

Technical Support Center: 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](#)

This guide provides troubleshooting advice and answers to frequently asked questions regarding the critical task of balancing loss terms during the training of Physics-Informed Neural Networks (PINNs).

Frequently Asked Questions (FAQs)

Q1: What is loss balancing in PINNs and why is it crucial?

In PINNs, the total loss function is a composite objective, typically a weighted sum of several distinct terms:

- PDE Residual Loss (ngcontent-ng-c4139270029="" _nghost-ng-c4104608405="" class="inline ng-star-inserted">

$$L_f$$

): Measures how well the network's output satisfies the governing partial differential equation (PDE) at various points in the domain.

- Boundary Condition Loss (ngcontent-ng-c4139270029="" _nghost-ng-c4104608405="" class="inline ng-star-inserted">

$$L_b$$

): Enforces the known conditions at the boundaries of the domain.

- Initial Condition Loss (ngcontent-ng-c4139270029="" _nghost-ng-c4104608405="" class="inline ng-star-inserted">

$$L_{ic}$$

): Enforces the known conditions at the initial time step for time-dependent problems.

The total loss is often expressed as: ngcontent-ng-c4139270029="" _nghost-ng-c4104608405="" class="inline ng-star-inserted">

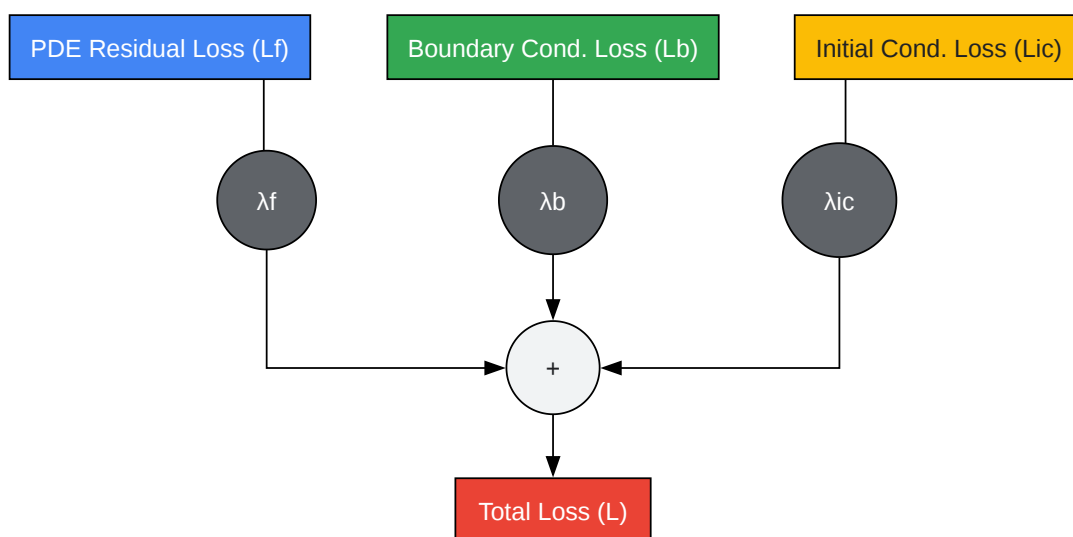
$$L = \lambda_f L_f + \lambda_b L_b + \lambda_{ic} L_{ic}$$

[1]

Loss balancing is the process of tuning the weights (ngcontent-ng-c4139270029="" _nghost-ng-c4104608405="" class="inline ng-star-inserted">

$$\lambda_f, \lambda_b, \lambda_{ic}$$

) to ensure that each loss term contributes appropriately to the total loss. This is crucial because these terms can have vastly different magnitudes, units, and gradient dynamics.[1][2][3] An imbalance can cause the training process to be dominated by the term with the largest gradient magnitude, leading the network to satisfy one physical constraint (e.g., the governing equation) while neglecting others (e.g., the boundary conditions), ultimately resulting in an inaccurate and non-physical solution.[2][4]



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

Figure 1: Structure of a typical weighted PINN loss function.

Q2: What are the common symptoms of imbalanced loss terms in my PINN training?

Identifying loss imbalance early can save significant computational resources. Key symptoms include:

- Stagnating Loss Components: One loss term (e.g.,

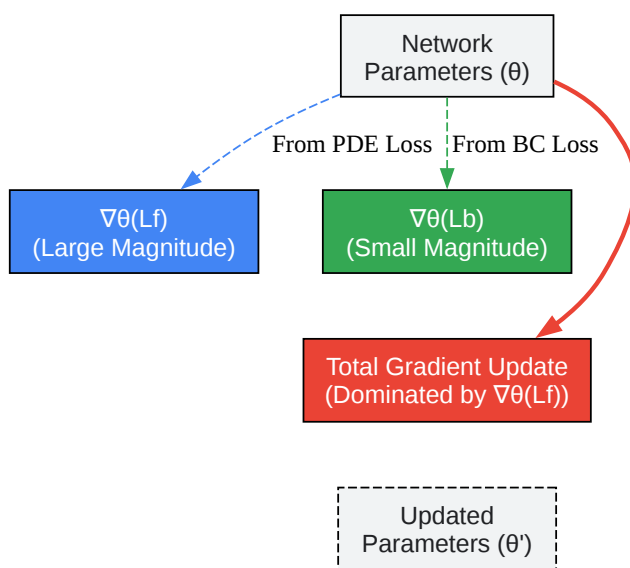
$$L_f L_f$$

) decreases steadily, while other terms (e.g.,

$$L_b L_b$$

) stagnate or even increase during training.[2]

- Violation of Physical Constraints: The trained model produces a solution that appears to satisfy the PDE in the interior of the domain but clearly violates the specified boundary or initial conditions.[2]
- Training Instability: The values of the loss components fluctuate wildly, indicating a lack of convergence.[5]
- Poor Overall Performance: Despite extensive training, the model fails to achieve an acceptable level of accuracy, often getting stuck in a local minimum.[6][7]



Illustrates how a large gradient from one loss term can dominate the update direction.

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

Figure 2: Conceptual diagram of imbalanced gradients.

Q3: What are the primary techniques for balancing loss terms in PINNs?

Loss balancing strategies can be broadly categorized into static and dynamic (or adaptive) methods.

- **Static Weighting:** This involves manually setting fixed scalar weights for each loss term before training begins.^[1] While simple, this approach is often suboptimal as it requires extensive, time-consuming trial-and-error and the ideal weights may change during the training process.^[6]
- **Adaptive Weighting:** These methods dynamically adjust the loss weights during training based on various heuristics, which is generally more effective.^[6] Several prominent techniques exist.

Q4: How do I choose the right adaptive loss balancing technique?

The choice of technique depends on the complexity of your problem and computational budget. Adaptive methods are strongly recommended over manual tuning.

| Technique | Core Principle | Key Characteristics |
|-------------------------------------|--|--|
| Learning Rate Annealing | Adjusts loss weights based on the magnitude of back-propagated gradients to balance their influence.[7][8] | A foundational adaptive method. Can be sensitive to the learning rate schedule.[6] |
| Gradient Normalization (GradNorm) | Dynamically tunes weights to ensure that the gradient norms from different loss terms remain at a similar scale.[4][9] | Aims for balanced training speeds across all objectives, improving stability.[4] |
| Relative Loss Balancing (ReLoBRaLo) | Aims for each loss term to decrease at a similar relative rate compared to its initial value at the start of training.[2][10] | Effective at ensuring all objectives make progress. Outperforms other methods in several benchmarks.[10] |
| Self-Adaptive PINNs (SA-PINNs) | Treats loss weights as trainable parameters that are updated via gradient ascent to maximize the loss, forcing the network to focus on harder-to-learn points.[11][12] | Uses a min-max optimization approach.[6] The weights act as a soft attention mask.[12] |
| Power-Based Normalization | Weights the loss terms to ensure their physical units are comparable, for instance, by ensuring all terms represent a form of power.[5] | A physics-based approach to ensure comparable magnitudes from the outset.[5] |

Troubleshooting Guides

Guide 1: Implementing a Self-Adaptive Weighting (SA-PINN) Protocol

The Self-Adaptive PINN (SA-PINN) is a powerful technique that treats the loss weights themselves as trainable parameters. The core idea is a min-max game: the network parameters are updated to minimize the loss, while the loss weights are simultaneously updated to maximize it.[6][11] This forces the model to pay more attention to the loss components that are currently largest.

Experimental Protocol:

- Define Learnable Weights: For each loss component (ngcontent-ng-c4139270029="" _ngghost-ng-c4104608405="" class="inline ng-star-inserted">

L_f, L_b, L_{ic} L_f, L_b, L_{ic}

), define a corresponding trainable scalar weight (ngcontent-ng-c4139270029="" _ngghost-ng-c4104608405="" class="inline ng-star-inserted">

$\lambda_f, \lambda_b, \lambda_{ic}$ $\lambda_f, \lambda_b, \lambda_{ic}$

). These are parameters, just like the network's weights and biases.

- Construct the Composite Loss: The total loss is the standard weighted sum: ngcontent-ng-c4139270029="" _ngghost-ng-c4104608405="" class="inline ng-star-inserted">

$L(\theta, \lambda) = \lambda_f L_f(\theta) + \lambda_b L_b(\theta) + \lambda_{ic} L_{ic}(\theta)$ $L(\theta, \lambda) = \lambda_f L_f(\theta) + \lambda_b L_b(\theta) + \lambda_{ic} L_{ic}(\theta)$

, where

θ

represents the network parameters.

- Establish Dual Optimizers: You will need two separate optimizers. A standard optimizer (e.g., Adam) for the network parameters

θ

, and another for the loss weights

λ

- Implement the Training Step: Within each training iteration, perform two distinct optimization steps:

- Minimize w.r.t. Network Parameters (

 θ

): Calculate the gradients of the total loss with respect to the network parameters

 θ

and apply a gradient descent step. This updates the network to better fit the physics and boundary conditions.

- $\theta_{k+1} = \theta_k - \eta_{\theta} \nabla_{\theta} L(\theta_k, \lambda_k)$

$$\theta_{k+1} = \theta_k - \eta_{\theta} \nabla_{\theta} L(\theta_k, \lambda_k)$$

[11]

- Maximize w.r.t. Loss Weights (

 λ

): Calculate the gradients of the total loss with respect to the loss weights

 λ

and apply a gradient ascent step. This increases the weight of the loss terms that are currently contributing the most error.

- $\lambda_{k+1} = \lambda_k + \eta_{\lambda} \nabla_{\lambda} L(\theta_k, \lambda_k)$

$$\lambda_{k+1} = \lambda_k + \eta_{\lambda} \nabla_{\lambda} L(\theta_k, \lambda_k)$$

[11]

- Monitor and Tune: Observe the evolution of both the loss components and the adaptive weights (

 λ

). The weights for more challenging objectives should increase, indicating the network is focusing its resources.

Figure 3: Workflow for a single training epoch in an SA-PINN.

Guide 2: My Model Still Fails With Loss Balancing. What's Next?

If implementing an adaptive weighting scheme doesn't solve your training issues, the problem may lie elsewhere in your PINN setup. Consider investigating the following:

- **Network Architecture:** PINNs are highly sensitive to architecture. Extremely deep networks can be prone to vanishing or exploding gradients, which is exacerbated by the need for higher-order derivatives.^[1] Recent studies suggest that shallow but wide networks often yield better performance.^[13]
- **Activation Functions:** The choice of activation function is more critical than in standard neural networks. The function must be differentiable to at least the order of the PDE. For a PDE with an n-th order derivative, the activation function must have at least n+1 non-zero derivatives.^[1] For example, while ReLU is popular, its derivatives quickly become zero, making it unsuitable for many PDEs. Smoother functions like tanh or swish are common choices.
- **Learning Rate Scheduling:** A fixed learning rate may not be optimal. Using a scheduler, such as ReduceLROnPlateau which reduces the learning rate when the loss plateaus, can significantly improve final performance.^[1]

- Input Normalization: As with most neural networks, normalizing the input coordinates to a standard range (e.g., [-1, 1]) is a crucial preprocessing step that can stabilize training.[1]
- Collocation Point Sampling: Uniformly sampling collocation points may not be efficient, especially for solutions with sharp gradients or multi-scale behavior. Consider adaptive sampling strategies that add more points in regions where the PDE residual is high.[8]

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. medium.com [medium.com]
- 3. emergentmind.com [emergentmind.com]
- 4. mdpi.com [mdpi.com]
- 5. Power-Based Normalization of Loss Terms to Improve the Performance of Physics-Informed Neural Networks (PINNs)[v1] | Preprints.org [preprints.org]
- 6. Dynamic Weight Strategy of Physics-Informed Neural Networks for the 2D Navier–Stokes Equations - PMC [pmc.ncbi.nlm.nih.gov]
- 7. Accuracy and Robustness of Weight-Balancing Methods for Training PINNs¹^{footnote 1}This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. [arxiv.org]
- 8. Self-adaptive weighting and sampling for physics-informed neural networks [arxiv.org]
- 9. researchgate.net [researchgate.net]
- 10. [2110.09813] Multi-Objective Loss Balancing for Physics-Informed Deep Learning [arxiv.org]
- 11. Review of Physics-Informed Neural Networks: Challenges in Loss Function Design and Geometric Integration | MDPI [mdpi.com]
- 12. repository.tudelft.nl [repository.tudelft.nl]
- 13. ijcai.org [ijcai.org]
- To cite this document: BenchChem. [Technical Support Center: Physics-Informed Neural Networks (PINNs)]. BenchChem, [2025]. [Online PDF]. Available at: [<https://www.benchchem.com/product/b15614678#techniques-for-balancing-loss-terms-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? [[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