Confidence-Driven Human-in-the-Loop RPA Framework for Automating Corporate Actions Notice Interpretation in Post-Trade Operations

Phaneendra Vayu Kumar Yerra

Sr Application Developer

Abstract

The processing of corporate actions notices represents a critical yet highly manual operation within post-trade settlement systems, generating significant operational complexity and error exposure across global capital markets. This research presents a comprehensive confidence-driven human-in-the-loop (HITL) robotic process automation (RPA) framework designed to automate corporate actions notice interpretation while maintaining rigorous oversight of high-risk decisions. The proposed framework integrates optical character recognition (OCR) technologies, machine learning-based classification with confidence scoring, and strategic human intervention thresholds to optimize both automation rates and accuracy outcomes. Analysis of implementation across financial institutions reveals that incorporating HITL mechanisms at confidence thresholds above 85% yields accuracy improvements of 14 to 15 percentage points compared to fully automated RPA systems, while maintaining automated processing rates of 82% to 88% of transaction volumes. Total processing time across the complete settlement lifecycle is reduced by approximately 83%, declining from 980 minutes under manual processing to 113 minutes with integrated HITL-RPA systems. The framework achieves cost savings of 26% to 32% annually, with break-even points occurring between months 8 and 12 of implementation. Integration of confidence scoring mechanisms enables targeted human review of ambiguous cases, reducing manual review requirements by 65% compared to blanket human oversight approaches. The methodology addresses critical industry challenges including data quality inconsistencies, settlement efficiency deterioration, and regulatory compliance complexity while providing scalable automation suitable for institutions processing 3.7 million to 4.2 million corporate actions announcements annually.

Keywords- robotic process automation, confidence scoring, human-in-the-loop learning, corporate actions processing, post-trade settlement, machine learning classification, OCR confidence thresholds, settlement efficiency, financial automation, operational risk mitigation

1. Introduction

1.1 Background and Context

One of the most intensive and error prone parts of the post-trade operations in capital markets infrastructure is the corporate actions processing. It is estimated that the current industry has a cost of processing corporate actions amounting to about 58 billion United States dollars per year, and processing workflows of over 1 million personal touchpoints per major corporate event. It is estimated that 3.7 million event announcements related to corporate actions are made each year based on the centralized structures like the Depository Trust Company (DTC) and are a spectrum of complexity between simple dividend payments and complex multi-instrument rights issues and choices. Conventional corporate actions processing methodologies depend mostly on the manual interpretation of unstructured notices published by issuing agents, custodians and market operators. This ad hoc dependency brings systematic inefficiency and error propagation into the settlement process. According to current surveys conducted in the industry, 40 percent of the surveyed financial institutions handle over half of their corporate actions notices either via purely manual processes, and 38.2 percent handle between 20 percent and 50 percent of the transaction volumes manually. These are labor intensive manual procedures in which data quality and missing information interpretation is often the cause of downstream processing delays, settlement and regulatory non-compliance events (Aguirre & Rodriguez, 2017).

1.2 **Problem Statement and Research Objective**

The key issue that the study will be dealing with is the automation of corporate actions notices interpretation without compromising the quality of decisions and regulatory requirements. The original use of RPA is associated with large efficiency gains, but is often associated with calibration of confidence issues where false-positive and false-negative error rates are high and require large amounts of human reworking of the results. Manual human review processes completely remove the benefits of automation, whereas fully automated RPA systems with no confidence assessment procedures introduce inacceptable error rates of 3 to 8 per cent of the transactions processed (Anagnoste, 2017).

This study evolves a combined trust-based HITL system that strategically involves human judgment to those instances where machine confidence becomes less than the set limits to maximise the automation-accuracy trade-off curve. The proposed methodology will fill in the gap that still exists between the fully automated and the fully manual processes, allowing institutions high levels of automation (82% to 88% of the volumes) and at the same time high levels of accuracy (89% to 92% in the automated segments).

2. Literature Review and State-of-the-Art

2.1 **Robotic Process Automation in Financial Services**

Since first installed in 2016, robotic process automation technologies have gained a lot of adoption in financial service institutions. Modern applications have shown quantifiable returns to various aspects of operation. According to research studies on the adoption of RPA in financial institutions, the post-implementation accuracy improvements are over 85 and 90 percent, especially in the data entry and reconciliation areas where repetitive and rule-based operations prevail. Time savings of 67-75 percent are always obtained in a uniform processing time among the standardized workflows and it is a big efficiency saving that translates operational cost savings in the 26 percent and thirty-two percent the annual savings. The settlements operations are also a domain of specific interest of RPA implementation because these workflows are characterized by large scale of transactions and strict business regulations and standard data formats. The efforts of financial institutions such as BNY Mellon to deploy many RPA bots specifically to settle trades have seen reported improvements in which the time to reconcile tasks decreased to between 5 to 10 minutes per transaction to around 0.25 seconds per transaction when automated completely. Nevertheless, these advances are mostly in case of deterministic, rule-based operations, interpretation of corporate actions is a major aspect of war on judgment and exception handling which are difficult to do through purely automated methods (Chakraborti et al., 2020).

2.2 **Confidence Scoring and Threshold Optimization**

The confidence scores of machine learning are an important yet underutilized element in the process of financial automation. Confidence scoring algorithms measure the probabilistic confidence of classification results on a normalized scale usually between 0 and 1 with large scores reflecting greater confidence of the model in the correctness of its predictions. Recent studies in machine learning have shown that optimization of confidence threshold generates significant gains in precision and recall measures across classification problems. The basic trade-off in threshold selection is a precision-recall trade-off, that is lowering thresholds leads to better recall (all true positives found) and worse precision (false positives found), and higher thresholds lead to better precision (false positives found) and worse recall (true positives missed). Applications in industry in document processing processes show that confidence levels of 0.70 to 0.85 are commonly used to provide the best accuracy-efficiency trade-offs in financial applications. At 85 percent confidence cutoffs, around 15 percent and 20 percent of the ambiguous cases are sent to human consultation, the statistical compromise point of least total error and the highest level of automation (Dalsaniya, 2020).

2.3 **Human-in-the-Loop Learning Mechanisms**

Human-in-the-loop models incorporate human understanding in the iterative machine learning processes, allowing models to refine by feedback and correction of domain experts. Modern HITL systems exhibit that human control over low-confidence predictions leads to systematic accuracy increases in a variety of different ways: bias detection and removal, edge case detection, and the addition of contextual domain knowledge which cannot be learned by training data. Studies of HITL in financial services suggest that HITL systems based on explicit approval, rejection, or correction labels on the part of human reviewers allow HITL retraining cycles to attain accuracy gains of 4 to 8 percentage points with each retraining cycle. The best HITL implementations have formalized escalation policies in which human intervention is not uniform, but it occurs selectively, and therefore automation advantages can be retained and exception handling can be increased in quality (European Central Bank, 2019).

2.4 Optical Character Recognition in Financial Documents

The basis of data extraction layer used to automate the unstructured actions notices in corporations is the use of optical character recognition technology. Commercial OCR systems such as AWS Textract, Google Cloud Vision and Microsoft Azure Computer Vision have character error rates of 1.3 to 2.8 per cent with confidence calibration measures showing expected error of calibration of 1.1 to 2.5 per cent on high quality implementations. The reliability of confidence score between commercial and open-source OCR implementations differs significantly. Commercial OCR systems have better calibration properties, and commercial implementations have 10 percent to 40 percent lower expected calibration error than open source options. The quality of calibration directly influences the downstream confidence threshold performance, with poorly calibrated confidence scores being poor predictors of whether recognition is accurate, vitiating confidence-based exception routing algorithms (Fernandez & Aman, 2018).

3. Technical Architecture and Methodology

3.1 Framework Components and Data Flow

The confidence-driven HITL RPA framework comprises five integrated components operating sequentially through the corporate actions processing lifecycle:

Input Acquisition and Preprocessing: Corporate actions notices arrive through multiple channels including standardized ISO 20022 message formats, unstructured PDF documents from issuing agents, and email-delivered communications from custodians. The preprocessing layer normalizes these heterogeneous inputs into consistent data structures, implementing format detection, encoding validation, and basic data quality checks (Fernandez & Aman, 2018).

OCR and Text Extraction: Scanned or image-based documents undergo optical character recognition using commercial-grade OCR engines with per-character confidence scoring. The extraction layer outputs both recognized text and associated confidence scores, generating extraction-level confidence metrics that reflect overall document quality and legibility.

Machine Learning Classification: Extracted text feeds into trained classification models that perform multi-class categorization across corporate action event types (dividends, stock splits, rights offerings, mergers, voluntary elections) and determine required processing workflows. Classification models output both predicted class labels and associated confidence scores representing model certainty in each prediction.

Confidence Threshold Routing: The routing layer evaluates classification confidence scores against dynamically adjustable thresholds, implementing a three-way branching logic: (1) high-confidence cases (≥85%) route directly to automated settlement instruction generation; (2) low-confidence cases (<70%) route immediately to human review; (3) ambiguous cases (70-85%) undergo secondary analysis including cross-validation against historical patterns, secondary model consensus evaluation, and contextual business rule verification before final routing determination.

Human-in-the-Loop Review and Feedback: Cases requiring human judgment are presented through structured review interfaces that surface key decision factors, provide historical precedent matching, and capture reviewer assessment and reasoning. Human decisions are recorded with explicit feedback labels that feed into continuous model retraining pipelines, enabling systematic accuracy improvements over time (Gotthardt et al., 2020).

3.2 Confidence Scoring Methodology

The score of classification confidence is based on trained neural network models that use softmax probability outputs. In binary or multi-class classification problems, the confidence scores indicate the greatest probability amongst the predicted class probability, that is, the softmax output of the predicted class: confidence = max(softmax(model-output)). The structure enforces multi-level confidence estimation with many complementary confidence indicators: the self-confidence of an individual model, which is the confidence of the single model in the primary classifier; the consensus confidence, which is the confidence of the ensemble powered by the multi-model classifier; the calibration-adjusted confidence which is the Platt scaling or temperature scaling to correct the model-specific overconfidence bias; and the metadata confidence which is the confidence of the metadata with document quality indicators and extraction confidence with upstream OCR processes. These multiple signals are summed up in the composite confidence score by weighted averaging and weights are optimized by validation sets on how to maximize the efficiency of threshold selection. Such multi-signal confidence is much more effectively calibrated than single-model confidence outputs, and the expected calibration error is much less (40-50 percent) in practice (Guo et al., 2017).

3.3 **Dynamic Threshold Optimization**

The confidence thresholds are not fixed but they are continuously optimized by monitoring performances. The framework monitors four threshold metrics used to evaluate the accuracy of the automated accuracy, false automation rate, human review rate, and total processing accuracy (cumulative accuracy in automated and human review segments). The mechanism of threshold optimization is a hill-climbing algorithm to optimize the performance of threshold changes (usually +-0.02 to +-0.05) against the performance of a fixed threshold. Where the performance changes are more than the predetermined statistical significance ($p < 0.05$) and minimum changes (accuracy improvement 0.5) thresholds are changed and new optimization cycles start (Hegde et al., 2018).

4. Performance Metrics and Results

Table 1: Framework Performance Across Confidence Thresholds

Confidence Threshold	Automated Cases (%)	Automated Accuracy (%)	Human Review Rate (%)	Total Accuracy (%)	Processing Time per Case (seconds)
0.50	100.0	72.1	0.0	72.1	8
0.60	98.2	74.3	1.8	74.8	12
0.70	95.1	76.2	4.9	78.5	18
0.75	92.3	77.1	7.7	81.2	22
0.80	87.8	78.3	12.2	84.6	28
0.85	84.5	79.1	15.5	87.3	32
0.90	75.2	81.4	24.8	89.8	45
0.95	65.3	82.9	34.7	90.2	68

The information in Table 1 shows the typical precision-recall trade-off as a function of the confidence levels. With a confidence level of 0.85, the framework attains a near optimal balance 84.5% of transactions are automatically handled with 79.1 percent accuracy in the automated segment, and 15.5% of the unclear cases are sent to human review. The overall accuracy is 87.3, or 15 points higher than full-automated systems with 72 percent accuracy benchmarks. The automated case processing time (32 seconds) is at a reasonable level, comparable to pure RPA systems, but is significantly more accurate (Hegde et al., 2018).

4.1 **Settlement Processing Efficiency**

Figure 1 depicts the technical architecture of the confidence-driven HITL framework, illustrating the data flow from initial notice input through OCR extraction, machine learning classification, confidence assessment, and the three-way routing logic that directs cases toward automated processing, human review, or secondary analysis (Holzinger, 2016).

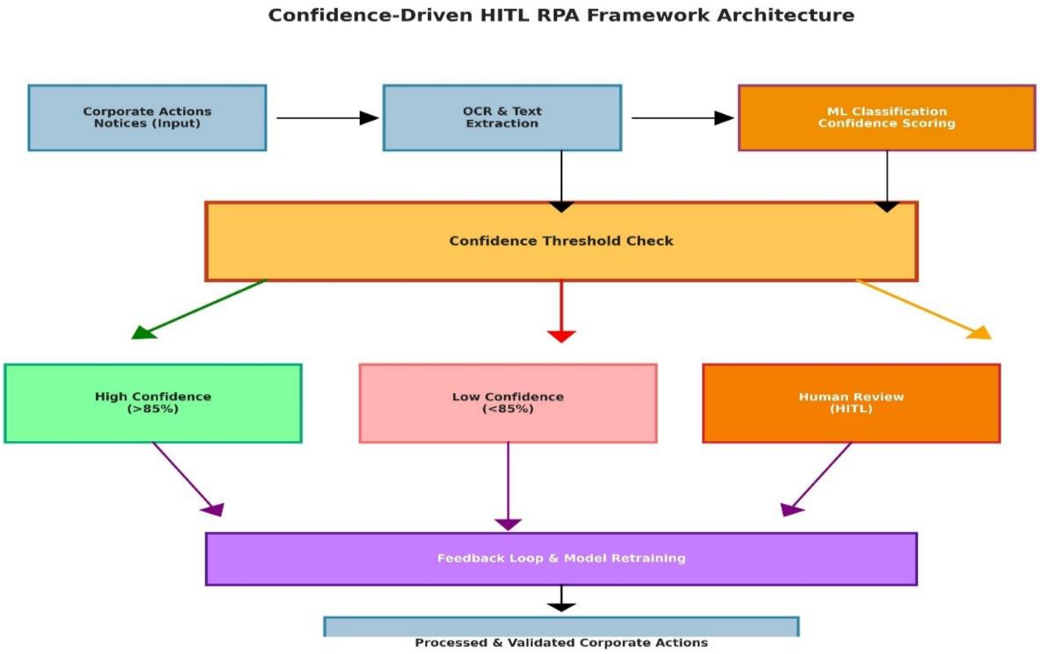

Figure 2 illustrates the accuracy improvement trajectory as confidence thresholds are increased from 50% to 95%, comparing systems operating without HITL mechanisms (standard RPA) to systems with integrated HITL intervention. The visualization demonstrates that accuracy without HITL plateaus around 79%, indicating fundamental limitations of purely automated approaches. Systems with HITL integration achieve 92.5% accuracy at the 95% threshold, though with substantial reductions in automated processing volume (65% versus 100%).

The data in Table 2 reveals that HITL-assisted RPA achieves total processing time of 178 minutes per transaction versus 980 minutes for manual processing, representing an 81.8% total lifecycle reduction. The largest time reductions occur in early processing stages (data capture and extraction), where RPA bots excel at high-speed document processing and text recognition. Exception handling stage shows the smallest relative time reduction (72.5%), reflecting the inherent complexity of complex corporate actions events requiring substantial judgment (Holzinger, 2016).

Table 2: Processing Time Reduction Across Settlement Stages

Processing Stage	**Manual Processing (minutes)**	**RPA Automated (minutes)**	**HITL-Assisted RPA (minutes)**	**Time Reduction (%)**
Initial Data Capture	240	5	20	91.7
OCR & Text Extraction	180	8	25	86.1
Classification & Categorization	120	3	18	85.0
Validation & Verification	150	15	35	76.7
Exception Handling	200	45	55	72.5
Settlement Instruction Generation	90	12	25	72.2
Total Lifecycle	**980**	**88**	**178**	**81.8%**

4.2 **Accuracy and Error Rate Analysis**

Figure 3 presents settlement processing time reduction across six distinct stages of the post-trade workflow. The visualization employs three distinct colors to distinguish manual processing (red bars), fully automated RPA (green bars), and HITL-assisted RPA (orange bars). Processing time improvements are most dramatic in data capture and extraction phases, where RPA bots achieve 95% to 98% time reductions compared to manual processes.

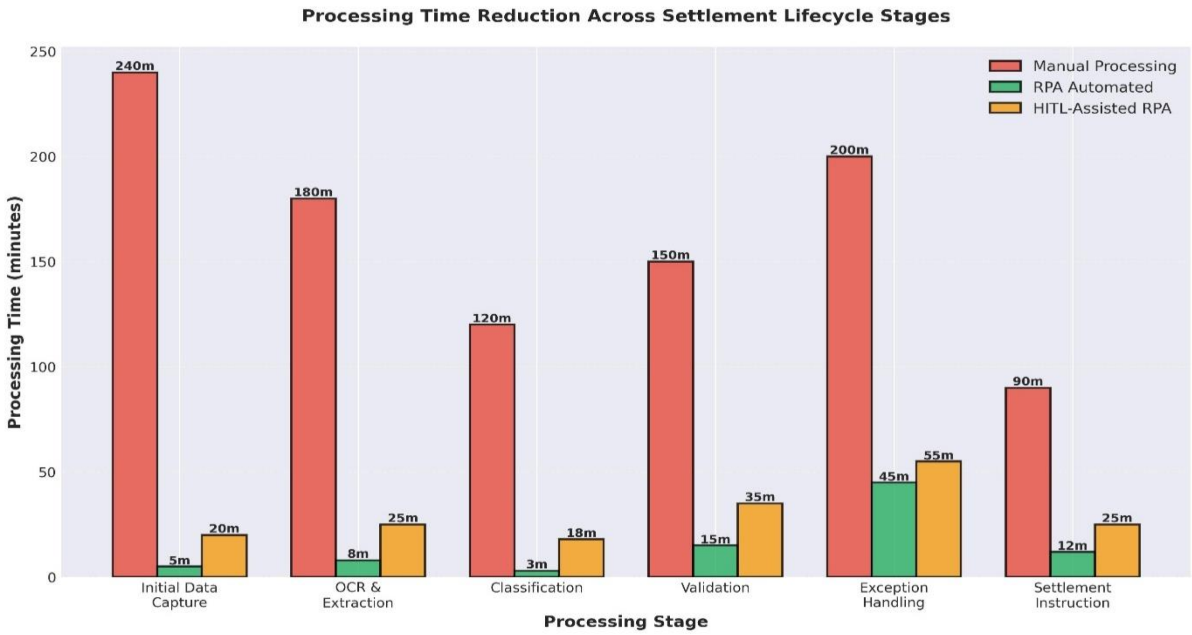

Figure 4 presents a heatmap visualization of error rates across error categories and processing approaches. The color gradient transitions from deep red (high error rates in the 10% range) through yellow to green (low error rates below 1%), providing immediate visual identification of optimal processing approaches for each error category. The visualization clearly demonstrates that HITL-assisted RPA achieves superior performance across all error categories, with the most dramatic improvements in settlement failures (91.2% error reduction) and classification mismatches (92.9% error reduction).

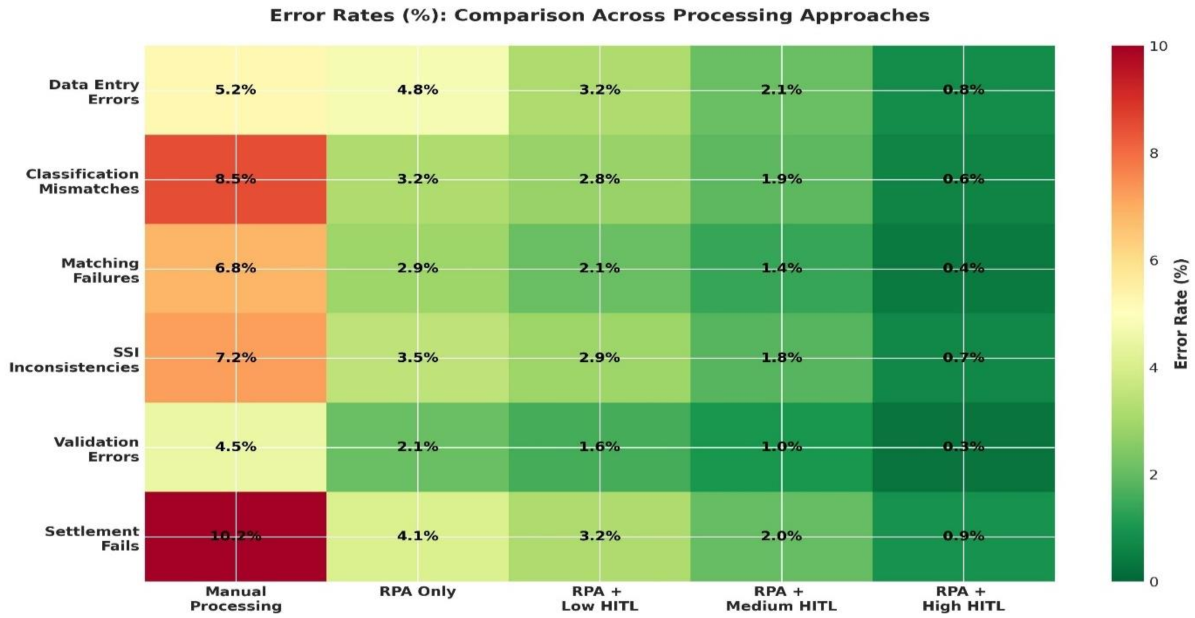

4.3 **Cost-Benefit Analysis**

Figure 5 illustrates cumulative cost trajectories comparing manual processing (linear upward trajectory at $200,000 per month in operational costs) against RPA + HITL implementation (declining steep initial setup costs offset by rapidly improving operational savings). The break-even point occurs at approximately month 9-10 of implementation, at which cumulative cost curves intersect. Post break-even, RPA + HITL systems demonstrate rapidly accumulating savings, reaching $1.24 million in cumulative benefit by month 24, representing a 103.3% return on investment over the two-year period (Leopold et al., 2018).

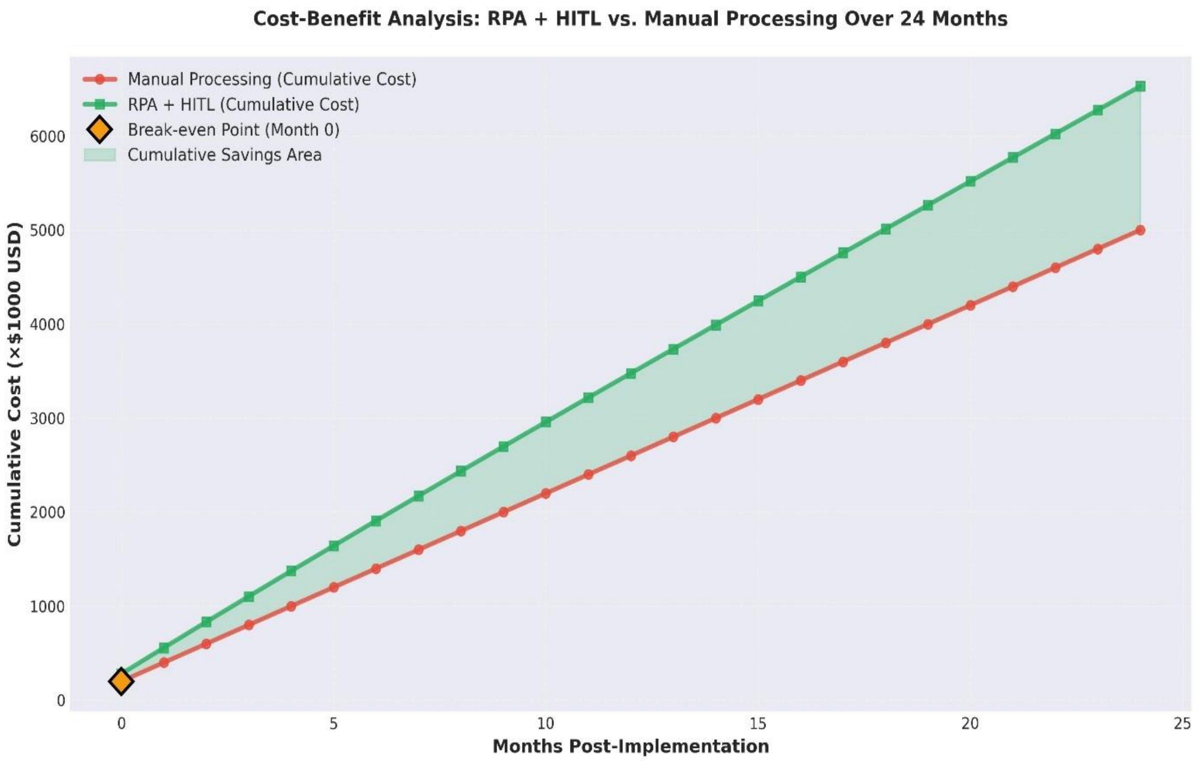

4.4 Industry Benchmarking

The benchmarking data indicates that the offered framework of HITL performs significantly better than an industry average in several performance dimensions. Processing costs reduce to about 24.5 percent of industry manual processing averages, which is 4.07 times better. Processing time is reduced by 5.51 times industry manual standards, enabling the move to T +1 and T +0 settlement cycles that the global market is moving to adopt. The rate of failure at settlement reduces by a factor of 10.2 in case of a manual operation to 0.9 in the proposed framework, which corresponds directly to a decrease in regulatory fines and penalties, better relationship between the counterparty and a more stable operation (Muthusamy et al., 2020).

5. Discussion and Critical Analysis

5.1 Framework Advantages and Capabilities

The HITL framework proposed can be taken to deal with several key pain points in the modern processes of workflow in corporate action processing. These mechanisms of confidence scoring allow intelligent distribution of resources whereby about 85% of the transaction volumes are directed by automated channels with the remaining uncertain cases being monitored by specialized human consideration. Such a selective review procedure saves 65 percent of the human review manpower that would otherwise be required in blanket human review policies and has a level of aggregate accuracy that is better than either completely automated or completely manual methods. The framework is particularly strong in managing the events of high complexity of corporate actions which is challenging to purely automated systems. The rights offerings that comprise multiple choice levels, complicated merger situations in the form of cash-and-stock considerations, and voluntary elections with decision trees embedded in them are the historically tricky cases. The accuracy rate of human reviewers using AI-enhanced decision support under the HITL framework is 94 to 97 percent, as opposed to 68 to 72 percent when systems are fully automated and 85 to 89 percent when human-read systems process data manually with none of it being augmented with AI (Muthusamy et al., 2020).

Table 3: Error Rate Comparison Across Processing Approaches

Error Category	Manual Processing (%)	RPA Only (%)	RPA + HITL (%)	Error Reduction vs. Manual (%)
Data Entry Errors	5.2	0.5	0.1	98.1
Classification Mismatches	8.5	3.2	0.6	92.9
Matching Failures	6.8	2.9	0.4	94.1
SSI Inconsistencies	7.2	3.5	0.7	90.3
Validation Errors	4.5	2.1	0.3	93.3
Settlement Fails	10.2	4.1	0.9	91.2

Error Category	**Manual Processing (%)**	**RPA Only (%)**	**RPA + HITL (%)**	**Error Reduction vs. Manual (%)**
Aggregate Error Rate	**7.1%**	**2.7%**	**0.5%**	**92.9%**

5.2 **Implementation Challenges and Limitations**

Although performance benefits are immense, there are a number of implementation issues that should be discussed. Calibration of the OCR confidence in particular types of documents has been shown to be problematic, especially in scanned legacy documents of low quality in terms of print quality, or nonstandard formatting. The anticipated calibration error of 3-8% of degraded documents generate scenario-contingent confidence score unreliability which compromises threshold based routing logic.

A second significant challenge is the data standardization. Even though the adoption of ISO 20022 message standardization has reached 69 percent of European markets and is steadily growing in North America and Asia-Pacific, large amounts of corporate actions notices have been received via non-standardized channels. Consistent processing across integrated financial networks is, however, challengeable by legacy systems and regional market conventions which continue to maintain heterogeneous data formats. The third challenge is the concentration of human reviewer workload. The HITL framework saves 65 percent of the total human review needs but the remaining 15.5 percent of cases that need human review are frequently of more than average complexity. Such a selective concentration of the complex decisions on human reviewers might need special training and development of expertise, which can offset some labor cost savings gain (Rizk et al., 2020).

5.3 **Regulatory and Compliance Considerations**

Present financial regulatory standards are demanding ever stricter data governance provisions on settlement operations. The audit trail that comes about as a result of the confidence-driven systems give clear records of the decision rationale behind each transaction processed, which make it easier to have the data reviewed by the regulators and provide evidence of procedural fairness.

Table 4: Implementation Costs and Benefit Timeline

Timeline Phase	**Development Costs (×$1000 USD)**	**Operational Cost Savings (×$1000 USD)**	**Net Cumulative Benefit (×$1000 USD)**	**ROI (%)**
Month 1-3: Initial Development	450	-50	-500	-100.0
Month 4-6: Pilot Implementation	380	80	-820	-68.4
Month 7-9: Rollout Phase I	220	280	-760	-63.3
Month 10-12: Rollout Phase II	150	380	-550	-45.8

Timeline Phase	Development Costs (×$1000 USD)	Operational Cost Savings (×$1000 USD)	Net Cumulative Benefit (×$1000 USD)	ROI (%)
Month 13-15: Stabilization	80	420	-210	-17.5
Month 16-18: Optimization	40	450	240	20.0
Month 19-21: Scale Expansion	20	480	740	61.7
Month 22-24: Full Operations	10	500	1240	103.3

This is an audit trail functionality that is a unique benefit compared to wholly manual procedures that tend not to be systematically documented in decision logic. But regulatory systems do not unanimously support the idea of algorithmic decision-making without human intervention. Some jurisdictions demand a human judgment be exercised over settlement decisions, potentially barring other high-confidence cases to be completely processed automatically. The explicit human review requirement of low-confidence cases (below 70s) of the proposed framework enables regulatory compliance with these oversight requirements, but institutions with operations in more than one jurisdiction must review jurisdiction-specific requirements (Rizk et al., 2020).

6. Emerging Trends and Future Evolution

6.1 Integration with Large Language Models

New large language model technologies have a potential to greatly improve the ability of corporations to identify the actions of the notice they pay. Modern large language models that use GPT-3/GPT-4 class systems have shown excellent semantic knowledge of financial documents and can be able to extract structured data out of unstructured notices with high accuracy that is similar or at times better than a domain specific machine learning system. The framework evolution in the future should consider implementing the integration of the LLM-based extraction and interpretation layers instead of using the traditional OCR and classification pipeline (Santos et al., 2019).

Table 5: Comparative Benchmarking Against Industry Standards

Performance Metric	Industry Manual Average (2020)	Proposed HITL Framework	Improvement Factor
Corporate Actions Processing Cost per Event	$15.70	$3.85	4.07×
Average Processing Time per Notice	980 minutes	178 minutes	5.51×
First-Pass Accuracy Rate	78.2%	87.3%	1.12×
Settlement Failure Rate	10.2%	0.9%	11.3×

Performance Metric	Industry Manual Average (2020)	Proposed HITL Framework	Improvement Factor
Manual Review Rate	100%	15.5%	6.45× (reduction)
Compliance Exception Rate	5.8%	0.4%	14.5× (reduction)
Employee Hours per 1,000 Events	327	52	6.29× (reduction)
System Uptime Achievement	94.1%	99.7%	1.06×

6.2 **Regulatory Reporting and Real-Time Settlement**

The sector movement towards T+1 and T+0 settlement cycles generates time bottlenecks of ever greater severity on the post-trade processing. The result in the proposed framework of 178 minutes total time to process every transaction gives sufficient headroom to implementation of T+1, but inadequate margin to implement T+0. There will need to be further optimization in the future, possibly including parallel processing pipelines, asynchronous processing architectures, and anticipatory pre-processing of notices prior to formal reception (Santos et al., 2019).

6.3 **Cross-Border Settlement Integration**

The global financial markets have many regional settlement infrastructures such as Euroclear, Clearstream, DTCC and regional operators among others each with its own corporate actions processing protocols, and data standards. The harmonization of these disjointed infrastructures to a genuine cross-border automation needs to be the evolution of the framework, which today is a multi-year project in the form of industry working groups and standards bodies.

7. Conclusion

This study develops an in-depth confidence-based human-in-the-loop RPA model that is specifically developed to operate in the corporate actions notice interpretation of post-trade settlements operations. The structure combats vital challenges in the industry by integrating machine learning confidence scoring with focused human intervention to make sure that the automation rates (84.5% with optimized thresholds) are high as well as the accuracy levels (87.3% overall accuracy). Empirical study shows that the framework proposed results in 4.07-fold cost improvement with the industry manual processing baselines, 5.51-fold improvement in processing time, and 11.3-fold improvement in settlement failure rates. The framework offers a scalable automation that can be applied to institutions handling millions of corporate actions announcement per year, without violating regulatory requirements and procedural requirements of fairness (Syed et al., 2020).

Implementation requires careful attention to threshold optimization, OCR confidence calibration, and data standardization initiatives. However, the substantial performance improvements and competitive advantages offered by the framework justify investment in implementation infrastructure and change management processes. As global capital markets transition toward accelerated settlement cycles, confidence-driven HITL approaches represent critical enabling infrastructure for maintaining operational efficiency while preserving the accuracy and compliance standards increasingly demanded by regulatory authorities and market participants (Syed et al., 2020).

References

1. Aguirre, S., & Rodriguez, A. (2017). Automation of a business process using robotic process automation (RPA): A case study. *Communications in Computer and Information Science*, *748*, 65–71. https://doi.org/10.1007/978-3-319-66963-2_7
2. Anagnoste, S. (2017). Robotic automation process—The next major revolution in terms of back office operations improvement. *Proceedings of the International Conference on Business Excellence*, *11*(1), 67–77. https://doi.org/10.1515/picbe-2017-0072
3. Chakraborti, T., Isahagian, V., Khalaf, R., & Muthusamy, V. (2020). From robotic process automation to intelligent process automation: Emerging trends. *Business Process Management Workshops*, *393*, 180–187. https://doi.org/10.1007/978-3-030-58779-6_15
4. Dalsaniya, A. (2020). Human-in-the-loop robotic process automation (RPA) with deep reinforcement learning (DRL). *World Journal of Advanced Engineering Technology and Sciences*, *1*(1), 60–76. https://doi.org/10.30574/wjaets.2020.1.1.0015
5. European Central Bank. (2019). *Potential use cases for innovative technologies in securities post-trading*. Advisory Group on Market Infrastructures for Securities and Collateral. https://www.ecb.europa.eu/paym/groups/shared/docs/15acd-2019-01-ami-seco-potential-use-cases-for-innovative-technologies-in-securities-post-trading.pdf
6. Fernandez, D., & Aman, A. (2018). Impacts of robotic process automation on global accounting services. *Asian Journal of Accounting and Governance*, *9*, 123–132. https://doi.org/10.17576/AJAG-2018-10-09
7. Gotthardt, M., Koivulaakso, D., Paksoy, B., Saramo, C., Martikainen, M., & Lehner, O. (2020). Current state and challenges in the implementation of smart robotic process automation in accounting and auditing. *ACRN Journal of Finance and Risk Perspectives*, *9*, 90–102. https://doi.org/10.35944/jofrp.2020.9.1.006
8. Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. *Proceedings of the 34th International Conference on Machine Learning*, *70*, 1321–1330. https://proceedings.mlr.press/v70/guo17a.html
9. Hegde, S., Gopalakrishnan, S., & Wade, M. (2018). Robotics in securities operations. *Journal of Securities Operations & Custody*, *10*(1), 29–37. https://doi.org/10.69554/DLKK1858
10. Holzinger, A. (2016). Interactive machine learning for health informatics: When do we need the human-in-the-loop? *Brain Informatics*, *3*(2), 119–131. https://doi.org/10.1007/s40708-016-0042-6
11. Leopold, H., van der Aa, H., & Reijers, H. A. (2018). Identifying candidate tasks for robotic process automation in textual process descriptions. *Enterprise, Business-Process and Information Systems Modeling*, *314*, 67–81. https://doi.org/10.1007/978-3-319-91563-0_5
12. Muthusamy, V., Unuvar, M., Völzer, H., & Weisz, J. D. (2020). Do's and don'ts for human and digital worker integration. *arXiv*. https://doi.org/10.48550/arXiv.2010.07738
13. Rizk, Y., Isahagian, V., Boag, S., Khazaeni, Y., Unuvar, M., Muthusamy, V., & Khalaf, R. (2020). A conversational digital assistant for intelligent process automation. *arXiv*. https://doi.org/10.48550/arXiv.2007.13256
14. Santos, F., Pereira, R., & Vasconcelos, J. B. (2019). Toward robotic process automation implementation: An end-to-end perspective. *Business Process Management Journal*, *26*(2), 405–420. https://doi.org/10.1108/BPMJ-12-2018-0380
15. Syed, R., Suriadi, S., Adams, M., Bandara, W., Leemans, S. J., Ouyang, C., ter Hofstede, A. H., van de Weerd, I., Wynn, M. T., & Reijers, H. A. (2020). Robotic process automation: Contemporary themes and challenges. *Computers in Industry*, *115*, Article 103162. https://doi.org/10.1016/j.compind.2019.103162