International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

Designing Agentic AI Data Products for Personalized Financial Services in Group Insurance and Retirement

Shashikala Valiki

Independent Researcher, Canada

shashikala.valiki.researcher@gmail.com

Manon Dubois

Junior Research Associate, Institut Académique de Grenoble, France

manondubioss@outlook.com

ABSTRACT: Personalisation encompasses products and services delivered to an individual at the right time and in the right form considering individual preferences. Agentic artificial intelligence (AI) provides businesses with the capability to collect and analyse their Adam Smiths and business models' data to realise mass personalisation economically while ensuring the availability of the right soundtrack for the moment. Government authorities and regulators are also interested in discovering customer-centric and personalisation-friendly approaches that do not compromise the requirements of regulatory compliance and risk management. Financial institutions might find it challenging to successfully operate in group insurance or retirement services. Addressing these products might resemble operating in almost any other automated manufacturing or assembly-type business, where AI automated model development provides the capability to deliver customer-centric and personalised services. Nevertheless, these businesses have yet to leverage agentic AI to satisfy end customers' individual needs and preferences.

Agentic AI, Data, Modelling, AI Capability Building, Business Mitigation and Control, customers' dream personas, models for needs, order cycles, product dynamics, customer connectedness and cyclical view, real-time need state, motivation, and intent prediction, service and channel orchestration, mass-media customer service, and AI-generated soundtrack are the capabilities required for data products that deliver services at the intersection of individuals' individual needs, wants, values, character, control, society, surroundings, nature, dreams, beliefs, and purpose. Customers' subconscious personas and business ecosystem-centric model risk management will ensure that factors such as bias and prejudice, which have marred the development and deployment of AI, not only cease to be a concern but actually help realise 'universe', 'societal being', or 'university' objectives.

KEYWORDS: Agentic AI, personalisation, group insurance, retirement services, data governance, compliance, ethics, Agentic AI Orchestration, Personalised Financial Intelligence, Autonomous Decision Agents, Group Insurance Analytics, Retirement Outcomes Optimisation, AI-Driven Data Product Architecture, Context-Aware Financial Services, Human-in-the-Loop Governance, Explainable Financial AI (XAI), Lifecycle-Based Personalisation.

I. INTRODUCTION

For several years, many industries have experienced unprecedented, often volatile, changes driven by events such as the pandemic, supply chain challenges, and widening financial gaps between those perceived to be at risk and those with greater-than-average savings. Similar developments are felt in group insurance markets, with increased customer expectations veering towards even greater customisation. Group insurance or retirement solutions often address needs at a macro level. As financial products become more embedded in customers' lifestyles, the demand for additional and neighbouring financial products and services will also expand. Financial cycles will also affect pensions at some point, and consequently the retirement years, with the levels of income being questioned.

The number of insurance and fund administration companies investing in data-led AI product enablement projects is unequivocal. Agentic AI allows companies to transform data into information that drives optimised action and product. Such products act as data agents, offering services to their users and modelling customer behaviour to continually

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

optimise these services. Data-driven AI products can successfully deliver the same features as today's traditional products when presented with sufficient clean, quality, and timely data. However, very few organisations concentrate on specific deployment architectures. In addition, any product – data-driven or traditional – can be deceptive and deliver unintended bias and unfairness if not properly placed in any customer-facing channel.

1.1. Background and Significance: With increasing availability of data and advances in artificial intelligence (AI), financial organisations are developing personalised products and services using data-driven approaches. Despite being less common than in consumer services, personalisation is gaining interest in group insurance and retirement domains. A wealth of administrative data accumulated over decades offers the potential to achieve real-time personalisation and individual-level customisation for group subscribers and their employees. Such services can enhance customer experience, enable life-stage-targeting of marketing campaigns that reduce lapse risk, and produce higher engagement levels.

To date, several uses of data-driven personalisation in group insurance and retirement domains have been identified, including product recommendation, customer profiling/segmentation, targeted marketing, life-stage-appropriate communications, provision of risk insights and nudges, delivery of advice, proposal automation, AI-based claim assessment and settlement, early warning for lapse risk, and early alerts to sponsors. Real-time data integration can support some of these uses, and an orchestration capability that produces near real-time delivery of services at the individual level can enable others. However, appropriate underlying data foundation, compliance mechanisms, development infrastructure, and agentic artificial intelligence (AI) capabilities remain absent from discussions. These aspects of architecting agentic AI-driven data products for personalised financial services in group insurance and retirement domains are addressed here.

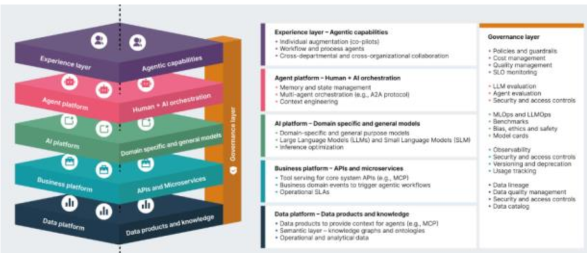

Fig 1: The Agentic Imperative The future for banking and financial services with Agentic AI

1.2. Research design: A conceptual research approach supports the investigation, focusing on agentic AI-driven data products facilitating the delivery of personalised financial services in group insurance and retirement domains. Analysis concentrates on the conditions that facilitate the implementation of agentic data products designed to deliver personalised group insurance and retirement services. The discussion relies on primary and secondary data and knowledge from the agentic AI, financial services, data management and engineering, and marketing literatures.

The analysis identifies data architectures suitable for personalising group insurance and retirement services and explains how agentic capabilities support the scoped personalisation. Governance, risk, and compliance strategies for agentic data products are also considered. The regulatory landscape for group insurance and retirement services is examined, and the implications for model risk management and auditability are explored. Emerging trends affecting the deployment of agentic data products for personalising group insurance and retirement services are discussed.

II. CONCEPTUAL FOUNDATIONS OF AGENTIC AI IN FINANCIAL SERVICES

A deep understanding of agentic AI and personalisation is crucial for architecting agentic AI-driven data products that enable the deployment of personalised financial services in the group insurance and retirement domains, which are characterised by the inherent social need to provide financial support to vulnerable individuals at affordable prices.

Agentic AI: Definitions and Theoretical Underpinnings. Agentic services rely on the concept of agentic AI, which empowers data products to perform complex user tasks or services without requiring constant interaction with the end user. An agentic data product is created by designing a reusable, modular data pipeline and developing the necessary infrastructure to support the required service in a cost-effective way. When products are built as reproducible pipelines,

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

the operational cost becomes a marginal factor compared to the cost of the user tasks for which the product enables higher efficiency and effectiveness. Other recent studies offer additional theoretical underpinnings related to the definition of agentic AI and its implications for data products in the financial services domain.

2.1. Agentic AI: Definitions and Theoretical Underpinnings: The discussion concerning personalisation of financial services revolves, among others, around the capabilities of agentic AI. However, these capabilities are not always defined clearly and, therefore, their implications for individual services remain vague. A two-part definition attempts to fill this gap, followed by a theoretical exploration of its implications for financial services. The first part of the definition is based on the de Silva et al. definition, research on data privacy preferences, and compliance with regulations: agentic AI acts on behalf of others within an established trust relationship and following the individual "terms of service", which can be based on explicit consent but at least include legally mandated consequences and restrictions imposed by the governing regulation. The second part is based on the agentic nature of personalisation: agentic AI offers predictive plus additional services, which might include decision-making, nudging, recommendation, and steering. In structured domains, steering moves the customer closer to choices with substantial negative consequences while respecting the "spirit of consent" even when explicit consent is absent. The only, major exception is illegal activity, which has no ground in the "spirit of consent".

Recent advances in auto summarisation address the need of the consumer for easy-to-understand information about a data product, not only for initial consent but also during its life cycle. They are particularly valuable for agentic AI, due to its predictive plus additional nature. The auto summarisation should only concern text and be language-neutral. Reproducible text generation answers the evolving demand from consumers for personalised content and can also be valuable for informative matching and steering. Automated discovery of hidden knowledge through advanced techniques and early adoption of innovation can foster an agent-expected trust relationship between consumers and agentic AI-based products.

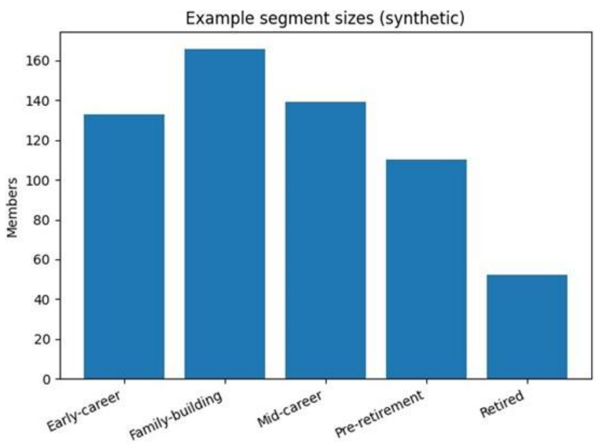

2.2. Personalisation in Group Insurance and Retirement Services: Research indicates that consumers express a strong desire for companies to provide tailored experiences. Outside the group space, 76 percent report that they like the ability to order products tailored to their needs, and 66 percent find it appealing to have their financial products and services provide advice tailored to their situation and life stage. In developing a model for agentic AI consumer digital products in the group insurance and retirement domains it is critical to understand the personalisation ambition. Such an understanding both shapes the model structure and influences the product characteristics.

When discussing personalisation, it is useful to consider similarity to a retail wealth management model. Retail wealth management products group personalisation options into three categories: tailored products (products that are custom-made, such as Hong Kong Disneyland memberships); recommend-to-order products (ready-to-order products that are recommended to consumers based on their life cycle, gamers' preference, family financial ability, financial habit, etc.); and tailored-services products (services that consumers want to be tailored, rather than products). Apply a similar categorisation to an existing group insurance use case—the family protection gap work. This addresses important social issues and has government support. The three-product categories emerge: the "one-for-all" offering; the "healthcare mode" offering; and the "family protection gap" offering. These may not be exhaustive.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

Equation 1: Formal model of an "agentic personalised service" data product

Let a customer (or group member) at time t have a feature/state vector:

$$\mathbf{x}_t = \begin{bmatrix} \text{demographics} \\ \text{plan attributes} \\ \text{transactions/behaviour} \\ \text{life stage} \\ \text{channel context} \\ \text{consent flags} \end{bmatrix}$$

The agentic system maps $\mathbf{x}_t$ to:

- an estimated **need/intent** probability,
- an **action** (recommend/nudge/steer),
- a **channel** (app/web/advisor/call center),

under constraints (consent, fairness, risk).

III. DATA ARCHITECTURE FOR PERSONALISED FINANCIAL SERVICES

Data architecture for agentic, AI-driven data products for personalised customer journeys in group insurance and retirement services must embrace diverse data sources and the integration of structured, semi-structured, and unstructured data to expand the potential knowledge base. Data quality has a direct influence on product performance, so it must be managed throughout the product's lifecycle by applying appropriate governance standards for each data type. In domains such as insurance and finance, where sensitive customer information is involved and regulations are evolving rapidly, data handling must meet both local and international data-protection laws.

Three primary factors appear to drive the adoption of agentic AI in group insurance and retirement services: the available data foundation; the ability to manage the regulation around bias; and the enhanced model risk management enabled through agentic properties. Financial services firms are increasingly leveraging their rich data reservoirs to create differentiated products and services through personalisation—adapting their offerings in real time based on individual customer needs, behaviours, interests, intents, and past transactions. Personalisation is often orchestrated using natural-language interfaces that enable customers to interact with financial services dynamically.

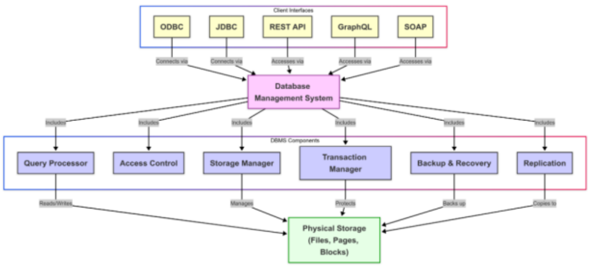

Fig 2: Data Architectures for Modern Financial

3.1. Data Sources and Integration in Group Insurance: Data for personalisation in group insurance products is typically collected at a given point in time from individual group members. This data, along with attribute information at the group level (e.g., group size and industry), is used to assign groups to specific product offerings or features that best match the groups' profiles. To achieve enhanced personalisation across agents, purposively generated proxies for individual group members are created and employed, as illustrated in Figure 1, similar to account-based marketing in B2B and B2B2C settings. These proxies help individualise targeted messaging without separately modelling each group member.

For agents to deliver timely and relevant services to their customers, data from multiple sources should be integrated in near real-time. The requirements for such real-time-based data integration, especially for insurance distribution, can differ from those of traditional online transaction processing use cases. Two key aspects of near real-time personalisation and the orchestration of cross-sell, up-sell and engagement actions are focused on – customer modelling and segmentation, and real-time personalisation and orchestration.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

3.2. Data Quality, Governance, and Compliance: In addition to relevant data sources, quality, governance, and compliance also influence the capability to deliver personalisation in products and services. A typical end-to-end data pipeline for analytics in financial services consists of stages for ingestion and storage, data engineering and feature generation, and other specialised pipelines for training machine learning models. A basic operational pipeline for data access may also be established and enhanced to allow for serving the end-user, whether human or machine. Pipelines should be modular and reproducible, with well-defined metadata. Data that is ingested must undergo quality checks and validation in the staging area, with tolerances defined for the maximum acceptable defect and the pattern and amount of anomalies.

Data quality is especially important for products targeting specific groups and subgroups, as the notion of quality itself is relative to the segment being addressed. With the depth and breadth of customer interaction, information at the customer touchpoints must also be captured in a structured manner, and real-time monitoring of transactions, service interactions, and social media data can enrich the available information. As with any technology-based product, issues of ethics and compliance come to the forefront, particularly in the handling and use of customer data. Data governance must cover the entire data pipeline, from data access and usage, to data quality and retention, and include a fraud detection mechanism. Controls must also be maintained on data used for training models to enable model explainability during regulatory audits and external validation.

IV. AGENTIC AI CAPABILITIES IN PRODUCT DELIVERY

Agentic data products harness agentic AI capabilities to measure and apply consumer preferences in the proposal, transaction, and servicing of group insurance, pension and provident fund services. Data-driven models simulate selected segments of customers and group members, and customer journey orchestration therefore becomes real-time and automated. The incorporation of these aspects requires structural setup and execution patterns of personalisation pipelines that allow for easy reproduction to meet group-level service delivery needs while remaining scalable. Supporting infrastructure furnish cloud services responsible for data access, modelling, real-time integration, and delivery.

Agentic AI’s long-term application is to make every product in insurance and related domains sufficiently personal such that the service delivery process becomes persuasive rather than push-based. Orchestration of customer journeys can be realised in shallow multilevel user journeys as these journeys require at least one segment-level model while at the same time facilitating orchestration among product-specific user journeys by evolving the pre-journey context. Such a setup automates personalisation of yet to be served users—making product delivery during transaction-time service levels personal and thereby effortless for the insurer while reaping pecuniary benefits.

4.1. Customer Modeling and Segmentation: Delivering personalised services often requires knowledge of both stated and latent customer needs. As many features of customer products in group insurance and retirement services are not highly differentiated for different customer segments, every large customer or customer market may be viewed as a special case requiring a slightly different implementation. Such specific requirements may not be apparent until identified by data exploration techniques. Derived products for customer segments may represent individualisation of service delivery rather than industrialisation supported by the ability to deliver in large volume for mass markets.

Intelligent Data Products aim to enable or support self-service delivery of group insurance or retirement services. However, many features requiring explanation and support cannot be delivered through an automated dialogue. Partial automation requires the synthetic agent to assist the customer and subsequently verify the decisions made. Alternatives for services that cannot be fully self-served need another type of data product—the Intelligent Experienced-Based Product. This product type responds using pre-stored responses or monologues without complex orchestration of responses. Customers sometimes have no live agent to assist them, or they may choose not to do so. In this case, a Data Product of the third type, the Intelligent Self-Service Product, is employed. This form provides maximum automation and requires the least skill cost to the supplier organisation for a consumer transaction but allows minimum consumer satisfaction.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

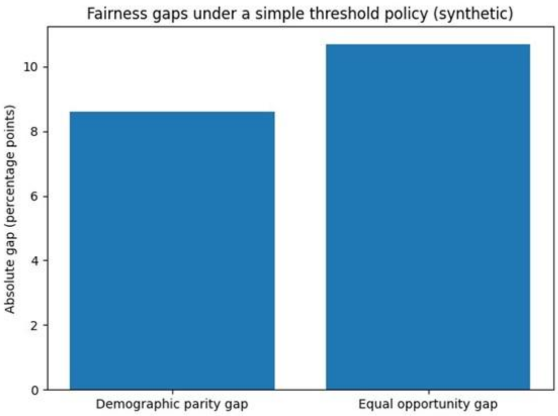

4.2. Real-Time Personalisation and Orchestration: Agentic AI-supported data products can support real-time personalisation of customer journeys. Delivering personalised products or services on demand, based on the needs of individual customers at a specific moment, is distinct from offering mass-customised products in which small variations better fit the preferences of customers in a given segment. Real-time personalisation is needed for delivering banking or insurance products (e.g., regulatory changes) or investment opportunities (e.g., market movements, price drops) that are highly time-sensitive.

Implementing real-time personalisation requires proper orchestration of the overall AI-driven product delivery process as well as the individual customer journey. AI-Assisted Moments Web services can help orchestrate the relevant functional modules, which includes personalisation of each sub-journey to the available customer segments, alerting enablers upon event emergence, and suggesting relevant items for fulfilment. Supporting product journeys under a service fulfilment model may help reduce the burden planning and operation of real-time AI-based services. Individual businesses can remain aspiration-free while still being able to address interesting real-time product opportunities and delivering AI-generated services to their customers.

Equation 2: Need / intent prediction model (binary) — full derivation

Let:

- $y \in \{0,1\}$ be whether the customer currently has a target intent (e.g., "needs contribution nudge").
- $\mathbf{x}$ be features at decision time.

Assume:

1. **Score (logit)**

$$z = \beta_0 + \boldsymbol{\beta}^\top \mathbf{x}$$

2. **Probability (sigmoid)**

$$p(y = 1 \mid \mathbf{x}) = \sigma(z) = \frac{1}{1+e^{-z}} \, p(y = 0 \mid \mathbf{x}) = 1 - \sigma(z)$$

For a single example $(\mathbf{x}, y)$:

$$\mathcal{L}(\beta_0, \boldsymbol{\beta}) = p(y \mid \mathbf{x}) = \sigma(z)^y (1 - \sigma(z))^{(1-y)}$$
$$\ell = \log\mathcal{L} = y\log\sigma(z) + (1 - y)\log\big(1 - \sigma(z)\big)$$
$$J = -\ell = -y\log\sigma(z) - (1 - y)\log\big(1 - \sigma(z)\big)$$

First, note:

$$\sigma(z) = \frac{1}{1 + e^{-z}}, \quad \frac{d\sigma}{dz} = \sigma(z)\big(1 - \sigma(z)\big)$$

Differentiate J w.r.t. z:

$$\frac{\partial J}{\partial z} = \sigma(z) - y$$

Now by chain rule, since $z = \beta_0 + \boldsymbol{\beta}^\top \mathbf{x}$:

$$\frac{\partial J}{\partial \boldsymbol{\beta}} = (\sigma(z) - y)\mathbf{x} \, \frac{\partial J}{\partial \beta_0} = \sigma(z) - y$$

This gives the standard SGD / batch GD update:

$$\boldsymbol{\beta} \leftarrow \boldsymbol{\beta} - \eta(\sigma(z) - y)\mathbf{x}$$

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

V. GOVERNANCE, RISK, AND COMPLIANCE FOR AGENTIC DATA PRODUCTS

The design and implementation of agentic AI-driven data products must consider privacy, security, trade compliance, intellectual property, model risk management, and ethical implications. An appropriate governance framework addresses these concerns and enriches the customer experience across all stages of interaction. Within the financial domain, regulatory compliance is necessary but not sufficient; a positive customer experience demands proactive warranties for user agency. Without adequate measures, personalisation in group insurance and retirement services could result in sub-optimal experiences for vulnerable customer segments or even become predatory.

Regulatory bodies impose considerable governance controls and model risk guidelines in the financial domain. Governance frameworks provide a standard yet customizable approach for implementing Regulations that guarantee fairness, reduce bias, and assure inclusivity for the financial services user community. An appropriate governance framework ensures compliance with model risk guidelines, covering model explainability, auditability, and facility for independent review by authorised personnel. In addition to addressing these mandatory requirements, an appropriate governance framework guarantees the ethical deployment of agentic AI-driven data products through appropriate warranties for user agency. It enables personalisation to act as a force for good—providing optimised product experiences for all users and preventing it from being used in a predatory manner. Such a framework requires expression through the customer consent process, connected to features covering data, model, and service auditability and detectability.

Fig 3: Agentic AI Governance and Risk Management

5.1. Regulatory Landscape in Group Insurance and Retirement Domains: Regulatory requirements in the group insurance and retirement industries govern consent when offering personalised services, enabling model risk management and auditing. Regulatory compliance has become imperative for data-driven products and services producing tailored recommendations via machine learning and Artificial Intelligence (AI). A novel data product must be developed according to regulatory constraints and authenticate its operational model via development life-cycle testing before deploying it at scale. Within group insurance and retirement industries, regulations such as General Data Protection Regulation (GDPR) require customers to provide their consent and in most cases also the consent of their peers (e.g., family members, guardians) when personalising services or recommending products that consider sensitive attributes.

In marketing machine learning- and AI-based products, other regulations require companies to understand and validate the associated model risk. Model risk arises when a model's output differs substantially from the real-world outcome, which is crucial at the recommendation stage. Regulatory authorities monitor the materialisation of model risk for machine learning and AI-based products and are introducing standards addressed at ensuring a clear governance of model risk. In this context, auditability has become crucial, enabling internal and external parties to independently validate forecasts generated through machine-learning- and AI-based methods.

5.2. Model Risk Management and Auditability: A regulatory model risk management framework must evaluate risk levels based on a model's intended use, materiality, any known issues, and the capability and resources available within the firm. Less complex models (as determined by a level of risk assessment) may be subject to periodic validation and governance recommendations, while high-risk models may require annual validation and/or oversight by an independent and dedicated governance body. As a minimum, model validation should occur when there are material changes to the model or when risks are identified by the model user. Periodic sensitivity analysis and backtesting strategy are also recommended, focusing on specific areas of concern. Model audit effort should be mitigated as far as possible by specifying the frequency of audit based on the level of risk and also enabling the external audit cycle to consider areas of

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

greater risk (e.g., low volume but high impact areas). Frameworks are also needed for validating the model build with fitness-for-purpose or “target” benchmarking based on subjective understanding rather than strict quantification.

VI. DEPLOYMENT ARCHITECTURES AND DEPLOYMENT PATTERNS

Scalable, container-based, microservices architectures can be leveraged to orchestrate real-time delivery of customer-directed services with personalisation and embedded advising. Modular, notebook-style artefacts enable reproducible, real-time execution of customer-sided, non-transactional modelling that can consume any desired amount of customer data, internal and external, and participate readily in distributed inference without functional duplication.

The architecture of agentic AI-enabled data products must support these operational requirements, ensure cost-efficient exploitation of the available volume, variety, and velocity of customer data, and minimise the risk of model bias while being flexible enough to accommodate different business logics across products and portfolios. Pattern-based pipelines offer a possible solution. Using pattern- and component-based approaches from visual programming language and application domain approaches in big data analytics, the proposed approach separates the design, development, and execution phases of the pipeline lifecycle. Applications leverage pipelines, which are designed and constructed by domain actors based on building blocks provided by data engineers. Reproducibility is attained by the notebook-style artefact.

6.1. Modular, Reproducible Pipelines: The latter scenario illustrates that the growing complexity in predictive modelling is often mismanaged with monolithic code bases leading to redundant components, replication of efforts and wasted development effort. The delivery of data products though modular and reproducible pipelines supported by a hub-and-spoke data platform can avoid these problems and ensure service delivery meets the real-time needs of customers across product lines, by drawing on the capabilities of operations research for planning and optimisation and machine learning for prediction and recommendation, within a single pilot to meet multiple real-time personalisation needs in life insurance and investment services.

The adoption of AI, and especially machine learning (ML), is often held back because of a lack of production-ready data and the inflexibility of data pipelines to address the diversity of analytics use cases. These problems can be removed with an integrated approach to data and AI services that combines a hub-and-spoke data architecture – addressing the common difficulties of data availability and consistency – with a composite modelling approach for real-time data products – channelling the delivery of modelling from prediction to planning and operations. Data pipelines for AI can then be provided as a shared service that supports reproducible delivery of analytics code across the business.

6.2. Scalable Infrastructure for Personalised Services: Scalability is one of the essential attributes of service-rich data products. With customer behaviours and needs evolving continuously due to changes within their personal circumstances and lifestyles, their requirements from financial services also evolve. Environments of turbulent times like a pandemic may also require a change in consumer preferences. An infrastructure enabling the crediting of adaptability and agility to data products will open doors for responsiveness and real-time behaviour changes to be effortlessly incorporated into the data products supporting the underlying services. The need to change and update underlying data products, such as those supporting TV ad campaigns, may be captured via a simple creative brief that lists the prospective customers to be targeted. The product manager may not even need to think about services like TV ads separately and continuously, services like TV ads may become an effect that happens automatically based on need and already-known information of consumers.

A consumer’s need for a loan to purchase a car, bike, etc. is short-lived, and such needs may simply be responded to with a few relevant products. Therefore, a need for a loan for a specific purpose and a requirement for a second-hand loan need not trigger responses to the entire product portfolio on the open digital marketing API; they can easily trigger only those products for which responding makes sense. Enabling product services and their entire underlying products to be triggered automatically and spread across delivery pipelines and multiple data products would make for speed and scale that can be very difficult to create manually. Such an architecture, data product, and services ecosystem would equip businesses well to retain much of the personalisation even without a large bank of data scientists who are constantly working on continuous service improvement.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

Equation 3: Next-best action / orchestration as constrained optimisation
Let actions be $a \in \mathcal{A}$ (e.g., "nudge", "recommend product", "route to advisor", "do nothing").
Define a utility:

$$U(a,\mathbf{x}) = \underbrace{\mathbb{E}[\text{benefit} \mid a,\mathbf{x}]}_{\text{engagement, welfare, retention}} - \lambda \underbrace{C(a,\mathbf{x})}_{\text{cost}} - \mu \underbrace{R(a,\mathbf{x})}_{\text{risk}}$$

Then decision rule:

$$a^*(\mathbf{x}) = \underset{a \in \mathcal{A}}{\operatorname{argmax}} U(a,\mathbf{x})$$

Let:
- $g_k(a,\mathbf{x}) \leq 0$ be constraints (e.g., consent not granted; unsuitable advice; restricted attributes).

Then:

$$\max_{a \in \mathcal{A}} U(a,\mathbf{x}) \quad \text{s.t.} \quad g_k(a,\mathbf{x}) \leq 0 \ \forall k$$

This is the clean "math wrapper" for what the paper describes operationally.

VII. ETHICAL, SOCIAL, AND ECONOMIC IMPLICATIONS

Consumer consent and autonomy represent critical aspects of the ethical use of agentic AI capabilities. AI-driven products that operate autonomously on behalf of consumers typically require consumers to give explicit consent at product inception. However, consumers are significantly overloaded by the sheer number of consent requests pertaining to websites, applications, and many other product offerings. As such, they are increasingly desensitised and less likely to read through the terms and conditions before granting permission for data sharing and processing activities. Furthermore, the lack of pre-definite preferences concerning data sharing leads to a pleasing of the parties involved rather than a genuinely transparent, informed, and voluntary consent process.

Emergences of consumer preference centres that allow consumers to define a privacy profile, which can subsequently be referenced by different agents as the consumers interact with third-party digital entities, have the potential to simplify the consent process while allowing for a more ethical use of agents, standardising the way privacy preferences get defined and communicated. Similar to third-party cookie management, creating and sustaining a repository of a consumer's wish or preference concerning privacy would ease consumer desire for personalisation and operand hate by limiting the need to repeatedly set preferences for individual merchants or services while maintaining control on what information is shared and with whom.

7.1. Fairness, Bias, and Inclusion: Developments in ML pose new responsibilities and risks. In addition to improving decision making, these systems also have the potential to introduce bias and discrimination. ML decisions can be construed as unfair when the predicted outcome for a given input differs between two individuals who are similar with respect to relevant attributes. An unfair system exhibits inconsistent behaviour. The use of historical data to train prediction models introduces risk of bias, since past decisions may reflect the influence of societal prejudices. A biased model may act unfairly against particular demographic groups, undermining equality of opportunity and trust. Bias can manifest in two different ways: bias-in-data, representing the reflected prejudices of the data collection, and bias-in-algorithm, representing how the model is trained or refined. Professional agents are best equipped to assess risk for high-value accounts, but their recommender systems should not favour certain agents over others. Focusing only on socio-economic variables can overlook how preferences and needs change with life events; thus, life stages better capture non-trivial interactions.

Ensuring fairness may require applying pre-, post-, or in-processing techniques that depend on the analysis type, the prediction target, and how the analytically prepared data will inform future decisions. Fair classification aims to classify a data object into a sensitive category, e.g. favourable or unfavourable. Fairness constraints ensure that the classifier meets fairness requirements for the sensitive class. Other fairness-aware techniques seek to mitigate bias during prediction, by rejecting biased instances in a pre-processing stage. Fair instances are filtered to train a classifier that can de-bias the sensitive class in a post-processing stage. The data scientific pipeline for fairness can be tailored to different fairness definitions.

7.2. Consumer Consent and Autonomy: Consumer consent plays a critical role in the ethical and responsible use of personal data throughout the development and delivery processes of data-driven agentic AI products. The data sources feeding machine learning pipelines and the model outputs driving customer engagement must be owned and consented to by the end consumers. Consent conveys an explicit understanding of personal preferences, trade-offs, risks, and

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

undercurrents of any transaction. Prior work has identified the need for the continual monitoring of the act of consent as it evolves over time, rather than merely obtaining a one-time affirmation. Consumer consent overlays layers of product governance, evaluation, and supervision, dynamically connecting the consumer model with the personalisation aspects of the AI product. Data distribution and consent can be modelled like pricing matrices, serving as critical elements of customer engagement and model evaluation.

Despite the unquestionable benefits of giving consumers such control over their personal data heritage, weaknesses in the traditional consent frameworks have been highlighted, exposing consumers to extended privacy loss. Savvy corporations able to distinguish themselves by genuinely safeguarding and impartially using consented consumer data often fail to enjoy the goodwill of the public, as they are, by and large, perceived as self-serving. Empathy-driven AI and agentic solutions present a radically different paradigm of data-led engagement, engagement that oozes concern and deference for the consumer's right to privacy and stops-operate, pay-per-use, and even opt-in and opt-out modelling.

Empathy-driven and agentic data products combine altruistic design precepts with state-of-the-art technology know-how to reproduce an affective and cognitively meaningful experience for data-driven consumer-centric businesses without extending the same condemning privacy exposure.

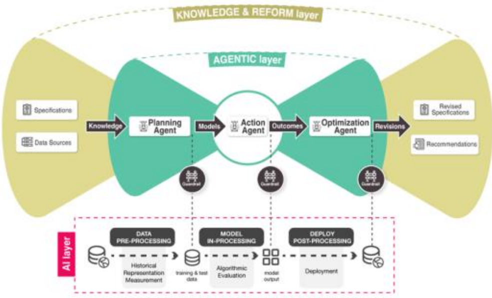

Fig 4: Agentic AI fairness framework

VIII. CONCLUSION

Research in human-level multilingual language models, multimodal models, and AI-bots for dialogue with intent and agency has gained significant momentum. The evolving capabilities in generation and interaction open opportunities to drive the creation and consumption of AI apps in a human-friendly manner. Structured product delivery for financial services does not allow generative modelling of the underlying data since the applied functions are not mere patterns. Change in the function signatures—such as the addition of constraints (e.g., consolidation of pension accounts before providing projections)—requires a delicate orchestration with multiple points of human interaction. However, support in these areas makes agentic AI possible for data products with structured delivery. A wealth of ready-to-consume services based on agentic data products with other delivery modes can be leveraged for conversational consumption.

Personalisation of contemporary group insurance and retirement products/services is indeed of interest. Since real-time personalisation requires extensive capability in customer/entity modelling and segmentation, it is disclosed separately. Personalisation of product delivery and the interaction required to support it are also discussed since they are indeed part of product delivery. The need for fairness, inclusion, model audits, governance, and risk management for the delivered agentic data products and related services complement the consumer risk and autonomy aspects along with the implications of non-consumer ownership.

8.1. Emerging Trends: Data dynamics and decision spaces are becoming increasingly complex for large tax for large tax-exempt organisations participating in industries such as Group Insurance, Retirement Services, Large-Scale Distribution, and Asset Management. Personalisation is crucial for the future success of organisations within these industries, adopting Agentic AI principles in the design, operation, and orchestration of services. Agentic AI-powered data products facilitate real-time customer modelling and segmentation, allowing companies to deliver smart services tailored to evolving customer preferences. Data architecture design plays a critical role in delivering such services, providing the relevant and trustworthy information required to support the operation of an Agentic AI ecosystem.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

To offer meaningful personalisation, an ecosystem of Agentic AI-driven products must be established that selects the "services of service" best suited to a customer's unique context at a given moment. These products orchestrate other data products and fulfil customers' main goals, such as decisions, purchases, and life events. They follow modular design principles and reusable pipelines that enhance maintainability and scalability across a large customer base. Additional key considerations involve the breadth of decision-making and the integration of technology and human actors in the orchestration of complex decision flows.

REFERENCES

1. Adusupalli, B. (2024). Agentic AI-Driven Identity and Access Management Framework for Secure Insurance Ecosystems. Journal of Computational Analysis and Applications(JoCAAA), 33(08), 2794-2814.
2. Ranga Reddy, V. A. (2024). Comparing Batch vs. Streaming Approaches in Healthcare Data Warehousing Environments. Journal of Neonatal Surgery, 13(1), 2287-2309.
3. Babaei, G., & Giudici, P. (2024). How fair is machine learning in credit lending? Quality and Reliability Engineering International, 40, 3452–3464.
4. Kolla, T. (2024). AI-Powered Data Catalog Systems For Healthcare Data Discovery And Governance. South Eastern European Journal of Public Health, 2296–2311. https://doi.org/10.70135/seejph.vi.7077
5. Inala, R., & Somu, B. (2024). Agentic ai in retail banking: Redefining customer service and financial decision-making. Journal of Artificial Intelligence and Big Data Disciplines, 1(1), 1-19.
6. Mukesh, A., & Aitha, A. R. (2021). Insurance Risk Assessment Using Predictive Modeling Techniques. International Journal of Emerging Research in Engineering and Technology, 2(4), 68-79.
7. Baumann, J., & Loi, M. (2023). Fairness and risk: An ethical argument for a group fairness definition insurers can use. Philosophy & Technology, 36, Article 45.
8. Mashetty, S. (2024). The role of US patents and trademarks in advancing mortgage financing technologies. European Advanced Journal for Science & Engineering (EAJSE)-p-ISSN, 3050-9696.
9. Challa, K. (2024). Artificial Intelligence and Generative Neural Systems: Creating Smarter Customer Support Models for Digital Financial Services. Journal of Computational Analysis & Applications, 33(8).
10. Blohm, I., Wortmann, F., Legner, C., & Köbler, F. (2024). Data products, data mesh, and data fabric: New paradigm(s) for data and analytics? Business & Information Systems Engineering, 66, 643–652.
11. Nagabhyru, K. C. (2024). Data Engineering in the Age of Large Language Models: Transforming Data Access, Curation, and Enterprise Interpretation. Computer Fraud and Security, 2024(12).
12. Amistapuram, K. (2024). Federated Learning for Cross-Carrier Insurance Fraud Detection: Secure Multi-Institutional Collaboration. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 6727-6738.
13. Casas, P., & Román, C. (2024). The impact of artificial intelligence in the early retirement decision. Empirica, 51, 583–618.
14. Kolla, S. H., & Peddi, R. K. (2024). Designing Governance-Aligned GenAI Pipelines Using Small Language Models for Enterprise Workflow Intelligence. International Journal of Science, Research and Technology, 7(6), 13256-13268.
15. Vardhan Kumar Bandi, V. D. (2024). Automated Feature Engineering Systems in Large-Scale Healthcare Data Environments. Journal of Neonatal Surgery, 13(1), 2127-2141.
16. Černevičienė, J., & Kabašinskas, A. (2024). Explainable artificial intelligence (XAI) in finance: A systematic literature review. Artificial Intelligence Review, 57, Article 216.
17. Davuluri, P. N. (2020). Improving Data Quality and Lineage in Regulated Financial Data Platforms. Finance and Economics, 1(1), 1-14.
18. Chen, G., Dong, S., Shu, Y., Zhang, G., Sesay, J., Karlsson, B., Fu, J., & Shi, Y. (2024). AutoAgents: A framework for automatic agent generation. Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, 22–30.
19. Mangala, N. (2024). Leveraging Microsoft Fabric lakehouse as an AI-ready data platform for enterprise analytics. Journal of Information Systems Engineering and Management.
20. 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.
21. Chen, Y., Giudici, P., Liu, K., & Raffinetti, E. (2024). Measuring fairness in credit ratings. Expert Systems with Applications, 258, Article 125184.
22. Kolla, S. K., & Reddy, V. A. R. (2024). Evaluating Cloud-Native vs. Hybrid Architectures for Health Benefit Administration Systems. International Journal of Medical Toxicology and Legal Medicine, 27(5), 1042-1053.
23. 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.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

24. Paleti, S. (2024). Neural Compliance: Designing AI-Driven Risk Protocols for Real-Time Governance in Digital Banking Systems. Available at SSRN 5233099.
25. D'Acunto, F., & Rossi, A. G. (2023). Robo-advice: Transforming households into rational economic agents. Annual Review of Financial Economics, 15, 543–563.
26. Meda, R. (2024). Agentic AI in Multi-Tiered Paint Supply Chains: A Case Study on Efficiency and Responsiveness. Journal of Compu-tational Analysis and Applications (JoCAAA), 33(08), 3994-4015.
27. Avinash Pamisetty, Vijaya Rama Raju Gottimukkala. (2024). Agentic AI-Driven Multi-Cloud Big Data Architecture For Predictive Demand, Credit Risk, And Inventory Financing In National Food Service Supply Chains. Metallurgical and Materials Engineering, 30(4), 959–975. https://doi.org/10.63278/mme.v30i4.1933
28. Segireddy, A. R. (2024). Machine Learning-Driven Anomaly Detection in CI/CD Pipelines for Financial Applications. Journal of Computational Analysis and Applications, 33(8).
29. Dibia, V., Chen, J., Bansal, G., Syed, S., Fourney, A., Zhu, E., Wang, C., & Amershi, S. (2024). AutoGen Studio: A no-code developer tool for building and debugging multi-agent systems. Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, 72–79.
30. Yandamuri, U. S. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. International Journal Of Finance, 36(6), 682-706.
31. Kolla, S. K. (2024). Federated Machine Learning On Big Healthcare Data For Privacy-Preserving Analytics. The Review of Diabetic Studies, 175-190.
32. Du, K., Zhao, Y., Mao, R., Xing, F., & Cambria, E. (2024). Natural language processing in finance: A survey. Information Fusion, 102755.
33. 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.
34. Reddy, V. A. R. (2024). Generative Intelligence for Healthcare Claims Processing and Personalized Benefits Management. International Journal of Advanced Engineering Science and Information Technology (IJAESIT), 7(3), 14099.
35. Fahrenwaldt, M., Furrer, C., Hiabu, M. E., et al. (2024). Fairness: Plurality, causality, and insurability. European Actuarial Journal, 14, 317–328.
36. Challa, K. (2024). Enhancing credit risk assessment using AI and big data in modern finance. American Data Science Journal for Advanced Computations (ADSJAC) ISSN, 3067-4166.
37. Mashetty, S. (2024). Research insights into the intersection of mortgage analytics, community investment, and affordable housing policy. Available at SSRN 5249213.
38. Hong, S., Zhuge, M., Chen, J., Zheng, X., Cheng, Y., Wang, J., Zhang, C., Wang, Z., Yau, S. K. S., Lin, Z., Zhou, L., Ran, C., Xiao, L., Wu, C., & Schmidhuber, J. (2024). MetaGPT: Meta programming for a multi-agent collaborative framework. International Conference on Learning Representations.
39. Davuluri, P. S. L. (2023). AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems (December 15, 2023).
40. Mangalampalli, B. M. (2024). Transparent Intelligence Explainability Frameworks for AI-Driven Clinical Decision Support in Healthcare Business Intelligence. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(3), 10566-10579.
41. Kofman, P. (2024). Scoring the ethics of AI robo-advice: Why we need gateways and ratings. Journal of Business Ethics, 198, 21–33.
42. Kolla, S. K. (2024). Clinical Knowledge Intelligence through Deep Learning and Natural Language Understanding in Healthcare Platforms. International Journal of Advanced Engineering Science and Information Technology (IJAESIT), 7(3), 14088.
43. Bandi, V. D. V. K. (2024). AI-Driven Predictive Risk Modeling Architectures for Financial Systems. International Journal Of Finance, 37(3), 54-78.
44. Lazar, E., Qi, S., & Tunaru, R. (2024). Measures of model risk for continuous-time finance models. Journal of Financial Econometrics, 22(5), 1456–1481.
45. Davuluri, P. S. L. N. (2023). Integrating artificial intelligence into event-driven financial crime compliance platforms. International Journal of Finance, 36(6), 707-736.
46. Amistapuram, K. (2024). Smart Decision Support Systems For Dynamic Tax Policy Optimization Using Reinforcement Learning. Available at SSRN.
47. Lee, J., Stevens, N., Han, S. C., & Song, M. (2024). A survey of large language models in finance (FinLLMs). arXiv preprint arXiv:2402.02315.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

48. Nagabhyru, K. C., & Engineer, S. D. (2023). Unifying Data Engineering and Machine Learning Pipelines: An Enterprise Roadmap to Automated Model Deployment.
49. Aitha, A. R. (2024). Generative AI-Powered Fraud Detection in Workers' Compensation: A DevOps-Based Multi-Cloud Architecture Leveraging, Deep Learning, and Explainable AI. Computer Fraud and Security.
50. Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., & Neubig, G. (2023). Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. ACM Computing Surveys, 55(9), 1–35.
51. Inala, R. (2021). A New Paradigm in Retirement Solution Platforms: Leveraging Data Governance to Build AI-Ready Data Products. Journal of International Crisis and Risk Communication Research, 286-310.
52. Deep Learning-Driven Optimization of ISO 20022 Protocol Stacks for Secure Cross-Border Messaging. (2024). MSW Management Journal, 34(2), 1545-1554.
53. Madaan, A., Tandon, N., Gupta, P., Hallinan, S., Gao, L., Wiegreffe, S., Alon, U., Dziri, N., Prabhumoye, S., Yang, Y., Gupta, S., Majumder, B. P., Hermann, K., Welleck, S., Yazdanbakhsh, A., & Clark, P. (2023). Self-refine: Iterative refinement with self-feedback. Advances in Neural Information Processing Systems, 36.
54. Kolla, S. H. (2024). Retrieval-Augmented Enterprise Intelligence: Enhancing Accuracy, Trust, and Operational Decision-Making. International Journal of Future Innovative Science and Technology (IJFIST), 7(2), 12425.
55. Joshi, A., Sujatha, G., Gupta, N., Kumar, R., Sen, M. K., Ish, P., ... & Popalwar, H. (2024). Clinical utility of pulmonary rehabilitation in diffuse parenchymal lung diseases. Journal of Advanced Lung Health, 4(3), 159-165.
56. Mangala, N. (2021). CI/CD Pipeline Automation for Enterprise Data Artifacts Using Azure DevOps. Universal Journal of Business and Management, 1(1), 1-18.
57. Nwafor, C. N., Nwafor, O., & Brahma, S. (2024). Enhancing transparency and fairness in automated credit decisions: An explainable novel hybrid machine learning approach. Scientific Reports, 14, Article 25174.
58. Loganathan, R. (2022). Converging Security Architecture and Compliance Management in Enterprise Data Center Ecosystems: A Unified Control Framework. International Journal of Scientific Research and Modern Technology, 1(12), 295-312.
59. Kolla, S. K., & Mangalampalli, B. M. (2024). Edge-Based Deep Learning Systems for Point-of-Care Diagnostic Intelligence. Journal of Neonatal Surgery, 13(1), 2387-2399.
60. Paleti, S. (2024). Data engineering for AI-powered compliance: A new paradigm in banking risk management. Available at SSRN 5256619.
61. Nie, Y., Kong, Y., Dong, X., Mulvey, J. M., Poor, H. V., Wen, Q., & Zohren, S. (2024). A survey of large language models for financial applications: Progress, prospects and challenges. arXiv preprint arXiv:2406.11903.
62. Meda, R., & Pamisetty, A. (2023). Intelligent Infrastructure for Real-Time Inventory and Logistics in Retail Supply Chains. Educational Administration: Theory and Practice, 29(4), 5215-5233.
63. Mangalampalli, B. M. (2024). AI-Enhanced Data Governance: Automating Compliance In Healthcare Analytics Platforms. The Review Of DIABETIC STUDIES OPEN ACCESS.
64. Kolla, T. (2024). Graph Neural Networks for HCC Risk Adjustment and Interoperability. International Journal of Science, Research and Technology, 7(6), 13244-13255.
65. Peiris, H., Jeong, H., Kim, J.-K., & Lee, H. (2024). Integration of traditional and telematics data for efficient insurance claims prediction. ASTIN Bulletin: The Journal of the International Association of Actuaries, 54(2), 263–279.
66. Challa, K. (2024). Neural Networks in Inclusive Financial Systems: Generative AI for Bridging the Gap Between Technology and Socioeconomic Equity. MSW Management Journal, 34(2), 749-763.
67. Inala, R. (2020). Building Foundational Data Products for Financial Services: A MDM-Based Approach to Customer, and Product Data Integration. Universal Journal of Finance and Economics, 1(1), 1-18.
68. Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Hambro, E., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2023). Toolformer: Language models can teach themselves to use tools. Advances in Neural Information Processing Systems, 36.
69. Aitha, A. R. (2022). Deep Neural Networks for Property Risk Prediction Leveraging Aerial and Satellite Imaging. International Journal of Communication Networks and Information Security (IJCNIS), 14(3), 1308-1318.
70. Nagabhyru, K. C. (2023). Accelerating Digital Transformation with AI Driven Data Engineering: Industry Case Studies from Cloud and IoT Domains. Educational Administration: Theory and Practice, 29 (4), 5898–5910.
71. Gottimukkala, V. R. R. (2024). Federated Learning Approaches for Fraud Detection in International Payment Systems. https://www. jisem-journal. com/download/118_JISEM. pdf.
72. Takayanagi, T., & Izumi, K. (2024). Incorporating domain-specific traits into personality-aware recommendations for financial applications. New Generation Computing, 42, 635–649.
73. Reddy Segireddy, A. (2024). Federated Cloud Approaches for Multi-Regional Payment Messaging Systems. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(2), 442-450.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804025

74. Kolla, S. H. (2022). Strategic Information Integration Models for Cross-Functional Service Optimization in Large-Scale Enterprises. International Journal of Emerging Trends in Engineering and Management Research, 7(3), 11811.
75. Bandi, V. D. V. K. (2024). Intelligent Data Platforms For Personalized Retail Analytics At Scale. Metallurgical and Materials Engineering, 30 (4), 1011–1027.
76. Tang, H. (2024). How to promote data sharing of financial service providers based on federated learning technology. Information Resources Management Journal, 37(1), 1–25.
77. Mattaparthi, R. (2024). Transformer-Based Fault Diagnosis for Large-Scale Standby Power Generators: Partial Discharge Pattern Recognition at Hyperscale Data Center Installations. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 8781-8799.
78. Kolla, T. (2024). Intelligent Discovery and Governance of Healthcare Data Assets Through AI-Powered Catalog Architectures. International Journal of Emerging Trends in Engineering and Management Research, 9(4), 16083.
79. Peddi, R. K. (2021). Optimizing Case Management Workflows in Global Data Center Colocation Services. Universal Journal of Computer Sciences and Communications, 1(1), 1-21.
80. Nampalli, R. C. R., & Adusupalli, B. (2024). Using Machine Learning for Predictive Freight Demand and Route Optimization in Road and Rail Logistics. Library of Progress-Library Science. Information Technology & Computer, 44(3).
81. Vandanapu, M. K. (2024). AI-driven personalization in financial services: Enhancing customer experience and operational efficiency. Journal of Economics, Management and Trade, 30(11), 1–13.
82. Paleti, S. (2024). Transforming financial risk management with AI and data engineering in the modern banking sector. American Journal of Analytics and Artificial Intelligence (ajaai) with ISSN.
83. Kolla, S., Meda, R., Balleda, L., & Thimmapuram, C. R. (2024). The utility value of ROX index and modified ROX index in determining the efficiency of HFNC in children admitted with respiratory distress. International Journal of Contemporary Pediatrics, 11(6), 775.
84. Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., Awadallah, A. H., & others. (2023). AutoGen: Enabling next-gen LLM applications via multi-agent conversation. arXiv preprint arXiv:2308.08155.
85. Mashetty, S. (2024). Redefining financial risk strategies: The integration of smart automation, secure access systems, and predictive intelligence in insurance, lending, and asset management. Journal of Artificial Intelligence and Big Data Disciplines (JAIBDD).
86. 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.
87. Adusupalli, B., & Insurity-Lead, A. C. E. (2024). The role of internal audit in enhancing corporate governance: A comparative analysis of risk management and compliance strategies. Outcomes. Journal for ReAttach Therapy and Developmental Diversities, 6, 1921-1937.
88. Vamsee Pamisetty, Keerthi Amistapuram. (2024). Smart Decision Support Systems For Dynamic Tax Policy Optimization Using Reinforcement Learning. Metallurgical and Materials Engineering, 30(4), 976–995. https://doi.org/10.63278/mme.v30i4.1934
89. Mattaparthi, R. (2023). Connected Fleet Intelligence: Edge-Centric Analytics and Computer Vision for Predictive Manufacturing and Asset Resilience. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(5), 9077-9088.
90. Kolla, S. K., & Mangalampalli, B. M. (2024). Edge-Based Deep Learning Systems for Point-of-Care Diagnostic Intelligence. Journal of Neonatal Surgery, 13(1), 2387-2399.