

Article

Not peer-reviewed version

Hybrid LVQ-Entropy Model for Robust Algorithmic Trading

[Florentin Șerban](#)* and Bogdan Petru Vrinceanu

Posted Date: 9 June 2025

doi: 10.20944/preprints202506.0649.v1

Keywords: Algorithmic trading; Machine learning; Shannon entropy; Learning Vector Quantization (LVQ); Entropy filtering; Quantitative finance



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Hybrid LVQ-Entropy Model for Robust Algorithmic Trading

Florentin Șerban ^{1,*} and Bogdan-Petru Vrînceanu ²

¹ Department of Applied Mathematics, Bucharest University of Economic Studies, 010374 Bucharest, Romania

² Master Business Analytics, Bucharest University of Economic Studies, 010374 Bucharest, Romania

* Correspondence: florentin.serban@csie.ase.ro

Abstract: This paper proposes a novel entropy-enhanced machine learning framework for algorithmic trading, designed to improve decision reliability and model interpretability under conditions of financial uncertainty. The core methodology integrates a Learning Vector Quantization (LVQ) classifier with a Shannon entropy-based filter that selectively activates trading signals based on their informational certainty. Mathematically, the system is formulated as a composite decision function constrained by entropy: only predictions satisfying $H(x_t) \leq \theta$ are executed, where $H(x_t)$ is the Shannon entropy of the class probability distribution at time t . The empirical evaluation spans three asset classes—cryptocurrencies, foreign exchange, and equities—demonstrating significant improvements in trade precision, Sharpe ratio, and robustness to market volatility. The entropy filter acts as a regularization mechanism, reducing overfitting and mitigating the propagation of noisy signals. Beyond its empirical success, the model contributes a formalized framework for entropy-regularized decision making, opening avenues for generalized entropic models such as Tsallis and Kaniadakis entropy in future research. This entropy-aware structure bridges statistical learning theory and real-world financial applications by introducing a quantifiable measure of confidence into the machine learning pipeline. The results underscore the model's potential as a robust, interpretable, and adaptive solution for algorithmic trading in complex and uncertain markets.

Keywords: algorithmic trading; machine learning; Shannon entropy; learning vector quantization (LVQ); entropy filtering; quantitative finance

MSC: 91G60; 94A17; 68T05

1. Introduction

The integration of machine learning (ML) techniques into financial forecasting and algorithmic trading has gained significant momentum in recent years, driven by both technological advancements and the increasing availability of high-frequency market data. In contrast to traditional rule-based systems, machine learning models are capable of identifying latent patterns in financial time series, allowing for more adaptive and data-driven decision-making frameworks [1,2]. However, a major challenge persists: many ML-based strategies remain vulnerable to overfitting and lack robustness when confronted with noisy or highly volatile market conditions [3,4]. To address this limitation, recent studies have investigated the role of information-theoretic measures—particularly entropy—as a means of quantifying uncertainty and filtering noisy signals in real-time financial environments [5,6].

Entropy, originally rooted in statistical mechanics and information theory, has found increasing applications in fields ranging from anomaly detection to portfolio optimization [7,8]. Specifically, Shannon entropy provides a rigorous yet intuitive metric for capturing the randomness or disorder present in a distribution of price movements, volatility regimes, or technical indicators [9,10]. This

versatility has made entropy-based approaches particularly appealing in financial contexts where uncertainty is pervasive and traditional statistical assumptions are often violated

In algorithmic trading, this has led to the emergence of entropy-based strategies aimed at filtering out low-confidence predictions and enhancing model interpretability [11–13]. When integrated with machine learning classifiers, entropy measures can guide the selection of high-confidence decisions, effectively reducing the incidence of false positives and improving risk-adjusted returns [14,15]. Recent research has demonstrated the effectiveness of entropy-enhanced trading systems across multiple ML architectures, including support vector machines (SVM), random forests, and deep learning models [16–18]. Among the variety of available classification methods, Learning Vector Quantization (LVQ) has proven particularly suitable for financial applications due to its simplicity, low computational overhead, and interpretability [19,20]. As a prototype-based supervised learning algorithm, LVQ excels in binary classification tasks such as trend detection or directional trading decisions.

Nonetheless, its performance can degrade under high uncertainty unless complemented by an additional filtering mechanism. This study aims to bridge this gap by proposing a hybrid decision-making framework that integrates Shannon entropy as a filtering mechanism within the LVQ classification pipeline, tailored to noisy and volatile market settings. By doing so, we not only improve prediction reliability and system robustness, but also contribute to the broader discourse on entropy-driven methodologies in algorithmic trading, where information theory and machine learning converge to offer novel, adaptive solutions for financial decision-making [21]. This direction aligns with recent explorations of hybrid decision-making architectures combining statistical filtering and AI models for automated agents, as proposed in [22].

The main objective of this study is to propose and empirically validate a hybrid framework that combines LVQ with entropy-based filtering to improve the reliability of trade signal generation. Unlike prior works that apply entropy in isolation or as a heuristic indicator, our approach integrates Shannon entropy as a formal probabilistic filter directly within the machine learning pipeline. The novelty of this framework lies in its ability to dynamically assess market uncertainty and suppress low-confidence signals before execution. The contributions of this paper are threefold:

1. We develop a mathematical framework that integrates entropy-based uncertainty modeling with prototype-based classification via LVQ.
2. We empirically demonstrate that entropy filtering significantly improves performance metrics—including Sharpe ratio, win rate, and risk-reward balance—across multiple asset classes and timeframes.
3. We position this model as a generalizable approach for robust financial decision-making under uncertainty, offering insights applicable to both algorithmic trading and broader ML-based forecasting problems.

Taken together, these contributions highlight the potential of entropy-informed learning systems to enhance both the interpretability and resilience of algorithmic trading models, especially in environments characterized by noise, volatility, and structural complexity.

The rest of the paper is organized as follows. Section 2 introduces the theoretical background of LVQ classification and Shannon entropy. Section 3 presents the empirical setup and experimental design. Section 4 reports the results and comparative evaluations. Section 5 discusses the implications, and Section 6 concludes with potential extensions for future research.

2. Materials and Methods

2.1. Learning Vector Quantization (LVQ)

Learning Vector Quantization (LVQ) is a supervised machine learning algorithm based on prototype-driven classification, first introduced as an extension of the Self-Organizing Map (SOM). Unlike black-box classifiers such as deep neural networks, LVQ operates through interpretable decision boundaries formed by representative vectors (prototypes) assigned to each class [23].

In the context of binary classification—such as predicting long versus short market positions—LVQ partitions the input space using a set of labeled prototype vectors $\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_k\}$ each associated with a specific class. For any input vector $\mathbf{x} \in \mathbf{R}^n$, the classification is determined by identifying the closest prototype in Euclidean distance: $\text{class}(\mathbf{x}) = \text{label}(\min_j \|\mathbf{x} - \mathbf{w}_j\|)$

During training, the algorithm iteratively adjusts the prototypes based on the labeled input data. If the input is correctly classified, the closest prototype is moved closer to the input vector. Conversely, if misclassified, the prototype is updated in the opposite direction. The update rule for a prototype \mathbf{w}_j at iteration t is: $\mathbf{w}_j^{(t+1)} = \mathbf{w}_j^{(t)} + \eta(t) \cdot \delta \cdot (\mathbf{x} - \mathbf{w}_j^{(t)})$ where: $\eta(t) \in (0,1)$ is the learning rate, which typically decays over time; $\delta = +1$ if the classification is correct, $\delta = -1$ otherwise. This learning dynamic enables LVQ to construct flexible, non-linear decision boundaries while maintaining low computational complexity. Moreover, LVQ is particularly suited for financial time series classification due to its ability to incorporate engineered features such as momentum indicators, volatility measures, and volume-based signals [23].

In practical trading applications, LVQ can be trained to distinguish between favorable and unfavorable market conditions using historical labeled data. Each input vector may consist of normalized technical indicators such as Relative Strength Index (RSI), Commodity Channel Index (CCI), Rate of Change (ROC), and volatility metrics. By learning prototypical configurations that precede profitable outcomes, LVQ facilitates directional trade decisions that are both data-driven and transparent.

However, as with most classification algorithms, LVQ is susceptible to degraded performance under high uncertainty or regime shifts. To address this, we propose enhancing the LVQ model by incorporating an entropy-based filtering layer—described in the next section—which selectively suppresses signals in disordered market environments where prediction confidence is low.

2.2. Shannon Entropy for Uncertainty Quantification

Shannon entropy, introduced in the context of information theory, provides a mathematically rigorous framework for quantifying uncertainty in probabilistic systems. Its core principle is that the level of uncertainty in a system is proportional to the unpredictability of its outcomes. Formally, for a discrete probability distribution $\mathbf{p} = (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n)$ over n possible events, Shannon entropy is defined as: $\mathbf{H}(\mathbf{p}) = -\sum_{i=1}^n \mathbf{p}_i \ln \mathbf{p}_i$

Entropy reaches its maximum value when the distribution is uniform—i.e., when all events are equally likely, indicating maximal uncertainty. Conversely, entropy is minimized (zero) when the outcome is known with certainty, i.e., when one of the \mathbf{p}_i equals 1 and the rest are zero. This property makes entropy particularly suitable for evaluating the reliability of machine learning classifiers, especially in binary classification tasks. In financial applications, market conditions are often characterized by varying degrees of randomness, noise, and structural instability. Traditional trading strategies may not account for these dynamics effectively, leading to overfitting, frequent misclassifications, and suboptimal decision-making.

To address this challenge, Shannon entropy can be leveraged as a filtering mechanism that suppresses trading signals generated under high-uncertainty regimes.

Specifically, when a classifier produces a set of posterior probabilities over possible trade directions—e.g., $\mathbf{p} = (\mathbf{p}_{long}, \mathbf{p}_{short})$ —the corresponding entropy quantifies the confidence level associated with the prediction. A highly uncertain output, such as $\mathbf{p} = (0.5, 0.5)$, yields maximum entropy and suggests that no dominant class has been identified. On the other hand, a confident prediction, such as $\mathbf{p} = (0.95, 0.05)$, generates low entropy, signaling that a clear directional bias has emerged.

In the present framework, we introduce an entropy threshold $\theta \in [0, \log n]$ to control the decision-making process. A trade signal is validated only when the entropy of the classifier output satisfies: $\mathbf{H}(\mathbf{p}) < \theta$

This entropy-filtering criterion enables the learning algorithm to become selectively active only under favorable market regimes. It acts as a protective layer, filtering out ambiguous or noisy inputs and thereby improving overall trade precision, reducing transaction costs, and mitigating exposure to drawdowns.

As will be demonstrated in the subsequent sections, incorporating entropy as a decision gate significantly improves risk-adjusted returns and helps stabilize performance across different financial instruments and market phases. The entropy-based filter is thus not merely a diagnostic tool but an integral component of a robust and adaptive classification architecture for algorithmic trading.

2.3. Hybrid Framework: LVQ with Entropy Filtering

The central contribution of this study lies in the development of a hybrid classification architecture that integrates Learning Vector Quantization (LVQ) with Shannon entropy-based filtering to improve decision quality in financial trading environments. While LVQ provides a simple yet effective prototype-based classifier suitable for noisy and non-linear financial data, the addition of an entropy filter allows the system to operate selectively by suppressing decisions made under high uncertainty.

LVQ functions by assigning class labels based on the proximity of input vectors to a set of representative prototypes. Let $\mathbb{X} \subset \mathbf{R}^d$ denote the input space of feature vectors derived from historical market data, and let $\mathbb{C} = \{\text{long}, \text{short}\}$ denote the binary classification space. During the training phase, the LVQ algorithm optimizes prototype vectors $\mathbf{w}_i \in \mathbb{X}$ such that for each new input $\mathbf{x} \in \mathbb{X}$, the predicted label corresponds to the class of the closest prototype under a chosen distance metric, typically the Euclidean norm.

The classification output can be represented probabilistically via soft assignments, yielding a probability vector $\mathbf{p} = (\mathbf{p}_{\text{long}}, \mathbf{p}_{\text{short}})$ reflecting the model's confidence in each possible trading decision. These probabilities serve as the input to the entropy filtering mechanism.

The filtering process is governed by a scalar threshold parameter θ , which is used to evaluate the Shannon entropy of the output distribution: $\mathbf{H}(\mathbf{p}) = -\sum_{i=1}^n \mathbf{p}_i \ln \mathbf{p}_i$.

If $\mathbf{H}(\mathbf{p}) < \theta$ the prediction is considered sufficiently confident and is passed to the execution engine; otherwise, it is discarded. This leads to a modular two-stage classification

1. Stage 1 – LVQ Classification: Input vector \mathbf{x} is classified using trained LVQ prototypes, resulting in a soft label \mathbf{p} .
2. Stage 2 – Entropy Filtering: The entropy of \mathbf{p} is computed. The decision is validated only if $\mathbf{H}(\mathbf{p}) < \theta$.

Formally, the output decision function $\mathbf{f}: \mathbb{X} \rightarrow \mathbb{C} \cup \{\emptyset\}$ can be written as:

$$\mathbf{f}(\mathbf{x}) = \begin{cases} \arg \max_i \mathbf{p}_i & \text{if } \mathbf{H}(\mathbf{p}) < \theta \\ \emptyset, & \text{otherwise} \end{cases}$$

where \emptyset denotes signal rejection due to excessive uncertainty.

This entropy-augmented framework enhances the selectivity of LVQ by acting as a gatekeeper that filters out ambiguous predictions. Such a design is particularly advantageous in high-frequency trading systems, where market noise can significantly degrade model performance if left unchecked. The architecture is general and may be extended to other probabilistic classifiers or entropy variants, such as Rényi or Tsallis entropy, to further tune the model's sensitivity to uncertainty. Experimental validation of this framework is presented in Section 3, demonstrating consistent improvements in classification precision, return metrics, and robustness across multiple assets and market regimes.

3. Results and Discussions

3.1. Experimental Setup and Data

To empirically validate the proposed hybrid classification framework, we conducted a series of experiments on diverse financial instruments that span multiple asset classes and volatility regimes. The selected assets include Bitcoin (BTC/USD), the Euro/US Dollar exchange rate (EUR/USD), and the Nasdaq 100 equity index (NDX). These assets were chosen for their contrasting market dynamics—ranging from the high volatility of cryptocurrencies to the relative stability of major forex pairs and equity indices.

The dataset covers the period from January 1st to March 31st, 2025, using 5-minute intervals for BTC/USD, 15-minute intervals for EUR/USD, and hourly data for NDX. Historical data were obtained via the [TradingView API], which provides high-frequency candle data and associated technical indicators. Each asset's time series was transformed into a structured dataset of labeled examples for classification purposes.

The input features used to train the LVQ classifier were selected based on their ability to capture trend, momentum, and volatility patterns. The final feature vector included the following indicators: Relative Strength Index (RSI) – 14-period; Rate of Change (ROC) – 10-period; Bollinger Band Width (BBW) – based on a 20-period moving average Stochastic Oscillator (%K, %D); Price-derived features – such as log returns, moving average crossovers, and volatility estimates. Each feature vector was labeled using a forward-looking return threshold over a prediction horizon of $h = 3$ candles (i.e., 15 min for BTC, 45 min for EUR/USD). A binary label was assigned as follows: “Long” if the asset's price increased by more than +0.3% within the horizon; “Short” if the price decreased by more than –0.3%. Cases within $\pm 0.3\%$ were discarded during training to reduce class overlap.

The dataset was split into 70% for training and 30% for testing. For each asset, we trained a separate LVQ model using the Kohonen algorithm with Euclidean distance and adaptive learning rate decay. The model parameters (number of prototypes, epochs, learning rate schedule) were selected via grid search using 5-fold cross-validation to maximize classification precision. The Shannon entropy filter was applied post-classification. The entropy threshold θ was empirically tuned based on the validation set to achieve an optimal balance between selectivity and coverage. Thresholds ranged from 0.35 to 0.75 depending on the asset and signal noise level. The experimental pipeline was implemented in Python using Scikit-learn, NumPy, and custom modules for entropy computation and trade simulation. Performance metrics including net return, Sharpe ratio, win rate, and maximum drawdown were computed using a fixed trade size per signal and zero leverage, assuming a 0.05% transaction cost per trade. This setup enables a robust comparison between baseline LVQ models and the proposed entropy-augmented framework under realistic market conditions. The results are discussed in the following subsections. With the experimental design and feature engineering stages completed, the next step involves evaluating the predictive performance and robustness of the proposed framework. To this end, we implemented a systematic comparison between the baseline LVQ classifier and the entropy-augmented variant across multiple financial assets. The results, presented in the following sections, offer a comprehensive perspective on the effectiveness of entropy as a filtering mechanism for improving decision reliability and trading outcomes.

3.2. Performance Metrics

The empirical evaluation focused on comparing the performance of the baseline LVQ classifier and the entropy-augmented LVQ+Entropy framework across three asset classes. Performance metrics included Net Return (%), Win Rate (%), Sharpe Ratio, Maximum Drawdown, and Precision (classification accuracy on test data). Results were averaged over five Monte Carlo runs with randomized training/test splits to ensure statistical robustness.

Table 1. Performance Metrics under Baseline and Entropy-Augmented LVQ Models.

Asset	Model	Net Return (%)	Win Rate (%)	Sharpe Ratio	Max Drawdown (%)	Precision (%)
BTC/USD	LVQ	−28.1	47.5	−0.41	22.4	53.2
BTC/USD	LVQ + Entropy	+14.3	62.5	0.94	11.7	67.8
EUR/USD	LVQ	−6.7	49.2	−0.18	8.9	50.4
EUR/USD	LVQ + Entropy	+4.6	58.1	0.51	4.2	61.2
Nasdaq 100	LVQ	−3.4	50.7	−0.09	6.7	52.9
Nasdaq 100	LVQ + Entropy	+6.9	59.0	0.63	3.5	64.7

Interpretation of Table

The empirical results summarized in Table X illustrate the substantial impact of entropy filtering on trading performance across various asset classes. For BTC/USD, the baseline LVQ model yielded a net loss of −28.1% and a Sharpe ratio of −0.41, indicating high volatility and poor return-to-risk efficiency. After entropy integration, the net return improved to +14.3%, with a Sharpe ratio of 0.94, reflecting a more balanced and profitable strategy. Similar improvements are noted for EUR/USD and Nasdaq 100, where win rate, precision, and maximum drawdown also benefited from the entropy-informed approach. These findings support the hypothesis that Shannon entropy acts as an effective probabilistic gate, enhancing both profitability and robustness in machine learning-based trading

- Remarks:
1. Profitability Reversal: On all assets, the baseline LVQ model produced either negative or near-zero net returns, whereas the entropy-filtered strategy consistently produced positive returns—most notably for BTC/USD, where the strategy shifted from a 28% loss to a 14.3% gain.
 2. Sharpe Ratio Improvement: Sharpe ratios improved significantly, with the BTC strategy rising from −0.41 to +0.94, indicating a substantial enhancement in risk-adjusted performance.
 3. Drawdown Reduction: Entropy filtering led to a meaningful decrease in maximum drawdown, highlighting the strategy’s ability to avoid false signals in volatile regimes.
 4. Precision vs Coverage Trade-off: Although the entropy filter reduced the number of executed trades by approximately 25–35%, the precision of the predictions increased by over 10% across all datasets, validating the efficacy of entropy as a confidence filter.
 5. Consistency Across Markets: The improvements were not isolated to a single asset class. Even in low-volatility environments like EUR/USD, the entropy-enhanced system yielded better performance across all metrics.

These findings underscore the utility of integrating entropy-based filters into ML pipelines, particularly in domains such as financial trading where signal quality and overfitting remain key challenges.

These quantitative outcomes lay the foundation for a deeper analysis of the model’s behavior under varying market conditions, which will be further explored in the following section

3.3. Discussion and Interpretation

The results obtained from the empirical tests highlight the critical role of entropy in enhancing the reliability and interpretability of machine learning models in financial decision-making. While

standard classifiers such as Learning Vector Quantization (LVQ) offer a framework for supervised learning in noisy time-series environments, they often suffer from overfitting and overtrading—especially when applied to high-frequency financial data.

By integrating a Shannon entropy-based filter, the proposed approach introduces a quantitative uncertainty control mechanism that evaluates the confidence of each classification before execution. The filtering mechanism acts as a *selectivity gate*, ensuring that only those predictions with sufficiently low uncertainty are allowed to generate trades. This leads to two key improvements: (i) an increase in the signal-to-noise ratio of the trading decisions, and (ii) a reduction in the number of suboptimal trades, which directly contributes to better performance and reduced drawdowns. The improvement in Sharpe ratios and precision metrics across all tested asset classes reinforces the hypothesis that entropy serves as a proxy for market clarity. In periods of high market turbulence or regime shifts, the entropy filter tends to block ambiguous signals, effectively mimicking a form of risk aversion. From a practical standpoint, this aligns with real-world trading behavior, where professional investors are more likely to avoid decisions under extreme uncertainty. Moreover, the observed consistency of results across three distinct markets—cryptocurrencies, foreign exchange, and equities—demonstrates the generality of the entropy-enhanced model. The framework does not rely on domain-specific assumptions or handcrafted rules, making it easily extendable to other financial instruments or time horizons.

The interpretability of the entropy-filtered decisions is another valuable feature. Since the model explicitly quantifies uncertainty, it provides an additional layer of transparency often missing in complex machine learning systems. This enhances the trustworthiness of the model's outputs, which is particularly important for real-world deployment in trading systems governed by regulatory or risk constraints.

In summary, the combination of LVQ and entropy filtering results in a hybrid classification system that is not only more accurate but also more robust and interpretable. These characteristics are crucial in financial applications, where false positives can be costly and explainability is increasingly demanded.

Taken together, these results highlight the mathematical viability and empirical strength of integrating Shannon entropy into the LVQ classification process. By functioning as a probabilistic constraint, entropy enhances the system's robustness against noise and overfitting, while ensuring a higher discriminative efficiency. The improvements in Sharpe ratio, precision, and drawdown reduction validate this hybrid architecture across different asset classes. These insights not only strengthen the theoretical rationale behind entropy-based filtering but also lay the foundation for broader algorithmic implementations. The following section concludes the study and outlines several promising avenues for future research and model extensions.

4. Conclusions and Future Work

This study proposed and evaluated a novel machine learning framework for algorithmic trading, which integrates Shannon entropy filtering within a Learning Vector Quantization (LVQ) classification pipeline. The core innovation lies in its capacity to selectively activate trading decisions based on the quantified uncertainty of predictions, thereby enhancing both robustness and interpretability.

Empirical evaluations across three distinct asset classes—cryptocurrencies (BTC/USD), forex (EUR/USD), and equities (Nasdaq 100)—demonstrated that the entropy-filtered LVQ model consistently outperforms the baseline version. Notable improvements were observed in Sharpe ratio, win rate, and risk-reward balance across various timeframes, supporting the model's generalizability in diverse market conditions.

Beyond numerical performance, the entropy component functioned as both a filter and an interpretability layer. It offered insight into the model's confidence in each prediction, enabling practitioners to make more informed and risk-sensitive decisions. These findings contribute to the

expanding body of research that combines machine learning with information-theoretic principles to design more adaptive and transparent trading systems.

Importantly, this approach addressed several structural weaknesses common in traditional classification-based trading frameworks. It reduced overtrading by suppressing ambiguous signals, mitigated drawdowns in volatile markets, and enhanced real-time applicability through quantifiable confidence scoring. Moreover, LVQ's prototype-based structure strengthened model interpretability—a growing priority in the field of explainable AI.

Theoretically, the paper offers two key contributions. First, it formulates entropy as a dynamic threshold mechanism for filtering within machine learning pipelines. Second, it demonstrates that entropy-based regularization enhances classification models by explicitly accounting for prediction uncertainty—particularly relevant in financial domains marked by noise, volatility, and structural instability.

Future Research Could Explore Several Extensions

-Generalized Entropic Models: Applying non-Shannon formulations such as Tsallis, Rényi, or Kaniadakis entropy could yield improved adaptability in nonlinear and heavy-tailed environments.

-Classifier Diversity: While LVQ was selected for its interpretability, testing this entropy filtering approach with alternative classifiers (e.g., SVM, XGBoost, DNNs) may uncover new synergistic behaviors.

-Multi-Objective Learning: Integrating entropy into optimization schemes balancing multiple objectives—such as accuracy, risk control, and trade frequency—may offer richer trading heuristics.

-Real-Time Implementation: Deploying the system in live markets, with adaptive learning and feedback mechanisms, would help validate its operational viability under latency and slippage constraints.

-Entropy as a Regime Identifier: Given its tendency to spike during structural shifts, entropy may serve as a market regime indicator, offering predictive value for portfolio rotation and rebalancing strategies.

In conclusion, this work lays both a conceptual and empirical foundation for entropy-informed algorithmic trading systems. By bridging statistical learning with probabilistic filtering, the proposed framework contributes to the evolution of interpretable, uncertainty-aware financial AI tools suitable for complex and dynamic markets.

References

1. Bao, Y.; Ke, B.; Li, B.; Yu, Y.J.; Zhang, J. (2020). Detecting Accounting Fraud in Publicly Traded U.S. Firms Using a Machine Learning Approach. *Journal of Accounting Research*, 58, 199–235. <https://doi.org/10.1111/1475-679x.12292>
2. Gerlein, E.A.; McGinnity, M.; Belatreche, A.; Coleman, S. (2016). Evaluating Machine Learning Classification for Financial Trading: An Empirical Approach. *Expert Systems with Applications*, 54, 193–207. <https://doi.org/10.1016/j.eswa.2016.01.018>
3. Dash, R.; Dash, P.K. (2016). A Hybrid Stock Trading Framework Integrating Technical Analysis with Machine Learning Techniques. *Journal of Finance and Data Science*, 2, 42–57. <https://doi.org/10.1016/j.jfds.2016.03.002>
4. Han, Y.; Kim, J.; Enke, D. (2023). A Machine Learning Trading System for the Stock Market Based on N-Period Min-Max Labeling Using XGBoost. *Expert Systems with Applications*, 211, 118581. <https://doi.org/10.1016/j.eswa.2022.118581>
5. Bat-Erdene, M.; Kim, T.; Park, H.; Lee, H. (2017). Packer Detection for Multi-Layer Executables Using Entropy Analysis. *Entropy*, 19, 125. <https://doi.org/10.3390/e19030125>
6. Chen, J.; Dou, Y.; Li, Y.; Li, J. (2016). Application of Shannon Wavelet Entropy and Shannon Wavelet Packet Entropy in Analysis of Power System Transient Signals. *Entropy*, 18, 437. <https://doi.org/10.3390/e18120437>
7. Hao, D.; Li, Q.; Li, C. (2017). Digital Image Stabilization Method Based on Variational Mode Decomposition and Relative Entropy. *Entropy*, 19, 623. <https://doi.org/10.3390/e19110623>

8. Kong, L.; Pan, H.; Li, X.; Ma, S.; Xu, Q.; Zhou, K. (2019). An Information Entropy-Based Modeling Method for the Measurement System. *Entropy*, 21, 691. <https://doi.org/10.3390/e21070691>
9. Pfleger, M.; Wallek, T.; Pfennig, A. (2014). Constraints of Compound Systems: Prerequisites for Thermodynamic Modeling Based on Shannon Entropy. *Entropy*, 16, 2990–3008. <https://doi.org/10.3390/e16062990>
10. Ibl, M.; Čapek, J. (2016). Measure of Uncertainty in Process Models Using Stochastic Petri Nets and Shannon Entropy. *Entropy*, 18, 33. <https://doi.org/10.3390/e18010033>
11. García-Martínez, B.; Martínez-Rodrigo, A.; Zangróniz Cantabrana, R.; Pastor García, J.; Alcaraz, R. (2016). Application of Entropy-Based Metrics to Identify Emotional Distress from Electroencephalographic Recordings. *Entropy*, 18, 221. <https://doi.org/10.3390/e18060221>
12. Rodríguez-Rodríguez, N.; Miramontes, O. (2022). Shannon Entropy: An Econophysical Approach to Cryptocurrency Portfolios. *Entropy*, 24, 1583. <https://doi.org/10.3390/e24111583>
13. Rasoul, S.; Samimi, A.J.; Paydar, M.M. (2021). Newcomers' Priorities in Portfolio Selection: A Shannon Entropy Approach. *Iranian Economic Review*, 25. <https://doi.org/10.22059/ier.2021.82849>
14. Serban F., Ștefănescu M.V., Ferrara M. (2011). Portfolio optimization and building of its efficient frontier. *Economic Computation and Economic Cybernetics Studies and Research*, 54(3), 125–137.
15. Wüstenfeld, J.; Geldner, T. (2022). Economic Uncertainty and National Bitcoin Trading Activity. *North American Journal of Economics and Finance*, 59, 101625. <https://doi.org/10.1016/j.najef.2021.101625>
16. Paiva, F.D.; Cardoso, R.T.N.; Hanaoka, G.P.; Duarte, W.M. (2019). Decision-Making for Financial Trading: A Fusion Approach of Machine Learning and Portfolio Selection. *Expert Systems with Applications*, 115, 635–655. <https://doi.org/10.1016/j.eswa.2018.08.003>
17. Ryś, P.; Ślepaczuk, R. (2019). Machine Learning Methods in Algorithmic Trading Strategy Optimization – Design and Time Efficiency. *Central European Economic Journal*, 5, 206–229. <https://doi.org/10.1515/ceej-2018-0021>
18. Cocco, L.; Tonelli, R.; Marchesi, M. (2019). An Agent-Based Artificial Market Model for Studying the Bitcoin Trading. *IEEE Access*, 7, 42908–42920. <https://doi.org/10.1109/access.2019.2907880>
19. Pourghasemi, H.R.; Gayen, A.; Lasaponara, R.; Tiefenbacher, J.P. (2020). Application of Learning Vector Quantization and Different Machine Learning Techniques to Assessing Forest Fire Influence Factors and Spatial Modelling. *Environmental Research*, 184, 109321. <https://doi.org/10.1016/j.envres.2020.109321>
20. Zhou, R.; Xiong, X.; Llacay, B.; Pfeffer, G. (2023). Market Impact Analysis of Financial Literacy among A-Share Market Investors: An Agent-Based Model. *Entropy*, 25, 1602. <https://doi.org/10.3390/e25121602>
21. Tucci, G.; Vega, M. (2016). Optimal Trading Trajectories for Algorithmic Trading. *Journal of Investment Strategies*, 5, 57–74. <https://doi.org/10.21314/jois.2016.065>
22. Șerban, F.; Vrînceanu, B.; Stana, A. (2023). The Dirty Little Robot: A Hybrid Framework for Automated Decision-Making under Uncertainty. *Journal of Applied Quantitative Methods*, 17(2–4), 64–76. http://jaqm.ro/issues/volume-17,issue-2,3,4/4_FS.php
23. Kohonen, T. (1995). *Learning Vector Quantization*. In: Self-Organizing Maps. Springer Series in Information Sciences, vol 30. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-97610-0_9

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.