gain something by adding that extra layer. The intuition behind ResNets is that one should only adds a new layer if you can get something extra out of adding that layer. ResNets ensure that the $N + 1$ th layer learns something new about the network by also providing the original input to the output of the $(N + 1)$ th layer and performing calculations on the residual of the two. This forces the new layer to learn something different from what the input has already encoded/learned [9].

A diagram of our specific DNN architecture is shown in Fig. 3. The input layer is the image captured by the CMOS camera. It is then successively decimated by 7 residual blocks of convolution + down-sampling followed by 6 residual blocks of deconvolution + upsampling, and finally 2 standard residual blocks. Some of the residual blocks are

comprised of dilated convolutions so as to increase the receptive field of the convolution filters, and hence, aggregate diffraction effects over multiple scales [38]. We use skip connections to pass high frequency information learnt in the initial layers down the network towards the output reconstruction, and have two standard residual blocks at the end of the network to finetune the reconstruction. At the very last layer of our CNN, the values represent an estimate of our input signal. The connection weights are trained using backpropagation (not to be confused with optical backpropagation) on the L_1 error between the network output and the nominal appearance of the training samples represented as:

$$\min \frac{1}{wh} \sum_{(m,n)} \| (Y(m,n) - G(m,n)) \|_1 \quad (1)$$

Here, w, h are the width and the height of the output, Y is the output of the last layer, and G is the ground truth phase value. $G(m, n)$ lies in the range $[0, -\pi]$.

We collected data from six separate experiment runs using training inputs from Faces-LFW or ImageNet and object-to-sensor distances of 37.5cm , 67.5cm , or 97.5cm . These data were used to train six separate DNNs for evaluation.

Fig. 1(iv) shows a sample DNN's output at the beginning of its training phase (*i.e.* with randomly initialized weights), and Fig. 1(v) shows the network output after training, for the same example object-raw image pairs. Training each network took ≈ 20 hours using Tensorflow and a Nvidia GTX 1080 GPU. We provide analysis of the trained DNN in Section 3.

Our testing phase consisted of: (1) sampling disjoint examples from the same database (either Faces-LFW or ImageNet) and other databases such as MNIST, CIFAR, Faces-ATT etc., (2) using these test examples to modulate the SLM and produce raw intensity images on the camera, (3) passing these intensity images as inputs to out trained DNN, and (4) comparing the output to ground truth.

3. RESULTS AND NETWORK ANALYSIS

The standard method of characterizing neural network training is by plotting the progression of training and test error across training epochs (iterations in the backpropagation algorithm over all examples). These curves are shown in Figure 5 for our network trained using the ImageNet database and tested using images from: (a) Faces-LFW (b) a disjoint ImageNet set, (c) images from an English/Chinese/Arabic characters database, (d) the MNIST handwritten digit database, (e) Faces-ATT, (f) CIFAR, (g) a constant-value "Null" image. Our ImageNet learning curves in Figure 5d show convergence to low value after ~ 10 epochs, indicating that our network has not overfit to our training dataset. We plot bar graphs for the mean absolute error (MAE) over test examples in the 7 different datasets for each of the 3 object-to-sensor distances in Figure 5. Lower MAE was reported for test images with large patches of constant value (characters, digits, Null) as their sparse diffraction patterns were easier for our DNN to invert. Notably, both our bar graphs and learning curves show low test error for the non-trained images, suggesting that our network generalizes well across different domains.

This is an important point and worth emphasizing: despite the fact that our network was trained exclusively on images from the ImageNet database – *i.e.*, images of planes, trains, cars, frogs, artichokes, etc., it is still able to accurately reconstruct images of a completely different class (*e.g.*, faces, handwritten digits, and characters from different languages). This strongly suggests that our network has learned a model of the underlying physics of the imaging system or at the very least a generalizable mapping of low-level textures between our output diffraction patterns and input images.




Fig. 4. Qualitative analysis of our trained deep neural networks for three object-to-sensor distances (37.5 cm, 67.5 cm and 97.5 cm) on different datasets. (i) Ground truth pixel value inputs to the SLM. (ii) Corresponding phase imaged calibrated by SLM response curve. (iii) Raw intensity images captured by CMOS detector at distance $d \sim 37.5\text{cm}$. (iv) DNN reconstruction from raw images when trained using Faces-LFW dataset. (v) DNN reconstruction when trained using ImageNet dataset. Columns (vi-viii) and (ix-xi) follow the same sequence as (iii-v) but in these sets the CMOS is placed at a distance of $\sim 67.5\text{cm}$ and $\sim 97.5\text{cm}$, respectively. Rows (a-f) correspond to the dataset from which the test image is drawn: (a) Faces-LFW, (b) ImageNet, (c) Characters, (d) MNIST Digits, (e) Faces-ATT, or (f) CIFAR.

A more pronounced qualitative example demonstrating this is shown in the columns (iv) (vii) and (x) of Figure 4. Here, we trained our network using images exclusively from the Faces-ATT database. Despite this limited training set, the learned network was able to accurately reconstruct images from the ImageNet, handwritten digits, and characters datasets. This is in contrast to results shown in [33], where an SVM trained on images of faces was able to accurately reconstruct images of faces but not other classes of objects.

How robust is our network to sensor displacement? Is it shift and rotation invariant? To answer these questions, we fed our trained network raw intensity images at different lateral and axial positions, relative to that of the training set images. Quantitative results of these perturbations are shown in Figures 6, 7, 8, and qualitative results for the networks trained at distance 37.5 cm are shown in Figures 11, 12 and 13. Qualitative results for the other 2 distances are in the supplement. The results show that our trained network is robust to moderate perturbations in sensor displacement and is somewhat shift and rotation invariant. As expected, the system fails when the displacement is significantly greater (Figure 9).

What exactly is our network learning? To get a sense of what the network has learned, we examined its maximally-activated patterns (MAPs), i.e., what types of inputs would maximize network filter response (gradient descent on the input with average filter response as loss function [41]). Our results are shown in Figure 10 and compared with the results of analogous analysis of a de-blurring network of similar architecture as well as an ImageNet classification DNN. Compared with MAPs of ImageNet and a Deblurring network, the MAPs of our phase-retrieval network show much finer/low-level textures at deep layers in the network. This suggests that the network is utilizing low-level textures (representative of a wide variety of localized diffraction

patterns) when learning how to invert our inverse problem.

4. CONCLUSIONS AND DISCUSSION

The architecture presented here was deliberately well controlled, with an SLM creating the phase object inputs to the neural network for both training and testing. This allowed us to quantitatively and precisely analyze the behavior of the learning process. Application-specific training, e.g. replacing the SLM with physical phase objects for more practical applications, we judged beyond the scope of the present work. Other obvious and useful extensions would be to include optics, e.g. a microscope objective for microscopic imaging in the same mode; and to attempt to reconstruct complex objects, *i.e.* imparting both attenuation and phase delay to the incident light. The significant anticipated benefit in the latter case is that it would be unnecessary to characterize the optics for the formulation of the forward operator—the neural network should “learn” this automatically as well. We intend to undertake such studies in future work.

FUNDING INFORMATION

This research was funded by the Singapore National Research Foundation through the SMART program (Singapore-MIT Alliance for Research and Technology) and by the Information Advanced Research Projects Agency (iARPA) through the RAVEN Program. Justin Lee acknowledges funding from the U.S. Department of Energy Computational Science Graduate Fellowship (CSGF) (DE-FG02-97ER25308).

ACKNOWLEDGMENTS

We gratefully acknowledge Ons M’Saad for help with the experiments, and Petros Koumoutsakos and Zhengyun Zhang for useful discussions




Fig. 5. Quantitative analysis of our trained deep neural networks for 3 object-to-sensor distances of (a) 37.5 cm, (b) 67.5 cm, and (c) 97.5 cm for the DNNs trained on Faces-LFW (blue) and ImageNet (red) on 7 datasets. (d) The training and testing error curves for network trained on ImageNet at distance 37.5 cm over 20 epochs.

and suggestions.

See Supplement for supporting content.

REFERENCES

1. A.N. Andrey Tikhonov, "On the stability of inverse problems," *Doklady Akademii Nauk SSSR*, 39 (5): 195–198, (1943).
2. N. Wiener, "The interpolation, extrapolation and smoothing of stationary time series," Report of the Services 19, Research Project DIC-6037 MIT, February 1942.
3. M. Bertero, and P. Boccacci, "Introduction to Inverse Problems in Imaging," IOP Publishing (1998).
4. E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE signal processing magazine*, 25.2: 21–30 (2008).
5. M. Minsky and S. Papert, "Perceptrons: An Introduction to Computational Geometry," The MIT Press (1972).
6. Y. LeCun, "Generalization and network design strategies," Technical Report CRG-TR-89-4, University of Toronto (1989).
7. A. Krizhevsky, I. Sutskever, and G. E. Hinton. "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, 1097-1105 (2012)
8. C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov,




Fig. 6. Quantitative analysis of the sensitivity of the trained deep convolutional neural network to the object-to-sensor distance. The network was trained on (a) Faces-LFW database and (b) ImageNet and tested on disjoint Faces-LFW and ImageNet sets, respectively.

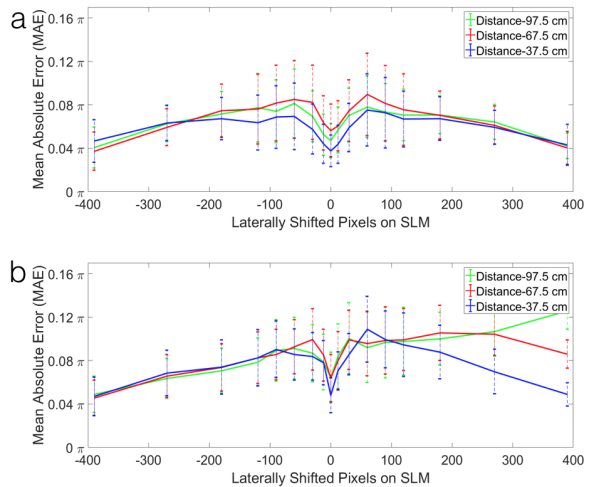


Fig. 7. Quantitative analysis of the sensitivity of the trained deep convolutional neural network to laterally shifted images on the SLM. The network was trained on (a) Faces-LFW database, (b) ImageNet and tested on disjoint Faces-LFW and ImageNet sets, respectively.

- D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going Deeper with Convolution," *IEEE Conference on Computer Vision and Pattern Recognition* (2015).
9. K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *IEEE Conference on Computer Vision and Pattern Recognition* (2016).
10. K. Fukushima, "Cognitron: a self-organizing multi-layered neural network," *Biological cybernetics* 20:121-136, 1975.
11. I. Goodfellow, Y. Bengio, and A. Courville, "Deep learning," (section 6.3) MIT Press, 2016.
12. D. Rumelhart, S. Hinton, and R. Williams, "Learning representations by back-propagation errors," *Nature* 323:533-536, 1989.
13. Y. LeCun, L. Bottou, G. B. Orr, and R.-K. Müller, "Efficient back-prop," *Lecture Notes in Computer Science LNCS 1524*, Springer-Verlag, 2010.

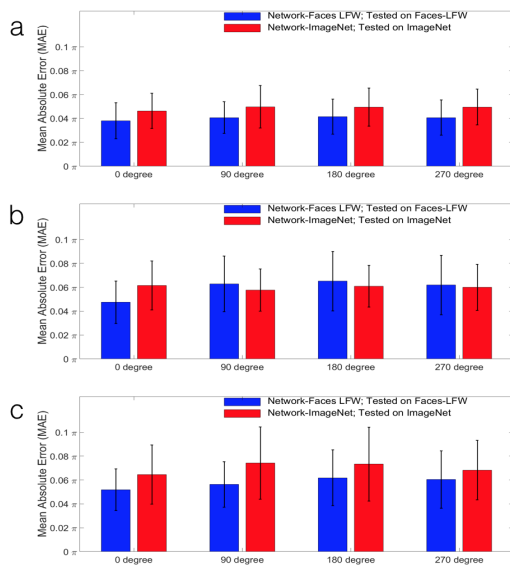


Fig. 8. Quantitative analysis of the sensitivity of the trained deep convolutional neural network to rotation of images on the SLM. The baseline distance on which the network was trained is (a) 37.5 cm, (b) 67.5 cm and (c) 97.5 cm, respectively.

14. D. Brady, "Optical imaging and spectroscopy", John Wiley and Sons (2009)
15. A. Dosovitskiy, J. T. Springenberg, M. Tatarchenko, and T. Brox, "Learning to Generate Chairs, Tables and Cars with Convolutional Networks," IEEE Conference on Computer Vision and Pattern Recognition, pp. 1538–1546. (2015).
16. A. Yevick, M. Hannel, and, D.G. Grier, "Machine-learning approach to holographic particle characterization," Optics Express (22), 26884–26890 (2014).
17. R. Horisaki, R. Takagi, and J. Tanida, "Learning-based imaging through scattering media," Optics Express 24, 13738-13743 (2016)
18. M. Teague, "Deterministic phase retrieval: a Green's function solution," Journal of the Optical Society of America 73, 1434–1441 (1983)
19. D. Paganin and K. Nugent, "Noninterferometric phase imaging with partially coherent light," Physical Review Letters 80(12),2586–2589 (1998).
20. E. D. Barone-Nugent, A. Barty, and K. A. Nugent, "Quantitative phase-amplitude microscopy I: optical microscopy," Journal of Microscopy 206(Pt 3), 194–203 (2002).
21. N. Streibl, "Phase imaging by the transport of equation of intensity," Optical Communications 49(1), 6–10 (1984).
22. C. J. R. Sheppard, "Defocused transfer function for a partially coherent microscope and application to phase retrieval," Journal of the Optical Society of America 21(5), 828–831 (2004).
23. V. Mnih, K. Kavukcuoglu, D. Silver, A. Rusu, J. Veness, M. Belle-mare, A. Graves, M. Riedmiller, A. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," Nature 518 (7540): pp. 529–533 (2015).
24. D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, "Mastering the game of Go with deep neural networks and




Fig. 9. Failure cases on networks trained on Faces-LFW (row *a*) and ImageNet (row *b*) datasets. (i) Ground truth input, (ii) calibrated phase input to SLM, (iii) raw image on camera (iv) reconstruction by DNN trained on images at distance 37.5 cm between SLM and camera and tested on images at distance 107.5 cm, (v) raw image on camera and (vi) reconstruction by network trained on images at distance 97.5 cm between SLM and camera and tested on images at distance 27.5 cm.




Fig. 10. (1) 16×16 inputs that maximally activate the last set of 16 convolutional filters in layer 1 of our phase retrieval network trained on ImageNet at distance of 37.5 cm, a deblurring network, and an ImageNet classification network. The image is downsampled by a factor of 2 in this layer. (2) 32×32 inputs that maximally activate the last set of 16 randomly chosen convolutional filters in layer 3 of: our network, a deblurring network, and ImageNet classification network. The raw image is downsampled by a factor of 8 in this layer.

- tree search," Nature 529, (7587): pp. 484–489, (2016).
25. Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature 521, 436–444 (2015).
26. Z. Cheng, Q. Yang, and B. Sheng, "Deep colorization," IEEE International Conference on Computer Vision, pp. 415–423 (2015).
27. C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," European Conference on Computer Vision, pp. 184–199 (2014).
28. L. Xu, J. S. Ren, C. Liu, and J. Jia, "Deep convolutional neural network for image deconvolution," Advances in Neural Information Processing Systems, pp. 1790–1798 (2014).
29. J. Sun, W. Cao, Z. Xu, and J. Ponce, "Learning a convolutional neural network for non-uniform motion blur removal," IEEE Conference on Computer Vision and Pattern Recognition, pp. 769–777 (2015).
30. J. Xie, L. Xu, and E. Chen, "Image denoising and inpainting

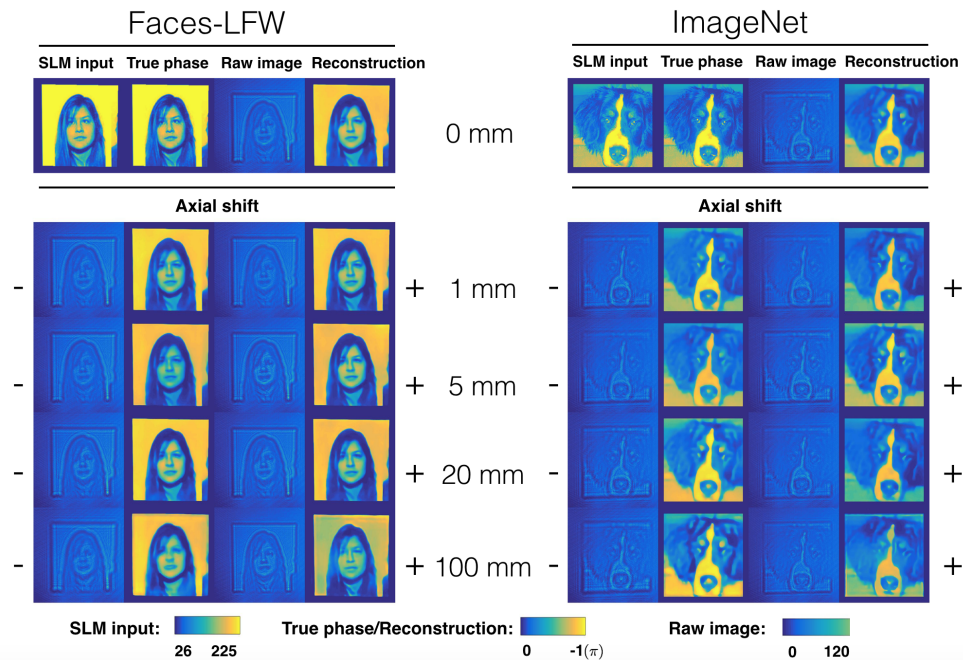


Fig. 11. Qualitative analysis of the sensitivity of the trained deep convolutional neural network to the object-to-sensor distance. The baseline distance on which the network was trained is 37.5 cm.

- with deep neural networks,” In *Advances in Neural Information Processing Systems*, pp. 341–349 (2012).
31. J. J. Hopfield, “Neural networks and physical systems with emergent collective computational abilities”, *Proceedings of the National Academy of Sciences of the USA*, vol. 79 no. 8, pp. 2554–2558 (1982).
 32. J. Jang, S. Jung, S. Lee, and S. Y. Shin, “Optical implementation of the Hopfield model for two-dimensional associative memory,” *Optics Letters* 13, 248–250 (1988).
 33. R. Horisaki, R. Takagi, and J. Tanida, “Learning-based imaging through scattering media,” *Optics Express* 24, 13738–13743 (2016).
 34. M.K. Kim, “Principles and techniques of digital holographic microscopy,” *Journal of Photonics for Energy*, (2010).
 35. G. B. Huang, M. Ramesh, T. Berg and E. Learned-Miller, “Labeled faces in the wild: A database for studying face recognition in unconstrained environments,” Technical Report, University of Massachusetts, Amherst, (2007).
 36. O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg and L. Fei-Fei, “ImageNet Large Scale Visual Recognition Challenge,” *International Journal of Computer Vision* 115, 211–252 (2015).
 37. L. Deng, “The MNIST database of handwritten digit images for machine learning research,” *IEEE Signal Processing Magazine* 29.6, pp. : 141–142 (2012).
 38. F. Yu and V. Koltun, “Multi-Scale Context Aggregation by Dilated Convolutions,” *ICLR* (2016).
 39. M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, “Spatial Transformer Networks,” *Advances in Neural Information Processing Systems*, pp. 2017–2025 (2015).
 40. L. Waller, L. Tian, and G. Barbastathis, “Transport of Intensity phase-amplitude imaging with higher order intensity derivatives,” *Optics Express* 18, 12552–12561 (2010).
 41. M. Zeiler, R. Fergus, D. Fleet, T. Pajdla, B. Schiele and T. Tuytelaars, “Visualizing and Understanding Convolutional Networks,”

European Conference on Computer Vision, pp. 818–833 (2014).




Fig. 12. Qualitative analysis of the sensitivity of the trained deep convolutional neural network to lateral shifts of images on the SLM. The baseline distance on which the network was trained is 37.5 cm.

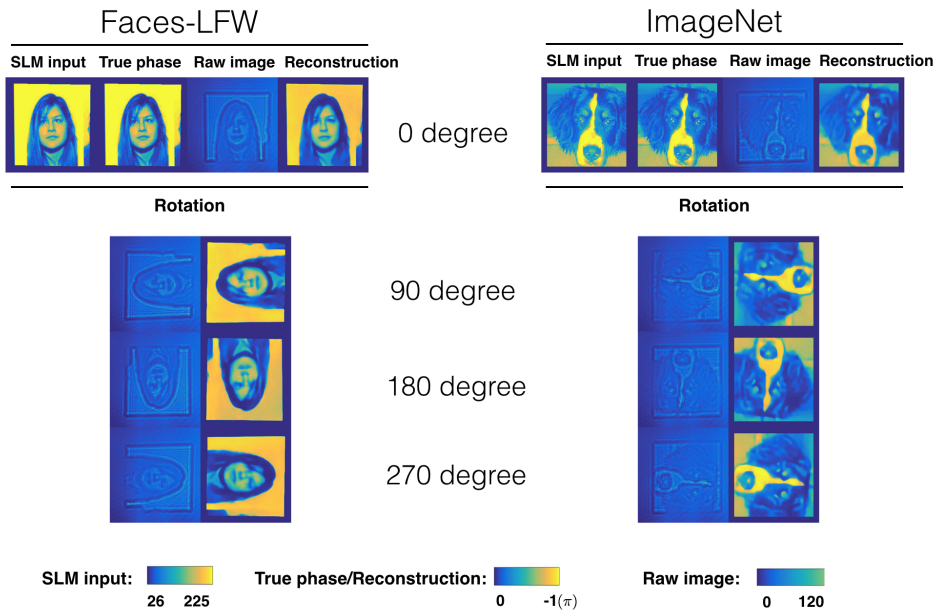


Fig. 13. Qualitative analysis of the sensitivity of the trained deep convolutional neural network to rotation of images in steps of 90. The baseline distance on which the network was trained is 37.5 cm.