Technical Support Center: Optimizing PINN Performance for Complex Geometries

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

Welcome to the technical support center for researchers, scientists, and drug development professionals utilizing Physics-Informed Neural Networks (PINNs). This resource provides troubleshooting guides and frequently asked questions (FAQs) to address specific issues you may encounter when applying PINNs to complex geometries.

Frequently Asked Questions (FAQs)

Q1: My PINN model is failing to converge or producing inaccurate results for a complex, non-rectangular domain. What are the likely causes and solutions?

A1: Failure to converge on complex geometries is a common challenge. The primary reasons often relate to how the network perceives the domain, how training points are sampled, and how the loss function is structured.

- Inadequate Geometric Representation: Standard Multilayer Perceptrons (MLPs) are defined on a Euclidean space and have no inherent knowledge of the domain's shape or topology.[1][2]

- Inefficient Sampling: Uniformly sampling collocation points is often inefficient for problems with complex solutions, such as those with steep gradients or multi-scale behaviors.[3][4]

- Loss Function Imbalance: The different components of the loss function (PDE residual, boundary conditions, initial conditions) may have vastly different magnitudes, leading to an imbalanced and difficult optimization landscape.[5][6]

Troubleshooting Steps:

- Enhance Geometric Input: Instead of feeding raw Cartesian coordinates, consider using techniques that encode the domain's geometry. One approach is to use a signed distance function (SDF) to represent the boundary.[7] Another advanced method is the Δ-PINN, which uses the eigenfunctions of the Laplace-Beltrami operator as a positional encoding.[1][8]

- Implement Adaptive Sampling: Move beyond uniform sampling to strategies that place more collocation points in regions where the PDE residual is high. This allows the network to focus on areas that are harder to learn.[9][10]

- Utilize Domain Decomposition: For very complex or large domains, break the problem down into smaller, simpler subdomains. A separate PINN can be trained on each subdomain, with continuity enforced at the interfaces.[11][12] This approach, used in Conservative PINNs (cPINNs) and Extended PINNs (XPINNs), can also aid in parallelization.[12]

Q2: How can I improve the enforcement of boundary conditions? My model satisfies the PDE in the interior but is inaccurate at the boundaries.

A2: This is a critical issue, as "soft" enforcement of boundary conditions via penalty terms in the loss function can be unreliable.[5] Here are several strategies to improve boundary condition satisfaction:

- Hard Constraint Enforcement: Modify the network's output formulation to satisfy the boundary conditions by construction. For example, using approximate distance functions (ADFs) and the theory of R-functions allows for the exact imposition of Dirichlet boundary conditions.[13]

- Boundary Connectivity Loss (BCXN): Introduce a novel loss term that provides a local structure approximation at the boundary. This helps the network better connect the interior solution to the boundary conditions, preventing overfitting in the near-boundary region, especially with sparse sampling.[14][15]

- Adaptive Loss Weighting: Employ a dynamic weighting scheme for the loss terms. This can automatically increase the importance of the boundary condition loss term during training if it is not being satisfied.[16][17]

Q3: What are the best practices for choosing a sampling strategy for collocation points?

A3: The choice of sampling strategy is pivotal for PINN performance. While uniform sampling methods can work for problems with smooth solutions, adaptive strategies are indispensable for more challenging systems.[3]

Comparison of Sampling Strategies:

Sampling Strategy Category	Specific Methods	Best Suited For	Key Advantage
Non-Adaptive Uniform Sampling	Grid, Random, Latin Hypercube, Sobol, Halton	Problems with smooth solutions and simple geometries.	Simple to implement.
Residual-Based Adaptive Sampling	RAD (Residual-based Adaptive Distribution), RAR (Residual-based Adaptive Refinement), FI-PINNs (Failure-Informed PINNs)	Problems with complex solutions (e.g., steep gradients, singularities).[3][4]	Dynamically concentrates points in high-error regions, significantly improving accuracy.[3]
Generative Adaptive Sampling	DAS-PINNs (Deep Adaptive Sampling)	High-dimensional problems and solutions with low regularity.[10]	Uses a generative model to generate new training points in regions of high residual.[10]
Sensitivity-Based Sampling	SBS (Sensitivity-Based Sampling)	Problems where the solution is highly sensitive to the location of training points.	Dynamically redistributes sampling probability to areas of high sensitivity.[18]

Experimental Protocol: Implementing Residual-based Adaptive Refinement (RAR)

- Initial Sampling: Begin with a set of uniformly distributed collocation points (e.g., using Latin Hypercube sampling).
- Initial Training: Train the PINN for a set number of epochs until the loss plateaus.
- Residual Evaluation: Generate a large set of candidate points within the domain and evaluate the PDE residual at each point using the current state of the network.
- Point Selection: Select a predefined number of new points from the candidate set that exhibit the highest residual values.
- Data Augmentation: Add these new, high-residual points to the existing set of collocation points.
- Retraining: Continue training the PINN with the augmented set of points.
- Iteration: Repeat steps 3-6 periodically throughout the training process.

Troubleshooting Guides

Guide 1: Diagnosing and Mitigating Stagnant Training and Vanishing Gradients

Problem: The training loss decreases initially but then stagnates at a high value, or the gradients become very small, hindering learning. This is common in deep networks or when dealing with stiff PDEs.[5][19]

Workflow for Diagnosis and Resolution:

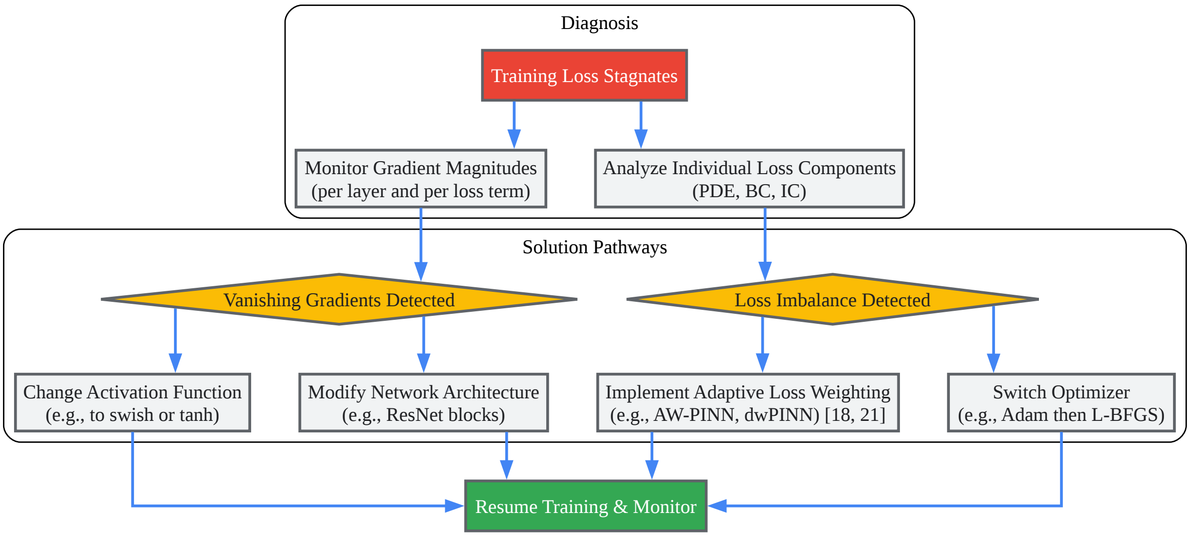
Diagnosis
Training Loss Stagnates
Monitor Gradient Magnitudes (per layer and per loss term)
Analyze Individual Loss Components (PDE, BC, IC)
Solution Pathways
Vanishing Gradients Detected
Loss Imbalance Detected
Change Activation Function (e.g., to swish or tanh)
Modify Network Architecture (e.g., ResNet blocks)
Implement Adaptive Loss Weighting (e.g., AW-PINN, dwPINN) [18, 21]
Switch Optimizer (e.g., Adam then L-BFGS)
Resume Training & Monitor

Click to download full resolution via product page

Caption: Workflow for troubleshooting stagnant PINN training.

Methodology:

- Monitor Gradients: During training, log the L2 norm of the gradients for the weights in each layer. If the gradients in the initial layers are consistently much smaller than in the later layers, you are likely experiencing vanishing gradients.
- Inspect Loss Components: Plot the individual values of the PDE residual loss, boundary condition loss, and initial condition loss over time. If one component is orders of magnitude larger than the others, it will dominate the gradient updates.[20]

- Change Activation Function: The choice of activation function can significantly impact gradient flow. While tanh is common, functions like swish can sometimes alleviate vanishing gradient issues.
- Implement Adaptive Weighting: Use an algorithm that dynamically adjusts the weights of each loss component. For example, an adaptive weighting scheme might update weights based on the magnitude of the backpropagated gradients to ensure all loss terms contribute to training.[16][17]
- Use a Hybrid Optimizer Strategy: Start training with a robust first-order optimizer like Adam to navigate the global loss landscape, then switch to a second-order method like L-BFGS for fine-tuning near a local minimum.[16]

Guide 2: Applying PINNs to Domains with Internal Discontinuities or Sharp Interfaces

Problem: The geometry is complex due to the presence of multiple materials or phases, leading to discontinuities in the solution or its derivatives across internal interfaces. A single PINN cannot effectively capture this behavior.

Logical Relationship for Domain Decomposition:

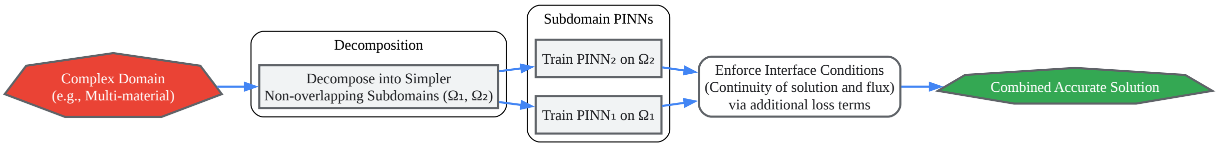
Complex Domain
(e.g., Multi-material)
Decomposition
Decompose into Simpler
Non-overlapping Subdomains (Ω_1, Ω_2)
Subdomain PINNs
Train PINN$_2$ on Ω_2
Train PINN$_1$ on Ω_1
Enforce Interface Conditions
(Continuity of solution and flux)
via additional loss terms
Combined Accurate Solution

Click to download full resolution via product page

Caption: Domain decomposition strategy for complex geometries.

Methodology: Implementing cPINN/XPINN

- Domain Partitioning: Decompose the complex domain Ω into a set of simpler, non-overlapping subdomains $\{\Omega_i\}$.[11]
- Network Allocation: Assign a separate neural network N_i to approximate the solution u_i within each subdomain Ω_i.
- Loss Function Formulation: The total loss function is a sum of the losses for each subdomain and the losses at the interfaces:
    - Subdomain Losses: For each network N_i, calculate the standard PINN loss (PDE residual and boundary conditions) using collocation points sampled only within Ω_i.
    - Interface Losses: For each interface between adjacent subdomains Ω_i and Ω_j, add loss terms to enforce the continuity of the solution ($u_i = u_j$) and the continuity of the normal flux. These are enforced by sampling points along the interfaces.[12]
- Training: Train all neural networks simultaneously by minimizing the total composite loss function. This allows the individual network solutions to be "stitched" together in a physically consistent manner.[12] This approach is particularly effective for multi-scale and multi-physics problems.[12]

References

- 1. openreview.net [openreview.net]
- 2. researchgate.net [researchgate.net]
- 3. emergentmind.com [emergentmind.com]
- 4. epubs.siam.org [epubs.siam.org]
- 5. mdpi.com [mdpi.com]
- 6. researchgate.net [researchgate.net]

- 7. Δ-PINNs: physics-informed neural networks on complex geometries [arxiv.org]
- 8. researchgate.net [researchgate.net]
- 9. pubs.aip.org [pubs.aip.org]
- 10. math.lsu.edu [math.lsu.edu]
- 11. Physics-informed neural networks - Wikipedia [en.wikipedia.org]
- 12. youtube.com [youtube.com]
- 13. Solving PDEs on Complex Geometries Using PINNs [dilbert.engr.ucdavis.edu]
- 14. openreview.net [openreview.net]
- 15. Fast-PINN for Complex Geometry: Solving PDEs with Boundary Connectivity Loss | OpenReview [openreview.net]
- 16. Physics-informed neural networks based on adaptive weighted loss functions for Hamilton-Jacobi equations [aimspress.com]
- 17. Dynamic Weight Strategy of Physics-Informed Neural Networks for the 2D Navier–Stokes Equations - PMC [pmc.ncbi.nlm.nih.gov]
- 18. skoge.folk.ntnu.no [skoge.folk.ntnu.no]
- 19. Exploring Physics-Informed Neural Networks: From Fundamentals to Applications in Complex Systems [arxiv.org]
- 20. researchgate.net [researchgate.net]
- To cite this document: BenchChem. [Technical Support Center: Optimizing PINN Performance for Complex Geometries]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b15614678#optimizing-pinn-performance-for-complex-geometries]

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]

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