2025 International Conference on Next Generation Computing Systems (ICNGCS) | 979-8-3315-9116-8/25/$31.00 ©2025 IEEE | DOI: 10.1109/ICNGCS64900.2025.11182987

Invoice Indexing and Business Expense Management using Generative AI

Devadarshini M, Karthikeyan A S, Manimegalai R
Department of Computer Science and Engineering
PSG Institute of Technology and Applied Research
Coimbatore, India
devadarshini10@gmail.com, askarthikeyan2003@gmail.com, drrm@psgitech.ac.in

***Abstract*—: Enterprise Resource Planning (ERP) tools are essential for both large and small, but growing organizations. ERP is basically business management information systems that support the day-to-day operations of an organization from end-to-end. Spend management or expense management is part of an ERP toolset, and can help companies reduce costs, streamline processes by providing all the necessary tools to control expenses, including travel and reimbursements. In addition to supporting the entire expense process, from receipt capture to card controls, these solutions also integrate with other systems, such as ERP accounting software. The in class integrations allow the data to flow directly from the expense management software into the ERP without data entry or multiple imports and exports. In this work, an application is developed to classify the invoices submitted by employees in an organization. The proposed system in this work uses a Small Language Model (SLM), making it ele-gant and easy the for employees to submit their expense reports. With the elimination of shifting between applications, this system therefore makes the process quicker and saves more time and effort in processing. A bot is developed as part of this work, to assist the user when navigating through submitting details such as the type-of-expense and the receipt images. In this regard, the SLM automatically categorizes expenses by separating travel, accommodation, and meal types. This intelligent categorization minimizes errors on the part of the system user and saves the burden of manual classification by the users themselves. Moreover, it is equipped with Optical Character Recognition (OCR), where the uploaded receipts are scanned to extract text from there. Notifications are sent, to speed up the approval, allowing managers to review and approve submissions quickly without unnecessary delays. In nutshell, the proposed system streamlines expense reporting by automating categorization and speeding up the approval workflow. This work minimizes manual data entry, eliminates errors, and increases the productivity of staff and management, allowing them to have time and space for high value work.**

***Index Terms*—Expense Management, Small Language Models (SLMs), Expense Categorization, Optical Character Recognition (OCR), Automation, Workflow Optimization**

I. Introduction

Reimbursement process in business is crucial to ensure that the business expenses are settled to employees and there is a proper record of business transactions done by employees during business trips. Initially, reimbursements were handled manually. Employees have to submit forms and attach the invoices as proof of payment to get the reimbursements. This is a cumbersome process often taking a lot of time and is prone to errors. This method also lacks transparency as the receipts can be easily manipulated. Approvals may be delayed as the documents have to be physically routed upwards in the hierarchy and there is also a chance of the receipts getting misplaced during the process.

In the 1990s, spreadsheet tools like Microsoft Excel were introduced. Employees and managers could now track their expenses digitally and email the expense data which made the claim submission process easier. But however this process is not fully automated and requires manual intervention for verification of documents and final approvals. Spreadsheet systems had limitations in scalability, security and auditability.

In the early 2000s, ERP tools like SAP, Oracle and Peo-pleSoft were introduced. These tools acted as a medium for centralised reporting of expenses and are often integrated with the accounting systems of the organization. They can also be used for policy enforcement and can have a predefined set of rules based on which approvals are made. The major drawback of ERP tools is that it's complex to use and it's expensive.

In the 2010s, the introduction of cloud based expense management platforms enabled mobile submissions, receipt scanning, and use of OCR technology to extract the invoice content. These ensured real time policy checks and faster approval process. Finance teams gained better visibility through dashboards and analytics, while employees experienced faster turnaround times and simpler workflows.

With the evolution of language models, the expense approval process can be fully automated. Text extracted from invoices via Optical Character Recognition (OCR) can be given as input to the language model which can make decisions and trigger agentic AI workflows . These agentic AI workflows can be seamlessly integrated with enterprise communication tools such as Slack and Microsoft Teams. This integration enables employees to submit expense claims through the tools used in day to day life, eliminating the need to use a dedicated expense management software.

The current cost management processes in place with most organizations presents several inefficiencies and challenges to employees and managers alike. It is as easy as using a couple of applications, all you need to do is capture the expense report, track their expenses and find what they actually expend those on spending most of it with own hands punching on mobile keyboard and chasing down approvers. Thus, its turnaround time is very long and the figures are inaccurate.

These inefficiencies lower the productivity of an organization, as the staff must spend more time on non-core tasks related to administration.

This tackles the issues by presenting an automated expense management system driven by a Small Language Model that has been hyperparameter-tuned. It looks forward to making the process of submitting an expense report less time-consuming and feasible, so that users may submit one without having to switch between applications too often. As a result, there is less chance of error and the expenditure submission process is completed more quickly.

Therefore, an expense bot that guides a user through the process of entering their expenses forms the core of this system. Important details like the type of expense, the amount paid, the date, and the ability to submit receipt photographs can all be entered. After the image is uploaded, OCR scans it to extract pertinent text and information. The SLM then analyzes this data to determine which classification actually exists. The SLM automates categorization processes, but it smartly classes up expenses appropriately into categories which include, for example, travel, lodging, and meals. To this extent, automation of categorization reduces the need for hand operations like data entry, consequently lowering the probability of human error by great percentages. Apart from the categorization process, the system collects living cost data that further improves on the correct classification as regards to expenses. Upon the submission of the expense report, the system automatically notifies the approver, reducing a lot of the hassle of manual routing. This eliminates the delays and administrative overhead. Employees obtain their reimbursements fast when the system sends the approvers notifications for approval and processing of expenses as well as routing.

The main advantages of the project are that it automates expense reporting easy, reduces the hours the team needs to input and re-input the same information thus leading to the efficiency of businesses as a whole. First, the system offers an integrated platform doing away with the need to go through multiple programs to file expenditure reports thus cutting time and frustration. Secondly, the expenditure bot and SLM classify the expenses automatically, limiting the quantity of manual input and lowering errors. Third, the automatic notification system helps in shortening the processing period since the report will be checked and approved immediately, thus reducing the time between reimbursement. Finally, the system provides efficiency benefits through the reduction of manual computation, which is one of the core errors contributing to inefficiencies, hence creating a more effective flow for the employees and managers.

Technical scope for this project will be the development of a backend using Python and Flask. The interaction between the user interface and the SLM has been controlled by the backend that is based on a fine-tuned T5 model optimized for expense categorization. The training data consists of Indian source-picked documents, including invoices and receipts. This system allows the focus on it to be capable enough for processing common types of expenses in that context, and the OCR engine is also trained to give precise text extraction. However, the system has its drawbacks. It has been developed to cater for expenses reports coming straight from within India, which, naturally enough will impact upon the processing of other regions' documents. There are differences in format and/or language. Thus, this automatic expense management system, with its limitations, is much more advanced than traditional processes. It facilitates faster, accurate, and simpler execution, thereby saving time as compared to the preparation of an expense report, reduces errors, and lifts the burden and sometimes stress from employees and managers in this respect.

II. Literature Survey

An AI based platform engineering approach for software development business is discussed in [1].The platform proposed in [1] uses a developer portal with integrated CI/CD pipelines to automate financial workflows and facilitate collaborative development and allows developer to focus on high quality code with minimum implementation overhead.

Manimegalai et.al [2] presented a comprehensive review that traces the evolution of software systems. The review focus on the importance of architectural progression starting from single-tier architecture to multi-tier architectures that incorporates cloud and distributed computing.The study also highlights the impact of CI/CD pipelines impact in streamlining the deployment process.

Energy efficient coding practices in cloud computing is discussed in [3].The study focuses on the impact of programming practices on power efficiency in cloud based environments. The study highlights the role of software level optimizations o achieve green computing goals.

Kushavaha, Kishan[4] says Optical Character Recognition According to him, it means robotic turn for the textbook information coming out from images translates the published or handwritten documents into something accessible on digital formats that may turn easier to store and process. The quality friction of characters and complexity in handwriting make it difficult for OCR, yet it is an essential part of data entry robotization, document digitization, and hundreds of diligence and operations.

Joshi and Pvt[5] mentioned Tesseract OCR, which referred to open- source textbook recognition machine for more than 100 languages. Out of these, Hindi and Bengali are significant for India. It can digitize published as well as handwritten documents. Compared to AI, it beautifully upgrades the data birth robotization process apart from document digitization across different operations. Exploring book summarization in natural language processing using convolutional T5 and Seq2Seq models.

Lubis et al.[6] discovered that several other variables other than case perceptibility affect the textbook appreciation. With these values, they could fine-tune the Cream criteria by readjusting the subcaste confines, learning rate, and Dropout parameters to reach a best Cream criteria value of 0.8 for Cream- 1, 0.83 for Cream- 2, and 0.8 for Cream- L using the same model.

TABLE I
CLASSIFICATION METRICS

	Precision	Recall	F1-score
miscellaneous	0.990	1.000	0.995
restaurant	0.994	0.983	0.989
stay	0.990	0.985	0.987
travel	0.982	0.986	0.984

Stealthily embedded at the system's technical heart is a two-stage pipeline of Optical Character Recognition (OCR) coupled with an optimized T5 Small Language Model (SLM).

Tesseract OCR is tasked with extracting text data from uploaded receipts, while the fine-tuned T5 model astutely classifies expenses into four categories: stay, travel, restaurant, and miscellaneous. The workflow of the process, as depicted in the system architecture, highlights how receipts are uploaded, processed, classified, and directed for approval, with results being stored securely in MongoDB. This end-to-end automation proves not only the technical viability but also the practical feasibility of such a system for industrial use cases.

The system's efficiency is emphasized through the experiment results. The confusion matrix in Figure 5 neatly shows that the fine-tuned T5 model works remarkably well in expense categorization, where the overwhelming majority of receipts are accurately labeled. Few misclassifications are noticed, including a handful of stay receipts being output as miscellaneous and one travel receipt being labeled as restaurant, which solidifies the generalization strength of the model even when trained on real-world data. The classification metrics also support the same, with precision, recall, and F1-scores all being above 0.98 for all classes. The miscellaneous class had a perfect recall of 1.0, so no receipts under this class were missed. The restaurant class exhibited the highest precision of 0.994 and indicates the performance of the system in terms of not missing false positives in that category. Stay and travel also had balanced and uniformly high F1-scores of greater than 0.98, affirming that the model is trustworthy for various classes.

These findings emphasize the importance of combining OCR and sophisticated language models in business processes. In addition to providing high accuracy in classification, the system makes meaningful contributions towards organizational productivity. By lowering the involvement of human intervention in categorization, the solution reduces errors on the part of humans that are prevalent in expense management. The real-time notification using Slack speeds up the approval process, while database integration ensures scalability and secure storage of high numbers of receipts. Combined, these capabilities ensure accuracy and usability, which makes the system feasible for deployment at scale.

Aside from its present implementation, the project leaves room for future scaling. Even though the system is developed first for the Indian market, its architecture is versatile enough to enable worldwide deployment. In addition to training on multilingual datasets, the system could be made to handle invoices in other languages and scripts. Incorporating cost-of-living information and currency conversion modules would enable companies with international operations to control costs better. The classification model itself might be extended to support more categories like office supplies, utilities, and subscription services to make it more useful. Secondly, embedding explainability functionalities within the classification pipeline would make the finance managers more transparent to the extent of knowing why each categorization is being made, which would make organizational trust in AI-based decision-making more robust.

In summary, the project creates a robust and scalable expense management system that integrates cutting-edge AI features, cloud storage for data, and business communication platforms. It demonstrates how merging OCR and language models fine-tuned can automate what has historically been a mind-numbing and error-prone process into an intelligent, streamlined, and highly accurate workflow. By minimizing administrative overhead, maximizing productivity, and guaranteeing dependable classification performance, the system plays a double role in contributing to employee convenience as well as organizational efficiency. With its current success and prospects of global applicability, this solution provides a solid platform for the next generation of intelligent cost management systems that can revolutionize financial processes for businesses globally.

Though the initial deployment targets the Indian market. Subsequent releases may expand the geographic coverage by integrating extra data feeds for receipts, invoices, and cost-of-living data in other regions. Further, the precision of the SLM's expense categorization may be refined even more by enlarging the training dataset to cover a greater variety of expense categories and international contexts. This project provides a good foundation for creating an effectively strong and world-relevant cost management system.

REFERENCES

[1] Srinivasan, V., Rajkumar, Manimegalai., Santhanam, S., & Garg, A. (2025). PlatFab: A Platform Engineering Approach to Improve Developer Productivity. Journal of Information Systems Engineering and Business Intelligence, 11(1), 79–90. https://doi.org/10.20473/jisebi.11.1.79-90

[2] Manimegalai R, Vishnu Charan, Venkateshwaran M, "Evolution of Software Systems from Incubation to Enterprise Solutions", In Proceedings of the First International Conference on Science, Engineering and Technology Practices for Sustainable Development, 2024, EAI, DoI: 10.4108/eai.17-11-2023.2342836

[3] R Manimegalai, Srivatsan Sandhanam, A Sunitha Nandhini, Priyam Pandia, "Energy Efficient Coding Practices for Sustainable Software Development", In Proceedings of the First International Conference on Science, Engineering and Technology Practices for Sustainable Development, 2024, EAI, DOI: 10.4108/eai.17-11-2023.2342635.

[4] Kushavaha, Kishan. (2024). OCR: Optical Character Recognition. International Journal for Research in Applied Science and Engineering Technology. 12. 3854-3856. 10.22214/ijraset.2024.60794.

[5] Joshi, Kartik & Pvt, Techsamvaad & Ltd,. (2021). Study of Tesseract OCR.

[6] Lubis, Arif & Safitri, Habibi & Irvan, Irvan & Lubis, Muharman & Al-Khowarizmi, A.. (2023). Improving Text Summarization Quality by Combining T5-Based Models and Convolutional Seq2Seq Models. Journal of Applied Engineering and Technological Science (JAETS). 5. 10.37385/jaets.v5i1.2503.

[7] O.PSharma, M.K Ghose, K.B Shah and B.K Thakur, “Recent Trends and Tools for Feature Extraction in OCR Technology.” International Journal of Soft Computing and Engineering (IJSCE), 2013, ISSN: 2231-2307, Volume-2, Issue-6.

[8] C. Patel, A. Patel and D. Patel,“Optical Character Recognition by Open Source OCR Tool Tesseract: A Case Study.” International Journal of Computer Applications (0975 – 8887) Volume 55– No.10, October 2012.

[9] Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J. Liu

[10] Velmurugan, A & Mayan, J. & Niranjana, P & Francis, Richard. (2020). Expense Manager Application. Journal of Physics: Conference Series. 1712. 012039. 10.1088/1742-6596/1712/1/012039.

[11] Mastropaolo et al., ”Studying the Usage of Text-To-Text Transfer Transformer to Support Code-Related Tasks,” 2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE), Madrid, ES, 2021, pp. 336-347, doi: 10.1109/ICSE43902.2021.00041.

[12] Wang, Mingye & Xie, Pan & Du, Yao & Hu, Xiaohui. (2023). T5-Based Model for Abstractive Summarization: A Semi-Supervised Learning Approach with Consistency Loss Functions. Applied Sciences. 13. 7111. 10.3390/app13127111.

[13] Rayhan, Abu & Kinzler, Robert & Rayhan, Rajan. (2023). NATURAL LANGUAGE PROCESSING: TRANSFORMING HOW MACHINES UNDERSTAND HUMAN LANGUAGE. 10.13140/RG.2.2.34900.99200.

[14] A., Bharathi. (2023). Natural Language Processing for Enterprise Applications. Ushus Journal of Business Management. 21. 10.12725/ujbm/61.2.

[15] Z. Bozkus, C. Bisson and T. Arsan, ”Analytical expense management system,” 2009 First International Conference on Networked Digital Technologies, Ostrava, 2009, pp. 527-532, doi: 10.1109/NDT.2009.5272148.