Neural networks for inverse control with system priors

Anonymous Author(s)
Affiliation
Address
email

Abstract

Controlling a physical system to behave in a desired manner requires a prior knowledge of the system, having access to a model or a collection of labeled data consisting of the desired inputs-outputs of the system. When the system is unknown or labeled data is not available or expensive to acquire, resorting to approaches that do not rely on the use of training data is inevitable. In this work, we propose an algorithm based on untrained neural networks that can be applied to a physical system in its most general form to obtain the required input that would result in a desired target output. We showcase the applicability of our algorithm to experimental phase-retrieval problems in the complex environment of a scattering medium whose input-output relation follows a nonlinear and slowly time-varying setting. We show that despite partial measurements of the system, comparable fidelity to that of fully-observed methods or supervised networks is achievable.

1 Introduction

Reconstructing the inputs of a physical system from measurements of its sensory outputs is a common practice in various disciplines such as microscopy Rivenson et al. [2017] and optical tomography Würfl et al. [2016], among others. Most learning approaches for information retrieval such as supervised learning methods, generative networks based on Generative Adversarial Networks (GANs) Mirza and Osindero [2014] or Variational Autoencoders (VAEs) Kingma and Welling [2013] and compressive sensing approaches Mousavi and Baraniuk [2017], Bora et al. [2017], Shah and Hegde [2018] rely heavily on labeled data to train deep neural networks that can faithfully recover the original inputs of the system. An interesting problem is finding inputs of the system that results in a desired target output. This is a common scenario when one is dealing with the problem of controlling a system when either the system is unknown or is too complex to be modeled, a setting in which no labeled input-output data from the distribution of target outputs might be available for supervised learning, *a priori*. In these settings, imposing a particular prior on the solutions of the inverse problem, i.e. the mapping from partially measured output to the input of the system, can encourage solutions whose resulting outputs lie within the desired part of the system's support Van Veen et al. [2018], Ulyanov et al. [2018], Heckel and Hand [2018]. Compatibility with the physics of the problem could be leveraged as a prior on the sought-after solutions.

We propose to use the system itself as a prior, i.e. to send the reconstructed inputs through the physical system or a model of the system and check if the resulting outputs agree with the original measurements. In cases where the physical system is difficult to model, the forward mapping could be *learned* as well. The learned forward and backward mappings together constitute a close loop where each one imposes a physical-compatibility prior on the other one. Here we show the applicability of the proposed algorithm for solving a real-world nonlinear problem with experimental results and compare its performance to that of the physics-based methods which require full characterization of

Submitted to 34th Conference on Neural Information Processing Systems (NeurIPS 2020). Do not distribute.

the system. It is concluded that the proposed algorithm could achieve the same level of fidelity as that of the gold-standard physics-based approach even despite partial measurements of the system. The contributions of this work are as follows:

- Assuming that a particular subsets of inputs exist that generates system's outputs belonging to a specific target distribution, we propose an algorithm for controlling the output of the system without the need to have access to labeled data from the distribution of desired target outputs. We show how one can fit the parameters of the network by using the physics-based compatibility as a prior.
- We compare the performance of the algorithm with that of classic non-learning based full-measurement methods as well as learning based approaches trained with labeled data in a supervised fashion. We show that our proposed algorithm can achieve the same level of fidelity as the other two approaches but without full-measurements of the system or use of expensive labeled data.

Related works: Untrained networks have recently been proposed to deal with scenarios in which labeled data is scarce or does not exists Ulyanov et al. [2018]. By fitting the parameters of the network, untrained methods impose a prior on the range of the network's outputs for various application such as denoising, super-resolution and inpainting. In these settings, the network is often optimized for solving single-image tasks, i.e. the network is unique to a single task and a single image, and needs to be re-optimized (often from scratch for a new image). Here, we show that our optimized network is applicable to a range of images belonging to a similar distribution. We also show the use of physics-compatibility priors as a regularization.

2 Methods

Assume that the forward mapping of a (non-)linear physical system, can be formulated as:

$$y = f(x) + \eta \tag{1}$$

where f, $x \in \mathbb{R}^n$ (or more generally $x \in \mathbb{C}^n$) and $y \in \mathbb{R}^m$, respectively, are the forward mapping function, input and output of the system and η is the random noise.

Inverse problem Denoting the target output and its sought-after corresponding input by y^t and x^*, respectively, the inverse problem of obtaining the required input that results in the target output is formulated as:

$$\hat{\theta}_{\mathbf{A}} = \underset{\theta_{\mathbf{A}}}{\operatorname{argmin}}\ \mathcal{L}\big(M\big(\hat{x}, \theta_M\big) - y^t\big) \tag{2a}$$

$$\hat{x} = A\left(y^t, \theta_A\right) \tag{2b}$$

in which, $\hat{\theta}_{\mathbf{A}}$ are the optimal parameters of the Actor network that yields the required input of the system, i.e. $x^* = A\left(y^t, \hat{\theta}_{\mathbf{A}}\right)$, $\mathcal{L}$ is a loss function, and M is the forward path of the physical system or a learned model of it. In the latter case, the parameters $\hat{\theta}_{\mathbf{M}}$ of network M are kept fixed while solving this optimization problem.

We also compare the efficacy of the proposed algorithm with that of a neural network trained with labeled data in a supervised fashion, thus the solutions are not explicitly constrained with physics-based compatibility priors. The solutions of the latter problem are obtained from:

$$\hat{\theta}_{\mathbf{G}} = \underset{\theta_{\mathbf{G}}}{\operatorname{argmin}}\ \frac{1}{N}\sum_{i=1}^{N}\left\|G(y_i^t, \theta_{\mathbf{G}}) - x_i^t\right\|_2^2 \tag{3}$$

where G is a neural network that is trained with input-output tuples (x_i^t, y_i^t) that are directly drawn from distribution of desired outputs and their corresponding inputs. Once trained, then a sample of

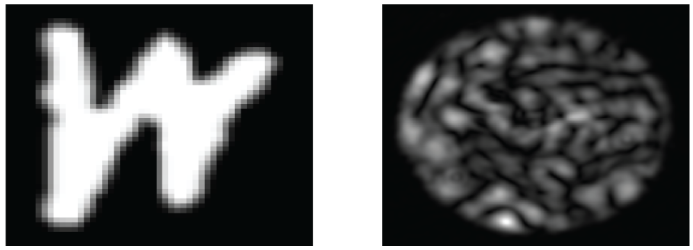

Figure 1: An example of the input (left) and partially measured (amplitude-only) output of the medium (right).

desired target output y^t from test dataset (also from the same distribution of desired targets) is given to the G network to obtain the estimated input control pattern $x^{**} = G(y^t, \hat{\theta}_{\mathbf{G}})$.

The fidelity of both algorithm 1 (A network) and algorithm 2 (G network) are calculated by sending x^* and x^{**} through the physical system and then comparing the resulting outputs.

Forward model Here we proposed our algorithm for cases where a forward model of the system is not readily available and needs to be learned. Applicability of the algorithm to cases where a forward model already exists is straightforward. With N input data points x_i^0 drawn from the distribution D (here the distribution of natural images from ImageNet) and their corresponding output measurements y_i^0, the forward path of the physical system can be modelled. In particular, the neural network *Model* is trained so that when fed with x_i^0, the output measurement y_i^0 is reproduced, i.e. $y_i^0 = M(x_i^0, \theta_{\mathbf{M}})$ where $\theta_{\mathbf{M}}$ are the parameters of the Model network. Solving the following optimization problem yields:

$$\hat{\theta}_{\mathbf{M}} = \underset{\theta_{\mathbf{M}}}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^{N} \left\| M(x_i^0, \theta_{\mathbf{M}}) - y_i^0 \right\|_2^2 \tag{4}$$

where $\hat{\theta}_{\mathbf{M}}$ are the sought-after optimal parameters of the Model network.

training The iteration of training M and A, in this order, could be repeated for multiple times each time with updated dataset for training of M. The updated dataset is comprised of the inputs x_i^j obtained from the trained Actor network in iteration j and their corresponding measurements outputs y_i^j, collected from the physical system through actual experiment. The intuition behind this is that training the Model network M with a dataset whose distribution is closer in some metrics to the distribution of the target outputs positively affects the reconstruction step.

3 Experiment and discussion

We showcase the applicability of the proposed algorithm for imaging in a diffuse medium. The latter is an example of a scattering media in which the input $x \in \mathbb{C}^{51\times51}$ is scrambled to give rise to random patterns known as speckles $y \in \mathbb{R}^{200\times200}$ (the system f, here, is effectively a random matrix). An example of an input-output of this medium is shown in Figure 1. For generating desired target outputs, we use the proposed algorithm, with Pearson coefficient Benesty et al. [2009] used as $\mathcal{L}$, to obtain the required input of the system that results in the target output. Figure 2 depicts examples of projected images through the diffuse medium. It can be inferred that our proposed approach has an fidelity on par with the supervised methods but without the need to rely on the labeled data. We used the same architecture of networks (two fully connected networks for the real and imaginary parts of the data with sigmoid non-linearity as the input x is a complex number) for both our method and the supervised learning approach. It should be noted that contrary to the untrained methods whose networks are optimized for a single image, we here optimize our network for 1000 images from EMNIST Cohen et al. [2017] dataset, simultaneously. For the sake of comparison, the network of supervised learning approach is trained with 1000 labeled dataset as well. The average fidelitiy (Pearson coefficient, PSNR) of our method and that of the supervised learning reads as

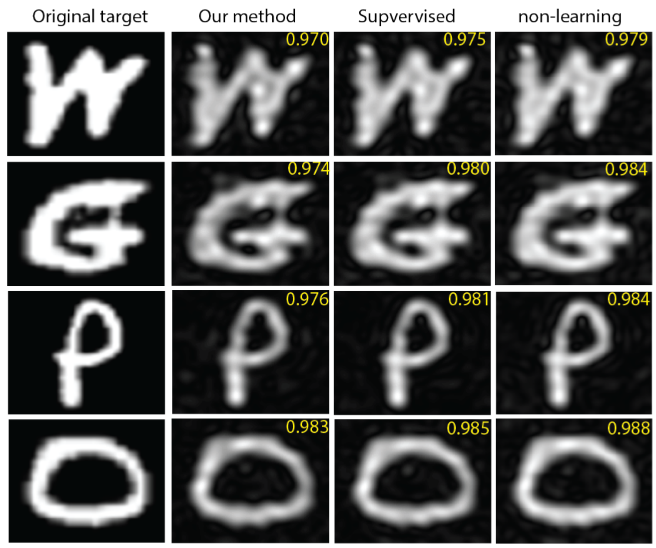

Figure 2: Qualitative results of the system output control obtained by our algorithm (second column), a neural network trained with labeled data (third column), full-measurement non learning methods (fourth column). The original targets are also shown (first column). The inset shows the fidelity (Pearson coefficient) of the generated output as compared with the original target.

$(0.9787, 20.17)$ and $(0.9823, 20.71)$, respectively. The fidelity of the full-measurement non-learning method reads as $(0.9850, 29.0486)$.

References

J. Benesty, J. Chen, Y. Huang, and I. Cohen. Pearson correlation coefficient. In *Noise reduction in speech processing*, pages 1–4. Springer, 2009.

A. Bora, A. Jalal, E. Price, and A. G. Dimakis. Compressed sensing using generative models. *arXiv preprint arXiv:1703.03208*, 2017.

G. Cohen, S. Afshar, J. Tapson, and A. Van Schaik. Emnist: Extending mnist to handwritten letters. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 2921–2926. IEEE, 2017.

R. Heckel and P. Hand. Deep decoder: Concise image representations from untrained non-convolutional networks. *arXiv preprint arXiv:1810.03982*, 2018.

D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.

A. Mousavi and R. G. Baraniuk. Learning to invert: Signal recovery via deep convolutional networks. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 2272–2276. IEEE, 2017.

Y. Rivenson, Z. Göröcs, H. Günaydin, Y. Zhang, H. Wang, and A. Ozcan. Deep learning microscopy. *Optica*, 4 (11):1437–1443, 2017.

V. Shah and C. Hegde. Solving linear inverse problems using gan priors: An algorithm with provable guarantees. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 4609–4613. IEEE, 2018.

D. Ulyanov, A. Vedaldi, and V. Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9446–9454, 2018.

D. Van Veen, A. Jalal, M. Soltanolkotabi, E. Price, S. Vishwanath, and A. G. Dimakis. Compressed sensing with deep image prior and learned regularization. *arXiv preprint arXiv:1806.06438*, 2018.

T. Würfl, F. C. Ghesu, V. Christlein, and A. Maier. Deep learning computed tomography. In *International conference on medical image computing and computer-assisted intervention*, pages 432–440. Springer, 2016.