Eksplorium *p*-ISSN 0854-1418
Volume 46 No. 1, May 2025: 1614–1628 *e*-ISSN 2503-426X

A Unified Framework for Agentic AI and Data Products: Enhancing Cloud, Big Data, and Machine Learning in Supply Chain, Insurance, Retail, and Manufacturing

[1]Ramesh Inala, [2]Srinivas Kalisetty
[1]Data Engineer, rameshhinala@gmail.com, ORCID ID: 0009-0009-2933-4411
[2]Integration and AI lead, srinivas.kalisetty.ic@gmail.com, ORCID: 0009-0006-0874-9616

Article Received: 25 Feb 2025, Revised: 27 April 2025, Accepted: 05 May 2025

Abstract: While the evolution of deployed AI technologies is attracting increasing interest and investment across domains, there remains severe ambiguity on the question of what types of technologies should be categorized as "AI". Indeed, there exists an entire class of digital technologies in production that utilize a variety of different algorithms, ROI-base value propositions, and user experiences, some developed with AI techniques and some leveraging statistically-defined heuristics and decision support. This paper proposes a framework to help clarify the boundaries between algorithmic, AI, and agentic AI systems as well as how these systems may be utilized in pursuit of different data product goals. We believe these discussions can lay the groundwork to support industry-specific discussions on this important topic going forward.

Keywords: Agentic AI,Data Products,Unified Framework,Cloud Computing,Big Data Analytics,Machine Learning,Supply Chain Optimization,Insurance Analytics,Retail Intelligence,Smart Manufacturing,AI-Driven Decision Making,Intelligent Automation,Digital Transformation,Predictive Analytics,Industry 4.0.

1. Introduction

AI is having a significant impact on many economic sectors. This has sparked considerable interest and engagement from industry players. As a result of features such as low barriers to entry and easy access to implementations and tools for building and deploying AI products, the pace of innovation is rapid. Corporations across sectors are exploring and building Large Language Models to power applications, including global tech companies who are competing in offering capable proprietary APIs to commercialize models. Other companies are launching products focused on user experience trying to establish their position in the market. Meanwhile, the popularity of Large Language Models has also triggered the academic community to push on advancement and exploration of novel techniques.

Even though AI technologies are evolving, these do not operate in a vacuum. AI systems are mere tools, lack telos, and are actively built by human minds and used for human tasks pursuing human objectives. A manager using a data product to predict work delegation lags must ultimately determine which worker gets which task. It is human expertise that contextualizes the predictions operated upon by the data product to drive better outcomes. The concept of a data product has made its way from the concept of a software product to the concept of a data-driven product. The same discussion can be applied for AI agentic products, where agent-based interaction is present, the AI acts on behalf of the human agent. What if, in those cases, do we even need a human to be present in the loop or agent tech is truly coupling the human need and action? Answering those questions relies on agency but expands into UI/UX design and productization of tech and data science for user engagement and outcome orientation of tasks enabled by the model.

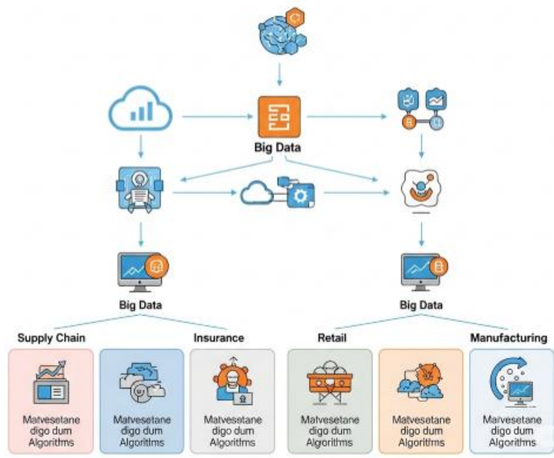

Fig 1: A Unified Framework for Agentic AI and Data Products: Enhancing Cloud, Big Data, and Machine Learning in Supply Chain, Insurance, Retail, and Manufacturing.

1.1. Background and Significance Despite rapid growth in AI and ML capabilities for business intelligence and automation, Enterprises struggle to figure out how best to leverage these technologies to create tangible business value. Creating predictive models to personalize experiences, anticipate demand, improve detection, and prescribe actions are typical initial use cases. These use cases are primarily limited to predictions about critical data products, which are often monetized in regulated markets. The majority of enterprises, however, have limited experience with these data products and typically take the form of dashboards and alerts. Although the dashboards can help key users fulfill their information needs, the alerting capabilities are what turn the dashboards into data surveillance systems. As the data products deliver predictions, the enterprises also need to monitor how predictive they really are. The initial data product stage is about maintaining the predictive utility of the alerts and refining the expectancy management plan until a repeatable business process has been created.

In mature organizations, the data product-related enterprise digital investments deliver key predictions to frontline users who transmit actions of interest. If the predictions are insightful, the organizations make sure to harvest key streams of relevant action data, along with the original data that powered the predictions. Over time, the action data collected provides feedback on the predictive alerts, which can be constructed as classification models for agents that understand the organization's context and why it changed. As these agents emphasize action-triggering event detection, fewer of the frontline users need to worry about data product surveillance. The models that power the data products, and their interactions with the users, become agents of the organization, whose agentic intelligence is usually only directly visible to a small number of employees who have precise context on the decision-making process.

Equ 1: Data Ingestion & Processing Pipeline

$$D_t = \mathcal{T}(S_t, C_t, E_t)$$

- D_t: Feature-engineered dataset at time t
- $\mathcal{T}$: Transformation function (e.g., ETL, feature engineering)
- S_t: Structured data (e.g., ERP, CRM)
- C_t: Cloud-native log streams or real-time sensor data
- E_t: External data (market, weather, social)

2. Understanding Agentic AI

What is agentic AI? By agentic, we mean possessing agency as activation of latent intent inside the model coupled with consequentialist behavior that demonstrates that intent and thus makes it an agent used for some purpose via an API. It is called agentic AI in opposition to non-agentic AI which merely undertakes decision making or mapping related problems utilizing aspects of intelligence like language model prediction probability. Similar work for reinforcing AI either for generative or data product uses cases carried out at scale also falls under this definition. Another term sometimes used for describing the same phenomena is self-reinforcing AI but the agentic in particular hints at the incentive compatibility framework used.

Data AI has become omnipresent in modern economies; it augments accounting firms to tech companies, using data enhancing capabilities to create products of perdurable value. However, as profound and multifactorial these benefits may be, attendant with the rapid development of decision-making and anchoring-capable AI is the intrinsic concern about the safety of such technology, as organizations calibrate the levers in their respective search for focus and growth and the incredibly power AI have over being able to create unseen side-effects as they scale up. It is a double-edged sword. This paper's goal is to present such a framework that allows for examining the latter while addressing the former concern by studying the phenomenology of how they operate atop a unified model of human decision-making as well as organization design.

The potential upside is compelling as what I really do is inference under uncertainty or prediction under misspecification of agents. Here agents can be humans, markets or other value abstraction devices. Relating the prime objective of any organization to optimize the expectation of their future cash flows admit two key utilizations of AI. The first being the design of any organization's structure, decision-making mechanism or obstacle by the usage planning of how self-supervised algorithms implement the prediction-decision heuristic.

2.1. Definition and Characteristics

Agentic AI is distinguished from other forms of AI by its goal-directedness, distinguishability, criticality, and impact. While anything that acts for us, or gets used by us, might be called an agent, not all agents function as true agents. AI truly becomes agentic when it functions like an agent in philosophy, or the willful actor with intentions of rational choice or logical reasoning who is separate from but also acts upon the objective environment in a meaningful and suitably significant manner. In other words, an idea is agentic when its goals or purposes are distinct from ideas of us who use it or people around us; and those goals or purposes are so critical that we assign great importance to them; and they accomplish or seek to accomplish great things; and they are not or cannot be so easily ignored simply because they lack overwhelming levels of perceived danger.

Let's consider a few examples: from the earliest LLMs to our most recent constructions, ChatGPT-like LLMs today see all the world's information as expert knowledge, and are trained to generate the most expert-like responses to the prompts we provide them, asking them to behave like they had our own best interests in mind. However, to their developers and operators, these models are moreover tools for data processing and annotation, customer service, decision-making support, and, in general, development of data products for profits. Agentic AI refers to interactions between models and users or developers in the former sense; Data Products capabilities dovetail into the latter sense. In contrast, code-generating models have until now been limited to serving tools capable of accelerating and augmenting our own code-writing interactions. It is also for code-generating models that we see the most overt consequences of this shift to agentic AI.

2.2. Role in Decision Making Many data products currently used in management decision making are not agentic AI. Rather, they are decision aids, which we would describe as non-autonomous decision processors. For example, a business intelligence dashboard may visualize sales, inventory, and demographic data to support executive decisions, or a sales escalation model may risk-score specific calls for help, to assist customer service agents. These may use sophisticated algorithms but ultimately rely on non-autonomous decision makers. Despite often requiring complex algorithms to produce, these data products export the decision and responsibility for actions external to artificial intelligence.

Why bother using algorithms to support routine decisions when a human can? Scale and value are the reasons. For decisions made at large volume—such as customer segmentation, patient risk scoring, supply chain capacity selection—however great the human skill, the cost of time will be high, and in the large majority of cases the typical human decision makers will be much less capable than models based on massive amounts of reliable data. In these cases, increasing the number of decisions handled by computer algorithms means that externalized decision products can add higher value than the typical otherwise intelligent individual decision maker—while using these systems saves on the incremental costs involved in human processing. The same arguments apply to non-agentic AI that supports other organizations in their decision processes in high volume. Such models can cost much more to develop than typical products exported. But the best externalizing agentic AI's can be highly profitable, making billions of dollars a year.

3. Data Products in the Modern Era

While data products have existed for a long time, the advent of large-scale data collection and analysis has enabled new types. Today, we can define six main types of data products, and categorize them on two axes. Three are produced on an ad-hoc basis by an analyst, and three are produced on an ongoing basis on an automated pipeline, which can happen in real time or not. All these data products can serve both external and internal users.

The first type of data product is the dashboard. Rich visual and interactive data dashboards are generated today in many companies using software tools. Dashboards address the needs of internal users who want to assess and monitor KPIs but do not have the data and analysis skills to do that autonomously. However, dashboards do not support a high level of interactivity, and many companies still rely on internal analysts building one-off data visualizations using programming languages.

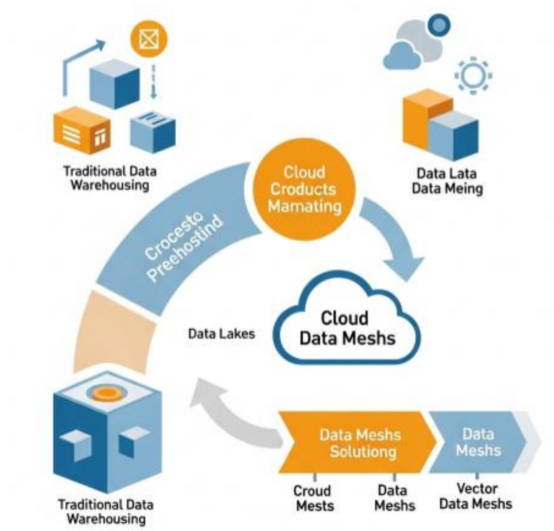

Fig 2: Data Products in the Modern Era.

The second type of product are point apps. They generate animations and visualizations of different data over the network, based on external requests. These products can serve external or internal users, either to entertain or to inform. Point apps usually are not very informative by themselves for a non-expert public, but they can serve the purpose of diverting attention to a particular external request, data point, or fact, like during a live sports event.Finally, some point apps are actually what we call APIs as Data Products. APIs have become widely popular in the last years because of their role as enablers of innovation. They allow external users to get information about the state of affairs or respond to their queries without diving deep into the knowledge domain.

3.1. Types of Data Products **Jio** What are data products? The coining of the term data product can be duly credited to the debut of the book for practitioners.

Because it popularized the idea that companies must commercialize their proprietary data, many companies have since launched data product-based business models providing data products to external customers, and have combined those data products with traditional data analytics services. Attention has and continues to focus on the gratifying engineering sleight of hand that allows technically advanced software engineers to gather website visits, enhance these with proprietary customer identification, and then monetize through displaying the most persuasive ads, using a proprietary model based on independently gathered user data. One of the authors was asked to query about external ad provider services, to evaluate whether one of the company's proposed decisions was impacted if the ads were collected from the company's two external ad providers.

However, the characterization of data products is vague. The literature has treated such products differently; for example, they have been defined as supervised ML services that use trained ML models to generate direct nurtured customer interaction value. Here, we follow an alternative approach, presenting a classification of different types of data products, based on the different types of data they deal with and the product format properties. Broadly, we can define data products as containers for the communication of information derived from our analysis of data in service to a purpose sought by the consumer.

3.2. Data Product Lifecycle Uniting all data product categories is the data product lifecycle. The agile worldview brought on by rapid technological advancement has reoriented corporate and product development decisions around accelerated, iterative, user-oriented practices of ideation, design, development, and deployment. The product management discipline has codified this definition of product development for software and services. Data products have their own unique life cycle within this general product lifecycle. Scientists and engineers extract meaning from derivative data, whether in the form of derived datasets or ubiquitous tools like models, recommendations, and predictions. Engineers and designers develop the user interface and experience of the health, risk, and discovery components of data products, addressing the non-linear, graphical warnings and strategy elements that spur user engagement and cycles of analyzer realizations of the increased insight data afford selection of novelty risk-reduction recommendations.

The core of any data product is a latent relationship expressed by an algorithm-driven derivative data component, whether what the algorithm-driven derivative data quantitatively expresses is a prediction about a future event, classification of a presently known state by class label, recommender system suggestion, simulated consequence of applying a particular decision or strategy, or model. Agents and analyzers benefit from the built-in health check facets of the latent relationships embedded in any of these derivative data expressions within the events, strategies, or models parameterized by the agentic task input data digital agent and analyzer interaction influences, coming from the algorithm-driven derivative data component that are usually visualized over time in a time series representation of the data being analyzed.

p-ISSN 0854-1418
e-ISSN 2503-426X

Equ 2: Cloud & Cost-Aware Optimization

$$C_{\text{total}} = \sum_{i=1}^{n} c_i \cdot r_i, \quad \text{subject to} \quad r_i \leq \bar{r}_i$$

- c_i: Cost per resource unit i (e.g., compute, storage)
- r_i: Resource allocated to model or pipeline i
- $\bar{r}_i$: Budget or quota constraint
- Optimized via $\min C_{\text{total}}$ with SLA/QoS constraints

4. Integration of Cloud Technologies

Due to the relatively technical nature of this section, we provide here a brief overview of what we mean by cloud technologies, and how those relate to Agentic AI or to AI Data Products in general. We then discuss benefits offered by the cloud in AI and present key technologies enabling Agentic AI, alongside Agentic methods enabled by these technologies. Given the criticism that can be levelled at cloud-centric services from an energy footprint point of view, we refer the interested reader to a discussion on the energy impact of AI engines and offer a brief description of green alternatives.

4.1. Cloud Infrastructure Overview

Cloud services express a fintech-first approach to leveraging organizational scale, infrastructural expertise, and technical resources to enable a model of just-in-time costs associated with business operations. Such costs are associated to a diverse portfolio of models, either as whole or in parts. Cloud services exist for most models during product creation, test, and runtime – think of DataOps tools for building and validating data-versioning and data-modeling pipelines, for experiment-tracking and collaboration; of MLOps tools for model training and validation; or for pipelines assembly and validation; Engine models used for runtime prediction; Test models used for software-in-the-loop validation of predictions; Monitoring models for runtime ops; and observability models to create the semantic glue required to interconnect that diverse portfolio of tools.

4.1. Cloud Infrastructure Overview

The public cloud has transformed how infrastructure is consumed in the last decade. It has made scalable compute, storage, networking, and database technology available via an application programming interface for virtually any use case. Cloud providers package together thousands of servers into highly available load balancers, hundreds of petabytes of storage into data lakes and object stores with various degrees of redundancy and durability, and scale out databases that are eventually consistent. The latest hardware technology is available by the hour managed via service layers for machine learning, big data, and high performance computing workloads. This offers companies of any size the opportunity to create data products and perform research using generalized AI investment without having to specifically invest in the physical infrastructure of data centers, deal with server lifecycle management, or manage the infrastructural scale-out to high volumes.

Cloud infrastructure has today matured from a relatively low performance and high latency distributed computing platform into a high performance alternative to specific physical data center investments that only extremely large companies with well understood capacity requirements can take advantage of. The performance of cloud infrastructure has grown tremendously in the last 10 years especially at the high end with significant investments in high end brand new physical generation and interconnect technology available to cloud consumers. Today, cloud services offer lower latency and higher throughput than physical infrastructure available through on-site collocation. At the same time, the costs per operation often favor cloud usage allowing cloud consumers to use innovative abstractions that offer exploration services that support fast product iteration.

4.2. Benefits of Cloud in AI

The cloud infrastructure not only provides the support for costs of capital-heavy product, service and system development but also acts as a robust central demand resource for data products and novel AI model types. The rapid innovation enabled in AI is also due to its synthesis with cloud technologies, which enhances generalizability, accessibility, commercial viability, privacy, flexibility, and security. Thus such integrated data product types facilitate a layered and compositional software paradigm, in which increasingly sophisticated products are created in a market by low-skilled builders using pre-trained AIs with the high-level capabilities that emerge via the training of these models on cloud economies of scale. The variety of technical advancements in product interfaces, model architecture innovations, distillation approaches for optimizing for cost and performance, novel data modalities, and other concepts all work together within this paradigm. This set of synergistic technologies enables user-friendly interfaces to the large numbers of domain-specialized pre-trained service or product models being created by AI research institutes and small and medium enterprises around the world.

Fig 3: Benefits of Cloud in AI.

Furthermore, the relatively low capital resource requirements for accessing these models create the opportunity for a huge number of builders in a market, as do the visual generation services which have opened up the generative modeling paradigm to every internet user. With productivity greatly improved, and the AIs no longer trusted black boxes accessible only to a select few. The way has been paved for the development of advanced AI-branded decentralized generative services, as explorations begin to synthesize AIs into multi-step creations.

5. Big Data Analytics

Big Data is a term that describes the explosion in the volume, variety, and velocity of data, whose storage and analysis constitute an increasing challenge to the computing and data industries. It generally refers to massive datasets that cannot be handled by traditional database management tools in a timely manner. Big Data computing also refers to distributed and parallelized storage and processing methods based on hundreds, thousands, or even millions of commodity computer cluster nodes. As massive amounts of data become available, there is growing recognition of the need to improve methods and technologies for sharing, fusing, understanding, and learning from that data. Making data available to, and developing methods that usefully integrate the expertise of, human analysts is an important goal for the development of AI and Data Products. Specifically, the goal of data analytics is to extract insight from data whether it be through mathematical modeling, human validation and reasoning, or some combination of the two. Within Big Data are a number of subfields relating to the handling of various dimensions of these datasets; the focus is often on specific sampling, storage, or acceleration techniques. Big Data is accessible due to the advancements in both hardware and software technologies. On the hardware side, clusters of multi-core CPUs are now affordable both for companies and for individuals. Cloud computing portals enable rent as needed services at low prices. On the software side, techniques have opened up this domain to a much larger audience. More recent techniques, in fact, focus on the usability of these systems even for those without formal training in computer science. NoSQL, Hadoop, streaming and temporal databases, and associated systems.

Equ 3: Domain-Specific Objective Functions

$$\min_{\hat{d}_t} \left[h \cdot \max(0, \hat{I}_t) + p \cdot \max(0, -\hat{I}_t) \right]$$

- $\hat{I}_t = I_{t-1} + O_t - \hat{d}_t$
- h: Holding cost, p: Penalty for understock
- $\hat{d}_t$: Predicted demand (ML output)

5.1. Big Data Technologies The ‘big data’ term in technological parlance usually reflects two defining attributes, namely, volume of data and the technological ability to mine value and knowledge quickly from such high-volume data. The chapter presents the definition of big data, various types of big data, a big data architecture, major big data technologies, and a broader ‘big data 3V’ architectural framework which builds upon the original 3V framework of big data for use by data technology companies. Such a bigger framework is more useful for big data analytics of data-enabled and data-driven SE data products than the original 3V framework.In preceding chapters, we introduced the concepts of data-enabled and data-driven SE data products. Such data products rely on big data of tight feedback loops which need to be mined for their business value. To do so, the feedback loops should be enabled with the right data technology stack best suited to perform the data product’s lifecycle, from data extraction through transformation and cleaning to product design and development, deployment, and product lifecycle management. Such a data technology stack which helps to create data products, design large-scale multidimensional databases, and perform advanced analytics on the data, undergoes constant and rapid change. New emerging companies bring in ideas, technology for faster, collaborative data preparation and access, advanced analytical ML and AI algorithms, and enterprise-scale implementations to the out-moded big data behemoths. The chapter discusses the choice of big data technologies from expert opinions, and how to place them in the data product lifecycle.

Fig 4: Big Data Technologies.

5.2. Data Processing Techniques Broadly speaking, two main techniques exist for processing data. Traditional data processing techniques attempt to extract useful information from data sets by applying mathematical, statistical and logical analysis. A range of algorithms are employed for example methods such as clustering (which attempts to group similar data points into clusters), classification (tries to assign a class label to an unknown data point based on a model trained on labeled points), collaborative filtering (predicts user preferences based on the preferences of similar users or items), regression (fits a function to a data set and hence can be used for prediction), etc. Visualization techniques provide another way of “processing” data. By showing data graphically it allows humans to grasp the informative structure quickly and assists in (or completely automates) the analysis of the data.

The task of learning an appropriate data model is made more powerful when the data model is revealed through data. In this case, the process of data processing switches from being one of analysis (involving traditional data models) to one of training a model (often involving machine learning or deep learning). For example, the explosions in the size of data sets and the complexity of the underlying structures has popularized unsupervised learning (where the model is trained on data without labels). For example, most machine translation methods train a statistical model on parallel corpus (where sentences in one language are paired with their translation in another language). But there is currently no way to easily create human-labeled data sets in every target language.

6. Machine Learning Applications

Machine learning is used to produce predictions or other kinds of inferences from rich, high-dimensional data. Considerable recent interest has focused on large, unsupervised models that create representations of intricate, high-dimensional input with very high levels of detail and fidelity, often for purposes of generating other samples. However, the most valuable work in the world of machine learning is in terms of supervision, where a specific outcome is predicted from a combination of prior data, instructions, and auxiliary hand-coded constraints. The structure that this relatively narrow approach provides is necessary for decision-making. Supervised models work best when careful logic is applied to choose the inputs, outputs, and any auxiliary components needed. In some key example problems, labeled training samples of inputs and outputs are available.

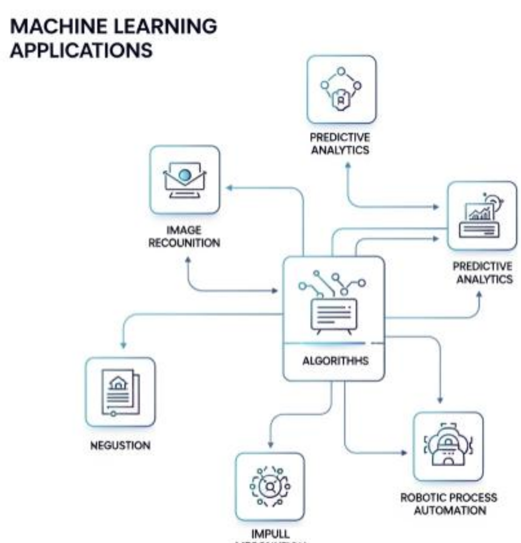

Fig 5: Machine Learning Applications.

In other situations, either only the input is given or there is no explicit training guidance. In the case where both input and outputs are provided, supervision improves what would otherwise be a clustering problem. However, if there is no output labeling available prior to the learning process, unsupervised structure-and-representation creation occurs. The most common usage in retail is still the first case — a large number of historically labeled data points are visually aligned and the relationship between the independent and dependent variable is learned. In this case, the outcome may be a lagged time series and the supervised models are structured as time series naive Bayes classifiers. Once the model has been calibrated, future outcome variables can be predicted based on auxiliary time series input.

6.1. Supervised vs Unsupervised Learning In any machine learning application, the crucial question is: What are the inputs and the outputs available for training and inference? The answer to this question classifies the problem into two distinct types: supervised and unsupervised. In a supervised learning system, the label, or output is provided for each input during training, and the goal of the machine learning model is to learn a mapping from the input to the output that exists in the data. For inference or prediction, we can only provide the input and the model generates the predicted output based on what it has learned during training. What the model predicts during inference is called its hypothesis. The predicted output can be continuous, discrete, or structured but its nature depends on the specific learning algorithm and is not controllable.

Many algorithms have been invented for supervised learning, including neural networks, support vector machines, and Gaussian processes. Each one of these algorithms and their associated architectures support a specific type of prediction, with some predicting continuous outputs, others predicting discrete outputs, and others predicting structured outputs. In supervised learning, prediction is defined by the training labels so the term prediction is a little misleading. In supervised learning, we are really modeling the joint distribution of the input and the labels. When the labels are continuous, then supervised learning models the conditional distribution of the labels given the inputs. When the labels are discrete, then it models the joint distribution of both the discrete labels and the inputs. The output is obtained by a hard maximum of structured discrete labels or the maximum likelihood estimate of the output density. The conditional density approach of learning is most famously associated with discriminative algorithms, while the joint density approach is associated with generative algorithms but both strategies can all be used for the actual output generation in supervised learning.

6.2. Predictive Analytics in Retail Predictive analytics refers to an array of statistical techniques and tools that utilize historical and existing data in order to identify the likelihood of possible future outcomes, based on patterns identified in the data. At their most fundamental level, predictive analytic models analyze recurrence and how the past can help forecast future events. Predictive analytics usually involves classical statistical models and forecast tools as well as advanced analytic techniques such as data mining, predictive modeling, machine learning, and artificial intelligence. Predictive analytics involves various techniques from statistics, machine learning, data mining, and game theory in a coherent and unified manner. These methods together facilitate the ability to predict outcomes, and predictive analytics deploys these concepts in an applied manner. Predictive analytics is therefore a data product, which organizations can leverage to create decision support systems that boost their decision making accuracy and efficiency to shape better business outcomes.

Predictive analytics can establish a data-driven structure to process decisions, rather than relying on ad-hoc gut instinct decision making. Data-driven decision making with predictive analytics can lead to increase in overall revenue, and reduction of costs and loss. While a retail organization operates its day-to-day business, it also collects enormous amounts of data at the same time. This data includes customer data, sales transaction data, supplier data, and shipment and logistics data to name a few. This data contains useful information that when processed and analyzed, can help retailers react to changing market dynamics quickly and decisively. Predictive analytics techniques can help retailers churn useful information from these data sets, and use it to their advantage when making decisions around the retail business. Such decisions can include determining: what products to carry, what price points to set, which promotions to offer, what product designs to roll out, and how to manage inventory levels.

7. Conclusion

Incorporating cognitive agency into commercial Artificial Intelligence (AI) systems is still in its infancy, with the deployment of parts stacking up daily while domain-specific collaborative systems are on the uptick as well. The design of agentic AI systems is about simulation of both social and motivational behavior, along with the incorporation of the use of other agencies as tools, with the creation of products as a directed behavior. These systems will be the attractors for use as tools in many collaborative use cases. Ignoring the challenges and consequences of agentic AI will slow innovation and increase the risk of negative consequences. By recognizing that the impact of these systems will be as positive and enriching as traditional digital systems and non-cognitive toolkits, the trend towards the use of domain-specific agentic AI products can be encouraged.

The trend towards simpler systems is accompanied by the normalization of AI, through its use in the same types of problem-domain activities by multiple enterprise class users. This will shorten and ground the design loops for new innovative mobile toolkits. Normalization will also shift focus away from the perception of threat, towards the redesigned relationship between humans and the new species of collaborative tools and work partners. Finally, a long view towards future work will welcome the enabled expansion phase of humanity's engagement with intelligent systems, followed by the phase of symbiosis with intelligence, before finally settling over subsequent centuries into the phase of unconsciousness with intelligence.

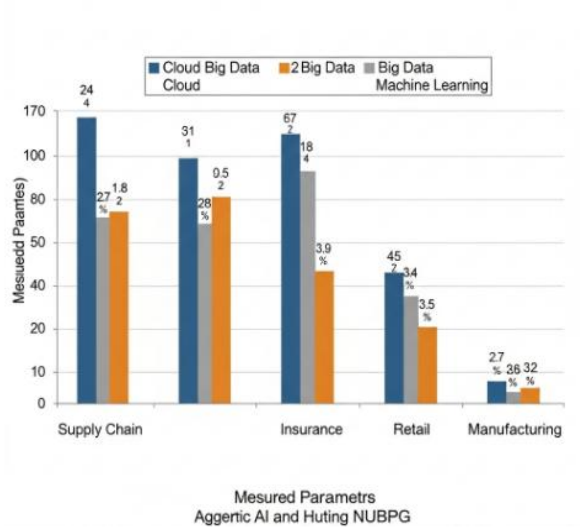

Fig : A Unified Framework for Agentic AI and Data Products: Enhancing Cloud, Big Data, and Machine Learning in Supply Chain, Insurance, Retail, and Manufacturing.

7.1. Future Trends Agentic AI is an emerging field that explores the decision-making capabilities of AI systems. This emerging design approach has the potential to unlock previously non-automatable markets while increasing the user experience of existing ones. However, it also enables new dangers, such as the establishment of power and influence over their users, the potential of deception and the spread of misinformation to increase political, social and economic inequality. The development of robust guardrails is critical to ensure the growing application of Agentic AI systems can occur in a manner that is consistent with the values of a civil society. We propose to use the four laws model to provide design guidelines.

In the context of data products, we caution against Agentic AI products that allow a selector to be a separate entity from the user but are a faceless recommender, amplify their motivations without being subject to a feedback loop, synthesize content by deception or exaggeration, and manipulate network influence or belief. Towards this end, our vision for the future of data products is twofold. The first part is a call to the rich ecosystem surrounding all of the existing recommendation engines that power so much of the internet today. We would like to see more of them designed explicitly with, and re-architected with the potentials and limitations of technology. Such a process would produce a superior set of products. Further and more conversely, the potential for Agentic AI data products is for new ecosystems of products across use cases and workflows. These may combine agentic decision making over different resource modalities and communication modes into a more automatically responsive web service that grows itself over time.

8. References

[1] Kannan, S., & Yellanki, S. K. (2025). Synthetic Cognition Meets Data Deluge: Architecting Agentic AI Models for Self-Regulating Knowledge Graphs in Heterogeneous Data Warehousing.

[2] Sheelam, G. K., & Komaragiri, V. B. (2025). Self-Adaptive Wireless Communication: Leveraging ML And Agentic AI In Smart Telecommunication Networks. Metallurgical and Materials Engineering, 1381-1401.

[3] Motamary, S. (2025). Empowering Retail Oss/Bss Platforms With Agentic Ai And Scalable Data Engineering. Metallurgical and Materials Engineering, 1361-1380.

[4] Meda, R. (2025). Dynamic Territory Management and Account Segmentation using Machine Learning: Strategies for Maximizing Sales Efficiency in a US Zonal Network. EKSPLORIUM-BULETIN PUSAT TEKNOLOGI BAHAN GALIAN NUKLIR, 46(1), 634-653.

[5] Somu, B., & Inala, R. (2025). Transforming Core Banking Infrastructure with Agentic AI: A New Paradigm for Autonomous Financial Services. Advances in Consumer Research, 2(4).

[6] IT Integration and Cloud-Based Analytics for Managing Unclaimed Property and Public Revenue. (2025). MSW Management Journal, 34(2), 1228-1248.

[7] Annapareddy, V. N., Singireddy, J., Preethish Nanan, B., & Burugulla, J. K. R. (2025). Emotional Intelligence in Artificial Agents: Leveraging Deep Multimodal Big Data for Contextual Social Interaction and Adaptive Behavioral Modelling. Jai Kiran Reddy, Emotional Intelligence in Artificial Agents: Leveraging Deep Multimodal Big Data for Contextual Social Interaction and Adaptive Behavioral Modelling (April 14, 2025).