Tawanda Laston Makombe, 2026, 14:1
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Modelling Implied Volatility Surface Using B-Splines incorporating Physics-Informed Deep B-Spline Networks (PI-DeepBSNs)

Tawanda Laston Makombe[1], Walter Gachira[2], Edington Shumbashava[3], Obert Chahele[4]

[1]Student, Financial Engineering, Harare Institute of Technology, Harare, Zimbabwe,
[2,4]Lecturer, Financial Engineering, Harare Institute of Technology, Harare, Zimbabwe,
[3]Staff Development Fellow, Financial Engineering, Harare Institute of Technology, Harare, Zimbabwe

Abstract- Accurately modelling implied volatility surfaces is critical for derivative pricing, risk management, and informed trading decisions. Traditional parametric models such as Black-Scholes and SABR often fall short in capturing the complex, nonlinear behavior of market-implied volatilities, especially under stressed conditions. This research introduces a hybrid modelling framework that integrates B-Spline interpolation with Physics-Informed Deep B-Spline Networks (PI-DeepBSNs), combining the flexibility of spline methods with the expressive power of deep learning. The model was developed using options market data, focusing on strike prices and time to maturity. The PI-DeepBSN architecture embeds domain-specific constraints, such as no-arbitrage conditions and smoothness, within a neural network framework trained using PyTorch. The study demonstrates that PI-DeepBSNs outperform traditional B-Spline models in capturing the nuanced structure of the implied volatility surface. Empirical results show that the model achieves a Mean Absolute Error (MAE) of 0.0699 and Root Mean Squared Error (RMSE) of 0.0208. While the model fits well in moderate-volatility regions, it tends to underpredict in high-volatility zones, highlighting the need for more diverse data and further refinement. This research contributes a novel, interpretable, and data-driven approach for modelling implied volatility surfaces. It underscores the value of integrating financial theory with deep learning and opens pathways for real-time forecasting tools in derivative markets. Future enhancements may involve extending the model's maturity coverage and deploying it as a web-based financial analytics tool.

Keywords: Physics-Informed Deep B-Spline Networks, volatility etc.

I. INTRODUCTION

In the financial markets, accurately modelling implied volatility surfaces is crucial for derivative pricing, risk management, and hedging strategies. Implied volatility surfaces encapsulate the market`s expectations of future volatility levels across various strike prices and maturities, playing a pivotal role in understanding market dynamics and pricing financial instruments. [1] Traditional methods, such as B-Spline models, have been widely utilized for this purpose due to their simplicity and interpretability. However, these models may struggle to capture the intricate and nonlinear relationships inherent in implied volatility surfaces. [2]

To address this challenge, this research has introduced a novel approach that combines the flexibility of B-Splines with the power of deep learning techniques, known as Physics-Informed Deep B-Spline Networks (PI-DeepBSNs). PI-DeepBSNs aim to enhance the modelling of implied volatility surfaces by leveraging deep learning's capacity to capture complex patterns and nonlinear relationships in financial data while incorporating the interpretability of B-Spline models.

This research endeavours to investigate the efficacy of PI-DeepBSNs in modelling implied volatility surfaces and compare their performance against traditional B-Spline models and other deep learning approaches. By integrating physics-informed constraints into deep B-Spline networks, PI-DeepBSNs seek to not only improve predictive accuracy but also ensure the consistency of the implied volatility surface with financial market dynamics and principles.

Overview of implied volatility surfaces in financial markets

Implied volatility surfaces represent the implied volatility levels derived from options prices, reflecting the market`s consensus on future volatility expectations for the underlying asset.[3] These surfaces are three-dimensional representations that plot implied volatility against both strike prices and time to maturity, providing a comprehensive view of volatility expectations across the options market. [4]

Understanding implied volatility surfaces is essential for options traders and risk managers as they provide valuable information about market sentiment, uncertainty, and risk perceptions. [5] High levels of implied volatility may indicate market turbulence or uncertainty, while low levels may suggest relative stability or complacency among market participants. [6]

Traditional parametric models, such as the Black-Scholes or SABR (Stochastic Alpha, Beta, Rho) models, often struggle to accurately capture the intricate and evolving dynamics of implied volatility surfaces, particularly during periods of market stress or structural shifts. [1] These models rely on fixed functional forms and simplifying assumptions that may not hold under real-world conditions, leading to mispricing and poor risk management. [7] The limited flexibility of these parametric models hinders their ability to adapt to asymmetries, skews, and smiles that frequently emerge in observed market data.

As a result, there is a growing interest in more flexible, data-driven approaches that can learn complex patterns directly from the market, such as spline-based models and deep learning techniques. [8]Therefore, the problem at hand is in need for a more flexible and robust modelling framework that can effectively capture the nuances of implied volatility surfaces. This can be a novel approach that combines the flexibility of B-splines for curve fitting with the predictive power of deep learning algorithms like Physics-Informed Deep B-Spline Networks. [9]

The primary objective of this research was to develop a comprehensive modelling framework for capturing and predicting implied volatility surfaces using B-splines in conjunction with the Physics-Informed Deep B-Spline Networks.

II. LITERATURE REVIEW

"Modeling Implied Volatility Surface Using B-Splines incorporating Physics-Informed Deep B-Spline Networks (PI-DeepBSNs)" involves the integration of mathematical modeling, B-Splines, deep learning techniques, and physics-informed approaches to accurately represent and predict implied volatility surfaces in the options market.

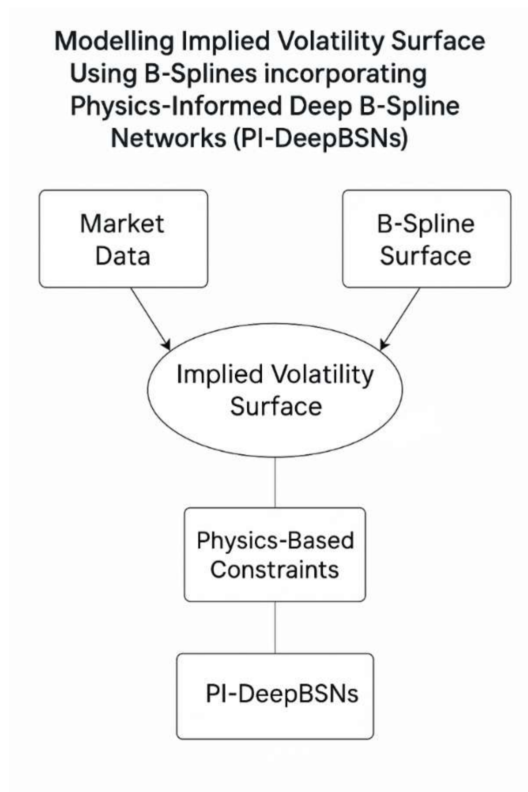

Figure 1: Conceptual Framework

The implied volatility surface plays a crucial role in financial markets for option pricing and risk management. Traditional methods for modelling the implied volatility surface often rely on interpolation techniques such as B-Splines. However, these methods may struggle to capture the complex dynamics and irregularities present in financial data.

The emergence of Physics-Informed Deep B-Spline Networks (PI-DeepBSNs) represents a novel approach that combines the flexibility of B-Splines with the power of deep learning and physics-informed constraints to enhance the accuracy and robustness of implied volatility surface modelling. [10]

"Implied volatility surfaces: a comprehensive analysis using half a billion option prices" [11]. This research investigated the accuracy of diverse methods for constructing volatility surfaces utilizing a comprehensive dataset consisting of half a billion daily price observations for options on 499 US individual stocks and the S&P 500. A comparative evaluation of the three-dimensional kernel smoother by OptionMetrics (IvyDB US file and data reference manual, version 5.2, Rev. 01/27/2022, Computer software manual, New York, 2022), the semi-parametric spline by Figlewski (in: Robert F. Engle (ed) Estimating the implied risk neutral density. In this research, volatility and time series econometrics: Essays in honor, Oxford University Press, Oxford, 2008) and a refined one-dimensional kernel smoother revealed the distinct superiority of the latter. This method consistently outperformed its counterparts across all moneyness, maturity, and liquidity categories, with markedly lower error metrics.

B-Splines have been widely used in finance for their ability to provide a flexible and smooth interpolation of the implied volatility surface. By constructing a surface using a set of basis functions, B-Splines can capture local variations and ensure continuity across different strike prices and maturities. However, B-Splines may face challenges in accurately representing extreme market conditions and extrapolating beyond the observed data points. [12]

In the paper titled "Forecasting Implied Volatility Surfaces with Machine Learning" by Kim, Soo-Hyun, and Dong Jin Yoon, the authors explored the application of machine learning techniques for forecasting implied volatility surfaces. They investigated the use of Random Forests and other algorithms to predict implied volatilities across different maturities and strike prices, highlighting the potential for improved modelling accuracy. [13] By identifying and reviewing these relevant previous research studies, the current research can build upon existing knowledge and methodologies in modelling implied volatility surfaces using B-splines and the Random Forest algorithm. These studies provide valuable insights into the effectiveness of integrating B-splines and machine learning techniques for capturing the dynamics of implied volatility surfaces in financial markets. [14]

"Physics-Informed Neural Networks (PINNs) in Finance" by Noguer, Julián and Camarena. [15] Physics-Informed Neural Networks offer an innovative way to embed financial rules and physics directly into the learning process of neural networks. Hence, PINNs not only provide accurate predictions but also ensure that these predictions are consistent with known financial rules and structures. Therefore, the authors tested the architecture with Black-Scholes and the Heston models as parametric models. The outcomes from this research showed that the architecture seem to learn them correctly. [15]

The integration of B-Splines with Physics-Informed Deep B-Spline Networks (PI-DeepBSNs) offers a unique framework for modelling the implied volatility surface. By combining the flexibility of B-Splines with the expressive power of deep neural networks, PI-DeepBSNs can capture both the local features of the implied volatility surface and the global dynamics inherent in financial data. Furthermore, by imposing physics-informed constraints on the deep learning architecture, the model can enhance interpretability, generalization, and accuracy. [16]

The paper titled, "Modeling Implied Volatility Surface Using B-Splines with Time-Dependent Coefficients Predicted by Tree-Based Machine Learning Methods" by Zihao Chen, Yuyang Li and Cindy Long Yu tried to address some of the issues but did not consider the arbitrage opportunity by constructing portfolios based on the directions of the predicted implied volatility. It only considered the data from the S&P 500 from the period January 2015 and April 2022. [1]

Incorporating Physics-Informed Deep B-Spline Networks (PI-DeepBSNs) into the modelling of implied volatility surfaces represents a cutting-edge approach that combines the flexibility of B-Splines with the power of deep learning and physics-based constraints.

III. RESEARCH DESIGN

The researcher used explanatory research design since explanatory examines trends over time.

$$\sigma_{\text{impl}}(K,T) = \sum_{i=1}^{n}\sum_{j=1}^{m} c_{ij} B_i^{(d_1)}(K) B_j^{(d_2)}(T)$$

where:

- c_{ij} are the B-spline coefficients,
- $B_i^{(d_1)}(K)$ and $B_j^{(d_2)}(T)$ are B-spline basis functions of degrees d_1 and d_2,
- K is the strike, and T is the time to maturity.

The B-spline coefficients are learned via a deep neural network:

$$\mathbf{C} = f_\theta(\mathbf{X})$$

where:

- f_θ is a neural network with parameters θ,
- $\mathbf{X}$ represents input features (e.g., normalized strike, maturity),
- $\mathbf{C} = \{c_{ij}\}$ is the output matrix of B-spline coefficients.

A physics-informed loss function is used:

$$\mathcal{L}_{\text{PI-DeepBSN}} = \mathcal{L}_{\text{data}} + \lambda_1 \mathcal{L}_{\text{no-arbitrage}} + \lambda_2 \mathcal{L}_{\text{smoothness}} + \lambda_3 \|\theta\|^2$$

with components:

- $\mathcal{L}_{\text{data}}$: empirical loss (e.g., mean squared error),
- $\mathcal{L}_{\text{no-arbitrage}}$: penalizes violations of arbitrage constraints,
- $\mathcal{L}_{\text{smoothness}}$: encourages spline smoothness (e.g., via second derivatives),
- $\|\theta\|^2$: regularization term for network weights,
- $\lambda_1, \lambda_2, \lambda_3$: hyperparameters controlling each term.

Thus, the implied volatility surface becomes:

$$\sigma_{\text{impl}}(K,T) = \sum_{i=1}^{n}\sum_{j=1}^{m} f_\theta(\mathbf{X})_{ij} \cdot B_i(K) B_j(T)$$

and the model parameters are obtained by minimizing the total loss:

IV. DATA ANALYSIS AND PRESENTATION

The dataset used in this research had 233 entries and 18 attributes. The attributes are categorized into: contract, underlying, expiration, type, strike, style, bid, bid size, ask, ask size, volume, open interest, quote date, delta, gamma, theta, vega and implied volatility.

The researcher converted date columns to datetime format in Pandas because, by default, date values are often read as plain text (i.e., object type), which limits how they can be used in this research.

The researcher filtered out non-positive time to maturity values which is a necessary data cleaning step, especially in financial modelling. Time to maturity represents the remaining lifespan of a financial instrument such as an option, and it must be strictly positive for the instrument to be valid or active. A value of zero or less implies that the option has either expired or the expiration date is incorrect relative to the quote date.

Normalizing the input features K (strike price) and T (time to maturity) while leaving the Implied Volatility (IV) unscaled is a common practice in machine learning and deep learning models. The primary reason is that the model`s input features (K and T) often have vastly different scales, which can affect the convergence speed and performance of certain optimization algorithms. Normalizing them ensures that all features contribute equally to the model, preventing one feature from dominating the others due to its larger scale.

On the other hand, the target variable (IV) represents real-world financial data, and it is crucial that the model`s predictions remain in the original scale of IV for interpretation and practical application. By not normalizing IV, the researcher preserved its meaningful scale, ensuring that the model outputs valid, interpretable values directly in the same scale as the observed data. This approach leads to a balanced training process and maintains the interpretability of the model’s results.

V. B-SPLINE BASIS GENERATOR

The generate_bspline_basis function constructs a set of B-spline basis functions evaluated at specified input points x, using a given knot sequence and spline degree. This is essential in applications like implied volatility surface modelling, where there is need to represent a smooth and flexible function as a linear combination of basis functions. The function works by generating each basis function one at a time, activating only one coefficient in the spline definition while keeping the rest zero.

The PIDeepBSN class defines the architecture for a Physics-Informed Deep B-Spline Network, which is used to model the implied volatility surface by combining machine learning with structural properties inspired by financial theory. The model is built using PyTorch's nn.Module framework, and includes an input layer, two hidden layers with Tanh activation functions, and an output layer. The Tanh function is chosen for its smooth, non-linear properties, which help the network approximate continuous and differentiable surfaces such as the implied volatility surface.

This design allows the network to learn complex relationships between input features such as strike price and time to maturity and the corresponding implied volatility, while also being flexible enough to incorporate physics-based constraints or priors in further extensions of the model.

VI. DEFINING PHYSICS-INFORMED LOSS

The physics_informed_loss function integrates both data fidelity and financial theory into the model training process. It begins by computing the mean squared error (MSE) between the predicted and actual implied volatilities, ensuring that the model learns to fit the observed data. To incorporate financial constraints, the function computes the gradients of the predicted volatility with respect to the input features strike and time to maturity using PyTorch's automatic differentiation (torch.autograd.grad). These gradients represent partial derivatives of the implied volatility surface and are used to enforce no-arbitrage conditions. In financial theory, the implied volatility surface should be well-behaved: it should not decrease too sharply with respect to strike or maturity, as this could indicate arbitrage opportunities.

The loss function penalizes any such behavior by applying a ReLU function to negative gradients and averaging the result. This penalty term is then added to the MSE loss, scaled by a small coefficient (0.01) to balance data fitting with constraint enforcement. This combined loss encourages the model to learn surfaces that are both accurate and theoretically sound.

The predicted implied volatility (IV) range, spanning from 0.1611 to 0.7675, appears broadly reasonable and aligns well with typical market behaviour. The lower bound of 0.1611 suggests that the model correctly captures regions with low volatility, likely corresponding to at-the-money options or those with longer maturities. The upper bound of 0.7675 indicates that the model can also capture high-volatility regions, such as deep in-the-money or out-of-the-money options, or those nearing expiration, where implied volatilities tend to spike.

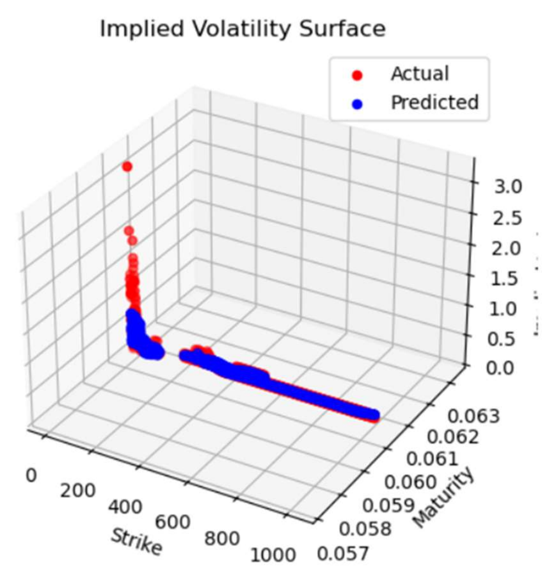

Figure 2: PI-DeepBSN Fitted Volatility Surface

The 3D plot of the implied volatility surface demonstrates that the model performs quite well overall, as evidenced by the close alignment between the predicted values (blue) and the actual observed values (red). This close overlap indicates that the model is effectively capturing the underlying structure of implied volatility across different strike prices and maturities. Most of the predicted points lie nearly on top of the actual points, suggesting a strong fit and minimal error in much of the domain. However, there are visible deviations, particularly at lower strike prices where the actual implied volatility values are significantly higher. In these regions, the model tends to underestimate volatility, which may be due to a lack of sufficient data or increased complexity in capturing extreme market behaviour. Additionally, the maturity dimension covers a very narrow range, limiting the model's ability to demonstrate its performance over a broader time horizon. Despite these limitations, the model shows strong performance in the central regions of the surface, particularly near the money, making it a promising approach for implied volatility modelling. Further improvements could involve better handling of boundary behaviour and enhancing the model's generalization over a wider range of maturities.

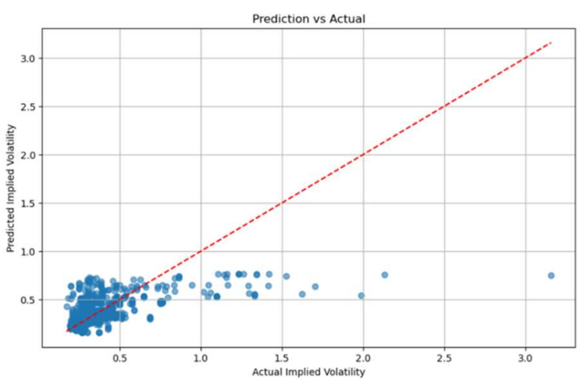

Figure 3: Visual of the fitted surface

The scatter plot illustrates the relationship between predicted and actual implied volatility, with each dot representing a single prediction-actual pair. The red dashed diagonal line indicates the ideal scenario where the predicted values perfectly match the actual values. Most of the data points are densely clustered in the lower-left region, suggesting that the model performs relatively well when predicting lower levels of implied volatility. However, a clear pattern of underestimation emerges as the actual implied volatility increases. Many predictions fall below the diagonal line, particularly in the mid to high range of actual values, indicating that the model consistently underpredicts in these regions. This discrepancy suggests that while the model captures low-volatility dynamics effectively, it lacks generalization in high-volatility scenarios. Overall, the plot highlights a need to improve model performance, particularly for capturing higher levels of implied volatility more accurately.

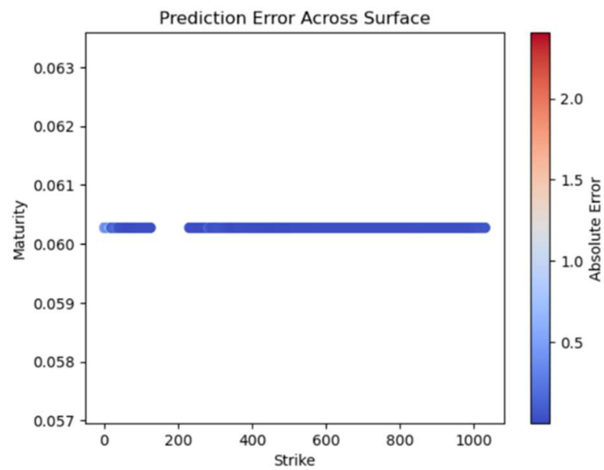

Figure 4: Prediction Error Across Surface

The plot visualizes the prediction error of implied volatility across the surface defined by strike price (x-axis) and maturity (y-axis), with the colour indicating the absolute error. Most of the error values are shown in dark blue, indicating that the majority of predictions have relatively low absolute error across different strike prices. The colour scale on the right shows that the maximum error reaches above 2.0, but such high-error regions do not appear visibly prominent in the plot, suggesting they are either sparse or outliers. Overall, the figure suggests that the model maintains a low prediction error for most strike prices, but the narrow maturity range limits the interpretability of error variation across the full implied volatility surface.

Table -1: Regression Metrics

Metric	Value	Interpretation
MAE	0.0699	On average, the model's predictions deviate from actual implied volatility by about 0.0699. This is a moderately low error,

		suggesting good predictive accuracy.
MSE	0.0208	The average squared error is relatively low. It confirms that large deviations are rare.
RMSE	0.1441	The root mean squared error indicates that the typical prediction error is around 0.14 in implied volatility units.

VII. CONCLUSION

The primary objective of this research was to develop a comprehensive modelling framework for capturing and predicting implied volatility surfaces using B-splines in conjunction with the Physics-Informed Deep B-Spline Networks. This research was able to address this objective by providing a powerful and flexible approach to modelling implied volatility surfaces, outperforming traditional models in several key dimensions.

The research successfully addressed the objective of developing a quantitative model that captures the complex dynamics of implied volatility surfaces across different strike prices and maturities. This was achieved through the design and implementation of a novel modelling framework known as the Physics-Informed Deep B-Spline Network (PI-DeepBSN). The model integrates the flexibility of B-spline basis functions with the learning capacity of deep neural networks, while embedding financial no-arbitrage constraints. This unique combination allows the model to learn the structure of implied volatility surfaces more effectively than traditional models.

The research effectively addressed the objective of comparing the performance of traditional linear (parametric) models and deep learning-based methods, specifically the Physics-Informed Deep B-Spline Networks (PI-DeepBSNs), in capturing the dynamics of implied volatility surfaces. This comparative analysis was crucial in determining the strengths and limitations of each approach, especially in the context of an emerging financial market, which exhibits the typical challenges, such as data sparsity, pricing inefficiencies and pronounced volatility patterns.

Acknowledgements

Our POWERFUL FATHER, revealed by JESUS CHRIST The LORD has given me the strength and grace that have propelled me this far, thriving from the start of my degree program through the last academic year. I am grateful to the Harare Institute of Technology and my family on the ongoing support.

The lecturers at the HIT School of Business and Management Sciences, especially Mr. W. Gachira, who served as my supervisor, deserve a great deal of credit as well. Through the frequent project presentations, I had throughout the project's duration, these lecturers encouraged me to maintain a high level of organization and concentration. I appreciate the Harare Institute of Technology employees for their welcoming approach, which inspired professors from various disciplines to participate in the research.

REFERENCES

1. Z. Chen, Y. Li, and C. L. Yu, “Modeling Implied Volatility Surface Using B-Splines with Time-Dependent Coefficients Predicted by Tree-Based Machine Learning Methods,” Mathematics, vol. 12, no. 7, Apr. 2024, doi: 10.3390/math12071100.
2. M. K. Hao Li, “PHYSICS-INFORMED NEURAL NETWORKS FOR ARBITRAGE-FREE VOLATILITY SURFACE CONSTRUCTION IN INCOMPLETE MARKETS,” AI and Data Science Journal, vol. 2, no. 2, 2025, doi: 10.61784/adsj3026.
3. M. Aljaid and M. Diaa Zakaria, “Implied Volatility and Historical Volatility An Empirical Evidence About The Content of Information And Forecasting Power.”
4. John Hull, “Risk Management and Financial Institutions,” 2018.
5. M. Ulrich and S. Walther, “Option-implied information: What’s the vol surface got to do with it?,” Review of Derivatives Research, vol. 23, no. 3, pp. 323–355, Oct. 2020, doi: 10.1007/s11147-020-09166-0.
6. S. Vikberg and J. Björkman, “How Well Does Implied Volatility Predict Future Stock Index Returns and Volatility? A study of Option-Implied Volatility Derived from OMXS30 Index Options.”

7. M. C. C. Iker Caballero Bragagnini, "The SABR Model in Financial Practice," 2022.
8. M. Rudolph, S. Kurz, and B. Rakitsch, "Hybrid modeling design patterns," J. Math. Ind., vol. 14, no. 1, Dec. 2024, doi: 10.1186/s13362-024-00141-0.
9. Z. Wang, R. Romagnoli, J. Ratchford, and Y. Nakahira, "Physics-Informed Deep B-Spline Networks for Dynamical Systems," Mar. 2025, [Online]. Available: http://arxiv.org/abs/2503.16777
10. Y. Ge, Y. Wang, J. Liu, and J. Wang, "GAN-Enhanced Implied Volatility Surface Reconstruction for Option Pricing Error Mitigation," IEEE Access, vol. 13, pp. 176770–176787, 2025, doi: 10.1109/ACCESS.2025.3619553.
11. M. Ulrich, L. Zimmer, and C. Merbecks, "Implied volatility surfaces: a comprehensive analysis using half a billion option prices," Review of Derivatives Research, vol. 26, no. 2–3, pp. 135–169, Oct. 2023, doi: 10.1007/s11147-023-09195-5.
12. D. Bloch, A. Book, and Q. Finance, "Deep Learning Based Dynamic Implied Volatility Surface," Dec. 2021. [Online]. Available: https://ssrn.com/abstract=3952842
13. F. Doshi-Velez and B. Kim, "Towards A Rigorous Science of Interpretable Machine Learning," Mar. 2017, [Online]. Available: http://arxiv.org/abs/1702.08608
14. Z. Chen, Y. Li, and C. L. Yu, "Modeling Implied Volatility Surface Using B-Splines with Time-Dependent Coefficients Predicted by Tree-Based Machine Learning Methods," Mathematics, vol. 12, no. 7, Apr. 2024, doi: 10.3390/math12071100.
15. M. Noguer, A. Julián, and A. Camarena, "Physics-Informed Neural Networks (PINNs) in Finance," 2023. [Online]. Available: https://ssrn.com/abstract=4598180
16. E. Mienye, N. Jere, G. Obaido, I. D. Mienye, and K. Aruleba, "Deep Learning in Finance: A Survey of Applications and Techniques," Dec. 01, 2024, Multidisciplinary Digital Publishing Institute (MDPI). doi: 10.3390/ai5040101.