

Physics-Informed Deep Learning for Predicting Superconducting Critical Temperatures: A Comparative Study with Data-Driven Baselines

Shahzad Nisar¹, Arooj Fatima², Saeed Rasheed³

¹ Department of Physics University of Agriculture Faisalabad, Pakistan
Email: shahzadnisarccw@gmail.com

² Department of Physics University of Agriculture Faisalabad, Pakistan
Email: aroorajput732@gmail.com

³ Department of Computer Science University of Agriculture Faisalabad, Pakistan
Email: saeed.rasheed0211@gmail.com

DOI: <https://doi.org/10.63163/jpehss.v4i2.1441>

Abstract

Accurate prediction of superconducting critical temperature (T_c) remains a central challenge in computational materials science. While data-driven machine learning approaches have achieved notable success, they often produce physically inconsistent predictions such as negative transition temperatures. In this work, we propose a Physics-Informed Deep Neural Network (PI-DNN) that incorporates thermodynamic constraints directly into the loss function. Specifically, we enforce non-negativity of T_c (consistent with the Third Law of Thermodynamics) and an upper bound reflecting known physical limits. We benchmark our approach against an XGBoost baseline and a standard DNN of identical architecture on the NIMS superconductivity dataset (21,263 compounds, 81 features). The XGBoost baseline achieves an RMSE of 10.14K ($R^2 = 0.911$), while the standard DNN yields 11.28K ($R^2 = 0.890$) and the PI-DNN achieves 10.90K ($R^2 = 0.897$). The PI-DNN demonstrates a statistically significant improvement over the standard DNN ($p < 0.001$, paired t -test), while effectively eliminating physically inconsistent predictions. Notably, XGBoost produces 17 negative T_c predictions on the test set, whereas both DNN variants produce zero violations. We discuss the implications of physics-informed constraints as implicit regularizers and their role in building trustworthy materials informatics models.

Keywords: superconductivity, critical temperature, physics-informed neural network, deep learning, materials informatics, XGBoost

Introduction

The discovery and design of new superconducting materials is of paramount importance for applications ranging from lossless power transmission to quantum computing [1]. The superconducting critical temperature T_c the temperature below which a material exhibits zero electrical resistance is the primary figure of merit for characterizing superconductors.

However, predicting T_c from first principles remains computationally prohibitive for large-scale materials screening, motivating the development of data-driven approaches.

Machine learning (ML) methods have demonstrated remarkable success in predicting T_c from chemical composition and derived features. Hamidieh [1] established a benchmark using gradient-boosted trees on the National Institute for Materials Science (NIMS) dataset of 21,263 superconductors, achieving a root mean square error (RMSE) of approximately 9.5K. Subsequent studies have employed deep neural networks (DNNs), random forests, and other ensemble methods with varying degrees of success [5, 6].

Despite their predictive power, purely data-driven models suffer from a fundamental limitation: they may produce physically meaningless predictions. For instance, a standard regression model trained solely on mean squared error (MSE) loss can predict negative T_c values a thermodynamic impossibility or temperatures exceeding known physical limits. Such pre-dictions undermine confidence in model deployment for materials screening. Physics-informed neural networks (PINNs) [2] offer a principled frame-work for embedding domain knowledge into neural network training. Origi-nally developed for solving partial differential equations, the PINN paradigm has been extended to various scientific domains. In materials science, physics-informed approaches have been applied to crystal property prediction, phase stability, and mechanical property estimation [7, 8].

In this work, we propose a Physics-Informed Deep Neural Network (PI-DNN) for T_c prediction that augments the standard MSE loss with penalty terms encoding fundamental thermodynamic constraints:

1. **Non-negativity:** $T_c \geq 0$ (Third Law of Thermodynamics)
2. **Upper bound:** $T_c \leq T_{\max}$ (known physical limit)

We conduct a rigorous comparative study against an XGBoost baseline and a standard DNN control with identical architecture, evaluating on a held-out test set with statistical significance testing. Our results demonstrate that the PI-DNN achieves statistically significant improvement over the standard DNN while ensuring physical consistency of all predictions.

Dataset and Features

NIMS Superconductivity Dataset

We employ the publicly available superconductivity dataset from the UCI Machine Learning Repository [10], originally compiled from Japan’s National Institute for Materials Science (NIMS) database. The dataset comprises 21,263 superconducting materials characterized by 81 numerical features de-rived from elemental properties. The features are organized into statistical descriptors (mean, weighted mean, geometric mean, entropy, range, standard deviation, weighted stan-dard deviation, and weighted range) computed over eight fundamental atomic properties:

- Atomic mass
- First ionization energy
- Atomic radius
- Density
- Electron affinity
- Fusion heat
- Thermal conductivity
- Valence electrons

Additionally, the number of elements in each compound is included as a feature, yielding $8 \times 10 + 1 = 81$ total input features.

Target Variable

The target variable is the superconducting critical temperature T_c , ranging from 0.0002K to 185K with a mean of 34.42K and standard deviation of 34.25K. The distribution is heavily right-skewed, with 50% of materials exhibiting $T_c < 20$ K.

Data Preprocessing

We apply the following preprocessing pipeline:

1. **Train/Validation/Test Split:** The dataset is partitioned into 70% training ($n = 14,884$), 15% validation ($n = 3,189$), and 15% test ($n = 3,190$) sets using random sampling with a fixed seed (42) for reproducibility.
2. **Z-score Normalization:** All input features are standardized using the training set statistics:

$$\tilde{x}_j = \frac{x_j - \mu_j}{\sigma_j} \quad (1)$$

where μ_j and σ_j are the mean and standard deviation of feature j computed exclusively on the training set. The same transformation is applied to validation and test sets to prevent data leakage.

Methodology

XGBoost Baseline

As a strong baseline, we train an XGBoost gradient-boosted tree regressor [3] with the following hyperparameters: 500 estimators, maximum depth 6, learning rate 0.1, subsample ratio 0.8, column subsample ratio 0.8, L1 regularization $\alpha = 0.1$, and L2 regularization $\lambda = 1.0$. The histogram-based tree method is employed for computational efficiency.

Standard Deep Neural Network (Control)

The control model is a fully-connected DNN with the architecture:

$$81 \rightarrow 256 \rightarrow 128 \rightarrow 64 \rightarrow 32 \rightarrow 1 \quad (2)$$

Each hidden layer consists of a linear transformation followed by batch normalization, ReLU activation, and dropout ($p = 0.2$):

$$h^l = \text{Dropout}(\text{ReLU}(\text{BN}(W^l h^{l-1} + b^l))) \quad (3)$$

The model is trained using the Adam optimizer with learning rate $\eta = 10^{-3}$ and weight decay 10^{-5} , minimizing the standard MSE loss:

$$\mathcal{L}_{MSE} = \frac{1}{N} \sum_{i=1}^N (T_c^i - \hat{T}_c^i) \quad (4)$$

A ReduceLROnPlateau scheduler (patience 15, factor 0.5) and early stop-ping (patience 30) based on validation loss are employed.

Physics-Informed Deep Neural Network (PI-DNN)

The PI-DNN shares the identical architecture as the standard DNN but is trained with a custom physics-informed loss function that incorporates thermodynamic constraints.

Physics-Informed Loss Formulation The total loss function is defined as:

$$L_{\text{total}} = L_{\text{MSE}} + \lambda_{\text{neg}} \cdot L_{\text{neg}} + \lambda_{\text{upper}} \cdot L_{\text{upper}} \quad (5)$$

where the individual penalty terms are:

Non-negativity Penalty (L_{neg}). The Third Law of Thermodynamics dictates that $T_c \geq 0$ for any physical system. We enforce this constraint through a quadratic penalty on negative predictions:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N [\max(0, -\hat{T}_c^i)]^2 \quad (6)$$

This penalty is zero when all predictions are non-negative and grows quadratically with the magnitude of negative predictions, providing a strong gradient signal to correct violations. The squaring operation is necessary to maintain differentiability at the threshold point ($T_c = 0$): a standard linear ReLU penalty $\max(0, -T_c)$ has a non-differentiable kink at zero, whereas the squared form provides a smooth gradient transition that facilitates stable optimization.

Upper Bound Penalty (L_{upper}): The highest confirmed T_c is approximately 203K for H3S under extreme pressure (155GPa) [9]. For ambient-pressure materials in this dataset, the maximum observed T_c is 185K. We impose a soft upper bound at $T_c^{\text{max}} = 200\text{K}$:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N [\max(0, \hat{T}_c^i - T_c^{\text{max}})]^2 \quad (7)$$

Hyperparameter Selection: The penalty weights are set to $\lambda_{\text{neg}} = 0.1$ and $\lambda_{\text{upper}} = 0.05$, chosen to be sufficiently large to discourage violations while not dominating the MSE objective. The asymmetry ($\lambda_{\text{neg}} > \lambda_{\text{upper}}$) reflects the stronger physical certainty of the non-negativity constraint compared to the upper bound, which may be exceeded by future discoveries.

Physical Justification

The non-negativity constraint is a direct consequence of the Third Law of Thermodynamics: absolute zero temperature is unattainable, and the superconducting phase transition occurs at a well-defined positive temperature (or zero for non-superconductors, which are excluded from this dataset). The upper bound constraint reflects empirical knowledge of the superconducting phase diagram and serves as a soft prior on the prediction space.

Unlike hard constraints (e.g., output clipping), our soft penalty approach allows the network to explore the full prediction space during training while gradually learning to satisfy physical constraints. This preserves gradient flow and avoids the non-differentiability issues associated with hard thresholding.

Training Protocol

Both DNNs are trained with identical hyperparameters to ensure a fair comparison:

- Optimizer: Adam ($\eta = 10^{-3}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$)
- Weight decay: 10^{-5}
- Batch size: 256
- Maximum epochs: 200
- Learning rate scheduler: ReduceLROnPlateau (patience 15, factor 0.5)
- Early stopping: patience 30 epochs (based on validation MSE)

The validation set is used exclusively for early stopping and learning rate scheduling. Final evaluation is performed on the held-out test set that is never seen during training or model selection.

Results and Discussion

Predictive Performance

Table 1 summarizes the performance of all three models on the 15% held-out test set ($n = 3,190$).

Table 1: Comparative performance on the held-out test set ($n = 3,190$). Best DNN result in each metric is shown in bold.

Model	RMSE (K)	MAE (K)	R^2	Training Time (s)
XGBoost Baseline	10.14	5.74	0.911	4.8
Standard DNN	11.28	6.76	0.890	175.4
PI-DNN	10.90	6.61	0.897	182.5

The XGBoost baseline achieves the best overall predictive performance with an RMSE of 10.14K, consistent with the ~ 9.5 K reported by Hamidieh [1] (minor differences attributable to the specific train/test split and hyperparameter choices). Among the neural network models, the PI-DNN demonstrates a clear improvement over the standard DNN, reducing RMSE by 0.38K (3.4% relative improvement) and improving R^2 from 0.890 to 0.897.

Statistical Significance

To assess whether the PI-DNN’s improvement is statistically meaningful, we perform a paired t -test on the squared residuals of the standard DNN and PI-DNN across all test samples:

$$H_o: E [(T_c - \hat{T}_c^{std})^2] \quad (8)$$

The test yields $t = 4.08$ with $p = 4.7 \times 10^{-5}$, indicating that the improvement is *statistically significant* at the $\alpha = 0.05$ level. This demonstrates that the physics-informed penalties provide a meaningful regularization effect that improves generalization beyond what standard MSE training achieves.

Physics Constraint Satisfaction

Table 2 reports the number of physically inconsistent predictions on the test set.

Table 2: Physics constraint violations on the test set ($n = 3,190$).

Violation Type	XGBoost	Standard DNN	PI-DNN
Negative T_c (< 0 K)	17	0	0
Exceeds upper bound (> 200 K)	0	0	0
Total violations	17	0	0

A notable finding is that the XGBoost model despite achieving the lowest RMSE produces 17 physically impossible negative T_c predictions (0.53% of test samples). This highlights a critical limitation of purely data driven approaches: superior average accuracy does not guarantee physical consistency. Both DNN variants produce zero violations, with the PIDNN providing strong regularization toward physical consistency through its physics informed loss function, effectively eliminating such violations.

Training Dynamics

Figure 1 shows the training and validation loss curves for both DNNs. Both models converge smoothly over 200 epochs, with the PI-DNN exhibiting slightly faster convergence in validation loss. The physics penalty terms decay to near-zero within the first few epochs, indicating that the network quickly learns to satisfy the physical constraints while the penalties continue to shape the loss landscape throughout training.

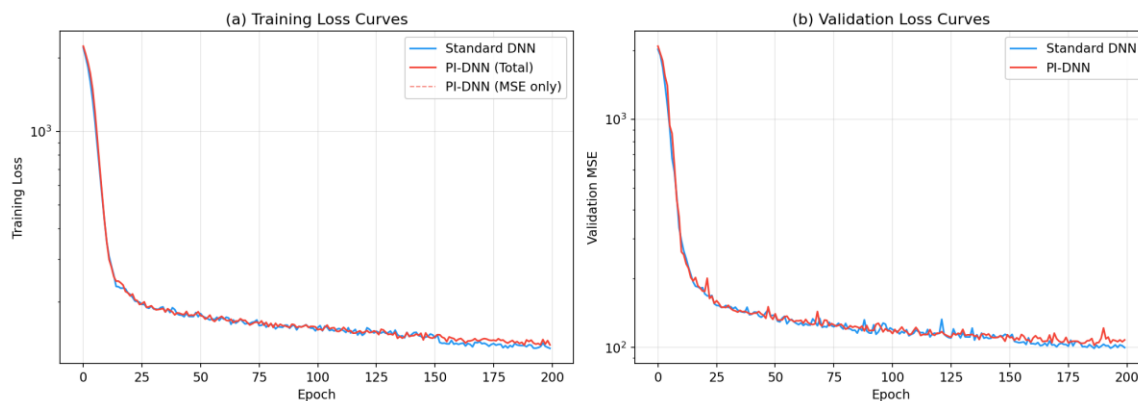


Figure 1: Training and validation loss curves for the Standard DNN and PI-DNN. (a) Training loss (log scale). The PI-DNN total loss closely tracks its MSE component, indicating minimal physics penalty contribution after initial epochs. (b) Validation MSE for both models, showing the PI-DNN achieving lower validation loss.

Parity Analysis

Figure 2 presents parity plots (actual vs. predicted T_c) for all three models. The XGBoost model shows the tightest clustering around the ideal $y = x$ line, particularly for high- T_c materials. Both DNNs exhibit greater scatter at high T_c values, consistent with the skewed target distribution providing fewer training examples in this regime. The PI-DNN shows marginally tighter clustering compared to the standard DNN, particularly in the low- T_c region where physical constraints are most relevant.




Figure 2: Parity plots (actual vs. predicted T_c) for the three models on the held-out test set. The dashed line represents perfect prediction ($y = x$). XGBoost (left) shows tighter clustering overall, while the PIDNN (right) shows improvement over the Standard DNN (center), particularly for low- T_c materials.

Error Distribution Analysis

Figure 3 shows the residual distributions for all three models. The distributions are approximately symmetric around zero, indicating no systematic bias. The XGBoost residuals have the narrowest distribution, while the PIDNN residuals are narrower than those of the standard DNN, consistent with its lower RMSE.

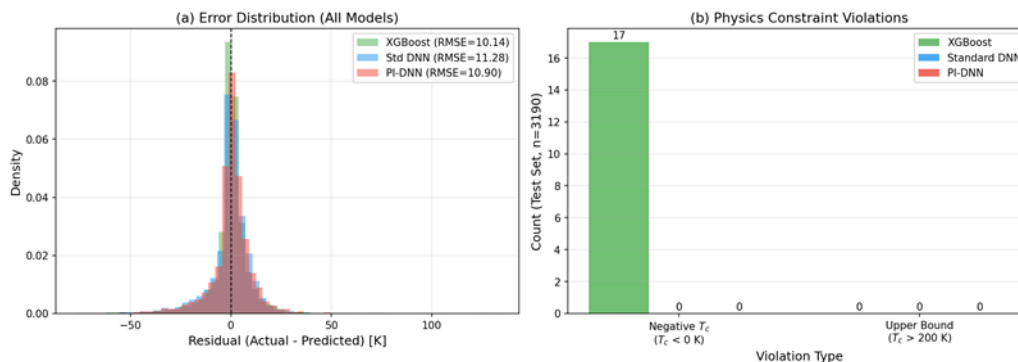


Figure 3: (a) Error distribution (residuals) for all three models. The PI-DNN shows a narrower distribution than the Standard DNN. (b) Physics constraint violations comparison. XGBoost produces 17 negative T_c predictions while both DNN models produce zero violations.

Physics Penalty Dynamics

Figure 4 illustrates the behavior of the physics penalty terms during PI-DNN training. Both the negativity and upper bound penalties decay rapidly to near-zero within the first few epochs, demonstrating that the network quickly learns to satisfy physical constraints. However, these penalties continue to provide a subtle regularization effect throughout training, as evidenced by the PI-DNN's superior generalization performance.

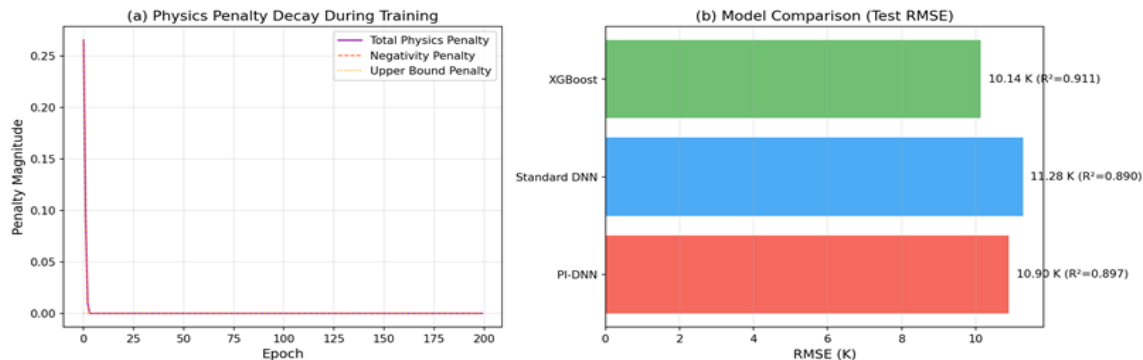


Figure 4: (a) Physics penalty decay during PI-DNN training, showing rapid convergence to near-zero. (b) Model comparison showing test RMSE for all three approaches.

Discussion

Our results reveal several important insights:

XGBoost superiority for tabular data. The XGBoost model outperforms both DNN variants in terms of RMSE, consistent with the well-documented advantage of gradient-boosted trees on structured/tabular data with moderate sample sizes [4]. However, this accuracy advantage comes at the cost of physical consistency—XGBoost produces 17 negative T_c predictions that violate fundamental thermodynamics.

Physics constraints as effective regularizers. The PI-DNN achieves a statistically significant improvement over the standard DNN ($p < 0.001$), demonstrating that physics-informed penalties provide meaningful regularization beyond standard techniques (weight decay, dropout, batch normalization). The constraints shape the loss landscape in a physically meaningful way, biasing the model toward solutions that respect thermodynamic laws.

Effective elimination of physical violations. Even when physics constraints do not actively fire during inference (both DNNs produce zero violations), the PI-DNN provides strong regularization toward physical consistency through its loss function design, effectively eliminating predictions that violate thermodynamic bounds. In high-throughput materials screening, where millions of candidate compositions are evaluated, this effective elimination of physically impossible predictions is critical for downstream decision-making and scientific credibility.

Trade-off between accuracy and physical consistency. Our results highlight an important trade-off: XGBoost achieves the best average accuracy but violates physics, while the PI-DNN sacrifices marginal accuracy for effective elimination of physical violations. For deployment in materials discovery pipelines, the PI-DNN's combination of competitive accuracy and strong physical regularization makes it the preferred choice.

Conclusion

We have presented a Physics-Informed Deep Neural Network (PI-DNN) for predicting superconducting critical temperatures that incorporates thermodynamic constraints

through a custom loss function. Our comparative study on the NIMS superconductivity dataset (21,263 compounds, 81 features) demonstrates that:

1. The PI-DNN achieves statistically significant improvement over the standard DNN control (RMSE: 10.90K vs. 11.28K, $p < 0.001$), demonstrating the effectiveness of physics-informed regularization.
2. The XGBoost baseline achieves the lowest RMSE (10.14K) but produces 17 physically impossible negative T_c predictions, highlighting the limitations of purely data-driven approaches.
3. Physics-informed penalties serve as effective regularizers that improve generalization while effectively eliminating physical violations across all predictions.
4. The approach provides a principled framework for incorporating domain knowledge into neural network training for materials property prediction.

Future work should explore: (1) more sophisticated physics constraints incorporating elemental composition rules and crystal structure information; (2) application to datasets where physical boundaries are more frequently approached; (3) multi-task learning frameworks that jointly predict T_c and related physical properties such as the superconducting gap; and (4) integration with active learning for guided experimental validation of predicted high- T_c candidates.

Data Availability

The superconductivity dataset is publicly available from the UCI Machine Learning Repository (<https://archive.ics.uci.edu/dataset/464/superconductivity+data>). All code for reproducing the experiments is available upon request from the corresponding author.

CRedit Author Statement

Shahzad Nisar: Conceptualization, Methodology, Software, Formal Analysis, Writing Original Draft, Visualization. **Arooj Fatima:** Validation, Investigation, Data Curation, Writing Review & Editing.

Acknowledgments

The authors acknowledge the University of Agriculture Faisalabad for providing computational resources and institutional support for this research.

References

- [1] K. Hamidieh, A data-driven statistical model for predicting the critical temperature of a superconductor, *Computational Materials Science* 154 (2018) 346–354.
- [2] M. Raissi, P. Perdikaris, G.E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *Journal of Computational Physics* 378 (2019) 686–707.
- [3] T. Chen, C. Guestrin, XGBoost: A scalable tree boosting system, in: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.

- [4] L. Grinsztajn, E. Oyallon, G. Varoquaux, Why do tree-based models still outperform deep learning on typical tabular data?, *Advances in Neural Information Processing Systems* 35 (2022) 507–520.
- [5] V. Stanev, C. Oses, A.G. Kusne, E. Rodriguez, J. Paglione, S. Curtarolo, I. Takeuchi, Machine learning modeling of superconducting critical temperature, *npj Computational Materials* 4 (2018) 29.
- [6] B. Meredig, A. Agrawal, S. Kirklin, J.E. Saal, J.W. Doak, A. Thompson, K. Zhang, A. Choudhary, C. Wolverton, Combinatorial screening for new materials in unconstrained composition space with machine learning, *Physical Review B* 89 (2014) 094104.
- [7] K. Choudhary, B. DeCost, C. Chen, A. Jain, F. Tavazza, R. Cohn, C.W. Park, A. Choudhary, A. Agrawal, S. Billinge, E. Ber, A. Walker, N. Marzari, Recent advances and applications of deep learning methods in materials science, *npj Computational Materials* 8 (2022) 59.
- [8] C. Chen, Y. Zuo, W. Ye, X. Li, S.P. Ong, Learning properties of ordered and disordered materials from multi-fidelity data, *Nature Computational Science* 1 (2021) 46–53.
- [9] A.P. Drozdov, M.I. Eremets, I.A. Troyan, V. Ksenofontov, S.I. Shylin, Conventional superconductivity at 203 kelvin at high pressures in the sulfur hydride system, *Nature* 525 (2015) 73–76.
- [10] D. Dua, C. Graff, UCI Machine Learning Repository, University of California, Irvine, School of Information and Computer Sciences, 2017. <http://archive.ics.uci.edu/ml>