

Replicating the Options Pricing Surface:

A Refined Deep Learning Approach

Will Duan¹, Prof. Ali Hirsra

¹Department of Industrial Engineering and Operations Research, Columbia University sd3578@columbia.edu



Abstract

Neural networks are known to provide fast and reliable pricing for various assets and derivatives such as options. We show that with careful tuning of the networks, including additional outputs and Physics-Informed Neural Network (PINN) inspired fine-tuning, we are able to both obtain accurate Greeks and improve significantly in out-of-sample tests. This results in the ability to create deep learning networks that more closely approximates the pricing surface, increasing reliability, robustness, and interpretability.

Introduction

Option pricing plays a fundamental role in quantitative finance, as it underpins valuation, hedging, and trading strategies across global markets. Pricing based on well-understood models such as Black-Merton-Scholes and Heston can be overly time-consuming or inefficient with established methods such as by analytic solutions or by Monte-Carlo methods. Instead, machine learning techniques, in particular artificial neural networks, have increasingly been explored as complementary tools for option pricing [1], including supervised [2] and unsupervised [3] approaches.

In trading and risk management, the derivatives of option prices, known as the Greeks, are often equally critical compared to the price because they guide hedging decisions and risk control. We contribute to this direction by leveraging neural networks to accurately replicate the options pricing surface and simultaneously producing accurate estimates of both option prices and their Greeks in an accurate and robust manner.

Data Generation

European call-option prices and Greeks were generated under six established models of the underlying dynamics. For each model, quasi-Monte Carlo sampling produced 2M rows from the main parameter domain, split 60/20/20 into training, validation, and test sets. Three additional test-only datasets (200K rows each) evaluated extrapolation to deep out-of-the-money options, long maturities, and their combination. Option prices were computed using FFT methods [4], with Greeks obtained through PyTorch automatic differentiation.

Table 1: pricing models and their characteristics

Pricing Model	Parameters	Features
Black-Merton-Scholes [5]	5	Widely used, stable, and relatively basic
Heston [6]	9	Introduces stochastic volatility
Variance Gamma (VG) [7]	7	Pure-jump process, though unstable and difficult to compute for certain Greeks
Variance Gamma with Stochastic Arrival [8]	10	VG with stochastic volatility, large number of parameters
CGMY [9]	8	VG with additional activity/variation parameter
Merton Jump Diffusion [10]	8	Contains both jump and diffusion processes

Methodology

In this supervised learning problem, we evaluate the neural networks we train based on accuracy (measured by R^2) of call price, and 8 representative Greeks:


Table 2: Greeks and Definitions

First-Order Greeks		Second-Order Greeks	
Name	Definition	Name	Definition
Delta (Δ)	$\partial V / \partial (S_0 / K)$	Gamma (Γ)	$\partial^2 V / \partial (S_0 / K)^2$
Vega (v)	$\partial V / \partial \sigma$	Vanna	$\partial^2 V / \partial (S_0 / K) \partial \sigma$
Theta (Θ)	$\partial V / \partial T$	Vomma	$\partial^2 V / \partial \sigma^2$
Rho (ρ)	$\partial V / \partial r$	Charm	$\partial^2 V / \partial (S_0 / K) \partial T$

where V is the value/price of the call option.

To produce accurate Greeks, we use two approaches. First, we attempt a multiple output network, in which not only the price is produced, but we also include additional outputs for Greeks. Two structures were tested, one with all 8 Greeks included (AllGreeks) and another with just Delta and Vega (2Greeks), so all Greeks can be computed by taking the gradient only and not the Hessian.

Figure 1: Structure of Network (Black-Merton-Scholes Parameters)



The other approach is to fine-tune a network by modifying the loss function to include loss of the partial derivatives as well. We introduce the first-order custom loss function for fine-tune training

$$Loss_1 = MSE(Call) + \sum_i \lambda_i MSE \left(\frac{\partial Call}{\partial p_i} \right)$$

and the second-order variant


$$Loss_2 = MSE(Call) + \sum_i \lambda_i MSE \left(\frac{\partial Call}{\partial p_i} \right) + \sum_{i,j} \lambda_{i,j} MSE \left(\frac{\partial^2 Call}{\partial p_i \partial p_j} \right)$$

To fine-tune the first-order loss (FT1), we start with the weights of a baseline network with one output, and ramp up the loss weighting parameters λ from 0.05 to 0.25 through 10 epochs, and train until convergence. The second-order loss (FT2) tuning starts with the weights from the first-order fine-tuned model, and ramp up $\lambda_{i,j}$ from 0.01 to 0.05 through 10 epochs. The first-order fine-tuning approach is also applied on the 2Greeks network (FT2Greeks), with the goal of obtaining more accurate first partial derivatives of Delta and Vega, which covers the second-order Greeks. Multi-Layer Perceptron (MLP) networks of various widths and depths were trained in order to find the best performing architecture.

Preliminary Results


We present the call price, and 4 representative first-order Greeks and second-order Greeks, as well as Merton and VG for the underlying for comparison.

Figure 2: Selected Data for Merton Jump Diffusion



We can see that all of the methodologies offer a significant improvement over the baseline in second-order Greeks for Merton. First-order fine-tuning (FT1), however, stays consistently accurate throughout different out-of-sample test sets.

Figure 3: Selected Data for Variance Gamma*



Due to a noisy dataset inherent with the characteristics of VG, especially for second-order Greeks, FT1 again offers robustness throughout different test sets, while networks with second-order Greeks outputs and those fine-tuning for second-order Greeks suffer in performance.

*: Negative values clipped at -0.1

Conclusion/Future Work

We have found that fine-tuning for first-order Greeks created the most significant improvements in the accuracy of both first and second order Greeks compared to a baseline model, while maintaining call price accuracy and out-of-sample robustness. The result holds for multiple pricing models as well, suggesting also robustness with respect to different underlying dynamics. Overall, replicating the option pricing surface with fine-tuned networks provides both speed and much-improved reliability in options pricing, hedging, and risk management. Further research could be done in several directions, including using real market data in place of synthetic data, and also on different types of options/derivatives such as American, Asian, Barrier, etc.

Citations

1. J. Raf and W. Wang (2020), "Neural Networks for Option Pricing and Hedging: A Literature Review," J. Comput. Finance 24(1), 1–45. doi:10.21314/JCF.2020.390.
2. A. Hirsra, T. Karatas, and A. Okou (2019), "Supervised Deep Neural Networks (DNNs) for Pricing/Calibration of Vanilla/Exotic Options Under Various Different Processes," arXiv:1902.05810. doi:10.48550/arXiv.1902.05810.
3. W. Fu and A. Hirsra (2022), "An Unsupervised Deep Learning Approach in Solving Partial Integer Differential Equations," Quantitative Finance 22(8), 1481–1494. doi:10.1080/14697688.2022.2057870.
4. A. Hirsra (2024), Computational Methods in Finance, 2nd ed., Chapman & Hall/CRC. doi:10.1201/9780429094743.
5. F. Black and M. Scholes (1973), "The Pricing of Options and Corporate Liabilities," J. Political Economy 81(3), 637–654. doi:10.1086/260062; R. C. Merton (1973), "Theory of Rational Option Pricing," Bell J. Econ. Manag. Sci. 4(1), 141–183. doi:10.2307/3003143.
6. S. L. Heston (1993), "A Closed-Form Solution for Options with Stochastic Volatility," Rev. Financial Stud. 6(2), 327–343. doi:10.1093/rfs/6.2.327.
7. D. B. Madan, P. P. Carr, and E. C. Chang (1998), "The Variance Gamma Process and Option Pricing," European Finance Rev. 2(1), 79–105. doi:10.1023/A:100703831535.
8. P. Carr, H. Geman, D. B. Madan, and M. Yor (2003), "Stochastic Volatility for Levy Processes," Math. Finance 13(3), 345–382. doi:10.1111/1467-9965.00020.
9. P. Carr, H. Geman, D. B. Madan, and M. Yor (2002), "The Fine Structure of Asset Returns," J. Business 75(3), 305–332. doi:10.1086/338705.
10. R. C. Merton (1976), "Option Pricing When Underlying Stock Returns Are Discontinuous," J. Financial Econ. 3(1–2), 125–144. doi:10.1016/0304-4052(76)90022-2.