

Technical Support Center: Training Physics-Informed Neural Networks (PINNs)

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

Compound of Interest

Compound Name: *Pdic-NN*

Cat. No.: *B15614678*

[Get Quote](#)

Welcome to the technical support center for Physics-Informed Neural Networks (PINNs). This resource is designed for researchers, scientists, and drug development professionals to provide troubleshooting guidance and answer frequently asked questions encountered during the training of PINNs.

Troubleshooting Guide

This guide provides solutions to common problems encountered during PINN training.

Issue 1: The training loss is not decreasing or is decreasing very slowly.

This is a common issue that can stem from several underlying problems, including an imbalanced loss function, vanishing or exploding gradients, or an inappropriate learning rate.

- Troubleshooting Steps:
 - Verify Loss Component Scaling: The different terms in your loss function (e.g., PDE residual, boundary conditions, initial conditions) might have vastly different magnitudes.[1][2][3] This can cause the optimizer to prioritize one term over the others.
 - Action: Implement a loss balancing strategy. Common techniques include using adaptive weights or annealing the learning rate for different loss components.[2][4]
 - Check for Gradient Pathologies: Unbalanced back-propagated gradients can stall training.[4][5] This "stiffness" in the gradient flow is a known failure mode for PINNs.[6]

- Action: Employ learning rate annealing algorithms that use gradient statistics to balance the interplay between different loss terms.[4] Consider using modified neural network architectures designed to be more resilient to these issues.[4][6]
- Adjust the Learning Rate: An inappropriate learning rate is a frequent cause of training problems.
 - Action: Experiment with different learning rates. A learning rate that is too high can cause the optimization to diverge, while one that is too low can lead to very slow convergence. Consider using a learning rate scheduler.
- Examine Network Initialization: Poor weight initialization can hinder the training process.[7]
 - Action: Use standard initialization techniques like Xavier or He initialization.

Issue 2: The model converges to a trivial or physically incorrect solution.

Even if the loss decreases, the PINN might learn a solution that is physically implausible or simply zero everywhere.

- Troubleshooting Steps:
 - Review Collocation Point Sampling: The distribution and number of collocation points are crucial for accurately enforcing the PDE residual.
 - Action: Ensure you have a sufficient number of collocation points sampled across the entire domain, including the boundaries.[8] A common practice is to have a similar number of points on the boundaries as inside the domain.[8] Consider re-sampling the collocation points at each iteration.[8]
 - Analyze the Loss Function Formulation: An incorrectly formulated loss function can lead the model to a trivial solution.
 - Action: Double-check the implementation of your PDE residual and boundary condition terms in the loss function. Ensure the relative weighting of these terms is appropriate.
 - Address Spectral Bias: PINNs can have difficulty learning high-frequency solutions, a phenomenon known as spectral bias.[9][10]

- Action: Consider using techniques like Fourier feature embeddings to help the network learn higher frequency functions.

Issue 3: The model shows good performance on the training data but fails to generalize to new data.

This is a classic case of overfitting, where the model has memorized the training data instead of learning the underlying physical laws.

- Troubleshooting Steps:
 - Increase the Number of Collocation Points: A denser sampling of the domain for the physics loss can act as a form of regularization.
 - Action: Increase the number of collocation points and ensure they are well-distributed.
 - Simplify the Network Architecture: An overly complex network is more prone to overfitting.
 - Action: Try reducing the number of layers or neurons in your network.
 - Introduce Regularization: While the PDE itself is a regularizer, explicit regularization techniques can sometimes be beneficial.
 - Action: Consider adding L1 or L2 regularization to the network weights.

Frequently Asked Questions (FAQs)

Q1: How do I balance the different terms in the PINN loss function?

Balancing the loss terms for the PDE residual, boundary conditions, and initial conditions is critical for successful training.^{[1][2][3]} These terms can have different magnitudes and physical units, leading to an imbalanced optimization problem.^[1]

- Answer: Several strategies can be employed:
 - Manual Weighting: Assign weights to each loss component. This often requires a trial-and-error approach to find suitable values.

- Adaptive Weighting: Use an algorithm that automatically adjusts the weights during training. Some methods are based on the magnitudes of the gradients of each loss term.
- Learning Rate Annealing: This technique involves adapting the learning rate for each loss component based on gradient statistics during training.[4]
- Self-Adaptive Loss Balancing: Methods like ReLoBRaLo (Relative Loss Balancing with Random Lookback) have been proposed to automate this process and have shown to improve training speed and accuracy.[2][3]

Q2: What is the best optimizer to use for training PINNs?

The choice of optimizer can significantly impact the training dynamics and final accuracy of a PINN.

- Answer: There is no single "best" optimizer for all PINN problems. However, a common and effective strategy is to use a combination of optimizers.[11] Many practitioners start with the Adam optimizer for a certain number of iterations to quickly navigate the loss landscape and then switch to a second-order optimizer like L-BFGS for fine-tuning.[11][12] This is because Adam is generally good at finding a good region of the loss landscape, while L-BFGS is efficient at finding the local minimum within that region.[12]

Q3: How do I choose the right neural network architecture?

The architecture of the neural network, including the number of layers (depth) and neurons per layer (width), is a critical hyperparameter.[7][13][14]

- Answer: The optimal architecture is problem-dependent. However, here are some general guidelines:
 - Start Simple: Begin with a smaller network and gradually increase its complexity if needed.
 - Hyperparameter Optimization: For complex problems, systematic hyperparameter optimization techniques, such as Bayesian optimization or neural architecture search (NAS), can be employed to find an optimal architecture.[13][14][15]

- Residual Connections: For deeper networks, incorporating residual connections can help with the flow of gradients and improve training.[8]

Q4: What activation function should I use?

The choice of activation function is more critical in PINNs than in standard neural networks because its derivatives are used to compute the PDE residual.[8]

- Answer: The activation function must be differentiable up to the order of the derivatives in your PDE.[8]
 - Common Choices: tanh and sin are popular choices as they are infinitely differentiable and have non-zero higher-order derivatives.
 - Avoid ReLU: Standard ReLU activation functions are not suitable for PINNs because their second derivative is zero everywhere, which can be problematic for solving second-order or higher PDEs.

Q5: How many collocation points are needed?

The number of collocation points used to enforce the physics loss is a crucial hyperparameter.

- Answer: There is no universal number, but some rules of thumb exist:
 - Sufficient Sampling: You need enough points to accurately represent the solution's complexity over the entire domain.
 - Boundary vs. Interior: A good starting point is to have a similar cumulative number of points on the boundaries as inside the domain.[8]
 - Adaptive Sampling: For problems with sharp gradients or complex behavior in certain regions, adaptive sampling strategies that place more points in these areas can be beneficial.

Quantitative Data Summary

The following table summarizes the impact of different optimizers on the final L2 error for various PDEs, as reported in a study on the PINN loss landscape.[12]

PDE	Optimizer	Mean Relative L2 Error
Convection	Adam	1.00e+00
L-BFGS	1.00e+00	
Adam + L-BFGS	8.12e-02	
Wave	Adam	1.00e+00
L-BFGS	1.00e+00	
Adam + L-BFGS	1.00e+00	
Reaction	Adam	1.00e+00
L-BFGS	1.00e+00	
Adam + L-BFGS	1.00e+00	

Note: The study highlights that while Adam + L-BFGS often performs best, achieving low error remains a significant challenge for certain PDEs.[\[12\]](#)

Experimental Protocols


Protocol 1: Combined Adam and L-BFGS Optimization

This protocol describes a common and effective training strategy for PINNs that leverages both a first-order and a quasi-second-order optimizer.[\[11\]](#)[\[12\]](#)

- Initialization: Initialize the neural network parameters.
- First-Order Optimization: Train the PINN using the Adam optimizer for a predefined number of iterations (e.g., 10,000 to 50,000 iterations). This phase is intended to quickly find a good region in the loss landscape.
- Second-Order Optimization: After the initial Adam training, switch to the L-BFGS optimizer for further training. L-BFGS is a quasi-Newton method that can converge faster and to a better minimum when in the vicinity of one.


- Termination: Continue training with L-BFGS until a convergence criterion is met (e.g., the change in loss falls below a threshold or a maximum number of iterations is reached).

Visualizations



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

Caption: A flowchart for troubleshooting common PINN training issues.



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

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. mdpi.com [mdpi.com]
- 2. Scientific Machine Learning | Physics-Informed Deep Learning and Loss Balancing - Michael Kraus Anton [mkrausai.com]
- 3. researchgate.net [researchgate.net]
- 4. epubs.siam.org [epubs.siam.org]
- 5. [PDF] Understanding and mitigating gradient pathologies in physics-informed neural networks | Semantic Scholar [semanticscholar.org]
- 6. Improved physics-informed neural network in mitigating gradient-related failures [arxiv.org]
- 7. publicacoes.softaliza.com.br [publicacoes.softaliza.com.br]
- 8. medium.com [medium.com]

- 9. When and why PINNs fail to train: A neural tangent kernel perspective | alphaXiv [alphaxiv.org]
- 10. openreview.net [openreview.net]
- 11. caeassistant.com [caeassistant.com]
- 12. arxiv.org [arxiv.org]
- 13. Auto-PINN: Understanding and Optimizing Physics-Informed Neural Architecture [arxiv.org]
- 14. researchgate.net [researchgate.net]
- 15. [2205.06704] Hyper-parameter tuning of physics-informed neural networks: Application to Helmholtz problems [arxiv.org]
- To cite this document: BenchChem. [Technical Support Center: Training Physics-Informed Neural Networks (PINNs)]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b15614678#common-challenges-in-training-physics-informed-neural-networks]

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

