YOLO-Based Detection of Buy and Sell Signals in Candlestick Charts with Moving Averages

Marco Antonio Sousa Santos[1] **Adriano Rivolli da Silva**[1] **André Roberto Ortoncelli**[1]

[1]Universidade Tecnológica Federal do Paraná (UTFPR)
Programa de Pós-Graduação em Informática

marcoantoniosantos@alunos.utfpr.edu.br

{rivolli,ortoncelli}@utfpr.edu.br

Abstract. *In stock trading, identifying trends is essential. Candlestick charts show price variations while moving averages help smooth these fluctuations and highlight trends. Some works focus on event detection in candlestick charts but without using moving averages. In this study, we investigate whether incorporating semantic information from moving averages into candlestick chart analysis can enhance the detection of buy and sell signals. The proposed approach includes image generation, data augmentation, labeling, and signal identification in short-term trading scenarios. We conducted experiments using four versions of YOLO (v3, v8, v9, and v11). Compared to previous work that did not use moving averages, our results were significantly better regarding F1-Score (up to 0.06 higher) and recall (up to 0.18 higher).*

1. Introduction

In recent years, Computer Vision (CV) has revolutionized various scientific fields by enabling automated visual data analysis. Deep Learning (DL)-based methods have shown excellent performance in audio, image, and video processing [Birogul et al. 2020]. Machine Learning (ML) has been adopted to examine the complex and nonlinear dynamics of the financial market, including asset pricing, derivative pricing, volatility forecasting, index tracking, portfolio selection, and price trend prediction [Zhang et al. 2023].

Regarding price trends, recent studies have explored using CV techniques to identify patterns in candlestick charts [Chen and Tsai 2020]. These charts display a financial asset's high, low, open, and close prices within a specific time frame. These visual patterns in Candlestick charts encapsulate valuable information about market sentiment and price movement. The world of candlestick technical analysis encompasses 103 recognized patterns, many of which are extensively cataloged and studied in depth in prominent works like Thomas Bulkowski's "Encyclopedia of Candlestick Charts" [Thomas 2012]. These formations serve as invaluable tools that traders leverage to pinpoint potential market entry and exit points. By leveraging pattern classification approaches, it becomes possible to automate the recognition of these patterns, reducing the need for manual analysis and enhancing decision-making in trading strategies.

Recent works have presented the use of DL in financial market charts, such as the use of Hybrid transformer-CNN architecture to predict eight patterns in candlestick charts, with an average accuracy of around 90% [Chen and Tsai 2020]. The method proposed by [Sezer and Ozbayoglu 2020] combines time series analysis with deep Convolutional

Neural Networks (CNNs), successfully outperforming a traditional buy-and-hold strategy by identifying hold, buy, and sell patterns in Dow Jones stock charts.

The YOLO (You Only Look Once) detector has also been applied to identify buy and sell decisions in candlestick charts from the Istanbul Stock Exchange between the years 2000 and 2018, achieving a prediction accuracy of 85%, which resulted in a total profit of 100% [Birogul et al. 2020]. In another related work, a CNN was used to predict price trends in charts, using attention mechanisms to focus on specific areas of the input images that are most relevant to prices, achieving performance with the analysis of chart images equivalent to those obtained with time series [Zhang et al. 2023]. Finally, we can highlight a study that customized the YOLOv8 model and achieved a mean Average Precision (mAP) of approximately 86% at an Intersection over Union (IoU) threshold of 0.50, in identifying four patterns used trend reversal formations observed in candlestick charts: Head and Shoulders, Reverse Head and Shoulders, Double Top, and Double Bottom [Thakur et al. 2024]

In this work, we also focus on identifying buy and sell signals by analyzing trends and patterns in candlestick charts. Although Temur et al. [2024] also address the detection of buy and sell signals, their approach relies exclusively on images of candlestick charts, without incorporating trend indicators. In contrast, studies such as Chen and Tsai [2020] and Birogul et al. [2020] focus exclusively on recognizing candlestick patterns, without establishing a direct connection with trading signals. Unlike these previous works, we integrate moving averages into candlestick charts to enrich the semantic content analyzed by the detection model.

To conduct the experiments, a dataset was created using MetaTrader 5 software[1], consisting of candlestick chart images with moving averages from NASDAQ stocks across three sectors (technology, communication, and consumer discretionary). A total of 519 distinct images were generated. The selected stocks are characterized by high volatility, a term that refers to frequent and significant price fluctuations, which are common in these sectors due to their sensitivity to innovation cycles, market sentiment, and economic trends. This high volatility implies greater investment risk but also creates opportunities for identifying buy and sell signals based on technical analysis. Through a process of data augmentation, the dataset was expanded to 1,356 images.

Our experimental results demonstrated notable improvements in recall when compared to related works. As an example, the recent work of [Temur et al. 2024], which also focused on identifying buy and sell signals, reported a recall of 0.69 using YOLOv3.

Our experimental results demonstrated notable improvements in recall when compared to related works. For example, the recent work of [Temur et al. 2024], which also focused on identifying buy and sell signals, reported a recall of up to 0.69. Our proposed approach achieved a relative improvement of over 18%, reaching a recall of 0.82 of up to 0.87 with YOLOv8. It is noteworthy that the databases used in the two works are distinct, so the works cannot be directly compared, but even so, the results obtained in this study indicate that the approach is promising, being feasible to have specific models trained for application in Candlestick Charts with Moving Averages.

The remainder of this work is organized as follows. Section 2 presents theoretical

[1]https://www.metatrader5.com/

concepts necessary for understanding this work. Section 3 describes the methodology used to produce the database and conduct and evaluate the experiments. The experimental results are in Section 4. Finally, Section 5 presents the conclusion and future work.

2. Theoretical Aspects

This section outlines the theoretical foundations essential for understanding this study. Subsection 2.1 discusses pattern analysis in candlestick charts and the use of moving averages in technical trading. Subsection 2.2 presents key concepts related to object detection in images using the YOLO architecture.

2.1. Candlestick and moving averages pattern analysis

Candlestick pattern analysis is one of the oldest and most visually intuitive methods of technical analysis, originating in 18th-century Japan with Munehisa Homma. This technique graphically represents the open, high, low, and close prices of assets, allowing traders to identify formations that reflect market sentiment and indicate potential trend reversals or continuations. Importantly, candlestick analysis is versatile and can be applied across various timeframes—including daily, weekly, monthly, or hourly intervals—making it suitable for both short-term and long-term trading strategies.

Several academic studies have examined the effectiveness of candlestick patterns. For instance, [Marshall et al. 2006] analyzed the applicability of such strategies in U.S. equity markets and concluded that, on average, candlestick patterns do not consistently generate abnormal returns, suggesting a high level of market efficiency in this context. However, other research has identified specific conditions where these patterns may provide value. [Thammakesorn and Sornil 2019] proposed a technique that generates trading strategies based on candlestick characteristics using the Chi-square Automatic Interaction Detector (CHAID) algorithm, demonstrating that such strategies can outperform popular indicators like the Moving Average Convergence Divergence (MACD) and the Relative Strength Index (RSI) in specific scenarios.

The formal definition of candlestick patterns has also been the subject of investigation. [Hu et al. 2019] developed first-order logic specifications for 103 known patterns to establish an unambiguous reference model, which can support future research in pattern classification and detection.

In the context of cryptocurrencies, [Cohen 2021] explored the optimization of candlestick-based strategies for Bitcoin trading systems, highlighting that the effectiveness of these patterns may vary significantly depending on the specific characteristics of the asset being analyzed. To better understand how these strategies are constructed, it is essential to revisit a candlestick's basic structure and interpretation, as illustrated in Figure 1.

Each candlestick has four prices: open, close, high, and low. The candle's body represents the range between the open and close prices, while the wicks (or shadows) show the high and low for the period. The structure of the candlestick body indicates price direction: a bullish candle typically has a hollow (unfilled) body, where the closing price is higher than the opening price, while a bearish candle has a filled (solid) body, indicating that the closing price is lower than the opening price.

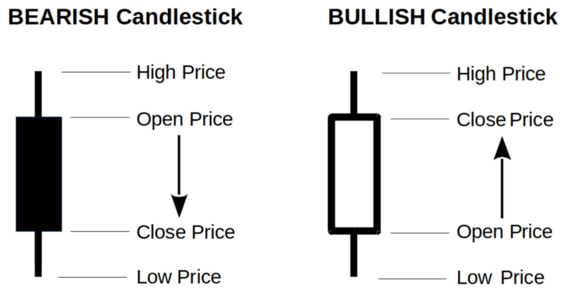

Figure 1. Bullish and bearish candlesticks

Candlestick patterns are formations that indicate potential trend reversals or continuations. Table 1 describes 11 common candlestick patterns illustrated in Figure 2.

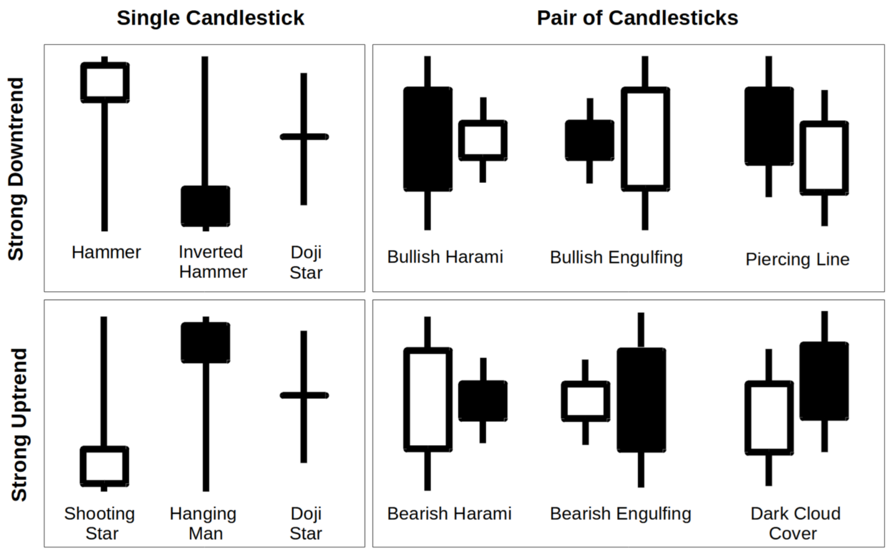

Figure 2. Common candlestick patterns

Moving averages are widely recognized as trend-following indicators that help smooth out price data over time, making it easier to interpret the market's general direction. A short-term moving average, such as a 10-period average, reacts quickly to price fluctuations and is particularly useful for identifying recent trends. In contrast, a longer-term moving average, like the 20-period, responds more slowly and is better suited for detecting broader, more established trends.

One of the most commonly used signals involving moving averages is the crossover. A bullish crossover, often called a ”Golden Cross,” occurs when the short-term moving average crosses above the long-term average, typically indicating the beginning of an uptrend. On the other hand, a bearish crossover, known as the ”Death Cross,” takes place when the short-term average crosses below the long-term average, suggesting the potential start of a downtrend.

Integrating candlestick patterns with moving averages improves the analytical depth and reliability of trading strategies. For instance, a bullish reversal pattern—such as a hammer or a bullish engulfing formation—gains strength when it emerges near a

Table 1. Descriptions of common candlestick patterns

Pattern	Description
Hammer	A small body at the top with a long lower shadow, found after a downtrend, signaling a potential bullish reversal.
Shooting Star	A small body with a long upper shadow near the lower end, found after an uptrend, suggesting a bearish reversal.
Inverted Hammer	A small body and long upper shadow, found at the bottom of a downtrend, indicating a potential bullish reversal.
Hanging Man	Resembles a hammer in appearance but forms at the top of an uptrend, signaling a possible bearish reversal.
Doji Star	A candle with little or no real body, where opening and closing prices are nearly equal. Reflects market indecision and, when following a strong trend, it may suggest an impending reversal.
Bullish Harami	Consists of a small bullish candle that fits entirely within the previous bearish candle, suggesting a possible bullish reversal.
Bearish Harami	The inverse of the Bullish Harami, which may indicate a bearish reversal.
Bullish Engulfing	A strong bullish candle that completely engulfs the previous bearish one, suggesting a potential upward reversal.
Bearish Engulfing	A bearish candle that fully engulfs the prior bullish candle, indicating a potential downward reversal.
Piercing Line	A bullish candle forms after a downtrend, opening below the previous low and closing above the midpoint of the preceding bearish candle. This formation signals a possible bullish reversal.
Dark Cloud Cover	A bearish candle appears after an uptrend. It features a bearish candle that opens above the prior high but closes below the midpoint of the previous bullish candle, potentially indicating a bearish reversal.

support level defined by a moving average, particularly when accompanied by a bullish crossover. Similarly, bearish patterns like the shooting star or dark cloud cover become more convincing if they appear after a bearish crossover, as this alignment supports the likelihood of a continuing downtrend.

Understanding that these patterns are most effective when used within clearly defined trends is crucial. The signals they generate in strong upward or downward markets tend to be more reliable. However, during periods of consolidation, when the market lacks a clear direction and moves sideways, candlestick patterns often lose effectiveness and may result in false signals or whipsaws.

The synergy created by combining visual price patterns with trend-based indicators like moving averages significantly reduces the likelihood of misleading signals. This integrated approach boosts confidence in trade entries and exits. Experienced traders and investors often rely on this method to refine their strategies and enhance decision-making in the stock market.

2.2. Object detection in images with YOLO

Object detection in images determines the category and location of objects of interest. CV methods for object detection have evolved driven by the increase in computational power and the progress of DL algorithms [Koirala et al. 2019], and can be divided into two periods: traditional object detection before 2014 and the DL-based object detection period after 2014, which explores CNN-based methods [Zou et al. 2023].

DL-based methods have achieved state-of-the-art performance in various object detection tasks. Among these methods, YOLO is one of the most prominent algorithms, with different applications in agriculture [Neto et al. 2023], livestock [Leal et al. 2024], and people monitoring [Pires et al. 2023], among others [Zhao et al. 2019].

Unlike traditional approaches, YOLO frames object detection as a single regression problem, predicting the spatial coordinates of bounding boxes and the associated class probabilities directly from the input image in a single forward pass of a neural network. The YOLO's entire detection pipeline is encapsulated in a single network, which can be optimized end-to-end based on detection accuracy. YOLO divides the input image into an $S \times S$ grid, where each grid cell detects objects whose centers fall within it. For each grid cell, the model predicts confidence scores for B bounding boxes, indicating both the likelihood that a bounding box contains an object and the accuracy of the predicted location.

The first version of YOLO (YOLOv1) resizes an image to 448 × 448 pixels as input to a single CNN, directly predicting bounding boxes and classes per grid cell [Redmon et al. 2016]. After the first version of YOLO, several new versions were proposed, generating a family of algorithms that grew significantly by releasing several algorithms that modified the architecture, seeking better performance in different situations. Among the latest versions of YOLO, versions v8, v9, and v11—used in this study—can be highlighted. YOLOv3 was also applied in the experiments, as it is the same version used by [Temur et al. 2024] for detecting patterns in candlestick charts.

YOLOv8 marks an important architectural evolution within the YOLO family. It introduces a modular, anchor-free design with a decoupled head for classification and localization tasks, enhancing training stability and inference speed. YOLOv8 also supports model scaling across multiple deployment contexts, from edge devices to cloud infrastructures [Ultralytics 2023].

Building upon the YOLOv8 foundation, YOLOv9 integrates transformer-based components and dynamic label assignment strategies to improve detection in visually complex environments. The inclusion of attention mechanisms allows YOLOv9 to better capture spatial dependencies, which is particularly useful in scenarios involving dense object layouts [Wang et al. 2023].

The most recent iteration, YOLOv11, explores advanced techniques such as multi-scale feature fusion and dynamic anchor mechanisms, aiming to further enhance detection performance in challenging visual conditions. YOLOv11 combines CNNs with hybrid transformer layers, improving both the robustness and interpretability of the model in specialized domains like financial chart analysis [Chen et al. 2024].

3. Methodology

This Section presents the experimental methodology, describing the process of producing the experimental database (Subsection 3.1), the labeling strategy (Subsection 3.2), the experimental setup used (Subsection 3.3), and the metrics used to evaluate the results (Subsection 3.4).

3.1. Dataset Generation

The experimental dataset comprises 519 images of highly volatile stocks, each depicting candlestick chart patterns combined with the crossover of two moving averages: a short-term average (10 periods), represented by a blue line, and a long-term average (20 periods), represented by a red line. These period settings are commonly used in technical analysis and are applicable across multiple timeframes, including daily, weekly, monthly, and hourly. All images were manually annotated using the LabelImg tool[2], with buy or sell signals identified according to established technical analysis criteria.

All 519 initial images were generated using the MetaTrader5 platform based on historical stock data from NASDAQ-listed companies, spanning from 2019 to May 2025. The selected assets covered three distinct sectors: technology (Apple Inc. — AAPL, AMD, Microsoft — MSFT, NVIDIA — NVDA, Intel Corporation — INTC), communication services (Alphabet — GOOGL, Meta — META), and consumer discretionary (Amazon — AMZN, Tesla — TSLA). These companies were chosen because their stocks exhibit high volatility, which motivated the use of relatively short moving average periods of 10 and 20. The temporal aspect does not affect the data preparation process because the 10- and 20-period moving averages are adjusted according to the timeframe used, whether daily, weekly, monthly, or hourly. Furthermore, the number of images generated for each company depends on the number of moving average crossovers identified for that company, with each crossover corresponding to one individual image.

The dataset emphasizes segments of candlestick charts highlighting moving average crossovers. For each moving average crossover of the selected stocks in the period considered, the candles preceding and succeeding the crossover were incorporated into the images to capture the full contextual information, as shown in Figure 3, which illustrates six images that make up the experimental database.

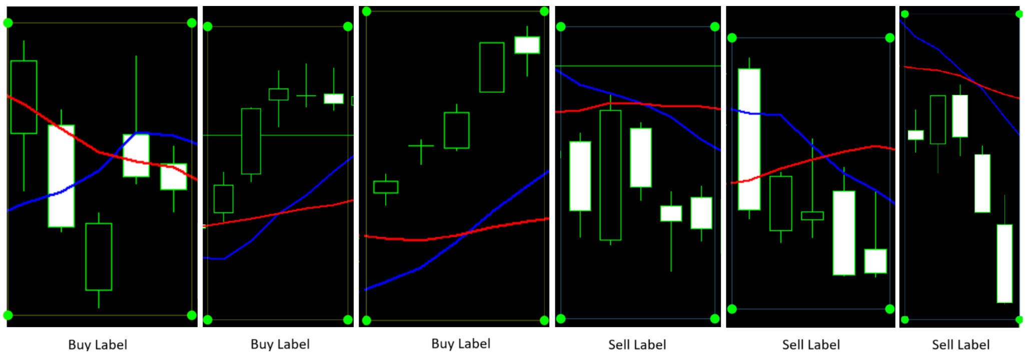

Figure 3. Labeled images of the experimental dataset

[2]https://pypi.org/project/labelImg/

The dataset was expanded using data augmentation techniques to improve model generalization and increase the variability of the training data, generating in 837 additional images and resulting in a total of 1,356 samples. Adjustments to brightness, light blurring, noise injection, compression, and proportional resizing were used to simulate variations that may occur during image generation or formatting. In contrast, transformations such as flipping and rotation were deliberately avoided, as they could distort the spatial orientation of candlestick patterns and lead to misclassification of trading signals.

3.2. Labeling Strategy Based on Technical Confirmation

The annotation of buy and sell signals integrates candlestick chart patterns with moving average crossovers, using a double-confirmation strategy to replicate the decision-making process commonly employed by traders in financial markets. Specifically, buy signals were labeled when a "golden cross"—meaning a short-term moving average crossing above a long-term moving average—coincided with a bullish candlestick pattern, such as a Doji Star. Conversely, sell signals were labeled when a "death cross"—where the short-term moving average crosses below the long-term moving average—was accompanied by a bearish pattern, such as a Dark Cloud.

By training different versions of YOLO with annotated images from the experimental database that follow the configuration shown in Figure 3, it is possible to apply the models to label a candlestick chart with multiple moving average crossovers, mirroring how traders conduct visual analysis in real-world trading scenarios, as shown in Figure 4, in which the labels "Buy 0.70" and "Sell 0.75" were automatically predicted by the YOLOv8 model trained on us our experiments. In Figure 4, supplementary observations were manually inserted (in yellow) to enhance contextual understanding and facilitate the interpretation of the prediction result.

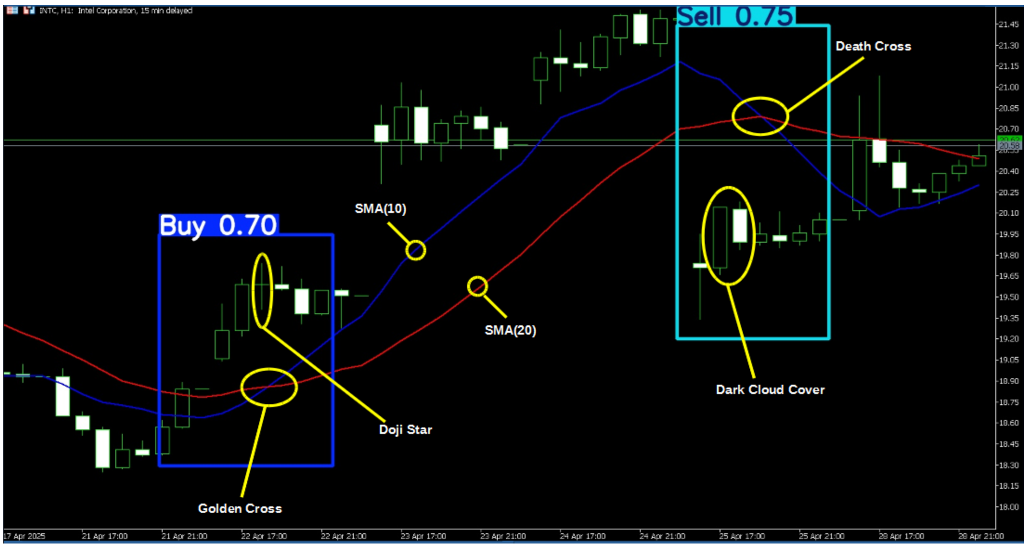

Figure 4. Prediction of candlestick patterns and moving average crossovers

3.3. Experimental Setup

For the experiments, the dataset was split into 80% for training and 20% for validation, ensuring a representative class distribution across both subsets to enable robust and reliable model evaluation. Each image in the training set contains information from only a limited portion of the graph (as in the six examples in Figure 3), ensuring that, regardless of how the images are distributed between the training and test sets, in the training process, the model is never exposed to data identical to that used for validation.

Data augmentation techniques were employed to generate diverse variants of the original images, enhancing the robustness and generalization capabilities of the models during both the training and testing phases. Subsequently, four versions of the YOLO architecture—YOLOv3 [Redmon and Farhadi 2018], YOLOv8 [Ultralytics 2023], YOLOv9 [Wang et al. 2023], and YOLOv11 [Chen et al. 2024]—were trained. These models were initialized with pre-trained weights and trained for 100 epochs using the following hyperparameters: input image size of 640×640 pixels, batch size of 16, and a learning rate of 0.001. No fine-tuning was applied. The choice of these YOLO versions reflects their progressive improvements in detection accuracy, computational efficiency, and generalization capabilities.

3.4. Experimental Metrics

For each trained model, considering the number of True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN), the following evaluation metrics were calculated:

- **Precision**: the proportion of correct TP detections among all predicted positives (TP + FP), reflecting the model's ability to avoid false alarms.
- **Recall**: the proportion of TP identified out of all actual positives (TP + FN), indicating the model's capacity to detect all relevant instances. Recall is a critical metric in financial contexts where missed buy or sell signals can cause direct losses for traders.
- **F1-score**: the harmonic mean of precision and recall, providing a single metric that balances both the ability to correctly detect relevant instances and to minimize false detections.

The selection of these metrics is justified by their widespread use in object detection research and their direct alignment with the YOLO evaluation pipeline, allowing for consistent, comparable, and meaningful analysis of model outcomes.

As in our work, the classification task considers two target classes—buy and sell—each evaluation metric was computed separately for each class. Additionally, macro-averaged values were calculated to summarize overall model performance across both classes. Since the classes are balanced and hold equal importance in the context of this study, macro-averaging provides an appropriate and fair means of evaluation. Moreover, it enables a consistent and meaningful performance comparison across different YOLO versions.

Additionally, to enable proper analysis of the results, we also used a Confusion Matrix for the results obtained with each version of YOLO. A confusion matrix is a table that compares the predictions made by the model with the real values of the classes, allowing a detailed analysis of the hits and misses made. In the matrices we used in this work, each column represents the real instances of a class, while each row represents the instances predicted as belonging to each class.

4. Results and analysis

Table 2 presents the experimental metrics calculated with the models trained applied in the training set. Each row refers to one of the YOLO versions. The first columns present the precision, recall, and f1-score results calculated for each of the two labels considered (buy and sell). Additionally, macro-averaged values were in the last columns to summarize overall model performance across both classes.

Table 2. Experimental metrics computed for each YOLO version

Model	Precision		Recall		F1-Score		Macro-Average		
	Buy	**Sell**	**Buy**	**Sell**	**Buy**	**Sell**	**Precision**	**Recall**	**F1-Score**
YOLOv3	0.81	0.83	0.80	0.83	0.80	0.83	0.82	0.82	0.82
YOLOv8	**0.90**	**0.86**	**0.83**	**0.92**	**0.86**	**0.89**	**0.88**	**0.87**	**0.88**
YOLOv9	0.89	0.86	0.83	0.91	0.86	0.89	0.88	0.87	0.87
YOLOv11	0.86	0.85	0.81	0.89	0.84	0.87	0.85	0.85	0.85

To complement the results analysis, Figure 5 has the confusion matrix of the results obtained with each version of YOLO explored, with YOLOv8 presenting the best results highlighted with a green bounding box.

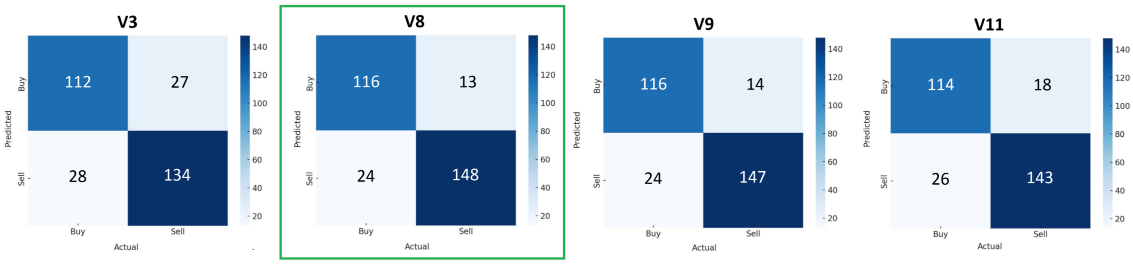

Figure 5. Confusion matrix for each YOLO version

Analyzing the results, it is possible to observe improvements across YOLO versions, highlighting that YOLOv8 always obtains the best results. Our work obtained superior results compared to other works that classify patterns in candlestick charts without moving averages. Compared to [Temur et al. 2024], which classified also buying and selling patterns in candlestick charts, our approach achieved recall values up to 0.18 higher and f1-scores up to 0.6 higher. Although the datasets employed in both studies differ — making a direct comparison unfeasible — the results obtained in our work suggest that the proposed method is promising. They highlight the potential viability of training specialized models for application in candlestick charts that include moving averages.

5. Conclusion

Integrating CV techniques with traditional technical analysis can increase decision-making agility in financial markets. Considering manual analysis's challenges, subjectivity, and inefficiency, especially in highly volatile environments, automatically detecting significant events on charts offers a solid basis for short-term trading strategies. It is important to develop strategies that further improve the results of CV methods in this context. In this work, we achieved promising results by evaluating the use of moving averages in candlestick charts, aiming to optimize the identification of buy/sell recommendation patterns on charts. In our experiments, YOLOv8 achieved the best results, obtaining an average precision of 0.88, recall of 0.87, and F1-score of 0.87.

While the results of the proposed approach were satisfactory to those found in previous state-of-the-art studies using similar models, its main contribution may lie in the methodology adopted. Specifically, using candlestick patterns known to signal trend reversals — when combined with moving average crossovers that indicate buying and selling opportunities—provided a robust framework for labeling and detection. We believe that this double-confirmation strategy increased the reliability of trading signals and contributed to improved model performance by providing a level of interpretability and alignment with trader behavior often lacking in purely data-driven approaches.

A limitation of the presented study is the limited experimental set. Although the study analyzed real data from nine NASDAQ-listed companies, the universe of companies on the stock exchange is much larger. It also considered only one stock-based chart formation, although different assets may exhibit varying patterns. The results indicate that training models using moving averages and candlestick charts can improve prediction results for charts that also use moving averages. However, it is still important to refine the experiments to confirm this hypothesis with certainty.

In addition to expanding the experimental database, future research could explore the integration of other technical strategies, such as Bollinger Bands or the Bigalow candlestick-based approach, to complement visual pattern detection and enrich the analytical framework for trading systems. Furthermore, synthetic data generation, model combinations, or transformer-based architectures could improve performance. Applying reinforcement learning techniques to simulate adaptive trading agents based on detected patterns is also a promising direction. Finally, deploying these models on real-time trading platforms, with latency analysis and continuous retraining, could bring this research closer to production-level applicability in the financial sector.

References

Birogul, S., Temur, G., and Kose, U. (2020). Yolo object recognition algorithm and 'buy-sell decision' model over 2d candlestick charts. *Journal of Artificial Intelligence and Data Mining*, 8(1):45–54.

Chen, J.-H. and Tsai, Y.-C. (2020). Encoding candlesticks as images for pattern classification using convolutional neural networks. *Financial Innovation*, 6(1):1–20.

Chen, L., Zhao, M., and Huang, R. (2024). Yolov11: Hybrid transformer-cnn architecture for robust object detection. *arXiv preprint arXiv:2403.01987*.

Cohen, G. (2021). Optimizing candlestick patterns for bitcoin trading strategies. *Review of Quantitative Finance and Accounting*, 57:987–1012.

Hu, Q., Xu, J., Luo, J., and Fan, Y. (2019). Formal specifications of candlestick patterns for technical analysis. *Applied Soft Computing*, 85:105700.

Koirala, A., Walsh, K. B., Wang, Z., and McCarthy, C. (2019). Deep learning–method overview and review of use for fruit detection and yield estimation. *Computers and electronics in agriculture*, 162:219–234.

Leal, I., Jayson, A., Lira, R., Veloso, R., Oseas, A., and Benjamin, M. (2024). Automatic cattle detection and counting system in aerial images using computer vision algorithms. In *Encontro Nacional de Inteligência Artificial e Computacional (ENIAC)*, pages 472–483. SBC.

Marshall, B. R., Young, E., and Cahan, R. H. (2006). Candlestick technical trading strategies: Can they create value for investors? *The Journal of Financial Research*, 29(3):305–316.

Neto, A. F. L., de Medeiros Santos, A., and Fernandes, S. (2023). Leaf detection using yolov4 for phytopathogenic diagnosis. In *Encontro Nacional de Inteligência Artificial e Computacional (ENIAC)*, pages 866–879. SBC.

Pires, V. C., Palmeira, E. S., and dos Santos, F. A. (2023). Application of deep learning techniques to depth images for person tracking and detection. In *Encontro Nacional de Inteligência Artificial e Computacional (ENIAC)*, pages 272–284. SBC.

Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. (2016). You only look once: Unified, real-time object detection. In *IEEE conference on computer vision and pattern recognition*, pages 779–788.

Redmon, J. and Farhadi, A. (2018). Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.

Sezer, O. B. and Ozbayoglu, A. M. (2020). Financial trading model with stock bar chart image time series with deep convolutional neural networks. *Financial Innovation*, 6(1):1–17.

Temur, G., Birogul, S., and Kose, U. (2024). Comparison of stock "trading" decision support systems based on object recognition algorithms on candlestick charts. *IEEE Access*, 12:35121–35132.

Thammakesorn, S. and Sornil, O. (2019). Candlestick pattern-based trading strategy using chaid. In *Journal of Physics: Conference Series*, volume 1195, page 012008. IOP Publishing.

Thomas, N. B. (2012). Encyclopedia of candlestick charts.

Ultralytics (2023). Yolov8 — ultralytics official documentation. Online; accessed 2025-05-14. `https://docs.ultralytics.com/models/yolov8/`.

Wang, C., Li, X., and Zhang, Y. (2023). Yolov9: Transformer-based label assignment for dense object detection. *arXiv preprint arXiv:2310.12345*.

Zhang, R., Zhao, C., and Lin, G. (2023). Interpretable image-based deep learning for price trend prediction in etf markets. *Quantitative Finance*.

Zhao, Z.-Q., Zheng, P., Xu, S.-t., and Wu, X. (2019). Object detection with deep learning: A review. *IEEE transactions on neural networks and learning systems*, 30(11):3212–3232.

Zou, Z., Chen, K., Shi, Z., Guo, Y., and Ye, J. (2023). Object detection in 20 years: A survey. *Proceedings of the IEEE*, 111(3):257–276.