

Physics-informed Neural Networks with Unknown Measurement Noise

Philipp Pilar

Department of Information Technology, Uppsala University, Sweden

PHILIPP.PILAR@IT.UU.SE

Niklas Wahlström

Department of Information Technology, Uppsala University, Sweden

NIKLAS.WAHLSTROM@IT.UU.SE

Editors: A. Abate, M. Cannon, K. Margellos, A. Papachristodoulou

Abstract

Physics-informed neural networks (PINNs) constitute a flexible approach to both finding solutions and identifying parameters of partial differential equations. Most works on the topic assume noiseless data, or data contaminated with weak Gaussian noise. We show that the standard PINN framework breaks down in case of non-Gaussian noise. We give a way of resolving this fundamental issue and we propose to jointly train an energy-based model (EBM) to learn the correct noise distribution. We illustrate the improved performance of our approach using multiple examples.

Keywords: physics-informed neural networks, energy-based models, non-Gaussian noise, system identification

1. Introduction

While the idea of using neural networks to solve partial differential equations (PDEs) dates back to the work of [Lagaris et al. \(1998\)](#), the field has received renewed attention ([Cuomo et al., 2022](#); [Karniadakis et al., 2021](#); [Markidis, 2021](#); [Blechschmidt and Ernst, 2021](#)) due to the seminal work of [Raissi et al. \(2019\)](#), in which they introduced physics-informed neural networks (PINNs). Both the forward problem, where the solution to a PDE is learned given the boundary conditions, as well as the inverse problem, where parameters of the PDE are to be inferred from measurements, can be solved with the PINN approach.

A multitude of applications for PINNs in science have already been considered: [Cai et al. \(2021a\)](#) review the application of PINNs to fluid mechanics, [Cai et al. \(2021b\)](#) review their application to heat transfer problems, [Yazdani et al. \(2020\)](#) utilize PINNs to infer parameters for systems of differential equations in systems biology, [Mao et al. \(2020\)](#) investigate the applicability of PINNs to high-speed flows, and [Sahli Costabal et al. \(2020\)](#) utilize PINNs to take into account wave propagation dynamics in cardiac activation mapping.

The main advantage of PINNs over traditional solvers lies in their flexibility, especially when considering the inverse problem ([Karniadakis et al., 2021](#)): as neural networks, they have the capacity for universal function approximation, they are mesh-free, and they can directly be applied to very different kinds of PDEs, without the need to use a custom solver. When dealing with the inverse problem, parameters are learned from data and the question as to the effect of noisy data on the quality of the estimates arises naturally. However, most of the existing work on PINNs assumes either noiseless data, or data contaminated with weak Gaussian noise. While some research has been done on the effects of noisy data in PINN training (see [Section 2](#)), they still consider either Gaussian noise or are focused on uncertainty quantification.

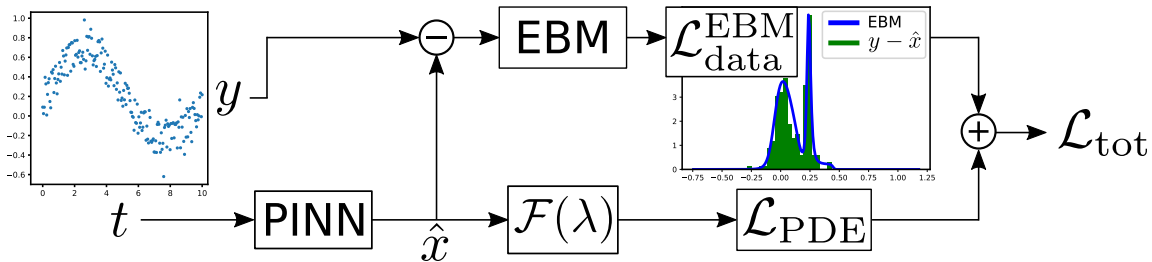


Figure 1: The PINN-EBM approach: the inputs t are passed through the PINN to obtain the PINN predictions \hat{x} . The differential operator $\mathcal{F}(\lambda)$ is then applied to \hat{x} to obtain the PDE residuals, which are subsequently utilized to calculate the PDE loss \mathcal{L}_{PDE} . At the same time, the residuals between the noisy measurements y and \hat{x} are formed, serving as noise estimates. The EBM is trained on these estimates in order to learn the noise PDF, which can in turn be utilized to compute the likelihood of the measurements serving as data loss $\mathcal{L}_{\text{data}}^{\text{EBM}}$. Finally, the two loss terms are combined to form the total loss \mathcal{L}_{tot} . Both PINN and EBM, as well as the PDE parameters λ , can be trained by backpropagating \mathcal{L}_{tot} .

In this work, we consider the inverse problem for the case of measurements contaminated with homogeneous non-Gaussian noise of unknown form. The least-squares loss, which is commonly employed as data loss in PINNs, is known to perform poorly in this case (Constable, 1988; Akkaya and Tiku, 2008). We give a way of mitigating this issue by suitably modeling the noise distribution. A high-level illustration of our method is given in Fig. 1: we employ an energy-based model (EBM) to learn the noise probability density function (PDF) jointly with the PINN. This PDF is utilized to estimate the likelihood of the measurements under our model, which in turn serves as data loss.

Non-Gaussian measurement noise of unknown form can appear in a variety of applications (Johnson and Rao, 1991). Measurements of geomagnetic fields may exhibit asymmetric, long-tailed noise (Constable, 1988), data in astrophysics or seismology may be contaminated with impulsive noise (Weng and Barner, 2005), and not accounted for systematic errors in the measurement procedure may give rise to bias in the noise (Barlow, 2002). For example, the noise distribution considered in Fig. 1 can be interpreted as resulting from the case where measurements from multiple sensors have been merged into one dataset. One of the sensors produces biased measurements, giving rise to the second peak away from zero. Then the PINN-EBM allows for solving the PDE correctly together with the simultaneous detection of previously unrecognized systematic errors in the data.

2. Related Work

Prior research has been done on PINNs in case of noisy measurements, although typically only Gaussian noise is considered. In Yang et al. (2021), the framework of Bayesian neural networks (Goan and Fookes, 2020) is combined with the PINN framework, in order to obtain uncertainty estimates when training PINNs on noisy data. Psaros et al. (2023) discuss different methods of uncertainty quantification and Zou et al. (2023) consider the case where both inputs and outputs are noisy. Bajaj et al. (2023) introduce the GP-smoothed PINN, where a Gaussian process (GP) (Rasmussen and Williams, 2006) is utilized to ameliorate noisy initial value data and make the PINN training more robust in this situation. In Chen et al. (2021), the PINN-SR method is introduced

which can be employed to determine governing equations from scarce and noisy data. In contrast to these works, we give a way of taking into account unknown, non-Gaussian noise in the PINN and provide a training procedure for this case.

While EBMs are commonly employed for classification (LeCun et al., 2006) and image generation (Du and Mordatch, 2019), they have also been successfully applied to regression problems; applications include object detection (Gustafsson et al., 2020a) and visual tracking (Danelljan et al., 2020). In Gustafsson et al. (2020b), different methods for training EBMs for regression are discussed. In our work we also consider regression tasks and to the best of our knowledge, our paper is the first to combine EBMs and PINNs. The standard EBM approach to regression would not take the physical knowledge in form of the differential equation into account.

3. Background

In this section we give a brief introduction to the two distinct methods which we combine in our work: physics-informed neural networks and energy-based models.

3.1. Physics-informed Neural Networks (PINNs)

The PINN-framework (Raissi et al., 2019) can be used as an alternate approach to solving PDEs, other than the standard, mesh-based solvers. The objective is to numerically determine the solution to a differential equation $\mathcal{F}x(t) = 0$, where \mathcal{F} denotes the differential operator defining the PDE, $x(t)$ the solution to the differential equation, and t the input; both x and t can be multidimensional. In the PINN approach, we now employ a neural network to parameterize the numerical solution $\hat{x}(t) = \hat{x}(t|\theta_{\text{PINN}})$ of the PDE, where θ_{PINN} denotes the weights of the PINN which are to be optimized.

In this paper, we consider the inverse problem: we are given a dataset $\mathcal{D}_d = \{t_d^i, y_d^i\}_{i=1}^{N_d}$ of N_d (noisy) measurements $y_d^i = x(t_d^i) + \epsilon$ of the PINN solution in addition to the parametric form of the differential operator $\mathcal{F}(\lambda)$, which may contain unknown parameters λ . Furthermore, we choose another set of N_c so-called collocation points $\mathcal{D}_c = \{t_c^i\}_{i=1}^{N_c}$. They can lie at arbitrary points in the input domain of the PDE and are used to take into account the PDE constraint. As long as collocation points are placed in the areas of interest, they can enable the PINN to extrapolate also to areas without measurements.

When training the PINN, two losses enter into the loss function: the data loss, $\mathcal{L}_{\text{data}}$, evaluating the fit of the PINN prediction with the data, and the PDE loss, \mathcal{L}_{PDE} , a measure of the fulfillment of the PDE by the PINN solution. In the standard PINN, the least-squares loss is used as data loss,

$$\mathcal{L}_{\text{data}}(\hat{x}, \{t_d, y_d\}_{\text{mb}}) = \frac{1}{N'_d} \sum_{i=1}^{N'_d} (\hat{x}(t_d^i) - y_d^i)^2, \quad (1)$$

and the squares of the PDE residuals, $f(t) = \mathcal{F}(\lambda)\hat{x}(t) \stackrel{!}{=} 0$, at the collocation points serve as PDE loss,

$$\mathcal{L}_{\text{PDE}}(\mathcal{F}, \hat{x}, \{t_c\}_{\text{mb}}) = \frac{1}{N'_c} \sum_{i=1}^{N'_c} f(t_c^i)^2. \quad (2)$$

Here, N'_d and N'_c denote the number of data points t_d and collocation points t_c , respectively, in the current mini-batch (mb). The PINN is then trained by minimizing the total loss $\mathcal{L}_{\text{tot}} = \mathcal{L}_{\text{data}} +$

$\omega \mathcal{L}_{\text{PDE}}$ with respect to the parameters θ_{PINN} and λ . Unless stated otherwise, the weighting factor $\omega = 1$. Utilizing this loss function, the PINN is optimized via some variation of gradient descent.

3.2. Energy-based Models (EBMs)

EBMs constitute a powerful method of learning probability densities from data (LeCun et al., 2006). They are frequently employed for image generation and modeling (Du and Mordatch, 2019; Nijkamp et al., 2020), and have successfully been applied to the domain of regression (see Section 2). The EBM can retain all of the flexibility of a neural network, via the following parametrization:

$$p(y|t, \hat{h}) = \frac{e^{\hat{h}(t,y)}}{Z(t, \hat{h})}, \quad Z(t, \hat{h}) = \int e^{\hat{h}(t,\tilde{y})} d\tilde{y}, \quad (3)$$

where $\hat{h}(t, y) = \hat{h}(t, y|\theta_{\text{EBM}})$ is the (scalar) output of the neural network with weights θ_{EBM} .

This expressive capacity, however, comes at the cost of a more complicated training procedure, since the partition function $Z(t, \hat{h})$ will typically be analytically intractable. In case of a high-dimensional y , it will also be expensive or infeasible numerically; Monte Carlo methods are often employed to approximate the intractable integral in (3). Various ways of training EBMs for regression are given in Gustafsson et al. (2020b).

In this paper, we will choose the approach of directly minimizing the negative log-likelihood (NLL), which can be written as

$$\text{NLL}(\{t_d, y_d\}_{\text{mb}}, \hat{h}) = -\log \left(\prod_{i=1}^{N'_d} p(y_d^i | t_d^i, \hat{h}) \right) = \sum_{i=1}^{N'_d} \log Z(t_d^i, \hat{h}) - \hat{h}(t_d^i, y_d^i). \quad (4)$$

In our case, the evaluation of the partition function to high accuracy remains tractable by utilizing numerical integration, since the noise distribution is one-dimensional. Higher-dimensional outputs are still possible, as long as the noise is uncorrelated. Also, note that this does not constitute a restriction on the dimensionality of the inputs t of the PDE.

4. Problem Formulation

We consider the following setup, consisting of the differential equation and a measurement equation:

$$\mathcal{F}(\lambda)x(t) = 0, \quad y(t) = x(t) + \epsilon, \quad (5)$$

where the parametric form of the differential operator $\mathcal{F}(\lambda)$ is given, as well as N_d measurements $y(t)$ of the corresponding solution $x(t)$ contaminated with homogeneous measurement noise ϵ . In this work, we aim to combine the PINN and the EBM framework in order to solve the inverse problem in case of measurements contaminated with non-Gaussian and non-zero mean noise.

4.1. PINNs in Case of Non-zero Mean Noise

To see why noise with non-zero mean is problematic in the standard PINN framework, consider the loss function $\mathcal{L}_{\text{tot}} = \mathcal{L}_{\text{data}} + \omega \mathcal{L}_{\text{PDE}}$ (compare Section 3.1):

$$\mathcal{L}_{\text{tot}} = \frac{1}{N'_d} \sum_{i=1}^{N'_d} (\hat{x}(t_d^i) - y_d^i)^2 + \frac{\omega}{N'_c} \sum_{i=1}^{N'_c} f(t_c^i)^2. \quad (6)$$

A major challenge when training PINNs, even in case of Gaussian noise, lies in the fact that data and PDE loss counteract each other, since the first term would encourage overfitting towards each data point, whereas the second term would push the solution towards fulfilling the PDE.

Let us first consider the case of a zero-mean noise distribution in the limit of infinite data generated according to (5), and assume that the PINN has sufficient capacity to represent the solution accurately. Then both loss terms would be minimal for $\hat{x}(t) = x(t)$: in the limit, the data loss would become $\lim_{N_d \rightarrow \infty} \mathcal{L}_{\text{data}} = \mathbb{E}[(\hat{x} - x - \epsilon)^2] = \mathbb{E}[(\hat{x} - x)^2] + \text{Var}_\epsilon$, and hence $\hat{x} = x$. The PDE loss would then obviously also be minimal (more precisely zero) for $\hat{x} = x$, due to (5).

For noise with non-zero mean, however, the problem of mismatch between the two loss terms is not simply an effect of finite data, which could be treated via an adequate regularization procedure, but has deeper roots. Considering again the limit of infinite data, we have $\lim_{N_d \rightarrow \infty} \mathcal{L}_{\text{data}} = \mathbb{E}[(\hat{x} - x - \epsilon)^2] = \mathbb{E}[(\hat{x} - x - \mu_\epsilon)^2] + \text{Var}_\epsilon$, and the minimizer of the data loss becomes $\hat{x} = x + \mu_\epsilon$, where μ_ϵ denotes the mean of the noise distribution. For the PDE loss, on the other hand, the minimizer remains $\hat{x} = x$ (assuming that the correct parameters λ of $\mathcal{F}(\lambda)$ are known).

The two losses are now counteracting each other and the optimization procedure will produce some compromise between them, even when the parameters λ of the PDE are known exactly a priori. In the case where the parameters of $\mathcal{F}(\lambda)$ are unknown, the optimization procedure will tend to converge to parameter values which result in a lower least-squares loss than the correct ones might.

4.2. Reconciling the Losses

A simple way of restoring consistency between the two losses in the limit of infinite data is to add an offset parameter θ_0 to the PINN prediction \hat{x} in the data loss term, which is supposed to learn the bias in the noise term. The data loss then reads $\lim_{N_d \rightarrow \infty} \mathcal{L}_{\text{data}} = \mathbb{E}[(\hat{x} + \theta_0 - x - \epsilon)^2] = \mathbb{E}[(\hat{x} + \theta_0 - x - \mu_\epsilon)^2] + \text{Var}_\epsilon$ and the optimum $\hat{x} = x$, $\theta_0 = \mu_\epsilon$ would be compatible with the PDE loss. Hence the optimization procedure could converge to the correct solution, in principle.

However, since the maximum likelihood estimator is asymptotically optimal only when using the correct likelihood (Wasserman, 2004), this data loss still does not constitute the best option (except for the case of Gaussian noise with non-zero mean). Furthermore, in the practical case of finite data, outliers may have an outsized effect when using the least-squares loss. Hence, a way of taking the non-Gaussianity of the noise into account explicitly would likely further improve the speed of optimization as well as the final learning outcome. In the next section, we demonstrate how EBMs can be used for this purpose.

4.3. Using EBMs to Learn the Noise Distribution

In order to also take the shape of non-Gaussian noise into account, we can choose to use a different data loss in (6). A natural choice would be the log-likelihood of the measurements, given the noise distribution. In case of Gaussian noise, this would again result in a least-squares loss term, plus a constant. However, since in our problem setup we do not know the form of the noise a priori, we need to determine the shape of the noise distribution jointly with the PINN solution.

Algorithm 1: Training the PINN-EBM

Input : N_{tot} ; N_{EBM} ; PINN $\hat{x}(\cdot|\theta_{\text{PINN}})$, EBM $\hat{h}(\cdot|\theta_{\text{EBM}})$, data points \mathcal{D}_d , collocation points \mathcal{D}_c , differential operator $\mathcal{F}(\lambda)$, weighting factor ω

Output: optimized θ_{PINN} , θ_{EBM} , λ

$i = 0$

while $i \leq N_{\text{tot}}$ **do**

Draw a minibatch of data points $\{t_d, y_d\}_{\text{mb}}$ and of collocation points $\{t_c\}_{\text{mb}}$

if $i < i_{\text{EBM}}$ **then**

| Calculate the data loss $\mathcal{L}_{\text{data}}(\hat{x}, \{t_d, y_d\}_{\text{mb}})$ according to (1)

else

if $i = i_{\text{EBM}}$ **then** initialize the EBM for N_{EBM} iterations

| Calculate the data loss $\mathcal{L}_{\text{data}}^{\text{EBM}}(\{y_d - \hat{x}(t_d)\}_{\text{mb}}, \hat{h})$ from (7)

end

if $i \geq i_{\text{EBM}}$ **then** $\omega' = \omega$ **else** $\omega' = 1$

Calculate the PDE loss $\mathcal{L}_{\text{PDE}}(\mathcal{F}, \hat{x}, \{t_c\}_{\text{mb}})$ according to (2)

Compute the total loss $\mathcal{L}_{\text{tot}} = \mathcal{L}_{\text{data}} + \omega' \mathcal{L}_{\text{PDE}}$

Calculate the gradient $\nabla_{\theta} \mathcal{L}_{\text{tot}}$ and update $\theta = \{\theta_{\text{PINN}}, \theta_{\text{EBM}}, \lambda\}$

$i += 1$

end

To this end, we train an EBM and employ it to obtain the negative log-likelihood of the data, given both our models. For training, we then utilize the following loss function:

$$\mathcal{L}_{\text{tot}} = \mathcal{L}_{\text{data}}^{\text{EBM}}(\{y_d - \hat{x}(t_d)\}_{\text{mb}}, \hat{h}|\theta_{\text{PINN}}, \theta_{\text{EBM}}) + \omega \mathcal{L}_{\text{PDE}}(\mathcal{F}, \hat{x}, \{t_c\}_{\text{mb}}|\theta_{\text{PINN}}, \lambda), \quad (7)$$

where $\mathcal{L}_{\text{data}}^{\text{EBM}}(\cdot, \hat{h}) = \frac{1}{N'_d} \text{NLL}(\cdot, \hat{h}) = \log Z(\hat{h}) - \frac{1}{N'_d} \sum_{i=1}^{N'_d} \hat{h}(\cdot)$, utilizing (4); note that \hat{h} and hence the NLL are now independent of t , because we consider homogeneous noise. Since we do not know the actual magnitude of the noise at each data point, we use the residuals $y_d - \hat{x}(t_d)$ between the current PINN prediction and the measurements as our best guess. These estimates of the noise values then serve as training data for the EBM. In other words, we employ an unconditional EBM to model the PDF of the residuals. With the loss (7), both models are trained jointly until convergence.

4.4. Training PINN and EBM Jointly

In Algorithm 1, the training procedure is summarized. Since both PINN and EBM need to be trained in parallel, the optimization procedure can be very challenging. In our experiments, it proved advantageous to start with the standard PINN loss in order to obtain a solution close to the data, and to only subsequently, after i_{EBM} iterations, switch to the EBM loss, in order to fine-tune the solution. Before the EBM loss is used for the first time, we initialize the EBM by training it for N_{EBM} iterations on the current noise estimates $y_d - \hat{x}(t_d)$, while keeping all parameters except θ_{EBM} fixed. When we choose to use a weighting parameter $\omega \neq 1$, we also wait until the EBM has been initialized before switching to the new value. In total, we train for N_{tot} iterations.

Initializing the EBM only later during the training process is advantageous for two reasons: firstly, the least-squares loss is still a sensible first guess, so initializing the EBM on the residuals

stemming from the pretrained network will let it start out with a reasonable form of the likelihood instead of a random one. Secondly, it will give us a better idea about the range in which the residuals between the PINN prediction and the data lie, allowing for a better normalization of the inputs to the EBM and hence more efficient training.

5. Experiments ¹

In our experiments, we compare the performance of the standard PINN from Section 3.1 to that of the PINN with learnable offset parameter θ_0 and the combination of PINN and EBM. We will refer to these models as PINN, PINN-off and PINN-EBM.

In our experiments, we consider homogeneous noise in a variety of shapes: Gaussian noise $p(\epsilon) = \mathcal{N}(\epsilon|0, 2.5^2)$, a uniform distribution $p(\epsilon) = \mathcal{U}[0, 10]$ and a Gaussian mixture of the form $p(\epsilon) = \frac{1}{3} (\mathcal{N}(\epsilon|0, 2^2) + \mathcal{N}(\epsilon|4, 4^2) + \mathcal{N}(\epsilon|8, 0.5^2))$. We also consider the same Gaussian mixture, shifted to have zero mean. We refer to these noise distributions as G, u, 3G, and 3G0. The values given here are scaled to the example of the exponential function in Section 5.1. For other experiments, we use rescaled versions of the same distributions.

When evaluating the performance of our models, the following metrics are considered: the absolute error in the learned values of the PDE parameters ($|\Delta\lambda|$), the root-mean-square error (RMSE) between \hat{x} and x in the range of the validation data, the negative log-likelihood (NLL) of the validation data according to the models, and the square values of the PDE residuals (f^2) in the range of the training data. In practice, the NLL is the most relevant performance metric: contrary to the RMSE and $|\Delta\lambda|$, it can also be calculated when the true solution is not known. While f^2 can also be calculated without knowledge of the true solution, it only measures how well the learned PDE is fulfilled and not necessarily the correct one.

Additional results are available in the extended report on arXiv (Pilar and Wahlström, 2022): the method is tested on another PDE, and the model performance is investigated as a function of various dataset properties and training parameters.

5.1. Toy Problem

As a first toy problem, we consider the following simple differential equation,

$$\dot{x}(t) - \lambda x(t) = 0, \quad (8)$$

where $\lambda = 0.3$ and with the true solution being the exponential function, $x(t) = e^{\lambda t}$.

In Fig. 2, the results for an individual run of PINN, PINN-off, and PINN-EBM are compared; blue dots represent training data and red dots validation data. From this plot, it is clear that PINN converges to a wrong solution, and PINN-off gives only a slightly better prediction, whereas the PINN-EBM result is very close to the true curve.

In Table 1, performance metrics of the different models as obtained for different noise forms and with $\omega = 1$ are given. The results have been averaged over ten runs with different realizations of the measurement noise.




Figure 2: The exponential differential equation (8) with Gaussian mixture noise (3G). Results for one individual run, where the blue dots represent training data and the red dots validation data.

1. The code for this project is available at <https://github.com/ppilar/PINN-EBM>.

We start by discussing the results for Gaussian mixture noise (3G). It is evident that PINN performs poorly on all metrics. While PINN-off learns the correct parameter λ better on average, the variance of its predictions is very high, thereby limiting its use. The PINN-EBM estimate of the parameter is significantly less variable than that of PINN-off, and is very close to the correct value of λ . In accordance with the quality of the parameter estimates, PINN-EBM clearly outperforms both PINN and PINN-off when considering the RMSE and NLL; it achieves the lowest RMSE as well as the lowest NLL on the validation data.

When considering f^2 , PINN needs to trade off accuracy in this metric with minimizing the data loss (1), due to the inconsistency between the losses discussed in Section 4.1.

Three other noise forms are considered in the table: for the uniform (u) noise distribution, the results are very similar to the Gaussian mixture. In case of Gaussian (G) noise, PINN performs best, which makes sense since it is implicitly using the correct likelihood. PINN-EBM, on the other hand, performs a bit worse as there is some overfit to the noise in the data.

For Gaussian mixture noise with zero mean (3G0), where the PDE and data loss are no longer misaligned for the PINN, PINN-EBM still outperforms PINN, although with a lesser margin. This shows that PINN-EBM can successfully mitigate the effects of non-Gaussian noise and significantly improves the quality of the parameter estimates, as well as that of the solution to the regression task. While PINN-off manages to improve upon PINN in the case of non-zero mean noise, the high variance in its predictions makes it less reliable than PINN-EBM.

In Fig. 3, the impact of various training parameters and dataset properties on the model performance is shown, when considering 3G noise (left plot). It is apparent that the weighting parameter ω can significantly impact the model performance. From the two plots on the right, we observe that the performances of the three models converge as either the number of training points or the noise strength is reduced; for higher values, the impact of non-Gaussian noise is larger and PINN-EBM outperforms. More in-depth discussions of these results can be found in the appendices of [Pilar and Wahlström \(2022\)](#).

noise		PINN-EBM	PINN-off	PINN
3G	100 $ \Delta\lambda $	1.2 \pm 1.1	4.0 \pm 3.6	10.2 \pm 1.5
	RMSE	1.9 \pm 1.4	6.1 \pm 6.8	8.3 \pm 1.9
	NLL	3.5 \pm 1.0	7.9 \pm 9.0	8.9 \pm 2.8
	100 f^2	0.5 \pm 0.3	30.9 \pm 11.5	48.8 \pm 17.1
u	100 $ \Delta\lambda $	1.3 \pm 0.3	3.3 \pm 2.5	12.3 \pm 0.7
	RMSE	1.9 \pm 0.8	4.0 \pm 3.7	9.1 \pm 1.1
	NLL	4.8 \pm 1.1	6.5 \pm 6.8	19.1 \pm 3.8
	100 f^2	0.7 \pm 0.3	16.3 \pm 4.8	37.9 \pm 12.6
G	100 $ \Delta\lambda $	2.5 \pm 1.8	2.0 \pm 1.3	1.1 \pm 0.9
	RMSE	3.4 \pm 2.1	3.2 \pm 2.7	2.8 \pm 2.5
	NLL	5.1 \pm 2.5	5.1 \pm 3.6	4.2 \pm 2.5
	100 f^2	0.4 \pm 0.3	10.8 \pm 6.0	11.7 \pm 6.6
3G0	100 $ \Delta\lambda $	1.2 \pm 1.0	5.3 \pm 3.2	2.2 \pm 2.3
	RMSE	2.2 \pm 1.4	7.9 \pm 6.9	5.0 \pm 5.1
	NLL	3.7 \pm 1.3	9.8 \pm 11.2	5.4 \pm 4.0
	100 f^2	0.6 \pm 0.4	32.2 \pm 15.1	34.0 \pm 15.5

Table 1: Results for the exponential differential equation (8). The entries in the table are averages obtained over 10 runs plus-or-minus one standard deviation. Bold font highlights best performance.




Figure 3: **Left:** The Gaussian mixture (3G) noise distribution, together with examples of the PDFs learned by the EBM. **Middle left:** Performance of the models as a function of the weighting factor ω . **Middle right:** Performance of the models as a function of the number of training points N_d . **Right:** Performance of the models as a function of the noise strength (the distribution in the left plot is scaled by the factor f_n). The standard values are $\omega = 1$, $N_d = 200$, and $f_n = 1$. The averages of 5 runs are depicted, plus-or-minus one standard deviation.

5.2. Navier-Stokes Equations

For our next experiment, we consider the example of incompressible flow past a circular cylinder and use the dataset from [Raissi et al. \(2019\)](#), to which we add Gaussian mixture (3G) noise. This setup can be described by the 2D Navier-Stokes equations,

$$u_t + \lambda_1(uu_x + vv_y) + p_x - \lambda_2(u_{xx} + u_{yy}) = 0, \quad (9a)$$

$$v_t + \lambda_1(uv_x + vv_y) + p_y - \lambda_2(v_{xx} + v_{yy}) = 0, \quad (9b)$$

where u and v denote the components of the velocity field and p the pressure; the subscripts indicate derivatives with respect to the different dimensions of the inputs $t = (t, x, y)$. Here, the PINN does not output u and v directly, but instead models the potential ψ and the pressure p ; then $u = \psi_y$ and $v = -\psi_x$. The true values of the parameters are $\lambda_1 = 1$ and $\lambda_2 = 0.01$.

The results are given in Table 2. It is apparent that PINN-EBM learns λ_1 better than PINN and PINN-off whereas the difference is smaller for λ_2 . PINN-EBM also performs better in terms of RMSE and NLL. Here, we have chosen $\omega = 50$ for PINN-EBM since the NLL on the validation data decreased when the PDE loss was weighted more strongly (see Fig. 4). For PINN and PINN-off, increasing the parameter did not lead to notable changes in the NLL and we kept $\omega = 1$.

5.3. Implementation Details

For both PINN and EBM, fully-connected neural networks with tanh activation function were used. Both inputs and outputs of the networks were normalized to their expected ranges. For the PINN, 4 layers of width 40 were chosen for the exponential equation; in case of the Navier-Stokes equations, 8 layers of width 20 were used instead. For the EBM, 3 layers of width 5 were used and a dropout layer with factor 0.5 was inserted before the last layer.

The PINN was trained for 40.000 iterations in case of the exponential equation (Table 1) and 100.000 iterations for the Navier-Stokes equations (Table 2). When evaluating the impact of training parameters (Figs. 3, 4), 50.000 iterations were used for both of them. The EBM was initialized for $N_{\text{EBM}} = 2000$ iterations, after $i_{\text{EBM}} = 4000$ and $i_{\text{EBM}} = 10000$ iterations, respectively, for the exponential equation and the Navier-Stokes equations. Convergence was checked by confirming




Figure 4: The NLL on the validation data as a function of the parameter ω , weighting the PDE loss. The averages of 5 runs are depicted, plus-or-minus one standard deviation.

noise		PINN-EBM	PINN-off	PINN
3G	$100 \Delta\lambda_1 $	1.19 \pm 0.67	2.92 \pm 0.47	23.10 \pm 0.11
	$100 \Delta\lambda_2 $	0.04 \pm 0.03	0.09 \pm 0.05	0.08 \pm 0.06
	RMSE	0.06 \pm 0.01	0.11 \pm 0.01	0.20 \pm 0.02
	NLL	-0.03 \pm 0.08	0.15 \pm 0.07	0.40 \pm 0.30
	$100f^2$	0.04 \pm 0.00	0.10 \pm 0.01	0.18 \pm 0.01

Table 2: Results for the Navier Stokes equations (9) in case of Gaussian mixture (3G) noise. The entries in the table are averages obtained over 5 runs plus-or-minus one standard deviation. Bold font highlights best performance.

that the values of the PDF tended toward zero on both ends of the distribution; if the check failed, the initialization was repeated with an increased number of iterations until it passed.

The datasets employed in Section 5.1 contained 200 training points, 50 validation points and 2000 collocation points. For the Navier-Stokes example in Section 5.2, 4000 training points, 1000 validation points and 4000 collocation points were used. The collocation points were generated on a uniform grid. The Adam optimizer with a learning rate of 0.002 was used for both PINN and EBM. In case of the Navier-Stokes equations, the learning rate was reduced by a factor of 0.3 after 80 000 iterations. The batch size for data points was 200 (i.e. full batches were used for the exponential equation) and the batch size for collocation points was 100. Training the PINN-EBM typically took 30-50% longer than the standard PINN.

6. Conclusions and Future Work

In this paper, we demonstrated that the standard PINN fails in case of non-zero mean measurement noise and we proposed the PINN-EBM to resolve this problem; utilizing an EBM to learn the homogeneous noise distribution allows the PINN to produce good results also in case of non-zero mean and non-Gaussian noise. Using several examples, ranging from a simple toy problem to the complex Navier-Stokes equations, we demonstrated the capabilities of our method and showed that it outperforms the standard PINN by a significant margin in case of non-Gaussian noise. In addition to determining the correct PDE solution, the PINN-EBM also allows for the identification of the true noise distribution, which may result in novel insights into the measurement procedure.

In the future, it would be interesting to investigate combining the PINN-EBM with other improvements to the PINN framework, such as adaptive weighting schemes (McClenny and Braganeto, 2020; Xiang et al., 2022), adaptive collocation point placement (Wu et al., 2023), scheduling approaches (Krishnapriyan et al., 2021; Wang et al., 2024) or additional loss terms (Yu et al., 2022). Extending the approach to certain kinds of heteroscedastic noise and investigating the applicability of the approach to other frameworks such as neural ODEs (Chen et al., 2018) constitute promising avenues for future research. Finally, it would be interesting to apply the method to suitable real-world datasets.

Acknowledgements

The work is financially supported by the Swedish Research Council (VR) via the project *Physics-informed machine learning* (registration number: 2021-04321) and by the *Kjell och Märta Beijer Foundation*. We would also like to thank Fredrik K. Gustafsson for valuable feedback.

References

- Ayşen D Akkaya and Moti L Tikunov. Robust estimation in multiple linear regression model with non-Gaussian noise. *Automatica*, 44(2):407–417, 2008.
- Chandrajit Bajaj, Luke McLennan, Timothy Andeen, and Avik Roy. Recipes for when physics fails: recovering robust learning of physics informed neural networks. *Machine Learning: Science and Technology*, 4(1):015013, 2023.
- Roger Barlow. Systematic errors: facts and fictions. *arXiv preprint hep-ex/0207026*, 2002.
- Jan Blechschmidt and Oliver G. Ernst. Three ways to solve partial differential equations with neural networks — a review. *GAMM-Mitteilungen*, 44(2):e202100006, 2021.
- Shengze Cai, Zhiping Mao, Zhicheng Wang, Minglang Yin, and George Em Karniadakis. Physics-informed neural networks (pinns) for fluid mechanics: a review. *Acta Mechanica Sinica*, 37(12):1727–1738, Dec 2021a.
- Shengze Cai, Zhicheng Wang, Sifan Wang, Paris Perdikaris, and George Em Karniadakis. Physics-Informed Neural Networks for Heat Transfer Problems. *Journal of Heat Transfer*, 143(6), 04 2021b. 060801.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Zhao Chen, Yang Liu, and Hao Sun. Physics-informed learning of governing equations from scarce data. *Nature Communications*, 12(1):6136, Oct 2021.
- Catherine G. Constable. Parameter estimation in non-Gaussian noise. *Geophysical Journal International*, 94(1):131–142, 1988.
- Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what’s next. *Journal of Scientific Computing*, 92(3):88, Jul 2022.
- Martin Danelljan, Luc Van Gool, and Radu Timofte. Probabilistic regression for visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. *Advances in Neural Information Processing Systems*, 32, 2019.
- Ethan Goan and Clinton Fookes. *Bayesian Neural Networks: An Introduction and Survey*, pages 45–87. Springer International Publishing, Cham, 2020.

- Fredrik K. Gustafsson, Martin Danelljan, Goutam Bhat, and Thomas B. Schön. Energy-based models for deep probabilistic regression. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 325–343, Cham, 2020a. Springer International Publishing.
- Fredrik K. Gustafsson, Martin Danelljan, Radu Timofte, and Thomas B. Schön. How to train your energy-based model for regression. *CoRR*, abs/2005.01698, 2020b.
- Don H Johnson and P Srinivasa Rao. On the existence of Gaussian noise (signal modelling). In *Acoustics, Speech, and Signal Processing, IEEE International Conference on*, pages 1673–1676. IEEE Computer Society, 1991.
- George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W. Mahoney. Characterizing possible failure modes in physics-informed neural networks. In *Advances in Neural Information Processing Systems*, 2021.
- Isaac E. Lagaris, Aristidis Likas, and Dimitrios I. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE transactions on neural networks*, 9(5):987–1000, 1998.
- Yann LeCun, Sumit Chopra, Raia Hadsell, Marc’Aurelio Ranzato, and Fu Jie Huang. A tutorial on energy-based learning. *Predicting structured data*, 2006.
- Zhiping Mao, Ameya D. Jagtap, and George Em Karniadakis. Physics-informed neural networks for high-speed flows. *Computer Methods in Applied Mechanics and Engineering*, 360:112789, 2020.
- Stefano Markidis. The old and the new: Can physics-informed deep-learning replace traditional linear solvers? *Frontiers in Big Data*, 4, 2021.
- Levi McClenny and Ulisses Braga-Neto. Self-adaptive physics-informed neural networks using a soft attention mechanism. *arXiv:2009.04544*, 2020.
- Erik Nijkamp, Mitch Hill, Tian Han, Song-Chun Zhu, and Ying Nian Wu. On the anatomy of MCMC-based maximum likelihood learning of energy-based models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:5272–5280, Apr. 2020.
- Philipp Pilar and Niklas Wahlström. Physics-informed neural networks with unknown measurement noise. *arXiv:2211.15498*, 2022.
- Apostolos F Psaros, Xuhui Meng, Zongren Zou, Ling Guo, and George Em Karniadakis. Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons. *Journal of Computational Physics*, 477:111902, 2023.
- Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

- Carl Edward Rasmussen and Christopher KI Williams. *Gaussian processes for machine learning*. MIT press, 2006.
- Francisco Sahli Costabal, Yibo Yang, Paris Perdikaris, Daniel E. Hurtado, and Ellen Kuhl. Physics-informed neural networks for cardiac activation mapping. *Frontiers in Physics*, 8, 2020.
- Sifan Wang, Shyam Sankaran, and Paris Perdikaris. Respecting causality for training physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 421:116813, 2024.
- Larry Wasserman. *All of statistics: a concise course in statistical inference*, volume 26. Springer, 2004.
- Binwei Weng and Kenneth E Barner. Nonlinear system identification in impulsive environments. *IEEE Transactions on signal processing*, 53(7):2588–2594, 2005.
- Chenxi Wu, Min Zhu, Qinyang Tan, Yadhu Kartha, and Lu Lu. A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 403:115671, 2023.
- Zixue Xiang, Wei Peng, Xu Liu, and Wen Yao. Self-adaptive loss balanced physics-informed neural networks. *Neurocomputing*, 496:11–34, 2022.
- Liu Yang, Xuhui Meng, and George Em Karniadakis. B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data. *Journal of Computational Physics*, 425:109913, 2021.
- Alireza Yazdani, Lu Lu, Maziar Raissi, and George Em Karniadakis. Systems biology informed deep learning for inferring parameters and hidden dynamics. *PLOS Computational Biology*, 16(11):1–19, 11 2020.
- Jeremy Yu, Lu Lu, Xuhui Meng, and George Em Karniadakis. Gradient-enhanced physics-informed neural networks for forward and inverse PDE problems. *Computer Methods in Applied Mechanics and Engineering*, 393:114823, 2022.
- Zongren Zou, Xuhui Meng, and George Em Karniadakis. Uncertainty quantification for noisy inputs-outputs in physics-informed neural networks and neural operators. *arXiv:2311.11262*, 2023.