



MIT Open Access Articles

On the interplay between physical and content priors in deep learning for computational imaging

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

Citation	Deng, Mo, Li, Shuai, Zhang, Zhengyun, Kang, Iksung, Fang, Nicholas X et al. 2020. "On the interplay between physical and content priors in deep learning for computational imaging." Optics Express, 28 (16).
As Published	10.1364/OE.395204
Publisher	The Optical Society
Version	Final published version
Citable link	https://hdl.handle.net/1721.1/138458
Terms of Use	Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



On the interplay between physical and content priors in deep learning for computational imaging

MO DENG,^{1,6,*} SHUAI LI,^{2,6} ZHENGYUN ZHANG,^{4,5} IKSUNG KANG,¹  NICHOLAS X. FANG,³  AND GEORGE BARBASTATHIS^{3,4}

¹Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02139, USA

²SenseBrain Technology Limited LLC, 2550 N 1st Street, Suite 300, San Jose, CA 95131, USA

³Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02139, USA

⁴Singapore-MIT Alliance for Research and Technology (SMART) Centre, One Create Way, Singapore 117543, Singapore

⁵Work performed while the author was with the Singapore-MIT Alliance Research and Technology (SMART) Centre, Singapore

⁶Equal contribution

*modeng@mit.edu

Abstract: Deep learning (DL) has been applied extensively in many computational imaging problems, often leading to superior performance over traditional iterative approaches. However, two important questions remain largely unanswered: first, how well can the trained neural network generalize to objects very different from the ones in training? This is particularly important in practice, since large-scale annotated examples similar to those of interest are often not available during training. Second, has the trained neural network learnt the underlying (inverse) physics model, or has it merely done something trivial, such as memorizing the examples or point-wise pattern matching? This pertains to the interpretability of machine-learning based algorithms. In this work, we use the Phase Extraction Neural Network (PhENN) [Optica **4**, 1117-1125 (2017)], a deep neural network (DNN) for quantitative phase retrieval in a lensless phase imaging system as the standard platform and show that the two questions are related and share a common crux: the choice of the training examples. Moreover, we connect the strength of the regularization effect imposed by a training set to the training process with the Shannon entropy of images in the dataset. That is, the higher the entropy of the training images, the weaker the regularization effect can be imposed. We also discover that weaker regularization effect leads to better learning of the underlying propagation model, *i.e.* the weak object transfer function, applicable for weakly scattering objects under the weak object approximation. Finally, simulation and experimental results show that better cross-domain generalization performance can be achieved if DNN is trained on a higher-entropy database, *e.g.* the ImageNet, than if the same DNN is trained on a lower-entropy database, *e.g.* MNIST, as the former allows the underlying physics model be learned better than the latter.

© 2020 Optical Society of America under the terms of the [OSA Open Access Publishing Agreement](#)

1. Introduction

1.1. Two unanswered fundamental questions of deep learning in computational imaging

Deep learning (DL) has been proven versatile and efficient in solving many computational inverse problems, including image super-resolution [1–7], phase retrieval [8–19], imaging through the scattering medium [20–22], optical tomography [23–25] and so on. See [26,27] for more detailed

reviews. Besides the superiority of performance over classical approaches in many cases, the DNN-based methods enjoy the advantage of extremely fast inference after the completion of the training stage. It is during the latter that the DNN weights are optimized as the training loss function between ground truth objects and estimated objects is reduced.

Despite great successes, two important questions remain largely unanswered: first, how well does a model trained on one dataset generalize directly to objects in disjoint classes; second, how well does the trained neural network learn the underlying physics model? These questions are well-motivated, since access to a large number of training data in the same category with those in the test set is not always possible and it would reduce the practicality of deep learning if the model trained on one set cannot reasonably well generalize to the other. Moreover, one major skepticism against deep learning is: has the algorithm actually learnt anything about the underlying physics or is it merely doing some trivial point-wise denoising, pattern matching, or, worse, just memorizing and reproducing examples from the training set?

In this paper, we recognize that these two questions are directly related: if the trained DNN were able to perfectly learn the underlying (inverse) physics law, which is satisfied unconditionally by all classes of objects, the prediction performance would not have degraded even if tested on very different examples. Conversely, if the DNN does not learn the model well and instead relies on regularization, i.e. priors on the examples, to reconstruct, then cross-domain generalization will be problematic.

Though training datasets typically contain at least hundreds, if not thousands, training examples, the DNNs are typically highly over-parameterized models; that is, the number of trainable parameters is larger compared to the number of the training examples. Such under-determined, nonlinear system necessitates measures be taken to ensure the learned DNN to correspond to the true physics model. Our investigation unearths one such important factor, the choice of training examples, using the Phase Extraction Neural Network (PhENN) [8] for lensless phase retrieval as an example. In fact, we discover that a trained DNN corresponds to the physics law better, and therefore cross-domain generalizes better, when the training set is more generic, e.g. the ImageNet [28], than if it is more constrained, e.g. the MNIST [29]. Therefore, when encountering insufficient training data in the same class as the test data, which is very common, the best compromise is to train the neural network on a less constrained publicly available standardized dataset such as the ImageNet, with reasonable confidence that it produces the reconstructions accurately in the domain of interest.

1.2. Phase retrieval and the weak object transfer function (WOTF) for lensless phase imaging

In the lensless phase imaging system (Fig. 1), the phase object is illuminated by collimated monochromatic light. The light transmitted through the phase object propagates in free space and forms an intensity pattern on the detector placed at a distance z away. Assuming that the illumination plane wave has unit amplitude, the forward model can be described as

$$g(x, y) = \left| \exp \{ i f(x, y) \} * \exp \left\{ i \frac{\pi}{\lambda z} (x^2 + y^2) \right\} \right|^2. \quad (1)$$

Here, $f(x, y)$ is the phase distribution of the object, $g(x, y)$ is the intensity image captured by the detector, λ is the wavelength of the illumination light, i is the imaginary unit and $*$ denotes the convolution operation.

Eq. (1) is non-linear. However, when the weak object approximation holds, $\exp \{ i f(x, y) \} \approx 1 + i f(x, y)$ [30], the forward imaging model may be linearized as

$$G(u, v) \approx \delta(u, v) + 2 \sin \left(\pi \lambda z (u^2 + v^2) \right) F(u, v). \quad (2)$$

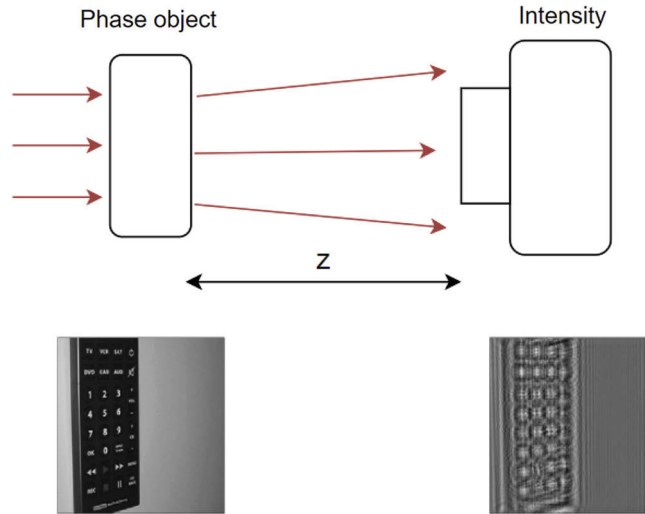


Fig. 1. Schematic plot of the lensless phase imaging system.

Here, $G(u, v)$ and $F(u, v)$ are the Fourier transforms of the intensity measurement $g(x, y)$ and the phase distribution of the object $f(x, y)$, respectively; $\sin(\pi\lambda z(u^2 + v^2))$ is the weak object transfer function (WOTF) for lensless phase imaging. The nulls of the WOTF are of particular significance as sign transitions surrounding these nulls in the frequency domain would cause a π phase shift in the measurement. We will address further on this effect in Section 3.2 and 4.3.

1.3. Phase Extraction Neural Networks (PhENN)

The Phase Extraction Neural Network (PhENN) [8] is a deep learning architecture that can be trained to recover an unknown phase object from the raw intensity measurement obtained through a lensless phase imaging system. Since PhENN was proposed, three general types of strategies have been followed to further enhance its performance. The first category focused on optimizing the network architecture or training specifics of PhENN, including the network depth, the training loss functions etc. [31] The second category focused on optimizing the spatial frequency components of the training data to compensate imbalanced fidelity in the high and low frequency bands in the reconstructions. The same rationale applies not only to phase retrieval but many other applications as well. Li *et al.* proposed a spectral pre-modulation approach [32] to amplify the high frequency content in the training data and experimentally demonstrated that the PhENN trained in this fashion achieved a better spatial resolution while at the cost of the overall reconstruction quality. Subsequently, Deng *et al.* proposed a learning to synthesis by deep neural network (LS-DNN) method [7], which is able to achieve a full-band and high quality reconstruction in phase retrieval and other computational imaging applications, by splitting, separately processing and recombining the low and high spatial frequencies. The third category is to make the learning scheme more physics-informed. Attempts are made where, unlike the original PhENN, the forward model is incorporated via model-based pre-processing [10,18,19], etc. and such strategy is proven particularly useful under most ill-posed conditions. However, since all these efforts are secondary to the main objectives of this paper, we choose not to implement them.

The PhENN architecture used in this paper is shown in Fig. 2. It uses an *encoder-decoder* structure with skip connections, a structure proven efficient and versatile in many applications of interest. The *encoder* consists of 4 Down-Residual blocks (DRBs) to gradually extract

compressed representations of the input signals to preserve high-level features; subsequently, the *decoder* component comprises of 4 Up-Residual blocks (URBs) and two (constant-size) Residual blocks (RBs) to expand the size of the feature maps and form the final reconstruction. To best preserve the high spatial frequencies, skip connections are used to bypass the feature maps from the *encoder* to the corresponding layers of the same size in the *decoder*. More details about the architecture of each functional blocks are available in Appendix A. Compared to the original PhENN [8], in this paper we implemented two modifications: first, we started from intensity patterns of size 256×256 , same as that of the objects, as only with the same pixel patch sizes and number of pixels, can the computation of the weak object transfer function (WOTF) (see Section. 1.2) be physically sound. Second, PhENN is trained with negative Pearson correlation coefficient (NPCC), a loss proven to be more beneficial for restoring fine details of the objects [21], instead of using the mean absolute error (MAE), a pixel-wise loss known to suffer from the oversmoothing issue. The exact form of NPCC is given later in Eq. (5).

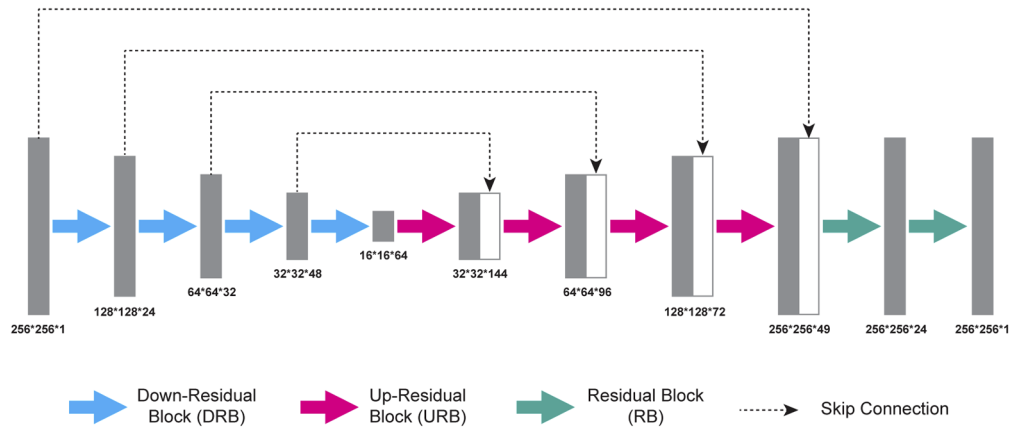


Fig. 2. The general architecture of PhENN.

1.4. Generalization error in machine learning

Previous works have addressed different aspects of the generalization errors of DNNs. Some [33,34] aimed at tightening bounds on the capacity of a model, *i.e.* the number of training examples necessary to guarantee generalization. Reference [35] provided a finer generalization bound for a subclass of neural networks in terms of the product of the spectral norm of their layers and the Frobenius norm of their weights. Yet, beautiful as the works are, the model capacity bounds and the generalization error bounds are not tight enough to provide practical insights. In [36], the authors provide insights into the explicit and implicit role of various regularization methods (not to confuse with the regularization effect imposed by the training data, as mentioned in this paper) to reduce generalization error, including using dropout, data augmentation, weight decay, *etc.* practices that are commonly applied in the deep learning implementation, including in PhENN. Other works examined the generalization dynamics of large DNNs using Stochastic Gradient Descent (SGD) [37], the generalization ability of DNNs with respect to their robustness [38], *i.e.* the ability of the DNN to cope with small perturbations of its input. As [39], a nice review on this general topic, points out, one of the open problems in this area is to understand the interplay between memorization and generalization. This work, unlike in previous works where generalization typically refers to generalization to unseen examples similar to the training dataset, attempts to provide some insights into how the choice of training dataset might affect the cross-domain generalization performance, by an empirical investigation.

2. Entropy as a metric of the strength of the regularization effect imposed by a training set

When a dataset is selected as the training dataset, its influence to the training is reflected by the regularization effect it imposes. For example, when examples with similar statistical distributions are presented to the DNN as the ground truth, the training, or the optimization of the weights in the DNN is susceptible to taking the shortcut of “memorizing” the examples so as to minimize the training loss, instead of learning the underlying physics law, making such trained DNN produce reconstructions similar to those training examples. The cross-domain generalization therefore becomes problematic. Under such scenarios, the overly constrained prior information represented by the training set imposes an unduly strong regularization effect to the training.

Intuitively, ImageNet [28] contains natural images of a broad collection of contents and scenes, and thus should present a more generic prior. On the other hand, MNIST [29], a database with the much more restricted set of handwritten digits only, would be expected to impose stronger regularization. However, to quantify the regularization effect imposed by a particular training dataset is not as straightforward. Here, we employ Shannon entropy [40,41] of the images in the dataset for that purpose: the higher the entropy, the weaker the regularization effect this dataset imposes to the training process.

Shannon entropy is a measure of the uncertainty of a random variable. Let X denote the random variable with the distribution $p(X)$. Without loss of generality, we assume that X is discrete and assumes a finite alphabet $\mathcal{X} = \{x_1, \dots, x_K\}$, then the entropy (in bits) of X , under the distribution $p(\cdot)$, $H_p(X)$, is defined as:

$$H_p(X) = - \sum_{k=1}^K p(x_k) \log_2 p(x_k), \quad (3)$$

where $p(x_k) := \Pr\{X = x_k\}$. The well-known Asymptotic Equipartition Theorem [40,41] states that, with the obvious constraint $\sum_k p(x_k) = 1$, the entropy is maximized at the equiprobable case, *i.e.* $p(x_k) = \frac{1}{K}$ for all k 's and the maximized entropy is $\log_2 K = \log_2 |\mathcal{X}|$. The higher the entropy, the higher the level of uncertainty is in the source. Conversely, if the distribution is completely deterministic (*i.e.* $p(x_i) = 1$ for some i and $p(x_j) = 0$ for $j \neq i$), the entropy of the source is 0, the lowest extreme.

For an image f of size $M \times N$, defined on the alphabet $\mathcal{X}^{M \times N}$, the empirical distribution on the pixel value $X \in \mathcal{X}$, is defined as

$$\hat{p}(x_k) = \frac{\sum_{i,j} \mathbb{1}\{f_{i,j} = x_k\}}{M \times N} = \frac{\text{number of pixels of value } x_k}{M \times N}, \quad (4)$$

and the entropy of this image f can be approximated by the entropy $H_{\hat{p}}(\cdot)$ according to this empirical distribution $\hat{p}(\cdot)$, according to Eq. (3). Note that such estimation of the Shannon entropy, though does not take into account correlations over pixels, turns out to be a good estimate and is widely adopted [42].

In Fig. 3, we show a histogram of our computed Shannon entropy of the images in ImageNet and MNIST, two representative standardized databases. 10000 8-bit images were selected from each set and a histogram of the entropy of the images is computed based on 1000 bins between 0 to 8 bits. The ImageNet images generally have higher entropy (mean 5.524 and standard deviation 0.561) and most images have their entropy between 5 and 6 bits; while the MNIST images have low entropy (mean 0.904 and standard deviation 0.159). Interestingly, this matches well with our anticipation that the MNIST images are approximately binary images – the entropy of the image will be 1 in the case of a perfect binary image with equal densities, *i.e.* $p(x_i) = p(x_j) = 0.5$, for $i \neq j$ and $p(x_k) = 0$ for all other k 's. The deviation of the observed entropy in MNIST images from the ideal entropy of perfect binary images comes from the fact that the

nonzero densities are not equal in the empirical distribution of MNIST images. (and later we will see that such narrow distribution of Shannon entropy found in MNIST makes the reconstruction of WOTF challenging as it offers very limited insight of out-of-focus images).

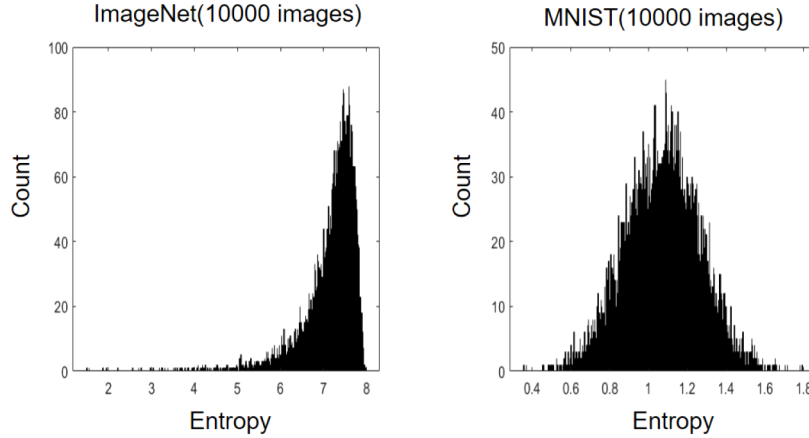


Fig. 3. Entropy histogram of ImageNet and MNIST. Each histogram is computed based on 1000 bins and 10000 images.

For a database, the connection between the entropy of its images and the strength of the regularization effect can be argued as follows: the higher the entropy of images in a database have, the closer the empirical distributions are to the uniform distribution, which is the extreme case where the weakest regularization effect can be imposed. Consider the extreme case where pixels on an image admit an i.i.d. uniform distribution on \mathcal{X} , the empirical distribution in Eq. (4) is a uniform one. In this extreme case, no regularization effect can possibly be imposed by the training examples since if the DNN were to “memorize” such training examples, nothing but the totally random distribution on pixel-values could be remembered. On the other extreme, if the training dataset contains zero-entropy images only (all pixels are identical on each image), the training loss can be minimized merely by the DNN learning to produce such uniform-valued images, totally neglecting the underlying physics model associated with the stochastic processes involved during the image formation. In the next section we show that the stark difference in entropy between ImageNet and MNIST shown in Fig. 3 strongly implies respective difference between the two datasets in terms of generalization ability.

3. Results on synthetic data

3.1. Performance of cross-domain generalization under different training datasets

In this section, we compare the cross-domain generalization performance of PhENN when trained on ImageNet, Face-LFW [43], IC-layout [10] and MNIST, respectively.

The intensity measurement is synthetically generated by the standard phase retrieval optical apparatus in Fig. 1, according to Eq. (1). All training and testing objects are of size 256×256 , with the pixel size of the objects and on the detector being $20\mu\text{m}$. The propagation distance is set to be $z = 150\text{mm}$. In order to qualify the weak object approximation, so that the learned WOTF can be explicitly computed from the measurements and the corresponding phase objects, the maximum phase depth of the objects is kept below 0.1π rad. In simulation, the weak objects are obtained from applying a heuristic one-to-one calibration curve, say a linear one, that maps from the 8 bit alphabet $\{0, 1, \dots, 255\}$ to one that lies within 0 and 0.1π rad, as the entropy of an image is invariant under one-to-one mappings and therefore, we can have the weak object approximation and the entropy distributions above hold concurrently.

During the training stage, pairs of intensity measurements (noiseless) and weak phase objects are used as the input and output of PhENN, respectively. The disparity between the estimate produced by the current weights in PhENN with the ground truth is used to update the weights using backpropagation. After PhENN is trained, in the test stage, the measurements corresponding to unseen examples are fed into the trained PhENN to predict the estimated phase objects. The loss function used in training is the negative Pearson correlation coefficient (NPCC), which has been proved helpful for better image quality in reconstructions [9–11,18,21,25,31]. For phase object f and the corresponding estimate \hat{f} , the NPCC of f and \hat{f} is defined as

$$\text{NPCC}(\hat{f}, f) \equiv \frac{-\sum_{x,y} (\hat{f}(x,y) - \langle \hat{f} \rangle) (f(x,y) - \langle f \rangle)}{\sqrt{\sum_{x,y} (\hat{f}(x,y) - \langle \hat{f} \rangle)^2} \sqrt{\sum_{x,y} (f(x,y) - \langle f \rangle)^2}}, \quad (5)$$

where $\langle \hat{f} \rangle$ and $\langle f \rangle$ are the spatial averages of the reconstruction \hat{f} and true object f , respectively. If the reconstruction is perfect, we have $\text{NPCC}(\hat{f}, f) = -1$. One caveat of using NPCC is that a good NPCC metric cannot guarantee the reconstruction is of the correct quantitative scale, as NPCC is invariant under affine transformations, *i.e.*, $\text{NPCC}(\hat{f}, f) = \text{NPCC}(a\hat{f} + b, f)$ for scalars a and b . Therefore, to correct the quantitative scale of the reconstructions without altering the PCC value, a linear fitting step is carried out based on the validation set and the learned a, b values are used to correct the test sets. The quantitative performance of the reconstructions on produced by PhENN trained with various datasets are compared in Tables 1 and 2. The quantitative metrics chosen considered both the pixel-wise accuracy through MAE and the structural similarity through PCC.

Table 1. Cross-domain generalization ability performance of PhENN trained with various datasets. The quantitative metric used is Pearson Correlation Coefficient. The defocus distance $z = 150\text{mm}$.

	Training datasets			
	ImageNet	Face-LFW	IC-layout	MNIST
Test on ImageNet	0.936 ± 0.043	0.909 ± 0.068	0.692 ± 0.184	0.586 ± 0.200
Test on Face-LFW	0.980 ± 0.011	0.982 ± 0.010	0.700 ± 0.078	0.747 ± 0.097
Test on IC	0.911 ± 0.021	0.891 ± 0.029	0.942 ± 0.018	0.835 ± 0.060
Test on MNIST	0.988 ± 0.005	0.948 ± 0.051	0.846 ± 0.111	0.9998 ± 0.0002
Training set entropy (bits)	5.524 ± 0.561	5.513 ± 0.372	2.753 ± 0.330	0.904 ± 0.159

Table 2. Cross-domain generalization ability performance of PhENN trained with various datasets. The quantitative metric used is Mean Absolute Error (MAE). The defocus distance $z = 150\text{mm}$.

	Training datasets			
	ImageNet	Face-LFW	IC-layout	MNIST
Test on ImageNet	0.033 ± 0.023	0.034 ± 0.022	0.048 ± 0.021	0.055 ± 0.021
Test on Face-LFW	0.013 ± 0.003	0.011 ± 0.002	0.043 ± 0.008	0.042 ± 0.008
Test on IC	0.021 ± 0.005	0.022 ± 0.004	0.016 ± 0.004	0.028 ± 0.006
Test on MNIST	0.005 ± 0.001	0.008 ± 0.001	0.012 ± 0.001	0.001 ± 0.0004
Training set entropy (bits)	5.524 ± 0.561	5.513 ± 0.372	2.753 ± 0.330	0.904 ± 0.159

From Tables 1 and 2, we make the following interesting observations:

- It is always ideal to train with examples that are in the same class with the test examples. For all test datasets, same-domain generalization outperforms cross-domain generalization. Among same-domain generalization, more constrained prior information (training dataset with lower entropy), and thus the stronger prior information, is beneficial, since it more strongly regulates the reconstructions to be similar to those examples from training, and fortunately, also those from the test.
- When access to training data in the same database as the test data is not possible, if the model is trained on higher-entropy datasets, *e.g.* ImageNet and Face-LFW, the cross-domain generalization to other datasets is satisfactory; however, the cross-domain generalization performance is poor, or even catastrophic, if the model is trained on a lower-entropy dataset, *e.g.* IC layout, MNIST. From this we conclude that, even though undoubtedly the spatial correlation structure of the datasets plays some role, the Shannon entropy of the training database has a strong effect in cross-domain generalization performance.

More interesting observations are available in Fig. 4, where we show representative examples of predicting objects in various classes when PhENN is trained with all four datasets, respectively. From these examples, the regularization effect imposed by the MNIST and IC layout databases is clear: PhENN has been forced to inherit the piece-wise constant features from the MNIST and IC examples and pass them to the reconstructed ImageNet, Face-LFW examples which caused distortions of different extent. It is also noteworthy that the IC objects share with the MNIST objects on the piecewise constant features (though differ from the latter on the sparsity priors), making the degradation caused by the regularization effect imposed by MNIST dataset to reconstructed ICs much less severe than that to ImageNet and Face-LFW [43] objects.

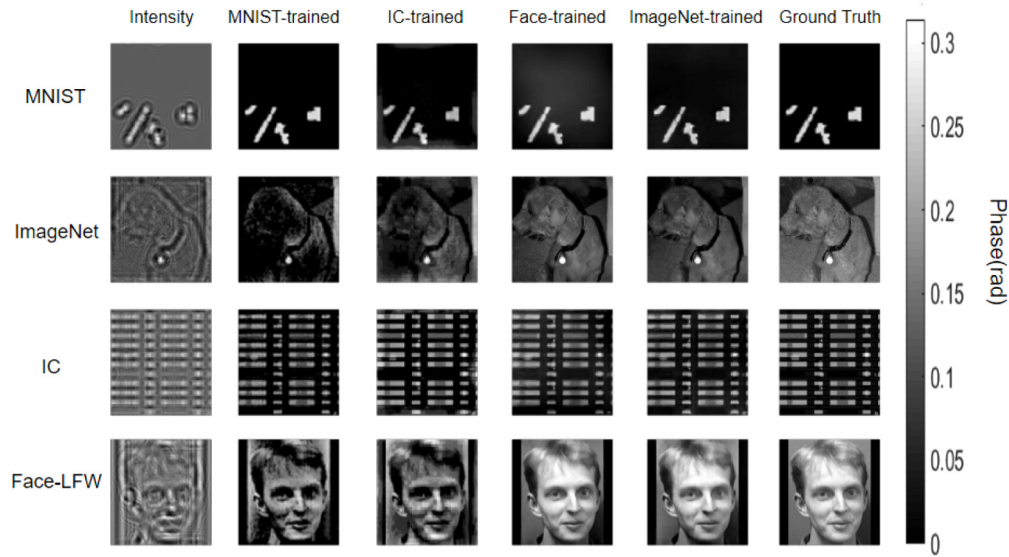


Fig. 4. Cross-domain generalization performance on synthetic data, when PhENN is trained with MNIST, IC, Face-LFW and ImageNet, respectively. The defocus distance $z = 150\text{mm}$.

To further investigate the role of Shannon entropy in determining the cross-domain generalization performance of the DNNs, we compare the cross-domain generalization performance when PhENN is trained with the original ImageNet, and two Saturated ImageNet (SImageNet) datasets, *i.e.* SImageNet-1.5 and SImageNet-2.0, where in SImageNet- q , we augment each pixel value by the factor q and clip the values at the upper bound 255. (*e.g.*, if the pixel value in

the original ImageNet is 220 and $q = 1.5$, then now the corresponding pixel value becomes $\min\{255, 1.5 \times 220\} = 255$. By doing this, the entropy of images in SImageNet is reduced as q increases. Subsequently, the same linear transformation (so that the entropy is preserved) is applied to map the SImageNet objects to weak phase objects with maximum phase depth 0.1π and the diffraction patterns at $z = 150\text{mm}$ are recorded. PhENN is trained to map each pair of intensity patterns to the corresponding weak objects and the cross-domain generalization performance under each training dataset is investigated. From Table 3, we can find that the cross-domain generalization performance does have a strong correlation with the Shannon entropy of the training dataset. Therefore, from now on, we will only present results obtained when PhENN is trained with ImageNet and MNIST, each representing the cases where a high-entropy and low-entropy database is used for training.

Table 3. Cross-domain generalization ability performance of PhENN trained with ImageNet, SImageNet-1.5 and SImageNet-2.0, respectively, for $z = 150\text{ mm}$ (synthetic data).

	Average PCC \pm std.dev		
	Train on ImageNet	Train on SImageNet-1.5	Train on SImageNet-2.0
Test on ImageNet	0.936 \pm 0.043	0.934 \pm 0.046	0.879 \pm 0.089
Test on MNIST	0.988 \pm 0.005	0.967 \pm 0.030	0.901 \pm 0.079
Test on IC	0.911 \pm 0.021	0.897 \pm 0.025	0.836 \pm 0.039
Test on Face-LFW	0.980 \pm 0.011	0.968 \pm 0.018	0.961 \pm 0.018
Training set entropy (bits)	5.524 \pm 0.561	5.080 \pm 0.949	4.102 \pm 1.255

3.2. How well has PhENN learned the physics model?

To quantitatively verify that ImageNet-trained PhENN learns the underlying physics law better than MNIST-trained PhENN, we compare the *learned* WOTF (LWOTF) of the ImageNet-trained PhENN and MNIST-trained PhENN.

Once the network has been trained, based on a test set of K test images, the LWOTF is computed as,

$$\text{LWOTF} = \frac{1}{K} \sum_{k=1}^K \frac{G_k(u, v) - \delta(u, v)}{\hat{F}_k(u, v)}, \quad (6)$$

where $G_k(u, v)$ and $\hat{F}_k(u, v)$ are the Fourier transforms of the intensity measurement $g_k(x, y)$ and the network's estimated phase $\hat{f}_k(x, y)$ for the k th testing object, respectively. For better generality, we split the test set of $K = 100$ images into four equally large subsets, that is, 25 test images from the ImageNet, MNIST, Face-LFW and IC layouts each. We denote the LWOTF of the ImageNet-trained PhENN and MNIST-trained PhENN as LWOTF-ImageNet and LWOTF-MNIST, respectively.

In Fig. 5, we show the 1D cross-sections along the diagonal directions of the WOTF (computed from the ground truth examples), LWOTF-ImageNet and LWOTF-MNIST, respectively. Also, we plot the theoretical WOTF $\sin(2\lambda z(u^2 + v^2))$, denoted WOTF-theory, under the same sampling rate as the detection process. WOTF-theory is indistinguishable from WOTF-computed, indicating that the weak object approximation holds well. For better visualization, values are cropped to the range $[-3, 3]$ (the values cropped out are outliers, all from LWOTF-MNIST). We find that ImageNet-trained PhENN indeed learned the WOTF better than the MNIST-trained PhENN. Also, we note the mismatch between the LWOTF-ImageNet and the WOTF-theory becomes larger at higher spatial frequencies, which is due to the under-representation of the high spatial frequencies in the reconstructions that has been extensively argued in [18,32]. In this paper, we choose not to overcome this limitation by the LS-DNN technique [7,18], so that the choice of training dataset is the only difference between ImageNet-PhENN and MNIST-PhENN.

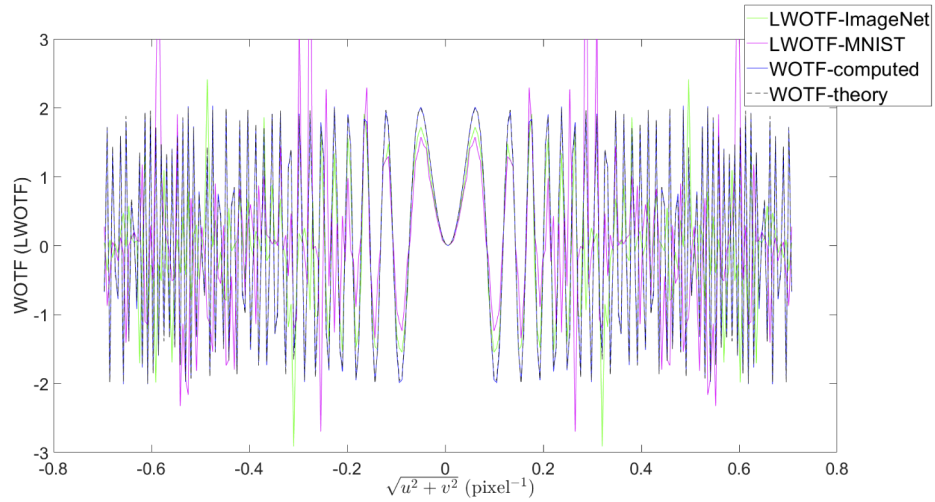


Fig. 5. Comparison of LWOTFs-ImageNet, LWOTF-MNIST, WOTF-computed and WOTF-theory. The values are cropped to the range of $[-3, 3]$ and outliers from LWOTF-MNIST have been cropped out.

Besides the direct WOTF comparison, we propose an alternative study where we verify that ImageNet-trained PhENN learns the propagation model better than the MNIST-trained PhENN. From Eq. (2), we see that there exist several nulls (the locations where the values of the transfer function are equal to zero) in the weak object transfer function and at those nulls, the sign of the transfer function switches, introducing a phase delay of π rad in the spatial frequency domain. As a result, the measured pattern at the detector plane will shift by half period in the spatial domain at those frequencies. We refer this phenomenon as the “phase shift effect”. Because of that, when we image a star-like binary weak phase object with P periods, the fringes in the measurement will become discontinuous (see Fig. 9 later). In particular, for a defocus distance z (in mm), the radii of discontinuity r_k for $k = 1, 2, \dots$, and the associated spatial frequency in (u_k, v_k) jointly satisfy

$$\lambda z(u_k^2 + v_k^2) = \lambda z\left(\frac{P}{2\pi r_k}\right)^2 = k. \quad (7)$$

If PhENN were doing something trivial, e.g. edge sharpening, it would be failing to catch the phase shifts due to (7). Therefore, the star pattern test also provides a way to test whether the physical model has been correctly incorporated. In Section 4.3 we show experimental results verifying that indeed ImageNet-trained PhENN incorporates the physics whereas MNIST-trained PhENN does not.

4. Experimental results

4.1. Optical Apparatus

The optical configuration of the experiments is shown in Fig. 6. Polarization angles of two linear polarizers (POL1 and POL2) were carefully chosen to minimize the maximum phase modulation depth of SLM (Holoeye LC-R 720) down to $\sim 0.1\pi$ (see Fig. 7). For the calibration process, we designed a Michelson-Young interferometer with the SLM, displaying a pattern with two values: one is the reference, fixed at 0, and the other varies from 0 to 255. This way we could calculate the phase shift by the SLM for each 8-bit pixel value. We captured several frames for each 8-bit value and took the average to suppress any artifacts (denoted as “Mean” in Fig. 7). The low-pass filter was applied upon the mean curve (“Low-passed mean” in Fig. 7) to further

reduce fluctuation in the curve because of very weak phase modulation being offered by the SLM, and the standard deviation was also computed along the mean (“Mean $\pm 1\sigma$ ” in Fig. 7). A $4f$ telescope of two lens ($f = 100$ mm and $f = 60$ mm) was used to transfer the image plane from SLM matching two different pixel pitches of the SLM and CMOS camera, followed by a defocus of $z = 150$ mm was given for capturing diffraction patterns with the camera. The diffraction patterns suffer from the normal contamination from the camera. Each experimental diffraction pattern was iteratively registered with the simulated one by applying affine transformation on the experimental one to the direction of maximizing NMI (Normalized Mutual Information) between the two using Nelder-Mead method [44,45]. More details on the pre-processing process is available in Appendix D.

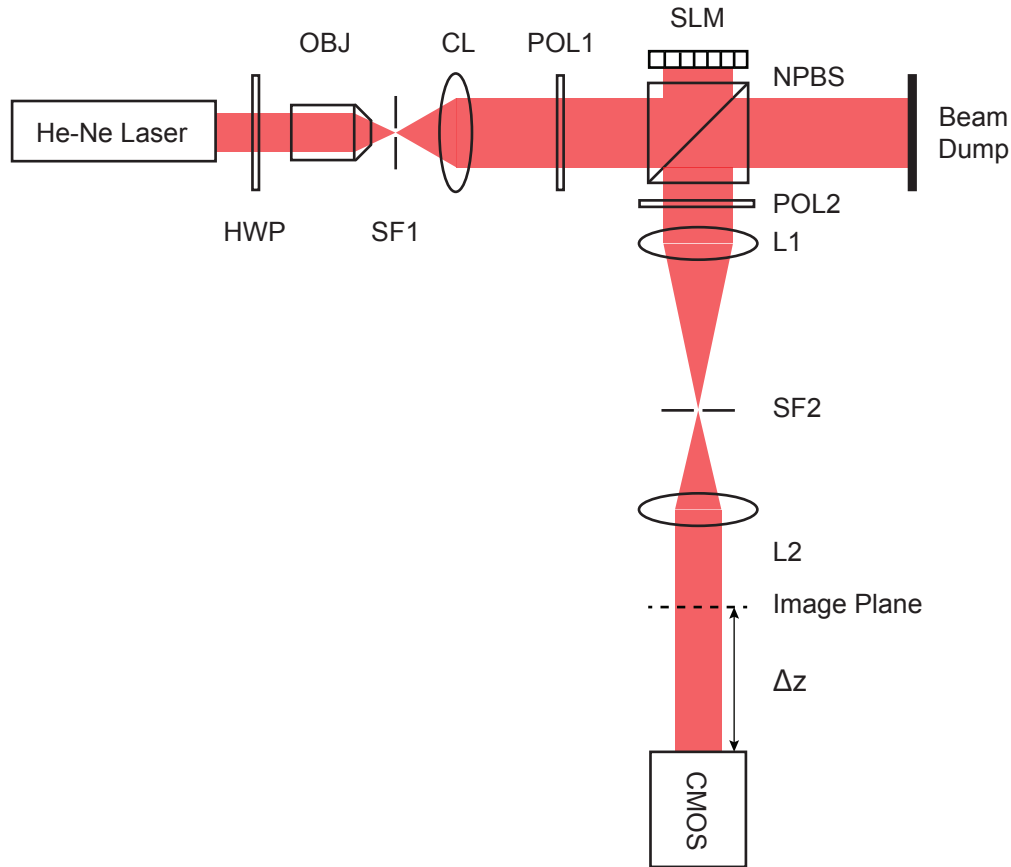


Fig. 6. Optical apparatus. HWP: Half-wave plate, OBJ: objective lens, SF: spatial filter, CL: collimating lens, POL: linear polarizer, SLM: spatial light modulator, NPBS: non-polarizing beamsplitter, L: lens.

4.2. Comparisons of the cross-domain generalization performance

Cross-domain generalization performance based on experimental data ($z = 150$ mm) for ImageNet-trained PhENN and MNIST-trained PhENN is shown in Fig. 8 and quantitatively compared in Table 4, from which we make similar observations as those from the synthetic data: cross-domain generalization performance of ImageNet-trained PhENN is generally good; however, such performance of MNIST-trained PhENN is poor. Moreover, the estimated phase objects produced by MNIST-trained PhENN display sparse or binarized features, resembling MNIST examples,

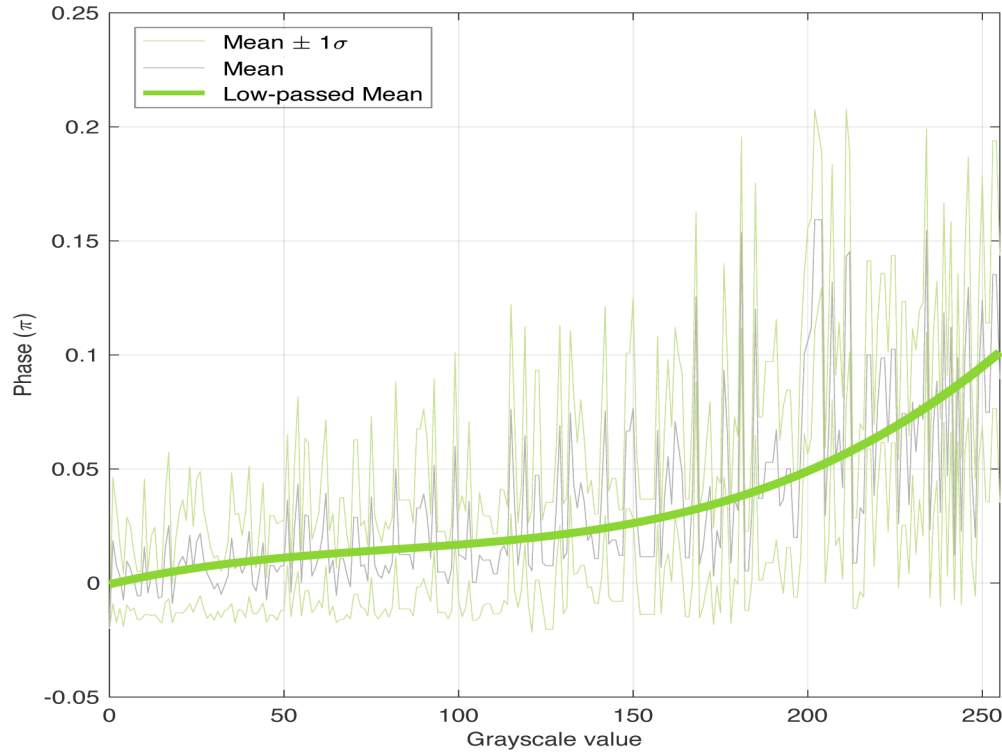


Fig. 7. Phase modulation vs. 8-bit grayscale value for the experiments.

which suggests that during training, MNIST-trained PhENN was encouraged to memorize the MNIST examples, whereas ImageNet-trained PhENN, without strong regularization effect imposed by its training examples, was able to learn the actual physics model better. Additional examples are shown in Fig. 12 in Appendix C.

Table 4. Cross-domain generalization ability performance of PhENN trained with ImageNet and MNIST, respectively, for $z = 150$ mm (experimental data).

	Average PCC \pm std.dev		Average MAE \pm std.dev	
	Train on ImageNet	Train on MNIST	Train on ImageNet	Train on MNIST
Test on ImageNet	0.908 ± 0.061	0.645 ± 0.233	0.008 ± 0.003	0.013 ± 0.006
Test on MNIST	0.956 ± 0.005	0.986 ± 0.007	0.005 ± 0.001	0.001 ± 0.0003
Test on IC	0.833 ± 0.035	0.736 ± 0.055	0.009 ± 0.001	0.012 ± 0.002
Test on Face-LFW	0.951 ± 0.012	0.805 ± 0.056	0.006 ± 0.001	0.011 ± 0.002

4.3. Star-pattern experiment to demonstrate the learning of the propagation model

From Eq. (7), for our experimental parameters $z = 150$ mm, $\lambda = 633$ nm, $P = 50$, the 2nd ($k = 2$) and 3rd ($k = 3$) discontinuity are at $0.0032 \mu\text{m}^{-1}$ and $0.0040 \mu\text{m}^{-1}$, respectively. From the measurement in Fig. 9(a), we see the two discontinuities are located at $r_2 \approx 2.44$ mm (red) and $r_3 \approx 2.16$ mm (blue), respectively, corresponding to spatial frequencies, $0.0033 \mu\text{m}^{-1}$ and $0.0037 \mu\text{m}^{-1}$, matching well with theoretical values, and indicating that the weak object approximation holds well. After PhENN was trained with ImageNet, a dataset drastically different from the appearance of the star-pattern, such discontinuity was corrected perfectly

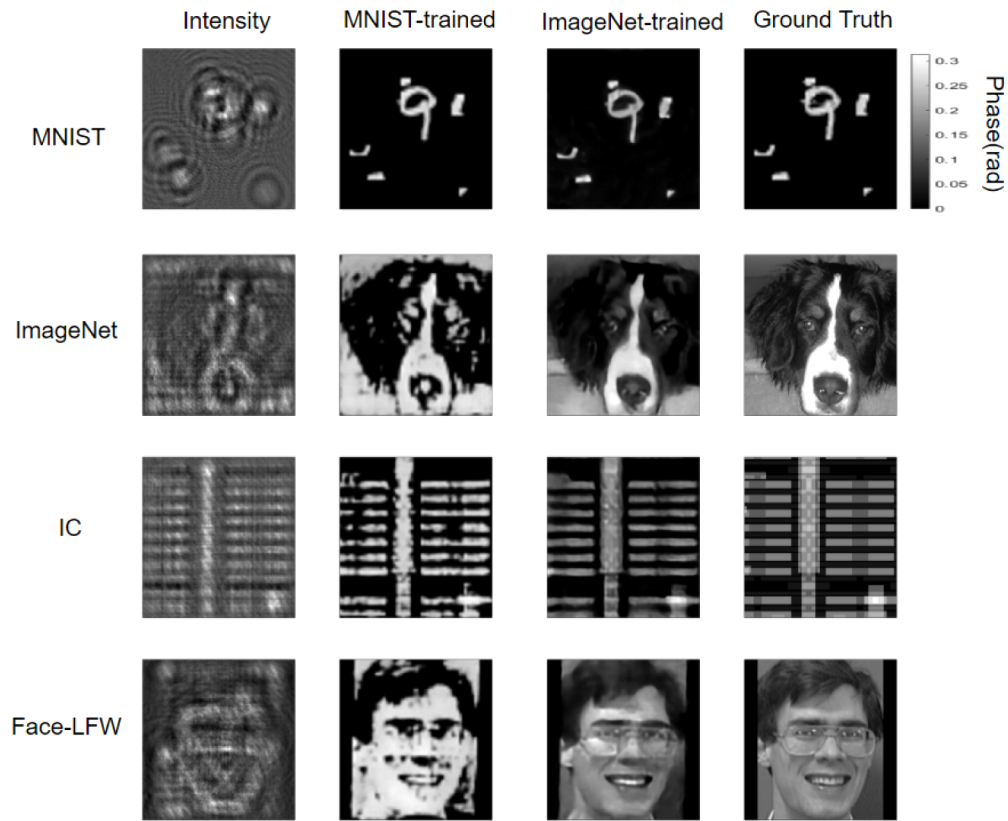


Fig. 8. Cross-domain generalization performances of ImageNet-PhENN and MNIST-PhENN on experimental data. The defocus distance is $z = 150\text{mm}$.

in the reconstruction (Fig. 9(c)), indicating that ImageNet-trained PhENN has learned the underlying physics (or the WOTF) while the MNIST-PhENN apparently failed (Fig. 9(d)). It is also noteworthy that there is still significant deficiency of high frequencies in the reconstruction of ImageNet-trained PhENN, which corroborates our observation earlier in Fig. 5 that even ImageNet-trained PhENN is not able to restore high frequencies very well. Using Learning-to-synthesize by DNN (LS-DNN) method [7,18] has been proven very efficient to tackle this issue [46], but we choose not to pursue it as it would deviate from the main emphasis of this paper.

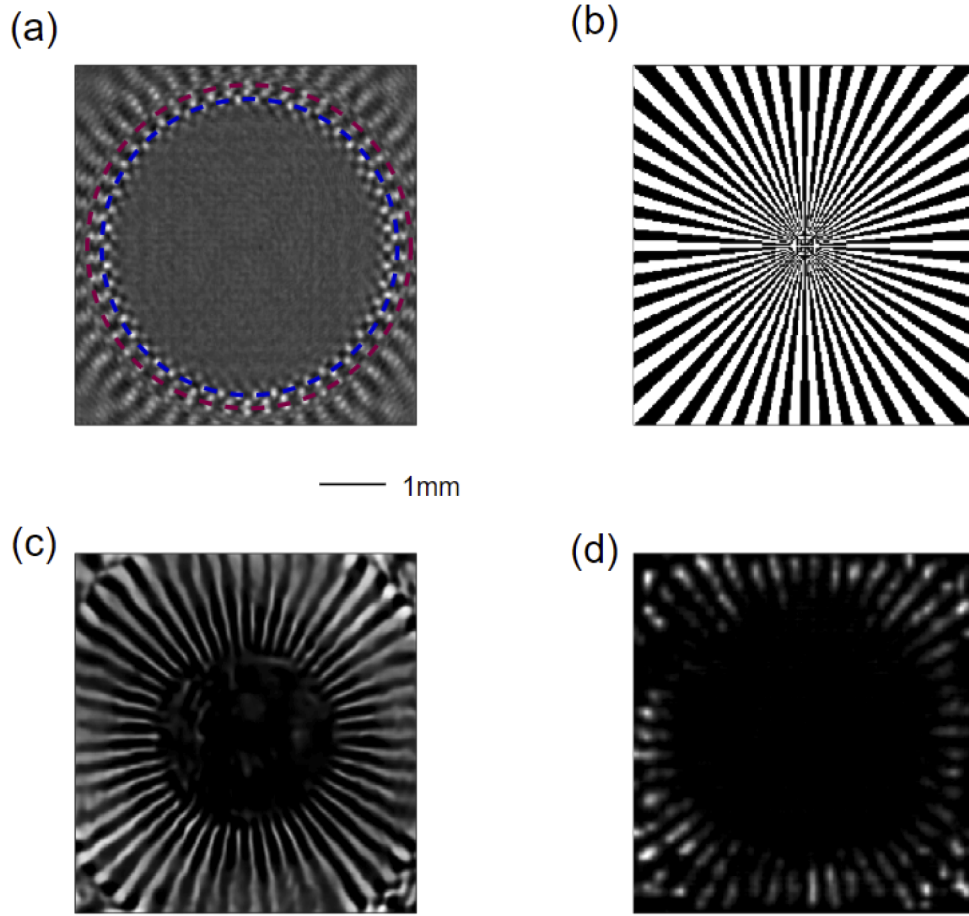


Fig. 9. Reconstruction of the star pattern. (a). Intensity measurement at $z = 150\text{mm}$. (b). star pattern (weak) object. (c). reconstruction by ImageNet-trained PhENN. (d). reconstruction by MNIST-trained PhENN.

5. Conclusions

In this paper, we used PhENN for lensless phase imaging as the standard platform to address the important question of DNN generalization performance when training cannot be performed in the intended class. This is motivated by the problem of insufficient training data, especially when such data need to be collected experimentally. We anticipate that this work will offer practitioners a way to efficiently train their machine learning architectures, by choosing the publicly available standardized dataset without worrying about cross-domain generalization performance.

Our work is suggestive of certain interesting directions for future investigation. A particularly intriguing one is to refine the bound of the (cross-domain) generalization error by incorporating the distance metric of the empirical distributions between the training and the test set, along with other factors that have been considered currently in the literature (recall Section 1.4). Moreover, though the chain of logic presented in the paper was centered on phase retrieval, it should be applicable to other domains of computational imaging subject to further study and verification.

Appendix A: More details of PhENN and the training specifics

In Section 1.3, we introduced the high-level architecture of PhENN (Fig. 2). In this Section, we provide in Fig. 10 details of the layer-wise architecture of each functional block in PhENN.

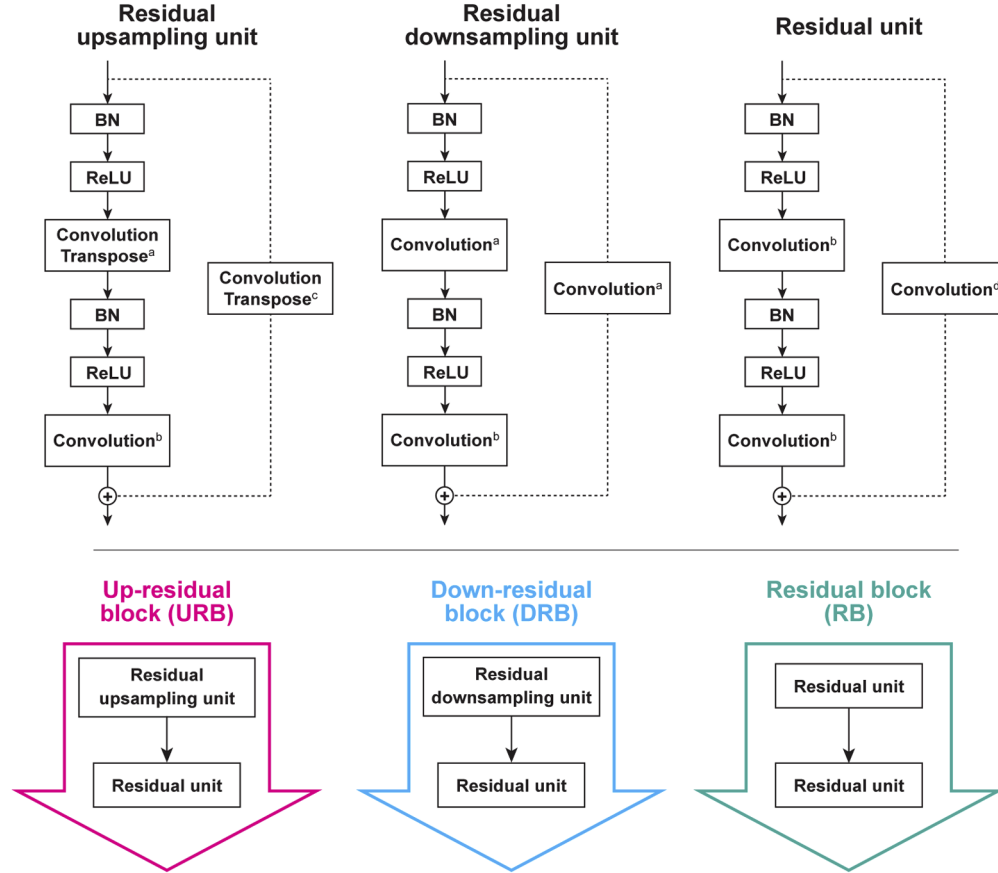


Fig. 10. More detailed architecture of PhENN. Superscripts a - d denote different kernel size and strides, listed as follows: a) Kernel size: (3, 3), strides: (2, 2). b) Kernel size: (3, 3), strides: (1, 1). c) Kernel size: (2, 2), strides: (2, 2). d) Kernel size: (1, 1), strides: (1, 1).

In PhENN, major functional blocks include Up-residual blocks (URBs), Down-residual blocks (DRBs) and Residual blocks (RBs). All convolutional (Conv2D) and convolutional transpose (Conv2DTranspose) kernels are 3×3 , except for the 2×2 kernels in the side-branch Convolutional Transpose in Residual upsampling units and the 1×1 (1D convolution) kernels in the side-branch of Residual units.

The simulation is conducted on a Nvidia GTX1080 GPU using the open source machine learning Platform TensorFlow. The Adam optimizer [47], with learning rate being 0.001 and exponential decay rate for the first and second moment estimates being 0.9 and 0.999 ($\beta_1=0.9$, $\beta_2=0.999$). The batch size is 5. The training of PhENN for 50 epochs, which is sufficient for PhENN to converge, takes about 2 hours.

Appendix B: More results on synthetic data

In this section, we show reconstruction examples with synthetic data at $z = 100\text{mm}$ in Fig. 11, and the quantitative metrics in Table 5, where we see MNIST-trained PhENN produced reconstructions

that are generally sparsified, as MNIST has imposed too strong regularization effect to the training and that gets passed onto the reconstructions. Superior cross-domain generalization performance of ImageNet-trained PhENN is again verified.

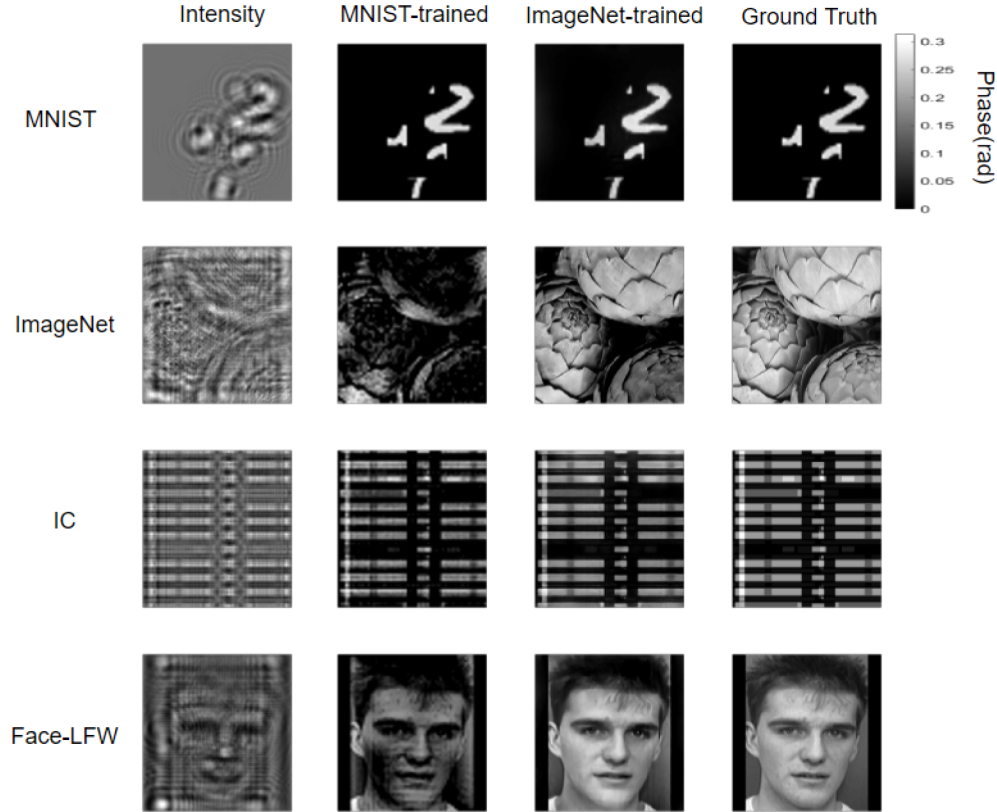


Fig. 11. Cross-domain generalization performances of ImageNet-PhENN and MNIST-PhENN on synthetic data. The defocus distance is $z = 100\text{mm}$.

Table 5. Cross-domain generalization ability performance of PhENN trained with ImageNet and MNIST, respectively, for $z = 100\text{ mm}$ (synthetic data).

	Average PCC \pm std.dev		Average MAE \pm std.dev	
	Train on ImageNet	Train on MNIST	Train on ImageNet	Train on MNIST
Test on ImageNet	0.932 ± 0.068	0.578 ± 0.213	0.032 ± 0.024	0.055 ± 0.021
Test on MNIST	0.964 ± 0.039	0.9998 ± 0.0002	0.006 ± 0.001	0.0008 ± 0.0004
Test on IC	0.942 ± 0.014	0.866 ± 0.051	0.018 ± 0.004	0.024 ± 0.006
Test on Face-LFW	0.969 ± 0.015	0.758 ± 0.104	0.016 ± 0.004	0.024 ± 0.008

Appendix C: More results on experimental data

In Fig. 12, we provide additional reconstructions of various classes of objects by ImageNet-trained PhENN and MNIST-trained PhENN ($z = 150\text{mm}$), based on experimental data. Consistent with previous observations, we see significant distortions in the reconstructions of non-MNIST objects by MNIST-trained PhENN. However, no significant distortions can be seen from reconstructions of non-ImageNet objects produced by ImageNet-trained PhENN.

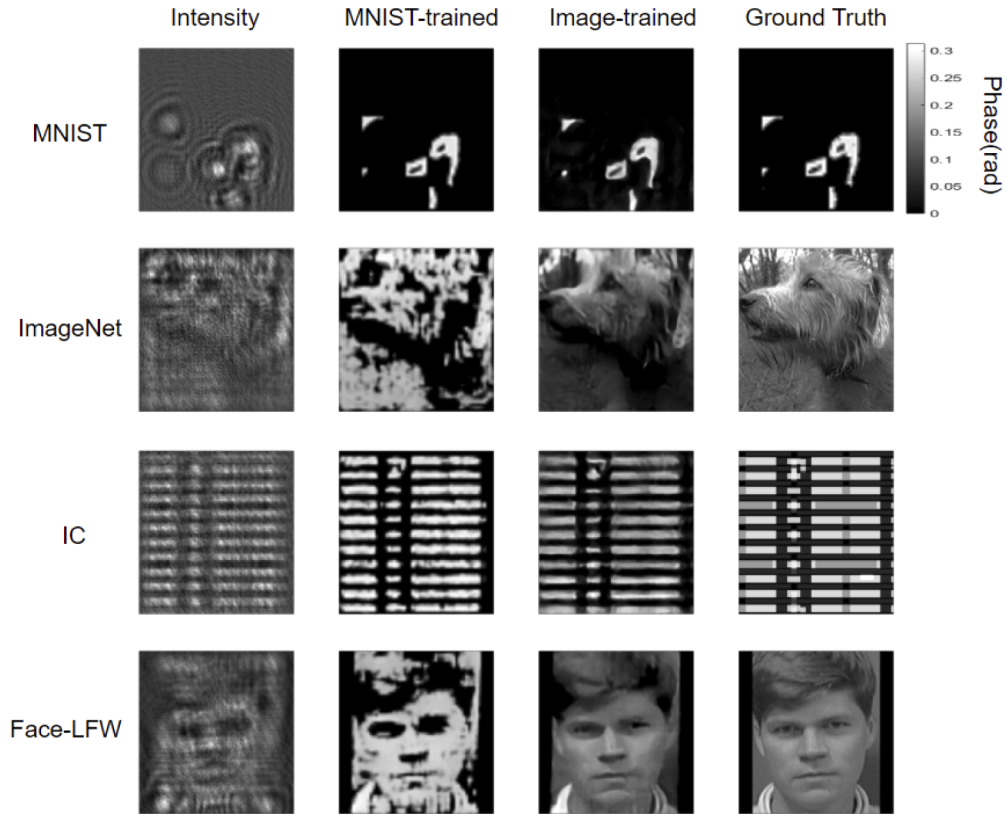


Fig. 12. More examples of cross-domain generalization performance on experimental data. The defocus distance $z = 150\text{mm}$.

Appendix D: Details on experimental data preprocessing

Two linear polarizers were used to achieve the maximum phase depth of the reflective SLM of $\sim 0.1\pi$. The training process with raw intensity measurements was preceded by the image registration on experimental diffraction patterns to match their shape with their corresponding simulated ones to resolve slight shape deformation during the experimental data acquisition, if any.

Using the calibration curve between 8-bit grayscale values and phase modulation depth in radian, input phase objects were simulated, followed by the generation of simulated intensity measurements without any deformation. Experimental measurements should match with the simulated ones if ideal. Image registration process finds an optimal affine transformation that brings an experimental measurement to its corresponding simulated one according to the optimization using Nelder-Mead method to the direction of minimizing the negative NMI (Normalized Mutual Information). Optimal affine transformation matrix was found for each dataset and applied accordingly.

Then, only center 256×256 pixels in each preprocessed intensity measurement were cropped to generate training, validation, and testing datasets, paired with labels or ground truth images. Thanks to the $4f$ system in the optical apparatus that matches the pixel size of the reflective SLM with that of the CMOS camera, only the same dimension of pixels as that of each displayed image on the SLM is needed.

Funding

Intelligence Advanced Research Projects Activity (FA8650-17-C-9113); National Research Foundation Singapore (015824); Korea Foundation for Advanced Studies.

Disclosures

The authors declare no conflicts of interest.

References

1. C. Dong, C. Loy, K. He, and X. Tang, "Learning a deep convolutional neural network for image super-resolution," in *European Conference on Computer Vision (ECCV) / Lecture Notes on Computer Science Part IV*, vol. 8692, (2014), pp. 184–199.
2. J. Johnson, A. Alahi, and Li Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *European Conference on Computer Vision (ECCV) / Lecture Notes on Computer Science*, vol. 9906 B. Leide, J. Matas, N. Sebe, and M. Welling, eds. (2016), pp. 694–711.
3. C. Ledig, L. Theis, F. Huczar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Zehan Wang, and Wenshe Shi, "Photo-realistic single image super-resolution using a Generative Adversarial Network," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (2017), pp. 4681–4690.
4. Y. Rivenson, Z. Gorocs, H. Gunaydin, Yibo Zhang, Hongda Wang, and A. Ozcan, "Deep learning microscopy," *Optica* **4**(11), 1437–1443 (2017).
5. H. Wang, Y. Rivenson, Z. Wei, H. Gunaydin, L. Bentolila, and A. Ozcan, "Deep learning achieves super-resolution in fluorescence microscopy," *bioRxiv*, <https://doi.org/10.1101/309641> (2018).
6. E. Nehme, L. E. Weiss, T. Michaeli, and Y. Shechtman, "Deep-STORM: super-resolution single-molecule microscopy by deep learning," *Optica* **5**(4), 458–464 (2018).
7. M. Deng, S. Li, and G. Barbastathis, "Learning to synthesize: splitting and recombining low and high spatial frequencies for image recovery," *arXiv preprint arXiv:1811.07945* (2018).
8. A. Sinha, Justin Lee, Shuai Li, and G. Barbastathis, "Lensless computational imaging through deep learning," *Optica* **4**(9), 1117–1125 (2017).
9. Y. Rivenson, Y. Zhang, H. Günaydin, D. Teng, and A. Ozcan, "Phase recovery and holographic image reconstruction using deep learning in neural networks," *Light: Sci. Appl.* **7**(2), 17141 (2018).
10. A. Goy, K. Arthur, Shuai Li, and G. Barbastathis, "Low photon count phase retrieval using deep learning," *Phys. Rev. Lett.* **121**(24), 243902 (2018).
11. T. Nguyen, Y. Xue, Y. Li, L. Tian, and G. Nehmetallah, "Deep learning approach for fourier ptychography microscopy," *Opt. Express* **26**(20), 26470–26484 (2018).
12. Ç. Işıl, F. S. Oktem, and A. Koç, "Deep iterative reconstruction for phase retrieval," *Appl. Opt.* **58**(20), 5422–5431 (2019).
13. H. Wang, M. Lyu, and G. Situ, "eholonet: a learning-based point-to-point approach for in-line digital holographic reconstruction," *Opt. Express* **26**(18), 22603–22614 (2018).
14. T. Pitkäaho, A. Manninen, and T. J. Naughton, "Performance of autofocus capability of deep convolutional neural networks in digital holographic microscopy," in *Digital Holography and Three-Dimensional Imaging* (OSA, 2017), p. W2A.5.
15. Y. Wu, Y. Rivenson, Y. Zhang, Z. Wei, H. Gunaydin, X. Lin, and A. Ozcan, "Extended depth-of-field in holographic image reconstruction using deep learning based auto-focusing and phase-recovery," *Optica* **5**(6), 704–710 (2018).
16. M. R. Kellman, E. Bostan, N. A. Repina, and L. Waller, "Physics-based learned design: Optimized coded-illumination for quantitative phase imaging," *IEEE Trans. Comput. Imaging* **5**(3), 344–353 (2019).
17. Z. Ren, Z. Xu, and E. Y. Lam, "Learning-based nonparametric autofocusing for digital holography," *Optica* **5**(4), 337–344 (2018).
18. M. Deng, S. Li, A. Goy, I. Kang, and G. Barbastathis, "Learning to synthesize: Robust phase retrieval at low photon counts," *Light: Sci. Appl.* **9**(1), 36 (2020).
19. M. Deng, A. Goy, S. Li, K. Arthur, and G. Barbastathis, "Probing shallower: perceptual loss trained phase extraction neural network (plt-phenn) for artifact-free reconstruction at low photon budget," *Opt. Express* **28**(2), 2511–2535 (2020).
20. R. Horisaki, R. Takagi, and J. Tanida, "Learning-based imaging through scattering media," *Opt. Express* **24**(13), 13738–13743 (2016).
21. S. Li, M. Deng, J. Lee, A. Sinha, and G. Barbastathis, "Imaging through glass diffusers using densely connected convolutional networks," *Optica* **5**(7), 803–813 (2018).
22. Y. Li, Y. Xue, and L. Tian, "Deep speckle correlation: a deep learning approach toward scalable imaging through scattering media," *Optica* **5**(10), 1181–1190 (2018).
23. U. S. Kamilov, I. N. Papadopoulos, M. H. Shoreh, A. Goy, C. Vonesch, M. Unser, and D. Psaltis, "Learning approach to optical tomography," *Optica* **2**(6), 517–522 (2015).

24. U. S. Kamilov, I. N. Papadopoulos, M. H. Shoreh, A. Goy, C. Vonesch, M. Unser, and D. Psaltis, "Optical tomographic image reconstruction based on beam propagation and sparse regularization," *IEEE Trans. Comput. Imaging* **2**(1), 59–70 (2016).
25. A. Goy, G. Rughoobur, Shuai Li, K. Arthur, A. Akinwande, and G. Barbastathis, "High-resolution limited-angle phase tomography of dense layered objects using deep neural networks," *Proc. Nat. Acad. Sci.* ((accepted) 2019).
26. G. Barbastathis, A. Ozcan, and Guohai Situ, "On the use of deep learning for computational imaging," *Optica* (2019).
27. M. T. McCann, K. H. Jin, and M. Unser, "Convolutional neural networks for inverse problems in imaging: A review," *IEEE Signal Process. Mag.* **34**(6), 85–95 (2017).
28. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database, in *2009 IEEE conference on computer vision and pattern recognition* (Ieee, 2009), pp. 248–255.
29. Y. LeCun, C. Cortes, and C. J. Burges, "MNIST handwritten digit database," AT&T Labs [Online]. Available: <http://yann.lecun.com/exdb/mnist> **2** (2010).
30. L. Tian and L. Waller, "Quantitative differential phase contrast imaging in an led array microscope," *Opt. Express* **23**(9), 11394–11403 (2015).
31. S. Li, G. Barbastathis, and A. Goy, "Analysis of phase-extraction neural network (phenn) performance for lensless quantitative phase imaging," in *Quantitative Phase Imaging V*, vol. 10887 (International Society for Optics and Photonics, 2019), p. 108870T.
32. S. Li and G. Barbastathis, "Spectral pre-modulation of training examples enhances the spatial resolution of the phase extraction neural network (PhENN)," *Opt. Express* **26**(22), 29340–29352 (2018).
33. B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro, "Exploring generalization in deep learning," in *Advances in Neural Information Processing Systems*, (2017), pp. 5947–5956.
34. B. Neyshabur, Z. Li, S. Bhojanapalli, Y. LeCun, and N. Srebro, "Towards understanding the role of over-parametrization in generalization of neural networks," *arXiv preprint arXiv:1805.12076* (2018).
35. B. Neyshabur, S. Bhojanapalli, and N. Srebro, "A pac-bayesian approach to spectrally-normalized margin bounds for neural networks," *arXiv preprint arXiv:1707.09564* (2017).
36. C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning requires rethinking generalization," *arXiv preprint arXiv:1611.03530* (2016).
37. M. S. Advani and A. M. Saxe, "High-dimensional dynamics of generalization error in neural networks," *arXiv preprint arXiv:1710.03667* (2017).
38. H. Xu and S. Mannor, "Robustness and generalization," *Mach. Learn.* **86**(3), 391–423 (2012).
39. D. Jakubovitz, R. Giryes, and M. R. Rodrigues, "Generalization error in deep learning," in *Compressed Sensing and Its Applications* (Springer, 2019), pp. 153–193.
40. C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.* **27**(3), 379–423 (1948).
41. T. M. Cover and J. A. Thomas, *Elements of information theory* (John Wiley & Sons, 2012).
42. V. K. Ingle and J. G. Proakis, *Digital signal processing using matlab: a problem solving companion* (Cengage Learning, 2016).
43. G. B. Huang, M. Mattar, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," in *Technical Report, University of Massachusetts*, (2007).
44. G. K. Matsopoulos, N. A. Mouravliansky, K. K. Delibasis, and K. S. Nikita, "Automatic retinal image registration scheme using global optimization techniques," *IEEE Trans. Inform. Technol. Biomed.* **3**(1), 47–60 (1999).
45. J. A. Nelder and R. Mead, "A simplex method for function minimization," *The computer journal* **7**(4), 308–313 (1965).
46. S. Li, "Computational imaging through deep learning," Ph.D. thesis, MIT (2019).
47. D. P. Kingma and J. Lei Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, (2015).