

## Addressing overfitting in physics-informed neural networks

**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 troubleshoot and address the common challenge of overfitting in their PINN experiments.

### Troubleshooting Guide: Is My PINN Overfitting?

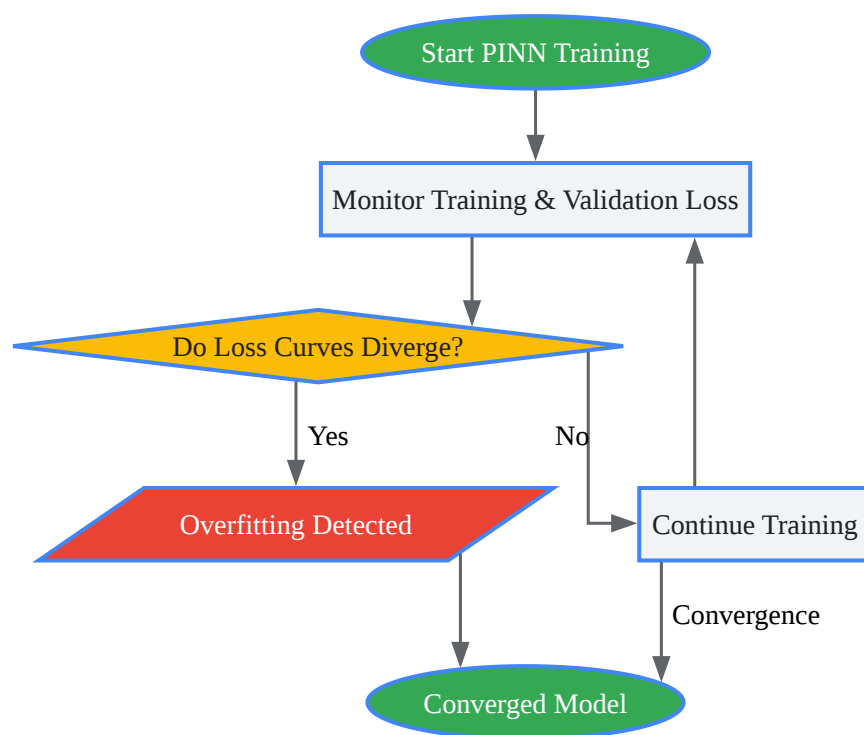
Overfitting is a critical issue where a PINN learns the training data too well, including noise and artifacts, leading to poor generalization and inaccurate predictions on new, unseen data. This guide provides a step-by-step approach to diagnose and mitigate overfitting.

#### Question 1: What are the common symptoms of an overfitted PINN?

An overfitted PINN will exhibit a significant discrepancy between its performance on the training data and on a validation or test set. Key symptoms include:

- **Low Training Loss, High Validation/Test Loss:** The model shows a very low error on the data it was trained on, but a much higher error when evaluated on data it has not seen before.<sup>[1]</sup>
- **Physically Inconsistent Solutions:** The model's predictions may violate the underlying physical laws in regions outside of the training data points, even if the physics-based loss is low.<sup>[2]</sup>
- **Sensitivity to Noise:** Small perturbations in the input data can lead to large, unphysical changes in the output.<sup>[3]</sup>
- **Poor Extrapolation:** The model fails to provide reasonable predictions for inputs outside the range of the training data.<sup>[4]</sup>

A common diagnostic workflow is to monitor the training and validation loss over epochs. A divergence in these two curves is a clear indicator of overfitting.<sup>[5]</sup>



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

Caption: Workflow for diagnosing overfitting by monitoring loss curves.

## Question 2: My PINN seems to be overfitting. What are the primary causes?

Overfitting in PINNs can stem from several factors, often related to the model's complexity, the training data, or the training process itself.

- **Excessive Model Complexity:** Deep neural networks with many layers or a large number of neurons per layer have a high capacity to memorize the training data, including noise.[6][7]
- **Insufficient or Poorly Distributed Training Data:** A small training dataset may not adequately represent the entire problem domain, making it easier for the network to overfit to the available points.[8] This is a known issue for PINNs, which can be susceptible to overfitting on boundary conditions if the number of collocation points is much larger than the number of boundary data points.[3]
- **Imbalanced Loss Function:** If the different components of the loss function (e.g., data loss, physics residual loss, boundary condition loss) have vastly different magnitudes, the training process might prioritize one term at the expense of others, leading to poor generalization.[9]
- **Overtraining:** Training for too many epochs can lead the model to start fitting the noise in the training data.[7]

## FAQs: Techniques for Mitigating Overfitting in PINNs

This section provides answers to frequently asked questions about specific techniques to combat overfitting.

### Data-Centric Approaches

Question 3: How can I leverage my data to reduce overfitting?

Answer:

- **Data Augmentation:** While traditional data augmentation techniques like rotation or flipping are common in computer vision, for PINNs, a "physics-guided data augmentation" (PGDA) approach is more suitable.<sup>[10][11]</sup> This involves generating new training data by leveraging physical properties of the system, such as linearity or translational invariance.<sup>[12]</sup> For example, if a PDE is linear, a linear combination of known solutions is also a valid solution and can be used as a new training sample.<sup>[12]</sup>
- **Adaptive Sampling:** Instead of uniformly sampling collocation points, adaptive sampling methods focus on regions where the PDE residual is high.<sup>[13]</sup> This forces the network to improve its accuracy in areas where it performs poorly, leading to better generalization. This can be more efficient than uniform sampling and can improve the convergence of the training process.<sup>[14]</sup>

## Architectural and Regularization Strategies

Question 4: How should I adjust my network architecture and apply regularization?

Answer:

Simplifying the model is often the first step in addressing overfitting.<sup>[10]</sup> Additionally, regularization techniques add a penalty to the loss function to discourage complex models.

Technique	Description	Impact on Overfitting
Reduce Network Complexity	Decrease the number of hidden layers or the number of neurons per layer. <sup>[5]</sup>	Reduces the model's capacity to memorize noise, forcing it to learn the underlying physical laws. <sup>[6]</sup>
L1 & L2 Regularization	Adds a penalty term to the loss function based on the magnitude of the network weights (L1: absolute values, L2: squared values). <sup>[15][16]</sup>	Penalizes large weights, leading to a simpler, more stable model that is less sensitive to small changes in the input. <sup>[5][11]</sup>
Dropout	Randomly deactivates a fraction of neurons during each training iteration. <sup>[10]</sup>	Prevents neurons from co-adapting and forces the network to learn more robust features. <sup>[15]</sup>
Physics-Informed Regularization	Incorporating physics-based terms into the loss function acts as a regularizer, constraining the solution space to physically plausible outcomes. <sup>[4][17]</sup>	Improves the model's extrapolation capabilities and ensures that the learned solution adheres to the governing physical laws. <sup>[4][18]</sup>
Ensemble Methods	Train multiple PINNs with different architectures or initializations and average their predictions. <sup>[19]</sup>	Reduces variance and improves robustness by combining the strengths of multiple models. <sup>[19]</sup>

## Experimental Protocol: Implementing L2 Regularization

- **Define the Loss Function:** The total loss for a PINN is typically a weighted sum of the data loss (ngcontent-ng-c4139270029="" \_ngghost-ng-c4104608405="" class="inline ng-star-inserted">

$$L_{data}$$

), the physics loss (ngcontent-ng-c4139270029="" \_ngghost-ng-c4104608405="" class="inline ng-star-inserted">

$L_{physics}$   $L_{physics}$

), and the boundary condition loss (ngcontent-ng-c4139270029="" \_ngghost-ng-c4104608405="" class="inline ng-star-inserted">

$L_{bc}$   $L_{bc}$

).

- Add the Regularization Term: Append the L2 regularization term to the total loss. This term is the sum of the squares of all the weights in the neural network, multiplied by a regularization parameter, ngcontent-ng-c4139270029="" \_ngghost-ng-c4104608405="" class="inline ng-star-inserted">

$\lambda$

.

$$L_{total} = w_{data} L_{data} + w_{physics} L_{physics} + w_{bc} L_{bc} + \lambda \sum_i ||W_i||_F^2$$

where

$W_i$

are the weight matrices of the network and

$|| \cdot ||_F$

is the Frobenius norm.

- Tune the Hyperparameter

$\lambda$

: The regularization parameter

$\lambda$

controls the strength of the penalty. A small

$\lambda$

will have a minor effect, while a large

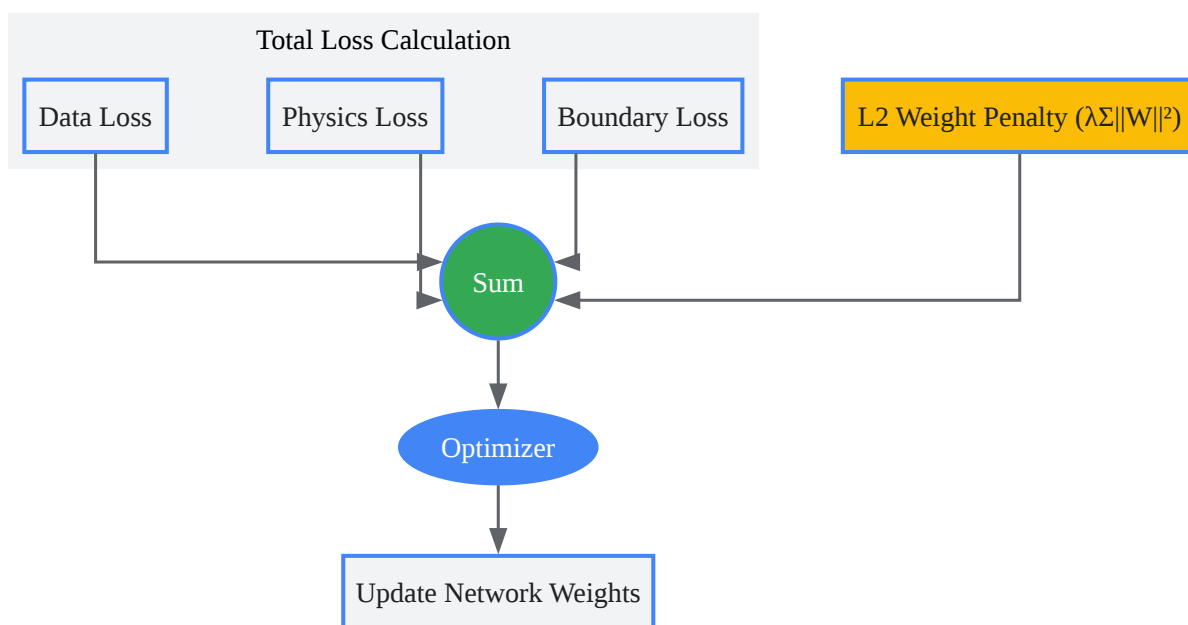
$\lambda$

can lead to underfitting. Use a validation set to find an optimal value for

$\lambda$

.

- Train the Model: Train the PINN using the modified loss function. The optimization process will now aim to minimize both the original loss components and the magnitude of the weights.



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

Caption: Logic of incorporating L2 regularization into the PINN loss function.

## Advanced Training Techniques

Question 5: Are there advanced training strategies to prevent overfitting?

Answer:

Yes, several advanced training strategies can help balance the different objectives in a PINN and improve convergence and generalization.

- **Adaptive Loss Balancing:** The magnitudes of the gradients from different loss components can vary significantly, causing training instabilities.<sup>[9]</sup> Adaptive weighting schemes dynamically adjust the weights of each loss term during training to ensure that they are on a similar scale.<sup>[20][21]</sup> This prevents the optimization from being dominated by a single term and promotes a more balanced training process.<sup>[9]</sup>
- **Learning Rate Scheduling:** Instead of using a fixed learning rate, a learning rate scheduler can be employed to decrease the learning rate as training progresses. A common approach is to start with a higher learning rate to quickly approach a minimum and then reduce it to fine-tune the model.<sup>[19]</sup>
- **Choice of Optimizer:** While Adam is a popular choice, combining it with a second-order optimizer like L-BFGS can be beneficial.<sup>[19]</sup> Adam is often used in the initial stages of training, and L-BFGS is used for fine-tuning once the loss has plateaued.<sup>[19][22]</sup>
- **Early Stopping:** This technique involves monitoring the performance of the model on a validation set and stopping the training process when the validation loss stops improving, even if the training loss continues to decrease.<sup>[1][10]</sup> This directly prevents the model from overtraining on the training data.<sup>[15]</sup>

Experimental Protocol: Adaptive Loss Balancing

- Initialize Loss Weights: Start with initial weights for each component of the loss function (e.g., `ngcontent-ng-c4139270029="" _ngghost-ng-c4104608405="" class="inline ng-star-inserted">`

$w_{data} = 1.0$

`ngcontent-ng-c4139270029="" _ngghost-ng-c4104608405="" class="inline ng-star-inserted">`

$w_{physics} = 1.0$

`ngcontent-ng-c4139270029="" _ngghost-ng-c4104608405="" class="inline ng-star-inserted">`

$w_{bc} = 1.0$

).

- Compute Gradients: During each training step, after computing the gradients of each loss component with respect to the network parameters, also compute statistics of these gradients (e.g., the mean or max).
- Update Weights: Update the loss weights based on a chosen heuristic. One common method is to update the weights inversely proportional to the magnitude of their respective gradients.<sup>[13]</sup> For example, for a loss term `ngcontent-ng-c4139270029="" _ngghost-ng-c4104608405="" class="inline ng-star-inserted">`

$L_i$

, its weight

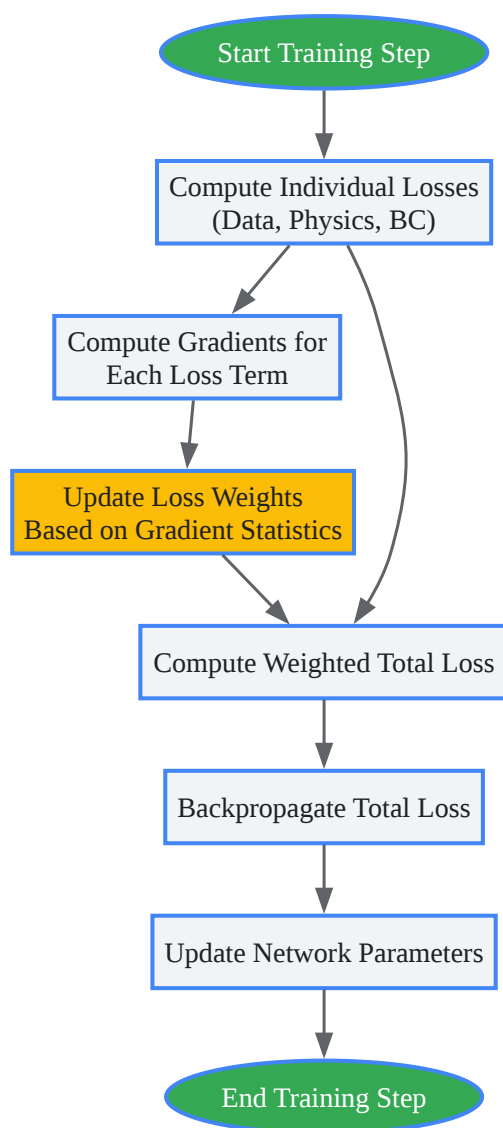
$w_i$

could be updated as:

$$w_i \leftarrow \frac{\text{mean}(\|\nabla_{\theta} L_{total}\|)}{\|\nabla_{\theta} L_i\|} w_i \leftarrow \frac{\text{mean}(\|\nabla_{\theta} L_{total}\|)}{\|\nabla_{\theta} L_i\|}$$

This aims to increase the influence of loss terms with smaller gradients.

- Normalize Weights: It is good practice to normalize the weights after each update to ensure they sum to a constant value.
- Apply Weighted Loss: Compute the total loss as the weighted sum of the individual loss components and perform the backpropagation step.



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

Caption: Workflow for adaptive loss balancing during PINN training.

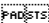
#### Need Custom Synthesis?

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

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

## References

- 1. What is Overfitting? | IBM [ibm.com]
- 2. mdpi.com [mdpi.com]
- 3. Recipes for when physics fails: recovering robust learning of physics informed neural networks - PMC [pmc.ncbi.nlm.nih.gov]
- 4. ml4physicalsciences.github.io [ml4physicalsciences.github.io]

- 5. medium.com [medium.com]
- 6. ijcsmc.com [ijcsmc.com]
- 7. kaggle.com [kaggle.com]
- 8. towardsai.net [towardsai.net]
- 9. medium.com [medium.com]
- 10. kdnuggets.com [kdnuggets.com]
- 11. analyticsvidhya.com [analyticsvidhya.com]
- 12. arxiv.org [arxiv.org]
- 13. Self-adaptive weighting and sampling for physics-informed neural networks [arxiv.org]
- 14. Ribbit Ribbit  Discover Research the Fun Way [ribbitribbit.co]
- 15. kaggle.com [kaggle.com]
- 16. towardsdatascience.com [towardsdatascience.com]
- 17. [PDF] Physics-Informed Regularization of Deep Neural Networks | Semantic Scholar [semanticscholar.org]
- 18. [1810.05547] Physics-Driven Regularization of Deep Neural Networks for Enhanced Engineering Design and Analysis [arxiv.org]
- 19. medium.com [medium.com]
- 20. aimspress.com [aimspress.com]
- 21. researchgate.net [researchgate.net]
- 22. [2402.01868] Challenges in Training PINNs: A Loss Landscape Perspective [arxiv.org]
- To cite this document: BenchChem. [Addressing overfitting in physics-informed neural networks]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b15614678#addressing-overfitting-in-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](mailto:info@benchchem.com)