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

Arka Daw[1], M. Maruf[1], and Anuj Karpatne[1]

[1]Department of Computer Science, Virginia Polytechnic and State University, Blacksburg, VA

Motivation

- **Perform Uncertainty Quantification (UQ)**

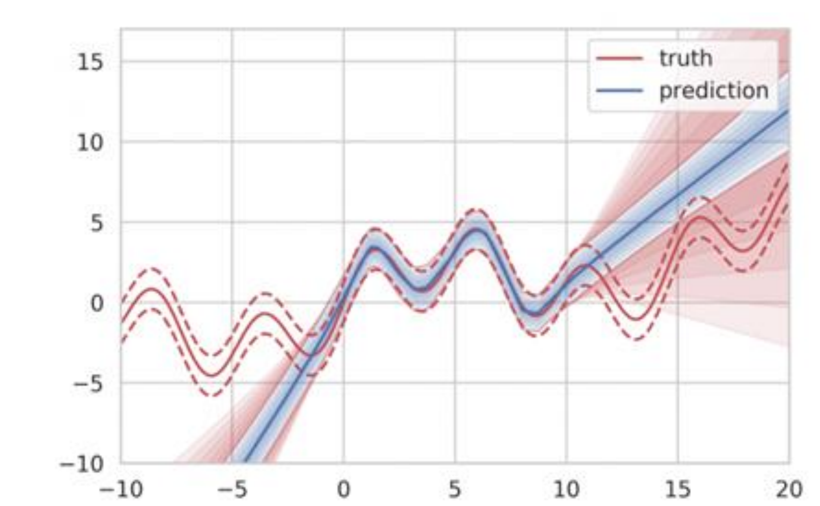

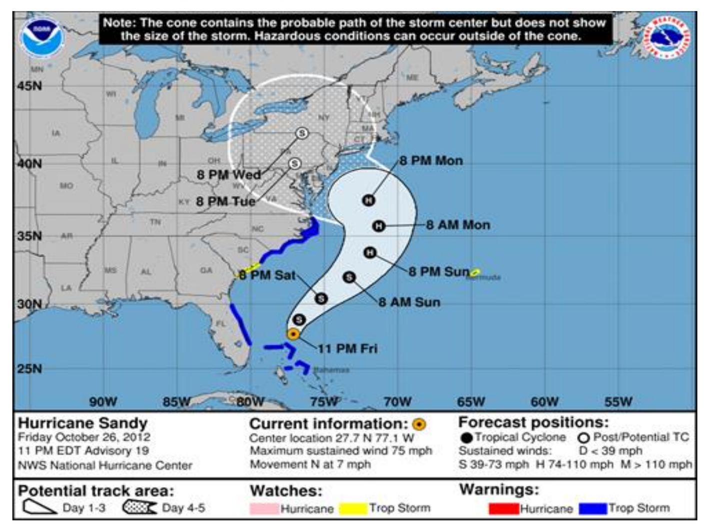

- **Ensure predictions obey known scientific (physical) laws.**

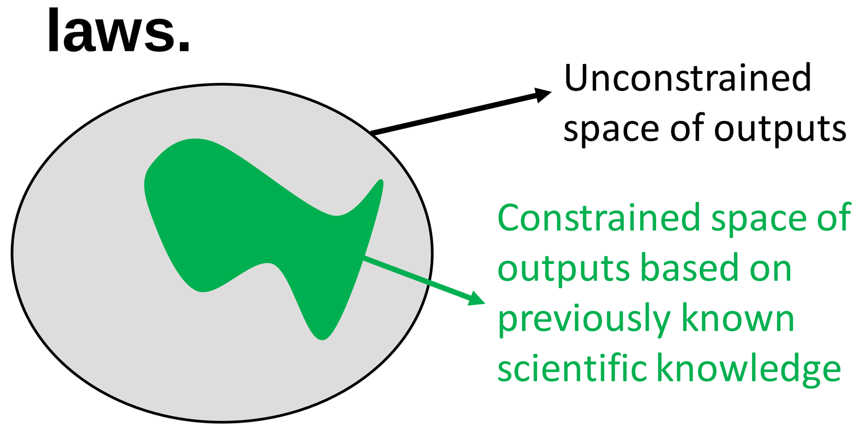

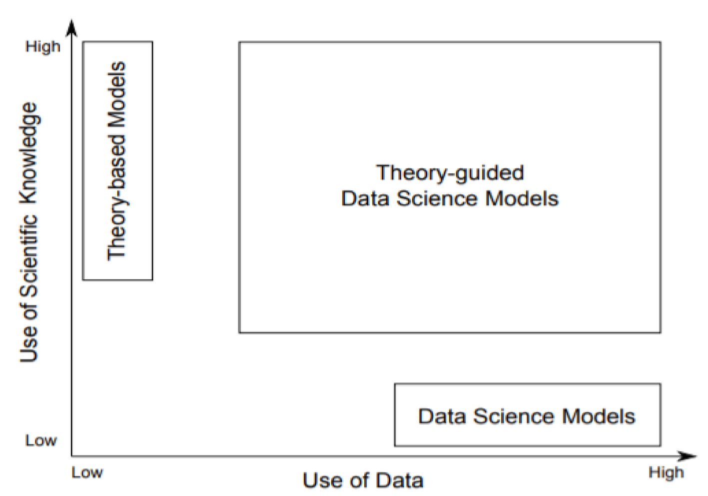

Problem Statement

Goal : To train a neural network $f_\theta(.)$ to learn the mapping $f_\theta: x, z \to y$ to approximate the posterior distribution of the outputs $P(f_\theta(x)|x, z)$, where z is a latent variable sampled from a prior distribution $z \sim p(z)$.

Baseline Methods :

1. Conditional Generative Adversarial Net (cGAN) :

$$\min_G \max_D \mathbb{E}_{z\sim p_z(z)}\left[\log\left(1 - D\left(x, G(x,z)\right)\right)\right] + \mathbb{E}_{y\sim p_{data}(y)}[\log D(x,y)]$$

2. Training cGAN with Physics Loss function (cGAN-PhyLoss) [1]:

$$\min_G \max_D \mathbb{E}_{z\sim p_z(z)}\left[\log\left(1 - D\left(x, G(x,z)\right)\right)\right] + \mathbb{E}_{y\sim p_{data}(y)}[\log D(x,y)] + \lambda_{phy}\mathrm{PhyLoss}(x,\hat{y})$$

Physics-Informed Discriminator (PID)

Key Idea : We add an additional input n_{phy} to the discriminator D.

$$n_{phy} = f\left(\lambda_{phy}\mathrm{PhyLoss}(x,\hat{y})\right)$$

Monotonic function.

The motivation of adding n_{phy} is to aid the discriminator such that it not only distinguishes between a real and a fake sample by learning the underlying data distribution but also uses the additional physical knowledge (n_{phy}) to make the distinction.

We use $n_{phy} = e^{-\lambda_{phy}\mathrm{PhyLoss}(x,\hat{y})}$

Proposed architecture of cGAN-PID:

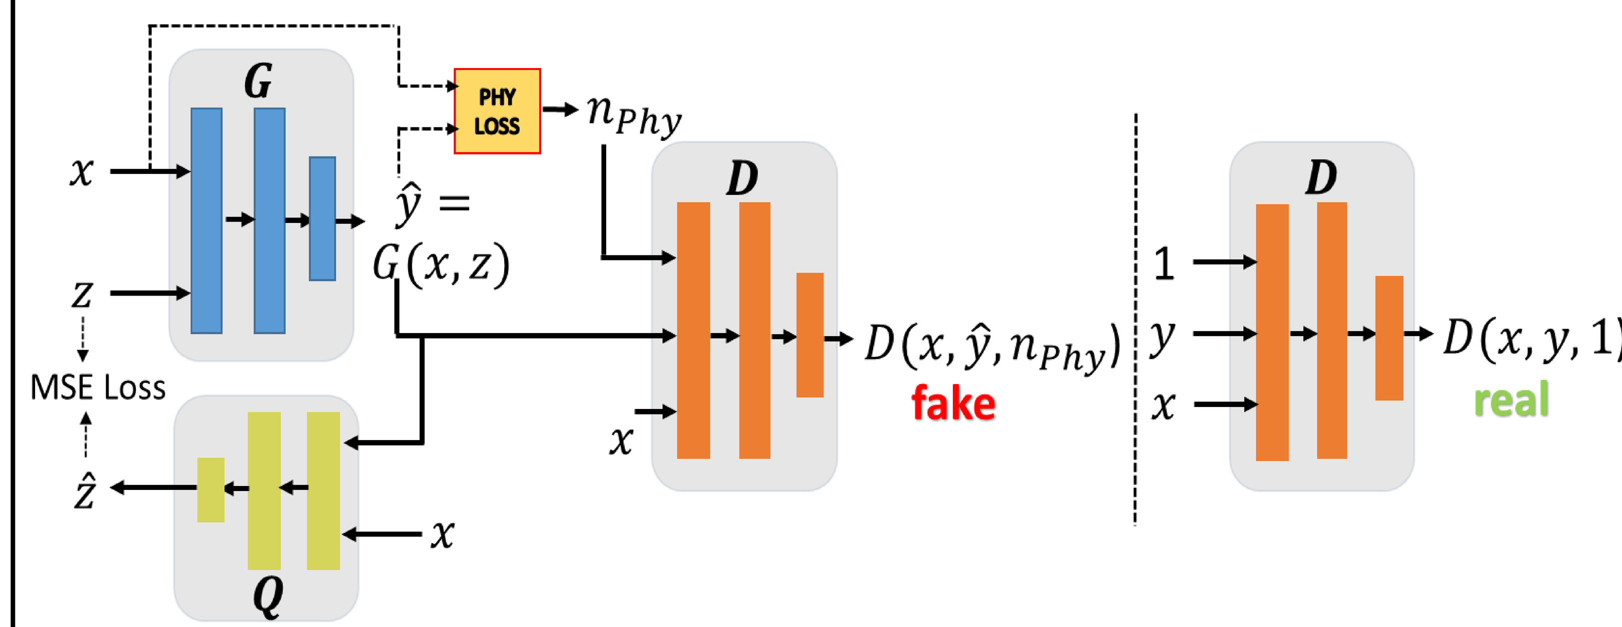

cGANs suffer from mode collapse, so we trained a network $Q: x, \hat{y} \to z$ The network Q delivers stability in the training process and provides a variational approximation over the latent variable z[2].

The overall objection function of cGAN-PID:

$$\min_G \max_D \mathbb{E}_{z\sim p_z(z)}\left[\log\left(1 - D\left(x, G(x,z), n_{phy}\right)\right)\right] + \mathbb{E}_{y\sim p_{data}(y)}[\log D(x,y,1)] + \lambda_Q \mathbb{E}_{z\sim p_z(z)}[(z - Q(x, G(x,z)))^2]$$

Datasets

- **Synthetic Dataset**

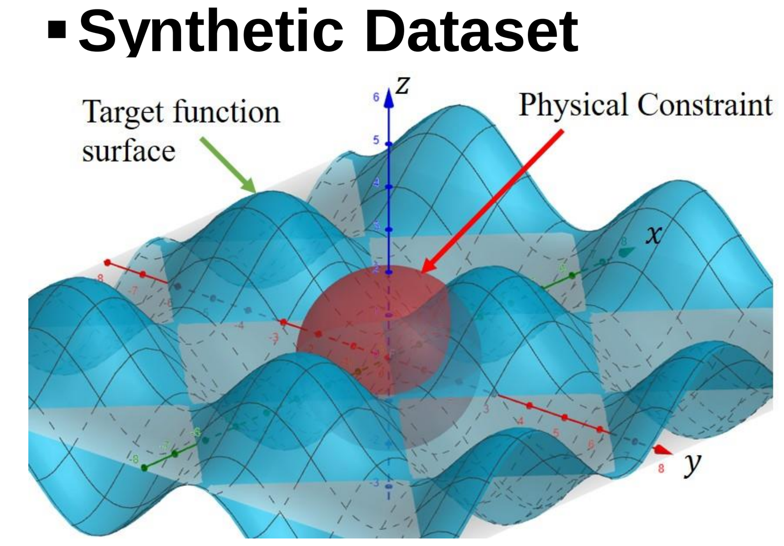

Target Function

$$z = \sin x + \sin y$$

$$x^2 + y^2 + z^2 = 4$$

Physical Constraint

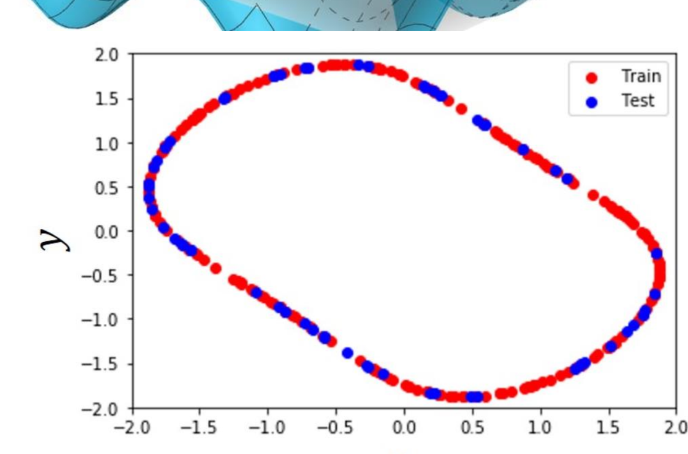

In-distribution Training/Test Split

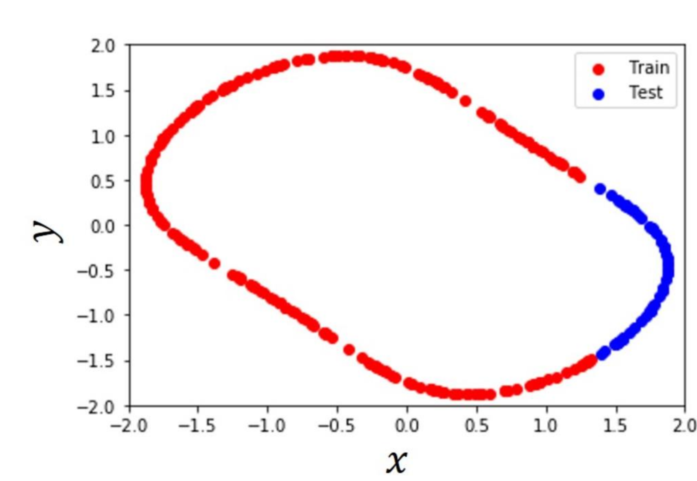

Out-of-distribution Training/Test Split

- **Collision Dataset**[3]

Momentum conservation:

$$m_a v_{a1} + m_b v_{b1} = m_a v_{af} + m_b v_{bf}$$

Energy conservation:

$$\frac{1}{2}m_a v_{a1}^2 + \frac{1}{2}m_b v_{b1}^2 = \frac{1}{2}m_a v_{af}^2 + \frac{1}{2}m_b v_{bf}^2$$

Results

- **Synthetic Dataset**

	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

- cGAN-PID outperforms all the baselines in terms of both test mean RMSE and test physical inconsistency.
- cGAN-PID also shows reduced sensitivity to random initializations.

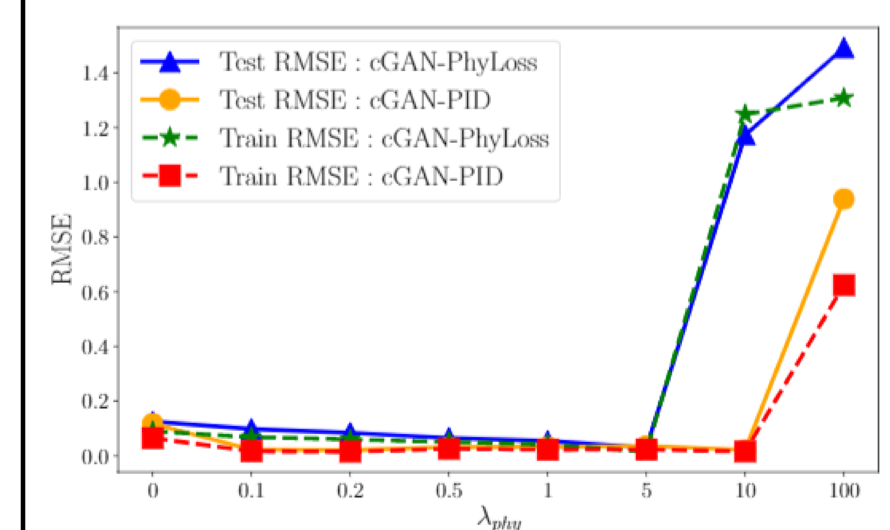

Sensitivity of hyper-parameter

- cGAN-PID is less sensitive to the choice of hyper-parameters.
- cGAN-PID uses physics loss in optimizing both the generator G and discriminator D.

- **Collision Dataset**

	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

- cGAN-PID again outperforms all the baselines in terms of both test mean RMSE and test physical inconsistency.

*All results are averaged over 10 random runs

Conclusions and Future Works

- We propose an alternative (perhaps more natural) way of incorporating physics into the adversarial learning framework
- cGAN-PID demonstrates lower sensitivity to the choice of hyper-parameters.
- Choice of the transformation function $f(.)$ can be explored.
- λ_{phy} can be adaptively changed during the training process.
- cGAN-PID can be extended to work with incomplete physics.

Acknowledgements

The authors would like to acknowledge the support from Discovery Analytics Center of the Computer Science department of Virginia Tech.

References

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

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

[3] Yunhao Ba, Guangyuan Zhao, and Achuta Kadambi. Blending diverse physical priors withneural networks.arXiv preprint arXiv:1910.00201, 2019