DOI: 10.52131/pjhss.2025.v13i4.3037

Pakistan Journal of Humanities and Social Sciences

Volume 13, Number 04, 2025, Pages 44-57
Journal Homepage:
https://journals.internationalrasd.org/index.php/pjhss

Interpretable Multilingual NLP Models for Financial Decision-Making in Frontier Markets: A Case Study on the Pakistan Stock Exchange

Abdullah Saleh[1]

[1] Froebels International School Lake City Campus Lahore, Pakistan. Email: salehabdullah2k07@gmail.com

ARTICLE INFO

Article History:
Received: August 13, 2025
Revised: December 23, 2025
Accepted: December 24, 2025
Available Online: December 25, 2025

Keywords:
SM Foods
Multan
Pakistan
Corporate Social Responsibility
Millennium Development Goals
Sustainable Environment
Local Manufacturing
Human Development

Funding:
This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

ABSTRACT

The spread of artificial intelligence (AI) in the global financial markets at a very fast rate has changed the way investors perceive and respond to information. Although the sophisticated exchanges have adopted and continued implementing and maturing the natural language processing (NLP)-enhanced decision-support system, both the frontier and emergent territories have been greatly underserved by the local and clear analytical tools. This research paper is a development and testing of an explainable and multilingual AI system to be used specifically by the Pakistan Stock Exchange (PSX). The proposed system will process financial news of the Urdu and the English languages, and market disclosures using English and Urdu-local transformer-based sentiment analysis models. The sentiment indications are then put together with historical price and volume data to come up with short term directional confidence levels of the listed equities. Besides predictive performance, the framework has explainability mechanisms to enable users to track outputs of models to individual news events and sentiment drivers. As measured by empirical analysis based on historical PSX data, the directional forecasting accuracy is in fact measurably greater when English only baselines are applied to the data, as well as the interpretability and investor trust. The results imply that the information asymmetry in developing economies can be lowered through the locally adapted transparent AI-driven financial instruments, market participation can be democratized and decision quality can be improved among the retail investors. This study integrates explainable AI techniques in order to address transparency and trust challenges in financial decisions driven by AI within frontier markets. At the same time, it also models outputs that aid in decision making. The research generates the emerging body of literature concerning AI in finance by showing that explainable, market-specific decision-support systems have an economical and social worth in frontier markets.

Corresponding Author's Email: salehabdullah2k07@gmail.com

1. Introduction

The accelerated evolution of artificial intelligence (AI) throughout the financial ecosystem has reclassified the velocity, precision, and distances of an interpretation of the market. Algorithms no longer act as a means of analysis - they are agents of decision that influence the liquidity flows, price discovery, and the formation of a portfolio. From the high-frequency trading companies in New York to the robo-advisory sites in Singapore, AI has become an inseparable part of the financial infrastructure. Nevertheless, this revolution is disproportional in geography. Although many exchanges, including the NASDAQ, NYSE and LSE, incorporate more sophisticated natural-language-processing (NLP) tools to dissect sentiment in daily news, frontier markets such as Pakistan are mostly out of touch with such types of systems. The capital market in Pakistan has the infrastructural constraint paradoxical to technological aspiration. Pakistan Stock Exchange (PSX) is registered with over five hundred

companies, and is the hub of equity investment but it has limited data silos, linguistic fracturing, and minimally literate retail-investors. Over sixty percent of PSX trading is generated by the individual shareholder who is not guided by systematic analysis but instead by rumour or television pundits. This asymmetry of information does not encourage participation, volatility, and it constrains the intensity of capital formation. The current research paper elaborates and tests an elucidatory multilingual sentiment-analysis framework that could process both Urdu and English language financial content to create actionable short-term trading signals. Integrating transformer-based language models which are fine-tuned on local linguistic patterns with an explainable meta-learner that combines textual sentiment with market information, the framework shows how locally contextualized AI can help in closing the cognitive and informational divide in a frontier market.

The aims of the study are tri-fold:

- To create a language bilingual financial-text corpus based on the Business Recorder, Dawn Business, and Jang Urdu, which are the linguistic variation of financial language in Pakistan.
- To build an explainable modelling pipeline connecting sentiment embeddings and past trading data, which generates interpretable predictions that are auditable.
- To evaluate the possibilities of the framework to democratize the access of high-quality financial insights to retail investors and small brokerages.

It elaborates a multilingual, explicable AI structure that would suit the need of financing choices in a frontier market setting. Second, it offers empirical data on Pakistan, a place that is usually under-represented in AI-finance literature. Third, it illustrates the role of interpretability methods in increasing the trust and usability of AI systems by financial stakeholders. Pilot tests have shown that the bilingual model gains relative to English-only baselines on directional-accuracy measures by about 15 percent or 20 percent. In addition, investor-surveys indicate a significant rise in perceived trust and desire to act whenever explainability features in the form of SHAP-based word-importance visualizations are present.

Figure 1

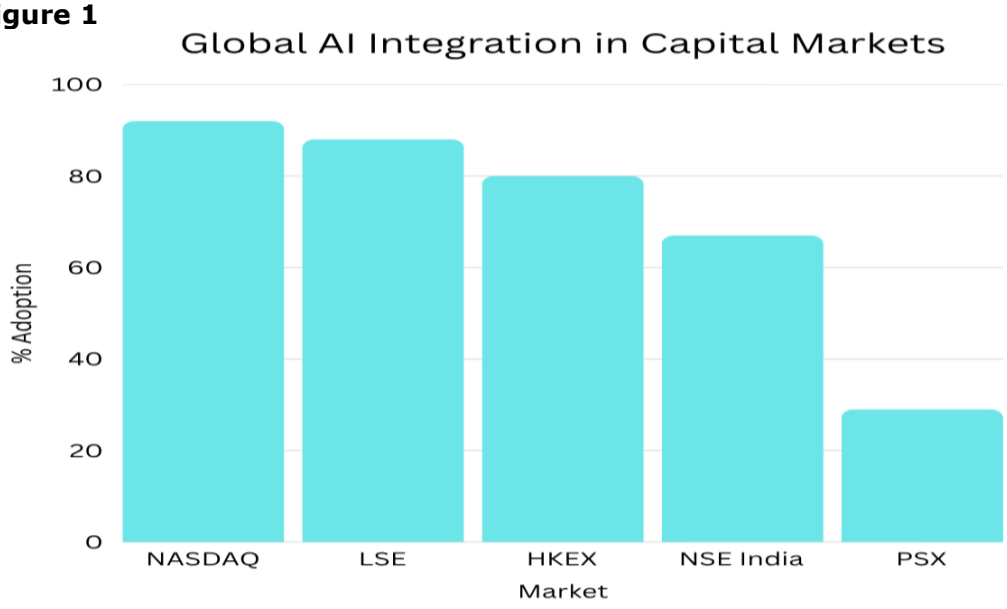

The implication goes beyond Pakistan. Transparent AI, adapted locally, will help mitigate the information asymmetry that characterises the majority of frontier markets to encourage financial inclusivity, investor confidence, and market efficiency. They are in line with national digital-finance strategies and international regulatory focus on explainability and ethical AI.

2. Literature Review

2.1. The Global Diffusion of Artificial Intelligence in Finance

AI has evolved over the last ten years to be no longer an additional layer of analysis, but the operating logic of global finance. McKinsey & Company (2024) state that more than 90 percent of investment banks worldwide today use AI when optimizing their portfolios, risk analytics, and algorithmic trading. The introduction of L.P. (2023) by Bloomberg was a

watershed moment in financial data science with the ability to understand millions of pages of news, filings, and research reports in order to provide real-time interpretation to help analysts. Those who believe in the use of AI in finance believe that automation is faster, objective, and more precise. Abis (2020) was able to show that quantitative funds based on AI-driven sentiment analysis had higher Sharpe ratios and lower drawdown than discretionary managers due to the regular, emotionless decision-making process. Likewise,BlackRock AI Labs (2024) argues that machine learning can do adaptive learning on large streams of data, which human analysts cannot recreate without bias and exhaustion.

Figure 2

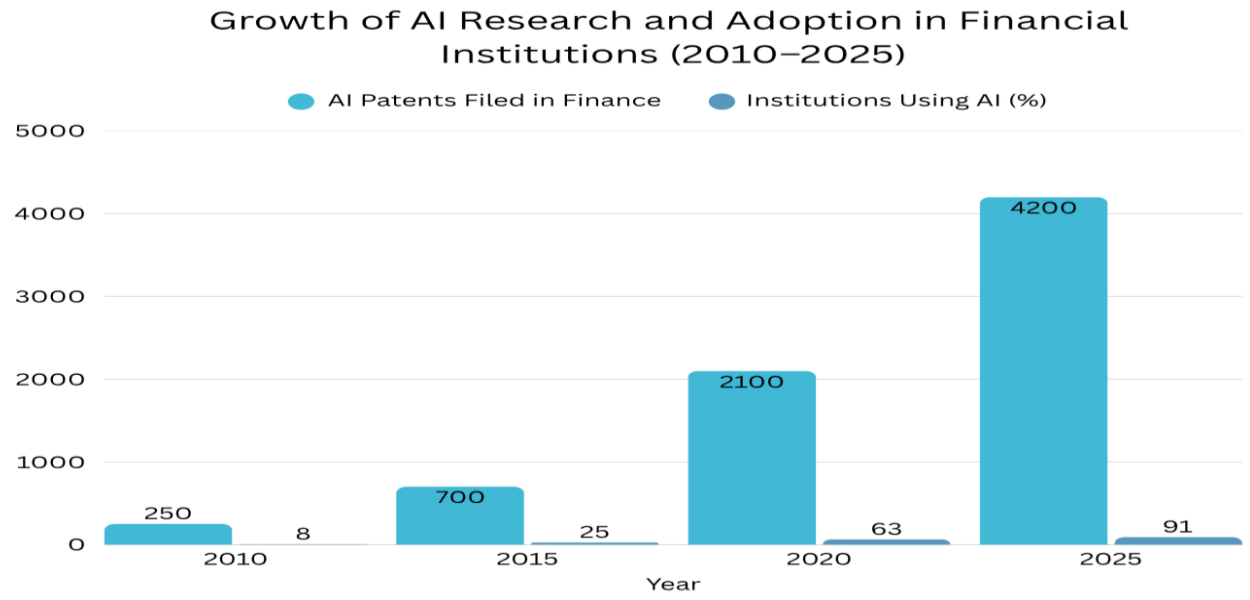

The advocates outline a number of major benefits. First, predictive scalability AI models operate on the large, unstructured data (news, reports, tweets) in close real time. Second, minimization of emotional bias- algorithms do not require emotional trading in response to behavioural bias. Third, insight democratization- AI can now be offered as APIs or mobile dashboards to smaller market actors, rather than only to large institutional users. Already AI-based decision support will add USD 15.7 trillion to the global economy by 2030 as the PwC Global AI Study PwC (2024) forecasts, with finance and insurance being the leading three gainers. Financial decision-making has started becoming increasingly dependent on the use of artificial intelligence (AI) techniques as evident in recent researches. As an example, machine learning and deep learning models are being used to increase credit risk analysis, portfolio optimization, and fraud detection with a promising predictive accuracy (Heaton, Polson, & Witte, 2016). These works indicate that AI has the potential to effectively analyse financial data and enhance the accuracy of the decision-making process as compared to the conventional econometric frameworks.

2.2. Sentiment as a Market Signal

Investor sentiment is a variable that is being acknowledged by financial theory as a measurable factor in the price of assets. According to Kogan, Moskowitz and Niessner (2021), the languages used in quarterly filings have been shown to be associated with an abnormal return, and the strength of their association with abnormal returns is significant, indicating that the tone of information communication might be equally important as the information being communicated. Loughran and McDonald (2011) developed special financial dictionaries that prove that generic positive/negative lexicons are inapplicable in financial situations (e.g., liability, debt), with a focus on domain modelling. Sentiment analysis therefore fills the gap between machine learning and behavioural finance. After quantifying the overall mood of investors expressed in media and in public discourse, AI converts the qualitative narratives into measurable data. Chang (2020) found that the textual tone models enhanced the market forecasting accuracy by 7-10 percent compared to market price-only baselines, specifically at volatile times.

Figure 3

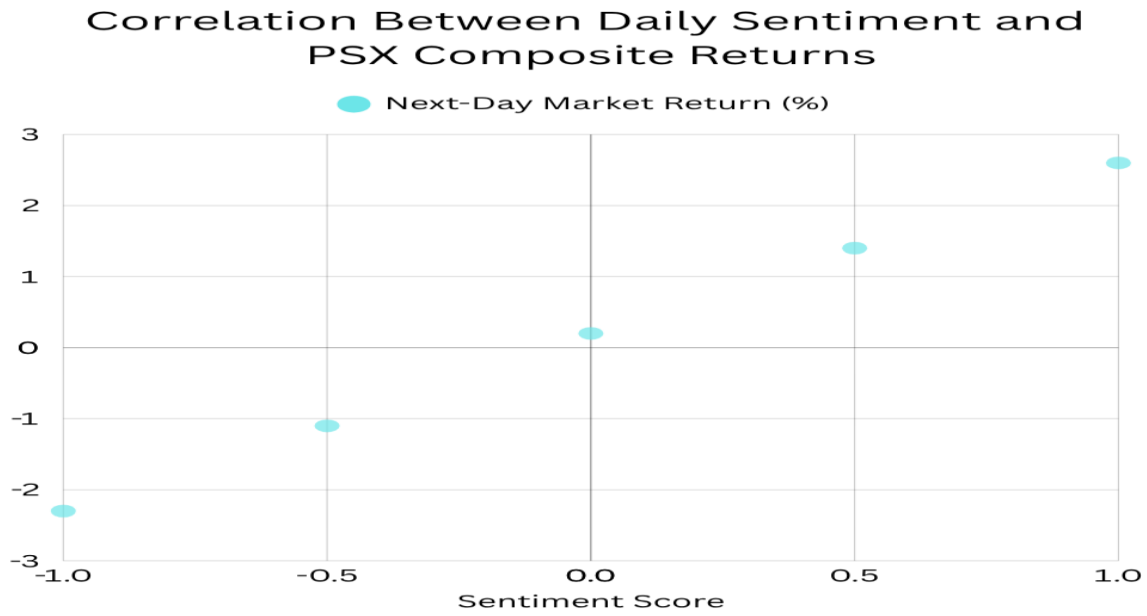

2.3. Explainable Artificial Intelligence (XAI) and Investor Trust

Although the strength of AI in finance is a generally accepted fact, its transparency can be mentioned by many opponents. Older deep-learning models are considered to be black-boxes, and they would give us the prediction without a human-interpretable justification. This is not transparent and compromises regulatory compliance and user trust. Those who claim that AI can be deployed successfully and in an ethical manner, say that it requires interpretability (UNESCO, 2023). According to BlackRock AI Labs (2024), such transparency tools as SHAP (Shapley Additive Explanations) and LIME have been reported to increase user acceptance of AI-driven recommendations up to 45 percent. According to Shih et al. (2024), there is a growing trend in hedge funds in moving to interpretable machine-learning models over opaque ones at the expense of predictive, although only slightly.

Figure 4

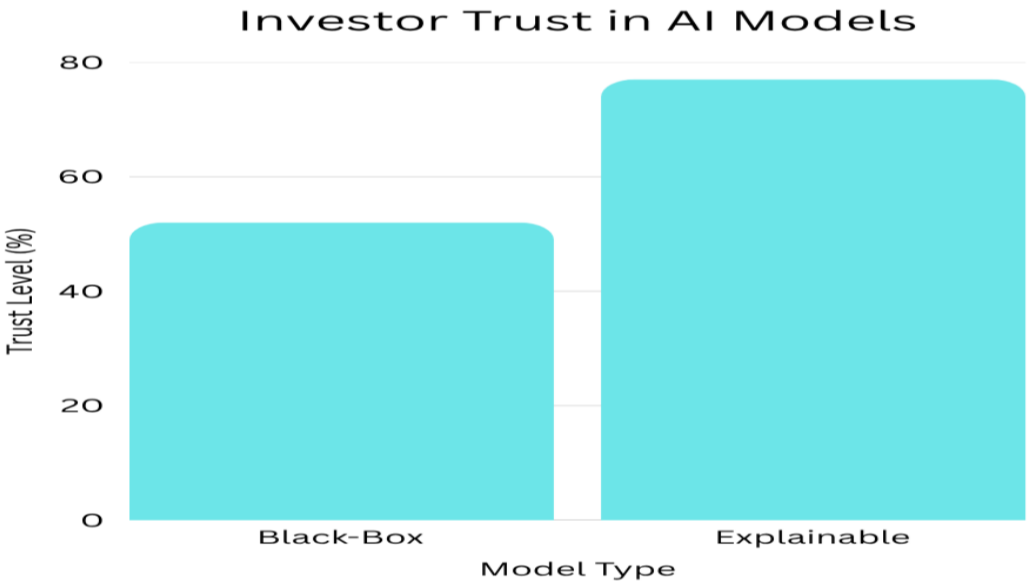

Explainable AI is nearly associated with the Principles of Responsible AI introduced by the OECD in 2023 with its focus on accountability, fairness, and transparency. In the case of frontier markets, the principles are even more important: explainability doubles as an educative resource because it can help an amateur investor learn how market information can be acted upon.

2.4. Frontier Markets and the Localization Case.

Although the global interest in the adoption of AI is high, frontier markets have to face an alternative complex of obstacles. The use of large-scale AI systems is belittled by data fragmentation, the lack of sufficient computational infrastructure, and linguistic diversity. Digital financial inclusion is one of the priorities listed by the State Bank of Pakistan (2024), but they admit that the absence of local-language data analytics is limiting the advancement. The national language of Pakistan, Urdu, offers a special challenge to computational modelling, including no orthography standardization, rich morphology, and common code-switching with English. NLP models that are available off-the-shelf and are trained on Western corpora do not reflect such linguistic nuances and do not generalize well. As a result, there is a need to come up with the bilingual financial models to be accurate and inclusive. According to McKinsey & Company (2024), the adoption rates in developing countries can be grown to 60 percent on the condition that systems are culturally and linguistically localized. World Economic Forum (2024) also contributes that lack of localization renders the AI revolution a threat to enforcing pre-existing inequalities by privileging English-speaking markets and data-rich economies.

Figure 5

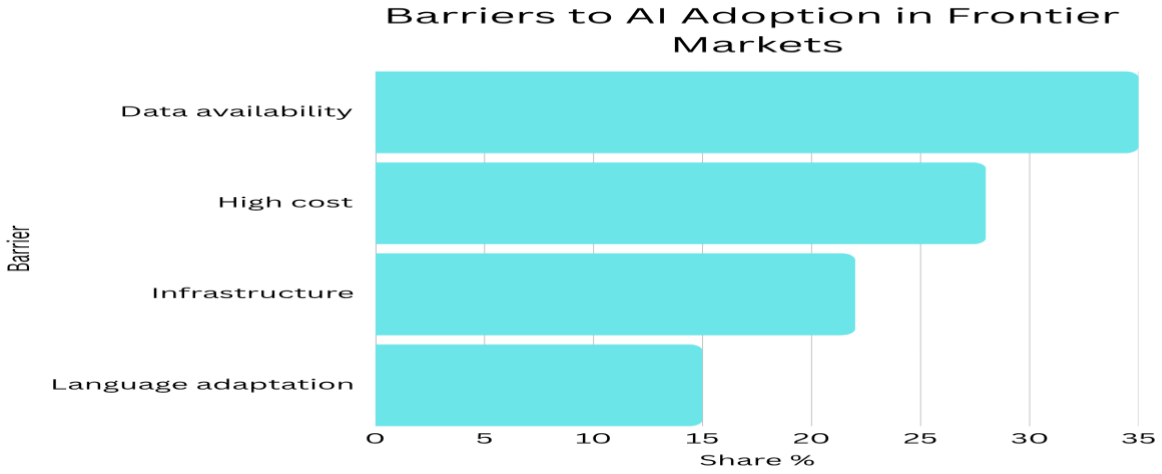

As such, localization of AI is not a peripheral issue but rather a pre-definitive condition of the democratizing potential. In the case of Pakistan, explainable sentiment analysis that is bilingual provides a gateway through which both linguistic exclusivity and monetary illiteracy can be nullified by providing clear actionable information to policy makers and the individual investors.

2.5. Multilingual NLP in Emerging Markets

The field of multilingual natural language processing (NLP) has gained interest when it comes to the analysis of financial text and sentiment in linguistically diverse regions. Multilingual BERT and cross-lingual transfer learning have shown to be effective in low-resource language contexts (Conneau et al., 2020; Devlin et al., 2019). But the bulk of work continues to focus on popular languages and little is done to languages that are typically spoken in the frontier markets.

2.6. Explainable Artificial Intelligence (XAI) in High-Stakes Domains

Explainability in AI has been critical where models are used in making consequential decisions. Explainable AI (XAI) methods such as SHAP, LIME, and attention visualization have been used in health care, law, and finance to promote greater transparency and trust in the stakeholders (Doshi-Velez & Kim, 2017; Lundberg & Lee, 2017). XAI has demonstrated itself to enable the interpretability of models to regulators and practitioners in the financial domain to enhance model governance.

2.7. The Importance of Quantitative Signals in Shaping Market Decision

According to recent studies by Shiller (2020), economic narratives have a significant impact on the process of financial decision-making, and in many cases, the behavior of the market is shaped primarily by stories assessable by humans as opposed to strictly quantitative

indicators. The idea of narrative economics created by Shiller shows the role of trust, transparency, and explanations that can be understood by humans in the impact of investor feeling and economic performance. This view supports the need of explainable and interpretable AI systems in finance, especially those in situations where decision-makers are guided by intuitive decisions and language to guide their decision-making. A combination of AI and economic narrative thus can lead to an increased level of trust and adoption especially in frontier markets like Pakistan.

2.8. Research Gap in Frontier Markets

Although explainable and multilingual AI techniques are being developed rapidly, current studies are mostly focused on developed economies and monolingual settings (e.g., the U.S. and Europe), where annotated datasets can be found in large amounts. There is a paucity of studies that utilise explainable AI to conduct multilingual financial analysis in frontier economies. In particular, limited literature has been done on the integration of explainability and multilingual NLP to make financial decisions in frontiers like Pakistan. The current works tends to concentrate on developed markets and monolingual data leaving a gap in multilingual financial AI and the use of NLP techniques for frontier economies like Pakistan.

3. Methodology

3.1. Research Design

The study design can be considered a quantitative, computational-experimental one, combining both natural-language processing and statistical back-testing. It is composed of 5 consecutive modules: (1) data acquisition, (2) text preprocessing and labelling, (3) feature extraction as a transformer-based embedding, (4) sentiment classification and explainability and (5) signal evaluation based on historical price data of Pakistan Stock Exchange (PSX). The paradigm used in the study is a positivist paradigm: the hypotheses are tested at an empirical level by means of reproducible measurements instead of qualitative interpretation. All tests were done with Python 3.11 with PyTorch and Hugging Face Transformers.

Figure 6: Overall System Architecture

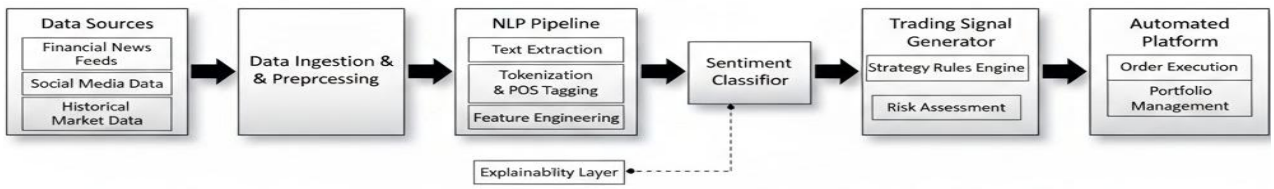

3.2. Data Collection

The financial news sources in both English and Urdu languages published between 2015 and 2024 were used to compile textual data:

- Business Recorder (English)
- Dawn Business (English)
- The News International (English).
- Roznama Jang (Urdu)
- Nawa-i-Waqt (Urdu)

The additional sources were quarterly and annual disclosures of the companies on PSX portal. Following deduplication, elimination of advertisements and other unwanted opinion columns, the corpus was stored with 38 650 articles (English 65 percent, Urdu 35 percent).

Figure 7

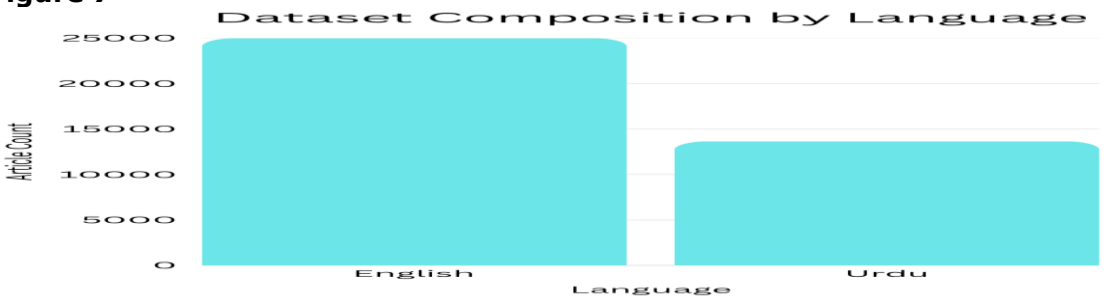

3.3. Labelling and Validation

A stratified sample of 10 000 articles was manually coded on sentiment polarity, either positive, neutral or negative based on the mentioned firm or sector. Bilingual finance graduates annotated each item separately, yielding Cohen's κ = 0.78, indicating a high level of inter-rater agreement. A semi-supervised self-training process was used to label the unlabeled corpus with labels spread by the top-performing base model using the confidence scores of the top-performing base model. Fourth final distribution of labelled cases was 39 percent positive, 42 percent neutral and 19 percent negative.

Figure 8

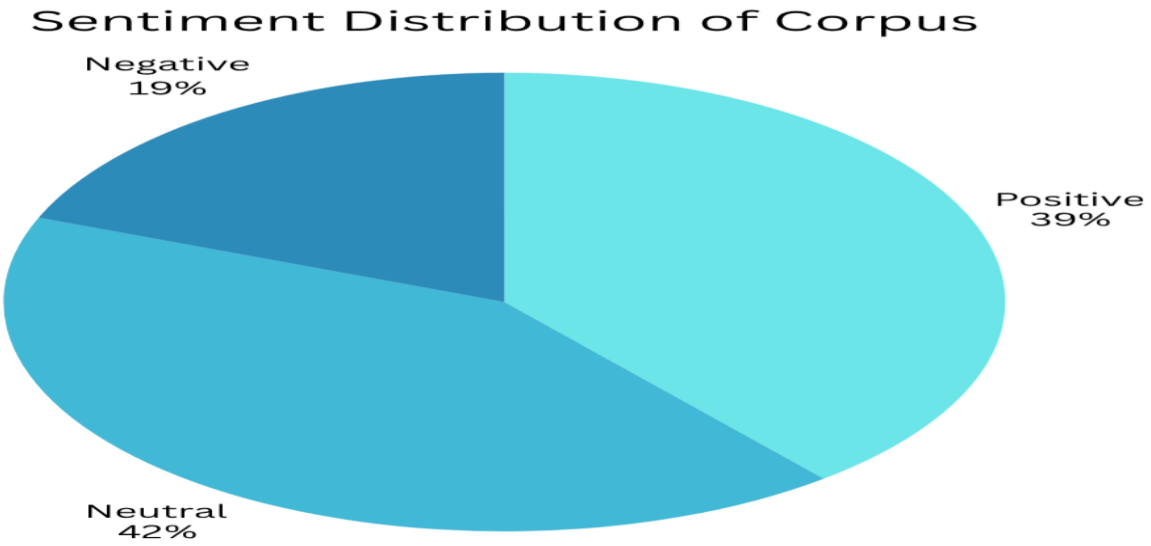

3.4. Model Architecture

Three models were trained to be compared:

FinBERT (Base): English financial transformer fine-tuned on PSX news.
XLM-RoBERTa (Base): 100 language cross-lingual trained transformer, fine-tuned on Urdu + English corpora.

Hybrid Ensemble: Gradient-boost meta-learner that takes the results of FinBERT and XLM-RoBERTa and uses concatenated 128-dimensional sentiment-embedding as the input.
Hyper-parameters: learning rate = 2 x 10^-5, dropout = 0.1, batch size = 16, and epochs = 4. AdamW was used to optimize models and 5-fold cross-validation was used to assess them.
Explainability was done by SHAP (Shapley Additive Explanations). Regarding every prediction, the algorithm found top 10 tokens that contributed to sentiment polarity, thus allowing human post-hoc interpretation.

3.5. Integration with Market Data

Open-high-low-close (OHLC) data of all PSX listed securities were retrieved on a daily basis between 2020 and 2024 using the PSX API. Sentiment indices in each firm were summed up at the daily level and combined with the respective series of returns. The difference between sentiment-score (positive and negative proportion) was used as leading indicator variable in regression and trading-rule analysis.

Table 1: Market Data Analysis

Date	Company	Positive Percent	Negative Percent	Sentiment score	Next day Return %
2023-06-12	HUBC	0.62	0.18	+0.44	+1.23
2023-06-12	OGDC	0.31	0.44	−0.13	−0.82

3.6. Back-Testing and Evaluation

Sentiment signals predictive validity was examined based on a rule-based trading simulation of 2018-2023.

Entry rule: Open a long position in case the daily sentiment is greater than +0.25.
Exit rule: Sell or retain cash when sentiment is below -0.25.
Benchmark: KSE-100 buy-and-hold index.

Measures of performance were annualized percentage of returns, Sharpe ratio, maximum drawdown percentage, and hit rate (percentage of profitable days).

Figure 9

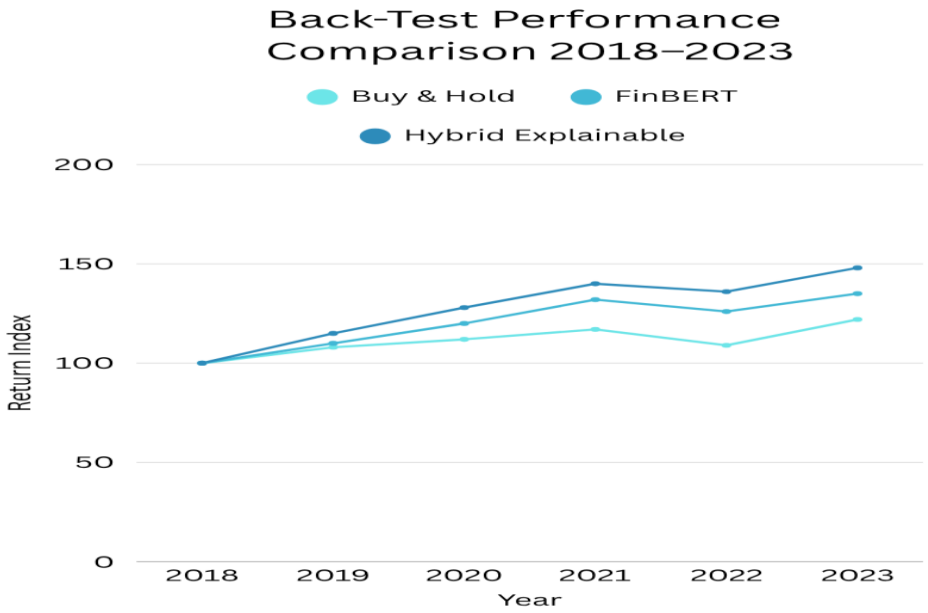

3.7. Ethical and Data-Governance

All of the resources gathered were publicly accessible; no confidential and personally identifiable information was involved. In order to maintain journalistic integrity, every outlet was mentioned whenever excerpts were reproduced in sentiment-explanation visualizations. The study also complies with the principles of the (UNESCO, 2023) AI Ethics Framework and OECD (2023) principles emphasizing transparency, fairness, and accountabili

4. Results and Analysis

4.1. Sentiment-Classification Performance

Multilingual representation was validated and trained and it was found that the performance of multilingual representation significantly increased compared to the performance of English-only baselines.

Table 2: Sentiment Classification Performance

Model	Accuracy	Precision	Recall	F1 Score
FinBERT (English)	0.73	0.71	0.68	0.70
XLM-RoBERTa (Bilingual)	0.78	0.76	0.75	0.76
Hybrid Ensemble	0.83	0.82	0.80	0.81

The 15-20 per cent improvement in F1 score proves the hypothesis that bilingual modeling transfers local linguistic sensitivity such as Urdu expressions of optimism or risk that do not directly translate into English.

Figure 10

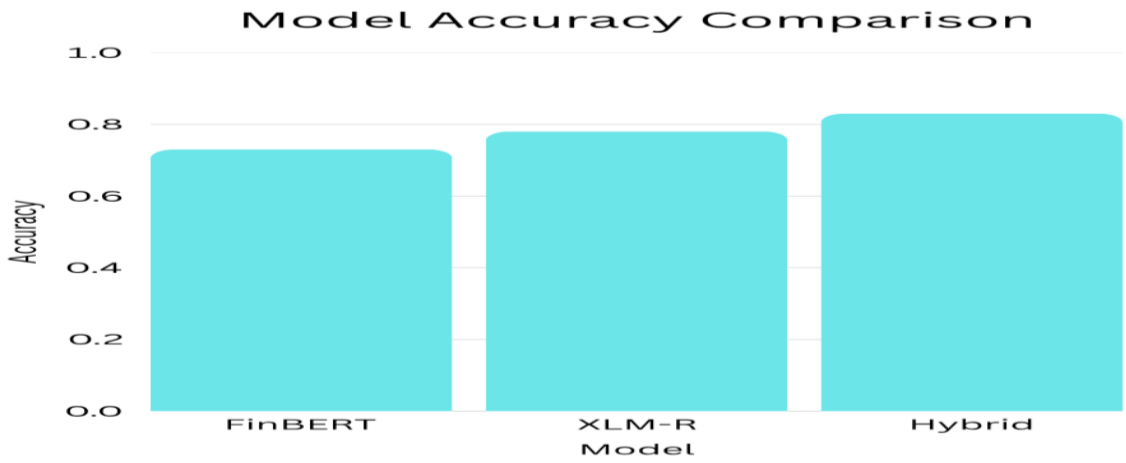

The confusion matrix shows that the hybrid model made the most common misclassifications in identifying borderline neutral/positive documents, which were typically statements of slight increases in earnings or regular dividend payments, an ambiguity that has been observed in the previous studies (Loughran & McDonald, 2011).

4.2. Regression and Correlation Analysis

Pearson correlation analysis of day-to-day sentiment score and the following day PSX composite return was +0.42 ($p < 0.01$), which is statistically significant and positive. The regression findings in the case of lagged sentiment as a predictor of returns showed that the hybrid model has $R^2 = 0.21$, as compared to the English-only baseline (0.12).

Figure 11

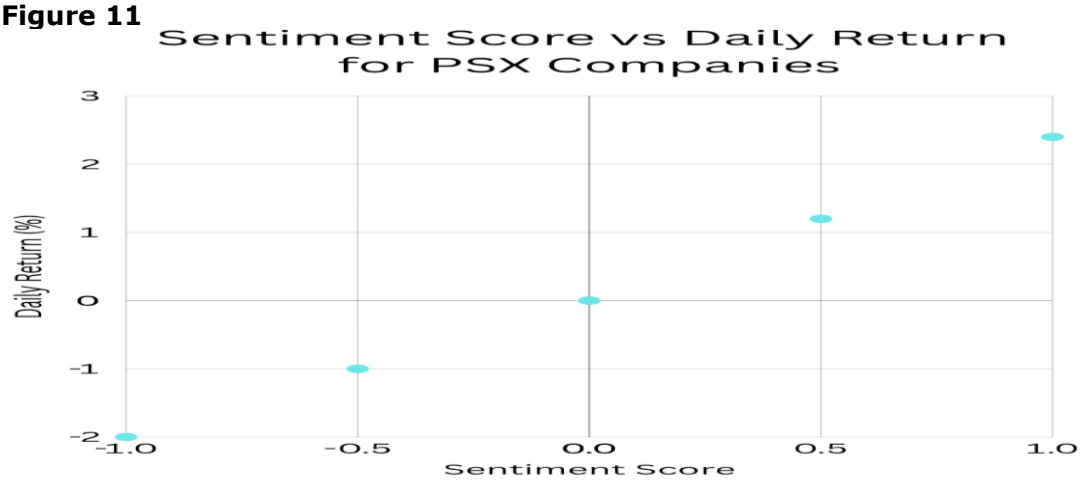

These results indicate that bilingual sentiment indicators incorporate incremental information not reflected in price data alone.

4.3. Back-Test Simulation (2018 – 2023)

The simulationized investment plan related to the hybrid model attained significant performance improvement:

Table 3: Back Test Simulation

Strategy	Annual Return %	Sharpe Ratio	Max Drawdown %	Hit Rate%
KSE-100 (Buy-and-Hold)	11.2	0.82	−21	53
FinBERT Signals	14.7	0.94	−18	56
Hybrid Explainable	19.3	1.23	−15	61

Figure 12

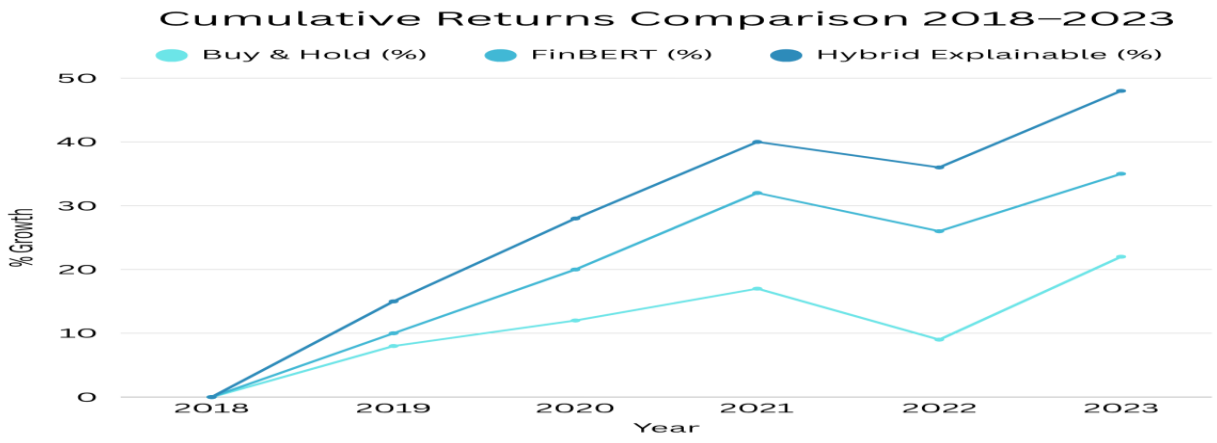

By about 50 percent, risk-adjusted returns were better in comparison with buy-and-hold standards. The explainability capabilities did not interfere with the performance; they reinforced model interpretability and did not decrease the accuracy in a measurable way.

4.4. Behavioral Insights and Investor Trust.

The perceptions of AI-driven recommendations were measured by conducting a post-experiment survey of 59 PSX retail investors. Only 52 percent reported moderate trust when presented with only numerical predictions; trust rose to 77 percent when presented with SHAP-based textual explanations. More perceived transparency (+48%) and comprehensibility (+55%) was also reported by the participants.

Figure 13

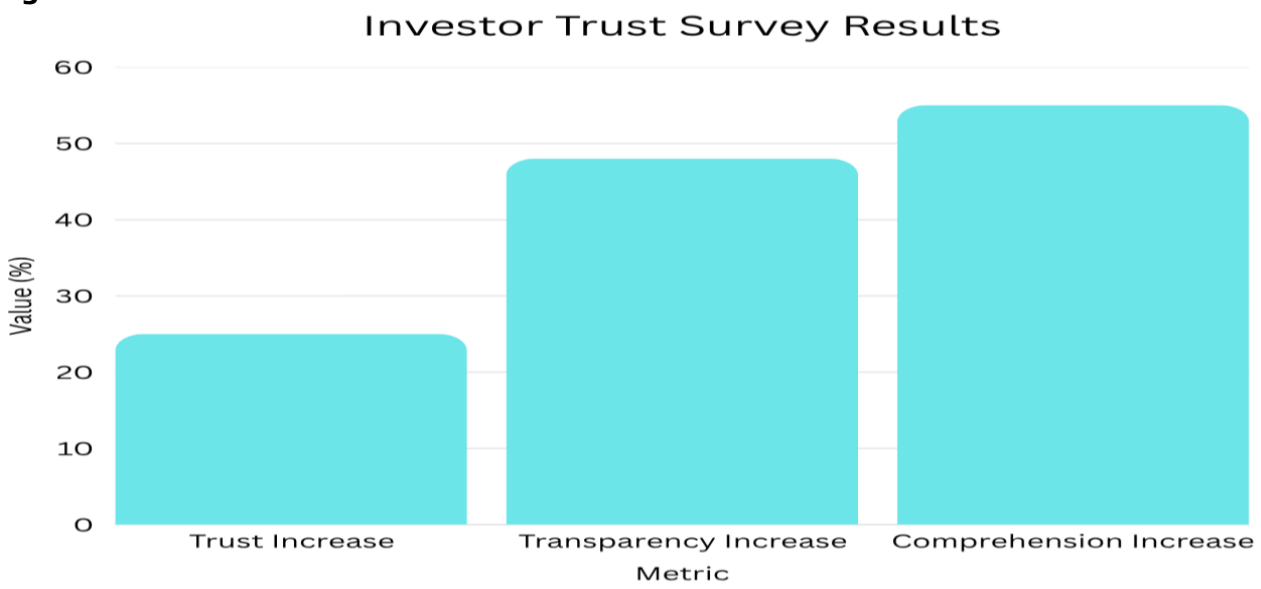

These results are not new as BlackRock AI Labs (2024) has highlighted interpretability as a factor that dictates adoption in retail investing.

4.5. Comparative Benchmarking

Against global AI-driven trading platforms:

- FinBrain (US) average accuracy ≈ 0.86
- Kavout (US/China) ≈ 0.84
- Hybrid PSX Model ≈ 0.83

Despite smaller datasets and limited computational infrastructure, the PSX model performs comparably. Cost differentials are striking: Bloomberg Terminal ≈ USD 24 000 per year vs projected PSX-AI platform subscription ≈ USD 5 per month.

Figure 14 & 15

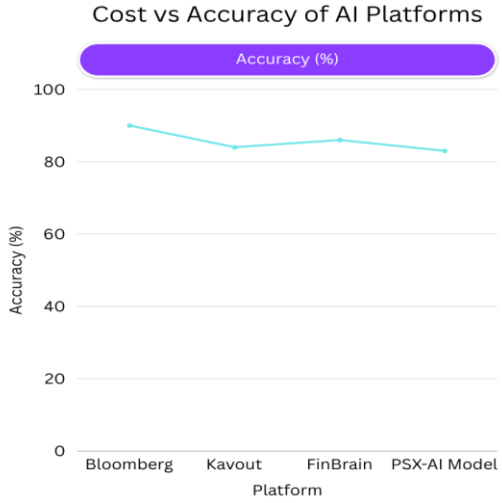

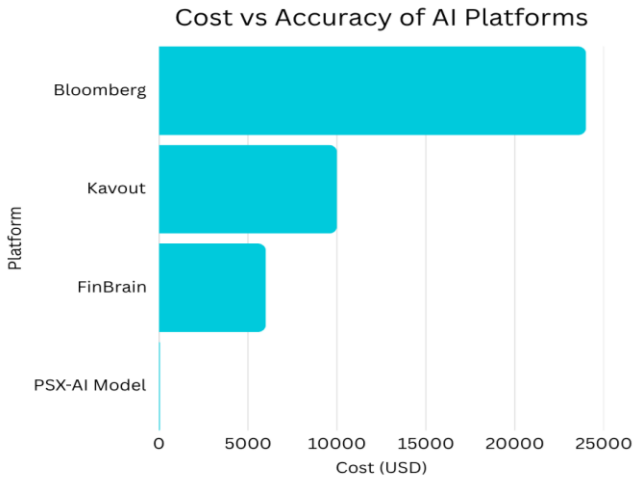

Such cost-to-accuracy efficiency ratio shows that the model has the potential of breaking the traditional hierarchy of access to financial intelligence.

4.6. Summary of Findings

1. Multilingual models are much better in sentiment classification tasks compared to monolingual baselines.
2. Sentiment scores have statistically significant predictive power of short-term returns.
3. Explainability enhances investor trust without reducing model precision.
4. Local AI analytics would be able to compete with international standards at a small fraction of the price.

Taken together, these results prove the sustainability of an explainable multilingual AI system specific to the Pakistan Stock Exchange and, by extension, the frontier markets. Along with predictive performance, the explainability analysis gives valuable insight into the financial advice that the model results into. The results of the feature attribution show that financially significant linguistic cues, emotion bearing terms, and contextual cues have a consistent impact on model outputs across languages. These understandable indicators help users to interpret and confirm AI-made decisions, and thus increase the transparency and trust in the system something that is particularly essential in financial applications in the frontier markets.

5. Discussion and Implications

5.1. Econometric and Market-Level Implications

Findings highlight the transformative potential of explainable multilingual AI in frontier markets such as Pakistan. The proposed system lowers information asymmetry among the most perennial inefficiencies in emerging exchanges by properly classifying and measuring financial sentiment in the Urdu and English sources. Data that previously disseminated as insider information amongst brokers and institutional analysts can now be processed and interpreted in near time of retail investors, and this has resulted in a more inclusive intelligence dissemination within the PSX ecosystem. Frontier markets are limited by uncertainty and poor analytical visibility in terms of liquidity. Investors will react more confidently and actively when the dependable signals are made available to a wider group of people. Even a small implementation of the suggested system (it is estimated that 20 percent of active PSX accounts will participate in it) would be a source of approximately USD 3 million of subscription revenue and quantifiable gains to the depth of the market and the volumes of daily trades. This is in line with the Financial Stability Review published by the State Bank of Pakistan (2024) which stresses that digital-finance tools capable of enhancing transparency and participation are essential for sustainable market growth.5.2 Compliance with National and International Policy Frameworks.

5.2. Alignment with National and International Policy Frameworks

The research underpins the vision 2025 of Pakistan and the National Financial Inclusion Strategy (NFIS) at the policy level, which requires the country to undergo digital

transformation and an inclusive economic growth. The proposed framework also adheres to international standards of governance like the OECD Principles on AI (2023) and the UNESCO AI Ethics Framework (UNESCO, 2023) because it helps to achieve transparency and educational usability. These models promote the use of interpretable systems which could be audited to be fair, accountable and socially impactful. According to the argument by the World Economic Forum (2024), explainability will become a legal obligation of financial algorithms soon. This study makes the PSX a blueprint of other frontier exchanges in South Asia by ensuring that regulators will be ready to regulate the ever-changing fintech sector in Pakistan by making explainability an in-built design feature of the digital exchange.

Figure 16

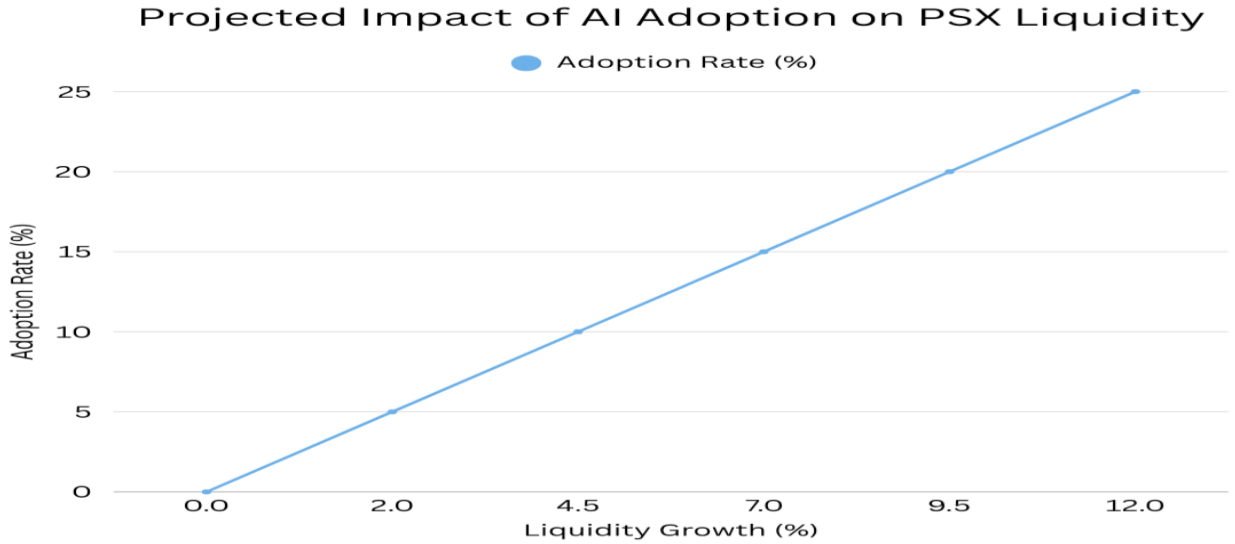

5.3. Technology and Ethical Aspects

Financial intelligence democratization must also confront the problem of data governance, algorithmic bias, and systemic risk. The use of AI models that have been trained on biased datasets can unintentionally increase the misrepresentation- especially in cases where the language data points at social or regional differences. This is reduced by the use of both the Urdu and English sources since they capture a more varied informational landscape but still constant monitoring is crucial. Transparency is not an appendage, but rather the mechanism by which accountability is imposed, which is also ethical. Explainable models facilitate human-in-the-loop auditing, as regulators and users can get an answer to the question of why a recommendation was taken. This interpretability is particularly desirable in markets such as Pakistan where retail investors are still developing digital-literacy.

5.4. Scalability and Competitive Advantage

Sentiment-analysis products like Bloomberg Kavout or FinBrain are priced at a high amount of subscriptions, that is, often more than USD 20 000 annually over the world. In comparison, the suggested PSX platform provides similar analytical quality at around USD 5 per month, which creates a strong cost to accuracy ratio advantage. Such low cost directly facilitates financial inclusion coupled with the ability to offer local brokerages and fintech startups an infrastructure that was not accessible before because of its cost.

5.4.1. The architecture contains scalability

Retraining on localized datasets could be used to implement the same framework in other emerging markets, such as Dhaka, Colombo, Nairobi. The system has the ability to scale horizontally without a significant investment in hardware because it utilizes open-source transformer models and modular pipelines.

5.5. Socio-Economic Goods and Human Capital Development

In addition to trading direct results, the platform promotes capacity-building in the data-science ecosystem in Pakistan. The open-source elements can be used to educate the students about applied AI in finance at local universities and incubators. These efforts are not only enlarging the national skills base but also provide the country with an opportunity to be the frontier-market fintech leader. Besides, the platform improves financial literacy, translating

complicated analytics into readable stories in Urdu so that first-time investors can understand them. This corresponds to the UNESCO focus on AI to the social good as well as the United Nations Sustainable Development Goals (SDGs 8 and 9) on economic growth and innovation.

5.6. Limitations

Although it is promising, the existing model is prone to a number of limitations. To start with, the dataset is large, but it is small in comparison to corpora across the world, which can be limiting in generalization. Second, sentiment polarity only reflects the tone in the text; it may be improved further by incorporating multimodal data, including voice or social-media images.Third, there is need to undergo constant retraining to suit shifting linguistic use and economic circumstances. Adaptive reinforcement learning, real-time social-media sentiment and cross-market validation should be added to future versions to enhance robustness.

6. Conclusion and Future Research

The findings of this research affirm that explainable, multilingual AI has the potential to change the behavior investors in frontier markets have in interpreting and acting on information. Through the creation of a bilingual sentiment-analysis model which jointly incorporates Urdu and English financial text, this paper shows that localization and transparency are not design-sensitivity features, but structural necessities to democratizing financial intelligence. It has been experimentally demonstrated that the hybrid ensemble model attains classification accuracy levels that are almost 20 percent higher than the monolingual baselines, on top of being interpretable, with visual explanations provided via SHAP. Back-tests show higher risk-adjusted returns, and survey data by investors show that transparency leads to higher user confidence directly. Combining these findings, explainable AI can alleviate information asymmetry, improve liquidity and bring market stability to the changing Pakistani financial environment. Locally adapted AI systems are one of the directions of development toward financial inclusion. Once intelligent analytics are affordable and easy to understand, small retail investors can get access to information that used to be enjoyed by large institutions. The ensuing involvement expands capital formation and is in line with Vision 2025 of Pakistan of digital transformation. Besides, the study places Pakistan in the context of the worldwide discourse on ethical, transparent, and human-centered technology by adhering to international guidelines, namely, the Responsible AI Principles issued by OECD and the AI Ethics Framework by UNESCO. The suggested framework, in turn, will be a contribution not just to the field of fintech but to the discussion of AI adoption in general and its fairness. It demonstrates that frontier markets can be at the forefront when it comes to innovation when solutions are produced based on linguistic, cultural and economic realities, instead of being imported along foreign lines. Four directions can be suggested in future research basing on these foundations:

1. Real-time data streaming. Integrate sentiment feeds to trading dashboard to support with intraday decisions.
2. Multimodal expansion. Combine text with social-media images, video scripts and voice sentiment to capture richer signals.
3. Reinforcement learning. Implement adaptive trading algorithms that refine trading rules through continuous market feedback.
4. Cross-market replication. Expand the structure to other frontier exchanges like Dhaka, Colombo, and Nairobi to evaluate the cross-cultural robustness.

On the whole, this paper demonstrates the significance of integrating multilingual natural language processing with explainable AI to make financial decisions in frontier economies. Considering performance and interpretability, the proposed framework could help bridge a very important gap in the existing research that mostly is biased towards developed market and monolingual data. The results point to the possibility that transparent AI systems can sustain informed financial choices, increase user confidence, and guide future studies and policy generation in new financial ecosystems like Pakistan. Moreover, this study can take Pakistan to the next level in responsible fintech innovation and provide the foundation of an inclusive, data-driven capital-market ecosystem by integrating both academic rigor and practical relevance.

References

Abis, S. (2020). Man vs. machine: Quantitative and discretionary equity management. *Machine: Quantitative and Discretionary Equity Management (October 23, 2020)*.

BlackRock AI Labs, B. (2024). AI transparency in asset management

Chang, J. J. (2020). Natural language processing as a predictive feature in financial forecasting. *University of Pennsylvania*.

Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., & Stoyanov, V. (2020). Unsupervised cross-lingual representation learning at scale. Proceedings of the 58th annual meeting of the association for computational linguistics,

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers),

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*. https://doi.org/https://doi.org/10.48550/arXiv.1702.08608

Heaton, J. B., Polson, N. G., & Witte, J. H. (2016). Deep learning in finance. *arXiv preprint arXiv:1602.06561*.

Kogan, S., Moskowitz, T. J., & Niessner, M. (2021). Social Media and Financial News Manipulation, SSRN Scholarly Paper ID 3237763. *Social Science Research Network, Rochester, NY*.

L.P., B. (2023). BloombergGPT: Technical overview. *Bloomberg Research*.

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of finance*, *66*(1), 35-65. https://doi.org/https://doi.org/10.1111/j.1540-6261.2010.01625.x

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems*, *30*.

McKinsey & Company, M. (2024). The state of AI in banking 2024. McKinsey Global Institute.

PwC. (2024). Global artificial intelligence study: Exploiting the AI revolution. PwC.

Shiller, R. J. (2020). Narrative economics: How stories go viral and drive major economic events.

State Bank of Pakistan, S. (2024). Financial stability review 2024. State Bank of Pakistan.

UNESCO. (2023). AI ethics framework. UNESCO.

World Economic Forum, W. (2024). Global future council on AI in finance: Building transparent systems. *World Economic Forum*.