

Abstract

In modern electronic financial markets, the closing auction is critical, with substantial volume and significant impact on portfolio valuation and benchmark pricing. Traditional linear models fail to capture the complexity of this process, especially during volatile periods of market close when algorithmic and high-frequency trading dominate.

Previous approaches have shown promise in the application of Hidden Markov Models (HMMs), in conjunction with Long Short-Term Memory (LSTM) networks, to forecast daily trading in stable environments; however faced struggles with overfitting and underperformance in more complex, dynamic settings. To address these predictive challenges, particularly during the closing auction period, we extend this methodology by introducing convolutional layers to better capture the nonlinearity of the problem. By combining convolutional filters and LSTM networks, our model is tailored to forecast closing auction volumes by leveraging both temporal and latent spatial dependencies in the market.

Data Description and Engineering

The dataset, obtained from the *Optiver - Trading at the Close Challenge 2023*, contains volume data across two-hundred anonymous NASDAQ stocks on ten second timescales throughout the closing auction period. There are four-hundred-eighty unique dates present in the data, ordered chronologically, the initial 25% is used as training data and the remaining 75% is used for model validation and testing.

- Reference price, matched size, imbalance strength, weighted average price, and total volume are obtained from the existing features
- Time indexing is cyclically encoded to denote the relative time of a sample point in the closing auction time frame
- Regime state via HMM fitted on weighted average price and total volume across all stocks during each closing auction day in the training
 - We fit the HMM using the Expectation-Maximization (EM) algorithm to maximize likelihood. The number of regimes, which is nine, was selected via the Bayesian Information Criterion (BIC).

Model Framework

Our model incorporates two main components:

1. Convolutional Neural Networks (CNNs) to capture local spatial and temporal patterns in the closing auction data
2. LSTM networks to model temporal dependencies over longer horizons

Initially, a stack of 1D convolutional layers is applied to the dataset, across increasing channel depths, to extract hierarchical representations. This acts as learnable smoothing filters that capture short-term and mid-range dependencies while preserving the original time resolution of the data. The outputs are fed into inception modules that apply multiple convolution filters of varying kernel sizes in parallel, allowing the model to learn feature representations across varying scales prior to the LSTM layer. The LSTM helps retain memory of previous regime states and allows the model account for order flow leading up to the auction, capturing how these patterns evolve over time.

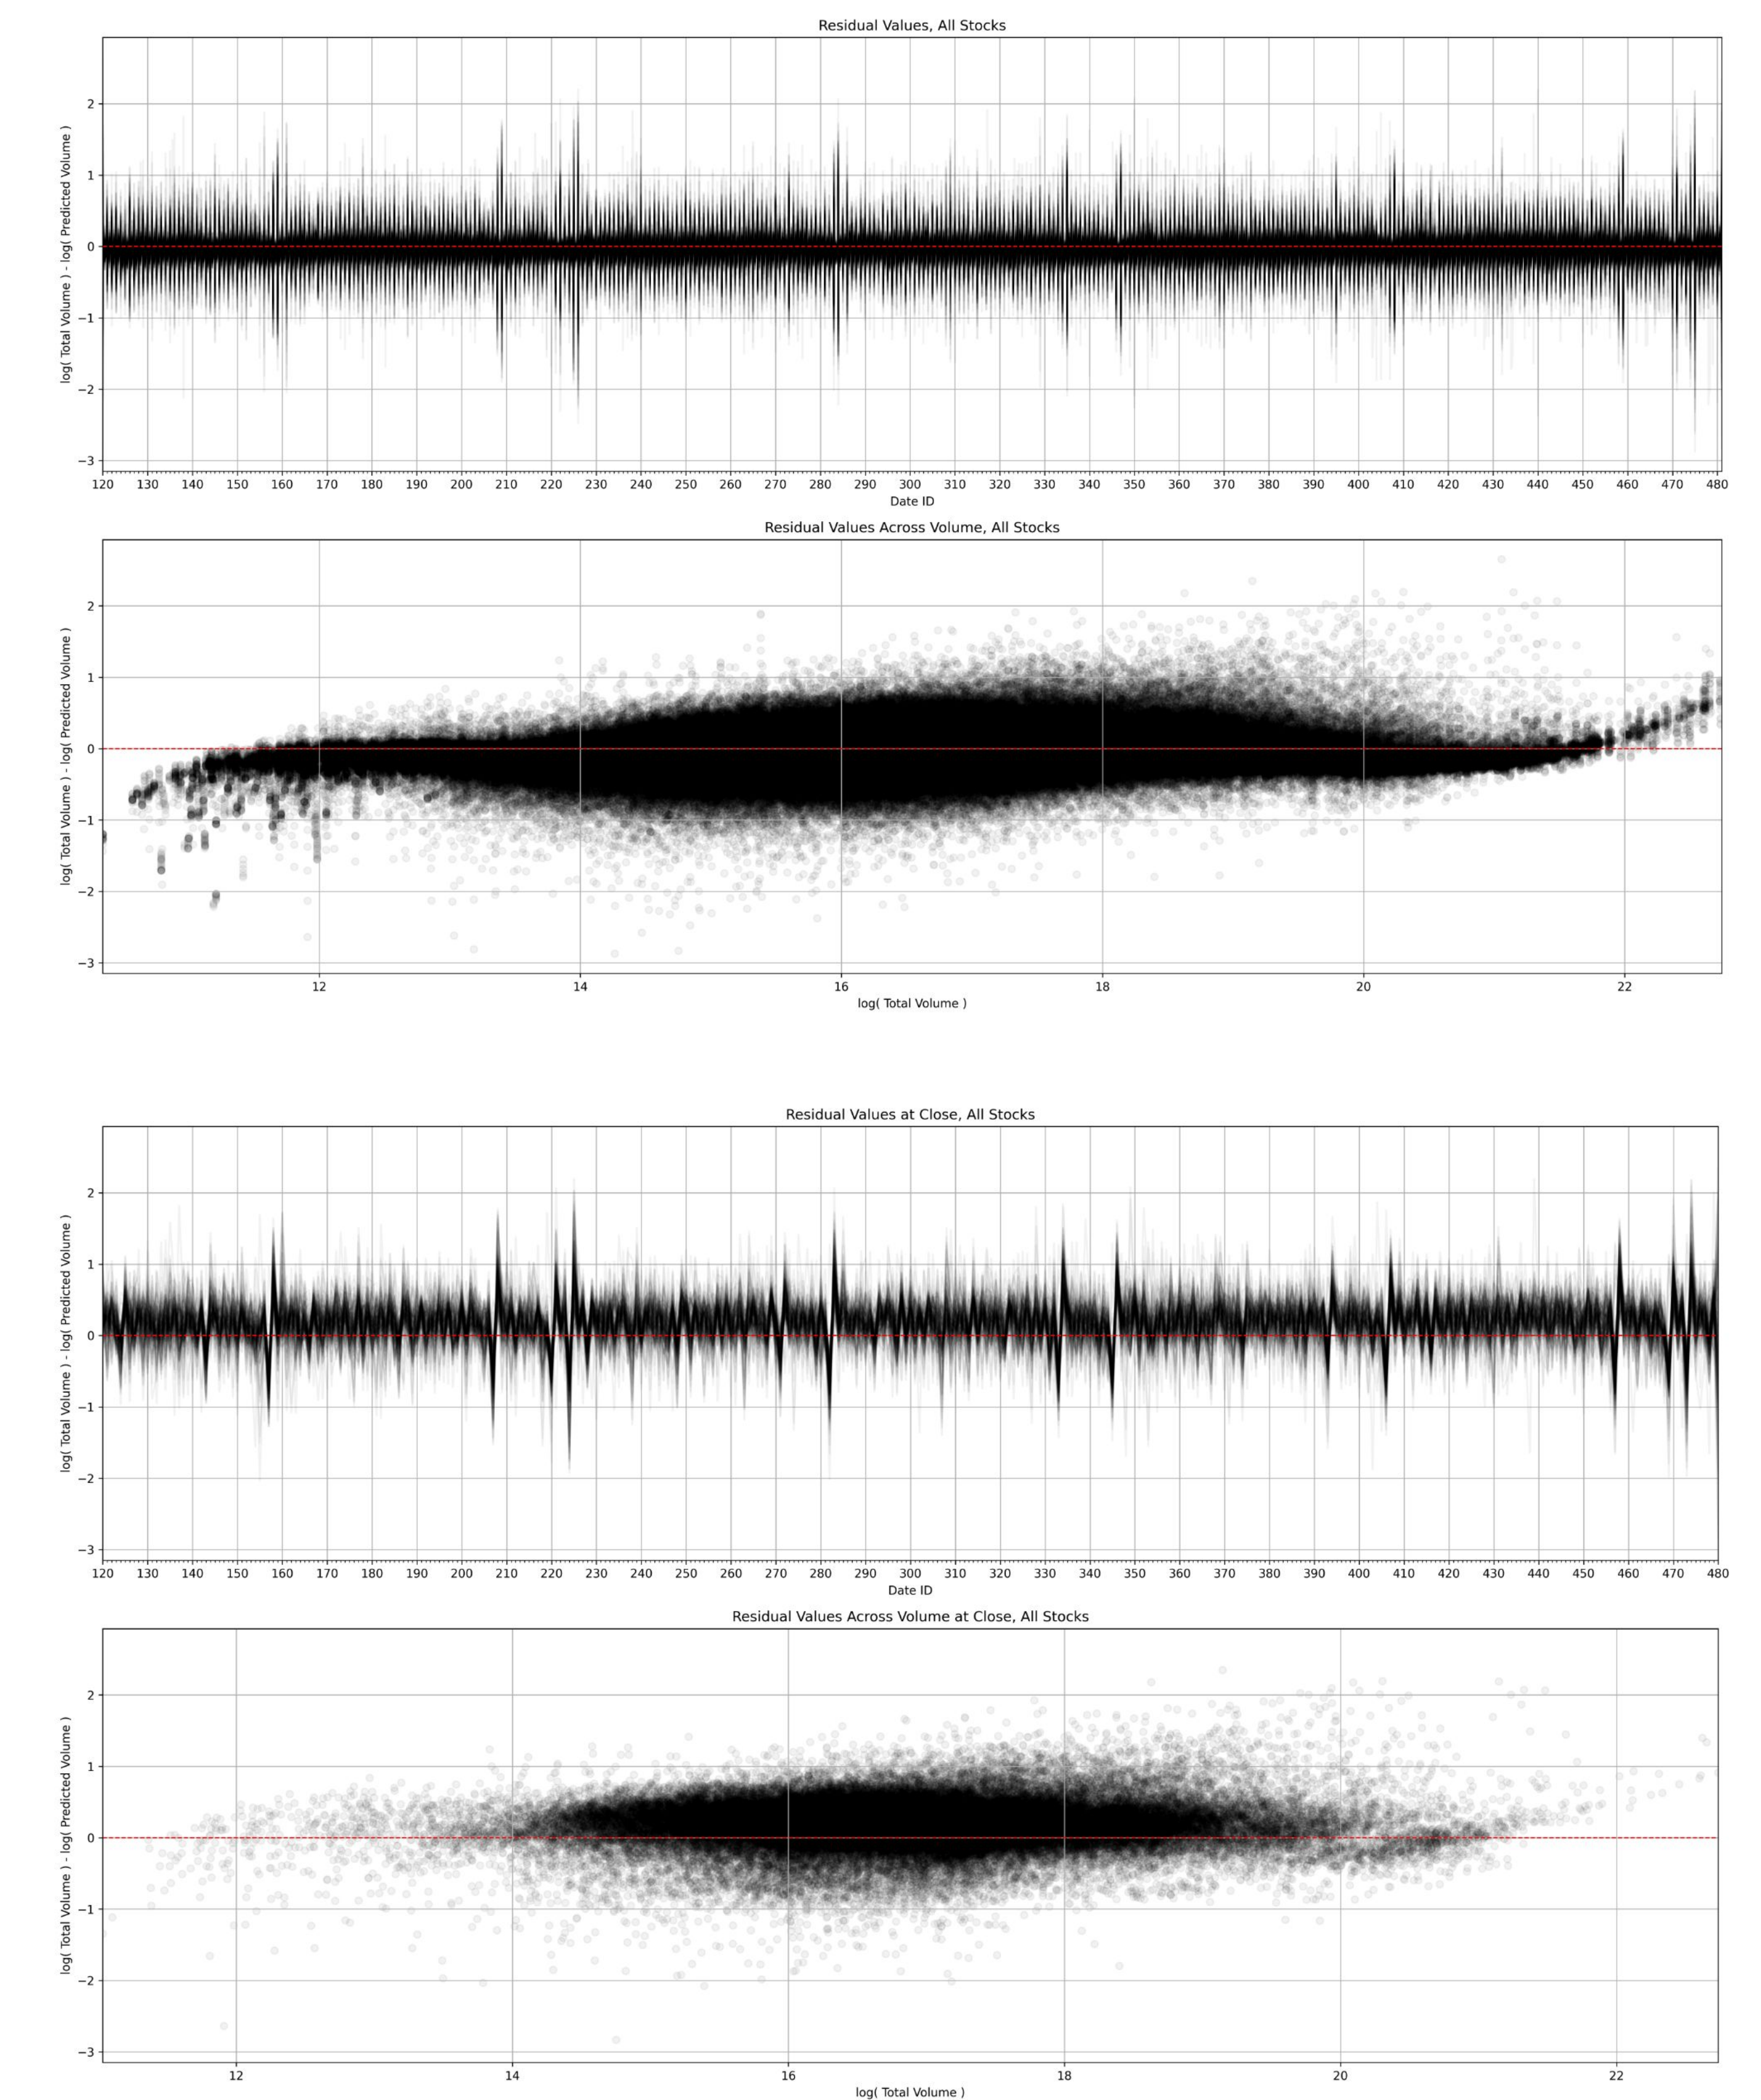
Model Performance

- Trained on the regime-informed data across twenty-five epochs. Empirically, no further gains were made using longer epochs as it trended asymptotically, and further evaluation led to overfitting
- The Mean Squared Error (MSE) loss function is used as it penalizes larger errors more heavily, encouraging the model to focus on minimizing large deviations and promoting stable convergence
- An Adam optimizer is chosen for faster convergence and adaptive learning rate due to the high dimensionality of the problem



Despite the low training percentage, our model is able to achieve great out-of-sample performance with minimal overfitting, demonstrating a high level of predictive power. There is minimal difference between that and the median volume approximation and the best-case volume approximation from an MSE perspective. Reflective of the MSE distribution across all stocks having a heavy positive skew, with the worst-case volume approximation still performing comparatively well, as residuals typically remain within an order of magnitude or less.

Residual Error Analysis



- Across all times, the density of the volume residual error remains relatively bounded within plus or minus 0.2 on a log-scale
- At lower and higher target volumes compared to the norm, the model starts to break down, increasing residual values
- At the close, the model under-predicts the volume by an average of 0.3, with higher variance compared to results at all other times

Conclusion

The inclusion of high-frequency features, combined with the utilization of convolutional layers to capture cross-sectional commonalities of stocks, are highly effective at capturing the volatility of the closing auction period. The convolution filters used in our CNN architecture, while increasing the computational complexity of the problem, capture the nonlinear nature of volume evolution accurately and can scale to provide a granular modeling of trading behavior.