

Technical Support Center: Improving Physics-Informed Neural Network (PINN) Stability

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: *Pdic-NN*

Cat. No.: *B15614678*

[Get Quote](#)

This technical support center provides troubleshooting guides and frequently asked questions (FAQs) to help researchers, scientists, and drug development professionals overcome common challenges encountered during the training of Physics-Informed Neural Networks (PINNs).

Troubleshooting Guide

This section addresses specific issues that can arise during PINN training, offering potential solutions and detailed experimental protocols.

Q1: My PINN training loss is stagnating or oscillating significantly. What are the first steps to troubleshoot this?

A1: Loss stagnation or oscillation is a common issue in PINN training, often pointing to problems with the learning rate, network architecture, or the balance of loss terms.

Initial Diagnostic Steps:

- **Adjust the Learning Rate:** A learning rate that is too high can cause the loss to oscillate, while one that is too low can lead to stagnation. Start with a moderately large learning rate (e.g., 0.1) and observe the loss evolution. If it oscillates, gradually reduce the learning rate. [\[1\]](#) A ReduceLROnPlateau learning rate scheduler, available in both TensorFlow and PyTorch, can be highly effective. This callback allows you to decrease the learning rate when the loss metric has stopped improving.[\[1\]](#)

- **Examine Network Architecture:** For many PDE problems, shallow and wide networks tend to perform better than deep and narrow ones.^[1] If you are using a very deep network, try reducing the number of hidden layers and increasing the number of neurons per layer.
- **Check Input Normalization:** Failing to normalize the input data to a consistent range (e.g., $[-1, 1]$) can significantly hinder convergence.^[1] It is crucial to incorporate this normalization step into the network architecture itself to ensure that the gradients are calculated correctly with respect to the original, un-normalized inputs.^[1]

Experimental Protocol: Implementing Input Normalization within the Network

- **Determine Domain Boundaries:** Identify the minimum and maximum values for each of your input dimensions (e.g., x_{\min} , x_{\max} , t_{\min} , t_{\max}).
- **Create a Normalization Layer:** Add a preliminary layer to your neural network model that scales the inputs. This can be a simple lambda layer or a custom layer.
 - TensorFlow/Keras Example:
- **Train the Network:** Proceed with training as usual. The normalization is now an integral part of the model, and the automatic differentiation will correctly handle the scaling.

Q2: My PINN converges to a trivial or physically incorrect solution. How can I guide the training towards the correct solution?

A2: This often indicates an imbalance in the loss function, where the network prioritizes minimizing one component of the loss (e.g., the PDE residual) at the expense of others (e.g., boundary or initial conditions), or it may be converging to an unstable fixed point of the system.

Potential Solutions:

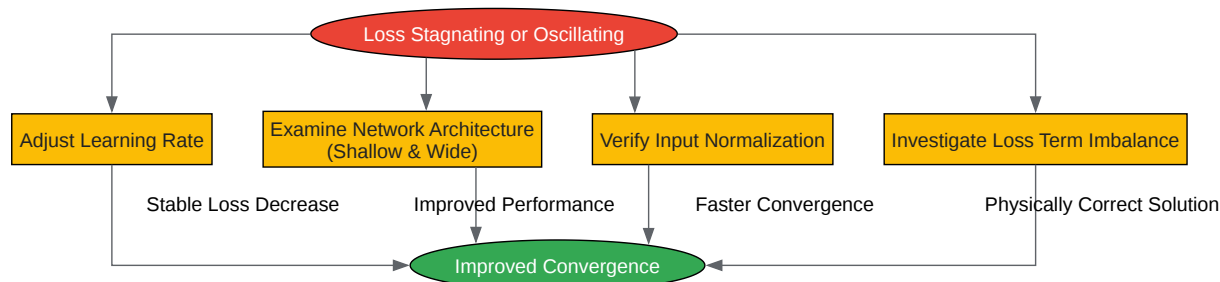
- **Loss Weighting:** Manually or dynamically adjusting the weights of the different loss components is a critical step. There is no one-size-fits-all set of weights, and the optimal values are problem-dependent. Recent research highlights the importance of proper weighting to balance data fitting and physics consistency.^{[2][3]}

- **Manual Weighting:** Start by assigning equal weights to all loss terms. If the network is failing to satisfy the boundary conditions, for example, increase the weight of the boundary loss term.
- **Dynamic Weighting:** More advanced techniques involve dynamically updating the weights during training. One approach is to use a method based on the Neural Tangent Kernel (NTK) to adaptively calibrate the convergence rate of different loss components.^{[4][5]} Another method, Dynamically Normalized PINNs (DN-PINNs), determines the relative weights based on gradient norms, which are updated during training.^[6]
- **Regularization for Dynamical Systems:** For problems involving dynamical systems, the PINN might converge to an unstable fixed point, which is a valid mathematical solution to the PDE but is not the physically correct one. A regularization scheme can be introduced to penalize solutions that correspond to unstable fixed points.^{[7][8]} This involves calculating the Jacobian of the system at collocation points and adding a penalty term to the loss if the eigenvalues indicate instability.^{[7][8]}
- **Adaptive Sampling:** Instead of a fixed set of collocation points, adaptively resample points in regions with high errors during training. This focuses the network's attention on the areas where it is struggling the most.^[9]

Experimental Protocol: Implementing a Simple Manual Loss Weighting Scheme

- **Define Individual Loss Components:** In your training script, calculate the loss for the PDE residual (`loss_pde`), boundary conditions (`loss_bc`), and initial conditions (`loss_ic`) separately.
- **Introduce Weight Hyperparameters:** Create trainable or tunable weight parameters (e.g., `w_pde`, `w_bc`, `w_ic`).
- **Combine the Losses:** The total loss is a weighted sum of the individual components:
$$\text{total_loss} = w_{\text{pde}} * \text{loss_pde} + w_{\text{bc}} * \text{loss_bc} + w_{\text{ic}} * \text{loss_ic}.$$
- **Tune the Weights:** Start with equal weights (e.g., `w_pde = 1.0`, `w_bc = 1.0`, `w_ic = 1.0`).
- **Iterate and Observe:** If, for instance, the solution at the boundaries is inaccurate, increase `w_bc` (e.g., to 10.0 or 100.0) and retrain. Monitor the individual loss components to see how they respond to the new weighting.

Logical Relationship for Troubleshooting Loss Stagnation



[Click to download full resolution via product page](#)

Caption: A flowchart for troubleshooting PINN training instability.

Frequently Asked Questions (FAQs)

Q: What is the best optimizer for training PINNs?

A: The choice of optimizer significantly impacts the performance of PINNs. While Adam is a popular and robust choice for initial training stages, quasi-Newton methods like L-BFGS can achieve more accurate results in fewer iterations.^{[10][11]} A common and effective strategy is a two-stage approach:

- Adam: Use the Adam optimizer for a large number of initial iterations to navigate the complex loss landscape and avoid saddle points.^[10]
- L-BFGS: Switch to the L-BFGS optimizer to fine-tune the solution and accelerate convergence to a sharp minimum.^{[10][11]}

Recent studies have also shown that advanced optimizers like SSBroyden with Wolfe line-search can be highly effective and reliable for training PINNs.^[12]

Optimizer	Strengths	Weaknesses	Recommended Usage
Adam	Robust, good for initial exploration of the loss landscape.	Can struggle to converge to sharp minima.	Use for the initial phase of training (e.g., first 1000-10,000 iterations). [12] [13]
L-BFGS	Incorporates second-order information for faster convergence to sharp minima. [12] [13]	More prone to getting trapped in local minima or saddle points if used from the start. [10]	Use after an initial training phase with Adam for fine-tuning. [10] [11]
Adam + L-BFGS	Combines the strengths of both optimizers. [10] [11]	Requires a two-stage training process.	A highly recommended state-of-the-art training scheme. [10]
SSBroyden	Strong convergence properties, effective for complex PDEs. [12]	Less commonly available in standard deep learning libraries.	For advanced users seeking optimal performance. [12]

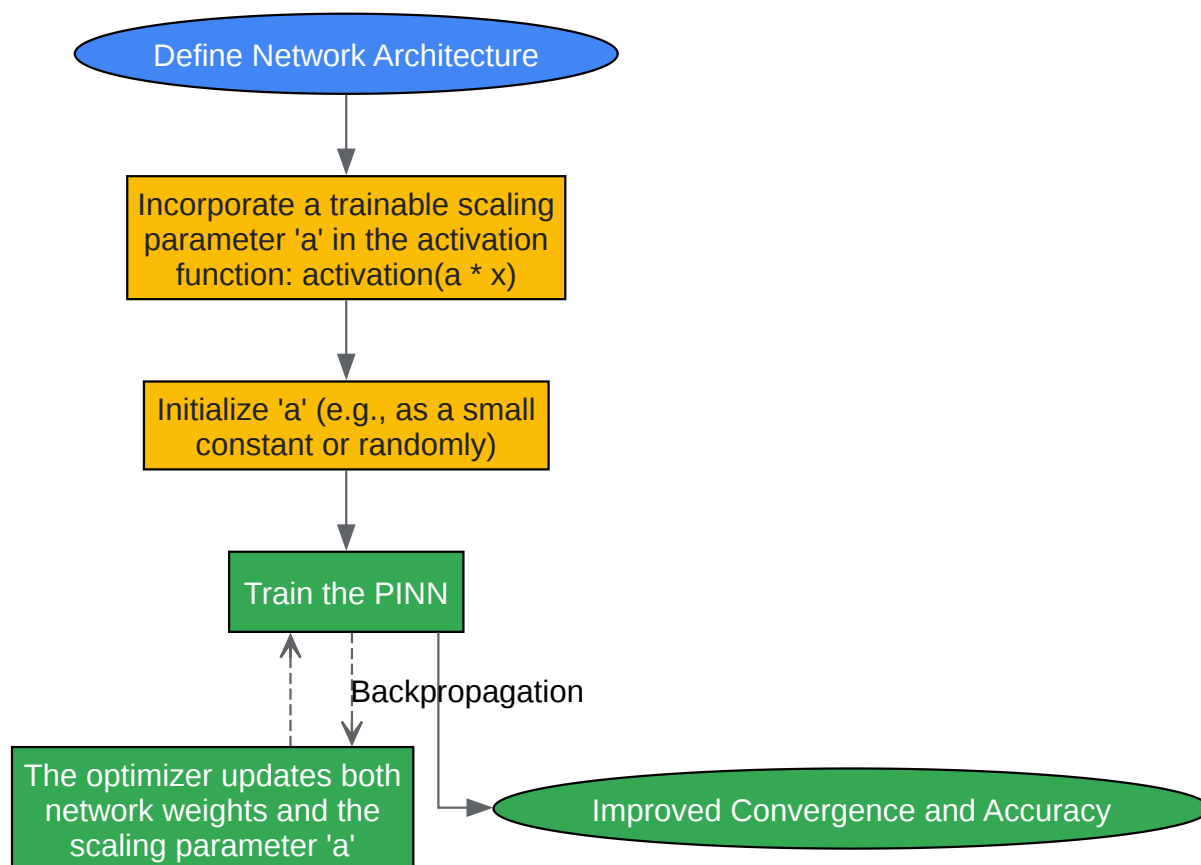
Q: How do I choose the right activation function for my PINN?

A: The choice of activation function is more critical in PINNs than in standard neural networks because the network's outputs are differentiated multiple times.[\[1\]](#)

- **Differentiability:** The activation function must be differentiable at least $n + 1$ times, where n is the order of the highest derivative in your PDE.[\[1\]](#) For example, if your PDE involves a second-order derivative, you need an activation function with at least three non-zero derivatives. This makes functions like tanh and sin suitable for many problems, while ReLU is often a poor choice due to its non-differentiable point at zero.
- **Adaptive Activation Functions:** Introducing a scalable hyperparameter within the activation function can significantly improve convergence rates and accuracy.[\[14\]](#)[\[15\]](#) This hyperparameter can be made trainable, allowing the network to learn the optimal activation

function shape for the specific problem.[14][16] For example, a common adaptive activation function is $\sigma(a * x)$, where a is a trainable parameter.

Experimental Workflow for Implementing Adaptive Activation Functions



[Click to download full resolution via product page](#)

Caption: Workflow for using adaptive activation functions in PINNs.

Q: My PINN is very sensitive to the network architecture. Are there any general guidelines for designing the network?

A: While the optimal architecture is problem-dependent, here are some general guidelines that have proven effective:

- **Shallow and Wide Networks:** As a rule of thumb, start with a network that is wider rather than deeper.^[1] For example, a network with 4 hidden layers and 50 neurons per layer is often a better starting point than a network with 10 hidden layers and 20 neurons per layer.
- **Ensemble Methods:** To improve stability and accuracy, consider using an ensemble of PINNs. Training multiple PINNs with different random initializations and averaging their predictions can help to avoid convergence to incorrect solutions.^[17]
- **Specialized Architectures:** For complex problems, more advanced architectures might be necessary:
 - **XPINNs (Extended PINNs):** These architectures use domain decomposition, breaking a large, complex problem into smaller, simpler sub-problems. This can be particularly useful for problems with discontinuities.
 - **MoE-PINNs (Mixture of Experts PINNs):** This approach uses a gating network to combine the predictions of several specialized PINNs, each potentially with a different activation function. This has been shown to work consistently well.^[1]

Comparison of Architectural Strategies

Strategy	Description	Best For
Shallow & Wide	Fewer hidden layers, more neurons per layer.	General starting point for most PDE problems. ^[1]
Ensemble PINNs	Averaging predictions from multiple independently trained PINNs.	Improving robustness and avoiding convergence to spurious local minima. ^[17]
XPINNs	Decomposing the computational domain into subdomains.	Problems with complex geometries or discontinuities.
MoE-PINNs	Using a gating network to weight the outputs of multiple "expert" PINNs.	Problems where different regions of the domain might benefit from different network properties. ^[1]

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- 1. medium.com [medium.com]
- 2. Impact of Loss Weight and Model Complexity on Physics-Informed Neural Networks for Computational Fluid Dynamics [arxiv.org]
- 3. mdpi.com [mdpi.com]
- 4. openreview.net [openreview.net]
- 5. When and why PINNs fail to train: A neural tangent kernel perspective | alphaXiv [alphaxiv.org]
- 6. GitHub - ShotaDeguchi/DN_PINN: Dynamic normalization with and without bias correction for physics-informed neural networks [github.com]
- 7. Stabilizing PINNs: A regularization scheme for PINN training to avoid unstable fixed points of dynamical systems [arxiv.org]
- 8. themoonlight.io [themoonlight.io]
- 9. pubs.aip.org [pubs.aip.org]
- 10. RUA [rua.ua.es]
- 11. researchgate.net [researchgate.net]
- 12. Which Optimizer Works Best for Physics-Informed Neural Networks and Kolmogorov-Arnold Networks? [arxiv.org]
- 13. Optimizing the Optimizer for Physics-Informed Neural Networks and Kolmogorov-Arnold Networks [arxiv.org]
- 14. [1906.01170] Adaptive activation functions accelerate convergence in deep and physics-informed neural networks [arxiv.org]
- 15. pubs.aip.org [pubs.aip.org]
- 16. [2308.04073] Learning Specialized Activation Functions for Physics-informed Neural Networks [arxiv.org]
- 17. medium.com [medium.com]

- To cite this document: BenchChem. [Technical Support Center: Improving Physics-Informed Neural Network (PINN) Stability]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b15614678#how-to-improve-the-stability-of-pinn-training]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [[Contact our Ph.D. Support Team for a compatibility check](#)]

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd
Ontario, CA 91761, United States
Phone: (601) 213-4426
Email: info@benchchem.com