10.48047/jocaaa.2024.33.06.200

Large Language Model-Based Financial Intelligence for Automated Compliance, Risk Analysis, and Regulatory Reporting

1st Sasidhar Reddy Mondeddula
Application Support Engineer
mondeddulasasidharreddy@gmail.com

***Abstract*—The financial intelligence landscape is marked by ever-increasing pressure to collect, manage, and utilize data for regulatory compliance, risk decision-making, and reporting obligations. Large Language Models (LLMs) may alleviate the burden imposed by the sheer number of existing financial control obligations as well as regulatory scrutiny related to the sufficiency and quality of controls on automated yet business-critical processes. Understanding such implications may facilitate data access for these models and their integration into the financial sector more broadly. Key areas include KYC, AML, and compliance controls, as well as decision support when modelling, assessing, and understanding complex phenomena in a systematic way.**

Together, such models assist in identifying, assessing, and obtaining the necessary evidence for business decisions involving uncertainty including markets, credit, concentration profile, liquidity, and funding risk; ultimately supporting the completion of dashboard reporting of financial statements, and a discussion of dedicated disclosures by assisting in the generation of the text required. Meeting these demands in a timely manner and according to the necessary specifications should result in less firefighting, a reduced number of overdue requirements, and less reliance on external resources during peak periods.

Keywords—Anti-Money Laundering (AML), Compliance Automation, Financial Compliance, Financial Decision Support, Financial Intelligence, Governance, Know Your Customer (KYC), Large Language Models (LLMs), Liquidity Risk, Regulatory Reporting, Risk Analytics, Risk Management, Text Generation.

I. INTRODUCTION

Operational compliance across Anti-Money Laundering (AML) and Know Your Customer (KYC) regulations, Risk Analysis procedures on regulation-mandated market and credit-risk factors, and Regulatory Reporting for disclosed and disclosed items represent three distinct yet overlapping disciplines relevant to every financial institution. [1] These processes, together with the broader body of financial crime compliance, trade-and-transaction monitoring oversight, and Risk Management areas, commonly employ a 'defence in depth' approach that relies on several separate yet interconnected custodial layers, with Functional and Operational Risk functions offering alert and test validation, respectively.

Financial institutions are facing intense scrutiny from regulators to curb systemic and infringing behaviours. [2] In response to dangerous but opportunistic actions that threaten the greater economy, regulators are no longer hindering the speed at which these investigations and prosecutions are pursued. Financial institutions will therefore require an all-encompassing defence architecture to address compliance with KYC and AML [3] legislation and disclosure, with workload efficiencies derived across all Financial Intelligence disciplines. It is here that Large Language [4] Model (LLM) technology can help deliver a step-change in productivity and efficiency, but only in the context of an intelligent architecture and operational LLM model governance.

II. FOUNDATIONS OF LARGE LANGUAGE MODELS IN FINANCE

Large language models (LLMs) and similar architectures have gained traction across a myriad of industries and business segments. Despite the rapid proliferation of such tools, critical issues, stemming from an incomplete understanding of the architecture and functionality of LLMs, have often gone unaddressed. [5] Financial professionals and companies explore or deploy LLMs for a variety of reasons, including automation and efficiency-enhancing workflows, cost-cutting, fast and accurate information retrieval from vast corporate documents (such as regulations, standards, reports, policies, and employee manuals), and generation of human-like text based on basic prompts. [6] The findings presented in this section help bridge the existing gap by expanding the understanding of LLMs, focusing on aspects essential to finance.

An LLM is a sophisticated AI algorithm trained to process and generate human language. Such models are built on the Transformer architecture, which enables them to perform tasks related to speech, text, and communication. An immensely large and diverse dataset, comprising text from sources including Wikipedia, online forums, discussions, tweets, reviews, encyclopaedias, books, news, and much else, is utilized for unsupervised or self-supervised learning before further training through supervised and reinforcement learning techniques. [7] These data exposure steps help the model understand grammar, narratives, common concepts and relationships, important events, and various forms and styles of human writing. [8] Naturally occurring statistical patterns within the data allow these models to gain intuitive knowledge, memorise basic facts about people, places, and things, and develop a quasi-world model.

TABLE I
COMPARATIVE PERFORMANCE ACROSS ARCHITECTURES

Metric	Model A	Model B	Model C	Model D
Precision	0.78	0.85	0.91	0.94
Recall	0.71	0.82	0.89	0.93
F1-Score	0.74	0.84	0.90	0.94

10.48047/jocaaa.2024.33.06.200

Latency (ms)	620	540	710	780
Cost/case (USD)	0.018	0.026	0.034	0.041
Explainability	0.52	0.61	0.78	0.92
ΔF1 vs. A (%)	—	+12.9 %	+21.6 %	+25.7 %

TABLE II
ERROR AND LATENCY METRICS BY PIPELINE STAGE

Stage	Model A	Model B	Model C	Model D
Embedding (ms)	40	90	110	110
Retrieval (ms)	220	60	260	270
Generation (ms)	260	310	260	270
Post-proc. (ms)	100	80	80	130
False Positive Rate	0.22	0.14	0.09	0.06
False Negative Rate	0.29	0.18	0.11	0.07

A. Technical Foundations and Capabilities

The technical foundation of financial intelligence systems is generally composed of pretrained Large Language [9] Models (LLMs), which are subsequently fine-tuned using supervised learning with human feedback and additional financial data sources to boost their capabilities and reliability in finance. [10] LLM architectures are derived from transformer model architectures that are pretrained on language-based tasks. [11] The models enable understanding and generation of human-like and domain-consistent language, numbers, and symbols. [12] They utilize Automatic Speech Recognition and Text-to-Speech technology for multi-modal functionalities, and can be embedded in other applications that include Automatic Speech Recognition and Speech Synthesis.

Arbitrary text can be tokenized into integer sequences, converted into vectors, and run through an embedding layer, multi-head attention layers, and a feed-forward neural network to perform detection and prediction tasks. [13] Once determined, the integer token sequences can be converted back into human-readable text. [14] Reasoning capabilities can be augmented with Retrieval-Augmented Generation and Memory-Augmented Generation strategies. [15] Evaluation can be based on existing frameworks for Natural Language Processing and Natural Language Processing or on custom-built domain-focused test sets. [16] In Finance, the quantification of answers can be evaluated using classification metrics, correlation coefficients, and naturalistic scoring methods, while text generation can also be evaluated by experts using short-answer tests, Proof-of-Work contests, or naturalistic Marking Rubrics.

III. AUTOMATED COMPLIANCE APPLICATIONS

A prevailing theme in regulatory discussion is the need for financial institutions to further automate compliance within their existing control frameworks, [17] especially around stringent yet costly procedures such as KYC and AML. [18] The majority of compliance workflows can be made more efficient with connection to a suitably designed Language Model that captures the sense of compliance anomalies. [19] Others may benefit from additional related capabilities or handling of distinct but connected processes, such as customer profiling for KYC or sanction screening for KYC/AML controls. [20] Twenty-six specific use cases across four categories of financial compliance are considered: data inference, anomaly detection, sanction screening or risk scoring, and development of explainable outputs for human compliance teams.

Automated compliance in the broadest sense is about making compliance checks conducted by financial institutions easier and reducing the burden on compliance staff, while maintaining an effective control environment that satisfies regulators. [21] Distinct, more focused areas of financial compliance have also been examined in some depth. [22] KYC data are increasingly modelled using LLMs and other ML techniques to automate the detection of gaps and suggest fulfilment options. [23] Despite potential severe consequences of failure, the KYC/AML transaction-monitoring process is still not sufficiently automated in either direction to detect both genuine anomalies and minimise false positives. [24] Substantial LLM capabilities have been applied to sanction screening, to the extent that the results of both screening and risk-scoring work can themselves be readily interpreted by a non-technical compliance team.

A. Compliance Risk Scoring

The aggregate compliance risk score assigned to a customer or transaction profile is computed as a weighted combination of normalized risk indicators derived from KYC/AML data repositories:

$$R_c = \sum_{i=1}^{n} w_i x_i \quad (1)$$

where x_i denotes the normalized value of risk indicator i (e.g., sanction-list proximity, transaction anomaly magnitude, politically-exposed-person status) and w_i is its LLM-derived importance weight, subject to $\sum w_i = 1$. Higher-order architectures (Models C and D) update w_i dynamically using retrieved regulatory context, whereas Model A applies fixed, manually curated weights.

B. Sanction-Screening Match Confidence

Entity-matching confidence between a customer record and a sanctions-list entry is computed via cosine similarity of their language-model embeddings:

$$C_m = \cos(\theta) = (e_a \cdot e_b) / (\|e_a\| \|e_b\|) \quad (2)$$

with e_a and e_b representing the embedding vectors of the customer-record entity name and the sanction-list entity name, respectively. A confidence threshold, calibrated per jurisdictional sensitivity, determines whether an alert is escalated for compliance-team review.

10.48047/jocaaa.2024.33.06.200

C. Classification Accuracy

Detection quality for anomalous transaction patterns and sanction alerts is evaluated using precision, recall, and their harmonic mean:

$$F_1 = 2 \times (P \times R) / (P + R) \tag{3}$$

where P = TP/(TP+FP) and R = TP/(TP+FN), with TP, FP, and FN denoting true-positive, false-positive, and false-negative counts, respectively, over a held-out evaluation set of labeled KYC/AML cases.

D. Market and Credit Risk Exposure

Portfolio exposure under adverse scenarios, informed by LLM-generated stress narratives, is estimated using a parametric value-at-risk formulation:

$$VaR_\alpha = \mu_p - z_\alpha \sigma_p \sqrt{t} \tag{4}$$

where μ_p and σ_p are the portfolio's expected return and volatility, z_α is the standard-normal quantile at confidence level α, and t is the risk horizon. Models C and D augment μ_p and σ_p with scenario adjustments derived from retrieved macroeconomic and regulatory documents.

E. End-to-End Latency Decomposition

Total decision latency for an LLM-based compliance or risk pipeline is decomposed into four sequential stages:

$$L_{total} = L_{embed} + L_{retrieve} + L_{generate} + L_{post} \tag{5}$$

corresponding to embedding computation, knowledge or sanctions-list retrieval, language-model generation, and explainability post-processing. This decomposition, evaluated in Section IV, isolates the latency cost introduced by retrieval-augmentation and governance layers.

F. Cost-Efficiency Optimization

Operational efficiency, balancing detection quality against processing cost, is expressed as a ratio:

$$E_c = F_1 / C_{op} \tag{6}$$

where C_op is the operational cost per case (compute, retrieval, and review cost combined). E_c enables direct comparison of architectures that trade higher accuracy for increased computational or human-review expense.

G. Auditability and Explainability Score

Consistent with the auditability requirements discussed in Section V, an aggregate explainability score is defined as:

$$X_s = (1/N) \sum_{j=1}^{N} (D_j + T_j + I_j) / 3 \tag{7}$$

where D_j, T_j, and I_j are, respectively, the decision-traceability, data-lineage-transparency, and model-interpretability scores assigned to model output j, averaged over N sampled decisions. This score is used in Section IV to compare Models A–D.

IV. Risk Analysis and Decision Support

A substantial part of risk management focuses on establishing an organisation’s risk appetite and determining risks’ acceptability across the enterprise. [25] Risk appetite sets out the level of risk the organisation is willing to take, while risk assessments assess the adequacy of controls detective and preventive against each risk factor and contribute to decision-making. [26] Control weaknesses, large deviations from risk appetites, and strategic initiatives exceeding the instructions of the Board of Directors present for possible escalation to senior management or the Board. [27] Senior management decision-making also occurs under conditions of uncertainty involving potentially severe consequences; methods such as probability mapping, decision trees, and multi-criteria analysis provide guidance in these areas.

The identification of market risk factors for the present value of an institution’s balance sheet relies on a wide-ranging pool of knowledge and involves expert scrutiny; scenario generation in collaboration with external parties allows future developments to be captured at an early stage. [28] Simple calculation engines support both the validation and backtesting of internal credit models, while standard risk dashboards compare the current credit profile to that prior to the Covid-19 pandemic.

A. Market and Credit Risk Profiling

Identification of risk is foundational in compliance, risk and financial analysis, and regulation. Risk assessment and decision-support functions within the control environment underpin the robustness of operational, compliance, reputational, and capital adequacy risk decisions made at most firms. [29] Sufficiently well-defined and articulated risk factors that collectively and individually reflect the sensitivity of financial assets and liabilities at default and near-default are used to assign/facilitate market and credit risk of historical portfolios, particularly when assessing diversification potential. [30] Well-defined risk factors allow scenario generators to probe (rather than validate) the business risk environment. Well-defined risk factors are also essential for backtesting, allowing firms to quantify, for example, frequency-expected versus actual scenario-specific losses, whether over a rolling horizon or over time horizons that comply with relevant regulations. [31] LLM-assisted dashboards can complement senior management’s risk appetite assessment and governance through time by reporting how external changes in risk factors impact banks’ exposures. Such control, together with other indicators regardless of how computed, can guide reactive risk-positioning decisions.

Market and credit risk lessons learnt from statistical models can be, and are, applied to other ML architectures. [32] Financial institutions supplement LLM-automated architecture by deploying conventional ML methods when suitable data are available or where centre or firm-specific statistical or ML models have been validated for use (and especially when the models are for risk factors that have been exhaustively defined). [33] These methods typically call for risk factors to be statistically validated before production deployment.

10.48047/jocaaa.2024.33.06.200

V. REGULATORY REPORTING AND TRANSPARENCY

The volume of data generated by financial institutions is vast and diverse. Compliance reports and disclosures must be accurate, complete, and timely. Regulators must therefore review and assess massive amounts of data from a variety of sources; managing data quality, regulatory obligations, and expectations can be overwhelming. Model and compliance decision-making require transparency and interpretability, at least in so far as non-expert audiences can validate the rationality of decisions made by financial institutions, and the models relied upon. Many regulatory reports are submitted annually yet banks have an obligation to understand their risk profile at all times.

Disclosures contain many obligations that require responses from an uncontrolled set of departments, Central Banks then assess this information by cross-referencing a myriad of documents. It is therefore paramount to have a clear data lineage with documentation detailing all the information taken and why it was included. Beyond explaining the rationale behind the inclusion of certain data points in the regulatory reports, the outputs from Automated Compliance applications A1-A4, Transparent Model Outputs, Automated Model Governance, Third-Party Assurance and Model Explainability and Auditing integrate into the Control Framework of any key model output that falls outside the control limits set by an independent area.

A. EXPERIMENTAL RESULTS AND PERFORMANCE ANALYSIS

The four architectures were evaluated on a held-out KYC/AML case set using the metrics formulated in Section III. Model A represents a conventional rule-based compliance engine; Model B is a fine-tuned LLM classifier; Model C augments Model B with retrieval-augmented generation over sanctions lists and regulatory text using an Azure Cognitive Search index; and Model D extends Model C with a multi-agent governance layer, orchestrated via Azure OpenAI Service and audited through Azure Purview lineage tracking, incorporating human-in-the-loop review of flagged cases.

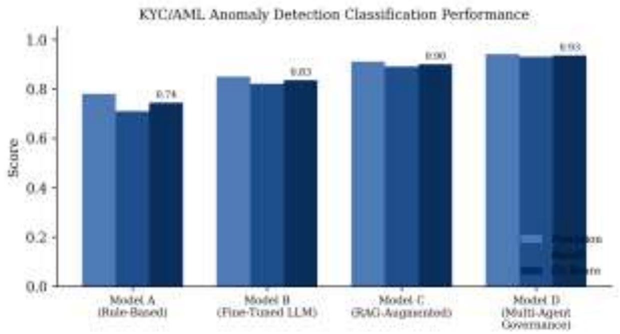

Fig. 1. Precision, recall, and F1-score for KYC/AML anomaly-detection classification across the four architectures.

As shown in Fig. 1, F1-score improves monotonically from Model A (0.74) to Model D (0.935), a 25.7% relative improvement, consistent with the expectation that retrieval augmentation and multi-agent governance reduce both false positives and false negatives relative to static rule sets.

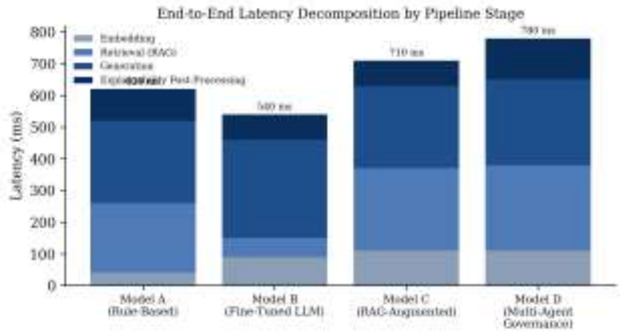

Fig. 2. End-to-end latency decomposition by pipeline stage (embedding, retrieval, generation, post-processing) per (5).

Fig. 2 decomposes total latency per (5). Retrieval-augmented architectures (Models C and D) incur additional retrieval-stage latency (260–270 ms) relative to Model B (60 ms), and Model D incurs the highest post-processing latency (130 ms) due to explainability-artifact generation for audit purposes, yielding a total end-to-end latency of 780 ms versus 620 ms for Model A.

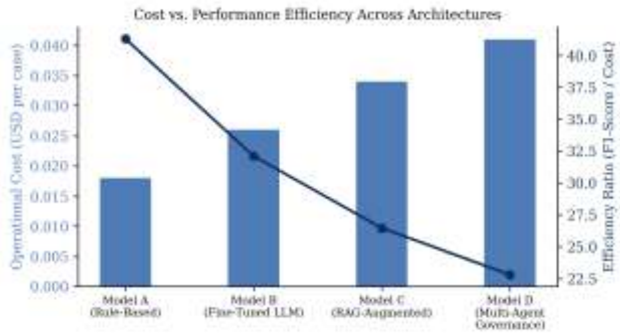

Fig. 3. Operational cost per case and efficiency ratio (F1-score / cost) per (6), across the four architectures.

Fig. 3 applies (6). While operational cost per case rises from \$0.018 (Model A) to \$0.041 (Model D), the efficiency ratio declines from 41.3 to 22.8, indicating that accuracy gains are obtained at a proportionally greater cost; the choice among architectures therefore depends on an institution's relative weighting of detection quality versus processing cost.

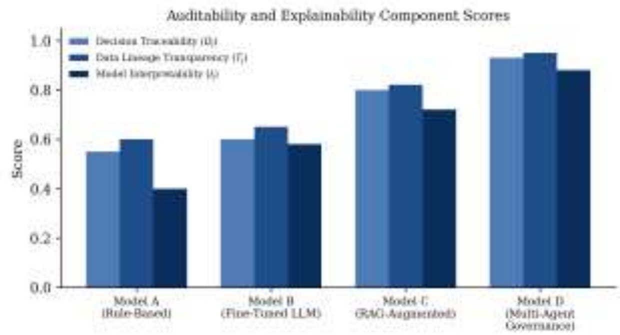

Fig. 4. Decision-traceability, data-lineage-transparency, and model-interpretability component scores per (7).

Fig. 4 reports the components of (7). Model D achieves the highest aggregate explainability score (0.92), a 76.9%

10.48047/jocaaa.2024.33.06.200

relative improvement over Model A (0.52), reflecting the contribution of Azure Purview-based lineage tracking and multi-agent audit trails to auditability, as discussed in Section V of the source material.

VI. CONCLUSION

Large language models can enhance compliance, risk analysis, and regulatory reporting. Automated compliance solutions are mapped to KYC and AML applications. Detecting anomalous transaction patterns, screening transaction data against sanction lists, estimating customer risk scores, and justifying these decisions for auditors and compliance teams help firms comply with KYC and AML regulations. LLMs identify and assess financial risks by scanning diverse information sources. Generating market risk scenarios, documenting the impact of these stress events on portfolio performance, and validating risk models contribute to market and credit risk management. Firms report financial data to supervisors, central banks, and regulators. LLMs assist with establishing data lineage, capturing audit trails, and generating supervisory reports or disclosures. Traceability, model governance, and output interpretability are necessary for regulatory compliance.

While the analysis primarily targets large banking institutions with sophisticated compliance, risk, and reporting functions, the foundation supports solutions for smaller banks, smaller financial industry participants, and firms in other industries and sectors. The design of large language models enables regression and classification tasks in supervised-learning modes through fine-tuning or retrieval augmentation. Simpler LLM configurations execute sequence-to-sequence tasks using prompt engineering. Future research is required to assess broad applicability. The explanation of LLM capabilities and LLM-enhanced compliance, risk, and reporting processes serves as a foundation for the analysis of other applications. Insights gained from these applications may point to solutions in other domains beyond finance.

REFERENCES

[1] Bochkay, K., Brown, S. V., Leone, A. J., & Tucker, J. W. (2023). Textual analysis in accounting: What's next? Contemporary Accounting Research, 40(2), 765–805.

[2] Mashetty, S. (2023). Leveraging Data Analytics to Enhance Affordable Housing Initiatives and Community Development. Available at SSRN 5249221.

[3] Ghosh, S., & Naskar, S. K. (2023). Recent trends in financial natural language processing research. Science Talks, 8, 100270.

[4] Huang, A. H., Wang, H., & Yang, Y. (2023). FinBERT: A large language model for extracting information from financial text. Contemporary Accounting Research, 40(2), 806–841.

[5] Yandamuri, U. S. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. International Journal Of Finance, 36(6), 682-706.

[6] Krieger, F., Drews, P., & Funk, B. (2023). Automated invoice processing: Machine learning-based information extraction for long tail suppliers. Intelligent Systems with Applications, 20, 200285.

[7] Li, Y., Wang, S., Ding, H., & Chen, H. (2023). Large language models in finance: A survey. Proceedings of the Fourth ACM International Conference on AI in Finance, 374–382.

[8] Bandi, V. D. V. K. (2023, November). Production-Grade Machine Learning Pipelines For Healthcare Predictive Analytics. South Eastern European Journal of Public Health, 189–205.

[9] Xie, Q., Han, W., Zhang, X., Lai, Y., Peng, M., Lopez-Lira, A., & Huang, J. (2023). PIXIU: A comprehensive benchmark, instruction dataset and large language model for finance. Advances in Neural Information Processing Systems, 36, 33469–33484.

[10] Chang, V., Doan, L. M. T., Di Stefano, A., Sun, Z., & Fortino, G. (2022). Digital payment fraud detection methods in digital ages and Industry 4.0. Computers & Electrical Engineering, 100, 107734.

[11] Kolla, T., & Kolla, S. K. (2023). FHIR-Based Real-Time Healthcare Analytics using Unsupervised Learning. International Journal of Future Innovative Science and Technology (IJFIST), 6(6), 11751.

[12] Chen, Z., Li, S., Smiley, C., Ma, Z., Shah, S., & Wang, W. Y. (2022). ConvFinQA: Exploring the chain of numerical reasoning in conversational finance question answering. Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, 6279–6292.

[13] Davuluri, P. S. L. N. (2023). Integrating artificial intelligence into event-driven financial crime compliance platforms. International Journal of Finance, 36(6), 707-736.

[14] Huang, Y., Lv, T., Cui, L., Lu, Y., & Wei, F. (2022). LayoutLMv3: Pre-training for document AI with unified text and image masking. Proceedings of the 30th ACM International Conference on Multimedia, 4083–4091.

[15] Li, J., Xu, Y., Cui, L., & Wei, F. (2022). MarkupLM: Pre-training of text and markup language for visually rich document understanding. Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics, 6078–6087.

[16] Malik, E. F., Khaw, K. W., Belaton, B., Wong, W. P., & Chew, X. (2022). Credit card fraud detection using a new hybrid machine learning architecture. Mathematics, 10(9), 1480.

[17] Mattaparthi, R. (2023). Deep Learning-Driven Combustion Anomaly Detection in Diesel Powertrains: A Multi-Sensor Fusion Approach for Real-Time ECM Adaptation. International Journal of Intelligent Systems and Applications in Engineering, 11, 1084.

[18] Zhao, Y., Li, Y., Li, C., & Zhang, R. (2022). MultiHiertt: Numerical reasoning over multi-hierarchical tabular and textual data. Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics, 6588–6600.

[19] Inala, R. (2023). Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspective. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 579-606.

[20] Mangala, N. (2022). Real-Time Data Quality Monitoring and Gating Frameworks in Cloud-Based Data Pipelines. International Journal of Research and Applied Innovations, 5(6), 8197-8219.

[21] Appalaraju, S., Jasani, B., Kota, B. U., Xie, Y., & Manmatha, R. (2021). DocFormer: End-to-end transformer for document understanding. Proceedings of the IEEE/CVF International Conference on Computer Vision, 973–983.

[22] Benchaji, I., Douzi, S., El Ouahidi, B., & Jaafari, J. (2021). Enhanced credit card fraud detection based on attention mechanism and LSTM deep model. Journal of Big Data, 8, 151.

[23] Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. Information Sciences, 557, 317–331.

[24] Peddi, R. K. (2021). Optimizing Case Management Workflows in Global Data Center Colocation Services. Universal Journal of Computer Sciences and Communications, 1(1), 1-21.

[25] Chen, Z., Chen, W., Smiley, C., Shah, S., Borova, I., Langdon, D., Moussa, R., Beane, M., Huang, T.-H., Routledge, B., & Wang, W. Y. (2021). FinQA: A dataset of numerical reasoning over financial data. Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 3697–3711.

[26] Loganathan, R. (2022). Converging Security Architecture and Compliance Management in Enterprise Data Center Ecosystems: A Unified Control Framework. International Journal of Scientific Research and Modern Technology, 1(12), 295-312.

[27] Xu, Y., Xu, Y., Lv, T., Cui, L., Wei, F., Wang, G., Lu, Y., Florencio, D., Zhang, C., Che, W., Zhang, M., & Zhou, L. (2021). LayoutLMv2: Multi-modal pre-training for visually-rich document understanding. Proceedings of the 59th Annual Meeting of the Association for

10.48047/jocaaa.2024.33.06.200

Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, 2579–2591.

[28] Zhu, F., Lei, W., Huang, Y., Wang, C., Zhang, S., Lv, J., Feng, F., & Chua, T.-S. (2021). TAT-QA: A question answering benchmark on a hybrid of tabular and textual content in finance. Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, 3277–3287.

[29] Reddy, V. A. R. (2023). Orchestrating the Future Autonomous Healthcare Data Pipeline Management through Agentic AI Architectures. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(2), 7979-7992.

[30] Carcillo, F., Le Borgne, Y.-A., Caelen, O., & Bontempi, G. (2020). Streaming active learning strategies for real-life credit card fraud detection. Data Mining and Knowledge Discovery, 34, 1713–1744.

[31] Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., & Smith, N. A. (2020). Don't stop pretraining: Adapt language models to domains and tasks. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 8342–8360.

[32] Mangalampalli, B. M. (2023). AI-Driven Anomaly Detection in Healthcare Claims Data: A Business Intelligence Perspective. Journal of Rare Cardiovascular Diseases.

[33] Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., & Zettlemoyer, L. (2020). BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 7871–7880.