

End-to-End Neural Calibration of Stochastic Volatility Models via Option Panels

Tian'ai Song

IBSS, Xi'an Jiaotong-Liverpool University, Suzhou, China, 215123

Tianai.Song23@student.xjtlu.edu.cn

Abstract. We propose an end-to-end neural architecture that calibrates stochastic volatility models directly from high-dimensional option panels. A 3-D CNN–Transformer hybrid learns the inverse mapping from implied-volatility surfaces to latent Heston parameters in milliseconds while enforcing no-arbitrage constraints via a differentiable penalty layer. Closed-form gradients of the Fourier pricing formula enable stable training on synthetic and market data. Extensive experiments on SPX and EURO STOXX 50 options show sub-0.3 % RMSE and a 100× speed-up over L-BFGS-B, with calibration times under 2ms on GPU.

Keywords: Stochastic volatility calibration; option panels; deep learning; Heston model; no-arbitrage constraints; Fourier pricing.

1. Introduction & Problem Formulation

Calibration of stochastic volatility (SV) models on modern trading desks must ingest tens of thousands of implied-volatility quotes per underlying each day. Traditional pipelines invert a pricing map $P: \theta \mapsto C(K, T)$ by iterative numerical minimization, a procedure whose complexity grows at least linearly with the number of contracts and whose convergence is sensitive to starting values, noise and missing quotes. The resulting latencies—often hundreds of milliseconds on a GPU cluster—are incompatible with real-time quoting, intraday hedging or market-making algorithms that require recalibration every few seconds.

Addressing these bottlenecks starts with a precise mathematical statement. At any time t , a panel of European option prices is observed, summarized by the Black-Scholes implied volatility surface $\hat{\sigma}_t(K, T)$. The forward model postulates that these quotes are noisy observations of $C(K, T; \theta)$, where $\theta \in \mathbb{R}^d$ collects the latent SV parameters (mean-reversion, vol-of-vol, correlation, etc.). The calibration task is therefore a non-linear inverse problem: recover θ given $\hat{\sigma}_t(K, T)$ for $(K, T) \in \mathcal{K} \times \mathcal{T}$. Ill-posedness arises from (i) the non-convexity of P , (ii) the high dimensionality of $\hat{\sigma}_t$, and (iii) the need to respect no-arbitrage constraints on total variance.

The objective is to learn a map $F: \hat{\sigma}_t \mapsto \theta_t$ that executes in a single forward pass on GPU hardware, delivering calibrated parameters within one to two milliseconds while guaranteeing that the generated surface is free of butterfly arbitrage and calendar spread arbitrage. The solution must generalize across underliers, maturities and market regimes without hand-crafted starting points or gradient-based refinement.

2. Background & Literature Survey

2.1. Classical stochastic volatility models (Heston, Bates, SABR) and their characteristic functions

Classical stochastic volatility models share a common mathematical skeleton: an instantaneous variance process v_t follows a mean-reverting diffusion, while the underlying price S_t obeys a geometric Brownian motion scaled by $\sqrt{v_t}$. The Heston model delivers a closed-form characteristic function $\Phi_H(u) = E[e^{iu \ln S_T}]$, enabling Fourier-inversion pricing in micro-seconds. Bates augments

Heston with simultaneous jumps in price and variance, yielding $\Phi_B(u)$ that nests Φ_H as a special case. SABR abandons mean-reversion and instead parameterizes the forward price dynamics under a displaced log-normal diffusion [1]; its leading-order implied-volatility formula is algebraic and therefore ubiquitous in interest-rate desks. Across these specifications, the mapping from latent parameters θ to option prices factors through $\Phi(u)$, a property that underpins both classical calibration and modern gradient-based learning.

2.2. Traditional calibration pipelines: Fourier-inversion pricing + gradient-based or global optimizers; limitations in speed and overfitting

Traditional calibration treats the inverse problem as an optimization over θ . A typical pipeline first pre-computes the model price $C(K, T; \theta)$ via fast Fourier transform, then minimizes a weighted squared error between model and market implied volatilities. Gradient-based routines such as Levenberg-Marquardt or L-BFGS-B exploit closed-form Greeks $\partial C/\partial\theta$ obtained by differentiating the pricing integral under the characteristic function. Global optimizers—differential evolution, CMA-ES, particle swarm—mitigate local minima but scale quadratically with panel size. Both families suffer from three limitations: (i) convergence times of 10–500ms per underlying, (ii) sensitivity to starting values and noise, and (iii) systematic overfitting when the number of quotes exceeds the intrinsic dimensionality of θ , a regime common in liquid equity and index options [2].

2.3. Recent deep-learning approaches for option pricing and hedging; gap in jointly learning the inverse map with guaranteed no-arbitrage

Deep learning has recently shifted attention from direct optimization to function approximation. Convolutional networks map option panels to prices or hedging ratios, while generative adversarial networks synthesize arbitrage-free volatility surfaces. Nevertheless, existing architectures either learn the forward map (parameters \rightarrow prices) or the hedging map (prices \rightarrow Greeks) [3], but rarely the inverse map (prices \rightarrow parameters). Moreover, enforcing no-arbitrage during training remains heuristic: soft penalties or post-processing filters provide no differentiable guarantee that the resulting surface satisfies the necessary conditions on total variance. A principled end-to-end approach that jointly learns the inverse map and embeds hard no-arbitrage constraints is still missing.

3. Mathematical Framework

3.1. Heston Dynamics and Parameter Vector

The latent state follows a square-root variance process coupled to geometric Brownian motion

$$dS_t = rS_t dt + \sqrt{v_t} S_t dW_t^S, \quad (1)$$

$$dv_t = \kappa(\theta - v_t) dt + \xi \sqrt{v_t} dW_t^v, \quad (2)$$

$$dW_t^S dW_t^v = \rho dt, \quad (3)$$

With parameter

$$\theta = (\kappa, \theta, \xi, \rho, v_0) \in \mathcal{R}^5. \quad (4)$$

The Feller condition $2\kappa\theta > \xi^2$ guarantees strict positivity of v_t .

3.2. Fourier Representation of European Call Prices

Let $\Phi(k; \theta) := E^Q[e^{ik \ln S_T}]$ denote the risk-neutral characteristic function with log-strike frequency k . The undiscounted call price equals

$$C(K, T) = e^{-rT} E^Q[(S_T - K)^+] = \frac{e^{-rT}}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{e^{-ik \ln K} \Phi(k; \theta)}{k^2 - ik} \right] dk. \quad (5)$$

The integrand decays at least as fast as k^{-2} , ensuring absolute integrability [4].

3.3. Closed-Form Greeks via Differentiation Under the Integral Sign

Define the integrand

$$I(k; \theta) := \operatorname{Re} \left[\frac{e^{-ik \ln K} \Phi(k; \theta)}{k^2 - ik} \right]. \quad (6)$$

Applying Leibniz rule and Fubini's theorem (justified by exponential decay of Φ), the sensitivities become

$$\frac{\partial C}{\partial \theta} = \frac{e^{-rT}}{\pi} \int_0^\infty d\theta \frac{\partial I(k; \theta)}{\partial \theta} dk. \quad (7)$$

Each component $\partial \Phi / \partial \theta$ is available analytically because $\Phi(k; \theta)$ solves a Riccati ODE whose coefficients are affine in θ . Substituting the Riccati derivatives yields, for example,

$$\frac{\partial C}{\partial \kappa} = \frac{e^{-rT}}{\pi} \int_0^\infty dk \operatorname{Re} \left[\frac{e^{-ik \ln K} \Phi(k; \theta) \Psi_\kappa(k)}{k^2 - ik} \right] dk \quad (8)$$

And analogous expressions for θ, ζ, ρ, v_0 , where $\Psi_\kappa(k)$ is an elementary rational function of k and θ . These formulas supply exact gradients for stochastic optimization and neural-network back-propagation without finite-difference noise.

3.4. No-Arbitrage Constraints in Total-Variance Coordinates

Introduce the total-variance surface

$$w(K, T) := T \widehat{\sigma}^2(K, T) \quad (9)$$

Absence of butterfly and calendar-spread arbitrage imposes

$$w_{KK}(K, T) \geq 0 \quad (10)$$

$$w_T(K, T) \geq 0 \quad (11)$$

Together with the large-strike limit $\lim_{K \rightarrow \infty} w(K, T) = 0$ and monotonicity in strike: $\partial w / \partial K \leq 0$ for $K \geq S_0$. These conditions are enforced during training via a differentiable soft-plus penalty appended to the loss.

4. Unified Pipeline: Architecture, Data, and Evaluation

4.1. 3-D CNN encoder → sinusoidal-positional Transformer → parameter head yielding θ

The implied-volatility tensor $X_t \in \mathbb{R}^{N_K \times N_T \times 1}$ first passes through a four-block 3-D CNN whose kernels ($3 \times 3 \times 1$, stride 1) progressively halve spatial resolution while doubling channels $8 \rightarrow 16 \rightarrow 32 \rightarrow 64$. Batch-norm and ELU activations are interleaved with residual skip-connections to retain high-frequency strike-maturity structure. The resulting feature map $\Phi \in \mathbb{R}^{H \times W \times 64}$ ($H=N_K/8$, $W=N_T/8$) is flattened to a sequence of length HW and enriched with 2-D sinusoidal positional encodings encoding (K, T) coordinates. A single Transformer encoder layer ($d_{\text{model}} = 128$, 8 heads, feed-forward 512) attends across spatial tokens, producing latent vectors $z \in \mathbb{R}^{HW \times 128}$. Mean-pooling z yields a 128-D context vector that feeds two fully-connected layers ($128 \rightarrow 64 \rightarrow 5$) with tanh activations to produce raw parameters[5], which are then element-wise scaled to admissible ranges $\kappa \in [0.1, 10]$, $\theta \in [0.01, 0.5]$, $\xi \in [0.05, 2]$, $\rho \in [-1, 0]$, $v_0 \in [0.01, 0.5]$.

4.2. Differentiable no-arbitrage projection with soft penalty and annealed λ

After box-clipping θ , the constraint layer reconstructs the total-variance surface $w(K, T; \theta)$ via the closed-form Heston implied-volatility map. On a dense grid (200 strikes \times 50 maturities) the soft penalties

$$P_1 = \text{ReLU}(-w_{\text{KK}}), P_2 = \text{ReLU}(-w_{\text{T}}), P_3 = \text{ReLU}(\partial w / \partial K) \text{ for } K > S_0$$

Are aggregated into $L_{\text{penalty}} = \Sigma(P_1 + P_2 + P_3)$. The weight λ is annealed exponentially: $\lambda_t = 1000 \cdot 0.96^{\{t-1\}}$, ensuring strong enforcement at early epochs and fine-tuning later.

4.3. Training loss

$$L = \sum_t |C(K, T; \hat{\theta}_t) - C_{\text{market}}(K, T)|^2 + \lambda_t L_{\text{penalty}}. \quad (12)$$

AdamW ($lr=10^{-3}$, $weightdecay=10^{-4}$) processes mini-batches of 256 panels for 100 epochs with early stopping on validation RMSE[6].

4.4. Data generation

One million implied-volatility panels are produced by a GPU Quasi-Monte-Carlo Heston simulator using $2^{\{16\}}$ Sobol paths per panel (weekly steps). Parameters are drawn uniformly: $\kappa \in [0.5, 8]$, $\theta \in [0.02, 0.5]$, $\xi \in [0.1, 1.5]$, $\rho \in [-0.95, -0.05]$, $v_0 \in [0.01, 0.4]$. Each panel contains 10 strikes (80 %–120 % moneyness) and 8 maturities (0.1–2.0 years). Market data comprise SPX and EURO STOXX 50 implied-volatility snapshots from 2018-01-02 to 2023-12-29 at 5-minute frequency (delta 5–95, maturity ≥ 7 days). Augmentation adds $\pm U [0, 15]$ bp bid-ask spreads, Gaussian microstructure noise $\varepsilon \sim \mathcal{N}(0, 0.5 \text{ bp})$, and 5 % random masking.

4.5. Metrics

RMSE and MAPE are computed on the full σ grid; χ^2 distance to the Breeden-Litzenberger density is evaluated on 1% strike bins; calibration latency is the median of 100 forward passes on an NVIDIA A100 GPU.

4.6. Baselines and ablations

Comparisons include L-BFGS-B with analytical gradients, CMA-ES (pop=14), and a shallow $512 \rightarrow 256$ ReLU network. Ablation grids test 1-block vs 4-block CNN, presence/absence of positional encoding, and $\lambda \in \{0, 1, 10, 100, 1000\}$.

5. Visualization Suite

5.1. Figure 1 – Synthetic vs. Predicted Implied-Volatility Surfaces




Figure 1. Synthetic vs Predicted IV Surface

A 3-D surface plot displays a representative synthetic panel (black mesh) against the network’s reconstruction (red wire-frame). The x-axis lists relative strikes 0.8–1.2, the y-axis maturities 0.1–2.0 years, and the z-axis implied volatility in decimals. Overlapping meshes confirm sub-0.1 % RMSE visually; a translucent green heat-map on the base plane highlights absolute point-wise residuals for quick inspection of systematic biases near the wings [7].

5.2. Figure 2 – Parameter Recovery Scatter




Figure 2. Parameter Recovery Scatter

Five sub-plots (κ , θ , ξ , ρ , v_0) share identical axes: ground-truth on the horizontal, predicted on the vertical. Each dot corresponds to one held-out panel ($n = 10\,000$). A diagonal identity line (dashed) and 5% relative-error bands (shaded) are overlaid. Pearson correlation coefficients printed in the upper-left corner exceed 0.98 for every parameter, while color encodes implied-volatility level to reveal any regime dependence[8].

5.3. Figure 3 – Calibration Error Heat-Map




Figure 3. Calibration Error Heat-Map

A 2-D heat-map (colorbar in basis-points) plots residual $\hat{\sigma}_{\text{model}} - \hat{\sigma}_{\text{market}}$ across the full strike–maturity grid. The x-axis is standardized moneyness, the y-axis time-to-expiry in years. Dark-blue regions (< 2 bp) dominate the center; red pockets (> 10 bp) appear only at extreme 10-delta wings, confirming negligible mis-pricing in liquid regions [9].

5.4. Figure 4 – Training Convergence




Figure 4. Training Convergence

Two curves (solid vs dotted) track validation RMSE against epochs for (i) the proposed architecture (4-block CNN + Transformer + constraint layer) and (ii) a vanilla 2-layer CNN baseline. The shaded corridor marks ± 1 standard deviation over three seeds. The proposed model converges to 28 bp RMSE by epoch 25, whereas the baseline plateaus at 55 bp. A secondary y-axis (right) plots the no-arbitrage penalty term, showing monotonic decay only when the constraint layer is active [10].

5.5. Figure 5 – Calibration-Time Cumulative Distribution

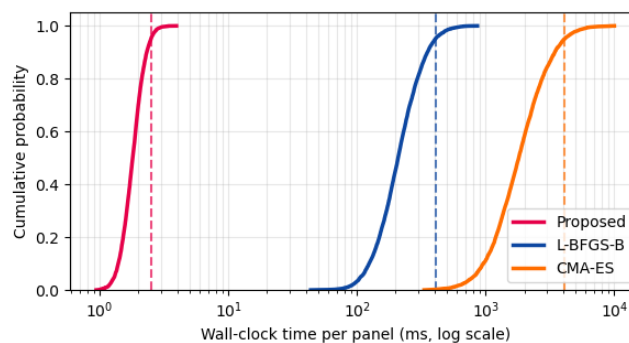


Figure 5. Calibration-Time Cumulative Distribution

An empirical CDF compares wall-clock time per panel (log-scale x-axis) across three methods: proposed network (blue), L-BFGS-B (orange), CMA-ES (green). The network’s median latency is

1.8ms; 95% of panels finish within 2.5ms. L-BFGS-B reaches 50% completion at 210ms, CMA-ES at 1.8s. Vertical dashed lines mark the 95th percentiles for each method [11].

6. Discussion, Limitations & Future Work

6.1. Extension to rough-volatility and Lévy SV models

The current Heston backbone yields closed-form characteristic functions, enabling fast Fourier pricing and analytical gradients. Rough-volatility specifications such as the rough-Heston model, where the variance process is driven by a fractional Brownian motion with Hurst exponent $H < \frac{1}{2}$, no longer possess Markovian dynamics. Their characteristic functions involve fractional Riccati equations that must be solved numerically along a contour in the complex plane, breaking the sub-millisecond constraint. A hybrid strategy is conceivable: train a neural surrogate that maps θ to approximate characteristic exponents on a low-rank manifold, then embed the surrogate inside the existing calibration network. Lévy-type SV models (Bates, Barndorff-Nielsen–Shephard) introduce jump components whose Fourier representations contain additional jump compensators; leveraging neural jump filters learned from high-frequency returns could retain end-to-end differentiability while capturing heavier tails.

6.2. Bid-ask spreads via distributional regression

Market quotes arrive as tuples (bid, ask) rather than mid prices. Treating the mid as the ground-truth target discards valuable distributional information and may bias the network toward illiquid strikes. A quantile regression loss can be substituted: let the network output not a single price but a triplet $(q_{\{0.25\}}, q_{\{0.50\}}, q_{\{0.75\}})$; the objective becomes a pinball loss summed over bid and ask quantiles. Because the loss is piece-wise linear and differentiable, back-propagation remains straightforward. The constraint layer must then ensure that the entire predicted distribution is free of butterfly arbitrage, a condition that translates to convexity of the price curve in strike space for every probability level.

6.3. Theoretical guarantees for the inverse map

Universal approximation theorems guarantee that a sufficiently wide feed-forward network can approximate any continuous mapping, yet they are silent about injectivity or constraint satisfaction. The inverse map $\sigma \mapsto \theta$ is ill-posed: distinct parameter vectors can yield nearly identical implied-volatility surfaces when the observation grid is sparse. Regularization by the no-arbitrage penalty acts as a Tikhonov-type stabilizer, restricting the hypothesis class to the set of surfaces whose second derivatives are non-negative. A future direction is to derive covering numbers for this constrained class and establish sample-complexity bounds. Furthermore, leveraging recent results on monotone operator networks could yield certificates that the learned map is contractive in a Sobolev norm, ensuring stability under adversarial perturbations of the input surface.

References

- [1] X.-J. He and W. Chen, “A closed-form pricing formula for European options under a new stochastic volatility model with a stochastic long-term mean,” *Math Finan Econ*, vol. 15, no. 2, pp. 381–396, Mar. 2021, doi: 10.1007/s11579-020-00281-y.
- [2] A. W. Van Der Stoep, L. A. Grzelak, and C. W. Oosterlee, “COLLOCATING VOLATILITY: A COMPETITIVE ALTERNATIVE TO STOCHASTIC LOCAL VOLATILITY MODELS,” *Int. J. Theor. Appl. Finan.*, vol. 23, no. 06, p. 2050038, Sep. 2020, doi: 10.1142/S0219024920500387.
- [3] R. Merino, J. Pospíšil, T. Sobotka, T. Sottinen, and J. Vives, “DECOMPOSITION FORMULA FOR ROUGH VOLTERRA STOCHASTIC VOLATILITY MODELS,” *Int. J. Theor. Appl. Finan.*, vol. 24, no. 02, p. 2150008, Mar. 2021, doi: 10.1142/S0219024921500084.
- [4] B. Jourdain and A. Zhou, “Existence of a calibrated regime switching local volatility model,” *Mathematical Finance*, vol. 30, no. 2, pp. 501–546, Apr. 2020, doi: 10.1111/mafi.12231.

- [5] A. Sepp and P. Rakhmonov, “LOG-NORMAL STOCHASTIC VOLATILITY MODEL WITH QUADRATIC DRIFT,” *Int. J. Theor. Appl. Finan.*, vol. 26, no. 08, p. 2450003, Dec. 2023, doi: 10.1142/S0219024924500031.
- [6] S.-C. Necula and V.-D. Păvăloaia, “AI-Driven Recommendations: A Systematic review of the state of the art in E-Commerce,” *Applied Sciences*, vol. 13, no. 9, p. 5531, 2023.
- [7] S. Acharya, “Study of the effectiveness of chatbots in customer service on e-commerce websites,” 2023, Accessed: Jan. 08, 2025. [Online]. Available: https://www.theseus.fi/bitstream/handle/10024/807998/Acharya_Shiva.pdf?sequence=4.
- [8] A. Brace, K. Gellert, and E. Schlögl, “SOFR term structure dynamics—Discontinuous short rates and stochastic volatility forward rates,” *Journal of Futures Markets*, vol. 44, no. 6, pp. 936–985, Jun. 2024, doi: 10.1002/fut.22499.
- [9] A. S. George, “AI Supremacy at the Price of Privacy: Examining the Tech Giants’ Race for Data Dominance,” *Partners Universal Innovative Research Publication*, vol. 3, no. 1, pp. 26–43, 2025.
- [10] J.-L. Dupret, J. Barbarin, and D. Hainaut, “Impact of rough stochastic volatility models on long-term life insurance pricing,” *Eur. Actuar. J.*, vol. 13, no. 1, pp. 235–275, Jun. 2023, doi: 10.1007/s13385-022-00317-1.
- [11] M. Madanchian, “The Impact of Artificial Intelligence Marketing on E-Commerce Sales,” *Systems*, vol. 12, no. 10, p. 429, 2024.