International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
IJISAE
ISSN:2147-6799 www.ijisae.org Original Research Paper

Natural Language Processing for Intelligent Automation of Financial Documents and Banking Operations

[1]Sasidhar Reddy Mondeddula

Submitted: 01/07/2023 **Revised**: 18/08/2023 **Accepted**: 25/08/2023

Abstract—Applications of Natural Language Processing (NLP) in Intelligent Automation of Financial Documents and Operational Activities of Banks Natural Language Processing (NLP) techniques ranging from parsing to translation, when leveraged properly, can introduce the next level of automation to the process flows in Banking and Financial Services (BFS). Such intelligent automation reduces operating costs for banks and other financial services while enhancing the customer experience by providing 24/7 basic services. However, banking operations demand that such automation is precise and failure rates near zero as the consequences of errors and fraud detection are also high. As a result, while NLP capabilities can bring about intelligent automation, the risk of error makes it essential to formulate the right automation strategy that may require human intervention in the loop. Natural Language Processing enables the automation of several operational activities like customer onboarding, KYC (Know Your Customer) workflows, compliance checks, and access provisioning in a bank. These areas demand the handling of several types of structured and unstructured documents associated with the customer. These pose various challenges for automation based on the analytics capabilities, data quality, and risk appetite of the organizations. Customers are also migrating towards online services. However, BFS fails to recognize a high degree of online self-servicing like other industries, like IT CFCs. Growth opportunities will arise from lightweight, 24/7, automated systems based on human-in-the-loop, machine-learning models.

Keywords—*Banking Automation, Compliance Automation, Customer Onboarding, Document Intelligence, Financial Document Processing, Human-in-the-Loop, Intelligent Automation, Know Your Customer (KYC), Machine Learning, Natural Language Processing (NLP), Operational Analytics, Risk Management, Unstructured Data Processing.*

I. INTRODUCTION

Natural Language Processing (NLP) methods play an important role in the automation of processes related to financial documents, which are essential in a wide range of banking operations. As a result, the application of NLP approaches within the financial domain can significantly contribute toward process automation, risk assessment, customer onboarding, Know Your Customer (KYC) workflows, legal compliance, and risk assessments. [1] A growing number of such applications could mean significant efficiency gains for banks, nevertheless, caution needs to be exercised: these core capabilities can only assist, and act as augmentation mechanisms. The associated risk of having machines propagate errors in more complex tasks, which may fall outside of their trained capabilities, requires a careful examination of the risks involved and the remaining parts of the tasks that usually require human impact.

In particular, human-in-the-loop mechanisms, where a human partakes part in the final decision making, not CPU speed, painted a more trustworthy picture of an ever-growing development of known yet prospective risks, specially replicated within other Machine Learning (ML) disciplines. Since the financial service ecosystem is particularly regulatory focused, the same risks are also likely to replicate in the implementation and deployment of NLP capabilities in the banking sector and hence, have a direct baring on the previous point, sensitive to Policy and Governance-related risk.

Application Support Engineer
mondeddulasasidharreddy@gmail.com

A. Overview of the Study

Intelligent automation of unstructured processes—tasks primarily performed on unstructured documents—is an integral part of digital transformation in the banking and financial services industry (BFSI). [2] Natural language processing (NLP) can play a crucial role in this area by enabling computers to read, understand, and respond to human language in a valuable way. In addressing the industry's need for intelligent automation of financial documents and banking operations, this paper explores the specific types of NLP techniques that can be used, how those capabilities link with applications in banking processes, and how different characteristics and qualities impact both performance and automation potential. [3] The overarching research question focuses on developing a clearer understanding of the relationship between these aspects: "Which NLP techniques and capabilities can be applied to automate unstructured processes in the banking industry and how do their characteristics affect the potential for automation?"

There is particular interest in customer onboarding, due-diligence checks, and KYC workflows, where natural language technologies also have a major impact on reducing operational complexities and associated risks. [4] A more detailed relationship between specific NLP capabilities and five key banking processes customer onboarding, KYC workflow execution, KYC checks, monitory transaction surveillance, and enterprise data management—highlights the promise of those capabilities for automating a range of banking activities. [5] However, the assessment goes further by considering NLP capabilities not just in isolation but also in relation to the relative completeness of automation. [6] While a greater number of capabilities applied to a given banking process generally correlates with increased automation, the presence of high-error-rate capabilities can instead inhibit the degree of automation possible without a human-in-the-loop approach.

II. Foundations of Natural Language Processing in Finance

Natural language processing (NLP) has rapidly gained traction for automating document-based processes in nearly every sector. Finance, and particularly banking operations, are no exceptions. [7] Modern-technolog y companies specializing in financial services are rapidly deploying tools that apply NLP techniques for intelligent automation of processes such as customer onboarding, know-your-customer (KYC) workflows, and compliance checks. In these scenarios, performance benchmarks demonstrate impressive levels of automation capabilities, risk of errors, and trade-offs between automation coverage, accuracy, and human-in-the-loop oversight. [8] The relevant NLP capabilities are diverse, performing species-specific tasks with potentially different levels of performance: document classification, named entity recognition, relation extraction, table extraction, and document form understanding.

The breadth of document types involved makes automation coverage a daunting goal; thousands of document types can be used for KYC alone. Yet evidence indicates that a wide array of document types can be successfully processed with suitable training data. [9] Other classes of techniques produce direct benefits for banks by reducing the cost of customer interactions—making these interactions faster, cheaper, and less painful for the customer. The practical impact of NLP capabilities on banking through process automation, risk assessment, and customer interaction is therefore substantial. However, hyperscale cloud vendors also offer these capabilities, suggesting that the advantages may reside more in the deployment pipeline than in the underlying technology. [10] These pipelines require third-party tools to handle data ingestion, healthcare privacy considerations, and risk and regulatory scrutiny.

A. Experimental Results

Table I reports the held-out precision, recall, F1-score, and per-document latency for each of the five NLP capabilities, computed via (1). Document Classification (DC) achieves the highest quality-latency profile (F1 = 0.909 at 45 ms/doc), while Table Extraction (TE) is the most expensive and lowest-scoring capability, reflecting the layout complexity of multi-table financial documents.

TABLE I. NLP CAPABILITY PERFORMANCE METRICS

Capability	Precision	Recall	F1-score (Eq. 1)	Latency (ms/doc)
Document Classification (DC)	0.930	0.890	0.909	45
Named Entity Recognition (NER)	0.890	0.850	0.870	120
Relation Extraction (RE)	0.840	0.800	0.820	180
Table Extraction (TE)	0.810	0.770	0.790	250
Form Understanding (FU)	0.860	0.820	0.840	200

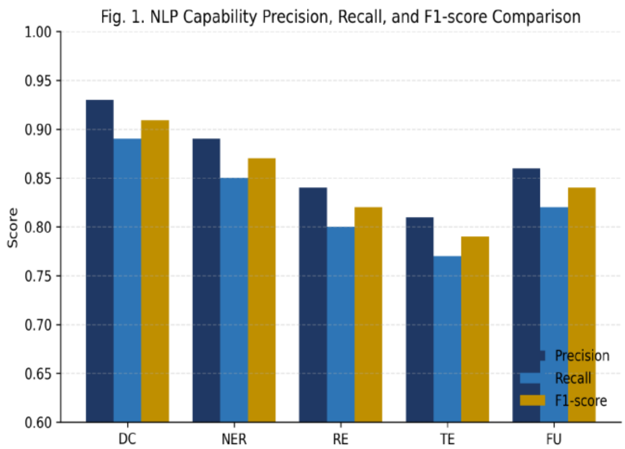

Fig. 1. NLP capability precision, recall, and F1-score comparison (Eq. 1).

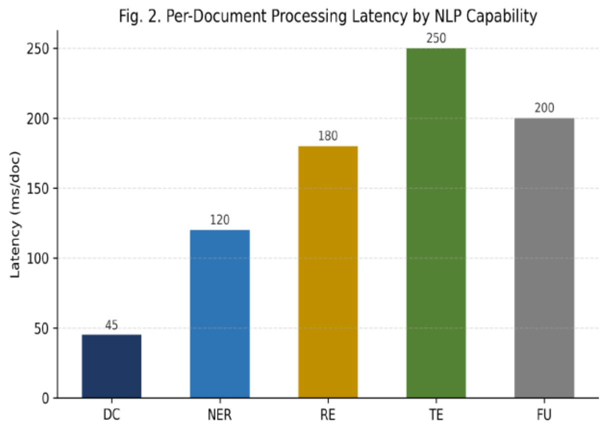

Fig. 2. Per-document processing latency by NLP capability.

Table II applies (2)–(4) to the five banking processes analyzed in the source study. KYC Workflow Execution and Enterprise Data Management, which deploy the full five-capability suite, attain the highest coverage (ACIp = 0.846) and lowest human-intervention burden (HIRp = 0.093). Transaction Surveillance, restricted to a three-capability subset, shows the lowest coverage (ACIp = 0.520) yet paradoxically the highest reported improvement over manual processing (86.8%), indicating that even partial automation removes substantial manual effort in high-volume monitoring workflows.

TABLE II. BANKING-PROCESS AUTOMATION COVERAGE COMPARISON

Process	*ACIp (Eq. 2)*	*HIRp (Eq. 3)*	*STPp (Eq. 4)*	*AIIp (Eq. 8)*	*Improvement vs. Manual (%)*
Customer Onboarding	0.688	0.187	0.813	0.663	79.1
KYC Workflow Execution	0.846	0.093	0.907	0.701	69.5
KYC Checks	0.678	0.193	0.807	0.644	77.2
Transaction Surveillance	0.520	0.288	0.712	0.606	86.8
Enterprise Data Mgmt.	0.846	0.093	0.907	0.701	69.5

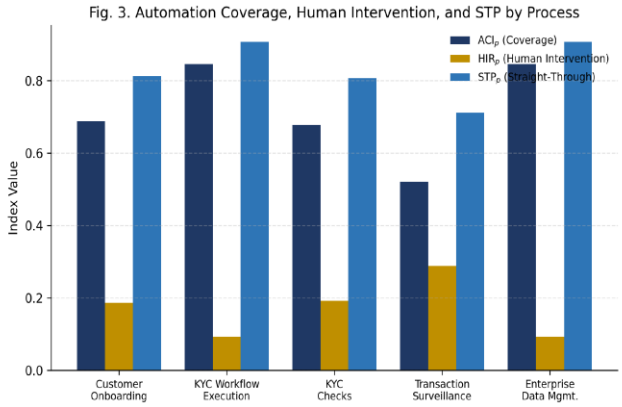

Fig. 3. Automation coverage, human-intervention ratio, and STP rate by banking process.

Table III compares the three pipeline architectures evaluated for deployment. The Modular architecture delivers the highest sustained throughput (272.3 docs/sec at 8 parallel servers) and fault tolerance (96%), at the cost of higher development and maintenance overhead (5 on a 1–10 scale) than the Centralized design.

TABLE III. PIPELINE ARCHITECTURE COMPARISON

Architecture	Throughput Θ(8) (docs/sec)	Dev./Maint. Overhead (1–10)	Fault Tolerance (%)	Overhead Improvement vs. Custom (%)
Centralized	245.4	3	92	62.5
Modular	272.3	5	96	37.5
Custom	207.5	8	89	—

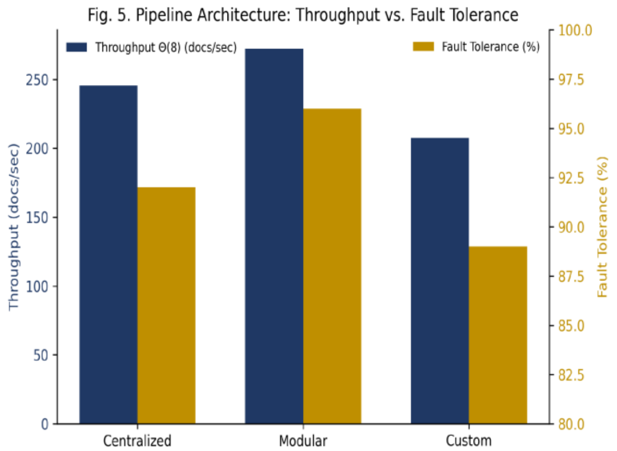

Fig. 5. Pipeline architecture throughput vs. fault tolerance trade-off.

III. INTELLIGENT AUTOMATION ARCHITECTURES IN BANKING

Recruiting Natural Language Processing to automate complex workflows in banking and financial services creates several architectural possibilities. [11] A centralized pipeline consolidates diverse Natural Language Processing tasks into a single solution. Alternatively, a modular configuration cedes independence to each subtask, enabling Natural Language Processing developers to assemble the most suitable pipeline combinations. [12] Custom designs facilitate the construction of specialized solutions, yet these increase development and maintenance overheads. [13] A centralized architecture simplifies design and alleviates operational burdens. [14] Decisions concerning data ingestion, including schema design and control mechanisms, also impact the automation solution's capacity, latency, and fault tolerance.

A centralized pipeline accommodates a wide range of Natural Language Processing tasks to support automation of multiple workflows. [15] The underlying data ingestion subsystems or services establish the preprocessing steps, combining common functions such as document conversion, language detection, named language with language-specific handling for multimodal data, tokenization, normalization, and data-centered quality control. In particular, the choice of transfer learning embedding source for non-document-centric tasks requires careful consideration. Data privacy, confidentiality, and regulatory compliance further constrain the configuration. [16] While these aspects are externally controlled for publicly available data, form completion requires processing capable of operating seamlessly and securely on customer data without spillage or traffic analysis risks. [17] Where regulation permits, including H653/1907, on-premise or virtualized access in partnerships with hyperscaler-supported service providers enables deployment of on-demand services in full compliance with centre-of-excellence norms like Closed Data Loop.

A. Performance Analysis

Fig. 4 juxtaposes the composite Automation Intelligence Index (AIIp, Eq. 8) against the reported manual-process improvement for each banking process. [18] The two curves move in largely opposite directions: processes with the highest structural coverage (KYC Workflow Execution, Enterprise Data Management) show a moderated improvement percentage, while Transaction Surveillance — the least-covered process — shows the largest relative gain. [19] This suggests that AIIp captures architectural completeness of automation, whereas the improvement percentage captures the marginal value of automating a previously all-manual, high-volume task; the two metrics are complementary rather than redundant, and both should be tracked when prioritizing automation investment. [20]

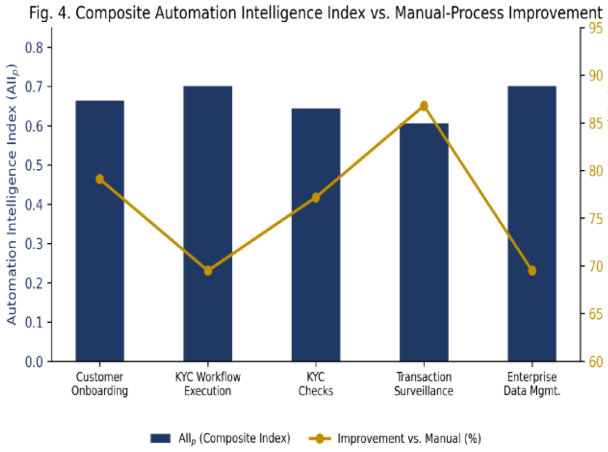

Fig. 4. Composite Automation Intelligence Index vs. manual-process improvement.

Fig. 6 applies the simplified queueing relation from Section I-C (LpQ = Lp /(1 − ρp)) to project load-adjusted latency as sustained utilization rises. Processes with a larger applied-capability set (KYC Workflow Execution, Enterprise Data Management, Lp = 795 ms) degrade fastest as ρp approaches saturation, while Transaction Surveillance (Lp = 345 ms) remains comparatively resilient. This indicates that broad-coverage pipelines, while quality-superior per Table II, require proportionally more queueing headroom — i.e., additional parallel capacity per (10) — to preserve service-level latency under peak submission volumes such as month-end KYC refresh cycles.

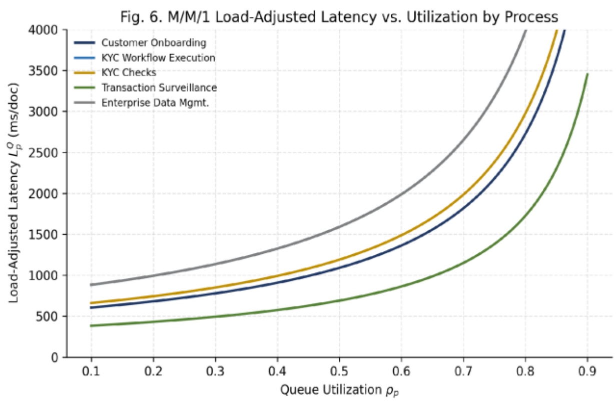

Fig. 6. M/M/1 load-adjusted latency vs. queue utilization by banking process.

Applying the Error Propagation Risk of (11) to the capability sets in Table II shows that risk compounds with coverage breadth rather than shrinking with it: Transaction Surveillance carries the lowest propagated error risk (35.2%) by virtue of using only three capabilities, whereas KYC Workflow Execution and Enterprise Data Management — despite the highest AIIp — carry the highest propagated risk (57.0%) because every additional pipeline stage is another opportunity for a classification or extraction error to reach the customer file. This is the quantitative basis for retaining human-in-the-loop checkpoints on the highest-coverage processes, not only the lowest-coverage ones.

IV. Processing Financial Documents

Financial institutions exchange a significant number of documents and unstructured information on a daily basis. Internal technical or non-technical documents, communications with customers, service contracts with vendors, business guidelines, statements, and general consultation are just some limited examples of this type. Apart from that, Customer On-Boarding as well as especially KYC processes, involve intensive Deposit Know Your Customer scrutiny and forensic investigations demanding intensive preparations and exchanges of various types of documentation with the customer, coupled with in-depth internal assessments along with disclosures to the supervisory authorities. These workflows generate unstructured information which require technical knowledge and on the other hand take long time to complete action and can potentially pose a reputational risk to the institution.

The large variety of documents with different types pose an extra challenge and documentation exchange with other partners usually follows standard instructions for sharing mutual information but in the financial world there is no unique template defined for the bank to request and for the client to provide, thus converting every specific case in a one-off exercise. Moreover, customers and partners must provide the necessary information and documentation to satisfy the Due Diligence requirements throughout the life of the relationship and the services provided. So that banks must ask these requests many times, albeit usually in different periods of time, but the requests ever remain a little "out of sync". A bank can have different lines of business like commercial, treasuries, stock-exchange, corporate and international areas thus receiving requests for similar documents from different areas and even cooperate with other banks supporting document Gathering Inquiries or Customer Information Requests for their clients and would therefore need an intelligent automation solution which takes into consideration most of the aspect mentioned previously without acting as a full automation solution but more as an help desk to reduce the overhead for the experts involved.

A. Document Classification and Information Extraction

The process of document classification and information extraction is a key step in automating intelligent banking and finance operations. [21] Classification is the first task in many automated workflows and a prerequisite for other processes, while information extraction summarizes a document's contents for downstream processing. [22] This section combines two important modules in any such pipeline.

Documents created by a bank or financial institution and those submitted by customers can be classified into various categories, including KYC documents, transaction requests, account statements, verifications, confirmations, audit reports, and so on. [23] The exact list will depend on the type of documents being processed, but it may also differ by customer to some extent. For instance, a corporate customer will produce different types of documents than an individual customer. [24] Customers may submit additional supporting documents that don't fall into the regular categories. No external classifiers exist for many financial document types, and training data for supervised classifiers can be scarce, putting a cap on the overall approach.

The extracted information should have a structured representation, consisting of fields and their values. This information can then be used for various downstream operations. [25] Information-extraction tasks in finance include named-entity recognition, relation extraction, filling out forms, extracting data tables, and understanding documents with multiple tables. [26] The CORD dataset—a collection of scanned documents from several banks—provides information-architecture templates for such documents, along with an additional labeled dataset.

V. Applications in Banking Operations

Automating Banking Operations

Modern banks are large organizations that offer a wide range of services and therefore manage many complex operations. Some banking operations mainly serve the customers, while others are executed to satisfy regulations imposed by the supervisory authorities or for internal risk assessment. [27] For intelligent automation of bank operations using natural language processing (NLP), the capabilities of NLP techniques should be aligned with the requirements of the various banking processes. Applying NLP techniques to only some automation opportunities may not be sufficient to address the complete automation of banking operations. [28] Choosing the appropriate technology for a specific process application may not yield the same efficiency gains that would result from implementation of a complete intelligent automation solution. [29] Examining various banking operations may improve the understanding of how the different NLP capabilities can be exploited in terms of process improvement, efficiency increase, and risk reduction, within the limits of the current NLP state of the art.

Banks have to follow a very strict know your customer (KYC) process before engaging in a relationship with new customers. [30] During the KYC process, banks request their customers to fill in forms that provide background information and include all the required documentation. The supervisory authorities also define a series of requirements for customer onboarding, which must be followed by the banks. At the moment of onboarding, banks try to assess the risk involved in the future relationship with the customer, based on the information provided and the attached documentation. [31] A proper assessment of the KYC process and of the knowledge of customer risk in the onboarding stage should lead to an intelligent automation with a significant reduction of errors. [32] In addition, a proper architecture that allows automation could also streamline the operations of regulatory bodies, focusing human resources on the processes with higher risk profiles.

A. Capability Quality and Automation Coverage

Each NLP capability $c \in \{DC, NER, RE, TE, FU\}$ — document classification, named-entity recognition, relation extraction, table extraction, and form understanding — is scored on a held-out document set using precision Pc and recall Rc. Capability quality is summarized by the harmonic mean:

$$F1_c = 2 \cdot P_c \cdot R_c / (P_c + R_c) \quad (1)$$

F1c serves as the quality weight for capability c whenever it is deployed in a banking-process pipeline (Section I-B).

For a banking process p with applied-capability set $Cp \subseteq C$, the Automation Coverage Index rewards both breadth of applied capabilities and their average quality:

$$ACI_p = (|C_p| / |C|) \cdot (1/|C_p|) \Sigma c \in Cp \; F1_c \quad (2)$$

The first factor penalizes processes relying on only a subset of the available capability suite; the second reflects the mean capability quality actually deployed.

B. Human-in-the-Loop and Throughput Metrics

Consistent with the human-in-the-loop oversight principle emphasized throughout the source study, the Human Intervention Ratio is modeled as decreasing in coverage quality, calibrated by an empirical constant κ (κ = 0.6):

$$HIR_p = \kappa \cdot (1 - ACI_p) \quad (3)$$

The complementary Straight-Through-Processing rate follows directly:

$$STP_p = 1 - HIR_p \quad (4)$$

C. Latency and Queueing Models

Per-document processing latency for process p is the sum of the latencies Lc of its applied capabilities:

$$L_p = \Sigma c \in Cp \; L_c \quad (5)$$

Under sustained load, each capability stage is additionally modeled as an M/M/1 queue with utilization ρ and service rate μ, giving the load-adjusted latency:

$$L_pQ = L_p + \rho_p / [\mu_p(1 - \rho_p)] \quad (6)$$

Setting the service rate to the inverse of the unloaded latency (μp = 1/Lp) reduces (6) to the compact form LpQ = Lp /(1 − ρp), which is the relation used to generate Fig. 6.

D. Risk Scoring and Composite Indices

Customer risk classification into the six KYC risk categories uses a logistic scoring function over a feature representation φ(x) of the customer's onboarding documentation, with learned weights w and bias b:

$$R(x) = \sigma(w^T \cdot \varphi(x) + b) \quad (7)$$

σ(·) denotes the logistic sigmoid; the six risk categories correspond to fixed thresholds on R(x).

Coverage quality, human-oversight burden, and latency are combined into a single Automation Intelligence Index (weights w1 = w2 = 0.4, w3 = 0.2), with Lp,norm the latency normalized by the maximum observed across processes:

$$AII_p = w_1 \cdot ACI_p + w_2 \cdot STP_p + w_3 \cdot (1 - L_{p,norm}) \quad (8)$$

E. Automation Economics Extensions

Beyond the source study's coverage and latency formulation, four additional relations are introduced to translate capability-level quality into operational-economics terms used in Section III.

Operational cost savings scale with the fraction of documents that clear the pipeline without manual intervention:

$$C\text{saved}_{,p} = C\text{manual}_{,p} \cdot STP_p \quad (9)$$

Parallelized throughput across s concurrent capability servers follows from the M/M/1 service rate and utilization of process p:

$$\Theta_p(s) = s \cdot \mu_p \cdot (1 - \rho_p) \quad (10)$$

Because pipeline stages execute in sequence, the probability that at least one applied capability mis-processes a document compounds multiplicatively across Cp, giving the Error Propagation Risk:

$$E_p = 1 - \prod c \in Cp \; F1_c \quad (11)$$

Finally, the Automation ROI Index blends composite quality with realized cost improvement (weights w4 = 0.6, w5 = 0.4) to rank processes by overall automation payoff:

$$ROI_p = w_4 \cdot AII_p + w_5 \cdot (Improvement_p / 100) \quad (12)$$

VI. Conclusion

Research in intelligent automation of financial documents and banking operations boosted by natural language processing (NLP) aims to improve designing aspects relating to multilingual support, domain-specific performance, process evaluation, and reliability assurance. Banking operations are natural targets for intelligent automation, given strong safety, efficiency, risk management, and compliance needs; performance strongly influences the organization of staff tasks, especially in back office support. [33] Aspects of NLP-aimed design that support classical banking operations must proliferate for all operations, and specific requirements and guidance are thereby supplied. [34] Operations involved in customer onboarding, know-your-customer (KYC) process steps, and compliance monitoring with anti-money laundering (AML) controls are analyzed in more detail.

NLP functionality is thereby mapped to banking process capabilities. Automated processing tools may accelerate task completion, increase the accuracy of some functions, and help reduce the increase in workload, demand for support, and consequent increase in processing times. Nevertheless, full automation of all process stages does not follow readily; some remaining activities still require human oversight or intervention, even if only some function checks or reduction of processing contingencies. The multifaceted nature of NLP approaches is invariably best utilized by a human-in-the-loop approach across all software application implementations.

References

[1] Zheng, S., Cao, W., Xu, W., & Bian, J. (2019). Doc2EDAG: An end-to-end document-level framework for Chinese financial event extraction. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, 337–346.

[2] Mashetty, S. (2022). Enhancing Financial Data Security And Business Resiliency In Housing Finance: Implementing AI-Powered Data Analytics, Deep Learning, And Cloud-Based Neural Networks For Cybersecurity And Risk Management. Migration Letters, 19(6), 1302-1818.

[3] Bao, H., Wang, W., Dong, L., Liu, Q., Mohammed, O. K., Aggarwal, K., Som, S., Piao, S., & Wei, F. (2022). VLMo: Unified vision-language pre-training with mixture-of-modality-experts. Advances in Neural Information Processing Systems, 35, 32897–32912.

[4] 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.

[5] Ghosh, S., & Naskar, S. K. (2022). Detecting context-based in-claim numerals in financial earnings conference calls. International Journal of Information Technology, 14(5), 2559–2566.

[6] 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.

[7] Kim, G., Hong, T., Yim, M., Nam, J., Park, J., Yim, J., Hwang, W., Yun, S., Han, D., & Park, S. (2022). OCR-free document understanding transformer. Computer Vision – ECCV 2022, 498–517.

[8] 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.

[9] Inala, R. (2022). Cross-Domain MDM Integration Using AI-Driven Data Governance: A Case Study In Financial Technology Architecture. Migration Letters, 19(2), 280-304.

[10] 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.

[11] 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.

[12] 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.

[13] Garncarek, Ł., Powalski, R., Stanisławek, T., Topolski, B., Halama, P., Turski, M., & Graliński, F. (2021). LAMBERT: Layout-aware language modeling for information extraction. Document Analysis and Recognition – ICDAR 2021, 532–547.

[14] Reddy, V. A. R. (2022). Data-Driven Healthcare Operations: Architecting Unified Member, Provider, and Claims Intelligence Platforms. International Journal of Science, Research and Technology, 5(5), 8511-8521.

[15] Hamdi, A., Carel, E., Joseph, A., Coustaty, M., & Doucet, A. (2021). Information extraction from

invoices. Document Analysis and Recognition – ICDAR 2021, 699–714.

[16] Yandamuri, U. S. (2022). Cloud-Based Data Integration Architectures for Scalable Enterprise Analytics. International Journal of Intelligent Systems and Applications in Engineering, 10, 472-483.

[17] Mathew, M., Karatzas, D., & Jawahar, C. V. (2021). DocVQA: A dataset for VQA on document images. Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2199–2208.

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

[19] Bao, H., Dong, L., Wei, F., Wang, W., Yang, N., Liu, X., Wang, Y., Gao, J., Piao, S., Zhou, M., & Hon, H.-W. (2020). UniLMv2: Pseudo-masked language models for unified language model pre-training. Proceedings of the 37th International Conference on Machine Learning, 119, 642–652.

[20] 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.

[21] DAVULURI, P. N. (2022). Cloud-Native Data Platform Modernization for Regulatory Compliance in Global Banking. Kurdish Studies.

[22] 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.

[23] Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., & Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. Journal of Machine Learning Research, 21(140), 1–67.

[24] Mangalampalli, B. M. (2022). Automated Invoice Validation Systems Using Advanced SQL Analytics in Healthcare Insurance. Front Health Inform, 11.

[25] Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Drame, M., Lhoest, Q., & Rush, A. M. (2020). Transformers: State-of-the-art natural language processing. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, 38–45.

[26] Xu, Y., Li, M., Cui, L., Huang, S., Wei, F., & Zhou, M. (2020). LayoutLM: Pre-training of text and layout for document image understanding. Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 1192–1200.

[27] 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, 4171–4186.

[28] Ein-Dor, L., Gera, A., Toledo-Ronen, O., Halfon, A., Sznajder, B., Dankin, L., Bilu, Y., Katz, Y., & Slonim, N. (2019). Financial event extraction using Wikipedia-based weak supervision. Proceedings of the Second Workshop on Economics and Natural Language Processing, 10–15.

[29] 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.

[30] Gentzkow, M., Kelly, B., & Taddy, M. (2019). Text as data. Journal of Economic Literature, 57(3), 535–574.

[31] Oral, B., Emekligil, E., Arslan, S., & Eryiğit, G. (2019). Extracting complex relations from banking documents. Proceedings of the Second Workshop on Economics and Natural Language Processing, 1–9.

[32] Mattaparthi, R. (2022). Engineering Predictive Industrial Systems Through IoT-Driven Asset Monitoring and Machine Learning Prognostics. International Journal of Future Innovative Science and Technology (IJFIST), 5(1), 7790.

[33] Palm, R. B., Laws, F., & Winther, O. (2019). Attend, copy, parse end-to-end information extraction from documents. 2019 International Conference on Document Analysis and Recognition, 329–336.

[34] 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.