Solving Kuramoto Oscillator Model using Physics Informed Neural Network

Anonymous Author(s)
Affiliation
Address
email

Abstract

Physics informed machine learning has been emerged as a powerful tool with the help of deep learning as the latter has been instrumental as a data-driven function approximator. Many recent works have been focusing on solving hard to solve differential equations with the help of physics informed neural network (PINN), a tremendously simple approach which blends physics and deep learning. We explore the application of PINN in solving Kuramoto system of coupled differential equations as well as in decision making problem of synchronization state of the system. The experimental results illustrate that PINN can not only be used to solve the coupled differential equations, but also be very handy when our objective is to figure out the synchronization capability of the oscillator system in consideration.

1 Introduction

The Kuramoto model [1] of oscillatory systems have been widely used for its capability to model the synchronized behaviour through a system of coupled differential equations. More recently, it is getting attention in the field of optimal experiment design [2, 3]. For Kuramoto system, the goal of the optimal experiment design is to reduce the effect of uncertainty on selecting a control oscillator to make the system synchronized. The uncertainty originates from the unknown interaction strengths between oscillators. If the coupling strengths are known accurately, then we would have chosen the optimum control oscillator. But, due to the lack of information, we suffer a cost by choosing a suboptimal solution. To quantify the effects of this uncertainty, we need to calculate the expected cost by solving a large number of Kuramoto systems which are generated by sampling from the uncertainty class for interaction strengths. Solving these large number of systems in numerical method requires a lot of computational time. To address this challenge, parallel computation of differential equation solver [2], surrogate model for estimating the average cost [3] are already proposed. In this project, we have tried to study the feasibility of the physics informed neural network (PINN) to tackle the computational complexity in optimal experiment design for controlling the Kuramoto system.

The PINN [4] has laid out the data-driven approach powered by the automatic differentiation to solve the system of differential equations. Although the training of PINN is not understood well enough, it is showing promising results [5], specially for the failure cases of numerical solvers. For the Kuramoto model, the numerical method works very well but takes long time. The PINN based approach may have potential to be used as accurate as well as fast solver. With this goal in mind, we have focused on whether the PINN solver can accurately approximate the synchronized and unsynchronized Kuramoto system.

In section 2, we provide an introductory description of the Kuramoto model and how we have applied the physics informed neural network to solve this coupled system. For each case of experiments under section 3, we have highlighted our findings that point out the successes as well as the existing challenges for PINN.

Submitted to 38th Conference on Neural Information Processing Systems (NeurIPS 2024). Do not distribute.

2 Methods

2.1 Kuramoto Model

A Kuramoto model of N oscillators can be represented by the system of coupled differential equations, 1 with the initial conditions in equation 2 for the angular position θ_i. The set of interaction strengths or coupling coefficients between each pair of oscillators, $a_{i,j}$ governs the dynamics of this system.

$$\dot{\theta}_i(t) = \omega_i + \sum_{j=1}^{N} a_{i,j} \sin(\theta_j(t) - \theta_i(t)), \quad i = 1, 2, \cdots, N \tag{1}$$

$$\theta_i(0) = \theta_i^0 \tag{2}$$

Along with the intrinsic angular frequencies, $\{\omega_i\}$, the coupling strengths can make a system synchronized or unsynchronized. A system is frequency synchronized when the instantaneous angular frequencies of all the oscillators converge to a single value. (equation 3)

$$\lim_{t \to \infty} |\dot{\theta}_i(t) - \dot{\theta}_j(t)| = 0, \quad \forall i, j \tag{3}$$

2.2 Synchronization condition for $N = 2$

For Kuramoto model with $N = 2$ oscillators, it is proved in [3] that equation 3 is satisfied if and only if $|w_1 - w_2| \leq 2a$ where $a = a_{1,2} = a_{2,1}$. We use this condition to build a synchronized and unsynchronized system of 2 oscillators and apply the PINN solver, described in the following section, to get the angular positions of each oscillator.

2.3 Solving with PINN

We model the solution provided by the PINN as $\boldsymbol{\theta}(t, \mathbf{W}) \in \mathbb{R}^N$ where $\mathbf{W}$ represents the parameters of the neural network. Also, we have $\boldsymbol{\theta}^0 \in \mathbb{R}^N$, the initial angular positions of the oscillators. We train the PINN by combining the loss at the initial time-point, $t = 0$ and the residual loss at N_r time-points, $t_1, t_2, \cdots, t_{N_r}$.

$$\mathcal{L}(\mathbf{W}) = \lambda_b \mathcal{L}_b(\mathbf{W}) + \lambda_r \mathcal{L}_r(\mathbf{W}) \tag{4}$$

$$\mathcal{L}_b(\mathbf{W}) = \left|\left|\boldsymbol{\theta}(0, \mathbf{W}) - \boldsymbol{\theta}^0\right|\right|^2 \tag{5}$$

$$\mathcal{L}_r(\mathbf{W}) = \frac{1}{N_r} \sum_{k=1}^{N_r} \left|\left|r\left(\boldsymbol{\theta}(t_k, \mathbf{W})\right)\right|\right|^2 \tag{6}$$

$$r\left(\boldsymbol{\theta}(t_k, \mathbf{W})\right) = \begin{bmatrix} \frac{\partial \theta_1(t_k, \mathbf{W})}{\partial t} - \omega_1 - \sum_{j=1}^{N} a_{1,j} \sin(\theta_j(t_k, \mathbf{W}) - \theta_1(t_k, \mathbf{W})) \\ \frac{\partial \theta_2(t_k, \mathbf{W})}{\partial t} - \omega_2 - \sum_{j=1}^{N} a_{2,j} \sin(\theta_j(t_k, \mathbf{W}) - \theta_2(t_k, \mathbf{W})) \\ \vdots \\ \frac{\partial \theta_N(t_k, \mathbf{W})}{\partial t} - \omega_N - \sum_{j=1}^{N} a_{N,j} \sin(\theta_j(t_k, \mathbf{W}) - \theta_N(t_k, \mathbf{W})) \end{bmatrix} \tag{7}$$

3 Results

3.1 Experiment for two oscillator system

We set the intrinsic frequencies as $\omega_1 = 0.5, \omega_2 = 0.9$. For synchronized system, we choose 1 as the coupling coefficient, and 0.1 in case of unsynchronized system. We use a PINN with 8 hidden layers, 20 neurons per layer to approximate the solutions of the system. $N_r = 2000$ points are randomly selected from $0 \leq t \leq 50$ seconds. We empirically set $\lambda_b = 1$ and $\lambda_r = 2$ and train the PINN for 4000 iterations. Figure 1a and 1b show the approximated angular positions (top row) by the fourth order Runge–Kutta (RK4) method and the trained PINN and also the corresponding instantaneous angular frequencies (bottom row). It appears that for synchronized system, the PINN's solution is very close to that from RK4 method. In unsynchronized system, the solution from PINN is way off from the RK4's solutions. To resolve this issue, we increase λ_r to 20, and the PINN's accuracy gets much better (Figure 1c).

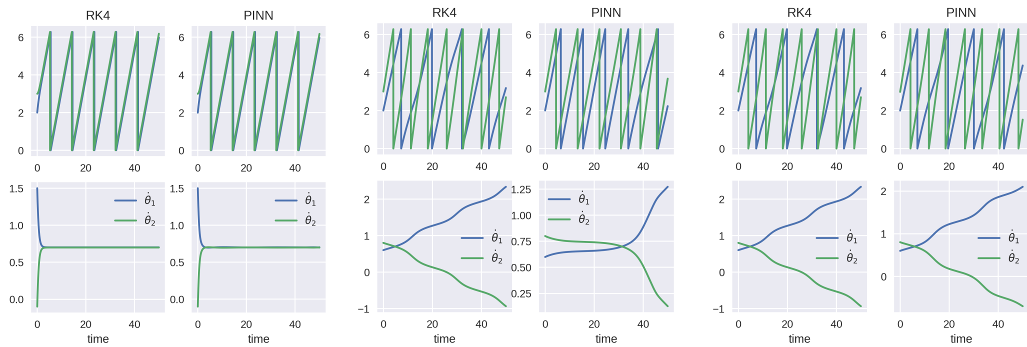

(a) Synchronized (b) Unsynchronized ($\lambda_r = 2$) (c) Unsynchronized ($\lambda_r = 20$)

Figure 1: Experiments with two oscillator Kuramoto network.

3.2 Effect of Coupling Coefficient on PINN's accuracy

We did the experiments with a Kuramoto model consisting of 10 oscillators. For simplicity, we assumed the whole Kuramoto model has a single coupling coefficient, $a_{i,j} = K/(N-1)$ instead of different coupling coefficients for each connection of the oscillator model. We used $\lambda_r = 10, \lambda_b = 1$ and $1,000$ collocation points. We trained the model for $20,000$ iterations for each case.
What we saw is the PINN can learn the solution with near perfection when the coupling coefficient of the Kuramoto model is low. As the coupling coefficient increases, the PINN fails to learn the true solution.
Figure 2a and 2b show the true solution and the predicted solution by PINN when $K = 0.1$. Figure 2c and 2d show the true solution and the predicted solution by PINN when $K = 2.0$.

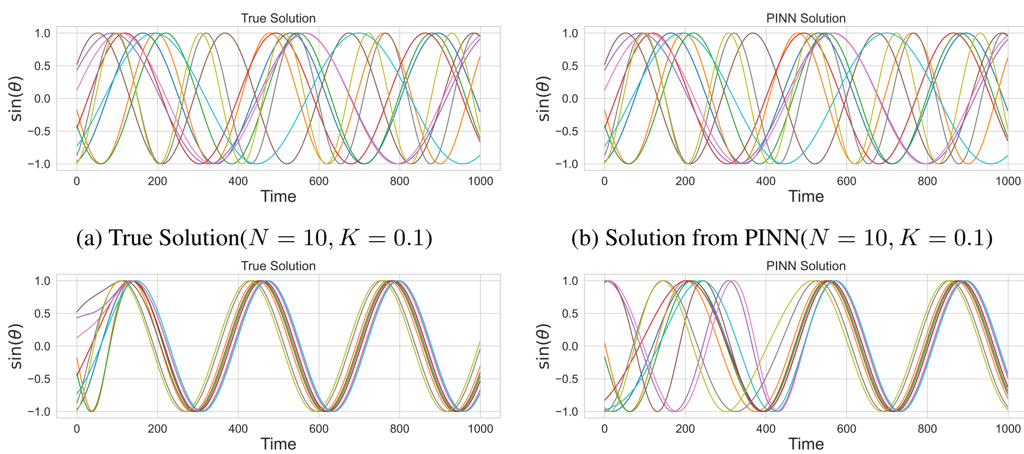

(a) True Solution($N = 10, K = 0.1$) (b) Solution from PINN($N = 10, K = 0.1$)

(c) True Solution($N = 10, K = 2.0$) (d) Solution from PINN($N = 10, K = 2.0$)

Figure 2: True and PINN solution for weakly and strongly coupled Kuramoto network.

3.3 Training PINN without Initial Condition

We trained PINN without giving the initial points with an objective of examining whether the PINN can just learn the synchronization capability of the Kuramoto model we are experimenting with. We experimented with two different coupling coefficients $K = 0.5, 2.0$. Here we only used $1,000$ collocation points to train the model.
Figure 3a and 3b shows the comparison of instantaneous frequencies between two different coupling coefficient values. For higher K, the Kuramoto model is supposed to be locked up to a single instantaneous frequency for all oscillators. For $K = 2$, the PINN can learn that instantaneous frequency quite successfully. As the initial phases are assumed by the PINN, they differ from the true solution. As time progresses the PINN predicts the locking frequency quite successfully. For $K = 0.5$ the true solution does not synchronise to a single frequency, neither does the PINN.

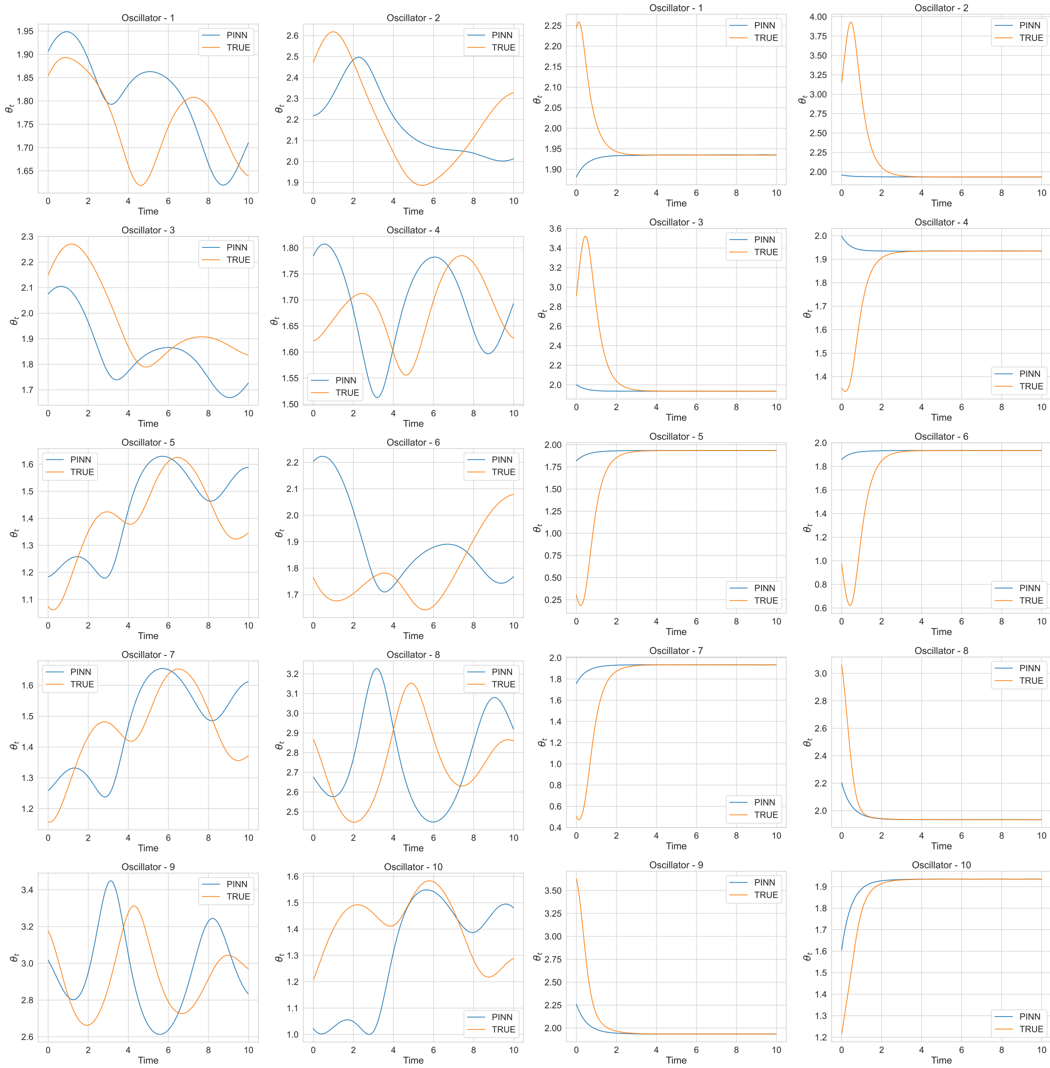

(a) $N = 10, K = 0.5$ (b) $N = 10, K = 2.0$

Figure 3: PINN solver predicted instantaneous frequencies. PINN was trained without providing the initial conditions.

4 Conclusion

In this work, we examined the potential of PINNs to solve coupled differential equations. Our main objective was twofold: firstly, examine whether the PINN can accurately solve the coupled differential equation when given all the necessary initial conditions, and secondly, whether the PINN can figure out whether a given Kuramoto system will synchronize or not without giving any initial conditions. We saw that PINN can learn the solution accurately with low coupling strength. With higher coupling strength, PINN seems to fail in following the abrupt change of some of the oscillators. This is well known phenomenon of PINNs to fail learn the high frequencies of the true solution. The self adaptive training approach [6] might be a solution to resolve this issue. For the second objective, PINN seems to be successful in achieving its goal. It can predict the locking frequency of a synchronized Kuramoto model. For models with higher number of oscillators, the figure shows that the true solution does not reach to a locking frequency whereas PINN does. This indicates the power of PINN that it can reach to the locking frequency even quicker than the actual solution with random initial phases. For future work, we intend to investigate the application of PINN for Kuramoto models with arbitrarily different interaction strengths for each pair of oscillators.

References

[1] Yoshiki Kuramoto. Self-entrainment of a population of coupled non-linear oscillators. *Lect. Notes Phys.*, 39:420–422, 1975.

[2] Hyun-Myung Woo, Youngjoon Hong, Bongsuk Kwon, and Byung-Jun Yoon. Accelerating optimal experimental design for robust synchronization of uncertain kuramoto oscillator model using machine learning. *IEEE Transactions on Signal Processing*, 69:6473–6487, 2021.

[3] Youngjoon Hong, Bongsuk Kwon, and Byung-Jun Yoon. Optimal experimental design for uncertain systems based on coupled differential equations. *IEEE Access*, 9:53804–53810, 2021.

[4] M. Raissi, P. Perdikaris, and G.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.

[5] Sifan Wang, Xinling Yu, and Paris Perdikaris. When and why pinns fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022.

[6] Levi D. McClenny and Ulisses M. Braga-Neto. Self-adaptive physics-informed neural networks using a soft attention mechanism. *CoRR*, abs/2009.04544, 2020.

A Quantification of Approximation Accuracy of PINN Solver

To quantify the the approximation accuracy of the PINN solver, we have used the metric of expected maximum relative L_2 error:

$$\text{Expected maximum relative } L_2 \text{ error} = \mathrm{E}\left[\max_{i\in\{1,2,\cdots,N\}} \frac{\sqrt{\sum_{k=1}^{N_q}\left(\theta_i(t_k,\mathbf{W}) - \theta_i^{true}(t_k)\right)^2}}{\sqrt{\sum_{k=1}^{N_q}(\theta_i^{true}(t_k))^2}}\right] \tag{8}$$

In equation 8, the true solution for each oscillator is obtained by numerical solver. The N_q test points are taken from the domain with uniform separation. For performing the expectation operation, we instantiate 5 PINNs with 8 hidden layers and 50 neurons per layer and take the average maximum L_2 error for Kuramoto systems with $N = 2, 10, 20$ oscillators and $K = 0.1, 1, 2$. Each PINN is trained with 1000 collocation points for 5000 iterations with $\lambda_r = 10$ and $\lambda_i = 1$. Figure 4 demonstrates that for Kuramoto model with larger number of oscillators, the PINN's accuracy decreases for network with stronger coupling.

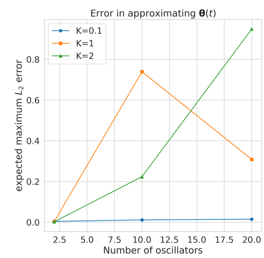

Figure 4: EMR L_2 error.

B Effect of Coupling Coefficient on PINN's accuracy

Figure 5a and 5b show the phases and instantaneous frequencies of each oscillator of a weakly coupled Kuramoto network. These figures show that the PINN quite successfully learns the true solution. Figure 6a and 6b show the phases and instantaneous frequencies of each oscillator of a strongly coupled Kuramoto network. These figures show that the PINN could not learn the true solution.

C Kuramoto Oscillator Network with 15 Oscillators

We created another Kuramoto oscillator network with 15 oscillators and solved it with PINN. Figure 7 shows the true and PINN solution for both a weakly and strongly coupled Kuramoto oscillator network with 15 oscillators. Figure 8 shows the phase and instantaneous frequencies of each oscillator for weakly coupled Kuramoto oscillator network. Figure 9 shows similar results for a strongly couples Kuramoto oscillator network.

D Kuramoto Oscillator Network with 20 Oscillators

We created another Kuramoto oscillator network with 20 oscillators and solved it with PINN. Figure 10 shows the true and PINN solution for both a weakly and strongly coupled Kuramoto oscillator network with 20 oscillators. Figure 11 shows the phase and instantaneous frequencies of each oscillator for weakly coupled Kuramoto oscillator network. Figure 12 shows similar results for a strongly couples Kuramoto oscillator network.

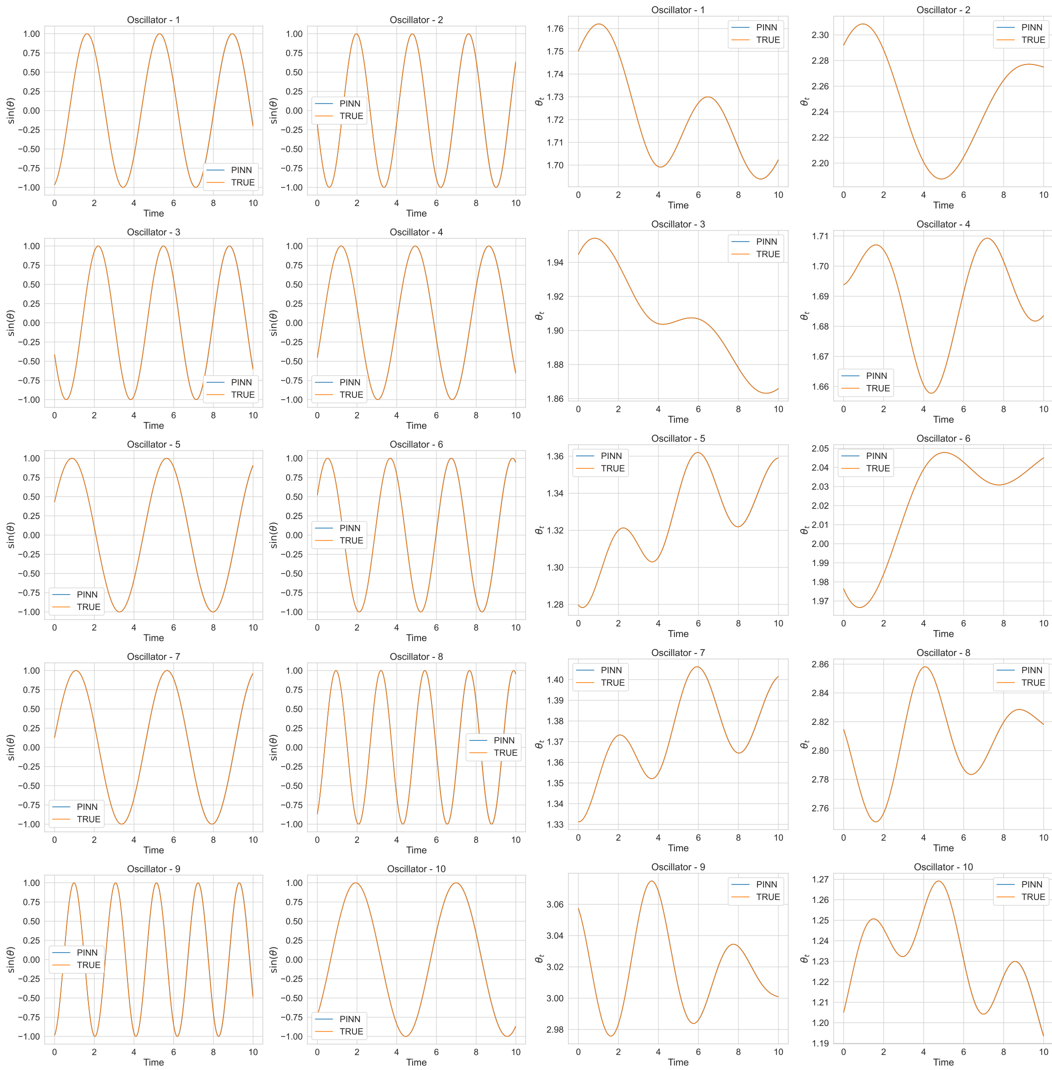

(a) Phase

(b) Instantaneous frequency

Figure 5: Phases and instantaneous frequency of each oscillator($N = 10, K = 0.1$)

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The paper proposes PINN as a solver for coupled oscillator network. PINN has been widely used for solving PDEs, but not so widely studied for coupled PDEs, like Kuramoto oscillator network. We showed that PINN can be used for coupled PDEs too.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.

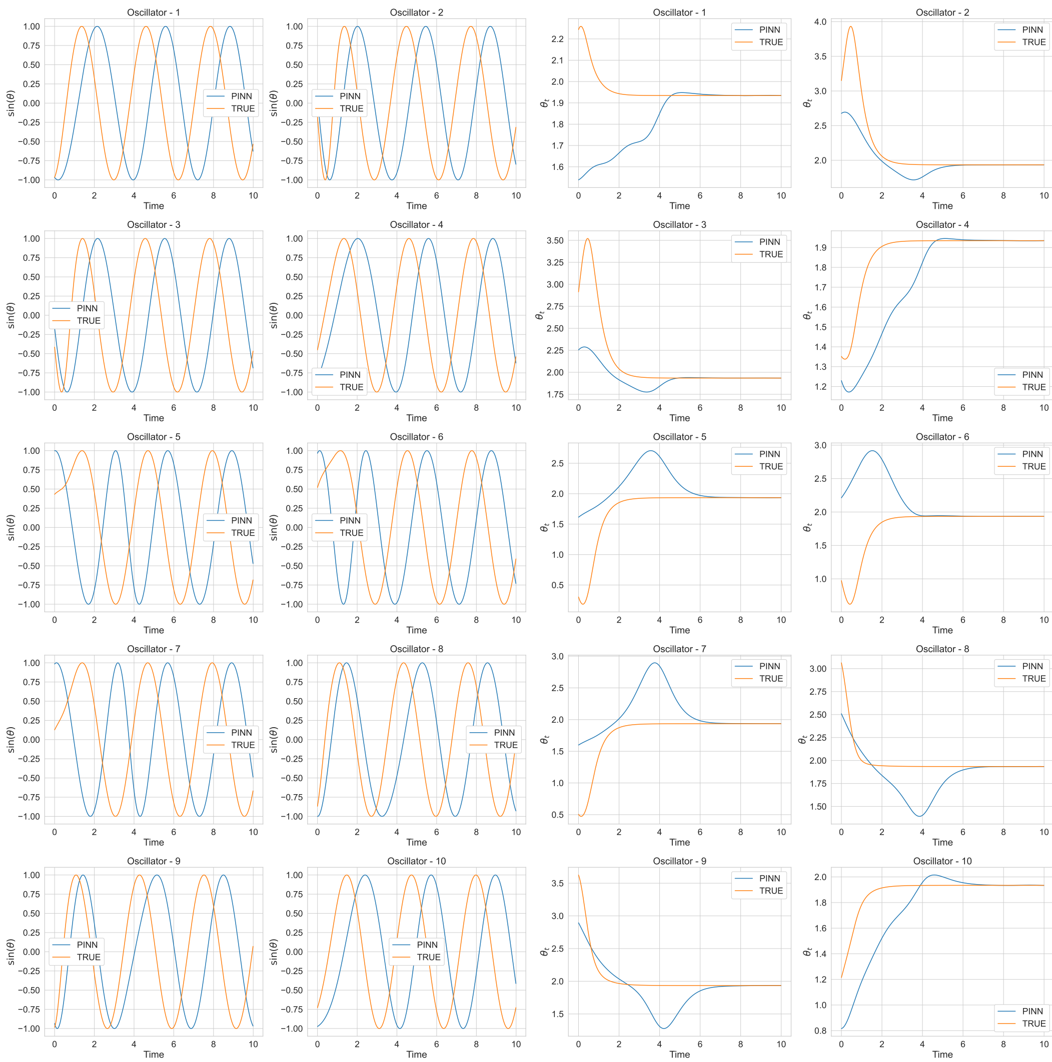

(a) Phase

(b) Instantaneous frequency

Figure 6: Phases and instantaneous frequency of each oscillator($N = 10, K = 2.0$)

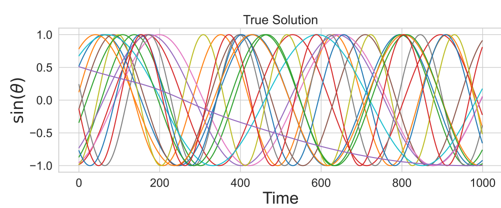

(a) True Solution($N = 15, K = 0.1$)

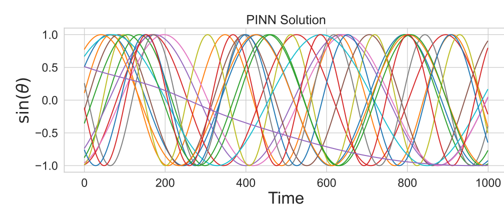

(b) Solution from PINN($N = 15, K = 0.1$)

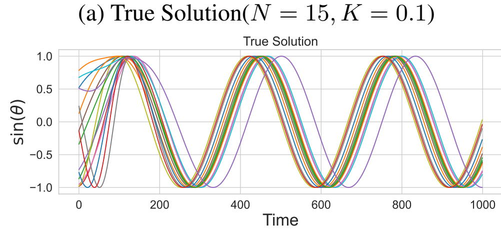

(c) True Solution($N = 15, K = 2.0$)

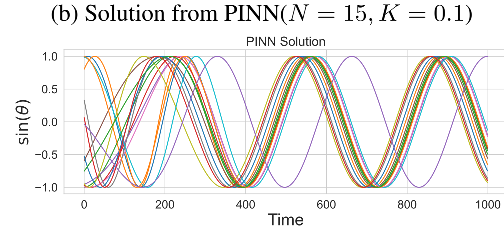

(d) Solution from PINN($N = 15, K = 2.0$)

Figure 7: True and PINN solution for weakly and strongly coupled Kuramoto network with 15 oscillators.

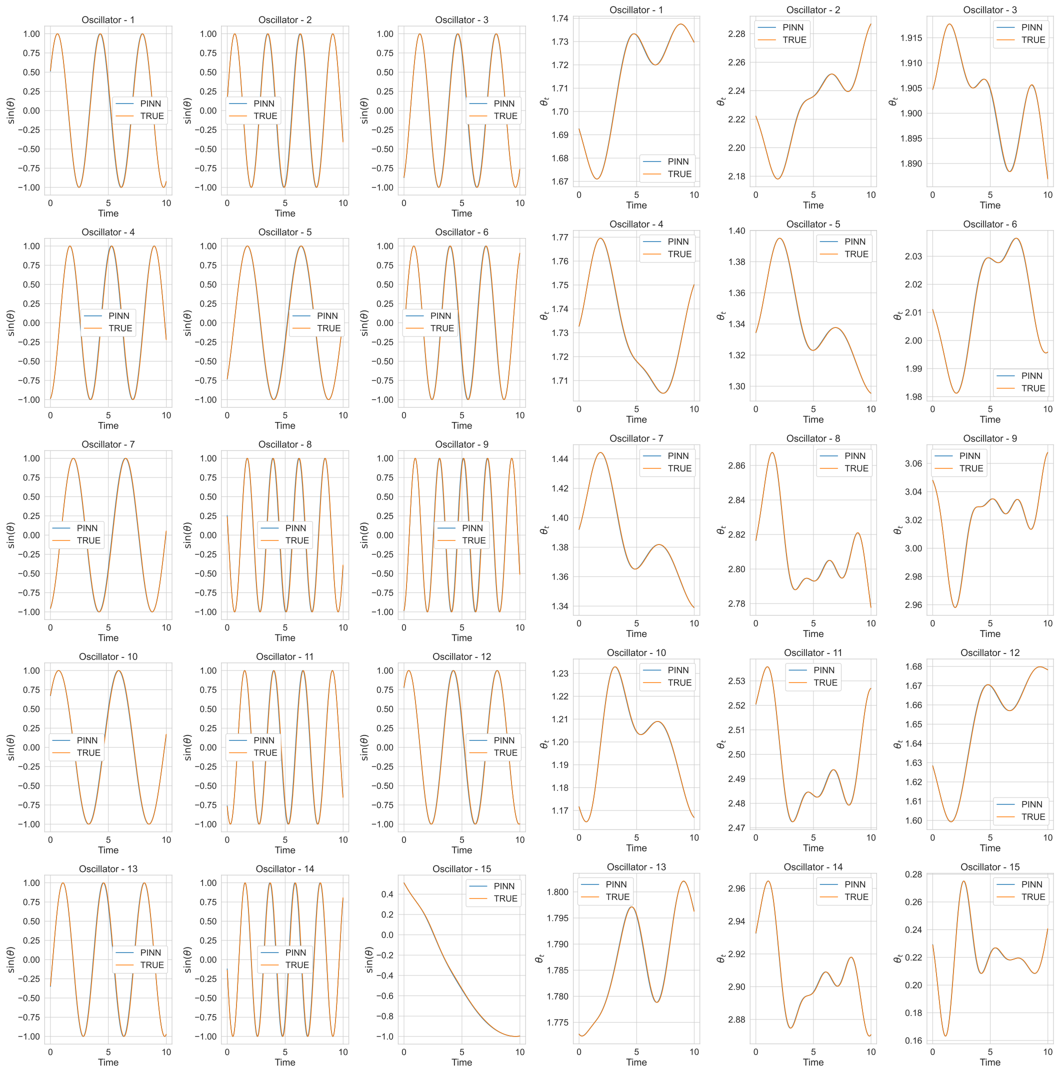

(a) Phase (b) Instantaneous frequency

Figure 8: Phases and instantaneous frequency of each oscillator($N = 15, K = 0.1$)

- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper shows that PINN struggles to solve strongly coupled Kuramoto oscillator network.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings,

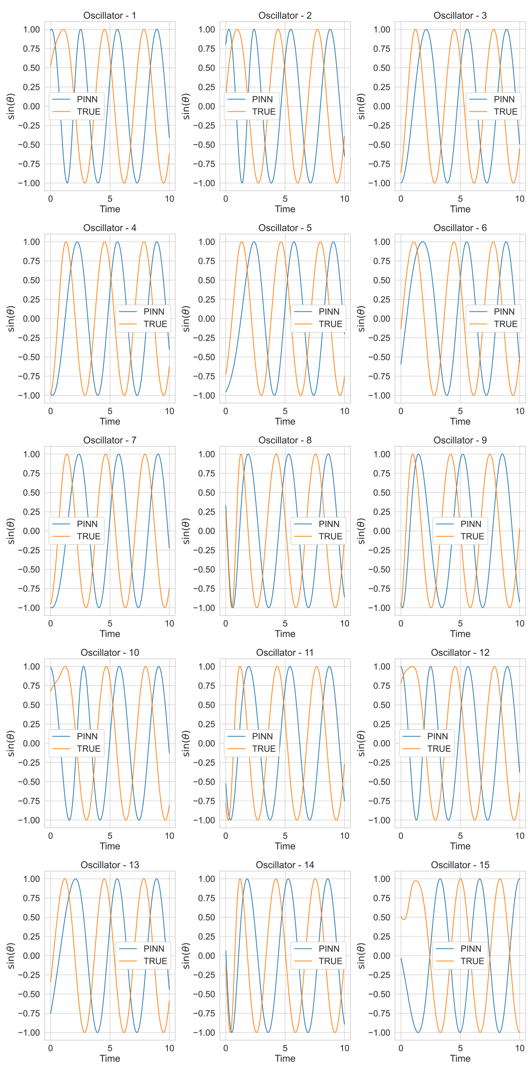

(a) Phase

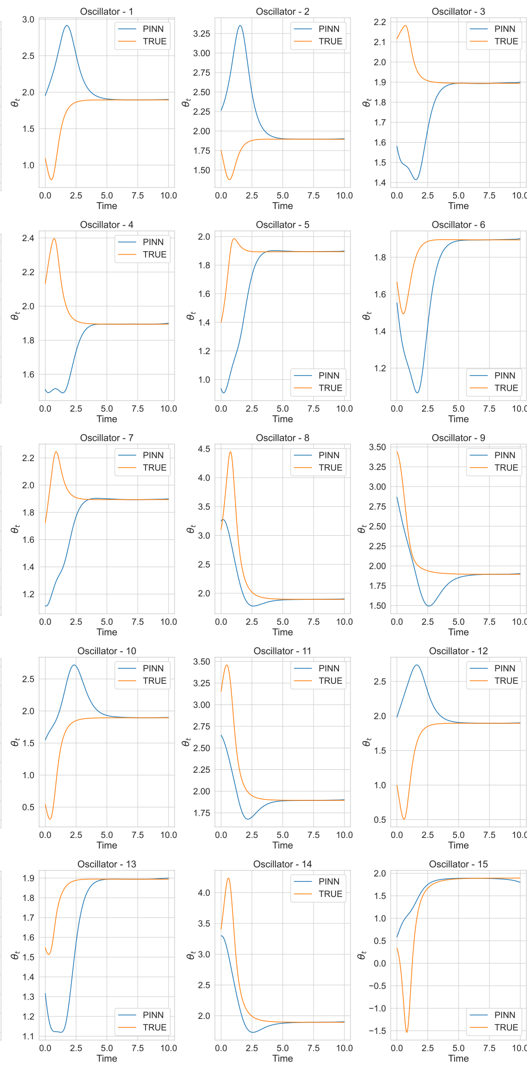

(b) Instantaneous frequency

Figure 9: Phases and instantaneous frequency of each oscillator($N = 15, K = 2.0$)

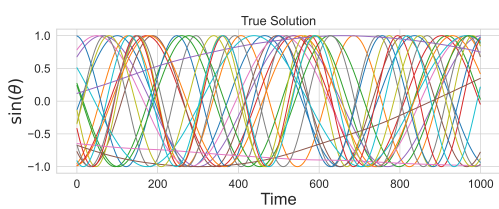

(a) True Solution($N = 20, K = 0.1$)

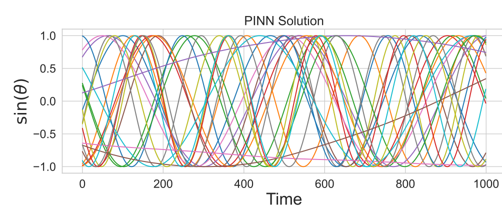

(b) Solution from PINN($N = 20, K = 0.1$)

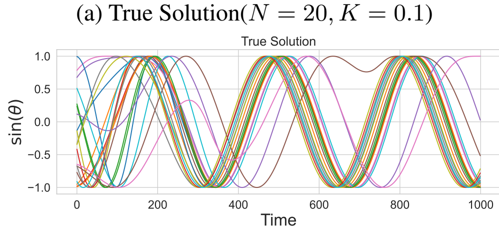

(c) True Solution($N = 20, K = 2.0$)

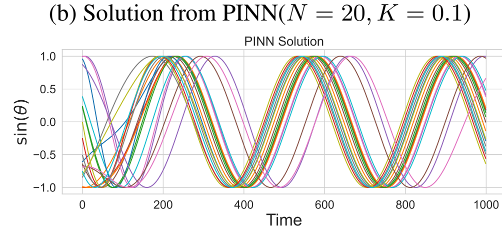

(d) Solution from PINN($N = 20, K = 2.0$)

Figure 10: True and PINN solution for weakly and strongly coupled Kuramoto network with 20 oscillators.

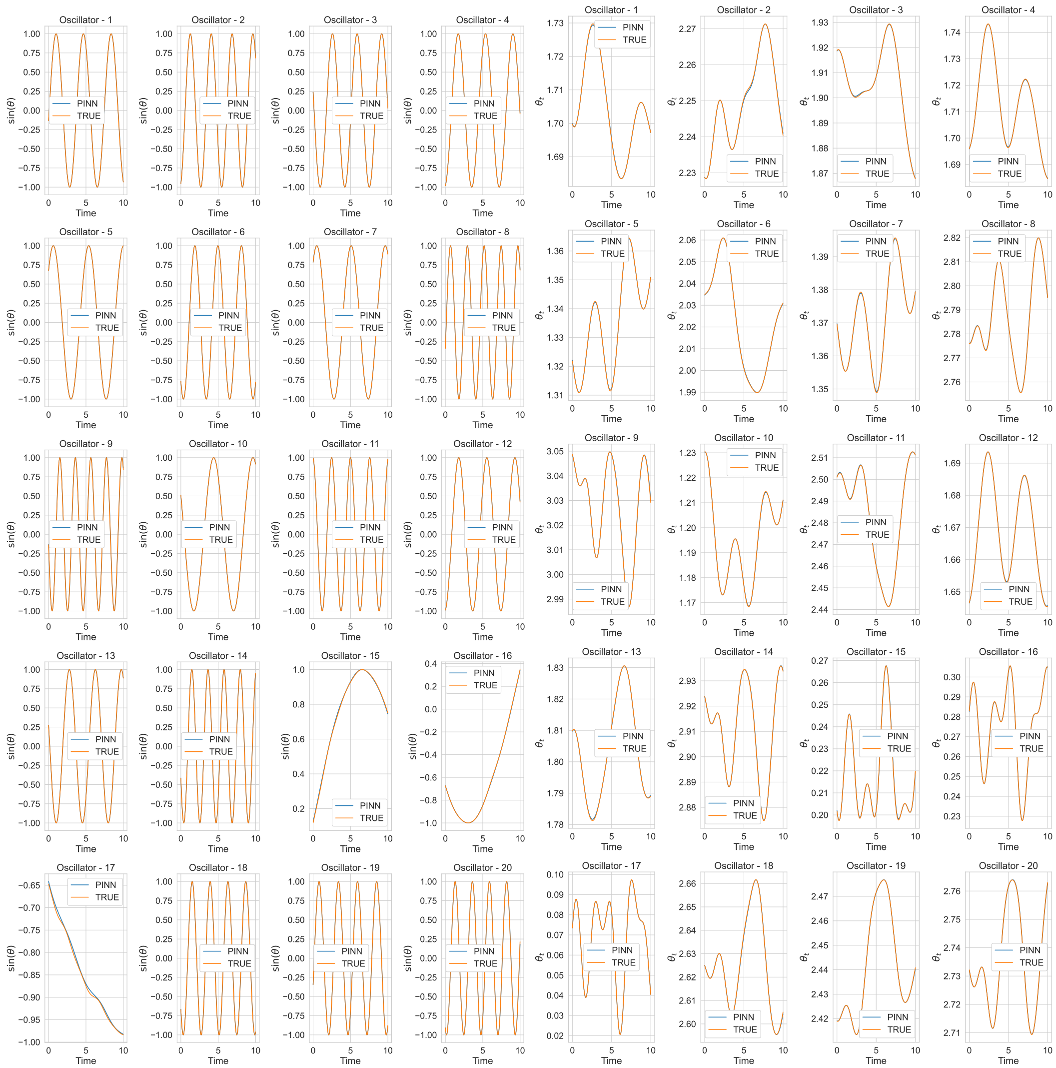

(a) Phase (b) Instantaneous frequency

Figure 11: Phases and instantaneous frequency of each oscillator($N = 20, K = 0.1$)

model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.

- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover

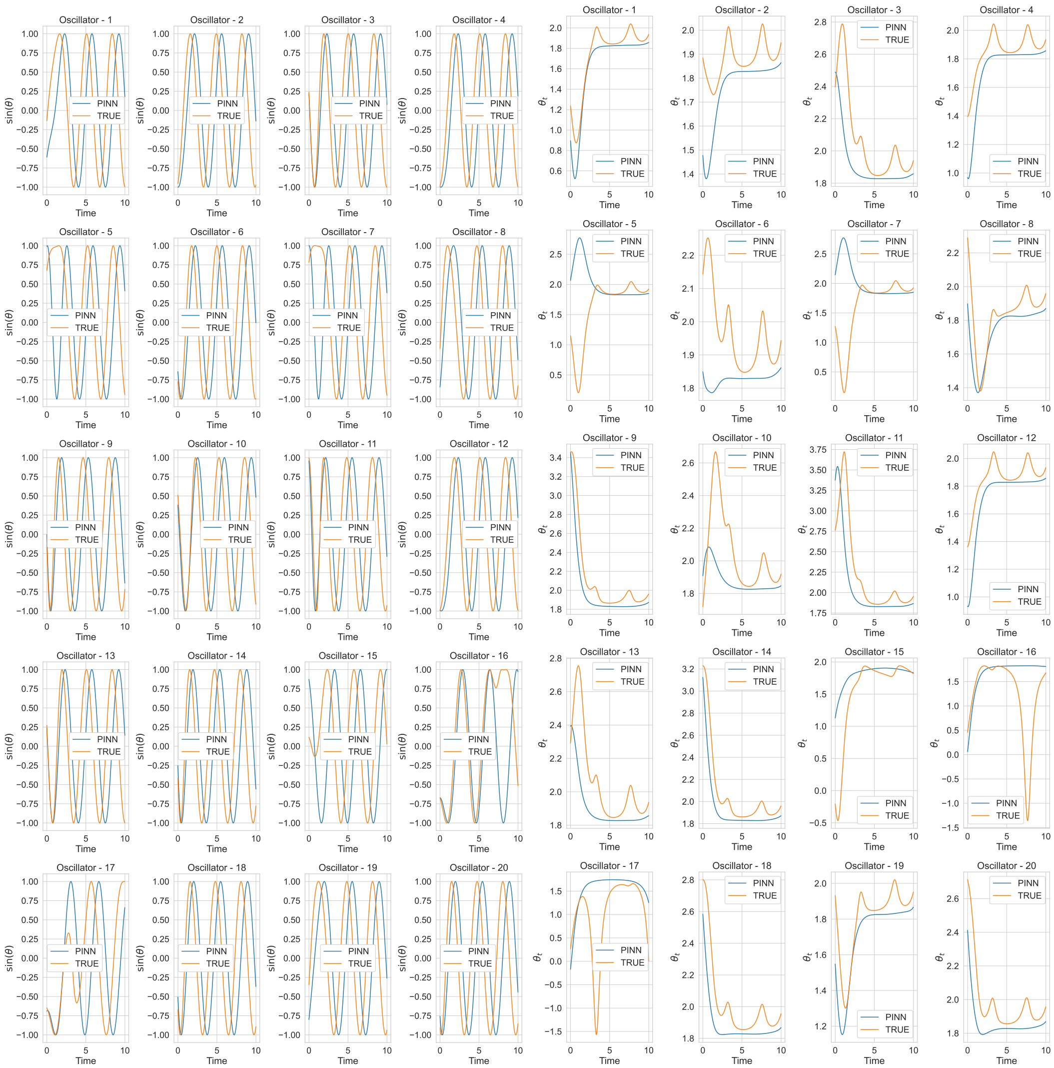

(a) Phase

(b) Instantaneous frequency

Figure 12: Phases and instantaneous frequency of each oscillator($N = 20, K = 2.0$)

limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: The theoretical background behind Kuramoto oscillaor network has been discussed in the paper. Also the optimization losses has been described in the paper.

   Guidelines:

   - The answer NA means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.

- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The implementation details has been discussed, and codes will be released upon the acceptance of this paper.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: The implementation details has been discussed, and codes will be released upon the acceptance of this paper.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.

- Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: The paper discusses all the training details and hyperparameters.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: The paper shows the PINN solver's performance for different number of oscillators.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.

- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The paper provides sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research conducted in the paper conforms, in every respect, with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The problem we solve in this paper does not leave any negative impact on society.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to

generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [Yes]

Justification: No safeguards were needed.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All sources have been properly credited with citation.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: Training details have been discussed and code implementations will be released upon the acceptance of the paper.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No human subjects involved.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No human subjects involved.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.