Solving Inverse-PDE Problems with Physics-Aware Neural Networks

Samira Pakravan*[a], Pouria A. Mistani[a], Miguel A. Aragon-Calvo[c], Frederic Gibou[a,b]

[a] *Department of Mechanical Engineering, University of California, Santa Barbara, CA 93106-5070*
[b] *Department of Computer Science, University of California, Santa Barbara, CA 93106-5110*
[c] *Instituto de Astronomía, UNAM, Apdo. Postal 106, Ensenada 22800, B.C., México*

Abstract

We propose a novel composite framework that enables finding unknown fields in the context of inverse problems for partial differential equations (PDEs). We blend the high expressibility of deep neural networks as universal function estimators with the accuracy and reliability of existing numerical algorithms for partial differential equations. Our design brings together techniques of computational mathematics, machine learning and pattern recognition under one umbrella to seamlessly incorporate any domain-specific knowledge and insights through modeling. The network is explicitly aware of the governing physics through a hard-coded PDE solver stage; this subsequently focuses the computational load to only the discovery of the hidden fields. In addition, techniques of pattern recognition and surface reconstruction are used to further represent the unknown fields in a straightforward fashion. Most importantly, our inverse-PDE solver allows effortless integration of domain-specific knowledge about the physics of underlying fields, such as symmetries and proper basis functions. We call this approach Blended Inverse-PDE Networks (hereby dubbed BIPDE-Nets) and demonstrate its applicability on recovering the variable diffusion coefficient in Poisson problems in one and two spatial dimensions. We also show that this approach is robust to noise.

Keywords: inverse problems, differential equations, deep learning, scientific machine learning, numerical methods

1. Introduction

Inverse differential problems, where given a set of measurements one seeks a set of optimal parameters in a governing differential equation, arise in numerous scientific and technological domains. Some well-known applications include X-ray tomography [11, 32], ultrasound [44], MRI imaging [20] and diffusion in porous media [19]. Another interesting area is in multi-scale modeling, such as coarse-graining theories in physics where one is concerned with discovering effective descriptions of complex systems at different scales of length [14, 51]. However solving inverse problems poses substantial computational and mathematical challenges to infer reliable parameters from limited data.

In this work, we focus on the case of physical phenomena that can be modeled by:

$$\nabla \cdot (\mathbf{a}(\boldsymbol{x})\nabla u) = f(\boldsymbol{x}), \quad \boldsymbol{x} \in \Omega, \tag{1}$$

$$u(\boldsymbol{x}) = u_0(\boldsymbol{x}), \quad \boldsymbol{x} \in \partial\Omega, \tag{2}$$

where u is the unknown and $\mathbf{a}(\boldsymbol{x})$ is any symmetric positive definite matrix of functions. The functions $f(\boldsymbol{x})$ and $u_0(\boldsymbol{x})$ are given. The inverse problem is to approximate the coefficients $\boldsymbol{\theta} \equiv (\mathbf{a}(\boldsymbol{x}))$, given the solution u.

Recently, deep neural networks have attracted considerable attention for data modeling in a vast range of scientific domains. For example deep neural networks have shown astonishing success in emulating sophisticated simulations [16, 50, 49, 7, 39], discovering governing differential equations

*Corresponding author: spakravan@ucsb.edu

from data [36, 6, 28, 38] as well as potential applications to study and improve simulations of multiphase flows [13]. However these architectures require massive datasets and extensive computations to train numerous hidden weights and biases. Therefore reducing complexity of deep neural network architectures for inverse problems poses a significant practical challenge for many real-world applications in physical sciences, especially where collection of large datasets is a prohibitive task [35]. One remedy to reduce the network size is to embed the knowledge from existing mathematical models [42] or known physical laws within a neural network architecture [27, 12]. Along these lines, *semantic autoencoders* were recently proposed by Aragon-Calvo (2019) [2] where they replaced the decoder stage of an autoencoder architecture with a given physical law that can reconstruct the provided input data given a physically meaningful set of parameters. The encoder is then constrained to discover optimal values for these parameters, which can be extracted from the bottleneck of this network after training. We shall emphasize that this proposal not only reduces the size of the unknown model parameters, it also adds interpretability to deep learning frameworks. Inspired by their work, we propose to blend traditional numerical solver algorithms with custom deep neural network architectures to solve inverse PDE problems more efficiently and with higher flexibility.

One may identify three general strategies for solving forward and inverse PDE problems using neural networks. In the first class, the neural network is treated as a component in the solution. An example is NETT (Network Tikhonov) of Li et al. [26], which seeks regularized solutions with the regularizers being trained using a neural network. The second type of algorithms uses the neural network to represent the full solution to the PDE system. These algorithms augment the cost function with terms that describe the PDE, its boundary and its initial conditions. Then the solution is obtained by optimizing the neural network. Recent literature has mostly adopted this approach with examples including [37, 40, 4]. In the third approach, albeit exclusively for inverse problems, one uses a neural network along with a traditional differential equation solver, with the network representing the unknown fields. Then training occurs to minimize the difference between the output of the differential solver and the given data of solution. Due to its lower computational cost, this approach has been used since the 90s [10]. However previous works in this direction relied on inventing special discretization schemes. Recently Xu and Darve (2019) [48] examined the possibility to directly use pre-existing discretizations schemes within the framework of neural networks. Particularly they analyzed accuracy, sensitivity and convergence of such schemes. Xu and Darve's work was inspired by Stuarte (2010) [43] who showed the advantages arising from deferring discretizations "*to the very end of algorithmic formulation*" in the context of Bayesian inference.

In the present work we generalize this idea to higher dimensional stationary PDEs, discuss computational challenges and present a novel neural network architecture that overcomes the computational overhead in higher dimensions by taking advantage of techniques developed in pattern recognition. Most importantly, our architecture learns the inverse transform to reveal the hidden fields only from data and in a self-supervised fashion. Moreover, the proposed framework is versatile as it allows straightforward consideration of domain-specific knowledge to further reduce the computational cost by constraining the space of possibilities. The advantages of this framework are:

- seamlessly blends existing efficient and reliable PDE solvers with established neural network architectures, such as those at the intersection of computer vision and deep learning,
- lowers the computational cost as a result of reusing classical numerical algorithms for PDEs during the learning process, which ensures best use of provided data, i.e. to infer the actual unknowns in the problem,
- enables the user to incorporate domain-specific physical knowledge about unknown fields within the proposed neural network, such as symmetries or specialized basis functions,
- benefits from a homogeneous design by treating the PDE solver as another layer that can be stacked with various neural network architectures. This point also improves the performance by allowing vectorized operations accelerated by GPUs, as well as alowing the deep learning backend to perform back propagations internally,

Procedure	The BIPDE-Net algorithm for inverse PDE problems
1	Construct a deep neural network for the unknown coefficients $\hat{\boldsymbol{\theta}}$. Note that the dimension of the output layer must match the sum of dimensions for all unknown coefficients.
2	Construct a traditional PDE solver in `TensorFlow` with the output of the previous layer as its coefficients.
3	Define the loss function as the difference between the data and the reconstruction.
4	Train to minimize the loss function by optimizing the weights and biases of DNN of coefficients.

Table 1: Proposed algorithm.

- is essentially mesh-free. Hence the proposed architecture can be easily generalized to accommodate custom data structures by adding interpolation layers, etc.

In section 2, we present the proposed architecture and show its performance on 1D inverse-PDE problems. In section 3, we extend this approach to 2D by embedding special basis functions that are well-suited for image reconstruction. We also demonstrate its performance on 2−dimensional case studies and illustrate that the trained encoder actually estimates the inverse transform operator. Finally, we conclude with discussions and enumerate future work in this direction.

2. Blended inverse-PDE architecture (BIPDE-Net)

We embed a numerical solver into a deep learning architecture to recover unknown functions in inverse-PDE problems. In this section we describe our proposed architectures for inverse problems in one and two spatial dimensions.

2.1. Deep neural networks

The simplest neural network is a single layer of perceptron that mathematically performs a linear operation followed by a nonlinear composition applied to its input space,

$$\mathcal{N} = \sigma\big(\mathbf{W}\mathbf{x} + \mathbf{b}\big) \tag{3}$$

where σ is called the *activation function*. Deep neural networks are multiple layers stacked together with some architecture. The simplest example is a set of layers connected in series, known as feedforward neural networks (FNN). Therefore the action of a FNN is simply the successive compositions of previous layer outputs with the next layers, i.e.,

$$\mathcal{N}_l = \sigma\big(\mathbf{W}_l\mathcal{N}_{l-1}(\mathbf{x}) + \mathbf{b}_l\big), \tag{4}$$

where l indicates the index of a layer. This compositional nature of NNs is the basis of their vast potential as universal function estimators that can approximate any arbitrary function on the input space $\mathbf{x}$.

Many architectures have been proposed, such as convolutional neural networks (CNNs) [25, 24], Long-short term memory networks (LSTM) [18], and many more. In this work we pay particular attention to CNNs, owing to their ability to extract complicated spatial features from high dimensional input datasets. Particularly we will use CNNs for 2D inverse problems, while a simple FNN is sufficient for 1D problems.

2.2. Blended classical-neural network architectures

The basic idea is to represent unknown coefficients with a proper DNN as input to a traditional PDE solver. This is in contrast to recent attempts in the literature where the whole PDE, its boundary and initial conditions are reproduced by the DNN through adding them to the loss function. The main motivation for this composite scheme is to re-use the existing knowledge developed in the scientific computing community to improve the efficiency and accuracy of the

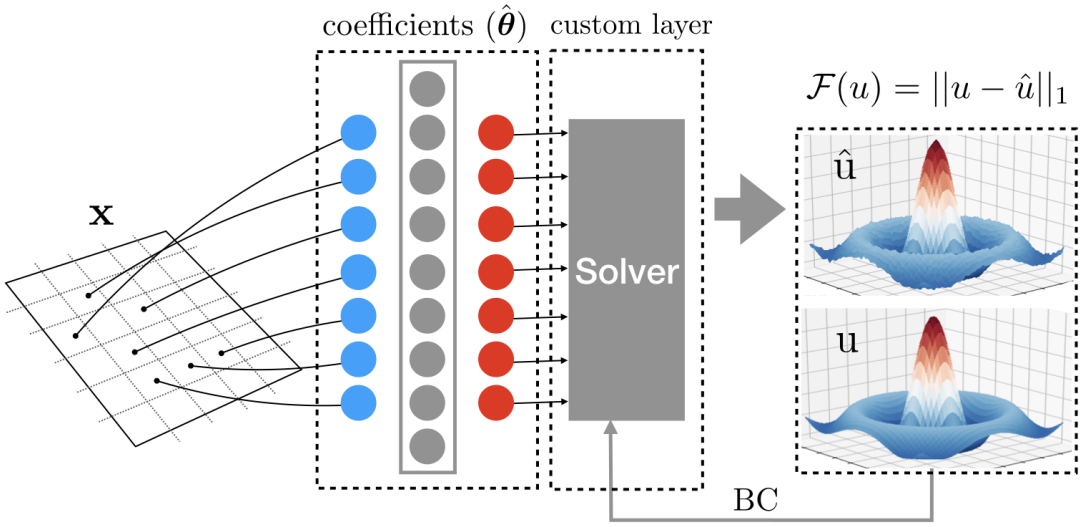

Figure 1: Two-stage architecture of BIPDE networks. The first stage learns the hidden fields and the second stage performs numerical computation defined within a `TensorFlow` custom layer. Note that depending on the problem, a variety of architectures, such as FNN or CNN, can be used for the first stage to represent the coefficients.

solution process. Our framework suggests a means to circumvent numerical errors that persist in the solutions provided by alternative existing deep learning frameworks that incorporate PDEs within the loss function of their optimizer, e.g. for example look at [37, 29].

The procedure of this composite approach is depicted in table 1 and illustrated in figure 1. As illustrated in figure 1, BIPDE-Net is a two-stage architecture, with the first stage responsible for learning the unknown coefficients and the second stage performing numerical operations as in traditional numerical solvers. To achieve higher performance it is essential to use the GPU-parallelism into the architecture. To leverage this capability, provided by the publicly available library `TensorFlow` [1], we implement our PDE-solver as a *custom layer* into our network using the `Keras` API [8]. Details of this includes vectorized operations to build the right-hand-side as well as the matrix of the linear system induced by the PDE discretization.

2.3. One-dimensional case studies

We consider the governing equation for diffusion dominated processes in heterogeneous media that is given by

$$\nabla \cdot \big(D(x)\nabla u(x)\big) + f(x) = 0, \qquad x \in \Omega \tag{5}$$

$$u(x) = u_0(x), \qquad x \in \partial\Omega \tag{6}$$

To solve the corresponding inverse problem, we consider a centered difference discretization in space for approximating the PDE portion in the interval $[x_{\min}, x_{\max}]$:

$$\frac{D_{i+1/2}(u_{i+1} - u_i) - D_{i-1/2}(u_i - u_{i+1})}{\Delta x^2} + f_i = 0,$$

$$D_{i+1/2} = \frac{D_i + D_{i+1}}{2},$$

where u_i refers to the numerical solution at grid point $x_i = x_{\min} + (i-1)\,\Delta x$ with $\Delta x = (x_{\max} - x_{\min})\,/(n-1)$, where n the number of discretization points. This discretization leads to a linear system $A\boldsymbol{u} = \boldsymbol{b}$ where the matrix A is tridiagonal. Here we show that shallow feedforward neural networks suffice to approximate the variable coefficient $D(x)$ accurately.

In figure 2, we depict the learned diffusion coefficients given a known source term $f(x) = \sin(\pi x)$. We considered four different cases that cover a broad range of possible functional forms for a variable coefficient function. The results indicate excellent agreement between predictions and ground truth. We emphasize the training is performed on a single snapshot with 100 uniformly distributed grid points along x-axis. These results corroborate the applicability of our proposed framework in learning complex unknown fields from limited data.

Here we used a feedforward neural network with 4 hidden layers composed of 20, 70, 40, 100 neurons respectively. The activation function of the first three layers is set to `ReLU`, *i.e.* a piecewise linear function that only outputs positive values $\mathrm{ReLU(x)} = \max(0, \mathrm{x})$ [15], while for the last layer we defined a custom activation function as a scaled sigmoid function given by,

$$\tilde{\sigma}(x) = \alpha\sigma(x), \qquad \sigma(x) = \frac{1}{1 + \exp(-x)}, \tag{7}$$

where α is a constant scalar value chosen to support the possible range of actual diffusion coefficients. This is to compensate the range of pure sigmoid function that is bound in $(0, 1)$. Even though there exist other activation functions such as `SoftPlus` that are essentially boundless and positive, we found that a scaled sigmoid function provides better performance as long as α is chosen properly to cover the possible range. Otherwise, we observed the predictions follow the true profile except around the peak where prediction values become saturated with values close to α. In each layer we also added an L^1-norm regularization to improve the results. We used `adam` optimizer [23] with mean absolute error loss function, and we trained the network for 100 epochs to obtain the results shown in figure 2.

3. Higher dimensions and a moment-based approach

For the two-dimensional case, we examined both a feedforward architecture and a convolutional neural network to recover the unknown diffusion map. In our experiments, we found that even though evaluating every pixel value of the hidden fields is computationally very expensive, a CNN is more robust than a FNN. This is because a CNN promotes construction of smoother manifolds by application of successive convolutions while a FNN suffers from noisy outputs, unless aided by proper regularization. Furthermore, to build a FNN we need to flatten the coordinates into large 1D arrays that substantially increase the training time and the number of unknown parameters. This is because in a convolutional layer one only optimizes for a single kernel per channel (as well as the bias parameters), which amounts to exponentially less unknown parameters. Overall we conclude that a direct reconstruction strategy is not the best strategy when treating inverse problems in higher dimensions. In what follows, we propose an alternative approach to tackle this challenge.

The previously discussed method relies on the idea to reconstruct every pixel value for unknown fields directly, however this task proves challenging in higher dimensions. Alternatively, one can decompose the hidden fields into a finite number of eigenfunctions and seek optimal coefficients for each term. This is also advantageous from a physics point of view, where pre-existing knowledge about hidden fields can be naturally formulated in terms of basis functions into this framework. One such family of series expansions are the moment-based methods that have been largely exploited in image reconstruction [22, 5, 33, 3]. In particular Zernike moments [45] provide a linearly independent set of polynomials defined on the unit circle/sphere in 2D/3D. Zernike moments are well-suited for such a task and are commonly used for representing optical aberration in astronomy and atmospheric sciences [34], as well as for image reconstruction and enhanced ultrasound focusing in biomedical imaging [9, 30, 21].

Zernike moments are advantageous over regular moments in that they intrinsically provide rotational invariance, higher accuracy for irregular patterns, as well as orthogonality that reduces information redundancy in different coefficients. Zernike polynomials capture deviations from zero mean as a function of radius and azimuthal angle. Furthermore, Zernike moments provide a complete set of orthogonal bases that can be obtained with lower computational precision from input images, which enhances the robustness of the reconstruction procedure.

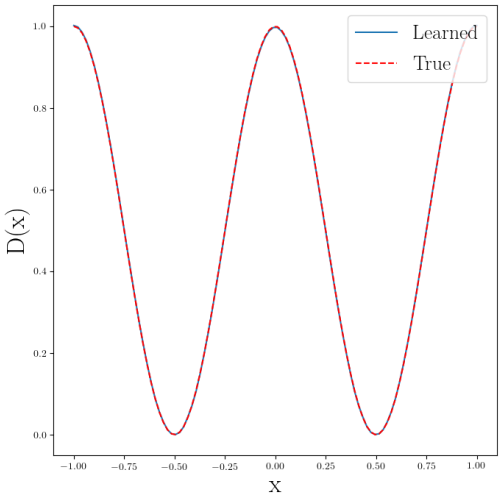

(a) $0.5\cos(2\pi x) + 0.5$

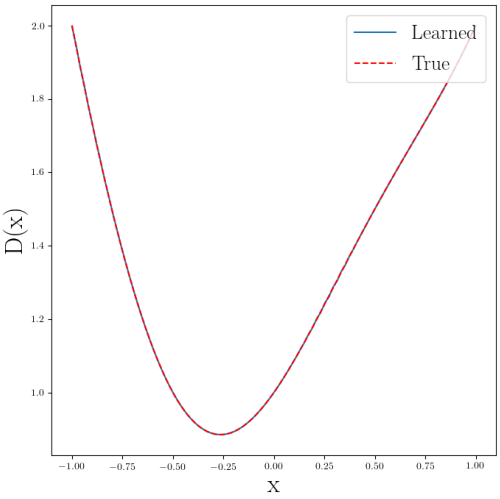

(b) $0.25\sin(\pi x) + 1 + x^2$

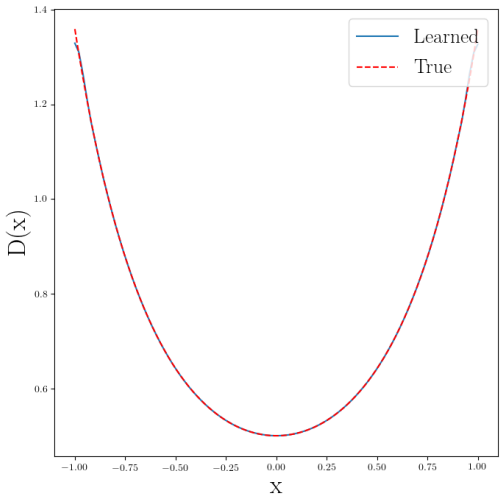

(c) $0.5\exp(x^2)$

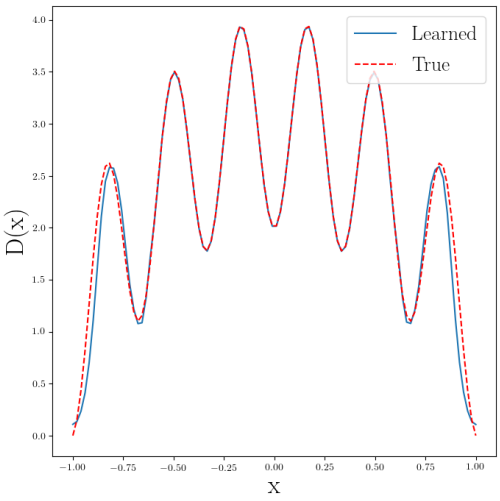

(d) $2(1+x)(1-x) + 2\sin(3\pi x)^2$

Figure 2: Discovered variable diffusion coefficients in the 1D case.

n	|m|	$\mathbf{R}_{nm}$	$\mathbf{Z}^o_{nm}$	$\mathbf{Z}^e_{nm}$	Aberration/Pattern
0	0	1	0	1	Piston
1	1	ρ	$\rho\sin(\theta)$	$\rho\cos(\theta)$	Tilt
2	0	$2\rho^2-1$	0	$2\rho^2-1$	Defocus
	2	ρ^2	$\rho^2\sin(2\theta)$	$\rho^2\cos(2\theta)$	Oblique/Vertical Astigmatism
3	1	$3\rho^3-2\rho$	$(3\rho^3-2\rho)\sin(\theta)$	$(3\rho^3-2\rho)\cos(\theta)$	Vertical/Horizontal Coma
	3	ρ^3	$\rho^3\sin(3\theta)$	$\rho^3\cos(3\theta)$	Vertical/Oblique Trefoil
4	0	$6\rho^4-6\rho^2+1$	0	$6\rho^4-6\rho^2+1$	Primary Spherical
	2	$4\rho^4-3\rho^2$	$(4\rho^4-3\rho^2)\sin(2\theta)$	$(4\rho^4-3\rho^2)\cos(2\theta)$	Oblique/Vertical Secondary Astigmatism
	4	ρ^4	$\rho^4\sin(4\theta)$	$\rho^4\cos(4\theta)$	Oblique/Vertical Quadrafoil

Table 2: First 15 odd and even Zernike polynomials according to Noll's nomenclature. Here, the ordering is determined by ordering polynomial with lower radial order first, cf. [47].

Odd and even Zernike polynomials are given by odd or even choice in azimuthal angle,

$$\begin{bmatrix} Z^o_{nm}(\rho,\theta) \\ Z^e_{nm}(\rho,\theta) \end{bmatrix} = R_{nm}(\rho) \begin{bmatrix} \sin(m\theta) \\ \cos(m\theta) \end{bmatrix},$$

with

$$R_{nm}(\rho) = \begin{cases} \sum_{l=0}^{(n-|m|)/2} \frac{(-1)^l (n-l)!}{l![(n+|m|)/2-l]![(n-|m|)/2-l]!} \rho^{n-2l} & \text{for } n-m \text{ even}, \\ 0 & \text{for } n-m \text{ odd}, \end{cases}$$

where n and m are integers with $n \geq m$, θ is the azimuthal angle and ρ is the radial distance between 0 and 1 measured from the center of image. A list of radial components are given in table 2 (from [46]). For an extensive list of Zernike polynomials in both 2D and 3D we refer the interested reader to [31].

Furthermore, each Zernike moment is defined by projection of the hidden field $f(x,y)$ on the orthogonal bases,

$$\begin{bmatrix} A_{nm} \\ B_{nm} \end{bmatrix} = \frac{n+1}{\epsilon^2_{mn}\pi} \int_x \int_y f(x,y) \begin{bmatrix} Z^o_{nm}(x,y) \\ Z^e_{nm}(x,y) \end{bmatrix} dxdy, \quad x^2+y^2 \leq 1,$$

where for $m=0,\ n\neq 0$ we defined $\epsilon_{0n} = 1/\sqrt{2}$ and $\epsilon_{mn} = 1$ otherwise. Finally, superposition of these moments expands the hidden field in terms of Zernike moments:

$$\hat{f}(x,y) = \sum_{n=0}^{N_{max}} \sum_m \left[A_{nm} Z^o_{nm}(r,\theta) + B_{nm} Z^e_{nm}(r,\theta) \right] \tag{8}$$

3.1. Semantic autoencoder for inverse-PDE problems

In order to identify the coefficients in the Zernike expansion (8) for hidden fields, we use a semantic autoencoder architecture as proposed by Aragon-Calvo (2019) [2] with Zernike moments being represented by the code at the bottleneck of the autoencoder. Figure 3 illustrates the architecture for the proposed semantic autoencoder. As a test problem, we consider the Poisson equation

$$\nabla \cdot \big(D(\mathbf{x})\nabla u\big) = -f(\mathbf{x}) \tag{9}$$

with Dirichlet boundary conditions.

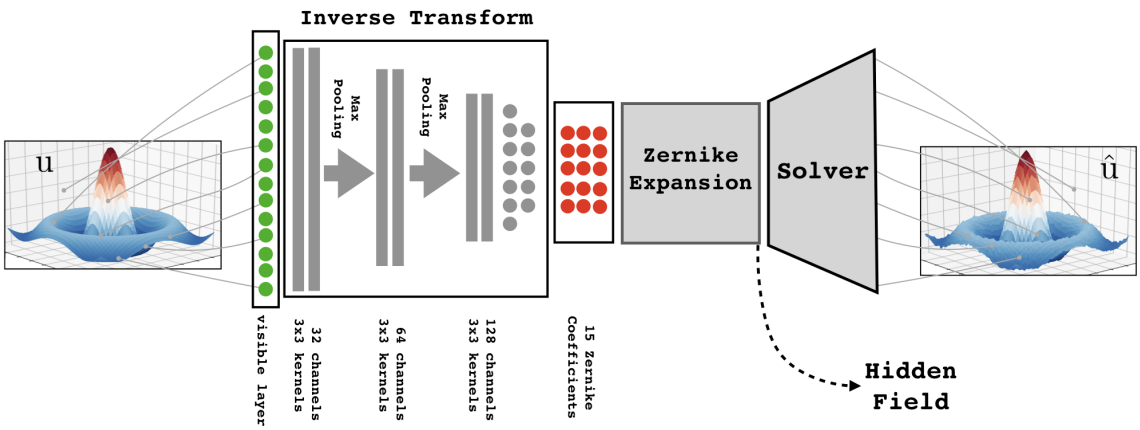

Figure 3: Architecture of the semantic autoencoder to infer hidden fields. Zernike moments are discovered at the bottleneck of the architecture.

Discretization. In our architecture we used the standard 5-point stencil finite difference discretization of the Poisson equation in the solver layer, i.e.

$$\frac{D_{i-1/2,j}u_{i-1,j} - (D_{i-1/2,j} + D_{i+1/2,j})u_{i,j} + D_{i+1/2,j}u_{i+1,j}}{\Delta x^2} + \\ \frac{D_{i,j-1/2}u_{i,j-1} - (D_{i,j-1/2} + D_{i,j+1/2})u_{i,j} + D_{i,j+1/2}u_{i,j+1}}{\Delta y^2} + f_{i,j} = 0,$$

and we use the linear algebra solvers implemented in `TensorFlow` to solve for the solution field.

Architecture. In the training process of a CNN, the kernels are trained at each layer such that several feature maps are extracted at each layer from input data. The CNN is composed of 3 convolutional blocks with 32, 64, 128 channels respectively and kernel size 3×3. Moreover, we use the `MaxPooling` filter with kernel size $(2, 2)$ after each convolutional block to downsample the feature maps by calculating the maximum values of each patch within these maps. We used `ReLU` activation function in the convolutional layers, followed by `Sigmoid` activation in dense layers and a scaled `Sigmoid` activation at the final layer,

$$\tilde{\sigma}(x) = D_{\min} + (D_{\max} - D_{\min})\sigma(x), \tag{10}$$

such that the actual values of diffusion coefficient are within the range $(\mathrm{D_{min}}, \mathrm{D_{max}})$, known from domain specific knowledge. After each dense layer we applied `Dropout` layers with rate of 0.2 to prevent overfitting [17, 41].

3.2. Solving 2D inverse-PDE problems

In this section we assess performance of the proposed solution in 2D for two test cases.

3.2.1. Case I. A tilted plane

In the first example we consider a linear diffusion model given by

$$D(x, y) = \frac{1}{\sqrt{2}} + 0.1(y - x)$$

with

$$u_{BC}(x, y) = 0.01\cos(\pi x)\cos(\pi y) \qquad \text{and} \qquad f(x, y) = \sin(\pi x)\cos(\pi y)$$

Figure 4 depicts the results obtained by the proposed scheme. The diffusion map is discovered

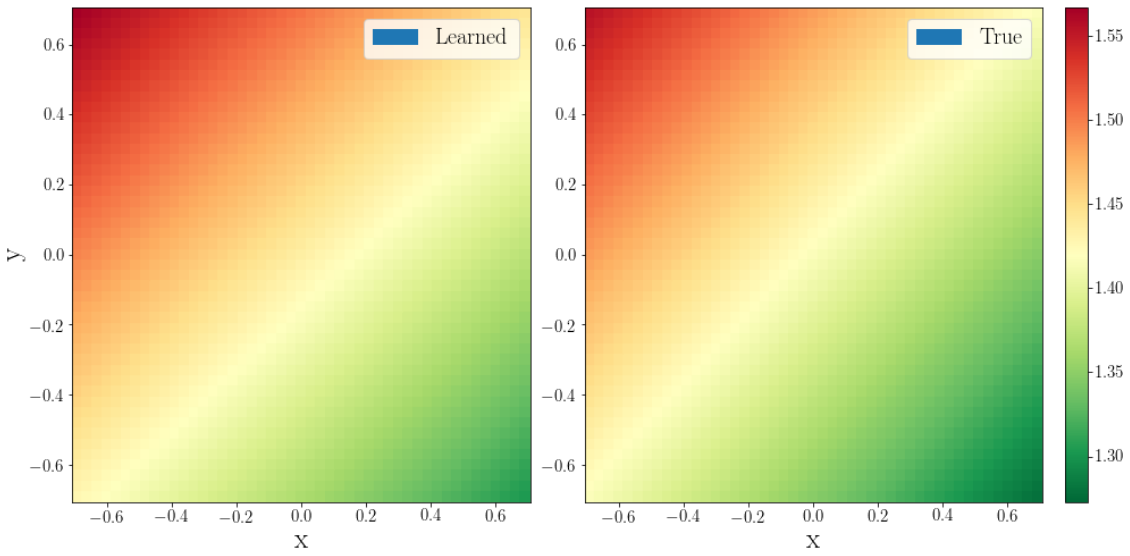

(a) Comparison of learned (left) versus true diffusion coefficient (right).

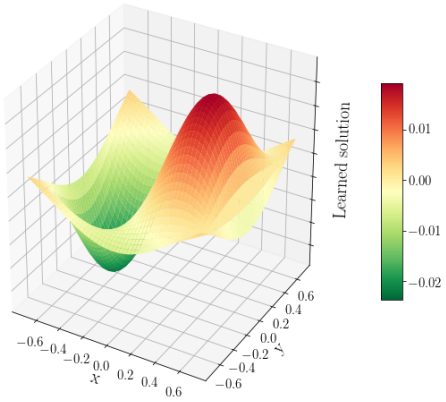

(b) Learned solution.

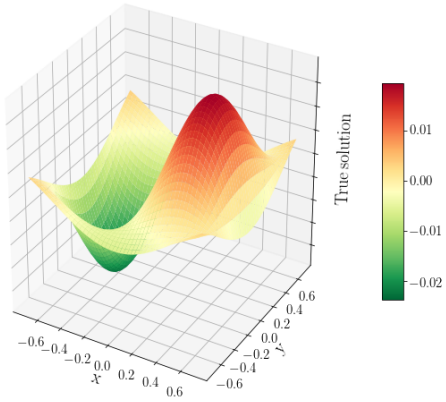

(c) True solution.

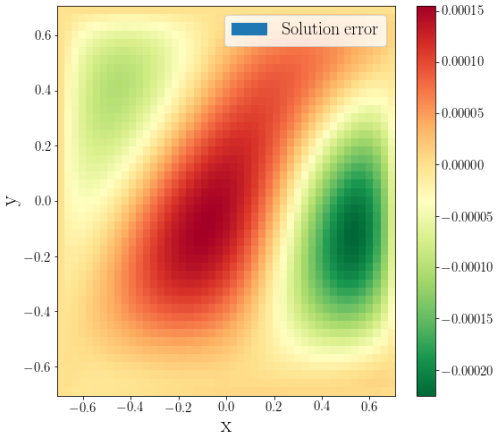

(d) Error in learned solution u − û.

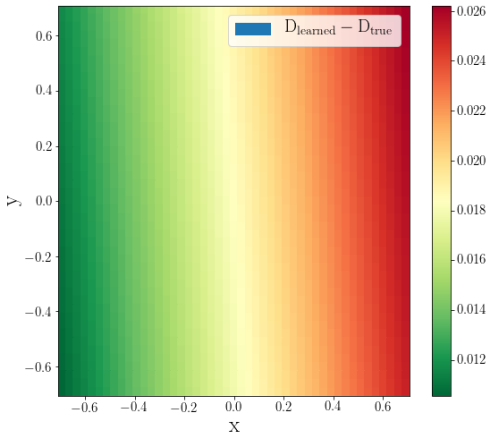

(e) Error in learned diffusion coefficient.

Figure 4: Results for the 2D linear case.

with a maximum relative error of only 2%, while the error in solution field is 1%. It is noteworthy to mention that the accuracy of the results in this architecture are influenced by the accuracy of the discretizations used in the solver layer. While we used a second-order accurate finite difference discretization in this work it is be possible to improve these results by using higher order discretizations instead. We leave such optimizations to future work.

3.2.2. Case II. superimposed Zernike polynomials

We consider a more complicated hidden diffusion field given by

$$D(x, y) = 4 + a_0 + 2a_1x + 2a_2y + \sqrt{3}a_3(2x^2 + 2y^2 - 1).$$

The boundary condition function u_{BC} and the source field f are given by

$$u_{BC}(x, y) = \cos(\pi x)\cos(\pi y) \qquad \text{and} \qquad f(x, y) = x + y.$$

Figure 5 illustrates the performance of the proposed Zernike-based network using a mean absolute error measure for the loss function. We trained the network for 100 epochs using an `adam` optimizer.

Resilience to noise. We also assess the performance of our scheme on noisy datasets. We considered a Gaussian noise with standard deviation 0.025 superimposed to the solution field. Figure 6 depicts the solution learned from a noisy input image. The network succeeds to discover the diffusion field with comparable accuracy as in the noiseless case. Interestingly, the network removed the added noise from the learned solution, a feature that is similar to application of low-pass filters on noisy images (which we attribute to the CNN).

4. Learning the inverse transform

In the previous sections, we have applied BiPDE to find the variable diffusion coefficient from a single input image. Another interesting feature of the proposed semantic autoencoder architecture, is to train neural networks to discover the inverse transform for the underlying hidden fields in a self-supervised fashion. In this view, the trained encoder learns the inverse transform function that outputs the hidden parameters given a solution field as input. In this scenario, the network learns the inverse transform, $\mathbf{T}^{-1}$, defined as $\hat{\mathrm{D}} = \mathbf{T}^{-1}(\mathrm{u})$. In this section we train an encoder that learns the hidden Zernike *moments*, given an ensemble of solution images.

4.1. 1D inverse transform

We build a 1D semantic autoencoder using 3 feedforward layers with 100, 40, and 2 number of neurons respectively. We used the `ReLU` activation function for the first two layers and a `Sigmoid` activation function for the last layer representing the hidden parameters. A linear solver is then stacked with the encoder as discussed earlier. However, the diffusion map is internally reconstructed using the hidden parameters before feeding the output of the encoder to the solver. We emphasize that in this example, we input the solution fields directly into the network instead of the coordinates as discussed earlier.

As a test problem we consider the 1D Poisson problem with a generic linear form for diffusion coefficient,

$$D(x) = 1 + a_0 + a_1x.$$

We consider identical left and right Dirichlet boundary conditions of 0.2 for all images and let $\mathrm{f(x)} = \sin(\pi \mathrm{x})$. We consider random diffusion coefficients a_0 and a_1 with a uniform distribution in $[0.25, 0.75]$ and we generate 1000 solutions over the domain $\mathrm{x} \in [-1, 1]$. We train a semantic autoencoder over 900 images from this dataset and validate its performance over the remaining 100

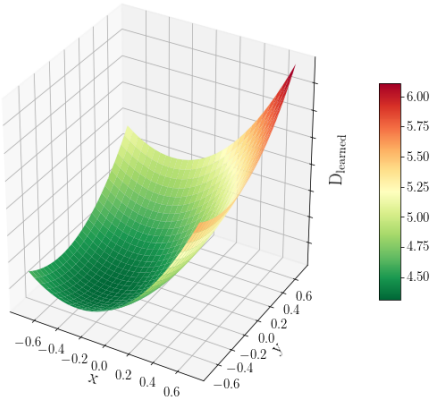

(a) Learned diffusion.

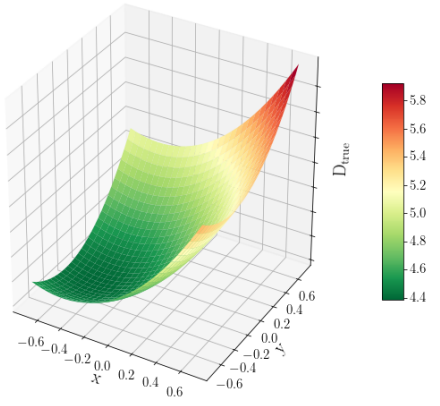

(b) True diffusion.

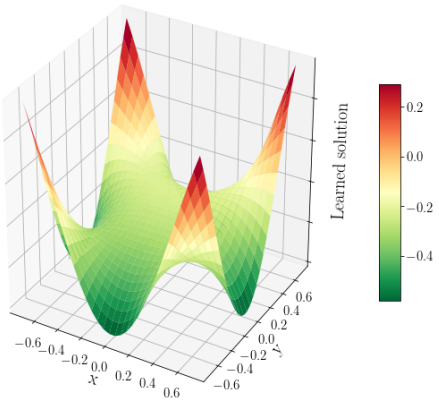

(c) Learned solution.

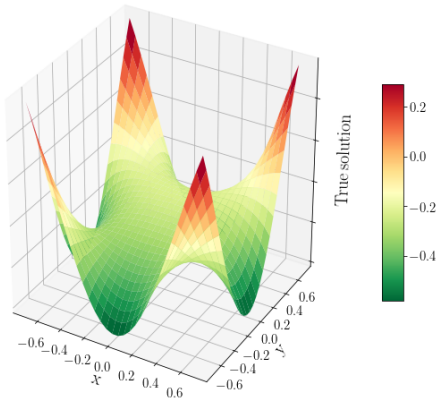

(d) True solution.

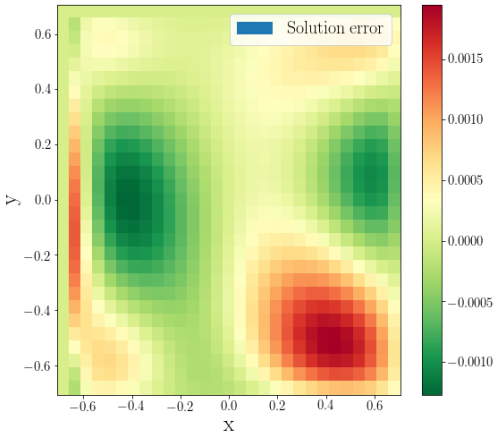

(e) Error in learned solution $u - \hat{u}$.

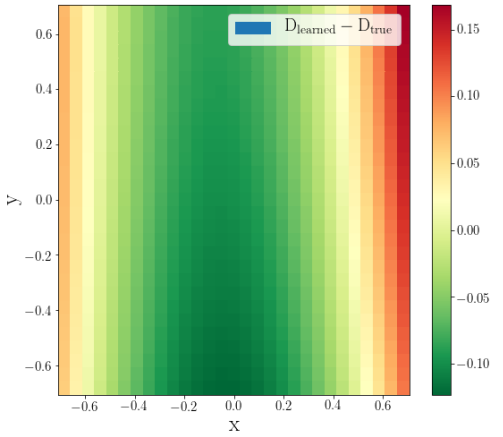

(f) Error in learned diffusion coefficient.

Figure 5: Results in the 2D parabolic case.

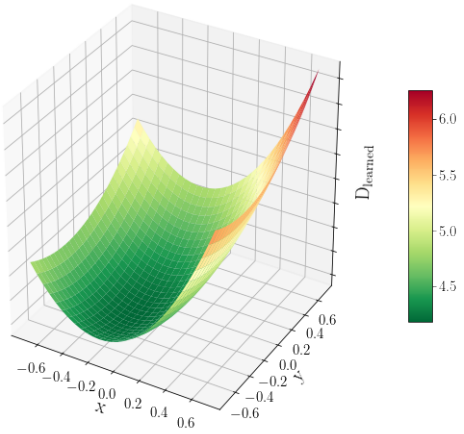

(a) Learned diffusion.

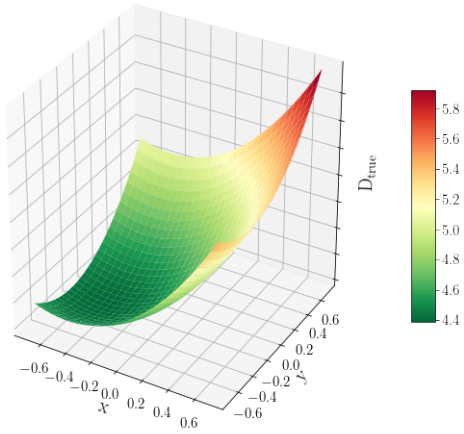

(b) True diffusion.

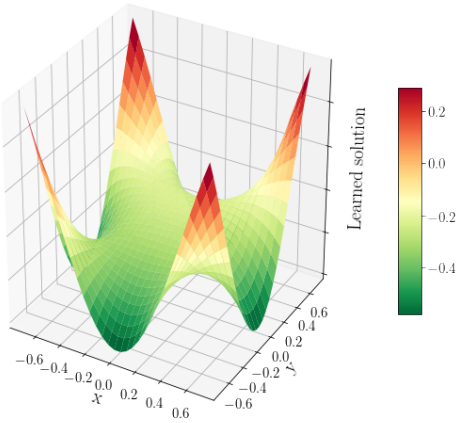

(c) Learned solution.

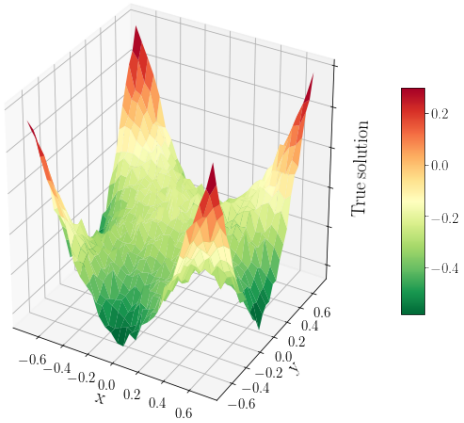

(d) Noisy input solution.

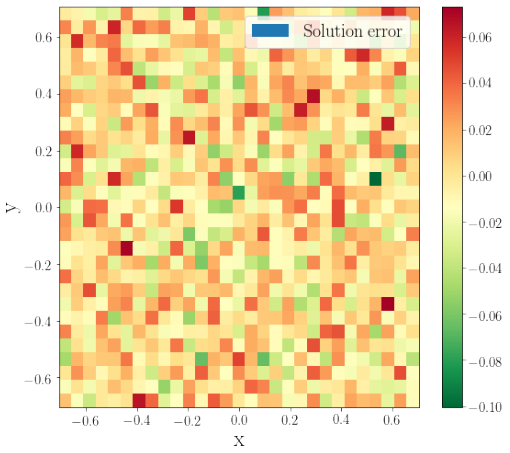

(e) Error in learned solution u − û.

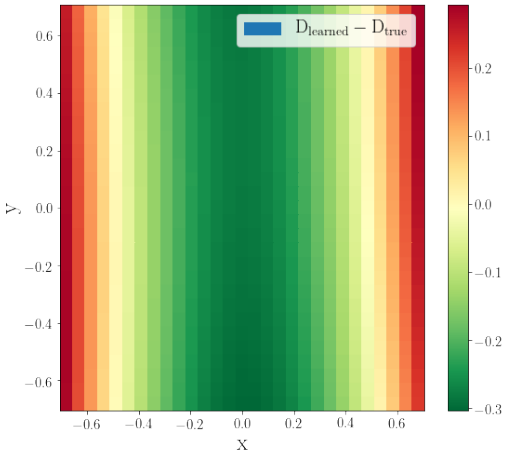

(f) Error in learned diffusion coefficient.

Figure 6: Results in the 2D case with added noise. After 300 epochs the network discovers the hidden diffusion field with a maximum relative error of 5%. Interestingly the learned solution is resilient to added noise and the network approximates a noise-free solution.

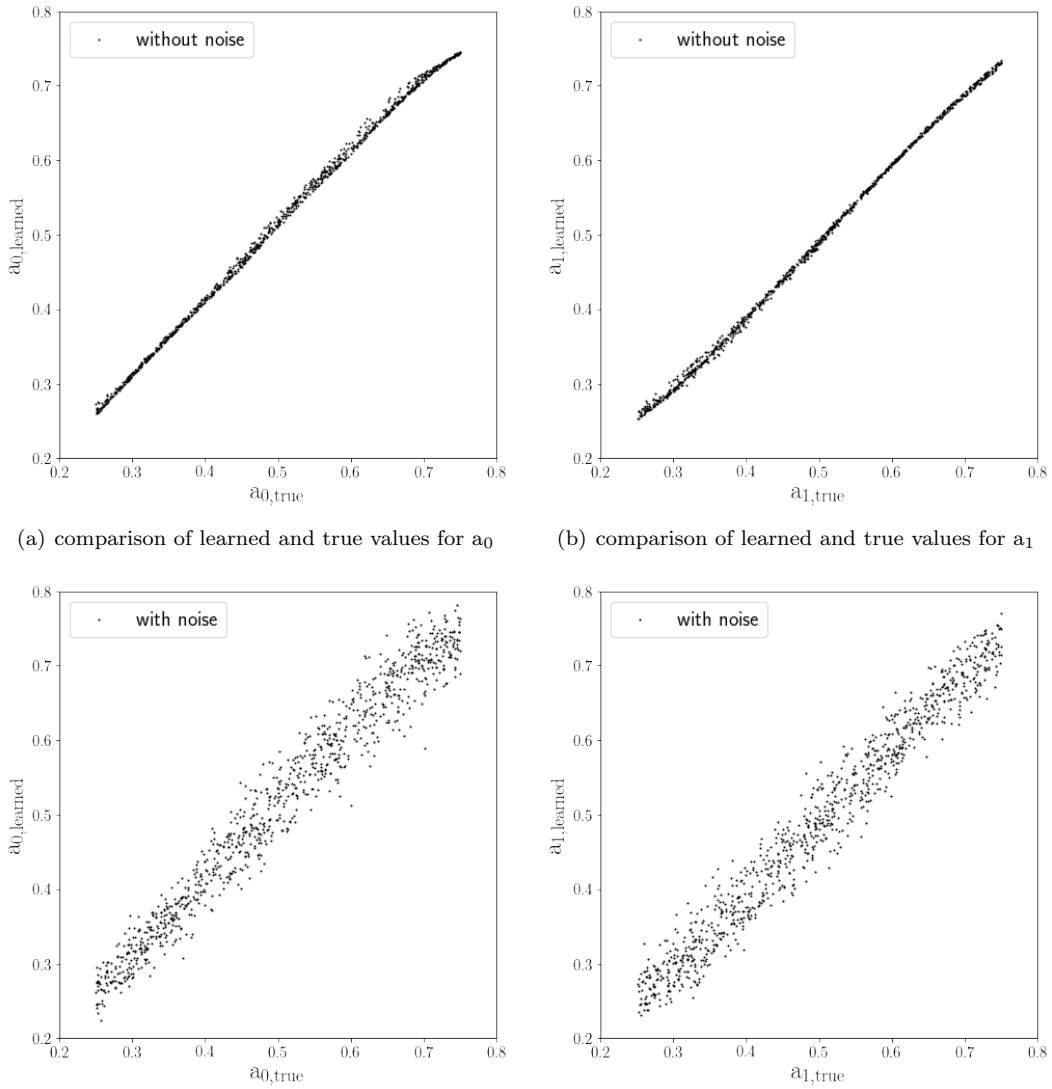

(a) comparison of learned and true values for a_0

(b) comparison of learned and true values for a_1

(c) comparison of learned and true values for a_0 with added noise

(d) comparison of learned and true values for a_1 with added noise

Figure 7: The top panel depicts the performance of the semantic-autoencoder over 1000 randomly chosen 1D images after 100 epochs. The bottom panel depicts performance of the trained network on the same dataset but with added Gaussian noise with standard deviation 0.025 to our test sample.

images using a mean absolute error loss function for 1000 epochs. We used a batch size of 100 in our experiments using the `adam` optimizer. Figure 7 compares the learned and the true coefficients over two independent test samples containing 1000 solutions, with and without a Gaussian noise of standard deviation 0.025, i.e. amounting to $< 13\%$ added noise over images. The observed discrepancies, even though at $\sim 10\%$ for a_0 and a_1, partially stem from relatively small sample size in our experiments and can be improved by enlarging the sample size. Furthermore, the predicted values for a_0 exhibit a systematic increase towards the lower and upper bounds on output of `Sigmoid` activation function, and hence may be improved by properly scaling the activation functions for the last layer as we discussed in previous sections. This hypothesis is also supported by the improved accuracy at the middle of output range, i.e. around a_0, $a_1 \sim 0.5$, indicative of influence of `Sigmoid` activation function used in the last layer.

4.2. 2D inverse transform

We consider an example of variable diffusion coefficients parametrized as $D(x, y) = 1 + a_1 x$, with a_1 randomly chosen in range $a_1 \in [0.05, 0.95]$. The equations are solved on a square domain $\Omega \in [-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}]^2$ governed by the Poisson equation:

$$\begin{aligned} \nabla \cdot \big([1 + a_1 x]\nabla u\big) + x + y = 0, \qquad & (x, y) \in \Omega, \\ u_{BC} = \cos(\pi x)\cos(\pi y), \qquad & (x, y) \in \partial\Omega. \end{aligned}$$

We trained our autoencoder over 900 generated solutions with randomly chosen parameter a_1 and validated its performance on 100 independent solution fields for 100 epochs using a mean absolute error loss function. Then we tested the trained model over another set of 1000 images with randomly generated diffusion maps independent of training dataset. Furthermore, we repeated this exercise over 1000 images with added Gaussian noise with standard deviation 0.025. In each case, the learned parameters are in good agreement with the true values, as illustrated in figure 8. We note that in two spatial dimensions, learning additional hidden parameters requires convolutional layers with higher number of channels to discover the independent signatures of each unknown parameter on the solution field; this subsequently makes training more expensive. Therefore, in this example we only considered a single unknown parameter to demonstrate the applicability of this approach in higher dimensions.

5. Conclusion

We introduced BIPDE-Net, a natural tool to infer effective properties of complex media given limited number of observations. The main advantage of BIPDE-Nets is the ability to incorporate physical knowledge about the problem. BIPDE-Net is a versatile algorithm that can be easily applied to arbitrary inverse-PDE problems, and is shown to be extremely data-efficient. We showed performance of this design on multiple inverse Poisson problems in one and two spatial dimensions and we showed how our framework is robust against added noise. We also showed applicability of this framework to discover inverse transforms for different inverse-PDE problems. There are many areas of research that could be further investigated, such as considering diffusion maps with discontinuities across substructures in the domain, using other network layers such as LSTM that are particularly suited to reproduce sequences of data, using higher order numerical solvers and finally tackle more complicated governing PDE problems with a larger number of unknown fields.

Acknowledgement

This research was supported by ARO W911NF-16-1-0136 and ONR N00014-17-1-2676.

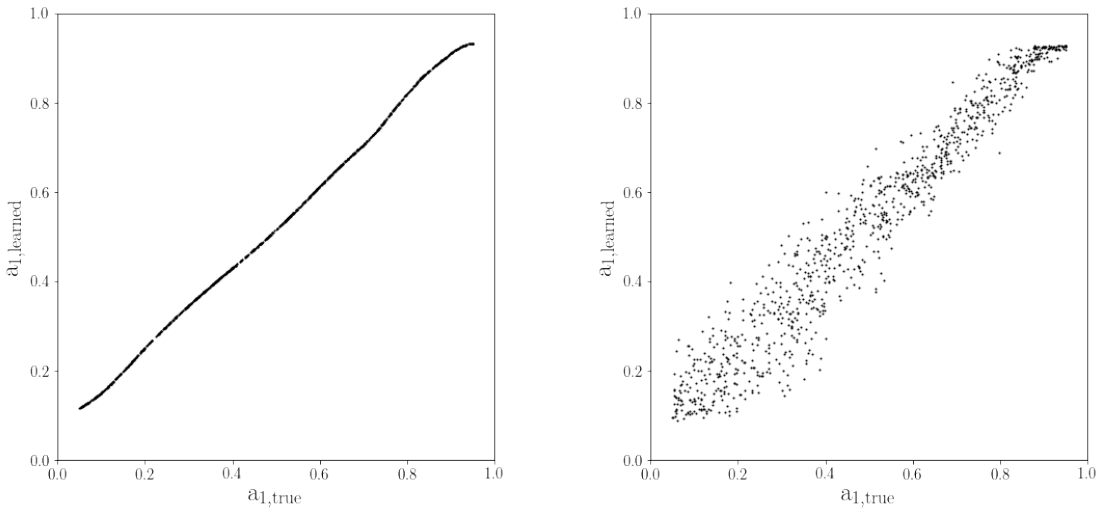

(a) Without noise.

(b) With Gaussian noise.

Figure 8: Left figure depicts performance of semantic-autoencoder over 1000 randomly chosen 2D images after 100 epochs, and the right panel shows performance of the same network on noisy images given a Gaussian noise with standard deviation 0.025.

References

References

[1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.

[2] M. A. Aragon-Calvo. Self-supervised learning with physics-aware neural networks i: Galaxy model fitting. *arXiv preprint arXiv:1907.03957*, 2019.

[3] R. R. Bailey and M. Srinath. Orthogonal moment features for use with parametric and non-parametric classifiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(4):389–399, 1996.

[4] L. Bar and N. Sochen. Unsupervised deep learning algorithm for pde-based forward and inverse problems. *arXiv preprint arXiv:1904.05417*, 2019.

[5] S. O. Belkasim, M. Shridhar, and M. Ahmadi. Pattern recognition with moment invariants: a comparative study and new results. *Pattern recognition*, 24(12):1117–1138, 1991.

[6] J. Berg and K. Nyström. Data-driven discovery of pdes in complex datasets. *Journal of Computational Physics*, 384:239–252, 2019.

[7] A. Chandrasekaran, D. Kamal, R. Batra, C. Kim, L. Chen, and R. Ramprasad. Solving the electronic structure problem with machine learning. *npj Computational Materials*, 5(1):22, 2019.

[8] F. Chollet et al. Keras, 2015.

[9] S. Dong, J. Kettenbach, I. Hinterleitner, H. Bergmann, and W. Birkfellner. The zernike expansion–an example of a merit function for 2d/3d registration based on orthogonal functions. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 964–971. Springer, 2008.

[10] I. Elshafiey, L. Udpa, and S. Udpa. A neural network approach for solving inverse problems in nde. In *Review of Progress in Quantitative Nondestructive Evaluation*, pages 709–716. Springer, 1992.

[11] C. L. Epstein. *Introduction to the mathematics of medical imaging*. SIAM, 2007.

[12] Z. Geng, D. Johnson, and R. Fedkiw. Coercing machine learning to output physically accurate results. *Journal of Computational Physics*, page 109099, 2019.

[13] F. Gibou, D. Hyde, and R. Fedkiw. Sharp interface approaches and deep learning techniques for multiphase flows. *Journal of Computational Physics*, 380:442 – 463, 2019.

[14] C. Grigo and P.-S. Koutsourelakis. A physics-aware, probabilistic machine learning framework for coarse-graining high-dimensional systems in the small data regime. *Journal of Computational Physics*, 397:108842, 2019.

[15] R. H. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, and H. S. Seung. Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature*, 405(6789):947, 2000.

[16] S. He, Y. Li, Y. Feng, S. Ho, S. Ravanbakhsh, W. Chen, and B. Póczos. Learning to predict the cosmological structure formation. *Proceedings of the National Academy of Sciences*, page 201821458, 2019.

[17] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.

[18] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[19] G. P. J.-P. Fouque, J. Garnier and K. Solna. *Wave Propagation and Time Reversal in Randomly Layered Media*. Springer, 2007.

[20] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, 2017.

[21] E. A. Kaye, Y. Hertzberg, M. Marx, B. Werner, G. Navon, M. Levoy, and K. B. Pauly. Application of zernike polynomials towards accelerated adaptive focusing of transcranial high intensity focused ultrasound. *Medical physics*, 39(10):6254–6263, 2012.

[22] A. Khotanzad and Y. H. Hong. Invariant image recognition by zernike moments. *IEEE Transactions on pattern analysis and machine intelligence*, 12(5):489–497, 1990.

[23] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[24] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[25] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[26] H. Li, J. Schwab, S. Antholzer, and M. Haltmeier. Nett: Solving inverse problems with deep neural networks. *arXiv preprint arXiv:1803.00092*, 2018.

[27] J. Ling, R. Jones, and J. Templeton. Machine learning strategies for systems with invariance properties. *Journal of Computational Physics*, 318:22 – 35, 2016.

[28] Z. Long, Y. Lu, X. Ma, and B. Dong. Pde-net: Learning pdes from data. *arXiv preprint arXiv:1710.09668*, 2017.

[29] L. Lu, X. Meng, Z. Mao, and G. E. Karniadakis. Deepxde: A deep learning library for solving differential equations. *arXiv preprint arXiv:1907.04502*, 2019.

[30] P. Markelj, D. Tomaževič, B. Likar, and F. Pernuš. A review of 3d/2d registration methods for image-guided interventions. *Medical image analysis*, 16(3):642–661, 2012.

[31] R. J. Mathar. Zernike basis to cartesian transformations. *arXiv preprint arXiv:0809.2368*, 2008.

[32] F. Natterer. *The mathematics of computerized tomography*. SIAM, 2001.

[33] R. J. Prokop and A. P. Reeves. A survey of moment-based techniques for unoccluded object representation and recognition. *CVGIP: Graphical Models and Image Processing*, 54(5):438–460, 1992.

[34] R. Ragazzoni, E. Marchetti, and G. Valente. Adaptive-optics corrections available for the whole sky. *Nature*, 403(6765):54, 2000.

[35] M. Raissi and G. E. Karniadakis. Hidden physics models: Machine learning of nonlinear partial differential equations. *Journal of Computational Physics*, 357:125 – 141, 2018.

[36] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Machine learning of linear differential equations using gaussian processes. *Journal of Computational Physics*, 348:683 – 693, 2017.

[37] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*, 2017.

[38] H. Schaeffer. Learning partial differential equations via data discovery and sparse optimization. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2197):20160446, 2017.

[39] A. V. Sinitskiy and V. S. Pande. Deep neural network computes electron densities and energies of a large set of organic molecules faster than density functional theory (dft). *arXiv preprint arXiv:1809.02723*, 2018.

[40] J. Sirignano and K. Spiliopoulos. Dgm: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339 – 1364, 2018.

[41] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.

[42] P. Stinis, T. Hagge, A. M. Tartakovsky, and E. Yeung. Enforcing constraints for interpolation and extrapolation in generative adversarial networks. *Journal of Computational Physics*, 397:108844, 2019.

[43] A. M. Stuart. Inverse problems: a bayesian perspective. *Acta numerica*, 19:451–559, 2010.

[44] R. J. van Sloun, R. Cohen, and Y. C. Eldar. Deep learning in ultrasound imaging. *Proceedings of the IEEE*, 2019.

[45] Z. von F. Beugungstheorie des schneidenver-fahrens und seiner verbesserten form, der phasenkontrastmethode. *physica*, 1(7-12):689–704, 1934.

[46] E. W. Weisstein. Zernike polynomial. 2002.

[47] J. C. Wyant and K. Creath. Basic wavefront aberration theory for optical metrology. *Applied optics and optical engineering*, 11(part 2):28–39, 1992.

[48] K. Xu and E. Darve. The neural network approach to inverse problems in differential equations. *arXiv preprint arXiv:1901.07758*, 2019.

[49] J. Zamudio-Fernandez, A. Okan, F. Villaescusa-Navarro, S. Bilaloglu, A. D. Cengiz, S. He, L. P. Levasseur, and S. Ho. Higan: Cosmic neutral hydrogen with generative adversarial networks. *arXiv preprint arXiv:1904.12846*, 2019.

[50] X. Zhang, Y. Wang, W. Zhang, Y. Sun, S. He, G. Contardo, F. Villaescusa-Navarro, and S. Ho. From dark matter to galaxies with convolutional networks. *arXiv preprint arXiv:1902.05965*, 2019.

[51] L. Zhao, Z. Li, B. Caswell, J. Ouyang, and G. E. Karniadakis. Active learning of constitutive relation from mesoscopic dynamics for macroscopic modeling of non-newtonian flows. *Journal of Computational Physics*, 363:116–127, 2018.