Physics-Informed Discriminator (PID) for Conditional Generative Adversarial Nets

Arka Daw
Department of Computer Science
Virginia Tech
Blacksburg, VA 24060
darka@vt.edu

M. Maruf
Department of Computer Science
Virginia Tech
Blacksburg, VA 24060
marufm@vt.edu

Anuj Karpatne
Department of Computer Science
Virginia Tech
Blacksburg, VA 24060
karpatne@vt.edu

Abstract

We propose a novel method of incorporating physical knowledge as an additional input to the discriminator of a conditional Generative Adversarial Net (cGAN). Our proposed approach, termed as Physics-informed Discriminator for cGAN (cGAN-PID), is more aligned to the adversarial learning idea of cGAN as opposed to existing methods on incorporating physical knowledge in GANs by adding physics based loss functions as additional terms in the optimization objective of GAN. We evaluate the performance of our model on two toy datasets and demonstrate that our proposed cGAN-PID can be used as an alternative to the existing techniques.

1 Introduction

Although state-of-the-art deep learning models (like Deep Neural Nets, Convolutional Neural Nets, Recurrent Neural Nets, Deep Generative Models, etc.) [8, 9, 11] have experienced enormous success over the last decade in numerous fields, they have often been questioned for their lack of robustness and interpretability. The ability of a model to perform uncertainty quantification (UQ) is often considered as a first step in the direction towards robust and interpretable model learning. Uncertainty quantification can be achieved in multiple ways [2, 4], one of them being, Monte Carlo (MC) Dropouts [5], which can be applied to any predictive deep learning model. Another more natural way of obtaining uncertainty estimates is through generative modeling [12], where a distribution on the outputs can be obtained by varying the built-in random noise vector z. An even challenging task, especially in scientific problems, is the ability of the model to adhere to previously known scientific knowledge. One of the promising lines of research in this direction, referred to as the *physics-guided learning* (PGL) paradigm [6, 7], focuses on guiding the learning algorithm by penalizing the model during training whenever scientific principles are violated in the model outputs, captured using physics loss (PhyLoss) functions. A primary motivation of this work is to incorporate scientific knowledge into deep learning models while performing UQ. Specifically, we are looking at a predictive modeling task, where we want to train a neural network $f_\theta(.)$ to learn the mapping from the inputs x to output y, $f_\theta : x \rightarrow y$ to get to the posterior distribution of the outputs $P(f_\theta(x)|x)$, while minimizing deviations from existing scientific knowledge, i.e., $\min_\theta \text{PhyLoss}(x, f_\theta(x))$.

A naïve appraoch to incorporate physical knowledge while performing UQ is to use the PGL paradigm in conjunction with MC Dropouts to estimate the posterior distribution $P(f_\theta(x)|x)$. However, as demonstrated in a recent work [3], this approach can fail miserably as a model trained using PGL

Third Workshop on Machine Learning and the Physical Sciences (NeurIPS 2020), Vancouver, Canada.

can easily become physically inconsistent with minor perturbations in model weights such as those introduced by MC Dropout. In case of generative models, physics loss has already been incorporated in the generator loss in existing works [12], where they consider the following optimization objective for generative adversarial networks (GANs): $\min_G \mathbb{E}_{z\sim p_z(z)}[\log(1 - D(G(x, z), x))] + \lambda_{phy}$PhyLoss(x,G(x,z)) and $\max_D \mathbb{E}_{z\sim p_z(z)}[\log(1-D(G(x, z), x))]+\mathbb{E}_{y\sim p_{data}(y)}[\log D(y, x)]$, for the generator G and discriminator D, respectively. However, this does not take the full advantage of the adversarial learning strategy that makes GANs such powerful tools. In particular, adding PhyLoss as another term in the objective function of generator can be interpreted as adding a separate auxiliary task to the primary task of adversarial learning by minimizing GAN-based loss terms. Note that this is a common characteristic of the PGL paradigm where PhyLoss is simply added as an auxillary task (with a constant trade-off hyper-parameter) to the primary task of minimizing data-driven loss fucntions such as the mean-squared error(MSE) or binary-cross entropy (BCE) loss.

In this paper, we explore a novel direction to incorporate physical knowledge directly into the adversarial learning strategy of conditional GAN (cGAN) models by using a transformation of the physics loss as an additional input in the discriminator along with the input x. We refer to our proposed approach as physics-informed discriminator for cGAN (cGAN-PID). We demonstrate that cGAN-PID shows good generalization performance on two toy datasets for both in-distribution and out-of-distribution predictive tasks, while showing reduced sensitivity to the choice of trade-off hyper-parameter controlling the importance of physics loss in the learning procedure.

2 Proposed Method

We use a conditional GAN (cGAN) formulation in this paper for predictive learning. In a cGAN setting, the generator G learns the mapping from the input vector x and some random noise vector z to y, $G : (x, z) \rightarrow y$. The generator is trained to make predictions $\hat{y} = G(x, z)$ which cannot be distinguished from "real" outputs by an adversarially trained discriminator D. The discriminator D, on the other hand, is trained to detect the generator's "fake" predictions as well as possible. We can also compute physical inconsistency PhyLoss$(x, \hat{y})$ of the generated predictions using the physical knowledge which are often available in the form of constraints.

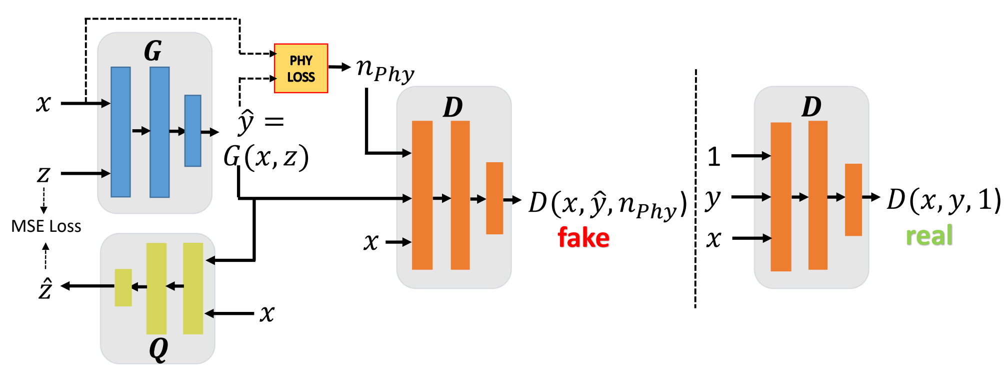

Figure 1: A schematic representation of the training procedure for the proposed Physics-informed Discriminator for cGAN (cGAN-PID).

We propose a novel more natural way of incorporating physics into the adversarial learning framework of cGAN where we provide a measure of physical consistency n_{phy} for each generated prediction as an additional input to the discriminator (See Figure 1). The motivation of adding physics into the discriminator is to aid the discriminator such that it not only distinguishes between a real and a fake sample by learning from the underlying data distribution but also uses the additional physical knowledge (n_{phy}) to make the distinction. A simple way of defining n_{phy} is as follows:

$$n_{phy} = f(\lambda_{phy}\text{PhyLoss}(x, \hat{y})), \quad (1)$$

where $f(.)$ is some monotonic transformation function and λ_{phy} is a trade-off hyperparameter.

In this paper, we use an exponentially decaying transformation function for $f(.)$, i.e., $n_{phy} = e^{-\lambda_{phy}\text{PhyLoss}(x,\hat{y})}$. Note that if PhyLoss$(x, \hat{y}) = 0$, i.e., the predictions completely obey the known physics, then $n_{phy} = 1$. Conversely, when PhyLoss$(x, \hat{y})$ is very large, i.e., the predictions have large deviations from the known physics, $n_{phy} \rightarrow 0$. This allows us to interpret n_{phy} as a likelihood score of the generated sample being physically consistent, a larger likelihood score indicates greater physical consistency.

Moreover, we adapt a similar formulation to [10], where we use an additional network Q to learn the mapping $Q : \{x, \hat{y}\} \rightarrow z$. This network delivers stability in the training procedure by mitigating mode collapse while providing a variational approximation to true posterior over the latent variables z.

The overall learning objective of the generator and the discriminator for the cGAN-PID is shown below:

$$\min_{G} \max_{D} \quad V(D,G) = \mathbb{E}_{y \sim p_{data}(y)}[\log D(x,y,1)] + \mathbb{E}_{z \sim p_z(z)}[\log(1 - D(x, G(x,z), n_{phy}))] + \lambda_Q \, \mathbb{E}_{z \sim p_z(z)}[(z - Q(x, G(x,z)))^2] \tag{2}$$

It must be noted that we use $n_{phy} = 1$ when we show "real" samples to the discriminator as we assume all "real" samples are physically consistent.

3 Dataset Description

We demonstrate the efficacy of our proposed approach on the following two datasets.

3.1 Synthetic Toy Dataset

First, we create a synthetic toy dataset to model a 2D-wave (See Figure 2(a)) given by the equation 3. We additionally impose a constraint on the output z, such that it should lie on the surface of a sphere of radius 2 units given by the Equation 4. We discretized both inputs x and y between the interval [-2,2] into 10,000 points and then sampled points from the intersection of the two curves with a tolerance value of 0.00001. As a result, we obtain a dataset with 300 points where both the Equation 3 and 4 are satisfied. We divide this dataset into two different training and test splits, to test the performance of the model on in-distribution and out-of-distribution test samples (See Figure 2(b) and 2(c)).

$$z = \sin(x) + \sin(y) \tag{3}$$

$$x^2 + y^2 + z^2 = 4 \tag{4}$$

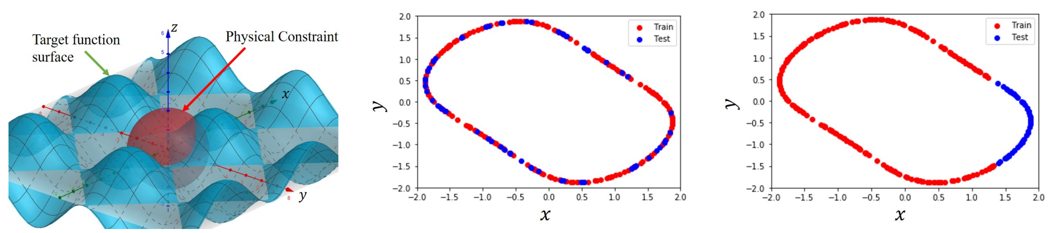

(a) Schematic Representation of Synthetic Dataset (b) In-Distribution Training-Test Split (c) Out-of-Distribution Training-Test Split

Figure 2: Synthetic Toy Dataset

3.2 Collision Dataset

We used a simulation dataset from the paper [1] to evaluate our proposed PID architecture. The dataset uses the speed of two objects at the two initial time stamps $\{v_{a_1}, v_{b_1}, v_{a_2}, v_{b_2}\}$, mass $\{m_a, m_b\}$, distance between the objects D as input x to estimate their speed after collision $\{v_{a_f}, v_{b_f}\}$. The physical constraints of momentum and energy conservation for this problem are shown in Equation 5 and Equation 6, respectively. In this problem, we computed the physical loss as the sum of inconsistencies in these energy and momentum equations.

$$m_a v_{a_1} + m_b v_{b_1} = m_a v_{a_f} + m_b v_{b_f} \tag{5}$$

$$\frac{1}{2} m_a v_{a_1}^2 + \frac{1}{2} m_b v_{b_1}^2 = \frac{1}{2} m_a v_{a_f}^2 + \frac{1}{2} m_b v_{b_f}^2 \tag{6}$$

4 Results

We use the standard conditional GAN (cGAN), cGAN trained with additional physical loss (cGAN-PhyLoss), and multi-layer perceptron (MLP) trained with additional physics loss (MLP-PhyLoss) as baselines to compare the performance of our proposed model (cGAN-PID). For the MLP-PhyLoss model we utilize dropout [5] to approximate the posterior distribution of the target variable (with dropout rate = 0.2). We use a sampling size of 10000 to obtain the distribution for both cGAN-based models and MLP-PhyLoss. For evaluation metrics, we use Mean Root Mean Squared Error (RMSE) and Mean physical inconsistency, i.e., we use the mean of our predicted distribution to calculate the RMSE and the physical inconsistency.

4.1 Synthetic Dataset

4.1.1 In-distribution and Out-of-distribution Results

On the synthetic dataset, we evaluate the performance of the various baselines for both in-distribution and out-of-distribution test setups. We observe that for both evaluation setups, using any form of physics supervision in the deep learning models helps to achieve better generalizability which is reflected through lower Mean RMSE as well as lower mean physical inconsistencies in the predictions. From Table 1, we see that cGAN-PID outperforms the existing baselines in terms of both Mean RMSE and mean physical inconsistency across both evaluation setups. It must also be noted that the variance of the mean RMSE for cGAN-PID is the least indicating less variability to random initialization, which is again desirable.

	In-Distribution Results		**Out-of-Distribution Results**	
	Test Mean RMSE	**Test Physical Inconsistency**	**Test Mean RMSE**	**Test Physical Inconsistency**
cGAN	0.134 ± 0.030	0.181 ± 0.046	0.579 ± 0.012	0.520 ± 0.027
cGAN-PID	$\mathbf{0.021 \pm 0.003}$	$\mathbf{0.026 \pm 0.006}$	$\mathbf{0.489 \pm 0.003}$	$\mathbf{0.428 \pm 0.040}$
cGAN-PhyLoss	0.048 ± 0.007	0.007 ± 0.016	0.549 ± 0.010	0.481 ± 0.015
MLP-PhyLoss	0.135 ± 0.010	0.173 ± 0.015	0.676 ± 0.013	0.596 ± 0.032

Table 1: In-distribution and out-of-distribution results on the synthetic toy dataset for 10 random runs. We choose $\lambda_{phy} = 0.2$ for cGAN-PID, $\lambda_{phy} = 0.5$ for cGAN-PhyLoss, and $\lambda_{phy} = 10$ for MLP-PhyLoss.

4.1.2 Sensitivity of trade-off hyper parameter

Note that the choice of the trade-off hyperparameter (weighting PhyLoss with other loss terms) plays a crucial role in the PGL paradigm. Here, we demonstrate that by varying the trade-off hyperparameter for both the cGAN-PID and cGAN-PhyLoss models. Figure 3 shows that for cGAN-PhyLoss, after a particular value of λ_{phy}, both the training and test RMSE skyrocketed. This can be explained using the analogy between the PGL paradigm and multi-task learning. Since, we can treat the PhyLoss as an auxiliary optimization task, so once we give too much weight on the auxiliary task by increasing the value of λ_{phy}, the model tries to optimize only for the PhyLoss rather than optimizing the primary objective which is the adversarial loss in this case. This results in the cGAN-PhyLoss model not converging, thus producing high test/train RMSE.

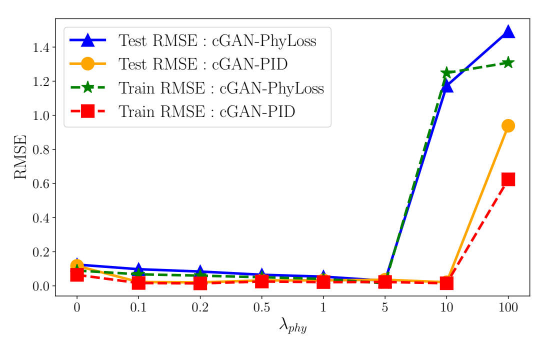

Figure 3: Plot to show the sensitivity of λ_{phy} for cGAN-PID vs cGAN-PhyLoss.

4.2 Collision Dataset

For the collision dataset, we can observe from Table 4.2 that cGAN-PID, cGAN-Phyloss, and MLP-PhyLoss perform better than cGAN, which re-establishes the idea that integrating physical knowledge into the learning algorithm improves the generalizibility of models on unseen test instances. The cGAN-PID again outperformes all other baselines by a significant margin in terms of Mean RMSE. Also, the cGAN-PID obtains the least physical inconsistency, which further demonstrates that by introducing an additional term in the discriminator we are indeed able to make the predictions more physically consistent and perhaps it might be even better than the conventional PGL paradigm of using loss functions in this special case of adversarial learning.

	Test Mean RMSE	**Test Mean Physical Inconsistency**
cGAN	2.318 ± 0.433	36.910 ± 6.520
cGAN-PID	$\mathbf{0.712 \pm 0.073}$	$\mathbf{14.008 \pm 3.300}$
cGAN-PhyLoss	1.048 ± 0.063	21.367 ± 2.016
MLP-PhyLoss	1.003 ± 0.027	26.178 ± 1.620

Table 2: Results on Collision Dataset for 10 random runs. We choose $\lambda_{phy} = 0.005$ for cGAN-PID, $\lambda_{phy} = 0.1$ for cGAN-PhyLoss, and $\lambda_{phy} = 0.001$ for MLP-PhyLoss.

5 Conclusion and Future Work

In this paper, we provide preliminary results in favor of a novel approach for incorporating physical knowledge as an additional input in cGANs, termed as cGAN-PID, which could be used as an alternative to the existing PGL paradigm. We also demonstrate that cGAN-PID shows lower sensitivity to the choice of the trade-off hyperparameter. There are numerous directions which could be explored from here on. The choice of transformation $f(.)$ used to obtain n_{phy} could be further explored. Also, efforts could be made to learn the trade-off hyper-parameter λ_{phy} or even adaptively change it during the training process. Furthermore, scientific laws are often based on oversimplified assumptions which are sometimes incomplete/incorrect. We will experiment the effectiveness of our cGAN-PID framework, in scenarios with imperfect scientific knowledge in future work. Further, we will also explore other ways of incorporating imperfect scientific knowledge in softer ways as additional inputs to the cGAN generator can be explored in future work. Future work can also demonstrate the effect of cGAN-PID on uncertainty quantification and extensions of PID to other GAN based formulations.

Broader Impact

Our work presents a novel way of incorporating physics in GAN formulations, which have several potential impacts on science and society. GANs are increasing being used in a number of commercial applications such as style transfer and photo-realistic image generation, as well as scientific applications such as particle physics and hydrology. In many of these applications, incorporating physics in GAN formulations can lead to better predictions than black-box GAN formulations.

References

[1] Yunhao Ba, Guangyuan Zhao, and Achuta Kadambi. Blending diverse physical priors with neural networks. *arXiv preprint arXiv:1910.00201*, 2019.

[2] Thang Bui, Daniel Hernández-Lobato, Jose Hernandez-Lobato, Yingzhen Li, and Richard Turner. Deep gaussian processes for regression using approximate expectation propagation. In *International conference on machine learning*, pages 1472–1481, 2016.

[3] Arka Daw, R Quinn Thomas, Cayelan C Carey, Jordan S Read, Alison P Appling, and Anuj Karpatne. Physics-guided architecture (pga) of neural networks for quantifying uncertainty in lake temperature modeling. In *Proceedings of the 2020 SIAM International Conference on Data Mining*, pages 532–540. SIAM, 2020.

[4] Meire Fortunato, Charles Blundell, and Oriol Vinyals. Bayesian recurrent neural networks. *arXiv preprint arXiv:1704.02798*, 2017.

[5] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.

[6] Anuj Karpatne, Gowtham Atluri, James H Faghmous, Michael Steinbach, Arindam Banerjee, Auroop Ganguly, Shashi Shekhar, Nagiza Samatova, and Vipin Kumar. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering*, 29(10):2318–2331, 2017.

[7] Anuj Karpatne, William Watkins, Jordan Read, and Vipin Kumar. Physics-guided neural networks (pgnn): An application in lake temperature modeling. *arXiv preprint arXiv:1710.11431*, 2017.

[8] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[9] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

[10] Chunyuan Li, Jianqiao Li, Guoyin Wang, and Lawrence Carin. Learning to sample with adversarially learned likelihood-ratio. 2018.

[11] Tomáš Mikolov, Stefan Kombrink, Lukáš Burget, Jan Černockỳ, and Sanjeev Khudanpur. Extensions of recurrent neural network language model. In *2011 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5528–5531. IEEE, 2011.

[12] Yibo Yang and Paris Perdikaris. Adversarial uncertainty quantification in physics-informed neural networks. *Journal of Computational Physics*, 394:136–152, 2019.