

Technical Support Center: Debugging Physics-Based Loss Functions in PINNs

Author: BenchChem Technical Support Team. **Date:** April 2026

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 assist researchers, scientists, and drug development professionals in debugging the physics-based loss function in Physics-Informed Neural Networks (PINNs).

Troubleshooting Guide

This guide addresses specific issues you might encounter during your experiments with PINNs, offering step-by-step solutions.

Issue 1: My PINN training is not converging, or the loss is stagnating at a high value.

Possible Causes:

- **Imbalanced Loss Terms:** The different components of your total loss function (e.g., PDE residual, boundary conditions, initial conditions) might have vastly different magnitudes, causing the optimizer to prioritize one term over the others.^{[1][2][3]}
- **Inappropriate Learning Rate:** The learning rate might be too high, causing oscillations, or too low, leading to slow convergence.

- **Poor Network Architecture:** The neural network may not have sufficient capacity (depth or width) to approximate the solution accurately.[\[4\]](#)
- **Challenging Loss Landscape:** The physics-based constraints can create a complex and non-convex loss landscape that is difficult for the optimizer to navigate.[\[5\]](#)[\[6\]](#)[\[7\]](#)[\[8\]](#)

Troubleshooting Steps:

- **Monitor Individual Loss Components:** Plot the evolution of each loss term (PDE residual, boundary conditions, etc.) separately during training. This will help you identify if one term is dominating or not decreasing.
- **Implement Loss Balancing Techniques:**
 - **Manual Weighting:** Start by assigning weights to each loss component and manually tune them. This is often a necessary first step to bring the losses to a similar order of magnitude.[\[3\]](#)
 - **Adaptive Weighting Methods:** Employ more advanced techniques that dynamically adjust the weights during training. Some popular methods include:
 - **GradNorm:** Normalizes the gradient magnitudes of different loss terms.[\[1\]](#)[\[9\]](#)
 - **ReLoBRaLo (Relative Loss Balancing with Random Lookback):** Aims to ensure that each loss term makes similar relative progress over time.[\[1\]](#)[\[9\]](#)[\[10\]](#)
 - **SoftAdapt:** Adaptively adjusts weights based on the rate of change of each loss component.[\[1\]](#)[\[9\]](#)
- **Tune the Learning Rate:**
 - **Learning Rate Schedulers:** Use a learning rate scheduler, such as ReduceLROnPlateau, which decreases the learning rate when the loss plateaus.[\[4\]](#)
 - **Experiment with Different Optimizers:** Start with an adaptive optimizer like Adam for the initial phase of training to navigate the complex loss landscape, and then switch to a

second-order optimizer like L-BFGS for fine-tuning, as it can be more effective in the later stages.[\[4\]](#)[\[7\]](#)[\[8\]](#)[\[11\]](#)

- Adjust Network Architecture:
 - Increase Network Capacity: Try increasing the number of hidden layers or the number of neurons per layer. Shallow but wide networks are often a good starting point for PINNs.[\[4\]](#)
 - Experiment with Activation Functions: The choice of activation function is crucial as its derivatives are used to compute the PDE residual. Ensure the activation function is sufficiently differentiable for the order of your PDE.[\[4\]](#)[\[12\]](#) Functions like tanh or swish are often preferred over ReLU for higher-order PDEs.[\[13\]](#)[\[14\]](#)

Issue 2: My PINN exhibits exploding or vanishing gradients.

Possible Causes:

- Deep Network Architectures: The repeated multiplication of gradients through many layers can cause them to grow exponentially (explode) or shrink to zero (vanish).[\[15\]](#)[\[16\]](#)[\[17\]](#)
- High-Order Derivatives: The computation of high-order derivatives in the physics-based loss can lead to noisy and unstable gradients.[\[4\]](#)[\[10\]](#)
- Stiff PDE Problems: Some partial differential equations are inherently "stiff," meaning they involve processes with widely different scales, which can lead to gradient issues.[\[2\]](#)[\[18\]](#)

Troubleshooting Steps:

- Gradient Clipping: Set a threshold for the maximum value of the gradients. If a gradient exceeds this threshold, it will be clipped, preventing it from becoming excessively large.[\[15\]](#)
- Use Residual Connections (Skip Connections): These connections allow the gradient to flow more directly through the network, bypassing some layers and mitigating the vanishing gradient problem.[\[4\]](#)[\[15\]](#)[\[17\]](#)

- **Batch Normalization:** Normalize the inputs of each layer to have a mean of zero and a standard deviation of one. This can help stabilize the training process and reduce the likelihood of exploding or vanishing gradients.[\[15\]](#)[\[17\]](#)
- **Choose Appropriate Activation Functions:** As mentioned before, the choice of activation function can impact gradient flow. Functions like ReLU can sometimes lead to dead neurons (zero gradients), while tanh and its variants can help maintain a healthy gradient flow.
- **Curriculum Regularization:** Start by training the PINN on a simpler version of the PDE and gradually increase the complexity. This can help the network learn the basic physics before tackling the more challenging aspects of the problem.[\[5\]](#)

Frequently Asked Questions (FAQs)

Q1: What is the role of the physics-based loss function in a PINN?

The physics-based loss function is a core component of a PINN. It embeds the governing physical laws, typically in the form of partial differential equations (PDEs), directly into the training process.[\[2\]](#)[\[19\]](#) The total loss function of a PINN is a combination of the mean squared error of the training data (data loss) and the mean squared error of the PDE residual (physics loss).[\[3\]](#)[\[20\]](#) By minimizing this combined loss, the neural network learns a solution that not only fits the observed data but also adheres to the underlying physical principles.[\[19\]](#)

Q2: How do I balance the different terms in my loss function?

Balancing the different loss terms (e.g., PDE residual, boundary conditions, initial conditions) is critical for successful PINN training.[\[1\]](#)[\[2\]](#)[\[3\]](#) These terms can have different scales and units, leading to an imbalanced optimization problem.[\[2\]](#)

Here's a summary of common approaches:

Method	Description	Pros	Cons
Manual Weighting	Manually assign constant weights to each loss term.[3][21]	Simple to implement.	Requires extensive trial and error; static weights may not be optimal throughout training.
Learning Rate Annealing	A form of adaptive weighting where the weights are treated as learnable parameters.[1][9]	Can automatically find good weights.	May introduce additional hyperparameters to tune.
GradNorm	Dynamically adjusts weights to balance the magnitudes of the gradients of each loss term.[1][9]	Helps prevent one loss term from dominating the training.	Can be computationally more expensive.
ReLoBRaLo	A self-adaptive method that aims for each loss term to have a similar relative improvement.[1][9][10]	Often leads to faster training and higher accuracy.[1]	The concept can be more complex to grasp initially.
SoftAdapt	Adjusts weights based on the rate of change of each loss term.[1][9]	Responsive to the dynamics of the training process.	Performance can be sensitive to its own hyperparameters.

Q3: What are collocation points and how many should I use?

Collocation points are points sampled from the domain (spatial and temporal) where the PDE residual is evaluated.[3] These points do not need to have corresponding measurement data, which is a key advantage of PINNs.[19] The number of collocation points is a hyperparameter that needs to be tuned. Too few points may not be sufficient to enforce the physics across the entire domain, while too many can increase the computational cost of training.[3] A common

practice is to start with a number of collocation points that is an order of magnitude larger than the number of data points and adjust based on the performance.

Q4: How does the choice of optimizer affect the training of a PINN?

The choice of optimizer can significantly impact the training dynamics and final performance of a PINN. The loss landscape of a PINN is often highly non-convex and challenging to navigate. [\[5\]](#)[\[6\]](#)[\[7\]](#)[\[8\]](#)


- Adam: This is a popular first-order optimizer that is generally a good starting point. It is effective at navigating complex loss landscapes in the initial stages of training. [\[4\]](#)[\[7\]](#)[\[8\]](#)
- L-BFGS: This is a quasi-Newton method that can be very effective for fine-tuning in the later stages of training. [\[4\]](#)[\[7\]](#)[\[8\]](#) It often converges faster and to a better minimum when the loss landscape is smoother. A common strategy is to train with Adam for a certain number of epochs and then switch to L-BFGS. [\[4\]](#)[\[11\]](#)

Q5: Why is my PINN accurate on the boundaries but not in the interior of the domain?

This is a common failure mode in PINNs and often points to an imbalance in the loss terms. [\[5\]](#) The optimizer might be prioritizing the minimization of the boundary condition loss at the expense of the PDE residual loss in the interior. This can happen if the magnitude of the boundary loss is significantly smaller than the PDE loss, or if the weights are not appropriately balanced. Refer to the troubleshooting steps for imbalanced loss terms to address this issue.

Visualizations


PINN Loss Function Structure



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

Caption: Structure of a typical PINN loss function, showing the combination of data and physics-based loss components.

Debugging Workflow for PINN Loss



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

Caption: A logical workflow for troubleshooting common issues with the PINN physics-based loss function during training.

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. Scientific Machine Learning | Physics-Informed Deep Learning and Loss Balancing - Michael Kraus Anton [mkrausai.com]
- 2. mdpi.com [mdpi.com]
- 3. medium.com [medium.com]
- 4. medium.com [medium.com]
- 5. proceedings.neurips.cc [proceedings.neurips.cc]
- 6. researchgate.net [researchgate.net]
- 7. Challenges in Training PINNs: A Loss Landscape Perspective [arxiv.org]
- 8. arxiv.org [arxiv.org]
- 9. [2110.09813] Multi-Objective Loss Balancing for Physics-Informed Deep Learning [arxiv.org]
- 10. medium.com [medium.com]
- 11. Reddit - The heart of the internet [reddit.com]
- 12. ijcai.org [ijcai.org]
- 13. Auto-PINN: Understanding and Optimizing Physics-Informed Neural Architecture [arxiv.org]
- 14. researchgate.net [researchgate.net]
- 15. David Oniani | Exploding and Vanishing Gradients [oniani.org]
- 16. cs.toronto.edu [cs.toronto.edu]
- 17. Vanishing gradient problem - Wikipedia [en.wikipedia.org]

- [18. Improved physics-informed neural network in mitigating gradient-related failures \[arxiv.org\]](#)
- [19. mathworks.com \[mathworks.com\]](#)
- [20. medium.com \[medium.com\]](#)
- [21. tandfonline.com \[tandfonline.com\]](#)
- To cite this document: BenchChem. [Technical Support Center: Debugging Physics-Based Loss Functions in PINNs]. BenchChem, [2026]. [Online PDF]. Available at: [\[https://www.benchchem.com/product/b15614678/docs#technical-support-center-debugging-physics-based-loss-functions-in-pinns\]](https://www.benchchem.com/product/b15614678/docs#technical-support-center-debugging-physics-based-loss-functions-in-pinns)

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?

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

Contact our Ph.D. Support Team for a compatibility check