

Zheng, Y., Yang, Y. and Chen, B. (2021) Incorporating Prior Financial Domain Knowledge into Neural Networks for Implied Volatility Surface Prediction. In: 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), Singapore, 14-18 Aug 2021, pp. 3968-3975. ISBN 9781450383325

(doi:[10.1145/3447548.3467115](https://doi.org/10.1145/3447548.3467115))

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

© 2021 Association for Computing Machinery. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in the Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), Singapore, 14-18 Aug 2021

<http://eprints.gla.ac.uk/243104/>

Deposited on: 1 June 2021

# Incorporating Prior Financial Domain Knowledge into Neural Networks for Implied Volatility Surface Prediction

Yu Zheng

<sup>1</sup> School of Finance  
Southwestern University of Finance  
and Economics

<sup>2</sup> Inboc Technologies Ltd.  
zhengyu@swufe.edu.cn

Yongxin Yang

<sup>1</sup> School of Informatics  
University of Edinburgh  
<sup>2</sup> ArrayStream Technologies Ltd.  
yongxin.yang@ed.ac.uk

Bowei Chen\*

Adam Smith Business School  
University of Glasgow  
bowei.chen@glasgow.ac.uk

## ABSTRACT

In this paper we develop a novel neural network model for predicting implied volatility surface. Prior financial domain knowledge is taken into account. A new activation function that incorporates volatility smile is proposed, which is used for the hidden nodes that process the underlying asset price. In addition, financial conditions, such as the absence of arbitrage, the boundaries and the asymptotic slope, are embedded into the loss function. This is one of the very first studies which discuss a methodological framework that incorporates prior financial domain knowledge into neural network architecture design and model training. The proposed model outperforms the benchmarked models with the option data on the S&P 500 index over 20 years. More importantly, the domain knowledge is satisfied empirically, showing the model is consistent with the existing financial theories and conditions related to implied volatility surface.

## KEYWORDS

Mathematical Finance; Implied Volatility Surface; Deep Neural Networks; Interpretable Machine Learning

### ACM Reference Format:

Yu Zheng, Yongxin Yang, and Bowei Chen. 2021. Incorporating Prior Financial Domain Knowledge into Neural Networks for Implied Volatility Surface Prediction. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21)*, August 14–18, 2021, Virtual Event, Singapore. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3447548.3467115>

## 1 INTRODUCTION

Machine learning algorithms are essentially data-driven models which mainly focus on producing accurate predictions. They are being used for a wider array of macro and micro level prediction tasks. According to Ernst and Young [14], machine learning applications in finance have become one of the hottest sectors globally,

with the expected direct investment growth of 63% from 2016 to 2022. Despite COVID-19, a recent survey from the Bank of England showed that there are still a third of banks would increase their investments in the number of planned or existing machine learning and data science projects [4]. Therefore, it is interesting for machine learning researchers and financial analysts to investigate this sector because it generates a lot of practical questions and challenges, and addressing them will result in positive economic and social consequence.

Although with excellent prediction performance, machine learning is usually used as the “black box” model in many financial applications. Compared with the well-developed models from mathematical finance, machine learning algorithms are less interpretable, e.g., features can be not understandable, and the learning process is not transparent or mathematically tractable. More importantly, they are not aligned with the well-developed financial theories. Therefore, many financial institutions are slow to adopt machine learning algorithms (particularly neural networks) into their major business operations. Developing interpretable machine learning models that are consistent with the existing financial markets and theories will resolve the bottleneck and will boost the applications of machine learning into finance.

In this paper, we propose a novel neural network model tailored for implied volatility surface prediction. Implied volatility is an important financial metric or indicator that captures the market’s view of the likelihood of changes in a given asset price. Technically, an *implied volatility* is defined as the inverse problem of option pricing, mapping from the option price of the asset in the current market to a single value [12]. When it is plotted against the option strike price and the time to maturity, it is called the *implied volatility surface*. Prior financial domain knowledge related to implied volatility surface includes: 1) the empirical evidence volatility smile; and 2) financial conditions such as the absence of arbitrage, the boundaries and the asymptotic slope. We use different ways to incorporate these domain knowledge. For the former, a new activation function that produces volatility smile is proposed, and it is used for the hidden nodes that process the underlying asset price. For the latter, financial conditions are embedded into the loss function for neural network training. In the experiments, we validate the proposed model with the option data on the S&P 500 index over a period of 20 years. Compared with the existing studies, our experimental settings are more challenging and this requires our model to be more robust and stable in producing convincing results. Our model outperforms the widely used state-of-the-art model in finance and other benchmarked neural network models on

\*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
KDD '21, August 14–18, 2021, Virtual Event, Singapore

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8332-5/21/08...\$15.00

<https://doi.org/10.1145/3447548.3467115>

the mean average percentage error in both training and test sets. In the meantime, the incorporated prior financial domain knowledge are met empirically.

Technology wise, our study makes a methodological contribution. We propose a framework of incorporating prior financial domain knowledge into neural network design and training. Therefore, the developed model is aligned well with the existing empirical evidence and financial theories related to implied volatility surface. This is an important step for interpretable machine learning, and we hope the framework can motivate many other investigations of machine learning applications in finance. On the other hand, from the application perspective, we develop a best-performing prediction model, and to the best of our knowledge, this is one of the very first neural networks tailored for implied volatility surface.

The rest of the paper is organised as follows. Section 2 reviews the related literature. Section 3 introduces our proposed model for predicting implied volatility surface. The used dataset, our experimental settings and results are presented in Section 4. Finally, we conclude the paper in Section 5.

## 2 RELATED WORK

Our research in this paper touches upon two streams of literature: mathematical finance and machine learning. For the former, we introduce the basic concepts and the related studies of option pricing and volatility modelling. For the latter, we review a number of recent applications of machine learning in finance, with special focuses on option pricing and volatility modelling.

### 2.1 Option Pricing and Volatility Modelling

In 1973, Black and Scholes [6] proposed an elegant closed-form pricing formula for the European style call options written on financial assets. Their model is simply called the *Black-Scholes option pricing model*, in which an underlying financial asset price is driven by a geometric Brownian motion [37] that contains a drift and a volatility, and the volatility term shows the small fluctuations of asset returns representing risk. The seminal work of Black and Scholes opened the floodgates of studying mathematical models in finance, and volatility models have soon become popular since then [15, 39].

Volatility models in finance can be classified into two groups [22]. The first group is called *indirect methods*, in which an implied volatility is driven by another dynamic model such as local volatility models, stochastic volatility models and Lévy models [21, 24, 27, 34, 38]. Models in this group usually have a limited number of parameters, and the volatility term is fitted by the market data along with the asset dynamics such as the geometric Brownian motion and the mean-reversion jump-diffusion process. These models exhibit mathematical elegance but are sometimes invalid empirically. Time-dependent parameters can be included but they will greatly increase computational time and optimisation difficulty in model calibration. The second group is called *direct methods*, in which an implied volatility is specified explicitly. Direct methods can also be divided into two types. The first type specifies the dynamics of an implied volatility surface and assumes it evolves continuously over time [9, 12]. The second type focuses on the static representation of implied volatility surface that uses either parametric or

non-parametric methods to fit an implied volatility surface and then for prediction [13, 23, 26]. Our proposed method is a static model. In this group, the stochastic volatility inspired (SVI) model is the most commonly used method [16]. It models the implied volatility slice for a fixed time to maturity. Gatheral and Jacquier then further improved the SVI model with a simpler representation on the conditions for no static arbitrage, and this improved SVI model is called the *surface SVI (SSVI)* model, which is the recent advance in mathematical finance and has been widely adopted by investors [17]. Therefore, we choose the SSVI model as one of the benchmarked models in this paper.

### 2.2 Machine Learning Applications in Finance

Applying machine learning to asset pricing and volatility prediction can be traced back to the late 1980's or early 1990's. In the early stages, single hidden layer neural networks were used to estimate option price [32] and to predict the volatility of the S&P 100 index [33]. Then, various machine learning algorithms were introduced, including ensemble methods [3, 18], kernel machines [11], Gaussian process [40], and deep learning models such as hybrid neural networks [28], gated neural networks [41], and recurrent neural networks [31]. In addition to the conventional financial data, several other recent studies developed models which use verbal, vocal and social information [35, 36, 42].

It should be noted that our research in this paper focuses on predicting the implied volatility surface when the underlying asset price and the time to maturity of the corresponding option quotes are given. This kind of prediction in essence is an inverse engineering of option pricing and it is not time series forecasting. Therefore, our scenario is different to the above mentioned studies. Our neural network architecture design is inspired by the work of [41] but has three significant differences. First, our model aims to predict the implied volatility rather than option price. Second, we design a new activation function that can incorporate volatility smile. Third, we embed the conditions related to implied volatility into neural network training.

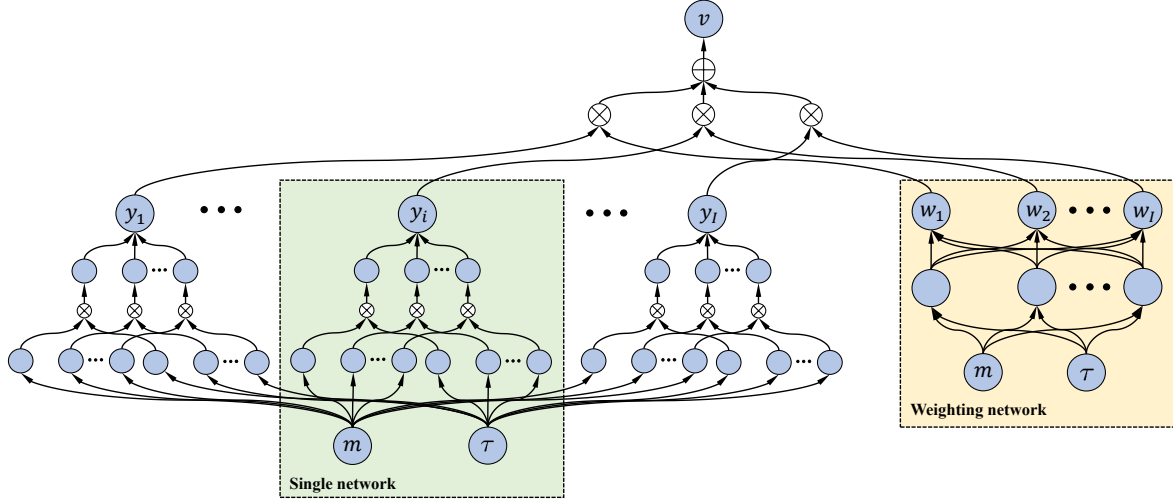
## 3 MODEL

In this section, we firstly introduce the preliminaries of implied volatility surface and lays down the mathematical settings. We then discuss our deep neural network architecture design and further explain how do we incorporate the prior financial domain knowledge into neural network.

### 3.1 Prior Financial Domain Knowledge

In mathematical finance, the spot price of an asset is usually modelled as a stochastic process  $(S_t)_{t \geq 0}$  that is defined on a filtered probability space  $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ , where  $t$  is the time index,  $\Omega$  is the sample space,  $\mathcal{F}$  is the sigma-field,  $(\mathcal{F}_t)_{t \geq 0}$  is the filtration, and  $\mathbb{P}$  is the probability space. The financial market is assumed to be arbitrage-free and a financial product's time to maturity (i.e., the time remaining until a financial contract expires) is always finite.

As mentioned previously, implied volatility is the inverse engineering of option pricing. It can be obtained by inverting the Black-Scholes option pricing model [6], in which one needs to



**Figure 1: Neural network architecture design tailored to implied volatility surface. The proposed multi model consists of several single models and their weights are determined by the weighting network. Bias terms are omitted;  $\otimes$  is the multiplication gate operator;  $\oplus$  is the addition gate operator.**

determine the constant interest rate and dividends from the market data. To avoid dealing with them, the forward measure can be used instead. Let  $(F_{t,T})_{t \geq 0}$  be the forward price of the asset with maturity date  $T$ , where  $0 \leq t \leq T$ . Then

$$F_{t,T} = \frac{S_t}{B(t,T)},$$

where  $B(t,T)$  is the price of a zero-coupon bond at time  $t$  which will pay one unit at time  $T$ . The absence of arbitrage ensures there exists an equivalent martingale measure in which  $(F_{t,T})_{t \geq 0}$  is a martingale [12]. In probability theory, a *martingale* is a stochastic process for which, at a particular time, the conditional expectation of the next value in the process is equal to the present value and regardless of all the previous values. So the log forward moneyness  $m$  can be defined and used as the underlying, where  $m = \log\{K/F_{t,T}\}$  and  $K$  is the strike price.

Another important variable is the time to maturity, which can be defined as

$$\tau = \frac{T-t}{A},$$

where  $A$  is the annualization factor. Therefore, in our mathematical setting, the implied volatility  $v$  can be written as a function of the log forward moneyness  $m$  and the time to maturity  $\tau$ .

**THEOREM 1.** Let  $d_{\pm}(m, \tau) = -\frac{m}{\sqrt{\tau}v(m, \tau)} \pm \frac{1}{2}\sqrt{\tau}v(m, \tau)$ ,  $n(\cdot)$  denote the density function of a standard normal distribution, and  $N(\cdot)$  denote its cumulative function. The following conditions are required to be met for the implied volatility  $v$ :

- 1) (Positivity) For  $(m, \tau) \in \mathbb{R} \times \mathbb{R}^+$ ,  $v(m, \tau) > 0$ .
- 2) (Twice Differentiation) For  $\tau > 0$ ,  $m \rightarrow v(m, \tau)$  is twice differentiable on  $\mathbb{R}$ .
- 3) (Monotonicity) For  $m \in \mathbb{R}$ ,  $\tau \rightarrow \sqrt{\tau}v(m, \tau)$  is increasing on  $\mathbb{R}^+$ , then

$$v(m, \tau) + 2\tau\partial_{\tau}v(m, \tau) \geq 0.$$

- 4) (Absence of Butterfly Arbitrage) For  $(m, \tau) \in \mathbb{R} \times \mathbb{R}^+$ ,

$$\left[1 - \frac{m\partial_m v(m, \tau)}{v(m, \tau)}\right]^2 - \frac{1}{4}\left[v(m, \tau)\tau\partial_m v(m, \tau)\right]^2 + \tau v(m, \tau)\partial_{mm} v(m, \tau) \geq 0.$$

- 5) (Limiting Behaviour) If  $\tau > 0$ , then

$$\lim_{m \rightarrow +\infty} d_+(m, \tau) = -\infty.$$

- 6) (Right Boundary) If  $m \geq 0$ , then

$$N(d_-(m, \tau)) - \sqrt{\tau}\partial_m v(m, \tau)n(d_-(m, \tau)) \geq 0.$$

- 7) (Left Boundary) If  $m < 0$ , then

$$N(-d_-(m, \tau)) + \sqrt{\tau}\partial_m v(m, \tau)n(d_-(m, \tau)) \geq 0.$$

- 8) (Asymptotic Slope) If  $\tau > 0$ , then

$$2|m| - v^2(m, \tau)\tau > 0.$$

Simply, Theorem 1 conditions 1-5 ensure the absence of arbitrage [20]; conditions 6-7 specify the boundaries [8]; and condition 8 is the asymptotic slope [29].

In addition to Theorem 1, implied volatility has an important empirical evidence (or stylised fact) called the *volatility smile* – for a given time to maturity, when the implied volatility is plotted against the strike price, it creates a line that slopes upward on either end, looking like a “smile” [12]. In the following discussion, we will present our neural network model by incorporating these mentioned financial conditions and empirical evidence related to implied volatility surface.

### 3.2 Deep Neural Network Architecture Design

Figure 1 presents a schematic view of our neural network architecture. The model input includes the log forward moneyness  $m$  and the time to maturity  $\tau$ , and the model output is the implied

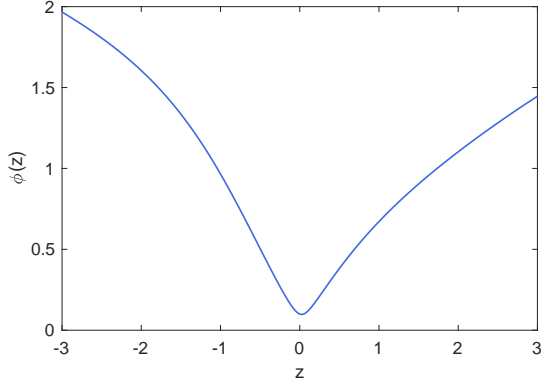


Figure 2: Smile function  $\phi(\cdot)$  where  $\epsilon = 0.01$ .

volatility  $v$ . The proposed neural network is constructed from sub-networks with two types of architectural structures: 1) several simple networks (called *single networks*) which predict the implied volatility separately; and 2) a weighting network which determines the “votings” of the predicted implied volatilities towards the final prediction. Similar to [41], multiplication and addition gate operators are used to process and merge information related to the log forward moneyness  $m$ , the time to maturity  $\tau$  and single networks.

**DEFINITION 1 (Smile Function).** For any  $z \in \mathbb{R}$ ,

$$\phi(z) = \left[ z \tanh\left(z + \frac{1}{2}\right) + \tanh\left(-\frac{1}{2}z + \epsilon\right) \right]^{1/2}, \quad z \in \mathbb{R}, \quad (1)$$

where  $\tanh(\cdot)$  is the hyperbolic tangent function and  $\epsilon$  is a small value to ensure numerical stability.

As shown in Figure 2, the defined smile function exhibits a skew pattern like volatility smile. We apply the smile function  $\phi(\cdot)$  to the nodes that correspond to  $m$  and the sigmoid function  $\psi(\cdot)$  for the nodes that correspond to  $\tau$ . It is not difficult to see that the positivity and twice differentiation conditions in Theorem 1 are met and the limiting behaviour condition can be proven theoretically by inverting the Black-Scholes option pricing model [20].

Our proposed neural network can be expressed as follows:

$$\hat{v} = \sum_{i=1}^I y_i w_i, \quad (2)$$

$$y_i = \sum_{j=1}^J \phi(m\tilde{w}_j^{(i)} + \tilde{b}_j^{(i)}) \psi(\tau\tilde{w}_j^{(i)} + \tilde{b}_j^{(i)}) e^{\tilde{w}_j^{(i)}} + e^{\tilde{b}_j^{(i)}}, \quad (3)$$

$$w_i = \frac{\exp\left\{\sum_{k=1}^K \psi(m\dot{w}_{1,k} + \tau\dot{w}_{2,k} + \dot{b}_k)\ddot{w}_{k,i} + \ddot{b}_i\right\}}{\sum_{i=1}^I \exp\left\{\sum_{k=1}^K \psi(m\dot{w}_{1,k} + \tau\dot{w}_{2,k} + \dot{b}_k)\ddot{w}_{k,i} + \ddot{b}_i\right\}}, \quad (4)$$

where  $I$  is the number of single networks,  $J$  is the number of layers in each single network, and  $K$  is the number of layers in the weighting network. For the  $i$ th single network,  $y_i$  is its prediction output. To ensure it is non-negative, we consider the exponential form of the weight  $e^{\tilde{w}_j^{(i)}}$  and the bias  $e^{\tilde{b}_j^{(i)}}$ . For the hidden layers,  $\tilde{w}_j$  and  $\tilde{b}_j$  are the weight and the bias of the  $j$ th hidden node that

corresponds to  $m$ , and  $\tilde{w}_j$  and  $\tilde{b}_j$  are the weight and the bias of the  $j$ th hidden node that corresponds to  $\tau$ . Therefore, there is a total of  $5J + 1$  parameter values for each single network. The weighting network predicts the weights of single networks towards the final prediction, where  $w_i$  is its prediction for the  $i$ th single network,  $\dot{w}_{\cdot,k}$  is the weight of the  $k$ th hidden node,  $\dot{b}_k$  is the bias of the  $k$ th hidden node,  $\ddot{w}_{\cdot,i}$  is the weight of the  $i$ th single network, and  $\ddot{b}$  is the bias of the  $i$ th single network. Since the dimensions of  $\dot{w}$ ,  $\dot{b}$ ,  $\ddot{w}$ ,  $\ddot{b}$  are  $2 \times K$ ,  $K \times 1$ ,  $K \times I$ , and  $I \times 1$ , respectively, the total number of parameter values in the our model is  $(5J + K + 2)I + 3K$ .

### 3.3 Embedding Constraints into Optimisation

Training the designed neural network solves an optimisation problem that finds for parameters which result in a minimum loss when evaluating the samples in the training data. We define a loss function tailored to implied volatility surface by embedding the related conditions from mathematical finance:

$$\ell = \ell_0 + \gamma\ell_1 + \delta\ell_2 + \eta\ell_3 + \rho\ell_4 + \omega\ell_5, \quad (5)$$

where  $\ell_0$  represents the data loss,  $\ell_1, \dots, \ell_4$  are the loss functions that incorporate financial conditions discussed previously, and  $\ell_5$  is the regularization term to avoid over-fitting.

The data loss  $\ell_0$  is defined as a joint loss from the mean squared log error (MSLE) and the mean squared percentage error (MSPE):

$$\ell_0 = \frac{\alpha}{N} \sum_{n=1}^N (\log v_n - \log \hat{v}_n)^2 + \frac{\beta}{N} \sum_{n=1}^N \left( \frac{v_n - \hat{v}_n}{v_n} \right)^2, \quad (6)$$

where  $N$  is the total number of the training samples,  $v_n$  is the ground truth implied volatility for the  $n$ th sample,  $\hat{v}_n$  is the predicted implied volatility for the  $n$ th sample,  $\alpha$  and  $\beta$  are hyperparameters that control the weights of MSLE and MSPE, respectively. We use the joint data loss here because it is efficient in dealing with sensitive data or high-dimensional feature spaces [19].

The monotonicity condition is specified by  $\ell_1$ , defined by

$$\ell_1 = \sum_{p=1}^P \sum_{q=1}^Q \max\{0, -a(m_p, \tau_q)\}, \quad (7)$$

where  $P$  and  $Q$  are the number of samples, and  $a(m, \tau) = v(m, \tau) + 2\tau\partial_\tau v(m, \tau)$ . The objective of  $\ell_1$  is to push  $a(m, \tau)$  to be non-negative. This can be achieved by randomly sampling  $P$  unique values from the domain of  $m$  and  $Q$  unique values from the domain of  $\tau$ . Penalty is added by  $\ell_1$  if  $a(m, \tau)$  is negative for the sampled  $(m, \tau)$  pairs.

The absence of a butterfly arbitrage condition is specified by  $\ell_2$ , defined as follows

$$\ell_2 = \sum_{p=1}^P \sum_{q=1}^Q \max\{0, -b(m_p, \tau_q)\}, \quad (8)$$

where

$$b(m, \tau) = \left(1 - \frac{m\partial_m v(m, \tau)}{v(m, \tau)}\right)^2 - \frac{(v(m, \tau)\tau\partial_m v(m, \tau))^2}{4} + \tau v(m, \tau)\partial_{mm} v(m, \tau).$$

The objective of  $\ell_2$  is to push  $b(m, \tau)$  to be non-negative. This can be achieved using the same way as  $\ell_1$ , by randomly sampling  $P$  unique values from the domain of  $m$  and  $Q$  unique values from the domain of  $\tau$ .

The left and the right boundary conditions are specified by  $\ell_3$ , defined as follows

$$\ell_3 = \sum_{p_1=1}^{P_1} \sum_{q=1}^Q \max\{0, -c_1(m_{p_1}, \tau_q)\} + \sum_{p_2=1}^{P_2} \sum_{q=1}^Q \max\{0, -c_2(m_{p_2}, \tau_q)\}. \quad (9)$$

where

$$c_1(m, \tau) = N(d_-(m, \tau)) - \sqrt{\tau} \partial_m v(m, \tau) n(d_-(m, \tau)),$$

$$c_2(m, \tau) = N(-d_-(m, \tau)) + \sqrt{\tau} \partial_m v(m, \tau) n(d_-(m, \tau)).$$

The objective is to push both functions to be non-negative. To achieve this, we sample  $P_1$  unique non-negative values from the domain of  $m$ ,  $P_2$  unique negative values from the domain of  $m$  and  $Q$  unique values from the domain of  $\tau$ .

The asymptotic condition is specified by  $\ell_4$ , defined by

$$\ell_4 = \sum_{p=1}^P \sum_{q=1}^Q \max\{0, -(g(m_p, \tau_q) - \epsilon)\}, \quad (10)$$

where

$$g(m, \tau) = 2|m| - v^2(m, \tau)\tau,$$

and  $\epsilon$  is a very small value. The aim of  $\ell_4$  is to ensure  $g(m, \tau)$  is positive. Therefore, similar to  $\ell_1, \ell_2, \ell_3$ , we sample  $P$  unique values from the domain of  $m$  and  $Q$  unique values from the domain of  $\tau$ . It is worth mentioning that the values of  $m$  and  $\tau$  in  $\ell_1, \dots, \ell_4$  can also be sampled from the training data. However, the trained neural network may fail to meet those conditions when the given values of  $m$  and  $\tau$  for prediction are out of the scope of the training data. If the training data have limited observations of input variables, creating synthetic data by sampling values from their domains is an effective way to train the model with good generalization capabilities [10, 30]

The regularization term  $\ell_5$  is defined by

$$\ell_5 = \sum_{i=1}^I \|\tilde{w}^{(i)}\|_F^2 + \sum_{i=1}^I \|\tilde{w}^{(i)}\|_F^2 + \sum_{i=1}^I \|\hat{w}^{(i)}\|_F^2 + \|\tilde{w}\|_F^2 + \|\hat{w}\|_F^2, \quad (11)$$

where  $\|\cdot\|_F^2$  is the squared Frobenius norms.

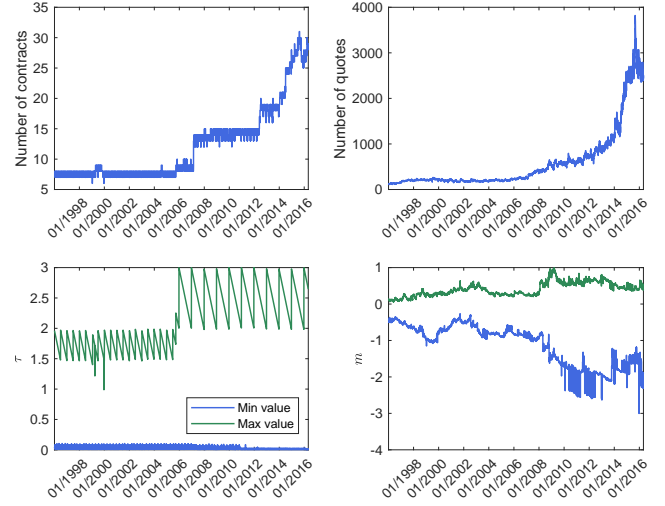
## 4 EXPERIMENTS

The used datasets are firstly introduced. We then provide the details of experimental design, including examined models and their training settings. The experimental results are finally presented and discussed.

### 4.1 Data

We use the option and the zero-coupon yield curve data from OptionMetrics, and the Overnight Index Swap (OIS) data from Bloomberg.<sup>1</sup> Our option data is for the S&P 500 index. It is one of the most commonly followed stock market indices which measures the

<sup>1</sup>OptionMetrics is a leading provider of historical implied volatility, greeks, and option pricing data for financial markets and Bloomberg is a premier financial services company.



**Figure 3: Time series plots of cleaned option contracts and their quotes from 04/01/1996 to 29/04/2016**

stock performance of 500 large companies listed on stock exchanges in the United States. The option data contains a total of 5,116 trading days, covering the period from 04/01/1996 to 29/04/2016. The zero-coupon yield curve shows the relationship between the level of the interest rate (representing the cost of borrowing) and the time to maturity. It is constructed based on the London Inter-Bank Offered Rate (LIBOR). However, after the 2008 financial crisis, the LIBOR-based zero curve is not risk-free [1]. Therefore, adjustments are performed for the period after 01/01/2008. Specifically, we extract the OIS data from Bloomberg and bootstrap the zero rate curve. In addition, we use the cubic spline to interpolate the risk-free rates in order to match the option maturity, and compute the forward price using the put-call parity [5].

The option data is further processed. Option quotes which are less than 3/8 are excluded because they are close to the tick size and can be misleading. The bid-ask mid-point price is calculated as a proxy for the closing price. The in-the-money option quotes are excluded because of the small transaction volume [7]. The existing studies usually do not analyse option contracts with the time to maturity of less than 7 days [2]. However, as these options are getting popular recently (e.g., weekly index options), we here analyse option contracts with a short time to maturity and only exclude the contracts with the maturity of less than 2 days. Analysing options with a short maturity is challenging because this requires our model with high robustness and stability. As shown in Figure 3, our prepared data finally contains 63,338 option contracts with 2,986,754 valid quotes. The quotes are then used to compute the ground truth implied volatility values by inverting the Black-Scholes option pricing formula.

### 4.2 Experimental Settings

Table 1 summarises the examined models in the experiments. In addition to our proposed model (simply denoted by Multi), we also deploy the SSVI model, and other neural network models. For the former, we aim to investigate if our proposed model can

Model	Description
SSVI	[17]
Multi	The proposed model specified in Eqs. (2)-(11).
Multi <sup>†</sup>	The Multi model trained without embedding $\ell_1, \ell_2, \ell_3, \ell_4$ .
Single	The single network model so there is no weighting network, and $  \dot{w}  _F^2$ and $  \ddot{w}  _F^2$ are not included in the regularization term $\ell_5$ for the model training
Single <sup>†</sup>	The Single model trained without embedding $\ell_1, \ell_2, \ell_3, \ell_4$ .
Vanilla	The neural network model with the simplest architecture – it has a single hidden layer which only uses the sigmoid activation function and the model's output is censored to be non-negative.
Vanilla <sup>†</sup>	The vanilla model trained without embedding $\ell_1, \ell_2, \ell_3, \ell_4$ .

**Table 1: Summary of the examined models.**

Model	Hyperparameter									
	$I$	$J$	$K$	$\alpha$	$\beta$	$\gamma$	$\delta$	$\eta$	$\rho$	$\omega$
Multi	4	8	5	1	1	10	1	10	1	5e-5
Multi <sup>†</sup>	4	8	5	1	1	0	0	0	0	5e-5
Single	1	32	-	1	1	10	1	10	1	5e-5
Single <sup>†</sup>	1	32	-	1	1	0	0	0	0	5e-5
Vanilla	1	32	-	1	1	10	1	10	1	5e-5
Vanilla <sup>†</sup>	1	32	-	1	1	0	0	0	0	5e-5

**Table 2: Hyperparameter settings of neural network models, where all models use the same learning rate 0.1 and the same number of iterations  $2e+4$ .**

achieve a better prediction performance than the state-of-the-art method from mathematical finance. For the latter, we want to see if the designed architecture (i.e., the ensemble of multiple single networks) can improve the prediction capability. To this end, the benchmarked neural networks include: 1) a single network (denoted by Single) where there is no weighting network; and 2) a simplest neural network (denoted by Vanilla) which has a single hidden layer, only uses the sigmoid activation function and the model's output is censored to be non-negative. Also, to further justify the importance of embedding financial conditions, all the neural network models are trained under a setting where  $\ell_1, \ell_2, \ell_3, \ell_4$  are removed from the total loss function  $\ell$ . To simplify the discussion, these neural network models are denoted with a superscript <sup>†</sup>.

Table 2 presents our hyperparameter settings of the examined neural networks. To avoid the size effect on model performance, the compared neural networks with the same architecture design are specified with the same model size and with the same hyperparameter values. As also mentioned earlier in Section 3, embedding  $\ell_1, \ell_2, \ell_3, \ell_4$  requires synthetic data. In our neural network training, the ratio of real market data and synthetic data is 1/6. The log forward money  $m$  is sampled in  $[-6, -3] \cup [3, 6]$  for the asymptotic

Model	Training		Test	
	Mean	STD	Mean	STD
Multi	1.74	0.50	3.34	2.18
Multi <sup>†</sup>	1.76	0.50	3.35	2.17
Single	2.15	0.67	3.60	2.12
Single <sup>†</sup>	1.82	0.52	3.38	2.16
Vanilla	3.21	0.98	4.46	2.07
Vanilla <sup>†</sup>	2.87	0.80	4.18	2.04
SSVI	2.59	0.85	3.73	2.18

**Table 3: Mean and standard deviation (STD) of the MAPEs for the predicted implied volatilities from the examined models.**

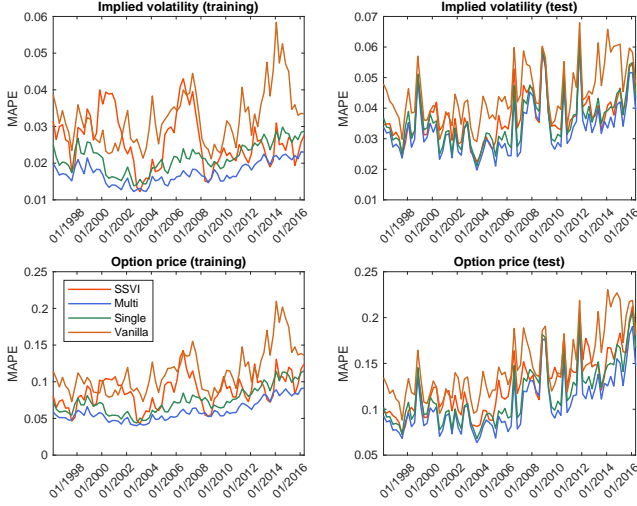
Model	Training		Test	
	Mean	STD	Mean	STD
Multi	5.97	1.86	10.64	6.72
Multi <sup>†</sup>	6.03	1.86	10.67	6.70
Single	7.38	2.57	11.64	6.68
Single <sup>†</sup>	6.20	1.91	10.77	6.67
Vanilla	11.31	3.57	14.61	6.42
Vanilla <sup>†</sup>	10.53	3.34	14.17	6.60
SSVI	8.71	2.72	12.74	6.74

**Table 4: Mean and standard deviation (STD) of the MAPEs for the option prices calculated using the predicted implied volatilities from the examined models.**

condition and in  $[-3, 3]$  for other conditions; and the time to maturity  $\tau$  is sampled in  $[0.002, 3]$ . These values are set based on the observations in Figure 3. Neural network models are trained using TensorFlow and we use Adam [25] for stochastic optimisation. We train the models using all option quotes from the previous trading day and test the models in the next day. We move the window of training and test split across all the trading days in the option data.

### 4.3 Results and Analysis

The summary statistics of the mean average percentage errors (MAPEs) of the predicted implied volatilities for the examined models in all trading days are presented in Table 3. To further investigate the effects or differences that the predicted implied volatilities can trigger in option pricing, we use them to compute the corresponding option prices and then report the summary statistics of the MAPEs of the option prices in Table 4. In both tables, the widely used SSVI model from mathematical finance underperforms the Multi, Multi<sup>†</sup>, Single and Single<sup>†</sup> models significantly but it still outperforms the simple neural networks like the Vanilla and Vanilla<sup>†</sup> models. Compared to conventional mathematical models in finance, data-driven deep learning models have shown great predictive capabilities but they should be with a proper architecture design and hyperparameters settings. Our study has been successfully validated with the results because the proposed Multi model is the best-performing prediction model in both training and test data

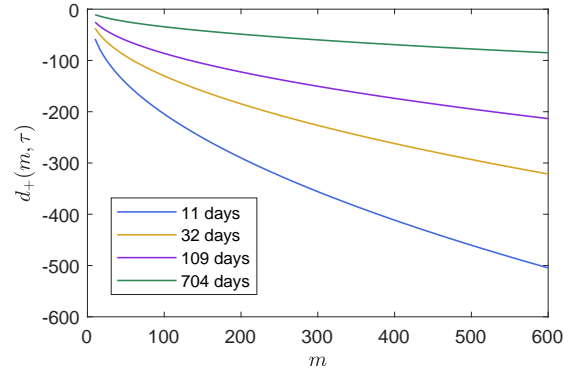


**Figure 4: Time series plot of the MAPEs for the SSVI model and the neural network models trained with full settings.**

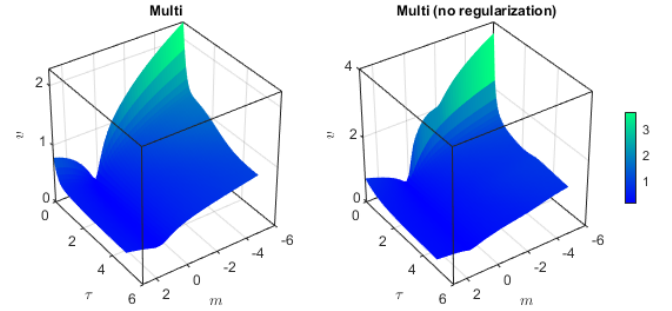
for implied volatility and option price, respectively. Figure 4 also compares the time series plots of the MAPEs of the SSVI model and the deep neural networks trained with the full settings on each sampled trading day, showing our proposed Multi model has the most stable outputs over time.

One interesting finding from Tables 3-4 is that excluding some financial conditions in training neural networks will not significantly decrease the models' prediction performance because the Multi<sup>†</sup>, Single<sup>†</sup> and Vanilla<sup>†</sup> models have a comparable performance with their counterparts which were trained with the full settings (i.e., the Multi, Single and Vanilla models). However, as discussed previously, incorporating the prior financial domain knowledge is mainly to ensure the model is consistent with the existing financial theories and assumptions rather than the model's prediction performance. In Table 5, we check if the monotonicity, boundary, absence of butterfly arbitrage and asymptotic slope conditions in Theorem 1 are empirically satisfied in the test data. Since we have discussed in Section 3 that the positivity and twice differentiation conditions are met in our neural network architecture design and the limiting behaviour condition can be proven theoretically, these three conditions are not checked by Table 5. It is not difficult to observe that the violation percentage of the test samples for the examined conditions of the Multi<sup>†</sup>, Single<sup>†</sup> and Vanilla<sup>†</sup> models are much higher than the corresponding Multi, Single and Vanilla models, showing the importance and necessity of embedding  $\ell_1, \ell_2, \ell_3, \ell_4$  loss functions. For illustration purpose only, Figure 5 demonstrates the limiting behaviour of the converted underlying forward contracts with 11, 32, 109, and 704 days duration, verifies the limiting behaviour condition.

Further, to get a sense of what an implied volatility surface looks like for our readers, Figure 6 shows the surfaces resulting from the Multi model and the Multi model trained without the regularization term for 11/01/2016. It is clear that the volatility smiles in the former are much smoother than those of the latter, which also verifies the effectiveness of regularization.



**Figure 5: Limiting behaviour of the converted forward contracts with 11, 32, 109, and 704 days duration.**



**Figure 6: Implied volatility surface on 11/01/2016 predicted by: (a) the multi model; and (b) the multi model without regularization.**

## 5 CONCLUSION

In this paper, we developed a novel neural network to predict implied volatility surfaces. Unlike many previous studies where the machine learning algorithms are mainly used as the “black box” in finance, our model is tailored to the unique characteristics of implied volatility surface. To the best of our knowledge, this is one of the very first studies which discuss a methodological framework that integrates the data-driven machine learning algorithms (particularly neural networks) with the related financial theories and empirical evidence. The proposed model framework can be easily extended and applied to solve other similar computational problems in finance and business analytics such as inventory pricing and revenue management.

In addition to the methodological contribution, we validated the proposed model empirically with the option data on the S&P 500 index. Compared with the existing studies, our experimental settings are more challenging because the used option data is over 20 years and the options with the short time to maturity are examined. Therefore, our model needs to be more robust in order to produce convincing results. As presented in the experiments section, our model outperforms the widely used SSVI model from mathematical finance and other benchmarked neural networks.

Model	Monotonicity	Absence of butterfly arbitrage	Left boundary	Right boundary	Asymptotic slope
Multi	0.00%	7.02e-6%	0.00%	0.00%	0.00%
Multi <sup>†</sup>	1.28%	4.87%	0.00%	14.06%	0.00%
Single	0.00%	5.56e-3 %	0.00%	0.05%	0.00%
Single <sup>†</sup>	0.00%	14.88%	0.95%	5.16%	0.00%
Vanilla	3.75e-3%	1.53e-2%	4.07e-3%	0.00%	0.00%
Vanilla <sup>†</sup>	5.32%	5.72%	14.63%	0.00%	0.54%

Table 5: Percentages of the test set samples which do not satisfy conditions 3,4,6,7,8 of Theorem 1.

More importantly, the conventional financial conditions and empirical evidence are met empirically, which resolve the bottleneck of data-driven machine learning applications in finance.

## ACKNOWLEDGMENTS

This research has been conducted with the support of: 1) the Financial Innovation Center of the Southwestern University of Finance and Economics; and 2) the Key Laboratory of Financial Intelligence and Financial Engineering of Sichuan Province. 3) the UK Economic and Social Research Council through the Impact Acceleration Accounts Business Booster Funding to the University of Glasgow. The first author also acknowledges the Imperial College Business School with the support of the high performance computing equipment for experiments during his PhD study [43]. The authors would also like to thank anonymous reviewers for their helpful comments on earlier drafts of the manuscript.

## REFERENCES

- AMETRANO, F., AND BIANCHETTI, M. Everything you always wanted to know about multiple interest rate curve bootstrapping but were afraid to ask. SSRN, 2013. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2219548](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2219548).
- ANDERSEN, T., FUSARI, N., AND TODOROV, V. Short-term market risks implied by weekly options. *Journal of Finance* 72, 3 (2017), 1335–1386.
- AUDRINO, F., AND COLANGELO, D. Semi-parametric forecasts of the implied volatility surface using regression trees. *Statistics and Computing* 20, 4 (2010), 421–434.
- BANK OF ENGLAND. The impact of Covid on machine learning and data science in UK banking. *Quarterly Bulletin* (2020).
- BILSON, J., KANG, S., AND LUO, H. The term structure of implied dividend yields and expected returns. *Economics Letters* 128 (2015), 9–13.
- BLACK, F., AND SCHOLES, M. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 3 (1973), 637–654.
- BLISS, R., AND PANIGIRTZOGLU, N. Option-implied risk aversion estimates. *Journal of Finance* 59, 1 (2005), 407–446.
- CARR, P., AND WU, L. Stochastic skew in currency options. *Journal of Financial Economics* 86, 1 (2007), 213–247.
- CARR, P., AND WU, L. A new simple approach for constructing implied volatility surfaces. SSRN, 2010. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1701685](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1701685).
- CHOE, J., PARK, S., KIM, K., PARK, J. H., KIM, D., AND SHIM, H. Face generation for low-shot learning using generative adversarial networks. *ICCV Workshops* (2017), 1940–1948.
- COLEMAN, T., LI, Y., AND WANG, C. Stable local volatility function calibration using spline kernel. *Computational Optimization and Applications* 55, 3 (2013), 675–702.
- CONT, R., AND DA FONSECA, J. Dynamics of implied volatility surfaces. *Quantitative Finance* 2, 1 (2002), 45–60.
- CORLAY, S. B-spline techniques for volatility modeling. *Journal of Computational Finance* 19, 3 (2016), 97–135.
- ERNST AND YOUNG. The future of FinTech and financial services: what’s the next big bet? *Technical Report* (2018).
- FRIZ, P., AND GATHERAL, J. Valuation of volatility derivatives as an inverse problem. *Quantitative Finance* 5, 6 (2005), 531–542.
- GATHERAL, J. A parsimonious arbitrage-free implied volatility parameterization with application to the valuation of volatility derivatives. *Presentation at Global Derivatives & Risk Management*, Madrid, 2004.
- GATHERAL, J., AND JACQUIER, A. Arbitrage-free svi volatility surfaces. *Quantitative Finance* 14, 1 (2014), 59–71.
- GAVRISHCHAKA, V. *Econometric Analysis of Financial and Economic Time Series*. Emerald, 2006, ch. Boosting-based Frameworks in financial modeling: application to symbolic volatility forecasting, pp. 123–151.
- GOODFELLOW, I., BENGIO, Y., AND COURVILLE, A. *Deep Learning*. MIT Press, 2016.
- GULISASHVILI, A. *Analytically Tractable Stochastic Stock Price Models*. Springer, 2012.
- HESTON, S. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6, 2 (1993), 327–343.
- HOMESCU, C. Implied volatility surface: construction methodologies and characteristics. arXiv, 2011. <https://arxiv.org/abs/1107.1834>.
- ITKIN, A. To sigmoid-based functional description of the volatility smile. *The North American Journal of Economics and Finance* 31 (2015), 264–291.
- KANG, C., KANG, W., AND LEE, J. Exact simulation of the Wishart multidimensional stochastic volatility model. *Operations Research* 65, 5 (2017), 1190–1206.
- KINGMA, D., AND BA, J. Adam: a method for stochastic optimization. *ICLR* (2015), 1–13.
- KOTZÉ, A., LABUSCHAGNE, C., NAIR, M., AND PADAYACHI, N. Arbitrage-free implied volatility surfaces for options on single stock futures. *The North American Journal of Economics and Finance* 26 (2013), 380–399.
- KOU, S. A jump-diffusion model for option pricing. *Management Science* 48, 8 (2002), 1086–1101.
- KRISTJANPOLLER, W., FADIC, A., AND MINUTOLO, M. Volatility forecast using hybrid neural network models. *Expert Systems with Applications* 41, 5 (2014), 2437–2442.
- LEE, R. The moment formula for implied volatility at extreme strikes. *Mathematical Finance* 14, 3 (2004), 469–480.
- LIU, T., CUI, Y., YIN, Q., ZHANG, W., WANG, S., AND HU, G. Generating and exploiting large-scale pseudo training data for zero pronoun resolution. *ACL* (2017), 102–111.
- LUO, R., ZHANG, W., XU, X., AND WANG, J. A neural stochastic volatility model. *AAAI* (2018).
- MALLIARIS, M., AND SALCHENBERGER, L. A neural network model for estimating option prices. *Applied Intelligence* 3, 3 (1993), 193–206.
- MALLIARIS, M., AND SALCHENBERGER, L. Using neural networks to forecast the S&P 100 implied volatility. *Neurocomputing* 10, 2 (1996), 183–195.
- MERTON, R. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3 (1976), 125–144.
- QIN, Y., AND YANG, Y. What you say and how you say it matters: predicting financial risk using verbal and vocal cues. *ACL* (2019).
- REKABSAZ, N., LUPU, M., BAKLANOV, A., DÜR, A., ANDERSSON, L., AND HANBURY, A. Volatility prediction using financial disclosures sentiments with word embedding-based IR models. *ACL* (2017).
- SAMUELSON, P. Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review* (1965), 41–49.
- SHIRAYA, K., AND TAKAHASHI, A. Pricing average and spread options under local-stochastic volatility jump-diffusion models. *Mathematics of Operations Research* 4, 11 (2018).
- SUNDARESAN, S. Continuous-time methods in finance: a review and an assessment. *Journal of Finance* 13, 4 (2000), 1057–1099.
- WU, Y., HERNÁNDEZ-LOBATO, J. M., AND GHAHRAMANI, Z. Gaussian process volatility model. *NIPS* (2014).
- YANG, Y., ZHENG, Y., AND HOSPEDALES, T. Gated neural networks for option pricing: rationality by design. *AAAI* (2017).
- ZHENG, J., XIA, A., SHAO, L., WAN, T., AND QIN, Z. Stock volatility prediction based on self-attention networks with social information. *CIFER* (2019).
- ZHENG, Y. *Machine Learning and Option Implied Information*. PhD thesis, Imperial College London, 2018.