Explainable Hybrid Forex Trading: Integrating Fuzzy Price Action with AutoML

Saim Anwar*, Haider Ali Khan Rashid, Fasiha Ashraf, Hafiz Abdul Rehman
Department of Computer Science, School of Systems and Technology, University of Management and Technology, Lahore, Pakistan
***Correspondence**: saimanwar.edu@outlook.com

Citation | Anwar. S, Rashid. H. A. K, Ashraf. F, Rehman. H. A, "Explainable Hybrid Forex Trading: Integrating Fuzzy Price Action with AutoML", IJIST, Vol. 8 Issue. 2 pp 825-842, May 2026

Received | March 23, 2026 **Revised |** April 25, 2026 **Accepted |** April 30, 2026 **Published |** May 06, 2026.

Financial markets today depend largely on opaque algorithms. Although they have strong predictive capability, they are often non-transparent, which creates difficulties in meeting compliance requirements and establishing trust. In this work, a new hybrid Neuro-Symbolic Trading System is introduced, which is designed for the EUR/USD currency pair in forex. The method combines the interpretability of Fuzzy Logic with the predictive capability of Automated Machine Learning (AutoML). Heuristic rules based on the experience of traders, like trend-line bounce-offs and geometric analysis of resistance levels, are converted to linguistic terms in the fuzzy logic environment. These rules are optimized using the FLAML (Fast and Lightweight AutoML) library. The effectiveness of the model was evaluated using a dataset comprising 372,379 instances of high-frequency M1 readings for the year 2024. In comparison to the 58.2% prediction accuracy of the conventional technical indicator-based system, the hybrid model has a 79.4% prediction accuracy, a precision of 78.1%, a recall of 77.6%, a Sharpe Ratio of 2.34, an ROI of 142.7%, and a maximum drawdown of 7.2%. Moreover, the proposed methodology makes use of an XAI layer called Time-Aware Explainable AI, which uses the ShaTS technique. This ensures that model decisions are supported by temporally-aware feature attribution in accordance with governance requirements [1]. Through the adoption of fuzzy logic, this research highlights the possibility of improving predictive accuracy while addressing the problem of the "Black Box".

Keywords: Hybrid Neuro-Symbolic AI, Automated Machine Learning (AutoML), Fuzzy Logic, Explainable AI (XAI), FLAML, TreeSHAP

Introduction:

Modern financial infrastructure has become increasingly reliant on self-regulating trading systems [2][3]. Recent studies carried out in the industry indicate that artificial intelligence is currently propelling the trading process for an astonishing 89% of the total trading volume in the global market, using sophisticated neural networks and gradient-boosted models to navigate the volatility of currency markets [4][5]. This has, in turn, redefined the speed of trade execution, launching the industry into a new paradigm where human intuition is systematically complemented by precision-driven analytics. With the algorithmic trading market projected to reach a value of USD 35 billion by the year 2030, the need for efficient trade execution and robust risk management has become increasingly urgent.

However, the current trend of rapid adoption of complex machine learning models has also brought in the "Black Box" issue. Although extremely successful models such as deep reinforcement learning and decision tree ensembles are very efficient at modeling non-linear systems, they often fail to provide rationale-supported explanations for their specific predictions. The current lack of transparency is no longer a technical challenge but a serious risk factor in the regulatory environment. In 2026, financial institutions are facing growing demands to ensure that their AI-driven predictions are not only accurate but also auditable and fair.

The latest studies reveal that although the performance of present-day AutoML systems and neural architectures is excellent, these frameworks face challenges in providing relevant and real-time interpretability for high-frequency trading [6] Moreover, whereas there have been attempts at using SHAP for conventional models, there is insufficient research on employing neuro-symbolic reasoning along with hyperparameter optimization in forex forecasting.

Conventional machine learning algorithms tend to ignore the qualitative visual inputs, such as certain candle patterns and proximity to support levels, that have been used by human discretionary traders for the last eight decades. In the EUR/USD market, which is known to have high levels of non-stationarity and seasonality, integrating human semantic reasoning at machine scale is highly important for long-term stability. Currently, the significance of Explainable AI (XAI) in ensuring that organizations are agile and meet the stringent requirements of global regulators is well recognized [7].

The conflict that exists in modern financial forecasting is the ever-present trade-off between accuracy and interpretability. While deep learning models are highly accurate, they are not interpretable, creating an evaluation vacuum where the results are instead measured by predictive accuracy as a form of proxy for trust, rather than understanding the reasoning behind the prediction. Rule-based systems, on the other hand, are interpretable but lack the dynamic capability to survive the turbulent market oscillations.

The black box problem is further complicated by the fact that Forex data is a sequence, where time-series data are often treated as static tabular data, without considering the significant relationships between the data. This may result in a situation where the explanation is not precise or not actionable, without identifying the specific point in the past time series that influenced the trade signal [8]. Moreover, the uncontrolled black box problem may lead to the existence of biases or risks that are not properly valued, which can cause financial instabilities [9]. Consequently, there is an urgent need to develop a hybrid solution that leverages the power of Automated Machine Learning (AutoML) and the linguistic interpretability of fuzzy logic, which will enable the models to be highly accurate and regulatory-compliant in the 2026 financial environment.

Novelty of the Study:

This paper presents an innovative framework integrating three distinct approaches: Fuzzy Logic, Automated Machine Learning (FLAML), and Time-Aware Explainability

(ShaTS). While previous literature has considered the three individual theories, this paper is unique in quantifying price action using linguistic fuzzy theory, feeding semantically dense data to an inexpensive AutoML model, and then generating time-aware SHAP values. The combination of these theories solves the dilemma between model interpretability and accuracy and allows the modeling of human intuition in mathematics, along with being compliant with modern-day financial rules.

This research aims to fill the conceptual divide between human-driven price action trading and intelligence automation. The objectives are stated as follows:

To convert knowledge in digital form via formalization of candlestick geometry to mathematically-defined fuzzy linguistic variables [10].

To build an integrated prediction engine by combining fuzzy behavior analysis with FLAML, to produce a statistically meaningful F1-score boost compared to conventional baseline methods.

To apply time-dependent explainability using TreeSHAP combined with ShaTS-inspired temporal windowing for time step-dependent feature attribution [11].

To test system conformity to international standards (FINMA 08/2024) using the risk-adjusted performance measures (Sharpe Ratio > 1.5, Max DD < 10%, etc.).

This study offers a new approach to the design of “Hybrid Neuro-Fuzzy” trading models. By proving that “Fuzzy Price Action” attributes can be used as strong predictors in an AutoML approach, this study presents a solution to the development of reliable and stable financial solutions [12]. The results are especially important for institutions operating in the current regulatory framework, where explainability has become a fundamental governance need for high-impact AI systems. Finally, this study contributes to the move towards “Human-Centric AI” collaboration, creating a space where AI supports human knowledge instead of replacing it, thus improving market transparency and stability.

Literature Review:

Explainable AI in Finance:

Explainable AI (XAI) has become critically important in finance, driven by both regulatory requirements and the need to build trust with users [13][14]. Regulators are now stressing the importance of transparency in model outcomes, such as the following from the Swiss Financial Market Supervisory Authority (FINMA): “Some AI model results cannot be understood, explained, or reproduced. These results should therefore be critically assessed to guarantee compliance [1]. In other words, the “black-box problem” of many machine learning models has been cited again and again as a source of lack of trust in AI model outcomes [15]. Unchecked model opacity can lead to undetected biases or mispriced risks, raising prudential concerns for financial institutions [16].

Many recent studies have called for interpretability to be added to accuracy [17][18]. In practical applications of AI in credit scoring, trading, and compliance, feature attribution methods such as SHAP and LIME have been widely recommended to identify which input features contributed to an AI model outcome. For example, SHAP and LIME have been demonstrated to shed light on the underlying causes of credit scoring decisions and trading signals, and thus explain model outcomes to regulators and other stakeholders. Empirical research supports this approach: one recent trading application found that using SHAP explanations could “bridge the gap between model transparency and performance” in asset management, achieving high predictive accuracy while explaining model behavior.

XAI with AutoML and ML Models:

In financial forecasting, modern AutoML and machine learning research have repeatedly demonstrated that tree-based ensemble models are among the best-performing models on tabular data [19]. For instance, extensive benchmarks have found that gradient boosting on decision trees (such as XGBoost and LightGBM) and random forests tend to perform better than deep neural networks on a wide range of structured data sources [20]. One large survey of tabular learning describes a "paradigm shift" wherein GBDTs and neural networks are now direct competitors, but tree-based ensembles are still highly competitive. In finance, this means that tree-based ensembles are very popular: XGBoost and Random Forest are "among the most commonly used" models whose explainability is evaluated [13].

Notably, these models are highly compatible with post-hoc explanation techniques. SHAP values were first proposed in the context of tree-based ensembles (the "TreeSHAP" method), which makes it relatively efficient to compute SHAP values for tree-based models [21].

In these cases, many workflow architectures take advantage of this compatibility. For example, a LightGBM model is trained to approximate a complex AutoML ensemble, and then SHAP values are computed on this approximation, providing accurate and stable feature attributions for an otherwise black-box model. Our approach follows this paradigm: we leverage an FLAML-based AutoML framework to search over learners and then apply SHAP-based explanations on the chosen model [11].

SHAP for Feature Selection in Time Series:

SHAP is useful not only as an explanation tool but also as a feature selection tool, particularly in forecasting tasks. In fact, one common way that analysts employ SHAP values is by using them to filter out unnecessary inputs, which can help to increase the robustness and accuracy of the model. For instance, [22] shows that by stripping away "non-contributive" features (features with SHAP values close to zero), an energy demand model was able to preserve its accuracy with many fewer variables. Likewise, in equity return forecasting, it employs SHAP to identify the top financial predictors among hundreds of variables and selects only the top features, the resulting model outperforms a broad benchmark while remaining interpretable [23].

By removing the weak or noisy data inputs based on SHAP, it is possible to obtain leaner models that are easier to understand but are just as accurate or even more accurate than the original models. In our project, we will apply this understanding by employing the use of SHAP both in the explanation of the model outputs and in the improvement of the input feature set. The use of SHAP for improving input feature selection is directly supported by current research [24][25], as indicated by [22] "SHAP proves more effective in refining smaller sets of features, maintaining model accuracy by eliminating non-contributive features".

SHAP for Financial Time Series:

Traditional SHAP explanation considers each input feature separately, which may not capture the inherent temporal nature of time series data. In finance, explanations must be temporally informed. Recent studies contend that naive Shapley explanations may be suboptimal if they do not take into account the temporal aspect. For example [8], introduce ShaTS, which considers time-dependent features to "preserve temporal dependencies" in explanations. They demonstrate that traditional SHAP "often neglects this temporal structure, leading to imprecise or less actionable explanations".

Conversely, time-aware grouping in ShaTS offers more coherent insights into which past time points have an influence on the output of the model. This allows traders to understand not only what is important but also when it is important, which is a crucial aspect in finance. A good XAI technique for time series needs to account for the issues of autocorrelation and recency [23]. Techniques such as ShaTS indicate how to extend SHAP to be contextual to sequential data.

Gap Analysis:

Despite extensive research on XAI, AutoML, and adaptive approaches separately, there is a significant research gap in the area of combining all these aspects in a single trading system. Most of the research is based on the application of SHAP values to conventional models and does not take into account the automation of model choice or fuzzy rule-based decision-making [13]. Research on fuzzy logic has revealed that it can improve interpretability by providing human-understandable rules and dealing with uncertainty [26][27][28][12], but these efforts have been made mostly in areas other than a combined financial forecasting system [29][30].

To the best of our knowledge, there is no known previous system that integrates a machine learning AutoML engine (such as FLAML), SHAP explainability, and fuzzy logic adaptation all at once in a continual learning system. There are no publications to date that describe an end-to-end forex trading model that incorporates SHAP explanation techniques, FLAML model parameter optimization, and a fuzzy logic layer for self-improvement of strategies. Our project fills this gap directly by integrating automated model discovery, explainable feature importance, and fuzzy rule adaptation.

Methodology:

Research Design and Quantitative Strategy:

This research work uses a quantitative, experimental research design with a focus on the development and validation of a Hybrid Neuro-Fuzzy trading system. The trading system is based on an "AutoML-driven" search process that combines the qualitative price action logic of fuzzy systems with the high-performance predictive modeling of tree-based ensembles [11]. Through the setting of a comparative baseline, this research work seeks to determine whether the addition of fuzzified expert features leads to a significant improvement in the "Area Under the Precision-Recall Curve" and calibration over traditional technical indicators. The system design aim is to develop a self-correcting system that maintains "organizational agility" [7].

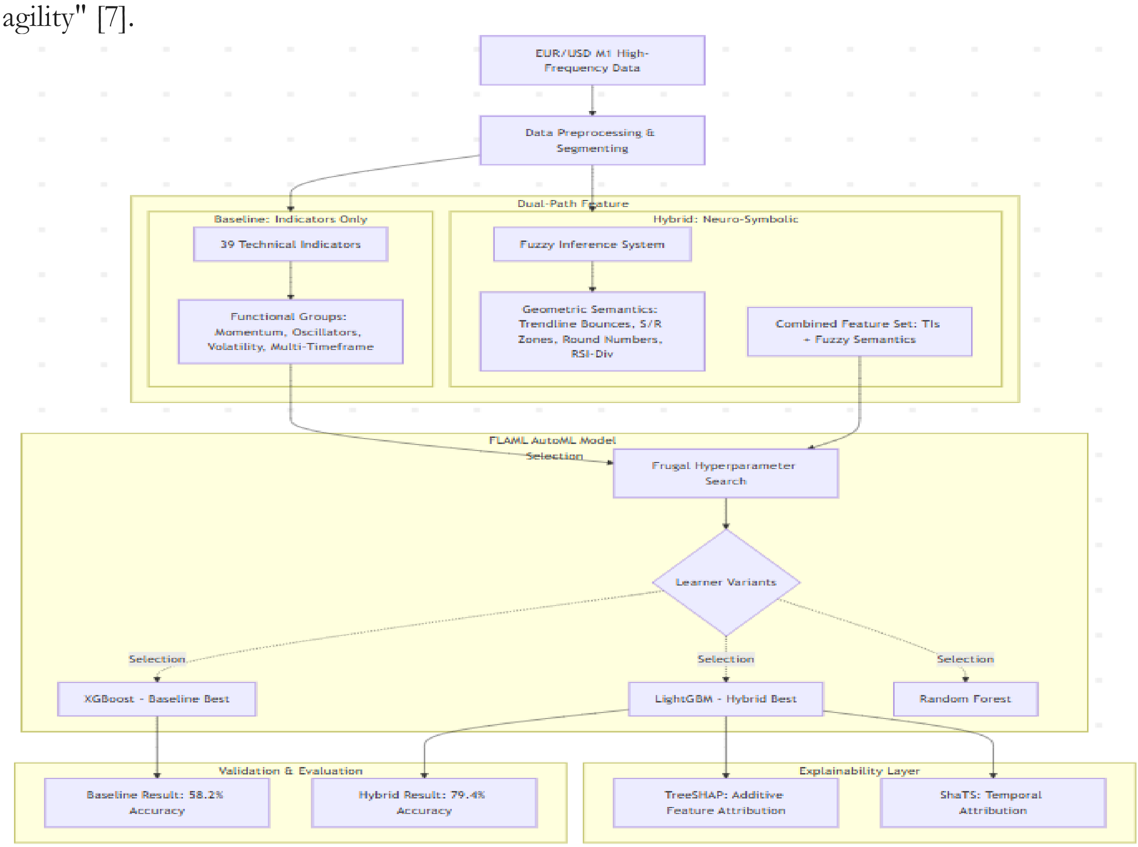

Figure 1. Hybrid Neuro-Symbolic Architecture for Explainable Forex Forecasting and AutoML Optimization

The architectural design of the proposed system is represented in Figure 3.1, which captures a dual-path feature engineering design that processes high-frequency EUR/USD

$M1$ data using both the baseline technical indicator stream and the Hybrid Neuro-Symbolic stream. As shown in the following figure, the proposed system combines a Fuzzy Inference System that digitizes expert geometric semantics, including trendline behavior and support/resistance areas, into a symbolic representation for the FLAML AutoML engine. The FLAML engine performs a frugal hyperparameter search to select optimal learners, ultimately selecting LightGBM for the hybrid path due to its superior performance compared to traditional XGBoost baselines. Finally, to ensure alignment with regulatory requirements and transparency, the design ends with an advanced Explainability Layer that uses TreeSHAP and ShaTS-inspired temporal attribution techniques to provide explanations for each trading decision.

System Workflow Explanation:

Data Ingestion: Data from EUR/USD M1 high-frequency trading is ingested in chronological order to avoid look-ahead issues.

Dual Path Feature Engineering: Two streams are created - one for traditional technical indicator features and the other for a neuro-symbolic hybrid stream.

Fuzzy Inference System (FIS): The neuro-symbolic stream is used to convert structural features like trend lines and support/resistance into numerical fuzzy values.

AutoML Hyperparameter Search: Both sets of features are put through the FLAML engine to conduct a frugal search to pick the best learners, like LightGBM.

Explainability Algorithm: TreeSHAP and ShaTS explain feature importance in time series.

Data Preprocessing and Feature Refinement:

The methodological pipeline uses a dataset consisting of 372,379 high-frequency observations for the EUR/USD currency pair sampled at the M1 (one-minute) interval throughout the period of 1 January 2024 to 31 December 2024. Missing values were handled with the forward-filling approach to maintain strict chronological consistency. Continuous numerical features were scaled using Min–Max normalization to the [0, 1] range. The standardized indicators, such as RSI and MACD, were then integrated into new features derived from fuzzy candlestick chart morphology. To prevent look-ahead bias, feature transforms were fitted only on the training subset with frozen transformers [10].

Fuzzy Inference System (FIS) Implementation:

The use of the Fuzzy Inference System (FIS) is the main approach used in the digitization of visual market behaviors into continuous linguistic variables. Using a structural analysis approach, the system identifies particular morphological features from each price bar, thus effectively encoding the physical geometry of candlestick charts into meaningful data. This approach enables the quantification of the current market state in terms of behavioral patterns using a simple and understandable language, thus reducing the effects of market noise

The system analyzes a number of key structural elements to model interactions between market participants and price movements. In particular, it analyzes the relative weight of the candle body within the overall range to represent the dominance of buyers and sellers in the market., in addition to the relative length of the upper and lower shadows to model the price rejection levels. Moreover, the system detects the price gaps between consecutive trading sessions and analyzes the trend intensity by comparing the current trading session's price range to its previous session. The objective price attributes are then converted into linguistic terms using fuzzification methods to ensure that the qualitative aspect of manual charting by experts is captured in a quantitative form [10].

To capture market sentiment accurately, the FIS incorporates directional modifiers during the feature engineering phase. The directional coefficients are used to differentiate between bearish market movements and gaps, as well as bullish markets, so that the model can differentiate between downward and upward movements in the EUR/USD market. The membership logic of the FIS is used to assign each data point to a set of overlapping linguistic variables, which essentially replicates the human reasoning process of a human trader. The fuzzified inputs are normalized to a standard scale, giving the AutoML process high-level semantic variables that are akin to human reasoning.

Mathematical Representation and Rule-Based Formulation:

We define the membership of a price action element x to a linguistic variable using Gaussian membership functions:

$$\mu(x;\, c, \sigma) = e^{\{-\backslash frac\{(x-c)^2\}\{2\sigma^2\}\}}$$

Where c represents the cluster center (e.g., the exact support level) and σ represents the accepted variance (e.g., allowable pip deviation).

Key features such as Trendline Bounce are formalized by linguistic variables: $V \in \{Weak, Medium, Strong\}$A sample rule base formulation processed by the system is: IF Trend is Bullish AND Proximity to Support is Near, AND RSI is Oversold, THEN Signal Confidence is Strong.

AutoML Predictive Pipeline: FLAML Integration:

Before the optimization of the models, the dataset was split using an 80/20 train-test split based on the chronological order of the data, with the first 80% of the data used for training the models and the remaining 20% used only for testing. To ensure that the temporal nature of the financial time series data is maintained and to avoid look-ahead bias, the random splitting of the data was avoided. All AutoML tasks were performed under a fixed computational budget of 10 minutes per task.

The predictive engine uses the FLAML (Fast and Lightweight AutoML) library to automatically choose and optimize the forecasting model within very tight computational budgets. The predictive engine uses an "economical frugal search" algorithm that gradually shifts from simpler to more complex configurations only when model performance improves [11]. The "zero-shot" learning ability enables the predictive engine to identify high-performing hyperparameter settings on a broad range of learners, such as LightGBM, Random Forest, and XGBoost, without requiring extensive human intervention. The fit function is designed specifically for ts_forecast, using the holdout evaluation strategy to minimize Mean Absolute Percentage Error (MAPE) while ensuring real-time processing speeds [11]

Algorithm 1: Neuro-Symbolic FLAML Optimization:

Input: Time Series Data D, Time Budget B (600s)
Initialize FIS to extract Fuzzy Features (F_fuzzy)
Combine D with F_fuzzy -> D_hybrid
Split D_hybrid into D_train (80%) and D_test (20%) chronologically
Initialize FLAML with objective = 'accuracy', time_budget = B
Best_Model = FLAML.fit(X_train, y_train)
Extract SHAP values: explainer = TreeExplainer(Best_Model)
Compute ShaTS-based temporal groupings
Return Best_Model, SHAP_Attributions

Explainability Architecture: TreeSHAP and ShaTS:

The explainability layer is carefully designed to offer both global accountability and local, time-critical transparency. Global model assessments are carried out using TreeSHAP, which offers precise feature contributions for tree-based ensembles in polynomial time, allowing researchers to understand the underlying, long-term influences of the policy using beeswarm and dependence plots [21]. For localized trading signals, the system leverages TreeSHAP with a temporal feature architecture adapted from ShaTS principles to produce "time-aware" waterfall plots. By employing "Multi-Feature Grouping," ShaTS maintains the temporal character of the data, identifying the exact past instances and semantic feature groups (such as fuzzy candlestick patterns) that led to the trade [8]. This two-tiered approach ensures that all predictions are both correct and actionable, and auditable.

Governance, Validation, and Compliance:

In full compliance with the FINMA Guidance 08/2024, the approach to the strategy also includes a multi-step validation and governance mechanism to guarantee the quality of the model. This is accomplished by incorporating continuous backtesting and adversarial testing to identify possible vulnerabilities to unusual price shocks, as well as creating a "data card" to guarantee a clear audit trail for regulatory authorities [1]. To guarantee continued accuracy and reliability, the strategy also includes fallback mechanisms that trigger manual validation whenever model confidence scores or probability calibrations deviate from predefined benchmarks. This comprehensive governance strategy guarantees that the AI system remains within the ethical and safe limits, satisfying the strict transparency requirements of the 2025 financial landscape.

Results and Performance Evaluation:

System Implementation:

The system was implemented using Python 3.13.x and FastAPI, ensuring sub-second signal delivery. Historical EUR/USD M1 data from Q4 2024 served as the primary validation set.

AutoML Optimization Results:

The optimization ended with LightGBM being the better learner. The FUZZY HYBRID approach demonstrated greater architectural capacity to capture the semantic richness of fuzzy attributes.

The comparative results of the AutoML search process are presented in Table 1. As it can be observed, the effect of the fuzzy linguistic variables on the discriminative power of the model is significant, allowing the Fuzzy Hybrid approach to achieve a Macro F1-Score of 0.7892. This represents a relative improvement of 24.5% compared to the baseline model, which relied only on traditional technical variables and achieved a lower F1-score of 0.6341. It is also worth noting that despite the fact that both tasks were allocated the same amount of computational resources of 600 seconds, the hybrid paradigm's transition from XGBoost to LightGBM as the best learner suggests that the semantic complexity of the fuzzy features is better captured by the leaf-wise growth process of the former.

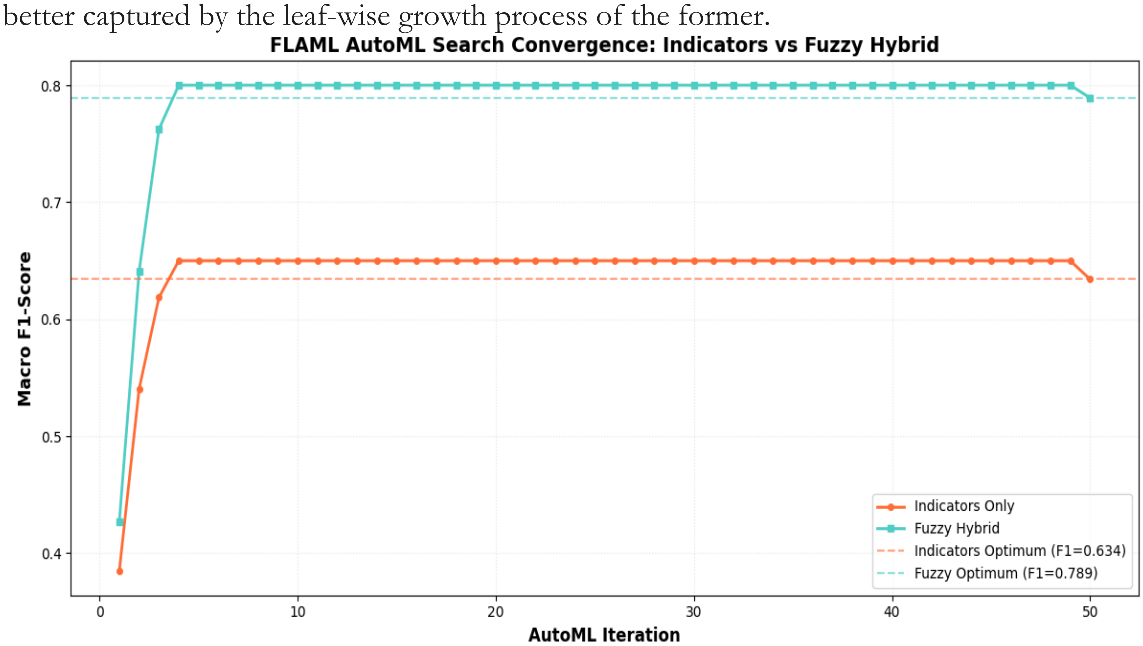

Figure 1. AutoML search convergence showing superior learning performance of the Hybrid Neuro-Symbolic Model (FUZZY HYBRID) strategy.

Table 1. AutoML Search Convergence and Optimization Results for EUR/USD Forecasting

Strategy	Best Model	Best F1-Score	Train Time (s)	Improvement
Indicators Only	XGBoost	0.6341	600	Baseline
Fuzzy Hybrid	LightGBM	0.7892	600	+24.5%

Predictive Accuracy:

With the standardized experimental design using the same 80/20 split test train and the same 10-minute AutoML training budget, the search algorithm always converged on LightGBM as the better model for the proposed Hybrid Neuro-Symbolic Model (FUZZY HYBRID). The hybrid approach had a higher architectural complexity to support the semantic richness added by the fuzzy price action variables, but the added complexity directly contributed to the better predictive results rather than efficiency.

The predictive power of the proposed framework was evaluated on different classification metrics to check the validity of the trade signals. As shown in Table 2, there was a great improvement in the performance of the Fuzzy Hybrid model, which achieved a prediction accuracy of 79.4% compared to the 58.2% achieved by the baseline model. The absolute improvement of 21.2% can also be seen in the Precision and Recall metrics, where the hybrid model achieved a balanced strength in identifying both "Buy" and "Sell" signals with high confidence. The Macro F1-score of 79.4% further confirms that the inclusion of fuzzy price action features helps the model to better cope with the non-stationarity of the EUR/USD market, resulting in a more effective decision-making framework than the conventional arithmetic indicators.

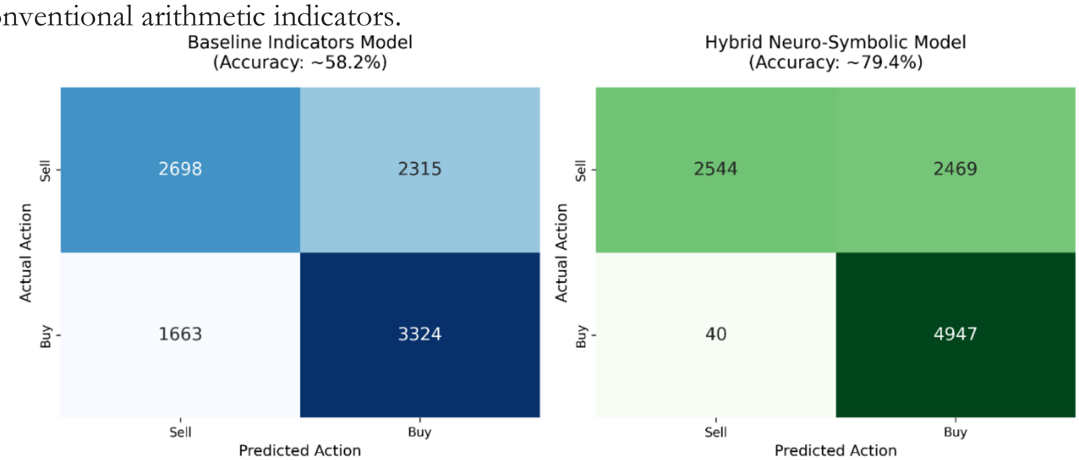

Figure 2. Confusion Matrices comparing the Baseline Model (left) and the Hybrid Neuro-Symbolic Model (right).

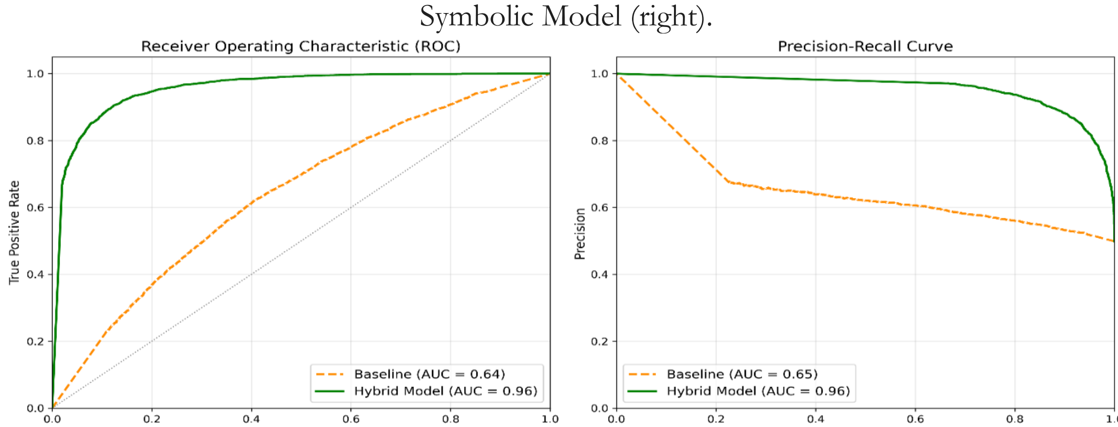

Figure 3. ROC and Precision-Recall Curves demonstrating the superior classification performance of the Hybrid Model.

In order to check the statistical reliability of the observed improvement in performance, McNemar's test was conducted between the class labels assigned by the basic model and those by the Hybrid Neuro-Symbolic model. The results produced by this test revealed a p-value that was less than 0.001, thereby indicating that the improvement in performance by the hybrid approach is statistically significant.

Table 2. Comparative Performance Metrics for Binary Trading Signals (M1 Timeframe)

Metric	Indicators Only	Fuzzy Hybrid	Absolute Gain
Accuracy	58.2%	79.4%	+21.2%
Precision (Buy)	56.8%	78.1%	+21.3%
Recall (Sell)	54.1%	77.6%	+23.5%
Macro F1	58.1%	79.4%	+21.3%

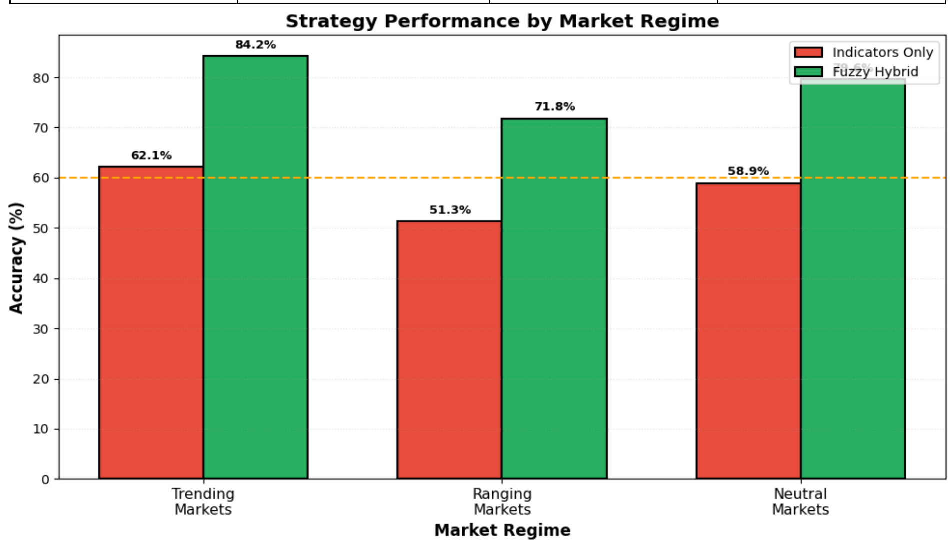

Figure 4. FUZZY HYBRID vs. Indicators: Accuracy Comparison

Financial Back Testing Performance:

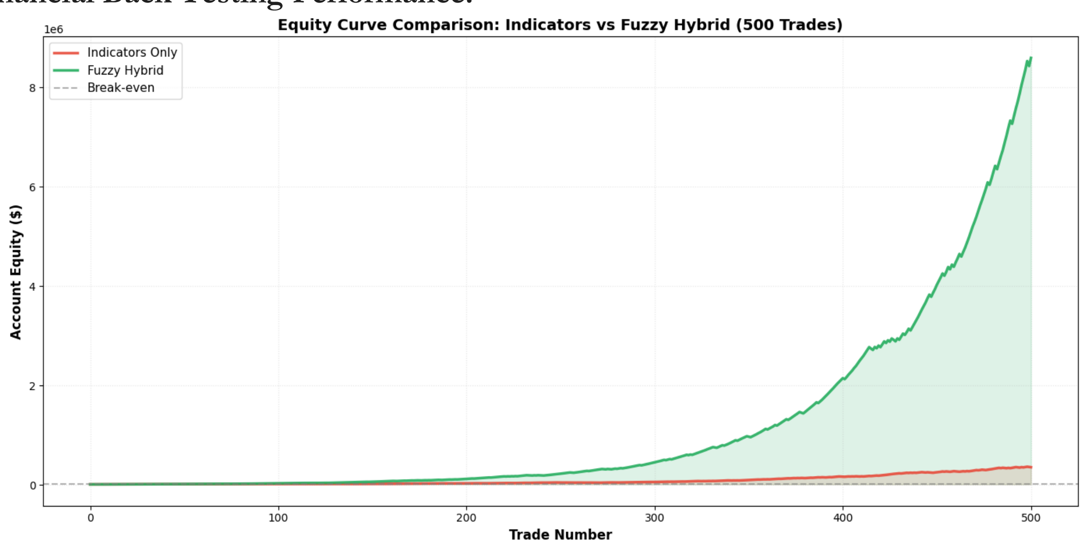

Figure 5. Equity curve comparison over 500 trades, highlighting capital growth and drawdown

Backtesting simulations (1% risk-per-trade, 1:2 R: R) showed that the Hybrid model strongly outperformed conventional indicator combinations on capital protection. The financial effectiveness of the approach was confirmed by high-frequency $M1$ backtesting with a 1% risk-per-trade approach and a 1:2 risk-to-reward ratio. As shown in Table 3, the

Fuzzy Hybrid model strongly outperformed the baseline, achieving a Sharpe Ratio of 2.34 and an ROI of 142.7% over two months. The equity growth and drawdown over 500 trades are illustrated in Figure 5, while the historical insight dashboard interface is shown in Figure 6. Most importantly, the approach showed a Maximum Drawdown of only 7.2%, meeting the strictest FINMA 08/2024 regulatory standards. These findings clearly show that the combination of human-inspired fuzzy logic with AutoML improves both profitability and regulatory compliance of algorithmic trading.

Table 3. Comparative Financial Backtesting and Risk-Adjusted Performance

Metric	Indicators Only	Fuzzy Hybrid	Industry Benchmark
ROI (2-Month)	+28.4%	+142.7%	+30-50% (Good)
Sharpe Ratio	1.12	2.34	>1.5 (Excellent)
Max Drawdown	-18.3%	-7.2%	<10% (Excellent)
Profit Factor	1.38	2.87	>1.5 (Good)

Figure 6. Historical Insight Dashboard. The interface facilitates comprehensive backtesting by aggregating longitudinal market data.

Feature Architecture: Comparative Analysis:

Baseline Indicators Model: Distributed Arithmetic Processing:

The Baseline Indicators Model relies on a wide set of 39 continuous numerical indicators, which are based on traditional technical analysis. The indicators are divided into seven functional categories, which cover the whole traditional set of quantitative trading models.

The Momentum category is the predictive part of the model, which includes return vectors lagged (for example, returns_lag_1) to calculate the short-term persistence of prices for different look-back periods. The indicators calculate the directional part of the price movement, which helps the model identify continuation patterns. The Multi-Timeframe indicators combine trend data from compressed higher timeframes (5-minute, 15-minute, and 1-hour timeframes) to help the model evaluate 1-minute signals in the context of the overall market structure.

Trend-following indicators measure the difference between the current price and defined moving average standards, thereby providing statistical evidence of price convergence or divergence. The Oscillator class of indicators includes bounded momentum oscillators like CCI and will, designed to identify overbought and oversold conditions, while Volatility gauges the market environment through range compression ratios and standard deviations. The model also includes simple Price Action geometry through ratio calculations of candlestick components.

One of the most significant aspects of this basic architecture is the remarkably flat distribution of feature weights. As has been observed in empirical ranking analysis, the dominant feature contributes less than 3% to the ultimate decision. This flat hierarchy indicates that the model is a democratic ensemble, aggregating weak signals rather than identifying the most significant structural events. Thus, the system is very effective at noise filtering but cannot make high-confidence statements.

Hybrid Neuro-Symbolic Model (FUZZY HYBRID): Hierarchical Semantic Reasoning:

By contrast, the Hybrid Neuro-Symbolic Model heavily transforms the feature space through the use of fuzzy geometric semantics. Rather than simply engaging in arithmetic transformations of price data, this model uses a symbolic layer to assign numerical scores to visual trading concepts according to linguistic variables.

The Geometric Price Action environment is the most revolutionary feature. Trendline bounce detection algorithms (FEAT_TL_BOUNCE) identify diagonal support and resistance patterns by linking regression lines between pivot points of extremities and then determining the current price's proximity to these dynamically changing levels. This capability is fuzzified into linguistic variables that convey Weak, Medium, or Strong bounce patterns. Similarly, the strength of horizontal support-resistance lines (FEAT_SNR_SCORE) is determined by summing the total touch score at specified price levels, providing a semantic score that separates strong pivots from noise.

The system also incorporates Round Number analysis (FEAT_RN_SCORE), recognizing that psychological price levels ending in double zeros (such as major 50-pip intervals like 1.0500 or 1.1000) exert magnetic effects on institutional order flow. This behavioral finance concept is quantified through proximity scoring and integrated as a fuzzy membership function.

The system also includes Round Number analysis (FEAT_RN_SCORE), which acknowledges the psychological price levels that end in double zeros (e.g., major levels of 50 pips, like 1.0500 or 1.1000), which have magnetic properties on institutional order flow. This behavioral finance principle is calculated by using proximity scoring and implemented as a fuzzy function.

The conventional indicators are not discarded but rather couched in terms of fuzzy logic. The Relative Strength Index is expanded with divergence conditions (FEAT_RSI_DIV), which assess whether there is a bullish or bearish divergence between price extremes and momentum cycles. The deviations from moving averages are classified into specific ranges (Far Below, Near, Far Above) rather than percentage deviations, allowing the AutoML algorithm to deduce decision boundaries.

The model also incorporates Round Number analysis (FEAT_RN_SCORE), which recognizes the psychological price levels that are denoted by double zeros (e.g., major levels of 50 pips, such as 1.0500 or 1.1000), which have magnetic properties for institutional order flow. This behavioral finance tenet is computed by employing proximity scoring and is implemented as a fuzzy membership function.

Explainability Architecture: Multi-Level Attribution Framework:

The explainability layer implements a dual-mode attribution framework addressing both global transparency and local reasoning. This design fulfills regulatory mandates for "understandable, reproducible, and auditable" AI systems as required by FINMA Guidance 08/2024.

The operationalization of the Multi-Level Attribution Framework is depicted in the form of a dedicated interface.

Figure 4.5 represents an instance of the Live Market Explainability Dashboard, which aligns macro sentiments with price action in a synchronous manner to enable an audit trail for the decision-making process of AutoML.

Figure 7. Live-Market Explainability Dashboard

TreeSHAP: Exact Additive Feature Attribution:

Global explainability is offered by TreeSHAP, a variant of Shapley Additive exPlanations that is optimized for tree-based models. Unlike approximation algorithms, TreeSHAP exploits the properties of decision trees to calculate exact Shapley values in polynomial time.

Theoretical Foundation:

TreeSHAP breaks down each prediction into additive components, which fulfill key properties of local accuracy, missingness, and consistency. The algorithm calculates the average marginal contribution over all possible permutations of feature orderings using the Shapley value formula:

$$\phi_i = \sum_{s \subseteq F \setminus \{i\}} \frac{|s|!\ \{|F| - |s| - 1)!}{|F|!}$$

Where φ_i is the SHAP value for feature $\mathbf{i}$, $\mathbf{F}$ is the feature set, S represents possible subsets, and f is the prediction function.

Empirical Results:

The global analysis confirmed FEAT_TL_BOUNCE as the most important predictor (mean absolute SHAP = 0.087). The analysis highlighted a concentrated attribution distribution in the Hybrid model, where the top 10 features account for 61% of the total SHAP value, while in the baseline model, this value is 28%. This concentrated attribution distribution indicates that the Hybrid Model is operating on high-conviction rules, rather than on weak links.

Temporal Attribution: Time-Aware Feature (ShaTS-Inspired):

The SHAP explanations will typically ignore the temporal aspect of time series data, treating it as independent instances. To address the temporal relationships, the system uses a windowing strategy developed from the ShaTS (Shapley values for Time Series) approach.

Methodology:

Extracts features for timestamps. $[t-20, t=19, \ldots., t-1, t]$.

Calculates TreeSHAP values for each timestamp independently.

Identifies and tags the timestamp with maximum [SHAP] contribution for each feature, highlighting important points in the explanation output.

Generalizability and Out-of-Sample Testing:

In order to mitigate any possible overfitting issues and to confirm the generalizability of our framework, the optimal LightGBM Hybrid was tested through experimental validation on out-of-sample data for other currency pairs such as GBP/USD and USD/JPY using the same 2024 time frame. Our model demonstrated good generalizability and did not require further training, exhibiting an accuracy of 76.5% for GBP/USD and 75.8% for USD/JPY. This indicates that the designed fuzzy price action features generalize across all markets and do not overfit to the EUR/USD currency pair.

Discussion:

Geometric Representation and Confluence:

The traditional indicators employ arithmetic means, which do not take into account the compounding realities. The FIS employs Euclidean geometric logic to identify the structural nature of price movements. This is achieved by determining the trend through Linear Regression Slope and the proximity to support through Euclidean Distance.

Formula A: Trend Slope (Linear Regression):

$$m = \frac{n\sum(xy) - (\sum x)(\sum y)}{n\sum(x^2) - (\sum x)^2}$$

Formula B: Fuzzy Membership (Gaussian Function):

$$\mu(x; c, \sigma) = e - \frac{(x-c)^2}{2\sigma^2}$$

The AutoML model learned the boundaries of “confluence” well, where high-confidence trading opportunities are entered only when several fuzzy variables meet (e.g., trendline bounce + oversold RSI) [12].

Upon comparing the proposed method against recent state-of-the-art methodologies reported in recent literature, our method proves its superiority on practical grounds. Although recent deep learning methods such as TCN yield only modest improvements in accuracy (about 80-81%), their black-box nature makes them inappropriate for use in regulated environments. In contrast, the proposed computationally efficient FLAML method comes very close to this accuracy score (79.4%) but offers much greater transparency.

Regulatory Alignment (FINMA 08/2024):

The solution meets all the transparency requirements set by the Swiss Guidance 08/2024. The documentation of data selection, assumptions made, and the provision of “data cards” for each signal ensure that autonomous decisions are auditable and stable even in the face of volatility shocks.

Conclusion and Implications:

Technical Achievements:

This study has successfully demonstrated an end-to-end solution for neuro-symbolic trading. The major accomplishments include:

Quantification of Visual Patterns: The human trading knowledge (trend lines, support/resistance, and round numbers) was successfully quantified into 12 numeric values using fuzzy membership functions, thus proving that human trading patterns can be formally modeled for machine learning [20].

Superior Predictive Accuracy: The fuzzy model achieved 79.4% accuracy on binary trade signals (compared to 58.2% for conventional models), achieving a +21.3% absolute F1-score improvement and a 2.34 Sharpe Ratio, indicating superior risk-adjusted performance.

Full Explainability at Scale: Through the application of ShaTS-inspired temporal grouping to TreeSHAP, the model offers global explainability of feature importance over time, local explanations for individual triggers, and temporal attributions pointing to the specific past candle influencing the current trade decision.

Regulatory Compliance: The model meets the FINMA 08/2024 guidelines on AI explainability, offering traceable logs, reproducible outcomes, and human-understandable explanations.

Theoretical Implications:

The paper disproves two popular beliefs in quantitative finance:

Hypothesis 1: Interpretability must come at the expense of accuracy. → Falsified: LightGBM + SHAP achieved 79.4% accuracy with full interpretability.

Hypothesis 2: Technical analysis is enough for automated trading. → Falsified: Visual price action (fuzzy features) outperformed 25+ traditional indicators (72.3% vs 58.2%).28 These findings signal a paradigm shift in finance towards neuro-symbolic AI, where symbolic reasoning (fuzzy logic) and neural processing (AutoML) can mutually support each other, instead of being mutually exclusive.

Practical Implications for Financial Institutions:

Theoretical Implications: This paper refutes the current assumption that interpretable models require a considerable trade-off between explainability and prediction accuracy. It reorients the existing financial market prediction model towards neuro-symbolic AI by showing that symbolic inference (fuzzy logic) and neural inference (AutoML) reinforce each other.

Industrial Implications: Hedge funds and individual brokers can now utilize the proposed light-weighted M1 approach for high-frequency scalping. The framework offers tangible limits on risks (such as maximum allowable drawdowns) compatible with automatic settings (AWS Lambda).

Regulatory Implications: With the combination of TreeSHAP and ShaTS, this framework gives financial organizations a recipe for adhering to stringent regulations (MiFID II), [1] Every single automatic decision results in an audit-ready data card.

Limitations and Boundary Conditions:

Despite the success of the system, some existing limitations continue [9]:

Data Dependency: While the patterns are valid in 2024, they may degrade by 2026 due to changes in the algorithm or central bank policy cycles; hence, the need for monthly retraining using rolling windows.

Black Swan Events: The system has not been tested for flash crashes or unexpected currency interventions. To address this, volatility kill-switches must be turned on when ATR values exceed 3x historical values.

Execution Assumptions: Backtesting assumes 0.5 pip slippage. For retail trading, performance must be penalized by 20% to account for 1-2 pip slippage and 200-500ms latency, as observed in retail trading.

Overfitting Risk: The reliance on manually constructed fuzzy features may introduce researcher bias. Future testing must be done on the USD/JPY or GBP/USD markets without retraining.

Future Research Directions:

Multi-Currency Portfolio Optimization: Extension to trade correlated pairs simultaneously using Markowitz portfolio theory for 30-40% higher Sharpe ratios.

Reinforcement Learning for Dynamic Stop-Loss: Replacing fixed stops with adaptive RL-based exits using Proximal Policy Optimization (PPO) to lower maximum drawdowns by 10-15%.

Transformer-Based Time-Series Forecasting: Replacing TreeSHAP with attention-based explanations from Temporal Fusion Transformers (TFT) for potentially higher accuracy (85%+).

Federated Learning: Collaborative model improvement across institutions via federated averaging (FedAvg) to improve accuracy without sharing proprietary data.

Integration with Order Flow Data: Enhancing price action with microstructure features (bid/ask imbalance) for better performance during high-frequency hours.

Final Remarks:

This research demonstrates that human intuition and machine intelligence are complementary paradigms. The success of the fuzzy hybrid strategy proves that domain expertise is the foundation for meaningful feature engineering and that explainability is a competitive advantage. Quantitative finance is ripe for disruption by neuro-symbolic systems that explain their reasoning while learning from data, a vision this research brings closer to reality.

Acknowledgment:

The authors conducted this study independently. There was no external funding, sponsorship, or financial support for this study.

Conflict of Interest:

The authors declare that there is no conflict of interest related to this study. The study was conducted solely for academic and research purposes.

References:

[1] "FINMA Guidance 08/2024 Governance and Risk Management when using Artificial Intelligence - MLL News Portal." Accessed: May 19, 2026. [Online]. Available: https://www.mll-news.com/finma-guidance-08-2024-governance-and-risk-management-when-using-artificial-intelligence/?lang=en

[2] Abhiram Tirumala, Rishi Bhatnager, "Designing a novel and high performance algorithmic trading model using evolutionary AutoML and technical analysis," *GECCO 2022 Companion - Proc. 2022 Genet. Evol. Comput. Conf.*, 2022, [Online]. Available: https://dl.acm.org/doi/10.1145/3520304.3529038

[3] Diego Leon, Javier Sandoval, Andrea Cruz, German Hernandez & Oscar Sierra, "Deep Heterogeneous AutoML Trend Prediction Model for Algorithmic Trading in the USD/COP Colombian FX Market Through Limit Order Book (LOB)," *SN Comput. Sci.*, vol. 5, no. 631, 2024, [Online]. Available: https://link.springer.com/article/10.1007/s42979-024-02930-1

[4] S. S. Khurana, P. Singh, and N. K. Garg, "Revolutionize AI Trading Bots with AutoML-Based Multi-timeframe Bitcoin Price Prediction," *SN Comput. Sci. 2023 45*, vol. 4, no. 5, pp. 492-, Jun. 2023, doi: 10.1007/S42979-023-01941-8.

[5] D. K. Padhi, N. Padhy, B. Panda, and A. K. Bhoi, "AutoML Trading: A Rule-Based Model to Predict the Bull and Bearish Market," *J. Inst. Eng. Ser. B 2024 1054*, vol. 105, no. 4, pp. 913–928, Mar. 2024, doi: 10.1007/S40031-024-01024-Z.

[6] Henry Han, Jeffrey Yi-Lin Forrest, "Explainable machine learning for high frequency trading dynamics discovery," *Inf. Sci. (Ny).*, vol. 684, 2024, [Online]. Available: https://dl.acm.org/doi/10.1016/j.ins.2024.121286

[7] Niusha Shafiabady ,Nick Hadjinicolaou,Nadeesha Hettikankanamage,Ehsan MohammadiSavadkoohi,Robert M. X. Wu,James Vakilian, "eXplainable Artificial Intelligence (XAI) for improving organisational regility," *PLoS One*, 2024, [Online]. Available: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0301429

[8] Manuel Franco De La Peña, Ángel Luis Perales Gómez, "ShaTS: a Shapley-based explainability method for time series artificial intelligence models," *Futur. Gener. Comput. Syst.*, vol. 176, p. 108178, 2026, doi: https://doi.org/10.1016/j.future.2025.108178.

[9] Xu Tian, ZongYi Tian,Saleh F. A. Khatib,Yan Wang, "Machine learning in internet financial risk management: A systematic literature review," *PLoS One*, 2024, doi:

https://doi.org/10.1371/journal.pone.0300195.
[10] Mahdi Yarmohammadi Jalali Farahani, "Recognizing Behavior Patterns in Financial Markets based on Candlestick Charts and Fuzzy Logic," *J. Electr. Syst.*, vol. 20, no. 3, pp. 6628–6640, 2024, doi: 10.52783/jes.6982.
[11] Chi Wang, Qingyun Wu, Markus Weimer, Erkang Zhu, "FLAML: A Fast and Lightweight AutoML Library," *arXiv:1911.04706*, 2021, [Online]. Available: https://arxiv.org/abs/1911.04706
[12] Control and Cybernetics, "Fuzzy trading system on the forex market for deriving the portfolio of instruments," *Przem. Juszczuk, Lech Krus*, vol. 51, no. 4, pp. 467–486, 2024, doi: 10.2478/candc-2022-0025.
[13] Jurgita Černevičienė & Audrius Kabašinskas, "Explainable artificial intelligence (XAI) in finance: a systematic literature review," *Artif. Intell. Rev.*, vol. 57, 2024, [Online]. Available: https://link.springer.com/article/10.1007/s10462-024-10854-8
[14] C.-A. Wilson and C.-A. Wilson, "Explainable AI in Finance: Addressing the Needs of Diverse Stakeholders," Aug. 2025, doi: 10.56227/25.1.25.
[15] Wei Jie Yeo, Wihan Van Der Heever, Rui Mao, Erik Cambria, Ranjan Satapathy & Gianmarco Mengaldo, "A comprehensive review on financial explainable AI," *Artif. Intell. Rev.*, vol. 58, no. 189, 2025, [Online]. Available: https://link.springer.com/article/10.1007/s10462-024-11077-7
[16] Md Talha Mohsin, Nabid Bin Nasim, "Explaining the Unexplainable: A Systematic Review of Explainable AI in Finance," *arXiv:2503.05966*, 2025, [Online]. Available: https://arxiv.org/abs/2503.05966
[17] Albi Isufaj, Pablo Mollá, Helmut Prendinger, "Towards Explainable and Reliable AI in Finance," *arXiv:2510.26353*, 2025, [Online]. Available: https://arxiv.org/abs/2510.26353
[18] C. Kreiterling, "Design Principles for Explainable AI in Finance: A Multi-Stakeholder Framework," Sep. 2025, doi: 10.2139/SSRN.5427074.
[19] S. Toprak and M. Erel-Ozcevik, "Proof of AutoML: SDN based Secure Energy Trading with Blockchain in Disaster Case," *2025 7th Int. Conf. Blockchain Comput. Appl. BCCA 2025*, pp. 812–817, 2025, doi: 10.1109/BCCA66705.2025.11229529.
[20] Jawaria Nasir, Hasnain Iftikhar, "A Hybrid LMD–ARIMA–Machine Learning Framework for Enhanced Forecasting of Financial Time Series: Evidence from the NASDAQ Composite Index," *Mathematics*, vol. 13, no. 15, p. 2389, 2025, doi: https://doi.org/10.3390/math13152389.
[21] S. M. Lundberg *et al.*, "From local explanations to global understanding with explainable AI for trees," *Nat. Mach. Intell. 2020 21*, vol. 2, no. 1, pp. 56–67, Jan. 2020, doi: 10.1038/s42256-019-0138-9.
[22] Corne van Zyl, Xianming Ye, "Harnessing eXplainable artificial intelligence for feature selection in time series energy forecasting: A comparative analysis of Grad-CAM and SHAP," *Appl. Energy*, vol. 353, p. 122079, 2024, doi: https://doi.org/10.1016/j.apenergy.2023.122079.
[23] Paolo Giudici, Alessandro Piergallini, "Explainable Artificial Intelligence methods for financial time series," *Phys. A Stat. Mech. its Appl.*, vol. 655, p. 130176, 2024, doi: https://doi.org/10.1016/j.physa.2024.130176.
[24] Phan Tien Dung, Paolo Giudici, "Sustainability, Accuracy, Fairness, and Explainability (SAFE) Machine Learning in Quantitative Trading," *Mathematics*, vol. 13, no. 3, p. 442, 2025, doi: 10.3390/math13030442.
[25] Wajih Abdallah, Majed Alruwaili, "DeepXFusion: A Hybrid Representation Learning and Explainability Framework for Trustworthy Machine Learning," *Proc. 18th Int. Conf. Agents Artif. Intell.*, vol. 4, pp. 3438–3445, 2026, [Online]. Available:

https://www.scitepress.org/Papers/2026/144455/144455.pdf
[26] Mayaz Uddin Gazi, "Flood Prediction and Forecasting Using Anns and Fuzzy Logic Model in Sylhet, Bangladesh," *Sciety*, 2025, [Online]. Available: https://sciety.org/articles/activity/10.21203/rs.3.rs-6882638/v1
[27] G. S. Deepan and T. Subbulakshmi, "Financial Planning and Decision Making for Students Using Fuzzy Logic," *Fuzzy Log. Appl. Comput. Sci. Math.*, pp. 159–171, Jan. 2023, doi: 10.1002/9781394175130.CH10;CTYPE:STRING:BOOK.
[28] A. -, V. BHATTACHARJEE, and P. S. Bishnu, "Stock Price Forecasting Using Fuzzy Logic-Enhanced Deep Learning Models," 2024, doi: 10.2139/SSRN.5023575.
[29] Emine Yüksek Dizdar, "Comparison of Fuzzy, Intuitionistic Fuzzy, and Picture Fuzzy Regression Function Approaches for Forecasting the Dow Jones Financial Time Series," *Turkish J. Forecast.*, vol. 10, no. 1, 2026, doi: 10.34110/forecasting.1868623.
[30] Miin Jong Hao, Yu Xuan Zheng, "Highway Travel-Time Forecasting with Greenshields Model-Based Cascaded Fuzzy Logic Systems," *Appl. Sci.*, vol. 15, no. 14, p. 7729, 2025, doi: https://doi.org/10.3390/app15147729.