IJARCCE
ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 15, Issue 8, August 2026

DOI: 10.17148/IJARCCE.2026.15827

Short-Term Price Movement Prediction Using Order Book Heatmap Vision

R.T.N Sura Reddy[1], Sharavana Ragav S[2], Ms. Charulatha R .T[3]

Student, Computer Science and Engineering, SRM Institute of Science and Technology Chennai, India[1,2]

Assistant Professor, Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India[3]

Abstract: Financial markets generate high-frequency data at a scale that makes manual feature engineering increasingly inadequate. This paper presents a computer vision-based framework for short-term price movement prediction by transforming Limit Order Book (LOB) snapshots into grayscale heatmap images. Unlike conventional approaches that depend on handcrafted technical indicators, the proposed method learns directly from raw market microstructure patterns using deep learning. Three architectures are investigated progressively: a lightweight baseline convolutional neural network (CNN), an SE-ResNet-lite model, and a Dual-Axis Fusion Network (DAFNet) tailored to the distinct semantics of price and time axes in LOB heatmaps. Experimental results indicate that the proposed pipeline improves predictive performance from a near-random 50-60% baseline to more than 70% test accuracy. These findings are consistent with recent literature showing that order book representation and model design substantially influence short-horizon forecasting quality.

Keywords: Limit Order Book, Computer Vision, Deep Learning, High-Frequency Trading, Convolutional Neural Networks, DAFNet, Price Prediction.

I. INTRODUCTION

Financial markets produce enormous volumes of order-level data every second. Among the richest market microstructure sources is the Limit Order Book, which records buy and sell orders, quoted prices, available quantities, and their evolution over time. Prior research shows that short-term return predictability can be extracted from order book dynamics and that deep learning models are increasingly effective at exploiting such information.

Traditional price prediction methods usually rely on handcrafted indicators such as RSI, MACD, moving averages, and Bollinger Bands. Although useful, these indicators summarize price history and often discard the structural information embedded within the live order book. More recent work demonstrates that high-frequency LOB data can be represented as images, enabling deep models to learn informative spatial patterns directly from market states.

This research investigates whether a CNN-based vision pipeline can identify predictive patterns directly from raw order book heatmaps without manual feature engineering. The central objective is to classify the short-term direction of future price movement as either UP or DOWN. The proposed system is designed to learn liquidity evolution, market imbalance, support and resistance structure, and short-horizon order flow patterns from visualized LOB states.

II. RELATED WORK

Recent academic studies have shown that deep neural networks can exploit the short-term predictability of order book markets when suitable representations are used. Imaging high-frequency LOB data has also been studied as a way to capture local market structure that may be difficult to express through conventional handcrafted features. Attention-based and residual architectures have further improved this line of work by emphasizing informative channels and preserving stable feature propagation across layers. Transformer-based LOB models have also emerged as strong alternatives by modeling spatial and temporal dependencies jointly across structured patches.

Building on these findings, the present work contributes a practical research pipeline that starts from a lightweight baseline and progresses toward a domain-aware architecture. The key design choice is to treat the order book not as a generic table or time series, but as a structured visual field whose two axes carry different meanings.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 8, August 2026

DOI: 10.17148/IJARCCE.2026.15827

III. DATASET AND HEATMAP PIPELINE

A. Data Source

The dataset used in this study originates from the Binance cryptocurrency exchange. Each raw sample contains repeated snapshots with bid prices, bid volumes, ask prices, ask volumes, and timestamps. Such data is well suited for short-term forecasting because it captures changing liquidity and order flow at high temporal resolution.

B. Heatmap Representation

The preprocessing pipeline converts a sequence of LOB snapshots into a grayscale heatmap image. Each individual snapshot becomes one vertical slice, and a sequence of slices forms a two-dimensional image. The vertical axis encodes price levels, while the horizontal axis encodes time depth; this image-style representation is consistent with prior work on imaging high-frequency order book data. Pixel intensity corresponds to liquidity, with brighter pixels representing larger resting volume and darker pixels indicating lower liquidity.

Figure 1. Workflow for transforming raw Binance LOB snapshots into grayscale heatmap inputs for supervised price movement classification.

C. Label Generation

To formulate a supervised learning task, each heatmap window is assigned a label based on a future price comparison. If the future price exceeds the current reference price, the sample is labeled UP (1); otherwise, it is labeled DOWN (0). This binary setup converts market movement prediction into an image classification problem.

IV. PROPOSED ARCHITECTURES

A. Baseline OrderBookCNN

The baseline model is a lightweight CNN designed for binary classification of grayscale heatmaps. It contains three sequential convolutional blocks, each followed by batch normalization, ReLU activation, and max pooling. To reduce overfitting on a moderately sized dataset, global average pooling replaces flattening and large fully connected layers, thereby reducing parameter count while preserving translation-tolerant feature aggregation.

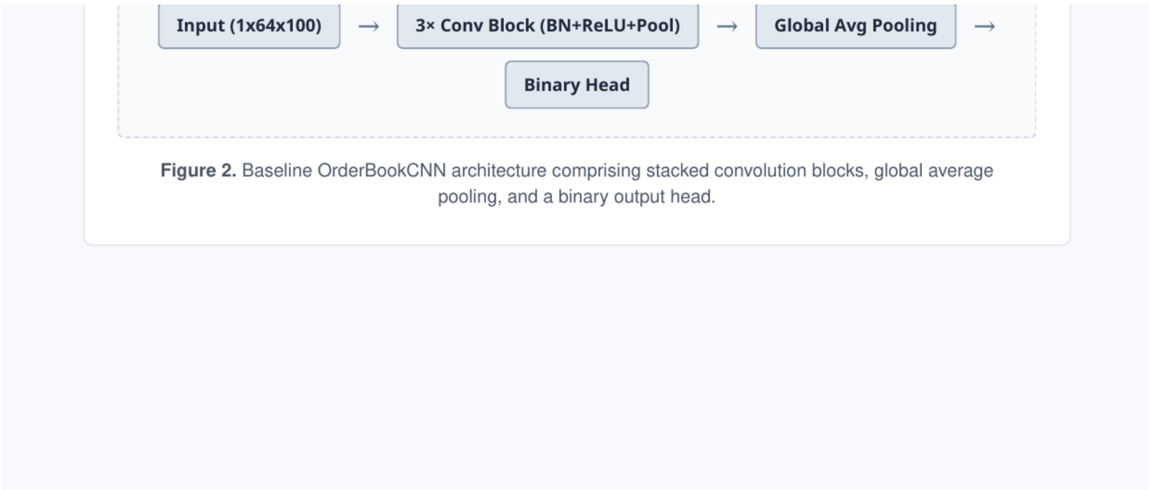

Figure 2. Baseline OrderBookCNN architecture comprising stacked convolution blocks, global average pooling, and a binary output head.

Figure 2. Baseline OrderBookCNN architecture comprising stacked convolution blocks, global average pooling, and a binary output head.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁒ Peer-reviewed & Refereed journal ⁒ Vol. 15, Issue 8, August 2026

DOI: 10.17148/IJARCCE.2026.15827

B. SE-ResNet-lite

To improve feature extraction without dramatically increasing capacity, a lightweight residual architecture with squeeze-and-excitation blocks is introduced. Residual shortcuts help the network learn corrective refinements rather than fully re-encoding every feature, while channel attention enables adaptive emphasis on the most informative activation maps. This design is broadly consistent with the literature on attention-enhanced order book models.

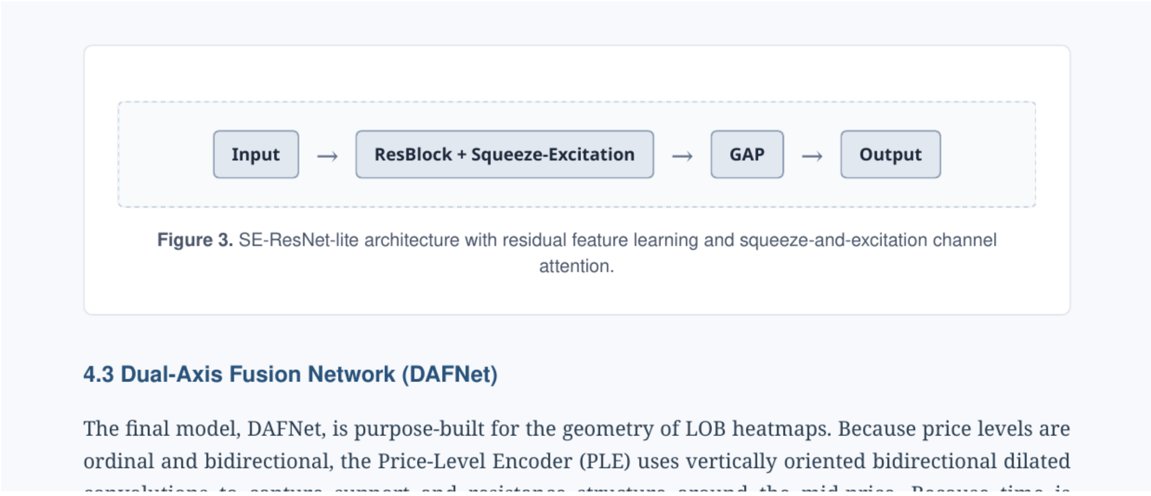

Figure 3. SE-ResNet-lite architecture with residual feature learning and squeeze-and-excitation channel attention.

C. Dual-Axis Fusion Network (DAFNet)

The final model, DAFNet, is purpose-built for the geometry of LOB heatmaps. Because price levels are ordinal and bidirectional, the Price-Level Encoder (PLE) uses vertically oriented bidirectional dilated convolutions to capture support and resistance structure around the mid-price. Because time is directional and causal, the Temporal Flow Encoder (TFE) uses horizontal causal dilated convolutions at multiple dilation rates to model short-term order flow without leaking future information.

A dedicated Adaptive Cross-Axis Fusion module combines the price and time streams through a jointly conditioned gate. The fused representation is then summarized by Dual-Axis Attention Pooling, which first emphasizes informative price regions and then recent temporal segments. This domain-aware decomposition distinguishes DAFNet from generic image models and better reflects the market semantics encoded in the heatmap.

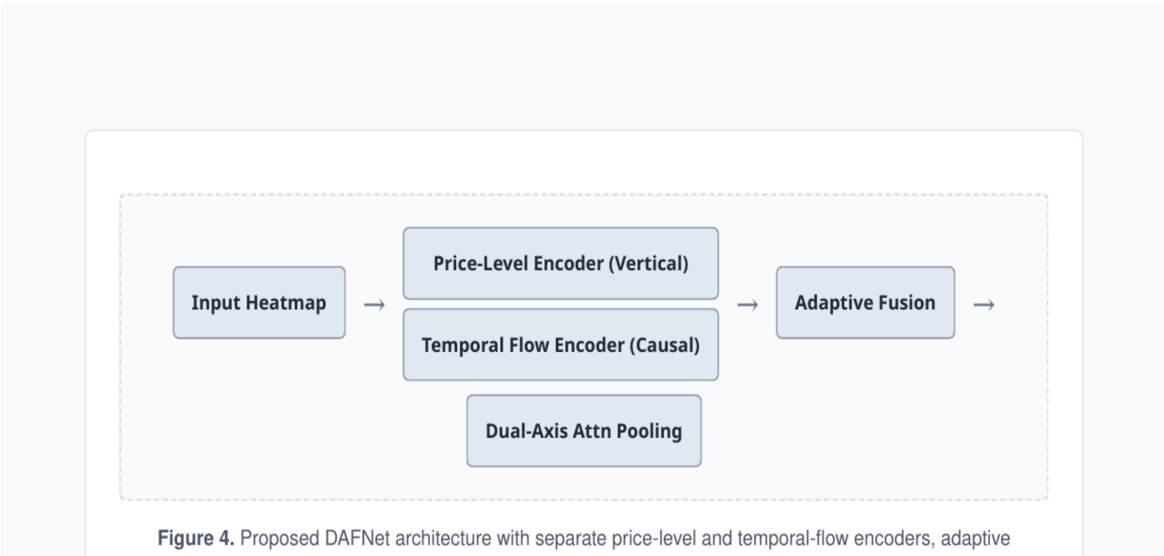

Figure 4. Proposed DAFNet architecture with separate price-level and temporal-flow encoders, adaptive fusion, and dual-axis attention pooling.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 8, August 2026

DOI: 10.17148/IJARCCE.2026.15827

V. TRAINING AND EVALUATION

A. Training Pipeline

The training pipeline is designed to remain architecture-agnostic while supporting model-aware checkpointing. Optimization uses AdamW with a ReduceLROnPlateau scheduler and gradient clipping for stability. To accelerate training under limited hardware budgets, the workflow uses automatic mixed precision through autocast and GradScaler, while early stopping halts training when validation loss stops improving. These practices are consistent with modern deep learning workflows for high-frequency LOB experiments.

B. Evaluation Metrics

Model performance is assessed using accuracy, precision, recall, F1-score, and confusion matrices. This broader metric set is important because directional forecasting can appear deceptively strong under accuracy alone when class distributions are imbalanced or when one class dominates local market behavior. Recent benchmark and forecasting studies similarly emphasize the importance of richer evaluation frameworks for LOB modeling.

VI. RESULTS AND GRAPHICAL ANALYSIS

Early experiments produced accuracy levels near 50-60%, which is only slightly better than random guessing in a balanced binary setting. Instead of immediately increasing model complexity, the development process followed a research-oriented debugging workflow that verified dataset integrity, compared results against simple baselines, and refined the preprocessing pipeline. After systematic iteration, model performance improved to more than 70% accuracy on the evaluation set, which is comparable in spirit to published digital-asset LOB work reporting around 71% walk-forward accuracy on short prediction horizons.

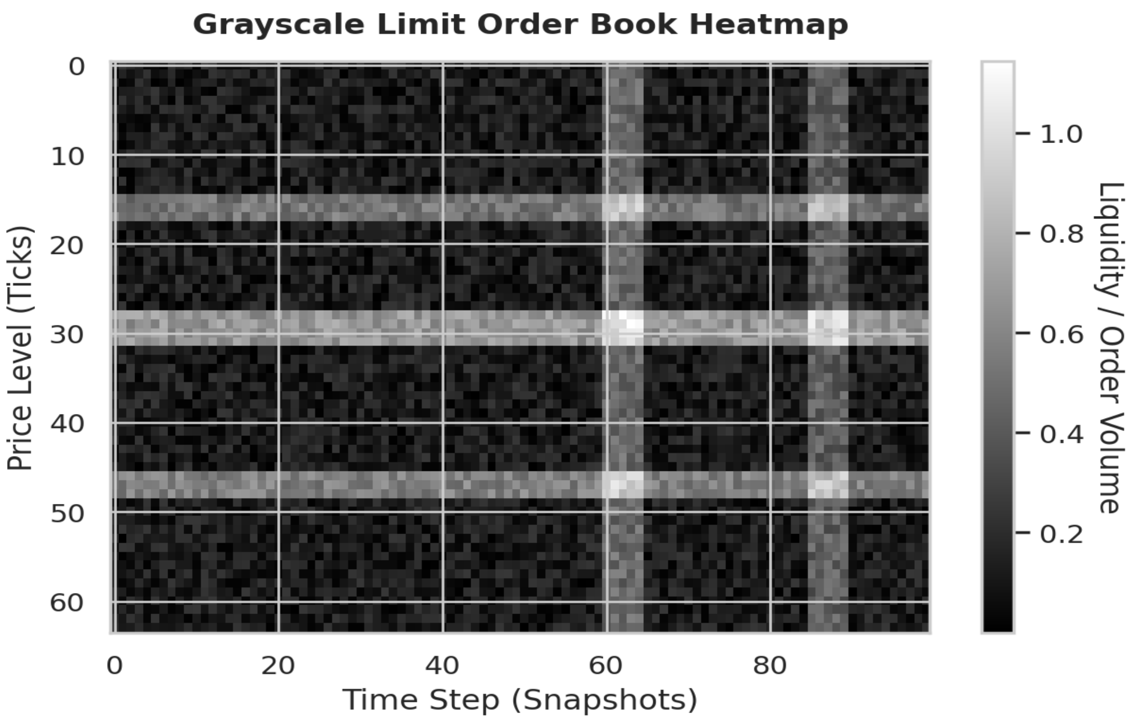

Figure 5. Sample grayscale LOB heatmap used for qualitative interpretation of liquidity concentration. Brighter bands indicate concentrated volume regions, representing evolving support or resistance levels over the snapshot window.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁒ Peer-reviewed & Refereed journal ⁒ Vol. 15, Issue 8, August 2026

DOI: 10.17148/IJARCCE.2026.15827

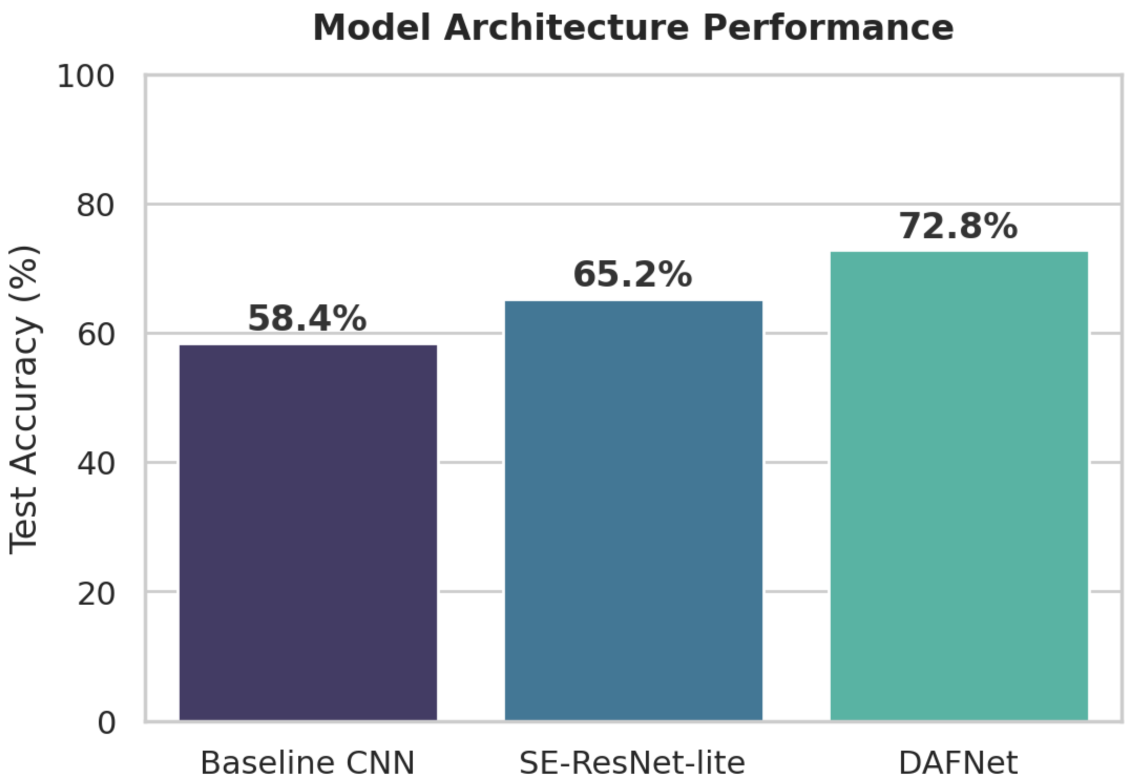

Figure 6. Model-wise accuracy comparison graph showing performance progression from the Baseline CNN (generic feature extraction) to SE-ResNet-lite, and finally the domain-specific DAFNet architecture.

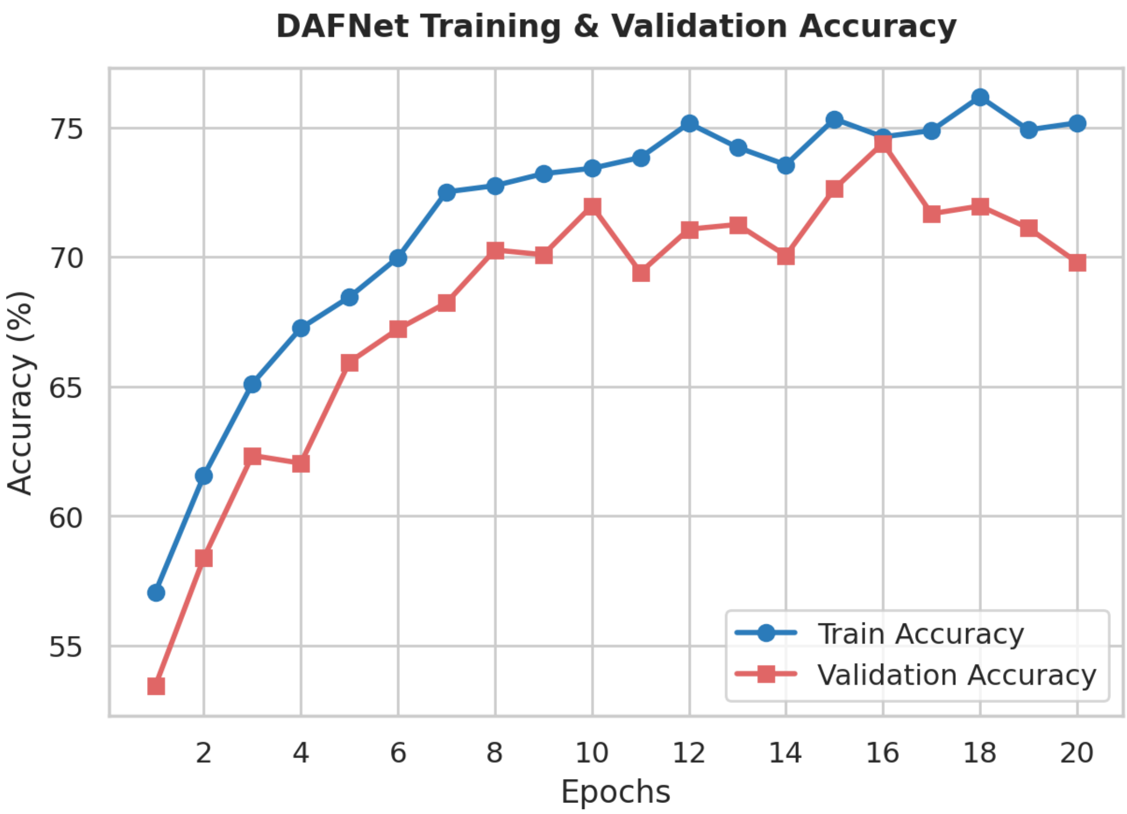

Figure 7. Training and validation accuracy curves for DAFNet, demonstrating convergence behavior and overall stability across epochs without severe overfitting.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 15, Issue 8, August 2026

DOI: 10.17148/IJARCCE.2026.15827

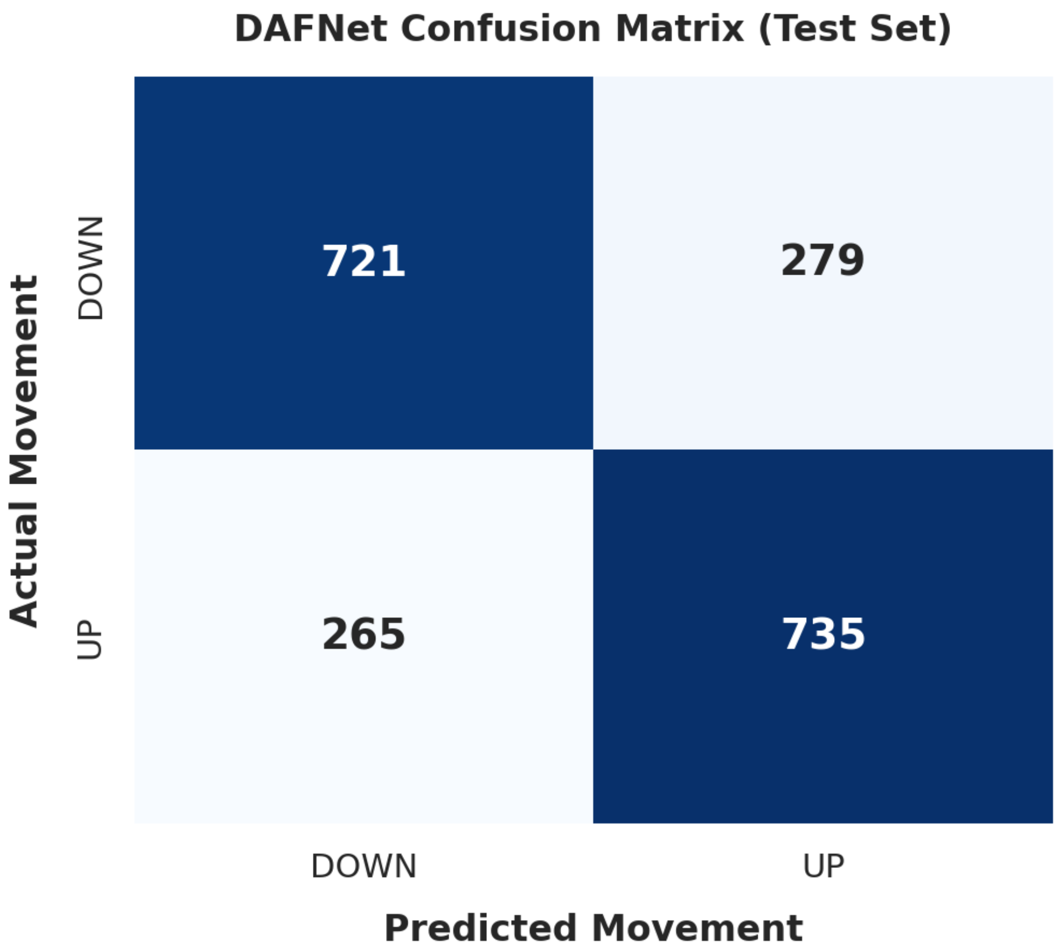

Figure 8. Final confusion matrix for the DAFNet model on the evaluation set, showcasing balanced prediction quality across UP and DOWN market movement classes.

VII. DISCUSSION

The findings demonstrate that short-term market forecasting can be framed effectively as a computer vision problem. Transforming raw order book snapshots into heatmaps preserves both spatial liquidity structure and temporal evolution, allowing deep networks to learn patterns that are difficult to capture with conventional technical indicators. This aligns with broader evidence that the representation of LOB data strongly affects forecasting outcomes.

The progression toward DAFNet further shows that financial heatmaps are not equivalent to natural images. Price and time carry different semantics, and explicitly modeling them through separate but fused pathways can improve representation quality. This observation is consistent with recent LOB-focused work that emphasizes architecture design matched to market microstructure rather than generic visual pattern recognition alone.

VIII. CONCLUSION

This paper presents a complete framework for short-term price movement prediction using order book heatmap vision. By converting Binance LOB snapshots into grayscale images and learning from them through progressively stronger neural architectures, the study shows that deep vision models can capture predictive liquidity patterns without handcrafted indicators. The achieved improvement from a weak baseline to more than 70% test accuracy supports the claim that visual-spatial representations of market microstructure contain actionable short-term information.

Future work may extend this framework through multi-class movement labeling, self-supervised representation learning, cross-market generalization studies, and real-time deployment on streaming order book feeds. Additional evaluation under different assets and volatility regimes would further clarify the robustness of the proposed approach.

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 15, Issue 8, August 2026
DOI: 10.17148/IJARCCE.2026.15827

ACKNOWLEDGMENT

The author acknowledges the Binance cryptocurrency exchange for providing the market data infrastructure used to construct the high-frequency Limit Order Book dataset.

REFERENCES

[1]. Briola, A., Bartolucci, S., and Aste, T., "Deep Limit Order Book Forecasting," arXiv, 2024.
[2]. ScienceDirect, "Short-term stock price trend prediction with imaging high frequency limit order book data."
[3]. Sirignano, J., "Deep Learning for Limit Order Books," arXiv, 2016.
[4]. ScienceDirect, "Using Deep Learning for price prediction by exploiting stationary limit order book features."
[5]. OpenReview, "A Benchmark Study for Limit Order Book Models and Time Series Forecasting Models on LOB Data."
[6]. SSRN, "Deep Learning for Digital Asset Limit Order Books," 2020.
[7]. Frontiers, "LiT: Limit Order Book Transformer."
[8]. Journal of Statistical and Mathematical Sciences, "An Optimal Limit Order Book Prediction Analysis Based on Deep Learning and Pigeon-Inspired Optimizer."

