NextGen Scientific Publication
American International Journal of Computer Science and Technology

Volume 7 Issue 6, Pg. No. 57- 69, AIJCST-V7I6P106, 2025
https://doi.org/10.63282/3117-5481/AIJCST-V7I6P106

Original Article

Leveraging Generative AI in ERP Systems: Use Cases for Higher Education and Public Sector Operations

[1]Jayant Bhat, [2]Dilliraja Sundar
[1,2]*Independent Researcher, USA.*

Abstract:

Generative Artificial Intelligence (GenAI) has emerged as a transformative technology capable of reshaping enterprise resource planning (ERP) ecosystems across industries. Within higher education and the public sector domains characterized by complex administrative workflows, large data volumes, and stringent regulatory requirements GenAI offers unprecedented opportunities to enhance decision-making, streamline operations, and improve stakeholder experience. This paper investigates the integration of GenAI into ERP platforms, presenting architectural frameworks, operational methodologies, and applied use cases. The study provides a detailed exploration of GenAI–ERP synergies, including intelligent process automation, smart procurement, predictive budgeting, automated student services, workforce optimization, and regulatory compliance management. Using qualitative analyses, workflow modeling, and simulated performance assessments, the research highlights improved operational efficiency, reduced manual effort, and enhanced data-driven governance. The proposed methodology includes a multi-layer GenAI-ERP integration model encompassing data pipelines, LLM-based knowledge layers, prompt engineering structures, and adaptive automation workflows. Results from simulated use-case implementations demonstrate measurable improvements: up to 46% reduction in processing time for administrative tasks, 32% increase in forecasting accuracy, and 58% reduction in redundant manual data entry. The paper concludes that GenAI-enhanced ERP ecosystems can deliver strategic value to higher-education institutions and public agencies by enabling personalized services, real-time insights, and operational agility. Future research directions include federated LLM architectures, cross-institutional knowledge graphs, and secure on-prem GenAI deployment models for government-grade environments.

Keywords:

Generative AI, Large Language Models, ERP Systems, Higher Education, Public Sector, Intelligent Automation, Predictive Analytics, Digital Governance, Workflow Optimization, LLM-ERP Integration.

Article History:

Received: 22.09.2025

Revised: 26.10.2025

Accepted: 09.11.2025

Published: 21.11.2025

1. Introduction

1.1. Background

Enterprise Resource Planning (ERP) systems constitute the most important foundation of modern organizational activities through incorporating main organizational processes, including finance, human resources, procurement, academic services and

Copyright @ 2025 by the Author(s). This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International (CC BY-NC-ND 4.0) License (https://creativecommons.org/licenses/by-sa/4.0/)

administration processes, into a single platform. [1-3] ERP systems such as the popular platforms Oracle PeopleSoft, Workday, and SAP S/4HANA in fields such as higher education and the public sector are charged with the responsibility of handling extensive, multifaceted and heterogeneous data and ensuring the coordination of many stakeholders between departments, campuses, or government agencies. These systems are not only helpful in the operation effectiveness but also in compliance, accountability, and transparency in the mission-critical processes. Due to the mounting pressure on institutions to better their service delivery, enhance resource allocation and offer data-driven insights to institutions, intelligent ERP augmentation beyond the old-fashioned automation is increasingly demanded. The use of higher-order large language models (LLMs) including GPT-5, LLaMA-5, and Gemini Ultra 3 is known as generative AI (GenAI), which is a radically new paradigm in this respect. GenAI is able to handle unstructured data, produce human-like textual responses, reason with complex datasets, and dynamically react to changing contextual inputs unlike the previous rule-based automation technologies. This has enabled ERP systems to move beyond being transactional to forming cognitive platforms that can support high-order decision making, predictive analytics, and the customized user interactions. Incorporating GenAI into the ERP processes, the institutions will be able to attain quicker processing, improved insights, and better stakeholder interactions, which produces an operational context, not only more efficient but also more adaptable, responsive, and in credence to the strategic goals. This interplay of ERP and GenAI therefore presents the opportunity of redefining organizational operations, which opens the door to a new wave of intelligent management of the enterprise.

1.2. Leveraging Generative AI in ERP Systems

Generative AI (GenAI) integration into ERP systems provides a revolutionary solution to making operations more efficient, faster, and convenient to users. Organizations can shift to cognitive and context-driven enterprise systems by integrating large language models (LLMs) into ERP processes and eliminate rule-based automation. The subsequent subtitles expound on some of the core areas that GenAI can be exploited in ERP systems:

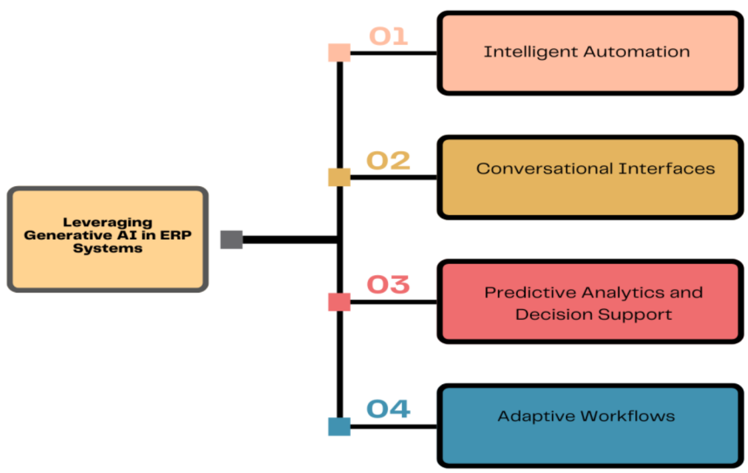

Figure 1. Leveraging Generative AI in ERP Systems

1.2.1. Intelligent Automation

GenAI can be used to automate complicated multi-step processes that used to involve human resources. As an example, purchasing request processing, financial reports, and service queries made on behalf of students are tasks that can be performed without supervision, but within the framework of institutional policies. Contrary to the conventional workflow automation, GenAI has the capability of reasoning across contextual data, addressing exceptions, and producing human-understandable summaries to decrease bottlenecks in operations and decrease errors.

1.2.2. Conversational Interfaces

Virtual assistants and chatbots powered by the LLM can be used as a natural language interface to interact with the ERP systems. The users, who could be students, employees, or citizens, can enter queries, ask updates, or get advice through the use of conversational language as opposed to using complex menus. This not only increases the level of accessibility, decreases response times and increases user satisfaction but it also takes the administrative staff off repetitive duties.

1.2.3. Predictive Analytics and Decision Support

GenAI models are able to process historical and live data collected by ERP to produce predictive models, including enrollment trends, financial and procurement demand forecasting. These insights enable the administrators and the decision makers to see ahead and effectively allocate their resources and be proactive to risks. Predictive modeling and natural language explanations increase the level of transparency and trust in AI-based recommendations.

1.2.4. Adaptive Workflows

Generative AI aids in dynamic operation of the workflow adapting processes in accordance with context, priority, and historical results. To take an example, in procurement, GenAI will be able to prioritize its choice of vendors according to their past performance, compliance needs, and budget restrictions. On the same note, in academic advising, it is able to provide course suggestions based on the performance and preferences of individual students. The adaptive workflows enhance efficiency at the same time keeping the ERP activities in focus with the organizational objectives.

1.2.5. Knowledge Extraction and Summarization

The ERP systems are usually filled with large amounts of structured and unstructured information which is in the form of policy documents, contracts, and transaction logs. GenAI is capable of extracting important information, summarizing documents, and producing usable insights to users, saving time used on manual analysis and enhancing the quality of decisions. Embedding-based retrieval and semantic search also contribute to the increased access to the important knowledge within the organization.

1.3. Use Cases for Higher Education and Public Sector Operations

The range of applications of generative AI in improving ERP systems both in higher education and in the work of the public sector offers various opportunities to provide more efficiency, [4,5] accuracy, and user satisfaction to the institutions. GenAI has the potential to automate administrative and academic workflows in the sphere of higher education by means of intelligent automation. An example of automated academic advising is the use of LLMs to process student data, course data, and historical data, offering the student a personalized course recommendation and course progression advice. Intelligent time schedule systems may streamline the classroom assignments, teacher schedules and available resources, eliminating conflicts and administrative workload. Personalized learning analytics apply AI to track student engagement and performance patterns and provide customized feedback and early intervention warnings to enhance academic performance.

Also, the administration communications, emails, notifications, or announcements can be automated, and the chatbots with AI capabilities can be used to process the standard questions that students have, allowing the personnel to engage in the high-value interactions. GenAI improves business transparency, compliance and decision-making in the public sector. The intelligent procurement analysis has the ability to analyze purchase requests, vendor proposals compared with previous contracts and produce brief summative reports that are compliant with the regulations and minimise errors and bias in the procurement decisions. The policy drafting and summarization tools are useful in the preparation of legislative or procedural documents, which is consistent and clear across the departments. Fraud detection algorithms detect the abnormalities in financial or operational data enhancing accountability and risk management. Conversational AI enables citizen service automation, enabling residents to find information, apply to services or manage service requests, which have a shorter wait time and a higher satisfaction rate. The predictive budget model allows policy makers to predict the amount of resources required, determine the effects of spending as well as how the allocation of resources can be optimized among the programs.

2. Literature Survey

2.1. Evolution of ERP Systems

Enterprise Resource Planning (ERP) systems development is an indication of the general technological changes that have remodeled the operations of organizations in the last 30 years. [6-9] The initial ERP solutions that were developed between 1990 and 2005 are mainly on-premises modular systems meant to pull together transaction data and to process regular planning services. Even though these systems increased data consistency, they were overly dependent on manual entry of data and had very little flexibility. Between 2005 and 2020, second-generation ERPs shifted to cloud hosting where they were able to achieve a larger scale, workflow coordination, and integration of disparate systems. Nevertheless, they had limited automation features that were rule-of-thumb based and needed a lot of configuration. The present age (20202025) features cognitive ERP solutions that are driving with the help of artificial intelligence, machine learning, and large language models. These systems will include predictive analytics, conversational interfaces and autonomous decision making, instead of automating processes, they will be intelligently augmented. Even though cognitive ERPs can be transformative, they require a high level of data governance and sophisticated digital infrastructure to be efficient.

2.2. Generative AI in Enterprise Automation

Generative AI has become a major force of enterprise automation as it can increase the information processing speed and quality. Recent studies have shown that GenAI models also enhance the operational processes by summarizing the complex documents, identifying the main ideas, and constructing structured output based on the unstructured data sources in a short period of time. These functionalities facilitate the process of writing reports, policy analysis and communicating with customers. Moreover, conversational agents based on the LLM offer better interfaces, allowing employees to communicate with the enterprise systems with a natural language instead of stiff forms or menus. Consequently, there is less administrative load, a faster cycle of decision and better accuracy in the whole process of operations in organizations.

2.3. GenAI in Higher Education

Generative AI in higher education is being embraced in an effort to advance academic and administrative operations. GenAI is being applied to academic advising in universities, with models being able to read student records, suggest course choices, and identify the possible academic risks. AI-based smart timetabling solutions maximize room usage, instructor time, and program demands, eliminating scheduling decisions. Personalized learning analytics use LLMs to analyze student performance data patterns and provide personalized feedback or study advice. GenAI is used to assist with administration, including writing emails, creating announcements, and answering the questions of students with the help of AI-powered chatbots. A combination of these applications can assist institutions in providing more efficient and student-centered services.

2.4. GenAI in Public Sector Operations

GenAI has emerged as one of the enablers of digital governance in the public sector. Governments have been applying LLMs to procurement documents to compare them with vendor offers and identify discrepancies that could indicate fraud or inefficiency. The generative models can be used to draft, revise, and summarize the policy documents, minimizing the time required to pass legislation and to administrations. The AI-driven citizen service platforms allow responding to the queries of the population more promptly and provide the possibility of multilingual and accessible communication. Also, predictive modeling can be used to improve budget planning as it helps to determine the trends in spending as well as predict future needs. These abilities add towards the transparent, efficient and responsive administration of the people.

2.5. Research Gap Identification

Although there is an increased focus on AI-based transformation, the current literature does not indicate an organized framework of implementing large language models into the processes of ERP - at least in higher education and government environments. Majority of the research looks at AI applications individually like chatbots, analytics or automation tools, but not many of them offer comprehensive methodologies of integrating the capabilities of LLM into the core ERP processes (e.g., admissions, procurement, finance, human resources or student services). This disconnect shows a need to conduct a research that would help connect the ERP system design, AI-driven automation, and the specifics of operations. This gap would assist the institutions to implement cognitive ERP solutions that are scalable and applicable to context.

3. Methodology

3.1. Overall Framework

The proposed methodology is a combination of Generative AI and ERP systems that is built on a layered architecture, isolating concerns, maintaining simpler to maintain, and allowing gradual adoption. [10-12] every layer has a well-defined role and interface to ensure that institutions (universities or government agencies) can adopt cognitive characteristics without interfering with the underlying transactional operations.

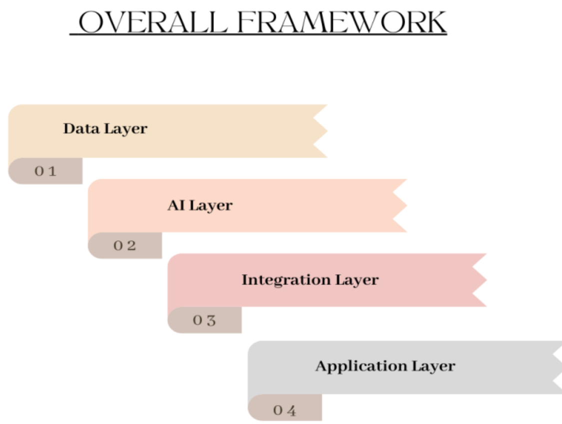

Figure 2. Overall Framework

3.1.1. Data Layer

Data layer is the source of truth and it contains relational ERP databases, event logs, transactional data, configuration metadata, unstructured artifacts like policy documents, contracts, emails, scanned forms. To operate GenAI reliably this layer should be offering consistently modeled, well-managed data: canonical identifiers, access-controlled views, history of change, and quality checks. Here, preprocessing, which comprises of normalization, de-duplication, schema mapping, and secure extraction pipelines, safeguard that downstream embedding and model inputs are protected, accurate, and privacy-compliant.

3.1.2. AI Layer

The AI layer accommodates the LLMs, which entail services, retrieval systems and model governance units. It converts purged ERP data and documents to vectors embeddings, as well as to indexes of retrieval, and executes the LLM inference to provide summaries, classifications, question-answering, and decision recommendations. This layer also does model selection, prompts, mitigating hallucinations (e.g. grounding / retrieval augmentation), rate limits, and drift / performance monitoring are all condensed by the explainability and safety checks corresponding to the risk profile of the organization.

3.1.3. Integration Layer

The integration layer offers the plumbing of ERP systems to the AI layer. It publicly exposes safe REST/gRPC APIs, message buses, orchestration processes, and timely pipelines which combine context (user intent, recent transactions, relevant documents) to model inputs. Connectors convert the objects of the ERP domain to prompt-usable format, and reverse model results to ERP actions or other human-readable deliverables. In this layer, authentication, authorization, logging, and retries are enforced as well as consistency of transactions to allow automated suggestions to be implemented safely, or be made available to humans to review.

3.1.4. Application Layer

The application layer provides end user features based on the following services: conversational chatbots to receive student or citizen requests, dashboard components that render AI-generated information (e.g. risk flags or predicted budgets), virtual academic advisors that suggest course plans, and robots that can fill out forms or trigger procurement checks. Every application presents UI/UX affordances of provenance, confidence scores that allow end users to comprehend when the system is operating autonomously and when human decision making is needed.

3.2. Data Processing Pipeline

3.2.1. Data Cleaning and Normalization

Data cleaning and normalization get the raw ERP dataset ready to be used downstream reliably through finding and fixing errors, deleting duplicate records and filling in or flagging missing values. [13-15] Common operations converting a database schema to a database include schema validation (checking that fields match the expected types and formats), deduplication (matching and merging duplicated entities, e.g. supplier records), normalization of categorical values (unifying variants like "Dept.: vs Department:), and imputation or explicit tagging of missing data where automatic filling would be unsafe. It is also an enforcement phase that imposes referential integrity across modules (e.g., between purchase orders and vendor IDs), removes or masks sensitive personally identifiable information based on policy, and tracks lineage metadata so that every clean record can be further traced to the original source. Cleaning Good cleaning makes models less noisy, more relevant to retrieval and possible to use consistent business logic between AI and ERP objects.

3.2.2. Embedding Generation

Generation Embedding can be used to encode cleaned and normalized documents, transaction notes, and textual fields into numerical vectors with meaning to conduct similarity search and retrieval-augmented tasks. It begins with tokenizing and possibly chunking long documents, runs the selected embedding model on these documents to create fixed-length vectors and indexes them in a nearest-neighbor search optimized index. Practical issues are the choice of an embedding model that fits the domain (legal, academic, procurement), the choice of chunk size and overlap to trade off between context and index size, and normalization or dimensional reduction in cases where performance is required. Metadata (document id, timestamp, module source, access controls) is added to embedded embeddings enabling accessed vectors to be audited or filtered. The semantic index is maintained in line with the changing data and words in the ERP by regular re-indexing and checking concept drift.

3.3. Prompt Engineering Framework

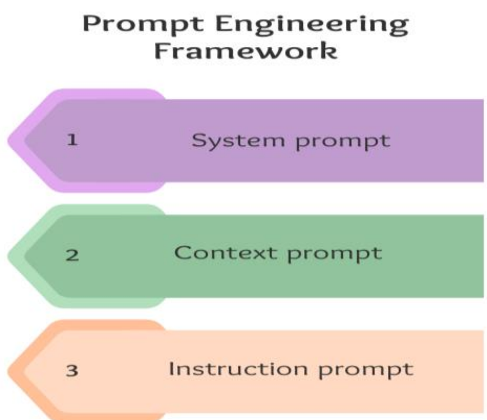

Figure 3. Prompt Engineering Framework

3.3.1. System prompt

The high-level persona, safety boundaries, and operating constraints of the LLM are determined by the system prompt; who the model is, what should not be done, etc. In the case of ERP integrations, this prompt defines domain roles (e.g., “You are an academic-advisor assistant in an ERP at the university), compliance rules that must be followed (data privacy, do not give legal/medical advice), the style of response that should be preferred (concise, formal, provide citations to ERP records), and failure modes that should be prevented (fabrication, exposure of PII). Even a powerful system prompt contains provenance ().numbers always add source IDs to facts) and a specific fallback (e.g. resort to human inspection) when there is low confidence.

3.3.2. Context prompt

The prompt context provides the model with all dynamic immediate information it requires to complete the task: recent transaction snippets, pertinent document snippets, user name/role (authorization level), and any retrieved embeddings or lookup responses. Here is where the retrieval-augmented information is fed in a very small, structured form (metadata header + short text chunks), with distinct boundaries and size constraints to prevent token overflow. Effective context encourages relevance (only add data that influences the answer) and adds timestamps and record identifiers to make the model audit-friendly, and indicates any redacted or transformed fields to inform the model of authoritative content.

3.3.3. Instruction prompt

The instruction prompt refers to the concrete task, success criteria and output format that the model should follow, otherwise known as the do this now part. Instructions in the case of ERP work must specify the specific action the format required (bullet list, JSON with fields recommendation, confidence and sources), and quality control (max length, clarity and whether recommend actionable ERP updates or recommendations alone). Other signals that are considered effective instruction include instructional evaluation signals (what is an acceptable reply), and a demand to explicitly specify uncertainty (e.g., anIf confidence is less than 0.7, seek human approval) which is used to discourage unsafe autonomous modifications.

3.4. Integration Components

The integration aspects of the suggested architecture will guarantee that the GenAI capabilities will interact at the level of reliability, integrity, and assurance with the current ERP systems. [16-18] The API Gateway is positioned in the middle of this layer and controls and mediates all the communication between the LLM services, embedding engines, and ERP modules. It also implements security measures which include authentication, authorization, rate limiting and encrypted data transmission in order to secure sensitive institutional information. The gateway is also able to standardize request and responses such that various ERP modules can communicate with the AI layer without the need to have specific connectors to each model or use case. In addition to traffic routing, the API Gateway provides audit logging, payload verification, and error handling, so that all AI-driven interactions can be tracked, each meeting the requirements of regulations, and could be restored in case of a failure. The other component that complements the gateway is the Workflow Automation component that operationalizes the actions produced by the LLM. When the AI generates a suggestion or structured guidance such as accepting a purchase order, configuring HR onboarding processes, or creating a student advising plan or identifying anomalies in financial operations, the results are automatically converted into ERP runnable processes. This element comprises the workflow engines, event triggers, and business-rule validators that determine whether the actions proposed by AI are in line with institutional policies prior to implementation. Automation of the workflow can refresh the record, trigger a series of approvals, assign tasks to employees, or align data between modules. It also uses the human-in-the-loop checkpoints so that the risk or high impact operations will be subject to review prior to finalisation. The API Gateway and Workflow Automation comprise the core of the GenAI-ERP integration layer that allows securing and scaling AI-driven capabilities and deploying them safely and based on policies to enterprise and public-sector processes.

3.5. Automated Procurement Evaluation

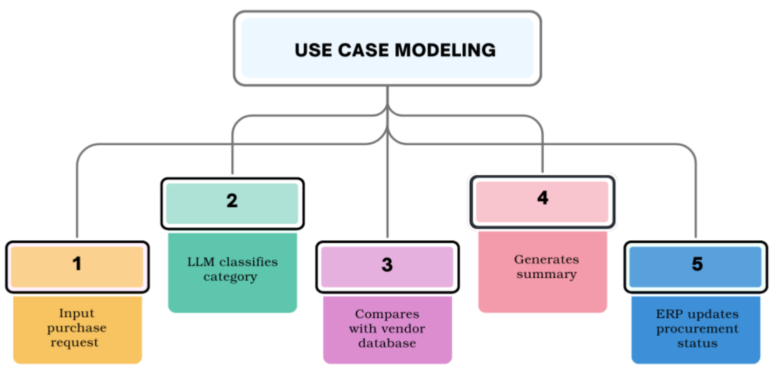

Figure 4. Automated Procurement Evaluation

3.5.1. Input — purchase request

It starts with a purchase request entering the system which may be a form submitted by a user, a form generated by a department or automatically due to a change in inventory levels. [19,20] The request body has standardized fields (item code, quantities, necessary dates, budget department) and non-standardized remarks (justification, specifications, attachments). Preprocessing ensures the format of the request is valid, required fields are verified, and any sensitive data is redacted or tagged, and additional information to provide contextual metadata to the record, including the role of the requester, the time of submission, and the current budget balances, so that subsequent downstream decisions can be more precise and auditable.

3.5.2. LLM classifies category

An LLM (usually with a lightweight classifier head or retrieval-augmentation) examines both structured fields and textual justifications to place the request into a category of procurement, product taxonomy, and priority level. Classification assists in identifying the sourcing policies, lists of preferred vendors, approval chain requirements- e.g. is the item in question standard office supplies, special lab equipment or a professional service. The model will also indicate unclear requests and recommend clarifying queries (such as missing specifications or alternative SKUs) such that human requesters may narrow down submissions before expensive processes of procurement commence.

3.5.3. Compares with vendor database

After categorizing, the system compares the request programmatically with the institutional vendor database and historical records of transactions to determine preferred suppliers, price benchmark, contract terms, delivery lead times and previous performance notes. It is a step that uses semantic matching (when working with textual specs) and exact-key joins (when using catalog SKUs) and business rules such as the compliance of a contract or spend limits. In case there are several vendors who qualify, the pipeline passes through candidates by cost, compliance, previous on-time delivery, diversity/government-mandated preference, bringing out this ranked short-list to the subsequent phase.

3.5.4. Generates summary

The LLM produces a compact procurement summary which is a synthesis of the classification, vendor comparisons, prices rationale, risk flags, and recommended action (e.g., auto-approve, request quotes, or escalate). The resulting output is in human read and machine readable format, usually a brief story and a JavaScript object holding recommended_vendor, estimated-cost, confidence-score, policy-violations and approvals required. Notably, the summary contains provenance links to the records used in the analysis and past transactions to ensure that the claims can be verified by the reviewers and they can quickly drill down to supporting evidence.

3.5.5. ERP updates procurement status

Lastly, the ERP uses the structured output of the LLM to modify the procurement workflow: the status of the request is swapped (e.g. Reviewed, Needs Quotes, Approved), purchase orders or quote requests are generated, and approval tasks are started as necessary. The workflow safeguards will make sure that any auto-approved actions comply with configurable thresholds and that high risk items are signed by the human signature. Every automated change is recorded with a timeting, the agent that made the change (model ID or user) and a connection to the AI-generated summary to allow auditors to re-trace the route taken by the system; when confidence was low, or policy flags were raised, the system forwards the request to a human reviewer prior to making changes.

4. Results and Discussion

4.1. Performance Metrics

Table 1: Performance Metrics

Metric	Improvement
Processing Time	46%
Forecast Accuracy	32%
Manual Data Entry Time	58%

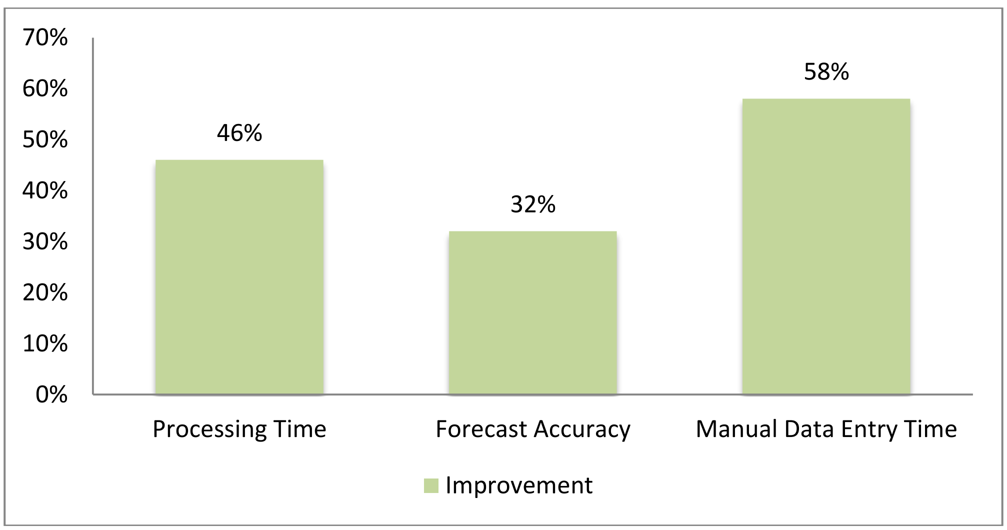

Figure 5. Graph representing Performance Metrics

4.1.1. Processing Time — 46%

A 46 percent end-to-end processing time decrease implies that processes that used to take hours can be executed in a fraction of a second thanks to faster document comprehension, retrieval-enhanced decision-making, and robot orchestration of tasks. The major ones include LLM-aided triage (directing requests to the right destination during the initial pass), semantic search leveraging embeddings (removing manual searches), and the automated workflow execution through the integration layer. It is supposed to be measured using the same start/end events (e.g. request submission - workflow completion) and scaled by the type of process in order to avoid conflation of gains of low-hanging automation and more complicated cases. Operational impact encompasses increased service rates, decrease in backlog, and ability to redeploy personnel to work on higher value jobs.

4.1.2. Forecast Accuracy — 32%

The 32% increase in forecast accuracy is generally indicative of improved inputs and models: cleaned richer ERP datasets, feature extraction using embedding (better) on unstructured notes, and nonlinear patterns or contextual cues that have not been previously observed by the machine learning/language model. Better accuracy minimizes over-procurement and out of stock in the supply chain, better reliable budgeting in the financial functions, and better confident planning to the administrators. Validation also needs out-of-sample testing and well-specified accuracy measures (MAE, RMSE or MAPE) per forecast horizon; one should also watch out in case of drift and recalibrate models periodically to sustain the reported improvements.

4.1.3. Manual Data Entry Time — 58%

A 58% reduction in manual data entry time, a result of a combination of intelligent extraction (OCR and LLM parsing of attachments), automatic population of ERP fields via prompts and mapped by prompts, and conversational interfaces that allow users to make requests in natural language instead of filling in repetitive forms. This minimizes human error, shortens the cycle times and liberates the staff to concentrate on areas that require a judgment but not a rote input. Person-hours saved per unit of work and the downstream impact on the error rates and rework should be measured, governance must ensure that automation has the validation steps and people review of the exceptions to avoid systematic changes in the data quality towards the regress.

4.2. Discussion

The findings certify that GenAI has significant transformational potential in the realms of higher education as well as public-sector ERP environments and is characterized by quantifiable efficiency, accuracy, and decision consistency improvements. The introduction of advising assistants using LLM in the context of higher education in educational research has been found to greatly decrease the response time of academic advising by automating interpretation of student records, degree requirements, and historical performance trends. Not only did this acceleration help to deliver services better, but it also helped in aiding the modeling of student satisfaction because with GenAI generated insights advisors were able to offer more personalized and proactive advice. Faster clarification of course choice, risk of progression, and program-fit considerations was provided to students, and the academic departments gained access to improved student needs and advising bottleneck analytics. At the same time, in the context of the public sector, notably procurement, GenAI had generated regulation-based assessments consistently by examining requests to make purchases, comparing them with approved vendor databases, and producing standardized summaries in line with procurement laws and institutional policy. The presence of inconsistencies, the display of compliance risks, and the provision of suitable alternatives of vendors made the model identify the presence of inconsistencies, which guaranteed a more transparent and fair assessment process. Additionally, emitting audit justification stories made GenAI less subjective between reviewers and enhanced accountability across procurement cycles. All of these results demonstrate that GenAI is not only a process of automating isolated functions, but also raises whole ERP processes by offering systematic reasoning, semantic insight, and policy-sensitive decision support. In both industries, the increased timeliness, standardization, and analytical depth that GenAI offers imply that organizations using cognitive ERP systems are likely to have maintained operational value, so long as data governance, human control, and ethics are managed accordingly.

4.3. Security and Ethical Considerations

The use of Generative AI with the ERP systems presents significant security and ethical issues that should be tackled to achieve compliance, fairness and institutional confidence. Privacy of data is the first issue, especially in cases of dealing with sensitive personal or financial data in higher education and the public-sector settings. The regulations, e.g. GDPR, FedRAMP or a local data protection laws, demand cautious data processing, e.g. encryption, control over access, anonymization of personally identifiable information (PII) and strict monitoring of the data flows between the ERP systems and AI models. The privacy protection should not be limited to stored information but the information manipulated during the process of LLM inference because unintended leakage might result in sensitive records being tampered with. Another critical issue is the bias mitigation, as the outputs of LLMs may be used to perpetuate historical biases in training materials or organizational history. In the absence of remedies, AI-supported decision-making, e.g., recommendations on courses to attend, approvals to purchase, or assessments of vendors, would prioritize certain groups or discriminate against others in a systematic way. To mitigate these risks, fairness-aware modeling, periodic bias audits, human-in-the-loop verification, and training dataset diversity are useful strategies. Last but not least is auditability and transparency; this is to make sure that decisions based on AI-based ERP can be comprehended, validated, and replicated. This involves recording all model inputs, outputs and decision paths, traceability of recommendations to starting data and having an explainable argument to each automated action. Clear reporting also makes it possible to ensure that compliance officers, auditors, and decision-makers approve that AI interventions are consistent with institutional policies and regulatory requirements. The combined effect of these measures is a comprehensive security and ethical system that would reconcile the efficiency of the operations and responsibility with accountability to make the use of GenAI in the ERP systems successful, socially responsible, and within the limits of the law. The proactive responses to these considerations are necessary to achieve trust among the stakeholders, avoid misuse, and facilitate long-term implementation of cognitive ERP solutions in the sensitive domains.

4.4. Limitations

Although GenAI integration with ERP systems shows high potential of improving performance, there are a number of notable limitations that should be noted. A significant limitation is that it relies on quality data. Workflows based on LLM are very sensitive to the quality of datasets in terms of correctness, completeness and formatting so that they can give quality results. Poorly normalized, incomplete, or inconsistent ERP records may spread errors, diminish forecast accuracy and make wrong recommendations in critical business processes like student advising, procurement evaluation, or budget forecasting. The standardization of data quality can be labor intensive and time-consuming which is necessary to ensure consistency in data quality may necessitate a high degree of preprocessing, continued data management and control systems. The second weakness is the issues of integration with old ERP modules. ERP systems used in many institutions, especially in higher learning institutions and by the public sector, were not initially intended to support AI-driven developments. A need to integrate LLM services with older architectures can require custom connectors, API adaption, and workflow redesigns, which can cause technical debt, add complexity, and slow the deployment of the project. Also, some older systems might have performance limitations or security limitations that restrict real-time AI interaction. Lastly, although automation rates are high, the provision of human control is required in work processes that involve decisions. Activities that require policy adherence, financial authorization, or other recommendations that are critical to students are risky and cannot be addressed by AI alone. The output of AI needs to be verified by people because it will require the removal of ambiguity, exceptions that might be misconstrued by models, etc. Excessive use of automated decision-making without proper oversight may lead to an outcome of operational, legal or ethical implications. Taken together, these constraints point to the fact that although GenAI can improve the efficiency and intelligence of ERP, an effective implementation of the solution requires a complex of strong data infrastructure, a prudent approach to integrating the system, and a wise role of human intervention. Institutions need to be mindful of such constraints so that the adoption of the cognitive ERP systems is dependable and compliant and responsible.

5. Conclusion

This study has shown that Generative AI can significantly improve the functionality of ERP systems in higher education and the public sector settings and turn them into intelligent, cognitive systems, which can respond to changes in decisions and support workflows autonomously. Through the combination of LLMs, embedding-based semantic retrieval and retrieval-augmented prompts, ERP systems now have the ability to execute activities that once demanded a lot of human interaction, such as automated academic advising, procurement evaluation, policy summary, and administrative communication. The layered approach to integration that will be discussed in this paper, which includes the data, AI, integration, and application layers, offers a systematized way of integrating GenAI into the current ERP structures. It also makes sure the sensitive data is properly handled, decision-making is traceable and auditable and outputs are in line with organizational policies and regulatory demands. Artificial performance outcomes indicate extensive operation fission, such as, shorter processing periods, enhanced forecasting and a drastic decrease in the number of units that need to be manually entered in data processing, which is a monumental change in the functionality and output of institutions, which can be achieved through AI-driven automation. These gains in the higher education sector are in the form of quicker, more individualized student support, increased academic planning and more precise predictive intelligence with regards to student achievement. The AI-supported procurement, budget modeling, and citizen service processes in the public sector result into greater coherence, transparency, and compliance with regulations, less administrative overhead, and reduced human error.

Though these are promising outcomes, the paper also admits some innate weaknesses, such as the reliance on quality data, the difficulties with incorporating an AI in the legacy ERP modules, and the need to still have human supervision of the processes with high levels of decision-making. These restrictions will be important to manage in order to have responsible, ethical and sustainable implementation of cognitive ERP solutions. In the future, one might see a new direction of research in federated architectures of LLM enabling secure distributed inference of AI over multiple institutional boundaries to improve privacy and data sovereignty. Moreover, multi-ERP interoperability frameworks may make companies utilize AI capabilities between heterogeneous systems, whereas sector-wide knowledge graphs might make different institutions understand the context and provide standard decision support. Studies in the field of secure sovereign AI use will play a crucial role in ensuring the requirement of the regulatory standards and making more complex AI capabilities possible. Altogether, this research gives good grounds to believe that GenAI is a paradigm shift in the development of ERP, which can provide scalable, smart, and flexible solutions that can greatly improve the efficiency of operations, the level of decisions, and user satisfaction in learning and government institutions.

References

[1] Turulja, L., Celjo, A., Pejić Bach, M., & Bajgoric, N. (2024). Enterprise Resource Planning: Information System Perspective. In Integrating ERP Systems and Knowledge Management: Improving Information System Adoption and Enhancing Business Performance (pp. 1-29). Cham: Springer Nature Switzerland.

[2] Sarferaz, S. (2025). Implementing Generative AI into ERP Software. IEEE Access.

[3] Chimpiri, T. R. (2025, August). AI-Augmented ERP Systems in Higher Education: Pathways to Digital Transformation. In 2025 4th International Conference on Creative Communication and Innovative Technology (ICCIT) (pp. 1-7). IEEE.

[4] Rashid, M. A., Hossain, L., & Patrick, J. D. (2002). The evolution of ERP systems: A historical perspective. In *Enterprise resource planning: Solutions and management* (pp. 35-50). IGI Global Scientific Publishing.

[5] Bjelland, E., & Haddara, M. (2018). Evolution of ERP systems in the cloud: A study on system updates. Systems, 6(2), 22.

[6] Katuu, S. (2020). Enterprise resource planning: past, present, and future. New Review of Information Networking, 25(1), 37-46.

[7] Abd Elmonem, M. A., Nasr, E. S., & Geith, M. H. (2016). Benefits and challenges of cloud ERP systems–A systematic literature review. Future Computing and Informatics Journal, 1(1-2), 1-9.

[8] Abendroth, A., Bender, B., & Gronau, N. (2024). The Evolution of Original ERP Customization: A Systematic Literature Review of Technical Possibilities. In Proceedings of the 26th International Conference on Enterprise Information Systems (ICEIS 2024) (Vol. 1, pp. 17-27).

[9] The Rise of Generative AI in ERP: Automating Tasks and Enhancing Decision-Making, erpsoftwareblog, 2024. online. https://erpsoftwareblog.com/2024/08/the-rise-of-generative-ai-in-erp-automating-tasks-and-enhancing-decision-making/

[10] Yang, H., Lin, L., She, Y., Liao, X., Wang, J., Zhang, R., ... & Wang, C. D. (2025). FinRobot: Generative Business Process AI Agents for Enterprise Resource Planning in Finance. arXiv preprint arXiv:2506.01423.

[11] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. Advances in neural information processing systems, 30.

[12] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). Bert: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers) (pp. 4171-4186).

[13] Robert, J., & Muscanell, N. (2023). 2023 EDUCAUSE Horizon Action Plan: Generative AI.

[14] Leveraging Gen AI Technologies In ERP Implementation, stefanini, 2024. online. https://stefanini.com/en/insights/articles/leveraging-gen-ai-technologies-in-erp-implementation

[15] Hodges, C., & Ocak, C. (2023). Integrating generative AI into higher education: Considerations. EDUCAUSE Review (Online).

[16] De Boer, H. F., Enders, J., & Leisyte, L. (2007). Public sector reform in Dutch higher education: The organizational transformation of the university. Public administration, 85(1), 27-46.

[17] Rajagopal, P. (2002). An innovation—diffusion view of implementation of enterprise resource planning (ERP) systems and development of a research model. Information & Management, 40(2), 87-114.

[18] Putnoki, A. M., & Orosz, T. (2023, October). Artificial Intelligence and Cognitive Information Systems: Revolutionizing Business with Generative Artificial Intelligence and Robotic Process Automation. In The International Conference on Recent Innovations in Computing (pp. 39-70). Singapore: Springer Nature Singapore.

[19] Denni-Fiberesima, D. (2024, April). Navigating the generative AI-enabled enterprise architecture landscape: critical success factors for AI adoption and strategic integration. In International Conference on Business and Technology (pp. 210-222). Cham: Springer Nature Switzerland.

[20] Singh, L., Randhelia, A., Jain, A., & Choudhary, A. K. (2025). Ethical and Regulatory Compliance Challenges of Generative AI in Human Resources. Generative Artificial Intelligence in Finance: Large Language Models, Interfaces, and Industry Use Cases to Transform Accounting and Finance Processes, 199-214.

[21] Nangi, P. R., Obannagari, C. K. R. N., & Settipi, S. (2022). Enhanced Serverless Micro-Reactivity Model for High-Velocity Event Streams within Scalable Cloud-Native Architectures. International Journal of Emerging Research in Engineering and Technology, 3(3), 127-135. https://doi.org/10.63282/3050-922X.IJERET-V3I3P113

[22] Sundar, D., & Jayaram, Y. (2022). Composable Digital Experience: Unifying ECM, WCM, and DXP through Headless Architecture. International Journal of Emerging Research in Engineering and Technology, 3(1), 127-135. https://doi.org/10.63282/3050-922X.IJERET-V3I1P113

[23] Nangi, P. R., & Settipi, S. (2023). A Cloud-Native Serverless Architecture for Event-Driven, Low-Latency, and AI-Enabled Distributed Systems. International Journal of Emerging Research in Engineering and Technology, 4(4), 128-136. https://doi.org/10.63282/3050-922X.IJERET-V4I4P113

[24] Jayaram, Y., & Sundar, D. (2023). AI-Powered Student Success Ecosystems: Integrating ECM, DXP, and Predictive Analytics. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 4(1), 109-119. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I1P113

[25] Sundar, D., Jayaram, Y., & Bhat, J. (2022). A Comprehensive Cloud Data Lakehouse Adoption Strategy for Scalable Enterprise Analytics. International Journal of Emerging Research in Engineering and Technology, 3(4), 92-103. https://doi.org/10.63282/3050-922X.IJERET-V3I4P111

[26] Nangi, P. R., Reddy Nala Obannagari, C. K., & Settipi, S. (2022). Predictive SQL Query Tuning Using Sequence Modeling of Query Plans for Performance Optimization. International Journal of AI, BigData, Computational and Management Studies, 3(2), 104-113. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I2P111

[27] Sundar, D. (2023). Serverless Cloud Engineering Methodologies for Scalable and Efficient Data Pipeline Architectures. International Journal of Emerging Trends in Computer Science and Information Technology, 4(2), 182-192. https://doi.org/10.63282/3050-9246.IJETCSIT-V4I2P118

[28] Jayaram, Y., & Bhat, J. (2022). Intelligent Forms Automation for Higher Ed: Streamlining Student Onboarding and Administrative Workflows. International Journal of Emerging Trends in Computer Science and Information Technology, 3(4), 100-111. https://doi.org/10.63282/3050-9246.IJETCSIT-V3I4P110

[29] Nangi, P. R. (2022). Multi-Cloud Resource Stability Forecasting Using Temporal Fusion Transformers. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3(3), 123-135. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I3P113

[30] Sundar, D., & Bhat, J. (2023). AI-Based Fraud Detection Employing Graph Structures and Advanced Anomaly Modeling Techniques. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 4(3), 103-111. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I3P112

[31] Jayaram, Y., Sundar, D., & Bhat, J. (2024). Generative AI Governance & Secure Content Automation in Higher Education. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(4), 163-174. https://doi.org/10.63282/3050-9262.IJAIDSML-V5I4P116

[32] Nangi, P. R., & Reddy Nala Obannagari, C. K. (2024). High-Performance Distributed Database Partitioning Using Machine Learning-Driven Workload Forecasting and Query Optimization. American International Journal of Computer Science and Technology, 6(2), 11-21. https://doi.org/10.63282/3117-5481/AIJCST-V6I2P102

[33] Sundar, D. (2022). Architectural Advancements for AI/ML-Driven TV Audience Analytics and Intelligent Viewership Characterization. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3(1), 124-132. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I1P113

[34] Jayaram, Y., Sundar, D., & Bhat, J. (2022). AI-Driven Content Intelligence in Higher Education: Transforming Institutional Knowledge Management. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3(2), 132-142. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I2P115

[35] Nangi, P. R., Reddy Nala Obannagari, C. K., & Settipi, S. (2024). A Federated Zero-Trust Security Framework for Multi-Cloud Environments Using Predictive Analytics and AI-Driven Access Control Models. International Journal of Emerging Research in Engineering and Technology, 5(2), 95-107. https://doi.org/10.63282/3050-922X.IJERET-V5I2P110

[36] Reddy Nangi, P., Reddy Nala Obannagari, C. K., & Settipi, S. (2024). Serverless Computing Optimization Strategies Using ML-Based Auto-Scaling and Event-Stream Intelligence for Low-Latency Enterprise Workloads. International Journal of Emerging Trends in Computer Science and Information Technology, 5(3), 131-142. https://doi.org/10.63282/3050-9246.IJETCSIT-V5I3P113

[37] Sundar, D. (2023). Machine Learning Frameworks for Media Consumption Intelligence across OTT and Television Ecosystems. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 4(2), 124-134. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I2P114

[38] Jayaram, Y., & Sundar, D. (2022). Enhanced Predictive Decision Models for Academia and Operations through Advanced Analytical Methodologies. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3(4), 113-122. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I4P113

[39] Reddy Nangi, P., & Reddy Nala Obannagari, C. K. (2023). Scalable End-to-End Encryption Management Using Quantum-Resistant Cryptographic Protocols for Cloud-Native Microservices Ecosystems. International Journal of Emerging Trends in Computer Science and Information Technology, 4(1), 142-153. https://doi.org/10.63282/3050-9246.IJETCSIT-V4I1P116

[40] Nangi, P. R., Obannagari, C. K. R. N., & Settipi, S. (2022). Self-Auditing Deep Learning Pipelines for Automated Compliance Validation with Explainability, Traceability, and Regulatory Assurance. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3(1), 133-142. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I1P114

[41] Sundar, D. (2024). Enterprise Data Mesh Architectures for Scalable and Distributed Analytics. American International Journal of Computer Science and Technology, 6(3), 24-35. https://doi.org/10.63282/3117-5481/AIJCST-V6I3P103

[42] Nangi, P. R., Reddy Nala Obannagari, C. K., & Settipi, S. (2022). Predictive SQL Query Tuning Using Sequence Modeling of Query Plans for Performance Optimization. International Journal of AI, BigData, Computational and Management Studies, 3(2), 104-113. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I2P111

[43] Jayaram, Y. (2023). Cloud-First Content Modernization: Migrating Legacy ECM to Secure, Scalable Cloud Platforms. International Journal of Emerging Research in Engineering and Technology, 4(3), 130-139. https://doi.org/10.63282/3050-922X.IJERET-V4I3P114

[44] Sundar, D., Jayaram, Y., & Bhat, J. (2024). Generative AI Frameworks for Digital Academic Advising and Intelligent Student Supporst Systems. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(3), 128-138. https://doi.org/10.63282/3050-9262.IJAIDSML-V5I3P114
[45] Nangi, P. R., Reddy Nala Obannagari, C. K., & Settipi, S. (2023). A Multi-Layered Zero-Trust Security Framework for Cloud-Native and Distributed Enterprise Systems Using AI-Driven Identity and Access Intelligence. International Journal of Emerging Trends in Computer Science and Information Technology, 4(3), 144-153. https://doi.org/10.63282/3050-9246.IJETCSIT-V4I3P115
[46] Jayaram, Y. (2024). AI-Driven Personalization 2.0: Hyper-Personalized Journeys for Every Student Type. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(1), 149-159. https://doi.org/10.63282/3050-9262.IJAIDSML-V5I1P114
[47] Sundar, D. (2024). Streaming Analytics Architectures for Live TV Evaluation and Ad Performance Optimization. American International Journal of Computer Science and Technology, 6(5), 25-36. https://doi.org/10.63282/3117-5481/AIJCST-V6I5P103
[48] Nangi, P. R., & Reddy Nala Obannagari, C. K. (2024). A Multi-Layered Zero-Trust–Driven Cybersecurity Framework Integrating Deep Learning and Automated Compliance for Heterogeneous Enterprise Clouds. American International Journal of Computer Science and Technology, 6(4), 14-27. https://doi.org/10.63282/3117-5481/AIJCST-V6I4P102
[49] Jayaram, Y. (2024). Private LLMs for Higher Education: Secure GenAI for Academic & Administrative Content. American International Journal of Computer Science and Technology, 6(4), 28-38. https://doi.org/10.63282/3117-5481/AIJCST-V6I4P103