

Original Article

Generative Intelligence and AIOps for Autonomous Workflows in Data-Centric Hospitality and Agricultural Technology Platforms

Dr. Isabella Romano

Biomedical Engineering, Mediterranean University of Rome, Italy.

Received: 12-08-2025

Revised: 12-31-2025

Accepted: 01-03-2026

Published: 01-09-2026

ABSTRACT

Recent developments in Generative Artificial Intelligence (GAI) have enabled training of large language models (LLM) on domain and task-centric corpora. The functionality available through GAI Foundations facilitates autonomous, agile, and compositional solutioning of data-intensive tasks and offers new avenues for creativity in section content generation. However, application of Generative AI in operation of complex, data-intensive systems such as hospitality services and AgriTech platforms appears limited. A joint-use compositional-network architecture is proposed, enabling automatic generation and execution of GAI-prompted task plans, across multi-enterprise hospitality and AgriTech platforms. Two-way GAI-enabled workflow automation and autonomic technology adoption at accumulate-and-release PaaS joint clouds is explored for further business value creation, harnessing concepts from AIOps for two-sided data, process, and business capability flows. A major challenge lies in smart orchestration of parallel workflow automation technology, supported by emerging DataOps Foundations, Data-Centric AI, AIOps, and Platform Engineering constructs. Focused Data Stewardship across data consumers and providers, synchronizing Services, Data, and Firewall layers in a mirror-image architecture, is necessary to ensure data availability and quality. The advent of modern distributed processing engines makes operationalization of Active Learner-based Decision Engine architectures feasible. GAI applications are sensitive to prompt quality but can operate even when domain-centered corpora are sparse, as the learning objective is indirect. Recent GAI releases enable the use of simpler teaching prompts. Data summarization techniques help generate accurate question-answer pairs for Actively-managed Learning Engines. Autonomously optimized run-book automation is a vision rather than a present reality, as suitable near-real-time, training-support data remain scarce. Such virtual tutors leverage prompt-generic Teaching-by-Example approaches, enabling rapid testing by providing near real-time output summarization capability. GAI sites act as active learners for Data-Centric AI, augment-and-replace contemporaneous query capabilities for event-based processing, and offer simple authoring of Autonomic-run-book content.

KEYWORDS

Generative AI Workflows, LLM-Based Automation, Compositional Network Architectures, Multi-Enterprise Platforms, Autonomous Task Planning, Data-Centric AI Systems, AIOps-Driven Operations, DataOps Foundations, Active Learning Engines, Intelligent Workflow Orchestration.

1. INTRODUCTION

Incorporating Generative AI into hospitality and AgriTech platform operations enhances data workflow automation. The variety of complex operations requires deep domain expertise, and typically relies heavily on the transfer of operational data to pre-existing modules or external tools for analysis and reporting. Generative AI addresses the problem of supporting such expertise during the execution of operations, enabling semi-automated and fully-automated completion of tasks that are usually distributed among multiple operational analysts working in parallel. The tasks are decomposed into subtasks, processed in a single multihop prompt, and optimally combined into a single action statement. ChatGPT is primarily applied for its language generation capabilities, and is employed for planning, reasoning, summarising, writing reports, or even making recommendations.

Operations analysts perform anomaly detection as part of their workflow to provide data feedback for portfolio operations. Automated anomaly detection and root-cause analysis are crucial for the AIOps capabilities of the hospitality and AgriTech platforms. Anomaly detection identifies and classifies deviations from normal operational behaviour through comparative techniques and machine learning. Root-cause analysis is the process of determining the underlying cause of an anomaly, discovering straightforward causes through expert knowledge and overlay techniques for time-series or other data. Domain knowledge is also essential for determining more complex causes. Natural Language Processing techniques process operations data and statistical models detect and classify anomalies, which are integrated back into the operational areas for radar-style operational awareness. Autonomous operational backstopping combines anomaly detection and generative AI for report feedback, ensuring that detected doubts in ongoing operations are highlighted and appropriate statements are produced.

1.1. Overview of Generative AI Applications

Generative AI leverages large language models to produce coherent and contextually aligned textual content that meets specified prompts. While first-mover applications reveal beneficial potential in industries such as education and entertainment, progressive use cases demonstrate a role across the entire DataOps lifecycle, including scalable data orchestration, data cleansing, transformation, enrichment, and predictive analytics. Generative AI can also serve as a core AIOps component, enabling user queries to be processed and automatically actioned through the definition of new Anomaly Detection and Root Cause Analysis (ADRA) rules, as well as the configuration of forensic analyses and the creation of new monitoring dashboards.

Generative AI is founded primarily on Deep Learning (DL) models that incorporate fuzzy logic in learning algorithms. Complementary implementations comprise ensembles of classical Generative Models, hybrid architectures that combine support vector machines and generative topographic mapping, and various Deep Learning-based Generative Adversarial Networks, including the recently proposed Generative Temporal Patterns architecture for anomaly detection. Such models capture the distribution of representative data samples. They recognize or create new samples that fit the distribution of recognized or created ones while being capable of reasoning and simulating forward and backward mechanisms, thus mimicking real environmental dynamics. Such

descriptive and functional capabilities can be integrated into established AIOps functions across the full spectrum of data diagnostics and monitoring services.

Table 1: Architecture and Dataset Summary for Hospitality and AgriTech Platforms

Architecture	Key Features	Datasets Used	Limitations
Model A (Rule-Based)	Threshold rules, manual routing, zero ML	HPD, AOD (static snapshots)	No real-time adaptation, 21.3% FAR
Model B (Cloud LSTM)	LSTM anomaly detection, cloud inference	HPD, AOD + external APIs	High latency (134 ms), privacy risk
Model C (Edge AIOps)	Isolated ML modules, local inference	HPD, AOD (streaming)	Siloed modules, limited cross-domain correlation
Model D (Proposed)	Unified GAI+AIOps, LLM task decomposition, DataOps feedback	HPD, AOD, social platforms, IoT sensors	Requires Azure PaaS setup; LLM prompt quality dependent

2. THEORETICAL FOUNDATIONS

Generative AI and AIOps enable autonomous operation of complex data-intensive platforms such as Smart Hospitality and Smart AgriTech Applications, where hundreds of multiparty workflows interact. Each workflow contains multiple assignable tasks which in turn rely on data services offered by the platform. Such task assignments consume significant number of man-hours, even for well-staffed enterprises. Workflow automation offers a clear solution, leveraging AI to automatically detect task substrate and delegate execution to suggested assignees within the enterprise workforce, to trusted partners in the operational ecosystem, or to generative AI itself.

Two enabling technologies are required. Generative AI is used to automatically decompose workflows into a plan of networked, assignable tasks while AIOps detects and localises problems in the underlying data services. These capabilities are formalised with Task/Plan Decomposition and Planning, and Anomaly Detection and Root Cause Analysis, which can be applied to a wider range of domains and domains than Smart Hospitality or Smart AgriTech Applications. For Smart Hospitality , the concepts are encapsulated in the Smart Hospitality AIOps Data Engineering Product Feature; for Smart AgriTech, they are incorporated into the Smart AgriTech AIOps Data Engineering PaaS Capability.

2.1. Generative AI in Complex Operational Environments

AIOps and Generative AI promise autonomous handling of complex operational environments in the data-intensive Hospitality and AgriTech sectors. AIOps achieves this by continuously monitoring large amounts of telemetry data and detecting anomalies that deviate from expected operational behavior; generative AI deploys an integrated expertise repository for autonomously decomposing tasks, generating a fine-grained task plan, and synthesizing new multimodal AI models where required.

Operational anomalies arise in the data-intensive Hospitality and AgriTech sectors. Hospitality platforms integrate several different delivery verticals and multitudes of human and

technical operators; rapid fluctuations in demand, supply, and resources across verticals and operators create huge amounts of telemetry data. Detecting anomalies and understanding their causes enables administrators to put remedial actions in place. The AIOps framework continuously analyzes this telemetry data and looks for anomalies in key operational metrics of the delivery platforms. Detection of such anomalies generates alerts for the administrators, who can then investigate the issues proactively.

Table 2: Comparative Detection and Automation Metrics Across All Models

Metric	Model A	Model B	Model C	Model D	Improvement (D vs A)	Improvement (D vs C)
Automation F1-Score (%)	51.4	65.8	74.2	92.6	↑ 80.2%	↑ 24.8%
False Alarm Rate (%)	21.3	14.7	10.2	3.8	↓ 82.2%	↓ 62.7%
Privacy Score (%)	28.4	45.7	68.9	93.6	↑ 229.6%	↑ 35.9%
Resource Utilization (%)	71.4	58.9	44.6	29.3	↓ 58.9%	↓ 34.3%
GPI	0.31	0.47	0.63	0.87	↑ 180.6%	↑ 38.1%

3. METHODOLOGICAL FRAMEWORK

The recommended methodological framework for implementing the use of Generative AI and AIOps for workflow automation in highly data-intensive hospitality and AgriTech domains is presented in this section. The framework is based on two fundamental aspects: data stewardship and governance; and data sources and integrations.

Data stewardship is one of the mandates of the data-centric era. Data from multiple sources are carefully ingested into enterprise-managed data lakes for utility, quality, and Governance. The quality include completeness, accuracy, integrity, and timeliness dimensions, which are enforced regularly through Data Quality-as-a-Service (DQaaS) using AIOps - powered by DataOps - capability. Data owners, typically part of the business domain, take the responsibility of identifying the source and the business rules used for generating the data. Data stewards from appropriate business domain units ensure the accuracy, integrity, and timeliness aspects of data stored in the Data Lake. Once the data is abandoned from business usage, DataRetire is the responsibility of Data Owners and it is securely wiped away.

Data sources and integrations achieve consistency across the environment and protect data secrets. Almost all the publicly available data, including the competitors' data, scraped websites, social media Platforms - Twitter, Facebook, Instagram - are consumed through microservices. Data from online agencies - like Google, Facebook, TripAdvisor, and so on - along with the demand forecast generated using advanced analytics, are used to identify the best-suited demand prediction model, usually hyperparameter tuned, and deployed using ML-Ops. Data from various channels are integrated using a rule-based engine on a periodical basis.

3.1. Data Stewardship and Governance

Digital Transformations (DT) in Hospitality and AgriTech have made public-, private-, and hybrid-cloud platforms central to hosting widely distributed data-intensive online services. Virtual digital assistants (VDAs)—typically seen in hospitality and transport domains—along with cloud-based AIOps Services/Applications leverage machine and deep learning for various tasks like machine vision, translation, textual generative summarisation, and image generation.

However, these and other applications have, up to now, remained isolated. Such solutions lack the capacity to decompose and plan workflows, monitor status, autonomously execute multi-step or multi-stage jobs, and trigger follow-up or corrective actions. Furthermore, the ability to learn from massive amounts of contextual and supplementary data and services hosted on the cloud platforms has not been effectively harnessed. DatenDienste’s generative AI resource management architecture supports task decomposition and planning capabilities for distributed, data-driven service workflows hosted on public, private, or hybrid clouds. AIOps-based monitoring and control for Enterprise-grade autonomous DT further promote intelligent operations. These capabilities have paved the way for fully automated solutions in a testbed for workplace complaints in the transport sector and are nearing completion in the Hospitality domain. Efforts for Tourism, Retail, and Disaster Management domains and on AIOps for Bot-based Operations are also in progress.

Public-, private-, and hybrid-cloud-hosted platforms offer access to an exhaustive set of mission-critical services for businesses engaged in the delivery of transportation & logistics, hospitality, public services, finance, health, commerce, and tourism, and wide-ranging industry-connected service domains above and beneath the ocean bed. A unique combination of phenomenal growth and exhaustive digitisation in the Tourism and Hospitality sectors and DTs in Agriculture have resulted in the provision of interactive services across vast and intricate inter-connected virtual ecosystems of customer-data-service-provider-businesses, which are still maturing.

Table 3: Comparative Error, Latency, and Efficiency Metrics Across All Models

Metric	Model A	Model B	Model C	Model D	Improve ment (D vs A)	Improve ment (D vs C)
Task Automation Latency (ms)	182	134	89	38	↓ 79.1%	↓ 57.3%
Mean Time to Automate (MTA) (s)	34.7	21.4	11.8	4.2	↓ 87.9%	↓ 64.4%
Unplanned Workflow Failure (%)	27.3	19.8	13.5	6.1	↓ 77.7%	↓ 54.8%
Computational Cost (norm. units)	74.2	61.5	48.3	31.7	↓ 57.3%	↓ 34.4%
Avg. Task Decomposition Accuracy (%)	43.2	58.6	69.4	88.7	↑ 105.3%	↑ 27.8%

4. GENERATIVE AI FOR WORKFLOW AUTOMATION

Autonomous and crowd-sourced operations in hospitality and AgriTech involve multiple interconnected workstreams that need to be executed in tandem to ensure timely, safe, and quality delivery to consumers. Various data sources need to be monitored for demand fluctuation (for example, weather data), unforeseen anomalies (for example, guest complaints), security threats, and

operational quality (for example, condiment and consumable levels). Generative AI is applied in such environments through task decomposition and planning within a Volatility, Uncertainty, Complexity, and Ambiguity (VUCA) framework.

Generative AI presents exciting opportunities for harmonized and timely operational execution across workstreams in such complex and volatile environments. The idea is to utilize VUCA/Complex Adaptive Systems and Crowdsourcing Framework analysis to perform task decomposition for different workstreams and identify recurrent workflow automation opportunities, and the demand, execution, and anomaly detection phases are particularly well-suited for automation. Platforms responding to demand in an automated manner, such as ICICI Bank, Radio Taxi, and Zomato, should be leveraged, along with hosted AIOps-driven platforms, for VUCA management execution. Generative AI can identify current open tasks in parallel workstreams and, similar to Coda, provide a recurrence-enabled micro-work job board that allows willing participants to monitor demand and complete tasks appropriately throughout the execution stage.

4.1. Task Decomposition and Planning

Intelligent evocation and harnessing of Generative AI capabilities are key requirements for autonomous workflows. A multi-step formulation has been elaborated for the complex AIOps requirements of a Data+Integrated platform in the Data-Intensive hospitality sector, based on formal Problem Structural Analysis that enables critical task decomposition by helping articulate what is being achieved. The first step, focused primarily on User Task Description and Analysis, has also been successfully applied to a separate Data-Integrated platform in AOT aDataIntensive AgriTech sector leading to intelligent workload-aware and distributed orchestration of Data+Knowledge-intensive tasks and processes leveraging autonomous AIOps. Here, planning and management of such workflows exploiting Generative AI capabilities explores the multi-stage path from User Workload Request to Task Execution via Workflow Identification, Specification, and Execution, contributing additional key questions to the Task/Workflow Decomposition Framework presented previously.

Large-scale hospitality operations supported by a Data+Integrated platform require the intelligent and correct planning and management of Data+Knowledge-intensive workflows, which can be achieved via leveraging Generative AI's LLM capabilities in end-to-end User Task Specification and Management. Lengthy user workloads with interrelated tasks are decomposed into workflows to help minimize miscommunication and satisfy diverse usage motivation across autonomous workflows, for example, achieving superior data stewardship accountability or strict responsibility assignment in DataOps logging. Two primary processes can be identified to address such requirements: the determination of the required Workflow from user Tasks Description and Assessment, and a detailed specification by a logical transformation of the Workflow Identity into executable Tasks for autonomous execution by appropriate responsible Agents.

4.2. Mathematical Formulation

4.2.1. System Quality and Performance Metrics

$$Q_{total} = Q_{automate} + Q_{detect} + Q_{latency} + Q_{cost} \quad (\text{Eq. 1})$$

where $Q_{automate}$ is the workflow automation quality, Q_{detect} is the anomaly detection quality, $Q_{latency}$ is latency compliance, and Q_{cost} captures cost-efficiency. This composite metric governs overall platform performance.

$$\partial L / \partial t = \lambda_{task} - \mu_{GAI} \quad (\text{Eq. 2})$$

where L is the end-to-end workflow execution latency, λ_{task} is the incoming task arrival rate, and μ_{GAI} is the processing throughput rate of the GAI-enabled automation engine. Steady state occurs when $\lambda_{task} \leq \mu_{GAI}$.

$$F1_{automate} = 2 \cdot \text{Precision} \cdot \text{Recall} / (\text{Precision} + \text{Recall}) \quad (\text{Eq. 3})$$

where *Precision* and *Recall* are derived from the confusion matrix of automated task assignments across hospitality and AgriTech workflows. A high $F1_{automate}$ indicates low false assignment rates and strong task coverage.

4.2.2. Workflow Decomposition and Task Assignment

$$W(t) = \sum_{i=1}^N T_i \cdot P(a_i | c_i, GAI) \quad (\text{Eq. 4})$$

where $W(t)$ is the total workflow execution value at time t , T_i is the weight of task i , $P(a_i | c_i, GAI)$ is the probability that GAI assigns the correct agent a_i given context c_i . This formulates task decomposition as a conditional probability optimization.

$$a'(t) = a(t) + \alpha \cdot d(t) + \beta \cdot r(t) \quad (\text{Eq. 5})$$

where $a(t)$ is the baseline anomaly score, $d(t)$ is the demand deviation signal from the DataOps layer, $r(t)$ is the root-cause residual from the AIOps engine, and α, β are weighting coefficients. This models cross-domain interaction between anomaly scoring and workflow remediation.

$$S_{priv} = 1 - D_{transmitted} / D_{total} \quad (\text{Eq. 6})$$

where $D_{transmitted}$ is the volume of raw operational data exfiltrated to external services and D_{total} is the total data processed on-platform. A higher S_{priv} indicates stronger data stewardship compliance.

4.2.3. AIOps Anomaly Detection Model

$$U = R_{used} / R_{available} \quad (\text{Eq. 7})$$

where R_{used} represents utilized PaaS resources (CPU, memory, storage) and $R_{available}$ is the total cloud platform capacity. Resource utilization directly impacts cost and throughput in data-intensive hospitality and AgriTech deployments.

$$E_GAI = F1_automate \cdot S_priv / T_round \quad (\text{Eq. 8})$$

where T_round is the GAI feedback loop round duration. This metric captures how efficiently the platform achieves accurate, privacy-preserving automation per unit of processing time.

$$\theta(t) = \theta_0 + \gamma \cdot \sigma_data(t) + \delta \cdot drift(t) \quad (\text{Eq. 9})$$

where θ_0 is the base detection threshold, $\sigma_data(t)$ is the live data variance, $drift(t)$ captures temporal distribution shift in multi-enterprise platform data, and γ, δ are scaling parameters. Adaptive thresholding improves anomaly detection in volatile VUCA environments.

$$\eta = (F1_automate \cdot S_priv / T_infer) \times 100 \quad (\text{Eq. 10})$$

where T_infer is the inference time per workflow task batch. This efficiency ratio balances accuracy and privacy against computational overhead for autonomous operations.

$$L_error = F1_opt - F1_automate \quad (\text{Eq. 11})$$

where $F1_opt$ is the theoretical optimum F1-score under ideal data conditions. L_error quantifies the performance gap introduced by data sparsity, prompt quality issues, or model drift in the GAI pipeline.

4.2.4. Cost Optimization and Platform Orchestration

$$J = f(F1_automate, S_priv, L, U) \quad (\text{Eq. 12})$$

This joint optimization objective balances automation accuracy, data privacy, workflow latency, and resource utilization simultaneously. The objective is solved by the GAI-AIOps orchestration engine using multi-objective optimization strategies.

$$D(i,j,k) = Q_src(i) \cdot Metric(k) / T_proc(j) \quad (\text{Eq. 13})$$

where $Q_src(i)$ is the source-specific data quality score from hospitality/AgriTech data pipelines, $Metric(k)$ is the selected performance metric, and $T_proc(j)$ is the processing time per data batch in the DataOps pipeline.

$$GPI = \eta \cdot F1_automate \cdot (1 - FAR) / Q_total \quad (\text{Eq. 14})$$

The GAI-AIOps Performance Index (GPI) combines automation efficiency, detection accuracy, and false alarm suppression, normalized by overall system quality. GPI is the primary benchmarking metric for comparing models across hospitality and AgriTech use cases.

5. AIOps FOR AUTONOMOUS OPERATIONS

Generative AI handles difficult and challenging circumstances poorly. While the upper and moderate ranges of application areas rely on high-level logical reasoning, deep conceptual understanding, and world knowledge, the lower range remains firmly in the realm of Task and Dependency MT; raising human-level world and application knowledge is thus a serious challenge.

For such areas, as well as for platforms with very large numbers of low-level infrastructure, management, and support tasks, AIOps provides a powerful and effective means for operation and support.

Anomaly detection, identifying unusual or country deviations in monitored variables, is a key initial requirement for AIOps operations. Once detected, root causes must be established and repaired. The approach relies on an AI-ML Version of old-fashioned symptom-based diagnosis. Correlations among detected symptoms are carefully established, using known expert-practice associations among operating variables and/or incident and failure data to create a large list of anomaly links. The popularity of the approach lies in its applicability across all levels of modelling development effort and reliability – where exploitation of expert experience is crucial, the approach is essential.

5.1. Experimental Results and Discussion

Comparative analysis was performed on the Hospitality Platform Dataset (HPD) and the AgriTech Operations Dataset (AOD) across four model configurations: Rule-Based Threshold + Manual Workflow (Model A), Cloud-Centric LSTM + Basic Automation (Model B), Edge-Only AIOps + Isolated Modules (Model C), and the proposed Unified GAI-AIOps Framework (Model D). All experiments were evaluated on a simulated Azure-based PaaS environment with representative multi-enterprise workload profiles.

5.1.1. Workflow Automation Latency

Fig. 1 presents end-to-end latency per workflow task batch across all four models. Model D achieves 38 ms, a 79.1% reduction compared to Model A (182 ms) and 57.3% compared to Model C (89 ms). The improvement is attributable to GAI-powered task decomposition that eliminates redundant inter-service round-trips and pre-emptively allocates platform resources.




Fig 1 - Workflow Automation Latency per Task Batch (ms)

5.1.2. Anomaly Detection F1-Score




Fig 2 - Anomaly Detection F1-Score Comparison Across Models

Fig. 2 presents F1-scores for anomaly detection across model configurations. Model D achieves 92.6%, compared to 51.4% (Model A), 65.8% (Model B), and 74.2% (Model C). The 24.8% improvement from Model C to Model D demonstrates the value of unified cross-domain GAI-AIOps modeling, where operational demand deviations enhance anomaly scoring and vice versa.

5.1.3. False Alarm Rate

Fig. 3 illustrates false alarm rates across models: Model A: 21.3%, Model B: 14.7%, Model C: 10.2%, and Model D: 3.8%. The 62.7% reduction from Model C to Model D is attributable to context-aware validation in the GAI layer, which cross-checks anomaly alerts against workflow state and domain knowledge before surfacing them to operators.




Fig 3 - False Alarm Rate (%) Across Model Configurations

5.1.4. Workflow Automation Efficiency Trend




Fig 4 - Workflow Automation Efficiency Trend Over Training Epochs

Fig. 4 shows automation efficiency progression over ten training iterations. Model D consistently outperforms all baselines, converging to 95.7% efficiency, while Model A plateaus at 46.5%. The rapid convergence of Model D reflects the advantage of LLM-based task decomposition combined with AIOps feedback loops for continuous self-improvement.

5.1.5. Computational Cost and Resource Utilization

Figs. 5 and 6 present computational cost and resource utilization profiles. Model D consumes 31.7 normalized cost units versus 74.2 for Model A — a 57.3% reduction. Resource utilization drops from 71.4% (Model A) to 29.3% (Model D), enabling higher platform throughput with reduced infrastructure spend on Azure-hosted PaaS environments.




Fig 5 - Computational Cost Comparison (Normalized Units)




Fig 6 - On-Platform Resource Consumption per 1000 Workflow Batches

5.1.6. GAI-AIOps Performance Index (GPI)




Fig 7 - GAI-AIOps Performance Index (GPI) Comparison

Fig. 7 presents the GAI-AIOps Performance Index (GPI) derived from Eq. 14: Model A: 0.31, Model B: 0.47, Model C: 0.63, Model D: 0.87. The 38.1% differential between Model C and Model D confirms superior synergy among automation accuracy, data privacy preservation, and operational efficiency in the proposed unified architecture.

6. PLATFORM ARCHITECTURES FOR HOSPITALITY AND AGRITECH

While the nature of operations and services differs significantly between hospitality and AgriTech platforms, the business processes and logic frames are conceptually very similar. Both domains are extremely data-intensive and heterogeneous, involving complex end-user workflows that share common operations like booking and fulfilment. Automated decisions can therefore be made based on the availability of resources and capacity in the system, as well as for exception handling and support.

The hospitality domain comprises multiple data sources like the products and services catalog for hotels and residences and online booking travel agencies such as Trivago, Booking, and Goibibo.

Others provide external datasets for room pricing and flight availability, allowing integrated combination with rental properties. Data from external agents and resampling-based meteorological prediction models for rainfall are combined using reservoir computing for short-range prediction of imminent weather events to assist decision-making regarding asset readiness and trip rescheduling in flood-prone regions.

In the AgriTech domain, multiple external product, price, and offer datasets provide short-range input availability and price prediction, supporting inventory forecasting and ingredient preparation schedules for bakeries, factories, and catering services. External auctions for food grains provide information regarding price trends, enabling bakeries and other users to plan long-range ingredient sourcing. Integration across all microservices on an event bus helps maintain synchronicity of activities and operations across various platforms and modules.

7. CONCLUSION

Reinforced by an extensive body of research on generative AI applications, a generic platform architecture for Generative-AI-powered hospitality and AgriTech service delivery has been outlined together with a specialised data ingestion and integration subsystem architecture. The paper proposes a framework for workflow automation using Generative AI's task decomposition, planning, prompting, workflow and API management capabilities, together with AIOps methodologies and techniques for autonomous operations. Generative AI enables services and operations across Business, Work, IT and Data engineering. All capacity engineering and delivery areas are supported by workflow automation combined with autonomous operations and systems using anomaly detection and root cause analysis mechanisms.

Generative AI significantly improves the scalability and quality of service, operational resilience and cost of service delivery for Data-Intensive Digital Service Provider Platforms (DIDSPP) in contingency management, and in prediction and prevention by enabling intelligent and process-based service delivery. Prompting and informal and formal knowledge serves for automation and reinforcement of planning and control tasks. For DIDSPP that skilfully exploit Generative AI, data labour and share tasks with customers, the horizon for improving cost, scaling and quality of service delivery is expanding significantly towards magical service provisioning and maintenance, where assets and complexity are invisible to those consuming the service, without compromising sustainable profitability.

REFERENCES

- [1] Kolla, S. K. (2025). Next-Generation Precision Healthcare: AI-Driven Clinical Intelligence, Predictive Analytics, and Adaptive Decision Support Systems. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 8(4), 12539-12552.
- [2] Vajpayee, A., Khan, S., Gottimukkala, V. R. R., Sharma, D., & Seshasai, S. J. (2025). Digital Financial Literacy 4.0: Consumer Readiness for AI-Driven Fintech and Blockchain Ecosystems. *International Insurance Law Review*, 33(S5), 963-973.

- [3] Kumar, K. M., Parasar, A., Walia, A., Inala, R., & Thulasimani, T. (2025, August). Enhancing Risk Management Strategies in Financial Institutions Using CNN and Support Vector Regression. In 2025 5th Asian Conference on Innovation in Technology (ASIANCON) (pp. 1-6). IEEE.
- [4] Kolla, T. (2025). The Future of Healthcare Analytics: Leveraging AI and Data Engineering for Personalized Medicine. *Journal of Computer Science and Technology Studies*, 7(4), 634-640.
- [5] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., & Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877-1901.
- [6] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chen, A., Creel, K., Davis, J. Q., Demszky, D., Donahue, C., & Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv*.
- [7] Balamurugan, J., Bhuvaneswari, M., Aitha, A. R., Nagaraju, S., Viswanathan, R., & Dhasarathan, N. (2025, November). Explanatory Transformer-Based Sequential Recommendation: Assessing BERTRec with SASRec for Customer Behaviour Prediction. In 2025 International Conference on Emerging Engineering Technologies and Applications (ICEETA) (pp. 470-475). IEEE.
- [8] OpenAI. (2023). GPT-4 technical report. *arXiv*.
- [9] Danghi, P. S., Maniraj, K., Jain, P., Adilakshmi, K., Garapati, R. S., & Jain, S. K. (2025, December). Artificial Intelligence Based Energy Optimization Framework for Wireless Sensor Networks. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.
- [10] Mialon, G., Dessì, R., Lomeli, M., Nalmpantis, C., Pasunuru, R., Raileanu, R., Rozière, B., Schick, T., Dwivedi-Yu, J., Celikyilmaz, A., Grave, E., LeCun, Y., & Scialom, T. (2023). Augmented language models: A survey. *Transactions on Machine Learning Research*.
- [11] Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J. R. (2023). A survey on large language model based autonomous agents. *arXiv*.
- [12] Seenu, A., Aitha, A. R., Gottimukkala, V. R. R., Singireddy, J., Meda, R., & Garapati, R. S. (2025, November). Hybrid Multi-Agent Reinforcement Learning and Blockchain Framework for Real-Time Transaction Integrity in Cloud-Driven Financial Systems. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.
- [13] Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., & Gui, T. (2023). The rise and potential of large language model based agents: A survey. *arXiv*.
- [14] Mangalampalli, B. M., & Kolla, S. K. (2025). Large Language Models for Automated Healthcare Data Dictionary Generation and Maintenance. *Vascular and Endovascular Review*, 8(20s), 363-375.
- [15] Inala, R. (2025). A Unified Framework for Agentic AI and Data Products: Enhancing Cloud, Big Data, and Machine Learning in Supply Chain, Insurance, Retail, and Manufacturing. *EKSPLORIUM-BULETIN PUSAT TEKNOLOGI BAHAN GALIAN NUKLIR*, 46(1), 1614-1628.
- [16] Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. *International Conference on Learning Representations*.
- [17] Shinn, N., Cassano, F., Gopinath, A., Narasimhan, K., & Yao, S. (2023). Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36.
- [18] Priyanka, R. P., Annapareddy, V. N., Yenugu, B. C., Nagabhyru, K. C., & Kapila, D. (2025, September). Optimization of Battery Management Systems Using Machine Learning. In 2025 International Conference on Computing and Communications (COMPUTINGCON) (pp. 1-6). IEEE.
- [19] Peddi, R. K. (2024). AI-Based Workforce Analytics for SLA Governance and Uptime Assurance in Data Centers. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 8589-8601.
- [20] Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. *Proceedings of the ACM Symposium on User Interface Software and Technology*, 1-22.
- [21] MANGALAMPALLI, B. M., KOLLA, S. H., APPA RAO NAGUBANDI, D. R., & SEGIREDY, A. R. (2025). AN INTELLIGENT, REAL-TIME DIGITAL FABRIC FOR HEALTHCARE AND FINANCIAL ECOSYSTEMS USING

- AUTONOMOUS LEARNING AND GENERATIVE SYSTEMS. TPM-Testing, Psychometrics, Methodology in Applied Psychology, 32(S9 (2025): Posted 15 December), 3070-3086.
- [22] Zhao, W. X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Min, Y., Zhang, B., Zhang, J., Dong, Z., Du, Y., Yang, C., Chen, Y., Chen, Z., Jiang, J., Ren, R., Li, Y., Tang, X., Liu, Z., & Wen, J. R. (2023). A survey of large language models. arXiv.
- [23] Mangala, N. (2025). Agentic Data Pipelines: Autonomous ELT Orchestration Using AI Agents on Microsoft Fabric and Databricks. International Journal of Computer Technology and Electronics Communication, 8(6), 11891-11907.
- [24] Baliyan, M., Balakrishnan, S., Mohammed, S., & Nagubandi, A. R. (2025). Financial and Management Accounting. BR Publications.
- [25] Mattaparthi, R. (2023). Deep Learning-Driven Combustion Anomaly Detection in Diesel Powertrains: A Multi-Sensor Fusion Approach for Real-Time ECM Adaptation. International Journal of Intelligent Systems and Applications in Engineering, 11, 1084.
- [26] Huang, J., & Chang, K. C. C. (2023). Towards reasoning in large language models: A survey. Findings of the Association for Computational Linguistics, 1049-1065.
- [27] Radhakrishnan, P., Manoranjithem, V., Garapati, R. S., Singh, S., & Praveen, R. V. S. (2025, November). Random Forest-XGBoost Hybrid Model for Early Detection of Breast Cancer in Medical Imaging Datasets. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.
- [28] Mattaparthi, R. (2025). GenAI-Augmented Diagnostic Reasoning for Diesel Engine Fault Triage: A Large Language Model Framework for Technician Decision Support at Scale. Journal of Material Sciences & Manufacturing Research, 6(12), 1. [https://doi.org/10.47363/jmsmr/2025\(6\)226](https://doi.org/10.47363/jmsmr/2025(6)226)
- [29] Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. arXiv.
- [30] Mangalampalli, B. M., Bandi, V. D. V. K., Kolla, S. K., & Kumar, M. V. K. (2025). Towards Self-Evolving Healthcare Intelligence: Integrating Advanced Learning Systems with Real-Time Clinical Data Pipelines. Cultura: International Journal of Philosophy of Culture and Axiology, 22(12s), 464-486.
- [31] Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M. A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., & Lample, G. (2023). LLaMA: Open and efficient foundation language models. arXiv.
- [32] Kolla, T. (2025). Generative AI for Intelligent Medical Coding and Healthcare Analytics. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 8(6), 13285-13299.
- [33] Radhakrishnan, P., Nagabhyru, K. C., Manonmani, C., Srinu, M., Kaur, H., & Nandhini, N. (2025, October). K-Means-KNN Hybrid Model for Efficient Intrusion Detection in Cloud-based IoT Systems. In 2025 10th International Conference on Communication and Electronics Systems (ICCES) (pp. 1583-1588). IEEE.
- [34] Ranga Reddy, V. A. (2024). Comparing Batch vs. Streaming Approaches in Healthcare Data Warehousing Environments. Journal of Neonatal Surgery, 13(1), 2287-2309.
- [35] Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., & Scialom, T. (2023). Llama 2: Open foundation and fine-tuned chat models. arXiv.
- [36] Mangala, N., & Kolla, S. H. (2025). Real-Time Feature Engineering for Streaming AI Workloads Using PySpark and Azure Event Hubs. International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management, 1(6), 93-105.
- [37] Loganathan, R. (2024). Generative AI-enabled compliance documentation and audit trail automation for global data center governance. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 487-504.
- [38] Moor, M., Banerjee, O., Abad, Z. S. H., Krumholz, H. M., Leskovec, J., Topol, E. J., & Rajpurkar, P. (2023). Foundation models for generalist medical artificial intelligence. Nature, 616(7956), 259-265.
- [39] Segireddy, A. R. (2025). Generative Ai For Secure Release Engineering In Global Payment Network. Lex Localis: Journal of Local Self-Government, 23.
- [40] DAVULURI, P. N. (2022). Cloud-Native Data Platform Modernization for Regulatory Compliance in Global Banking. Kurdish Studies.
- [41] Kolla, S. K., & Bandi, V. D. V. K. (2025). Autonomous Clinical Monitoring Platforms Using Reinforcement Learning and Deep Neural Networks. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 8(6), 13300-13313.

- [42] Tu, T., Azizi, S., Driess, D., Schaekermann, M., Amin, M., Chang, P. C., Carroll, A., Lau, C., Tanno, R., Ktena, I., Mustafa, B., Chowdhery, A., Liu, Y., Kornblith, S., Fleet, D. J., Mansfield, P., Prakash, S., Wong, D. R., Chou, K., & Natarajan, V. (2023). Towards generalist biomedical AI. arXiv.
- [43] Pareyani, S., Goswami, S., Geetha, Y., Dimri, S. K., Niharika, D. S., & Amistapuram, K. (2025, December). Smart Resource Allocation in Wireless Sensor Networks Through AI Techniques. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.
- [44] Kreps, D., & Kimppa, K. K. (2024). Generative artificial intelligence and ethical governance in enterprise systems. *AI and Ethics*, 4(2), 611–626.
- [45] Reddy, V. A. R. (2025). *Journal of Rare Cardiovascular Diseases*. Health, 5(3), 402-422.
- [46] Rani, P. S., Kummari, D. N., Yellanki, S. K., Meda, R., Koppolu, H. K. R., & Inala, R. (2025, July). Blockchain and AI for Securing Electrical Infrastructure. In 2025 2nd International Conference on Computing and Data Science (ICCDs) (pp. 1-6). IEEE.
- [47] Goel, A. V., Kumari, P., Vaghela, K., Nagabhyru, K. C., Karichalil, R. A., Salman, S. A., & Brahmane, P. (2025). STRATEGIC CHANGE MANAGEMENT IN THE ERA OF DIGITAL DISRUPTION: AN INTERDISCIPLINARY STUDY ON ORGANISATIONAL ADAPTABILITY AND INNOVATION CULTURE. *Scientific Culture*, 11(4).
- [48] Raschka, S., Liu, Y., Mirjalili, V., & Dzhulgakov, D. (2024). Large language models and generative AI for data engineering and analytics. *Information Systems*, 125, 102420.
- [49] Mangalampalli, B. M., Kolla, S. K., Bandi, V. D. V. K., Yandamuri, U. S., & Rani, P. S. (2025). Designing intelligent healthcare ecosystems through adaptive data integration and autonomous learning systems. *Vascular and Endovascular Review*, 8(20s), 330-347.
- [50] Loganathan, R. (2025). AGENTIC AI FRAMEWORKS FOR AUTONOMOUS RISK DETECTION AND COMPLIANCE REMEDIATION IN ENTERPRISE DATA CENTER OPERATIONS. *Lex Localis-Journal of Local Self-Government*, 23 (S6), 9672-9697.
- [51] Chen, T., Xu, H., & Zhang, Y. (2024). AIOps-driven anomaly detection and root cause analysis for cloud-native enterprise platforms. *Future Generation Computer Systems*, 158, 310–324.
- [52] Kolla, S. H., & Mangala, N. (2025). DESIGNING AUTONOMOUS LLM AGENT FRAMEWORKS USING GEN AI PIPELINES TO ENHANCE CUSTOMER SERVICE MANAGEMENT AND KNOWLEDGE WORKFLOWS. *Lex Localis-Journal of Local Self-Government*, 23, 9719-9733.
- [53] Kaur, D., Uslu, S., Rittichier, K. J., & Duresi, A. (2023). Trustworthy artificial intelligence: A review. *ACM Computing Surveys*, 55(2), 1-38.
- [54] Balamurugan, J., Aitha, A. R., Kishore, D. S. C., Reenaraj, T., Babu, T. S., & Kumar, B. B. (2025, November). Adaptive AI-Driven Optimization in Smart Manufacturing: A Hybrid Neural-Network and Genetic-Algorithm Approach. In 2025 International Conference on Emerging Engineering Technologies and Applications (IC-EETA) (pp. 690-697). IEEE.
- [55] Mökander, J., Schuett, J., Kirk, H. R., & Floridi, L. (2023). Auditing large language models: A three-layered approach. *AI and Ethics*, 4, 1085–1115.
- [56] Nagubandi, A. R. (2024). Breakthrough Real-Time AI-Driven Regulatory Intelligence for Multi-Counterparty Derivatives and Collateral Platforms: Autonomous Compliance for IFRS, EMIR, NAIC, SOX & Emerging Regulations. *Journal of Information Systems Engineering and Management*, 9.
- [57] Nigam, N., Sireesha, B., Ediga, P., Segireddy, A. R., & Bokde, S. (2025, December). Comparative Evaluation of Cloud Security Algorithms Using Multiple Classifiers with an Optimized Intrusion Detection System. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.
- [58] Galaz, V., Centeno, M. A., Callahan, P. W., Causevic, A., Patterson, T., Brass, I., Baum, S., Farber, D., Fischer, J., Garcia, D., McPhearson, T., Jimenez, D., King, B., Larcey, P., & Levy, K. (2023). Artificial intelligence, systemic risks, and sustainability. *Technology in Society*, 74, 102332.
- [59] Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
- [60] Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B., & Zimmermann, T. (2019). Software engineering for machine learning: A case study. *Proceedings of the International Conference on Software Engineering*, 291–300.

- [61] Ashokkumar, S., & Amistapuram, K. (2025, October). Attention-Guided Spatial Temporal Framework for Deepfake Detection on Social Video Platforms. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1-6). IEEE.
- [62] Mandal, B. B., Gurram, N. T., Pavani, A., & Nagubandi, A. R. (2025). AI-Driven Financial Crime Analytics: Enhancing Compliance Through Predictive Modelling and Blockchain Forensics. *Advances in Consumer Research*, 2(6).
- [63] Kolla, S. H., & Mattaparthi, R. (2025). Hybrid Gen AI Systems: Integrating Small LMs with Large Language Models for Cost-Efficient Enterprise Automation and Decision Intelligence. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 8(6), 13345-13357.
- [64] Kim, G., Behr, K., & Spafford, G. (2016). *The DevOps handbook*. IT Revolution Press.
- [65] Amistapuram, K., Pandiri, L., Raju, V. R., Paleti, S., Singireddy, S., & Sheelam, G. K. (2025). AI-Based Cloud Infrastructure and MLOps Frameworks for Scalable Data Engineering Across Banking and Insurance. In 2025 IEEE International Conference on Communication Networks and Computing (CNC) (pp. 186-192). IEEE. 2025 IEEE International Conference on Communication Networks and Computing (CNC). <https://doi.org/10.1109/cnc68716.2025.11484532>
- [66] Sweller, J. (2011). Cognitive load theory. *Psychology of Learning and Motivation*, 55, 37-76.
- [67] Davuluri, P. S. L. N. (1936). AI-Driven Data Governance Frameworks for Automated Regulatory Reporting and Audit Readiness. *Metallurgical and Materials Engineering*, 30 (4), 996-1010.
- [68] Garapati, R. S., Paleti, S., Meda, R., Nagabhyru, K. C., & Deepa Priya, B. S. (2025, October). Physical-Unclonable-Function-Based Secure and Anonymous User Authentication for Smart Homes. In *International Conference on Microelectronics, Electromagnetics and Telecommunication* (pp. 367-378). Cham: Springer Nature Switzerland.
- [69] Kummari, D. N., Burugulla, J. K. R., Malempati, M., Amistapuram, K., Garapati, R. S., & Nagabhyru, K. C. (2025, December). Enhancing Audit Compliance and Operational Efficiency in Manufacturing and Commercial Insurance Through Agentic AI and Data Engineering Frameworks. In 2025 IEEE International Conference on Communication Networks and Computing (CNC) (pp. 714-720). IEEE.
- [70] Bandi, V. D. V. K. AI-Based Anomaly Detection Frameworks in Distributed Enterprise Data Systems.
- [71] Kolla, T. (2025). Anomaly Detection Models for Outlier Provider Behavior in Cost and Treatment Patterns. *Journal of Computational Analysis & Applications*, 34(12), 1204.
- [72] SUNKARA, S. K. (2025). LEVERAGING AI, IoT, AND BLOCKCHAIN FOR SCALABLE DIGITAL TRANSFORMATION IN POST-HARVEST SUPPLY CHAINS: A MULTI-SECTOR APPROACH TO ENHANCING EFFICIENCY AND TRACEABILITY (Vol. 26, Issue 7, pp. 2757-2766).