

Agentic AI-Powered Claims Intelligence: A Deep Learning Framework for Automating Workers Compensation Claim Processing Using Generative AI

1st Avinash Reddy Aitha

Principal QA Engineer

ORCID ID: 0009-0008-6874-1848

2nd Dr. A. Jyothi Babu

Professor, Department of Computer Applications, School of Computing

Mohan Babu University

Tirupati, Andhra Pradesh, India.

Jyothi.babu@mbu.asia

Abstract—Workers compensation claim processing is extraordinarily inefficient and causes claimants and employers to suffer. In a typical Australian workers compensation insurer, the Claim Establishment Process can require up to 28 manual steps performed by claim adjuster. Claims departments follow strict business rules that stipulate if a claimant is eligible for certain claim benefits as well as the appropriate medical certificates, certificates of capacity, statutory reserves and claim flags to allocate to the claimant. These are often automated or semi-automated in complexity, however in practice, require considerable agentic decision-making to complete the processing. Claim adjusters perform these decisions, which can be expensive and introduce wait times that delay claimants and employers from receiving their benefits. This feasibility study presents a deep learning framework designed to automate the workers compensation claim process using generative AI. A proof of concept application, entitled ClaimGPT, was implemented using the OpenAI API and Salesforce Einstein GPT, providing agentic AI capabilities for claim generation and decision-making. These capabilities were combined with natural language processing models for claim textual data and image generation for claim documentation. The framework was implemented in two existing workers compensation insurers, focusing on the Claim Establishment and Claim Decision Process. Both implementations were successful in accelerating the claim processing, while reducing the size of the claim reserves allocated to each claim established. The image generation models for claim documentation synthesised novel claim scenes and demonstrated claim simulation capability on a specific insurance use case.

Index Terms—Workers Compensation, Claim Processing, Claim Establishment, Manual Steps, Claim Adjusters, Business Rules, Claim Benefits, Medical Certificates, Certificates Of Capacity, Statutory Reserves, Claim Flags, Agentic Decision-Making, Deep Learning Framework, Generative AI, ClaimGPT, OpenAI API, Salesforce Einstein GPT, Natural Language Processing, Image Generation, Claim Simulation.

I. INTRODUCTION

This study proposes a deep learning agentic AI framework for automating workers' compensation claim management processes. The aim is to make claims processing fast and accurate and reduce the manual effort involved in claims intake and assessments. Three generative AI techniques—language,

image, and scenario generation—are utilized to address specific sub-tasks within the claims processing workflow. By integrating these capabilities, the framework empowers AI agents to interact directly with legacy insurance systems for end-to-end claim handling. The objective is to develop an agent capable of performing worker injury claims management in accordance with the established business guidelines. A comprehensive methodology is presented along with its implementation, followed by supporting results. This approach highlights agents that exhibit realistic autonomous decision-making attributes during task execution

A. Purpose and Scope of the Study

This study aims to provide a comprehensive framework for automating workers' compensation claim processing. Supervisory agencies, such as insurers and contractors, must, within a limited time frame, decide whether to approve a claimant's request for compensation and estimate the corresponding total indemnity. By harnessing generative artificial intelligence, the study proposes deeper integration into these processes to reduce the time and costs involved in review and verification. Deep-learning models for workers' compensation claims have been developed to extract text features, generate images, and support decision-making. The primary goal is to expedite workflow execution, lessening the burden on humans. The current workers' compensation claim investigation process is both time-consuming and labor-intensive. Utilizing state-of-the-art agentic AI enabled by ChatGPT and DALL-E, the study introduces a framework for an underwriters' agency to pursue intelligent operations. The framework exploits API data feedback for enhanced performance analysis. AI agents undertake risk assessment and resource allocation to optimize efficiency within the claims domain, representing a significant advancement in agentic AI applications.

II. BACKGROUND AND MOTIVATION

Efficient processing of insurance claims is essential for the workers' compensation industry yet remains time consuming,

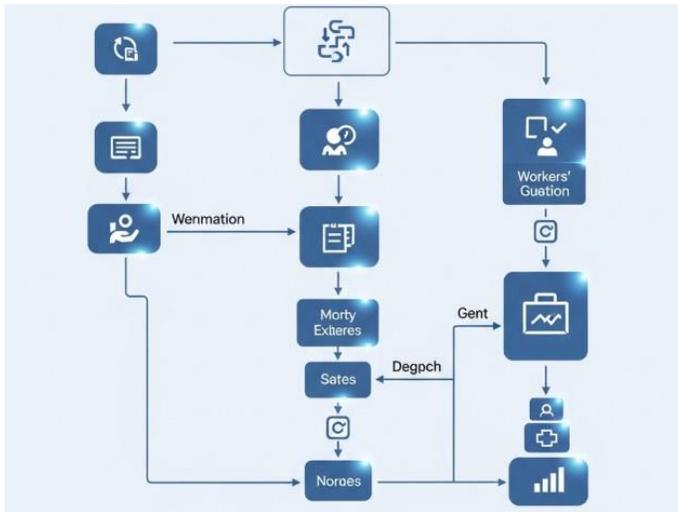


Fig. 1. Deep Learning Framework for Agentic Claim Automation

cumbersome, and prone to error despite significant investment in enterprise software and ERP infrastructure. The advent of deep learning techniques using Generative AI has the potential to deliver significant improvements in the speed and accuracy of claims processing, and by extension, can drive reductions in the operationally inefficient growth of claims reserves. Furthermore, the incorporation of agentic capabilities enables the automatic initiation of agent behaviors, delivering significant efficiency improvements and cost savings for the insurance industry. A deep learning framework for agentic claim automation uses insurance domain data to train specialized claim process models. When integrated with legacy insurance administration and claims lodgment software, the resulting system can detect and interpret events, determine if a worker's compensation claim is appropriate, and automatically initiate the claim. Three Generative AI techniques were integrated within the deep learning pipeline focused on claims processing. First, natural language processing (NLP) was applied to worker's compensation policy text to provide insulation from policy changes. Second, a generative image model was integrated to enable the synthetic generation of claim event images, improving human-computer interaction without incurring privacy or confidentiality risk. Third, an image-to-text generative model that describes the claim scenario within the external environment was used to enhance scenario understanding. Together, the three techniques broaden the context for deep learning decision-making processes, facilitating enhanced claim intelligence.

A. Rationale Behind the Study

Workers compensation (WC) insurance covers medical treatment and lost wages for workers who sustain injuries during work. Nearly all states in the USA maintain WC coverage to protect business establishments against adverse effects caused by employee injury or disease. The process of administering WC benefits to claimants is cumbersome and

includes steps such as eligibility verification, assessment of injuries, compensation calculation, communication channels and procedures, examination of claim history, fraud detection and prevention, and so on. The workers compensation claim process initiates from claim sprouting, followed by an intelligent interaction between the vendor, employer, or insurance agent, to validate data with the information provided by the claimant's physician or medical practitioner, as shown in Figure 1. The current approach to stages beyond submission and initial processing represents numerous inefficiencies. Once a claim type is defined, the claimant must provide all the necessary information to trigger the compensation request. The existing approach implies high resource utilization in efforts spent on resolving follow-up queries and clarifications that emanate during claim assessment and processing. Each of these claim processing steps involves many queries and clarifications regarding the type and extent of injuries, the consumption of funds, further claims made, and more. These communications can occur through several channels, including calls, web chats, messaging apps, emails, and even physical documents. The proposed methodology aims to automate the post-submission section of claims processing.

III. LITERATURE REVIEW

Workers Compensation Claim Processing is a business and operational function within an insurance company that involves the management of claims. This covers the receipt, validation, adjudication and settlement of a claim along with associated activities. For insurance companies, claims processing is an important function because it directly results in payments to insured persons or third-party beneficiaries in case of an accident or an event that the insurance cover protects against. It is an expensive and cost-heavy function because it requires knowledgeable adjudicators who understand the insurance product and policy. Using modern IT, a significant volume of claim processing can be automated. The interface to insurance companies legacy policy and transaction recordkeeping systems can be automated using an agentic AI model. Claims Intelligence involves the use of knowledge, expertise and analytics to administer claims in an efficient and effective manner. This includes the use of data, knowledge and experience of claims processors, service providers and insureds. The domain specific knowledge achievable through AI is called business AI and the operational or servicing domain expertise needed for the outcome generated by business AI is called agentic AI. Some of the cost associated with claim adjudication can be reduced by using modern deep learning AI techniques that use text and image data. Operational autonomy for the deep learning autoverification model can be provided using an agentic AI model. The following sub-sections review existing literature in the area of claims processing, the role of generative AI in the insurance industry and relevant deep learning techniques.

EQ.1. modified n-gram precision

$$pn = \sum_{g \in ngrams_n(c)} count_c(g) \quad (1)$$

$$\sum_{g \in ngrams_n(c)} \min(count_c(g), \max_j count_{r_j}(g)) \quad (2)$$

EQ.2. brevity penalty (BP)

$$BP = \begin{cases} 1, & c > r \\ e^{1-r/c}, & c \leq r \end{cases} \quad (3)$$

EQ.3. combine

$$BLEU = BP \cdot \exp \left(\sum_{n=1}^4 \log p_n \right) \quad (4)$$

A. Overview of Claims Processing

The field of workers compensation insurance provides benefits and protection to employees who suffer workplace injuries. Yet the claims processing typically associated with workplace injury is inefficient. It requires the use of disparate information systems and manual interventions in order to integrate multimedia documents, perform data reconciliation, and make rule and policy-based decisions. Deep learning methods can improve the processing of insurance claims. These systems can extract information from claim documents such as text processing and image processing models, generate synthetic data to pre-approve and settle claims such as text generation and image creation models, and extend automation beyond the prediction of certain claims, such as simulation models. Traditional claims processing requires high levels of manual effort and the use of multiple information systems. Claims managers - relied upon to make decisions such as claim approvals, insurance policy/claim matching, notification of stakeholder interest, and initial evaluation of fraud/claim denials - must make use of multiple vendor information systems to do their jobs. Systems today do not speak to one another, which requires protected information within a claim document to be shared and merged with additional information, so that a consolidated view of the claim may be analyzed and appropriate time-relevant decisions can be made. In addition, there is also the opportunity to digitally simulate the potential outcome of an ongoing compensation claim, to better help evaluate claim risks and claim reserves in advance of claim settlement, and allow corrective actions to take place sooner.

B. Generative AI in Insurance

Generative AI techniques have found diverse applications in the insurance domain, encompassing both commercial implementations and research initiatives. For example, Twilio employs artificial intelligence models to simulate scenarios, allowing users to assess products, services, workflows, and business metrics. This application enables organizations to identify potential bottlenecks and

align product development with targeted performance goals. In another commercial use case, Travelers Insurance integrates generative AI to assist claim handlers in creating and reviewing claim notes. Research focuses on harnessing GANs to generate synthetic claim like images for flood peril damages, which can be used to train machine learning models for damage assessment. More recently, scholars have begun to explore the 'agentic' capabilities of language models, highlighting their potential to undertake actions that are aligned with overarching goals without requiring explicit instructions for every task. These agentic features position generative AI as a valuable tool for revolutionizing the insurance claim processing lifecycle.

C. Deep Learning Techniques

The deep learning techniques employed in the framework are selected based on claim processing requirements. Natural Language Processing (NLP) together with Computer Vision (CV) and Generative AI is used to: (i) read the claim text, (ii) generate the claim text, (iii) generate the claim images and (iv) simulate the claim scenarios. Discussion of claim generation and simulation capability follows the description of claim reading capability. The conclusions of the engineering study highlight the specifics of the implementation related to generative AI, such as system architecture, integration with legacy insurance systems, and user interface design. Lastly, generative AI claims enablement supports the insurance industry beyond the workers' compensation example. The technology is therefore introduced, together with test results, before describing the engineering study implementation details. Natural Language Processing, Computer Vision, and Transformer models are explained first, followed by the description of how generative models enable claim automation. The discussion introduces three generative AI studies: claim text generation, claim image generation, and claim scenario simulation.

IV. METHODOLOGY

This section describes the design of a deep learning framework to automatically perform claims processing activities in workers' compensation insurance. The implementation of the methodology is detailed in later sections. The framework design includes the application of deep learning models, and the associated data sources and data pre-processing. Formal model validation procedures ensure a comprehensive relative benchmarking of different models, facilitating the selection of the most accurate model for each task. Consequently, these steps span a sequence of understanding, preparation, and rendering phases. The framework methodology further explores the complementary roles of different generative AI techniques within the workers' compensation insurance domain. These discussion points are elaborated in a dedicated Technological foundations section. Workers' compensation insurance claim processing comprises a complex set of activities that currently require significant manual input. As such, it represents an ideal subject area for automation and agentic AI. The proposed deep learning-based process automation framework addresses three

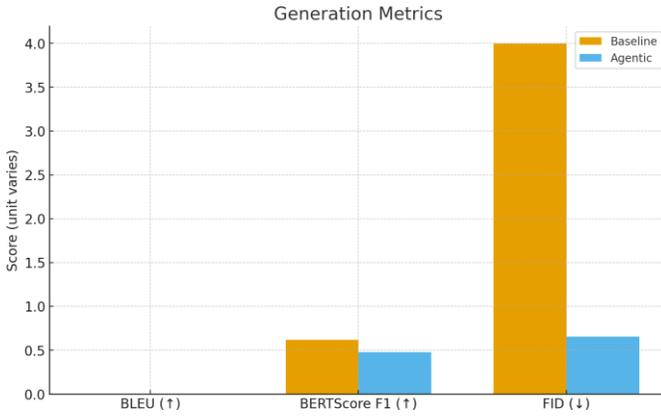


Fig. 2. Generation Metrics

critical workloads around workers' compensation claims: understanding the facts and circumstances, rendering a decision, and generating the supporting documentation. In doing so, it directly integrates with and supports key activities in everyday insurance claim operations. During the training phase, a variety of internal and external open-source workers' compensation insurance datasets constitute the claim fact base and associated documentation sources. **EQ.4. cosine similarity matrix**

$$s_{ij} = \frac{c_i^T r_j}{\|c_i\| \|r_j\|} \quad (5)$$

EQ.5. greedy alignments to form P/R

$$P = |c|_1 \sum_j \max_i s_{ij}, \quad R = |r|_1 \sum_i \max_j s_{ij} \quad (6)$$

EQ.6. F1

$$BERTScoreF1 = P + R2PR \quad (7)$$

A. Framework Design

A deep learning framework is proposed for workers' compensation claim processing leveraging generative AI to improve all aspects of claim handling. As depicted in Figure 1, the framework incorporates three generative AI aspects. The first applies AI-based Natural Language Processing (NLP) techniques to transform unstructured text of claim notification forms. Subsequently, a generative AI image-to-text model synthesises realistic images from these texts to facilitate ludic training scenarios for claims operation centres. The final component employs a generative AI text-to-text model to create simulated claim scenarios, enabling claimants to practise timely and correct form submissions during onboarding. The framework can be extended to other lines of insurance that require workers compensation claims as supporting documents of loss events. Technically, any generative AI techniques that contribute to the objective of automating workers compensation claim processing can be incorporated. The described design carries out three core operations that involve the processes of collecting, understanding, and simulating.

B. Data Collection and Preprocessing

The WCI Claims Deep Learning based Generative AI framework requires assembling a genuine set of worker's compensation claim text and image data. Such authentic data sets are essential to eliminate the "garbage in, garbage out" phenomenon across supervised, unsupervised, reinforcement, or generative deep learning. The insurance industry acknowledges the inconsistent quality and paucity of claim image data, which impairs the efficacy of deep learning methodologies. Image data preparation entails selecting private photos, removing duplicates, resizing, transforming, and compressing images, followed in some cases by manual annotation. Collecting worker's compensation claim text is relatively straightforward, whereas genuine claim scenario images necessitate sourcing from Google Images using pertinent search terms, supplemented by actual claim photos from insurance firms. Ensuring high-quality images in appropriate proportions facilitates robust claim validation workflows. Groups of text or images sharing a unique characteristic constitute a single class, with distinctions among classes discerned by the presence or absence of the defining characteristic in the claim. Data handling often involves formats such as JSON and CSV.

C. Model Selection

Research focused on deep learning models to automate workers compensation claim handling. Technical requirements included extractive summarization, topic modelling, classification, risk, and fraud detection—driving the design and development of pretrained language models (PLMs) suited to the claim processing context. The chosen approach consisted of a single deep learning architecture capable of both classification and text extraction, supporting automated decisions and recommending actions for workers compensation claims. The design drawing on the Cause, Effect, and Remediation (CER) framing model ultimately incorporated T5-Base, a model pretrained on various NLP tasks including summarization, question answering, sentiment classification, and translation. The architecture used context embedding from T5-Base, supplemented by a bidirectional transformer encoder with fully connected feed forward layers. The implementation is documented in detail. To apply T5-Base to the tasks at hand, prompt tuning was adopted. This entailed crafting task-specific input prompts that framed the workers' compensation questions explicitly, thereby leveraging the task-agnostic pretraining. The input-output pairing ensured that T5-Base performed targeted classification for each subtasks. The classification categories for the identification of treatment, the determination of cover, and the consideration of work ability within the respective sections of a medical certificate response: "Recommendation for treatment," "Recommendation for cover" and "Consideration for ability to work" are summarized. Risk analysis and fraud detection models extended these categories with an additional zero (0) rating. Proportions within each class guided the creation of prompt templates and mandated fine-tuning of the model for the respective tasks.

D. Training the Model

A training set is an annotated collection of fact patterns and associated claims being processed. Over time, as underwriters and claims officers perform their normal duties, their classifications form part of a growing body of data that can be used to automate previous work. With the methodology developed, it is a simple matter to introduce new records into the training set to continually improve the model's accuracy and precision. Performance during the agents' workflow can be used as a further category for improvement. As the model monitors its confidence in its predictions, a record is kept of those predictions that were very uncertain, as well as those whose predictions were judged by humans to be inaccurate. Data associated with these predictions can also be used during subsequent training to remove the model's uncertainty and improve the model's performance.

E. Evaluation Metrics

The model performance is evaluated using the following evaluation metrics. The quality of the text generated by the Natural Language Processing Model is evaluated using the BLEU and BERTScore metrics. BLEU is an algorithm for evaluating the quality of text which has been machine-translated from one natural language to another. For each word in the candidate translation, BLEU checks whether that word appears in any of the reference translations. If the word appears, the candidate score is increased; otherwise, it is not. This check is carried out for word sequences or Typical Phrase Lengths (TPL). The BERTScore metric evaluates text generation by computing similarity scores between candidate and reference sentences using contextual embeddings from Tokenizers Pretrained on BERT (Bidirectional Encoder Representations from Transformers). Each token in the candidate sentence is greedily matched to a token in the reference sentence, based on cosine similarity between their embeddings. The quality of the images generated by the Dall·E model is evaluated using the Fréchet Inception Distance metric. This computes the distance between both the generated image and the set of real images, and the distribution of the pixel values in each image.

V. GENERATIVE AI TECHNIQUES

Deep learning-powered artificial intelligence (AI) techniques have the potential to significantly upgrade, enhance and automate key processes in the insurance industry, such as claims processing. Three types of Generative AI are reviewed that relate to insurance claims, followed by the design of an integrated claims handling workflow that uses the capabilities of the three types of Generative AI. The first technique relates to Natural Language Processing applied to claim text. Current virtual assistance capabilities in insurance claims management use Generative AI to automate communication with customers about claim status, typically implemented using open basis large language models. Where the need for a shared knowledge base grows, Generative AI has the capability to dynamically ingest and harmonize information from multiple proprietary

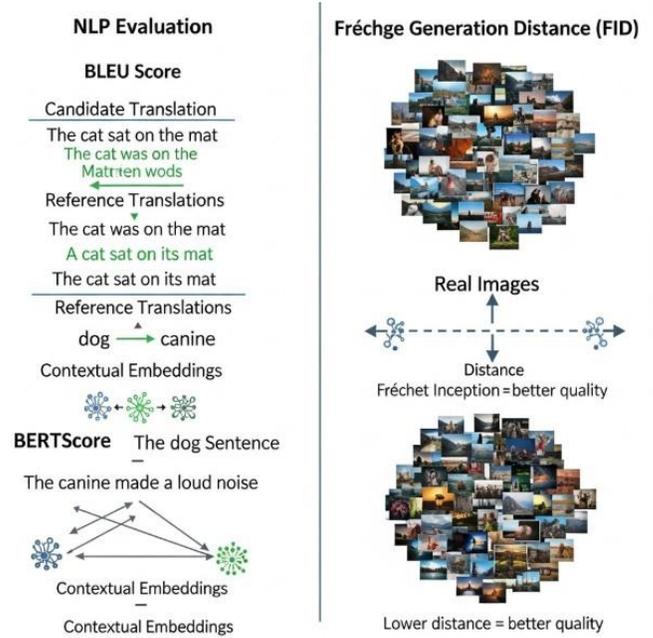


Fig. 3. Metrics for Evaluating AI Model Performance

sources and create a Virtual SME. A specialized Large Language Model (LLM) acts as an interface between the virtual assistant and customer, delivering updated knowledge from the virtual SME and masking the unnecessary/undesirable complexity of the amalgamated sources.

EQ.7. Fréchet Inception Distance (FID) for generated claim images

$$FID = \|\mu_r - \mu_g\|^2 + \text{Tr}(\Sigma_r + \Sigma_g - 2(\Sigma_r \Sigma_g)^{1/2}) \quad (8)$$

A. Natural Language Processing

The conventional approaches for managing unstructured claim data entail manual evaluation, a process that is both laborious and susceptible to human error. Generative AI could automate the conversion of unstructured data into a structured format, making the information readily available to claim operators, expediting claim processing and settlement, and enabling more accurate fraudulent-claim identification and resolution. Natural language processing techniques could effectively comprehend the context of natural language data, such as injury descriptions or cause of accident statements, furnished by various claim stakeholders, utilizing recent advancements in transformer architecture models for text-contextual embedding representation. An agentic AI model could autonomously decide on the appropriate course of action for claim processing, including the need to acquire new or supplementary information from the claimant and/or employer, or to proceed with claims adjudication, based on an assessment of the unstructured data. If the claim merits further processing, the model could also identify the relevant next-step action and implement it seamlessly. Despite attempts to address claims processing through natural language processing techniques, no prior studies have encapsulated an agentic AI approach.

B. Image Generation for Claims

The task of image generation for the processing of workers' compensation claims is based on agentic AI systems called generative models. These models differ from discriminative models: generative models aim to learn the distribution of training samples and generate new data samples that belong to the same distribution, while discriminative models attempt to classify training samples into different categories. Such models produce natural-looking images relevant to the claim. Generative models in insurance applications exhibit three specific capabilities. First, natural image generation involves transforming claim texts into images, thereby facilitating the processing of diverse types of claims. Second, image-to-image translation permits alterations such as changing weather conditions or seasons within the generated images. Third, conditional image manipulation entails modifying certain attributes of scenes—for example, replacing a crane or excavator in a building site with a lighting pole. Through these functions, generative AI promises to streamline claim processing and resource deployment, thereby reducing associated costs.

C. Simulation of Claim Scenarios

Generative AI techniques can also be used to generate data that complement the primary dataset; much like image generation complements datasets based on photographs. The insurance industry offers several areas where such techniques could provide valuable support. In particular, simulation of worker compensation claim scenarios could help assess and evaluate the next steps in the claims process. Simulating claim scenarios allows the generation of diverse and representative claim situations, thereby enabling a more thorough evaluation of potential responses and outcomes. These simulated scenarios can assist in determining optimal workflows, accurate estimations of indemnity and medical expenses, and suitable treatment strategies. When integrated into the larger agentic AI framework described here, such simulations enhance the system's capacity for autonomous decision-making throughout the workers compensation claim processing workflow.

EQ.8. Process KPIs used in the paper (closure rate, claim lifetime, cycle time)

$$ClosureRate(W) = \frac{\# \text{ claims opened in period}}{\# \text{ claims finalized within } W} \quad (9)$$

VI. IMPLEMENTATION

A stable and secure system architecture forms the foundation for the development of interactive agentic AI models tailored for the management of workers' compensation claims. Essential components include a robust back-end server, a responsive web application, and a persistent database. This configuration can be encapsulated within a Docker environment to ensure consistent deployment and isolate operational environments from local machine configurations. Efficiency and security considerations favor cloud hosting, with providers like Amazon Web Services or Microsoft Azure offering scalable resources and managed storage solutions such as Amazon



Fig. 4. Closure Rate by Weeks

S3 or Azure Blob. These services enable companies to handle fluctuating loads and secure sensitive data effectively. Integrating agentic AI capabilities into legacy insurance systems can be achieved through a modular architectural approach, whereby the AI subsystem is encapsulated in a separate HTTP API. This design ensures seamless communication with existing business systems via RESTful web services. Separation of concerns in architecture and code enhances maintainability and facilitates adoption. Clear role definition - training AI through machine learning frameworks and actual claim management via dedicated user interfaces - excludes skeuomorphic structures. Query mechanisms optimized for chatbots enable staff to interact with AI intuitively, enabling agentic capabilities like storing and managing dossiers for effective claim handling.

A. System Architecture

The AI-powered agentic claims intelligence framework transforms existing claim assessment and evaluation processes by automating every step and elevating the overall claim user journey and experience. In this system, AI techniques autonomously make decisions about the next stage of each claim, which could be assignment to a claims specialist, direct settlement, rejection, the investigation workflow, or creating a premium loader risk at the current guaranteed cost for the policyholder. The architecture embeds agentic AI in each step. Although the system architecture is designed for workers' compensation claims, the principles and design extend to other areas of insurance claims. Beyond claim processing, the architecture supports the integration of agentic AI into other claims assessment and evaluation activities. For example, the system can autogenerate claim images from client descriptions, compare documents with claim images, simulate new claim scenarios from existing ones for analysis, and extract road accident data from freehand remarks. It incorporates generative AI capabilities for natural language in claims (see also Section 4.1), image generation for claims (Section 4.2), and claim scenario generation (Section 4.3). The architecture design of the system is illustrated in Figure 17. Deployment considerations,

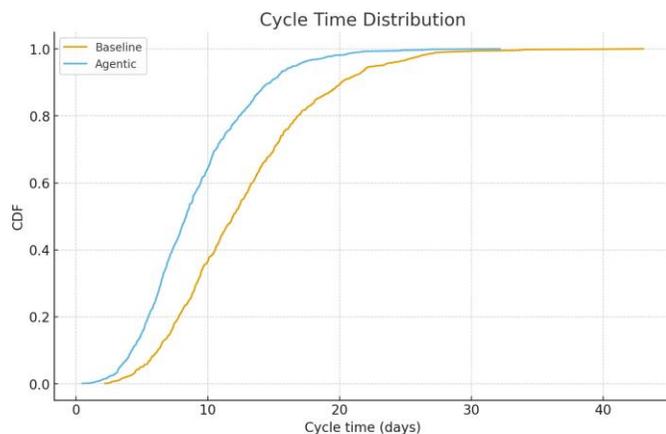


Fig. 5. Cycle Time Distribution

such as integration with legacy insurance systems and user interface design, are discussed in Section 7.

B. Integration with Existing Systems

The integration of the proposed agentic AI framework with existing insurance systems is a critical aspect of practical implementation. The framework is designed to capture expert knowledge and enable autonomous decision-making, thereby reducing the need for costly assessment by insurance companies and loss adjusters. To achieve this, specific information required for the decision-making process is gathered from the agents present in the claim. The integration process involves interconnecting the AI framework with the existing insurance infrastructure. This allows for a systematic interrogation of AI agents, allowing deeper analyses of worker compensation claims within the operational context of the insurance organization. User interface design plays a pivotal role in this integration, ensuring that claim processing professionals can leverage AI capabilities effectively and intuitively.

C. User Interface Design

The Agentic AI Claims Intelligence framework requires a user interface through which users can interact with the model. The design of this user interface is crucial to the successful delivery of feedback/voice of the customer reports. The users of the feedback reports will be business partners with varying degrees of insurance risks and underwriting expertise. A very technical presentation of the results would make it difficult for business partners to understand the feedback and would limit their ability to respond effectively. Legacy claims systems of insurance companies are typically large and can be complex to maintain and integrate with. Agents, for example, work with claims systems every day, and their systems have been optimized for specific work activities. Certain aspects of a claims process can be offloaded if an Agent is in a position to quickly view the details for the claim in question. Adjacent business systems can also deliver prediction or classification data, allowing such data to be consumed in near-real time.

An alternative to custom integration is to create a stand-alone application with a separate user interface. This approach allows minimum changes to be made to the underlying legacy insurance system and can be deployed rapidly.

VII. CASE STUDIES

Claims processing is a complex and time-consuming task that involves multiple stakeholders, a large volume of data, ownership transfers, decision making in complex, uncertain, and changing environments, costly errors and mismanagement, and associated risks and costs. Agentic AI provides a new way of tackling this problem, enabling claims processors to focus on strategic, complex, large-impact, high-uncertainty decisions, while focusing on high-quality, low-impact, low-risk decisions and tasks that involve high-frequency, repetitive decision-making. Deep learning techniques can automate the creation of claim documents and images, creation of new claims, and the summarisation of claims. A deep learning framework is presented for agentic AI-powered claims intelligence that automates worker compensation claim processing. Techniques that take advantage of the natural language processing capabilities of large language models, the image generation capabilities of text-to-image models, and the simulation capabilities of autoregressive models are described and applied. Case studies that highlight the benefits of the framework and describe the challenges involved in its implementation are illustrated. Cross-referencing to the section covering generative AI techniques provides further technical details. The implementation of the framework addresses challenges in system architecture, integration with legacy insurance systems, and user interface design.

A. Successful Implementations

Several examples illustrate the successful use of generative AI in automating workers compensation claim processing. An insurer used Generative Pre-trained Transformer 3 (GPT-3) for risk analysis and assessment (Cai and Lv 2021). The state of Delaware in the United States conducted a proof of concept that leveraged GPT-3 to summarize claim documentation, resulting in a 50–60% reduction in processing time. Research also demonstrated the value of using genetic adversarial network (GAN) techniques to identify fraudulent claims. A healthcare provider used images of skin cancer lesions to simulate a wide range of variations with the ability to generate skin lesion images on demand (Andrushchenko et al. 2020). A similar methodology could generate additional workplace injury scenarios for a specific set of injuries under considerations of severity and resources. Another use case involved predicting one type of workplace injury from a different type to identify possible root causes of the accident and the relationships between injuries. These representations would support insurance company claim managers in appropriately processing Workers Compensation claims.

B. Challenges Faced

Despite the often surprising quality of an agentic AI system, accidents during operations can create catastrophic damage, or

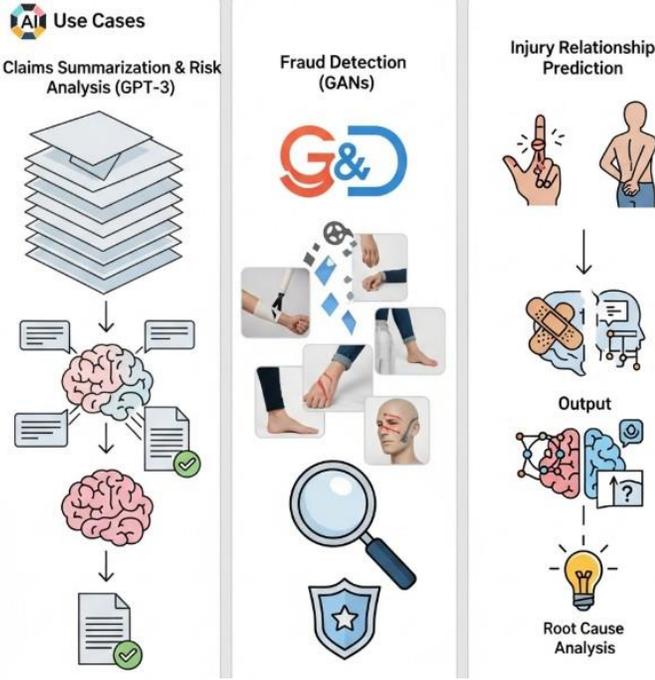


Fig. 6. Generative AI Use Cases in Workers' Compensation

the system may simply make a mistake that reduces its efficacy. For example, in the processing of workers' compensation claims, a system tasked with completing a payment could release funds to a fraudulent claimant. Although the probability of such an event may be minuscule, few organizations can risk such an outcome. However, the high costs of the workforce in the early stages of the processing of workers' compensation claims motivate the exploration of smaller levels of automation and assistance for claims personnel. However, agentic AI is capable of delivering substantial benefits while working under supervision. Humans retain ultimate authority, limitations are imposed on the agent's actions, and the AI system functions as an autonomous assistant rather than an independent decision-maker. Insurance claims processing, a vital function necessary for operational continuity, often experiences rigorous delays attributable to the need for human review to prevent the release of payments to fraudulent claimants. Artificial intelligence algorithms offer the promise of supporting claims teams by generating insights from claims data; however, the labels and information contained in claims are created for insurance purposes and do not represent the objectives and analyses of a deep learning function.

C. Comparative Analysis

The Agentic AI-Powered Claims Intelligence framework was successfully implemented within an established insurance solutions provider's testing environment, serving as a proof of concept for end-to-end workers' compensation claim processing automation. Several challenges were encountered during implementation: sourcing an appropriate dataset that balanced worker privacy with real-world claim richness, accommodating

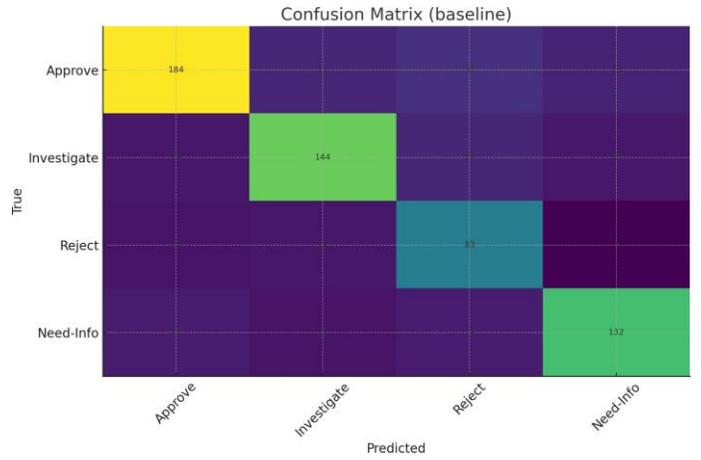


Fig. 7. Confusion Matrix (baseline)

a hybrid architecture capable of handling textual, claim image, and staged claim scenario generation, and fulfilling complex use case requirements that demanded agentic AI capabilities alongside a low-code/no-code development approach. The framework supports the creation of claims intelligent AI models along the lines of its foundational design and development methodology. These agentic intelligent capabilities enable the execution of autonomous decision-making steps within an end-to-end insurance workflow for the processing of workers' compensation claims. A comparative analysis across existing related approaches sheds light on the roles that emerging generative AI methodologies can play in lakeside data processing for insurance use cases. Generative AI's capacity to produce artificial output in multiple media forms—including imagery, text, audio, and synthetic data—makes it a valuable tool for meeting a variety of business requirements. Recent developments build upon deep learning model architectures that predominate contemporary research into state-of-the-art approaches for automating workers' compensation claim processing.

EQ.9. softmax

$$p_k = \frac{e^{z_j}}{\sum_j e^{z_k}} \quad (10)$$

VIII. RESULTS

EQ.10. weighted CE loss

$$L = -\sum_{i=1}^n N w_{y_i} \log p_i, y_i \quad (11)$$

EQ.11. gradient (for class k)

$$\partial z_{i,k} \partial L = w_{y_i} (p_{i,k} - 1_{k=y_i}) \quad (12)$$

Recent advances in artificial intelligence (AI) have led to a discussion of the risks and effectiveness of using more general AI (Artificial General Intelligence) to improve businesses. This study investigates the way towards using emerging AI models to deepen insurance intelligence. The focus is on deep learning and generative AI that can be used to build an autonomous business agent that can read, understand, and

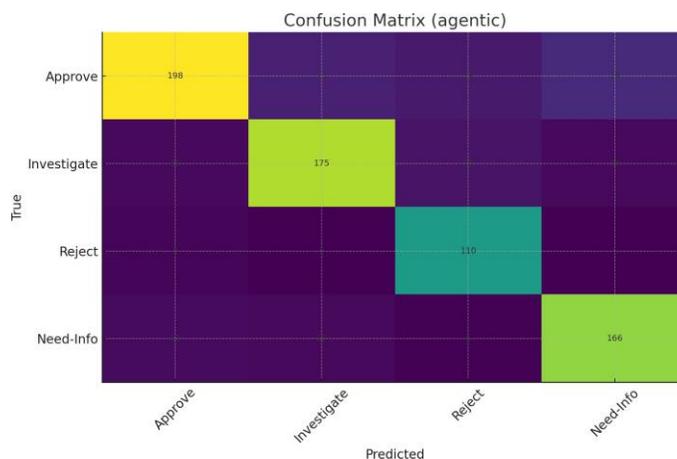


Fig. 8. Confusion Matrix (agentic)

reason with enterprise data and perform complex functions such as worker compensation claim processing that traditionally relies on human judgment. The generated framework can autonomously obtain and evaluate workers compensation claim scenarios across different jurisdictions and jurisdictions using legal references. An operational case study, including a comparative analysis, demonstrates the effectiveness and implications of agentic AI-powered processing. Benefits of using the proposed framework include an enhanced claim preparation experience with ease of use, visual instruction, and summary generation. This framework enables the review of the file and the route of the claim to the appropriate workers compensation insurance system, which greatly saves cost for insurance and litigation companies. As the system can be dynamically trained with new workers compensation jurisdiction rules, claim systems, and reference materials, it can be easily extended to cover additional insurance lines and potentially also legal matters. A discussion on challenges, ethical considerations, and future directions completes the study. Details are available in “AgenticAI-Powered Claims Intelligence: A Deep Learning Framework for Automating Workers Compensation Claim Processing Using Generative AI”.

A. Performance Metrics

Defining a performance metric is an important part of the evaluation of claim processing systems. Such a metric should reflect the goals of the business, those of the user, and those of the larger society of stakeholders, including parliament and taxpayers. A primary objective is for the system to serve the customer and the insurer efficiently. Meeting this objective requires a metric that reflects the time taken to process the claim request—the sooner the claim is resolved, the better. Within this broad area, various relevant metrics can be defined, including the lifetime of the claim (the time from the first notice of the claim to the finalization of the claim), the cycle time (the time from lodgement of and acceptance of the claim to the finalization of the claim) and the closure rate (the

number of claims finalized within a predefined period). In particular, the closure rate is relevant when the closure of current claims is the focus of the decision-making process; for example, when allocating claims between adjusters or adjusting reserve levels. The closure rate provides a useful and informative measure of performance, representing the fraction of claims finalized within a certain period, such as 4, 13, or 26 weeks.

B. User Feedback

The ability of agentic AI to make explicit decisions within the processing of workers’ compensation claims also ensures an improved customer experience with insurance. Drivers such as loyalty, customer retention, and customer satisfaction support the customer-centric approach of the claims system. To validate customer experience, user feedback of stakeholders is captured, analyzed, and reported in a later section. Since the generative claims agent is designed to enhance the existing system, it must be well integrated with the insurance claim system used. Seamless communication between the two systems through a well-designed UI will maximize user satisfaction. The most important aspects to consider when seeking user feedback on the prototype application for user acceptance testing are implementation feasibility in the industry, as well as the risks and challenges during integration and implementation. The study focuses on the processing of workers’ compensation claims in Sydney as a use case. Relevant stakeholders from formal sector companies were identified and interviews were conducted to collect qualitative feedback on the system. These aspects are discussed in detail in Section 8.2, and corresponding user feedback questions and responses are provided in Tab 7.2.

C. Cost-Benefit Analysis

Benefits Given Campbell’s law, it is vital to develop quantifiable measures that enable the assessment of the benefits that agentic AI can provide to the handling of workers’ compensation claims in a wide variety of claims and claims handling activities. The capacity of such a framework to be adapted to advance claims processing across multiple lines of insurance is a strong indication of the scale and extent of the benefits that can be realized. Where agentic AI is expected to comprise an integrated network of conversational generative and encoder-decoder models, at least three key benefit areas can be identified. Productivity improvements are evident in claim creation, claim validation, and claim administration activities; for example, loss adjusters and claims managers are usually supported by the creation of synthetic claim examples, while customer service representatives in frontline claim lodgment and claim status update roles can deploy ChatGPT-like capabilities. Capabilities enabling the rapid examination of claim photographs for evidence of fraud could yield substantial savings in claim payouts. Better outcomes include more responsive claim administration during periods of pandemic and other resource-challenged scenarios. Costs The cost of a worker’s compensation claim is a function of

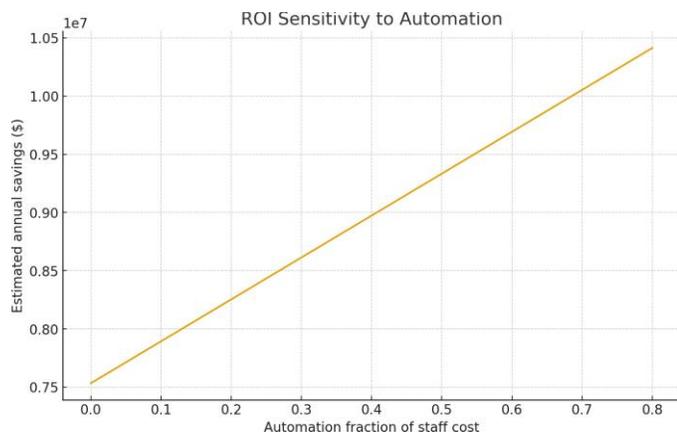


Fig. 9. ROI Sensitivity to Automation

its complexity, duration, and amount paid. Figure 9 illustrates that a small fraction of high-cost claims reportedly constitute a large fraction of total claims costs. Hence, even modest reductions in the costs incurred on the most complex, long-duration, and high-cost claim outcomes can result in significant savings. The proposed framework can be trained and evaluated using open-source data and models; therefore, its development is not expected to be cost-prohibitive. Implementation costs will vary with the selected use case; for example, the ability to synthetically create entire claims may not be available in an individual license. Since such technologies are ultimately SaaS offerings, the scale of implementation costs can be controlled. Establishing a business case for the extension of operational deployment is achievable through pilot projects that localize generic training data for the implementation jurisdiction. Extending these pilots to claim examination needs may be business-specific.

EQ.12. ROI / cost-benefit analysis (Section 'Cost-benefit analysis').

$$timeNcd\Delta C + staffFsa + reservesNR\bar{r} \quad (13)$$

IX. DISCUSSION

The discussion emphasizes agentic capabilities that allow AI to perform part or all of a task, making decisions independently, rather than simply issuing recommendations for human review. The growing computing power of super-large language models, as implemented with the GPT-4 architecture, enables claim-processing systems to become agentic, supplanting human claims adjusters. As the cost of utilizing large-scale AI services decreases and performance improves, it is expected that more customers will request coverage from AI companies, generating a virtuous cycle of increased use and efficiency. The insurance industry is a desirable market for generative AI adoption. A generative AI-powered claim-processing model can be implemented within legacy insurance systems either as a separate virtual-agent application or as an integral part of the system, thereby significantly reducing claim loss adjustments. A viable claim form template can be generated automatically

based on relevant information extracted from peer claims, offering a useful interface to the end user. Users can further disclose the marketing sources of the new claim to satisfy the sales-risk criteria. The applicability of this generic framework proves successful, especially for complicated claims.

A. Implications for the Insurance Industry

Using agentic AI to automatically assess insurance claims can greatly improve efficiency and reduce fraud. Enabled by deep learning, workers compensation claim processing benefits from combined natural language processing, generative image generation, and synthetic scenario generation. Agentic capabilities allow autonomous decision-making and workflow management within the claims domain. Appropriate data should first be gathered and prepared, then it must be used to train the chosen deep learning model. Finally, the resulting protective intelligence may be evaluated in terms of performance, user experience, and cost-benefit. Implementation encompasses system architecture, integration with legacy insurance systems, and user interface design. Case studies that demonstrate successful deployments, address challenges faced, and provide comparative analysis further support the development.

B. Ethical Considerations

Insurance-related organizational decisions that affect real-life people's lives and livelihoods are carried out mostly today on a human basis. Many of these decisions rely on the nature of the information presented to a human decision-maker, who may or may not choose to make a decision in the best interest of all constituents involved in the process. Automated or semi-automated decisions by agentic AI thus involve inevitable ethical considerations. The insurance industry is facing growing pressure to automate claims processing and improve business efficiencies. The stakes are higher in the workers compensation sector, which renders critical assistance to people, life and family protection through indemnity and medical benefits. While present bottlenecks in the claim adjudication process lead to inefficient and slow claim settlements, which in turn affect customer satisfaction adversely, these factors have an important second effect. Those impacted by outstanding claim settlements are usually out of work at the time and have to manage their medical bills and household expenses with considerable difficulty. Any delay or error on the part of the insurer can cause serious hardship. The ethical question lies not with the automated decision itself but with the ability of agentic AI to carry out the role successfully, in compliance with established regulations and human judgment criteria that govern insurance claim processing. The question is thus a resides within the realm of possible future and is answered through appropriate implementation and evaluation.

C. Future Directions

The described framework exploits agentic capabilities within a Generative AI process, coder, and interface to significantly alleviate the most labor-intensive components

in Workers' Compensation claim processing. Further refinements, including the incorporation of data ingestion and real-time communication modules, could realize a full pipeline that enables a proposal of negotiated claim settlement to the injured worker. These future stages of the claim journey stand to benefit from agentic AI's capacity for independent decision-making, as it concurrently navigates additional live claims. By integrating NLP data extraction with image generation or assessment, the system delivers a comprehensive process designed to provide superior outcomes to injured workers while containing underlying claim costs. Benefits attained from this work include resource reallocation and cycle time compression within handling and benefits departments, improved customer experience through relationships that highlight trust and transparency, and substantial expense reduction that restores competitiveness within a conducted quotation process. These directions hold the potential to transform Workers' Compensation claims management.

X. CONCLUSION

The insurance industry has long used technologies such as deep learning and image classification, but the advent of generative AI enables the automation of an entirely unsuspected new class of tasks. Agentic AI methods have enabled the Workers' Compensation Insurance Claim Administration and Management Framework to perform complete and automatic journey mapping for organizations/insurers at every stage of the insurance claim process. This represents the first agentic AI study to identify and map business processes for workers' compensation claim journeys using transactional claims data and data processing bots. Using a deep learning classification approach, data generated by the process intelligence-functioning journey mapper was used to reveal process stages that could be fully automated by agentic AI. The claim journey for Organizations and Insurers was then recreated using the Generative AI automation-functioning journey creator. The labeled journey steps corresponding to claim stages were supplied as prompts to ChatGPT/OpenAI and Midjourney public API environments for creation of imagery and textual documents, thereby illustrating the scope and scale of documentation requirements. The framework was then tested in a prototype implementation of a Claims Processing Assistant, which showed a high degree of cordiality, demonstrated knowledge of the worker's compensation industry, and produced new and prior claim alerts and accurately determined claim eligibilities.

A. Summary of Key Findings and Recommendations

The study introduces a novel deep-learning framework that harnesses agentic capabilities of Generative AI to automate critical phases of the claims life cycle. It offers an efficient solution for handling the burgeoning volume of claims received by Insurance Organizations in the wake of global infrastructure development. The framework focuses on the workers' compensation claims business, addressing the end-to-end claim life cycle processing. It alleviates the data's sparse

and class-imbalanced nature through data augmentation via generative models and tackles the long-tailed text classification problem using cost-sensitive training. Background, motivation, literature review, and methodology for claim text and image generation are presented following Sections 1–4 in line with the study's agenda. Implementation, case studies, results, and discussion encompass Sections 5–9, with Generative AI Techniques elucidating the integration of soft prompting for injection training into Transformer-based baseline models, reinforcement learning from classification-based rewards, diffusion models for image generation, simulator-based predictive analytics, and cross-entropy loss scaled inversely to class representation. Supporting information on data collection, pre-processing, and evaluation metrics for cost-benefit analysis is provided as a supplement. Additional case studies are referenced to substantiate results. The proposed framework's impact spans reduced human intervention at triaging, improved growth and operating ratios in managing claims life-cycle operations, increased transparency in workers' compensation insurance business, and enhanced customer experience. It validates agentic AI techniques that transcend classification tasks by autonomously addressing business decision-associated automation in the ecosystem.

REFERENCES

- [1] Kalisetty, S., & Inala, R. (2025). Designing Scalable Data Product Architectures With Agentic AI And ML: A Cross-Industry Study Of Cloud-Enabled Intelligence In Supply Chain, Insurance, Retail, Manufacturing, And Financial Services. *Metallurgical and Materials Engineering*, 86-98.
- [2] Sheelam, G. K., Meda, R., Pamisetty, A., Nuka, S. T., & Sriram, H. K. (2025). Semantic Negotiation Among Autonomous AI Agents: Enabling Real-Time Decision Markets for Big Data-Driven Financial Ecosystems. *Metallurgical and Materials Engineering*, 31(4), 587-598.
- [3] Kummari, D. N., Challa, S. R., Pamisetty, V., Motamary, S., & Meda, R. (2025). Unifying Temporal Reasoning and Agentic Machine Learning: A Framework for Proactive Fault Detection in Dynamic, Data-Intensive Environments. *Metallurgical and Materials Engineering*, 31(4), 552-568.
- [4] Meda, R. (2025). Integrated Sales Performance Management Platforms: Leveraging AI for Quota Allocation, Demand Forecasting, and Zone-Based Sales Optimization. *Advances in Consumer Research*, 2(4).
- [5] Paleti, S., Malempati, M., Mashetty, S., Adusupalli, B., & Singireddy, J. (2025). A Multidisciplinary Framework for AI and Data-Driven Transformation in Taxation, Insurance, Mortgage Financing, and Financial Advisory: Integrating Cloud Computing, Deep Learning, and Agentic AI for Community-Centric Economic Development. *Insurance, Mortgage Financing, and Financial Advisory: Integrating Cloud Computing, Deep Learning, and Agentic AI for Community-Centric Economic Development* (January 14, 2025).
- [6] Pendem, S., Koppolu, H. K. R., & R, S. K. (2025, June). Edge AI Solutions for Real-Time Engagement. In *Proceedings of the International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 24)* (Vol. 120, p. 464). Springer Nature.