

## RESEARCH ARTICLE

# Agentic AI Systems for Autonomous Healthcare Data Pipeline Orchestration and Optimization

Triveni Kolla <sup>1\*</sup>, Bindu Madhavi Mangalampalli <sup>2</sup><sup>1</sup>Senior Business Intelligence Developer ORCID ID: 0009-0009-2761-5287<sup>2</sup>Data Engineering Architect Team Lead ORCID ID: 0009-0001-1070-3856

\*Corresponding Author

Article History

Received: 10.09.2025

Revised: 14.10.2025

Accepted: 05.11.2025

Published: 20.12.2025

**Abstract:** **Background:** Data is undoubtedly the most relevant asset in the digital transformation of businesses. In healthcare, data pipelines enable the automated orchestration and optimization of data flow among diverse systems and stakeholders. Properly designed, they provide a mechanism to abstract and fulfil the resource-related needs of the processes they support, including ingestion, analysis, sharing, and publishing. Nevertheless, healthcare data pipelines are still implemented manually or with ad-hoc automation, which threatens their reliability, flexibility, and efficiency. Integrating Agentic AI (Artificial Intelligence) appears to be the solution for their automation and optimization. Agentic AI systems reproduce the autonomous, self-organizing, decision-making, and multi-role properties of human archaogenic agentic behaviours. In the context of healthcare data pipelines, these properties enable automatic data ingestion and sharing whenever and wherever needed; participation as data providers, consumers, or orchestrators; dynamic adaptation to changes in resource-related needs, costs, or availability; and multi-agent collaboration with associated role allocation. Supporters of this approach claim that Agentic AI will favour businesses in their race towards becoming data-driven by allowing managers to focus on their core business and outsource resource-related problems to self-organizing control loops operating over business processes. Nevertheless, Agentic AI systems have yet to be implemented and tested. This gap deserves attention, especially in a domain as heterogeneous and complex as healthcare, where patients' lives depend on the accurate and reliable availability of data. The following presents an in-depth analysis of healthcare data pipelines and identifies how Agentic AI can support their automation and optimization.

**Keywords:** Agent-based Systems, Autonomous Agents, Agent-oriented Software Engineering, Healthcare Data Orchestration, Healthcare AI.

## INTRODUCTION

The journey from data collection to analysis and interpretation in modern healthcare involves a series of interdependent steps, including the monitoring of resources, state maintenance, problem detection, and action triggering. Within an organizational context, these steps are orchestrated by humans and assistive technologies that often lack awareness of their environment and interdependencies. While data management and pipeline technologies are widely deployed to assist with these tasks, unforeseen upstream or downstream events often lead to pipeline failures, unexpected delays, or service-level agreement violations.

The introduction of intelligent agent technologies, capable of continuously monitoring ongoing processes, recognizing changes, re-planning resource allocation in real time, and executing corrective actions autonomously, provides an opportunity to increase process resiliency and reduce human supervision effort. These research efforts focus on emerging agent-based approaches that employ autonomous intelligence at the orchestration level for the explicit purpose of autonomous data pipeline operations. The application of intelligent autonomous agents to orchestrate the flow of heterogeneous data among interacting resources creates agentic data pipelines. The term agentic refers to the ability to adaptively sense environmental changes and, when required, take autonomous corrective actions. Agentic data pipelines are characterized by continuous operation and planning and managerial awareness actions that exploit interaction knowledge for real-time performance improvement and task scheduling.

The emergence of intelligent agent technologies is transforming the way data-intensive systems are designed and managed. By continuously monitoring processes, detecting environmental and operational changes, dynamically reallocating resources, and autonomously executing corrective actions, these agents significantly enhance process resilience while reducing the need for constant human oversight. At the orchestration level, agent-based approaches introduce autonomous intelligence that enables data pipelines to operate in a self-managing and adaptive manner. When applied to the coordination of heterogeneous data sources, services, and computational resources, this paradigm gives rise to *agentic data pipelines*. The term *agentic* emphasizes the system's capacity to perceive contextual shifts and proactively respond through informed decision-making. Such pipelines operate continuously, combining real-time planning, scheduling, and managerial awareness to optimize performance. By leveraging knowledge of interactions across distributed components, agentic data pipelines enable more robust, efficient, and scalable data operations in dynamic environments.

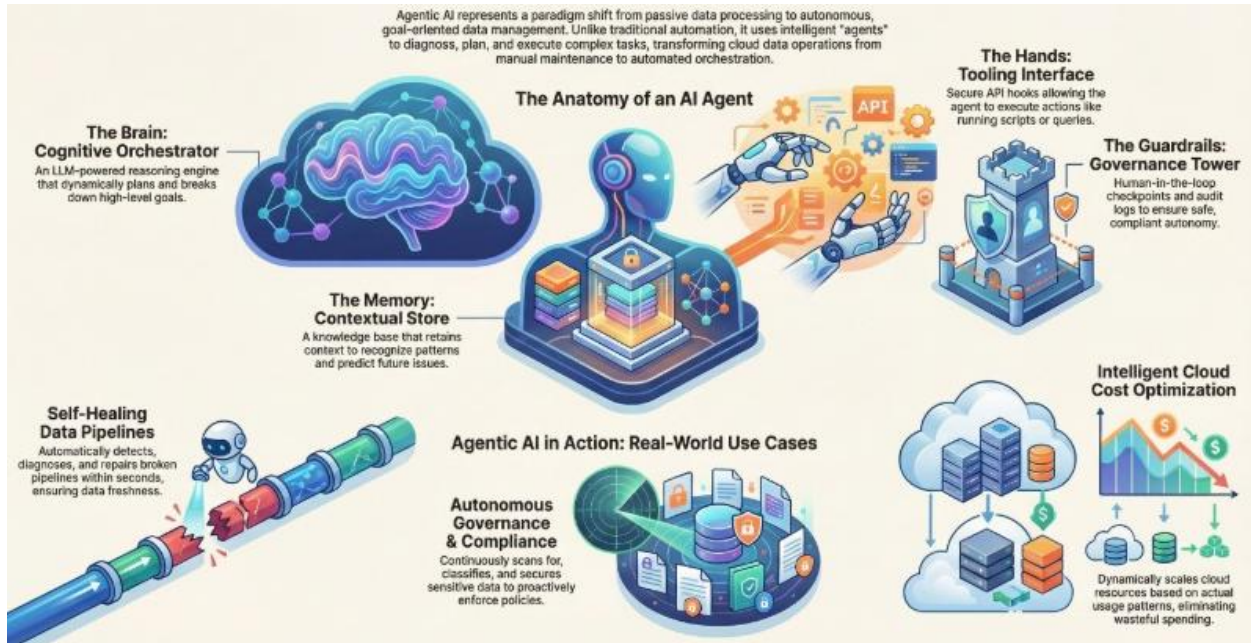


Fig 1: Agentic AI in Cloud Data Management

### 1.1. Background and Significance

continuous and equal access to digital devices and health data that characterize the current era has opened new opportunities for individuals and companies operating in the digital health market. Digital health encompasses a variety of tools, products, and services, including mobile, software, hardware, and services related to disease prevention, diagnosis, and pathology treatment, therapy, prognosis, and management. These solutions capitalize on the growing demand for digital Health 2.0 solutions, allowing users more agency, control, and active participation in caring for their health. While the 2.0 phase focused on Web 2.0 technologies and user-community engagement, a more current Health 3.0 phase is maturing, stimulated by the advances of IoT and sensor-based technologies.

At the same time, there is the need for scalable digital transformation across health delivery systems to ensure health equity and resilience. At this level, Smart Health projects support the automated, remote, and continuous monitoring of patients with chronic health conditions and other patients released from hospitals. Patients should be able to take ownership of their own health conditions, receiving timely alerts and education to avert negative health events. Possessing adequate health data increases the chances of having better health outcomes. From a business perspective, enabling better health service delivery results in lower operational costs and increases the potential for new business opportunities. A wide market of third-party developers can create and deploy new health applications, since data are shared through open APIs. However, despite the emergent demand and proposals, there are a lack of fully operational Smart Health services in practice and a lack of commercially-oriented solutions supporting the creation of Smart Health services by third parties.

#### Equation 1: Formalizing an agentic healthcare data pipeline as a scheduling problem

A data pipeline can be modeled as a directed acyclic graph (DAG):

- Tasks:  $\mathcal{T} = \{1, 2, \dots, N\}$
- Precedence edges:  $(i \rightarrow k) \in \mathcal{E}$  meaning task  $i$  must finish before  $k$  starts.

Let:

- $s_i$  = start time of task  $i$
- $d_i$  = execution duration of task  $i$

Then “precedence constraints” are:

$$s_k \geq s_i + d_i \forall (i \rightarrow k) \in \mathcal{E}$$

#### Derivation (step-by-step):

1. If  $i \rightarrow k$ , task  $k$  cannot start until  $i$  completes.
2. Completion time of  $i$  is  $s_i + d_i$ .
3. Therefore  $s_k \geq s_i + d_i$ .

### 1.2. Research design

Two main lines of inquiry are pursued. First, current pipelines and frameworks are reviewed. A second line of research looks for evidence indicating agentic systems are feasible for delivering autonomy while still respecting prevailing data regulation.


Investigations into the use of agents for orchestrating the components of healthcare data pipelines highlight the design and function of a set of primitives, describe the mechanisms for allocating, controlling, and optimizing groups of agents, and assess the sensitivity of decision support, risk, and scheduling considerations to computational resource availability. The results suggest the orchestration role can be fully automated.

## 2. Foundations of Agentic AI in Healthcare Data Pipelines

The medical data landscape is characterized by diverse sources, types, and structures, complicated further by the full range of well-known interoperability issues pertaining to different clinical, radiology, biomarker, genomic, registry, and patient-generated data. Infrastructures modelled on the data lake, data island, or data hub metaphors promise acquisition from any number of heterogeneous, possibly remote sources, after which the focus shifts to orchestrating processing and storage in an optimized manner for a plethora of currently required, emerging, and future uses.

While such multi-platform data-integration scenarios exercised the multi-agent systems (MAS) research community a decade ago, the climate crisis and its associated risks with respect to health and well-being now demand a shift from a predominantly supply-side focus on ingestion and accessibility toward a more demand-driven approach. This entails not only a responsive system that acts automatically when an evolving need manifests but also one that allocates different processing and analysis roles to an ensemble of collaborating agents. Such an inversion of the instantiation-deployment paradigm means the question of data availability recedes into the background, with the capacity to pull together whatever data are needed—when and for what—becoming the challenge.

Building on this inversion, a demand-driven paradigm reframes multi-agent systems not as passive orchestrators of pre-integrated data streams but as adaptive sense–reason–act collectives that co-evolve with the problems they are meant to address. Rather than beginning with what data are already available, the system begins with an articulated or inferred decision need, decomposes it into analytical sub-tasks, and dynamically recruits agents with complementary competencies—ranging from data discovery and provenance assessment to modeling, uncertainty quantification, and ethical oversight. In this way, data acquisition becomes a consequence of reasoning rather than its prerequisite. Such architectures privilege timeliness, contextual relevance, and interpretability over sheer volume, enabling MAS to operate as living infrastructures that continuously reconfigure themselves around emerging climate-related risks and societal priorities, ultimately supporting more anticipatory, equitable, and actionable responses.



**Fig 2: Agentic AI in Healthcare**

### 2.1. Data Ingestion and Interoperability

Agentic systems are designed to orchestrate the healthcare data integration process, enabling automatic monitoring, remediation, and optimization capabilities. However, in real-world scenarios, human intervention is still required due to errors related to data ingestion and source inter-communication, especially when the match between sources is poor. In this sense, the early stage of Data Pipeline formation continues to be manual and complex. New AI technology must enable a more efficient architecture for this process. A remote healthcare patient monitoring system that analyzes sensory data and

provides semantic alerts is used as an illustrative case study. The system architecture requires communication between two cloud-edge infrastructures. A data pipeline deployed on a Cloud environment is responsible for ingesting, processing, and storing sensory data generated on the Edge.

Communication is established using Web Services, with sensitive data being stored in a Private Cloud environment. The proposed AI-based technology detects dynamic changes in data sources and pipelines. When a healthcare actor joins the monitoring process, a semantic data pipeline is automatically instantiated. The system performs data supervision and ontology-verification tasks. Agent technology and natural language processing are employed to check actors' qualifications for pipeline participation. When the data exchange between two infrastructures has poor semantics, human intervention is required, and an agent recommends an appropriate algorithm to process the data. The solution has been validated in a real patient monitoring system and has the potential to ease the pipeline establishment of heterogeneous IoT cloud-edge infrastructures.

## Equation 2: Agent assignment and “resource-aware scheduling”

Let:

- Agents (workers):  $\mathcal{A} = \{1, \dots, M\}$
- Binary assignment variable:

$$x_{ij} = \begin{cases} 1 & \text{if task } i \text{ is assigned to agent } j \\ 0 & \text{otherwise} \end{cases}$$

Each task must be assigned to exactly one agent:

$$\sum_{j \in \mathcal{A}} x_{ij} = 1 \forall i \in \mathcal{T}$$

### Derivation:

1. A task must run somewhere (some agent).
2. “Exactly one agent” means choose one  $j$ .
3. Binary selection across all agents sums to 1.

### 2.2. Semantic Annotation and Provenance

Semantic

annotation links content to machine-actionable semantics, crucial for interpretability and retrieval. Resource-aware pipelines timely deploy computing resources across the entire pipeline, ensuring efficiency and efficacy. Knowledge graphs digitize domain expertise to implement proper semantic annotations and enable provenance generation. Data provenance tracks ingested data sources, pipeline processing, and produced output. Data custodians monitor provenance updates to validate semantic annotations.

Agentic systems integrate feedback loops and control mechanisms for constant operational adjustment. Agent action and reaction modules execute pipeline tasks, utilizing the knowledge graph to manage problem-specific requirements encompassed in Quality of Service templates. QoS monitors, modeled as discrete-event systems, guarantee adherence to user-defined pipeline execution constraints, serving as reactive assistant agents.

## 3. Architectural Frameworks for Agentic Pipelines

Architectural frameworks for agentic healthcare data pipelines provide the structures and organization needed to support agentic orchestration. They enable execution across multiple organizations in response to changing circumstances, including the usage of intermittent, non-continuous, or irregularized services. Multi-agent collaboration arises from the roles of the agents involved, covering all aspects of both the data-pipeline operation and external interactions with governance authorities. Institutionally based control loops capture simple agentic weather and scheduling information for use by role-allocation services, while time-driven deployment aligns business and data-generation cycles.

Learning-based detection of external events or conditions enables feedback to influence the orchestration and behavior of waterways and associated gardening resources. The same concepts, when extended across the organization boundary, facilitate the discovery, compliance-checking, and reuse of data across multiple organizations. Role-based operation ensures that the externally visible pipeline remains static, that interactions with partners remain clear and well-defined, and that appropriate decisions on partner selection can be made at schedule- and now-time.

Architectural frameworks for agentic healthcare data pipelines establish a cohesive foundation for orchestrating intelligent, adaptive data flows across complex, multi-organizational ecosystems. By defining clear roles, responsibilities, and interaction patterns, these frameworks enable autonomous and semi-autonomous agents to coordinate data ingestion, transformation, governance, and delivery even when services are intermittent, irregular, or dynamically provisioned.

Embedded control loops capture contextual signals—such as operational “weather,” resource availability, and scheduling constraints—to inform role allocation and time-driven deployment strategies that align analytics workloads with clinical, operational, and business cycles. Learning-based mechanisms further enhance resilience by detecting external events or environmental changes and feeding this intelligence back into orchestration logic, allowing pipelines to adjust behavior proactively. Extending these concepts across organizational boundaries supports secure discovery, compliance validation, and reuse of shared data assets, while role-based operation preserves a stable external interface and well-defined partner interactions, ensuring consistent governance, predictable integration, and informed partner selection at both planning and execution time.




Fig 3: A Guide to Building Agentic AI Workflows

### 3.1. Multi-Agent Collaboration and Role Allocation

Despite the availability of agent-based programming frameworks, a clear architectural blueprint for agent-based data pipeline control has yet to be established. Feedback kinelligence. Multi-agent frameworks usually support the implementation of different components in distributed settings, allowing for a collaborative approach to data pipeline orchestration. A specific architecture proposed for data pipelines focuses on centralized rollouts for performance or unit commitment optimization but does not explore their control architecture. In the context of traditional control loops—for example, control theory—the distributed optimization problem uses the gradient of a loss function to dynamically incentivize agents toward the central solution without requiring a centralized orchestrator.

The use of multi-agent systems for resource scheduling constitutes a prevailing paradigm in data pipeline control, naturally lending itself to QoS monitoring and control. However, data pipelines differ from traditional grid scheduling in that the data-generating processes and the data haven metadata/annotation providers are independent, and row-by-row construction covers all potential flows. The requirements these data-generating processes impose on datacentric computing environments necessitate a feedback mechanism that teachers monitoring agents how to better serve such requests. A dedicated monitoring agent also enables QoS assurance in terms of read access delay, bandwidth, and response time; the additional role once again activates these features when they are down or underperforming.

#### Equation 1: Objective functions: QoS-aware, QoI-aware, resource-aware scheduling

- **QoS-aware scheduling** (optimize travel cost + response time)

Agentic AI Systems for Autonomo...

- **QoI-aware scheduling** (minimize integration cost remember minimum QoI)

Agentic AI Systems for Autonomo...

- **Resource-aware scheduling** (respect budgets/constraints allocated by higher-level manager)

Agentic AI Systems for Autonomo...

A unified multi-objective (common in practice) can be written as a weighted sum:

$$\min_{x,s} \quad \underbrace{\alpha \sum_i \sum_j x_{ij} \ell_{ij}}_{\text{data travel cost}} + \underbrace{\beta \sum_i \sum_j x_{ij} \frac{w_i}{v_j}}_{\text{compute/response time}} + \underbrace{\gamma \sum_i C_i^{\text{grant}}}_{\text{privacy/access granting cost}} - \underbrace{\delta \sum_i Q_i}_{\text{data quality (QoI)}}$$

Where:

- $Q_i$  is the achieved QoI contribution of task  $i$
- $C_i^{\text{grant}}$  reflects the paper's point that privacy-preserving scheduling can include **cost of data-access granting**.

### 3.2. Control Loops and Feedback Mechanisms

Multi-

agent systems enable autonomous execution of healthcare data pipelines in a decentralized manner. Such a system requires a defining architecture and set of control loops that support its various control objectives. Multi-agent systems require roles, which can either be defined statically or determined dynamically at runtime.

Control loops realize feedback mechanisms for the underlying functions of the agentic systems and link the agents to the layers of a healthcare data pipeline. Feedback loops regulate shorter-term decisions such as scheduling, alter resource and node configuration, aggregate multiple scheduling constraints in a resource-aware manner, and delegate these resource allocation and scheduling tasks to the feeder agents whose function is primarily resource management. Longer-term or periodic optimization of larger parts of the pipeline for QoS, reliability, compliance, and candidate provider selection for invocation can also be integrated. Conceptual Cloud technologies enable Cloud environments to be utilized for the hosting, management, and execution of workloads of various nature, and application domains such as SaaS, PaaS, and DaaS have emerged.

## 4. Decision-Making and Optimization

Successful operation of autonomous pipelines necessitates a suite of online decision-making and optimization techniques. The scheduling of tasks and control loops adversely affecting resources use and data quality must be managed. Facilities for respecting quality-of-service and compliance constraints must also be provided. If resource-aware scheduling implemented is implemented correctly, it can also indirectly optimize the quality of data produced, even if data quality is not explicitly considered in the pipeline workflows. These methods complement other aspects of agentic pipelines, yet these aspects do not require agentic AI support. However, scheduling solutions depend not only on the agentic nature of pipelines but also on the designated application domain.

The scheduling of multi-task pipelines is an essential consideration for autonomous multi-agent data pipelines operating without human intervention. Such systems typically incorporate a two-level architecture, where an upper-level manager allocates execution budgets to lower-level response agents. Resource-aware scheduling attempts to optimize the monitoring phase of a multi-task monitoring pipeline while satisfying the planned monitoring costs assigned by the higher-level resource manager. The aim of resource-aware scheduling is to maximize the quality of the data being collected while satisfying user-defined quality-of-service constraints and completing the task within the assigned monitoring budgets. Scheduling solutions in agentic pipelines are resource-aware because their decisions respect the resource allocations made at a higher decision level.

Resource-aware scheduling assumes that both task execution and data delivery processes can be managed and planned independently. The master agent monitors the statuses of the data-delivery agents and allocates them to response tasks while considering users' quality-of-service requirements. The allocation of data-delivery agents to monitoring tasks can be seen as a resource-assignment problem. Alternatively, the quality of service of data-delivery agents can be viewed as a type of cost that monitoring tasks need to incur to fulfill their data-delivery requests. Properly managing this quality-of-service assignment aspect can also facilitate the optimization of data quality without requiring an explicit optimization criterion in the pipeline task.




Fig 4: Decision-Making and Optimization of Agentic AI Systems

#### 4.1. Resource-Aware Scheduling

Rescheduling of awoken Data AgEntry Roles must be based on the Quality of Service criteria guaranteed to the respective Data Pipelines. Timeliness requirements can be expressed in the form of a scheduling discipline. The orography of the communication network can be exploited by Lane Agent Interactivation, by means of Path Calculation and Optimization, to minimize delay in time-critical Data Pipelines.

The scheduling policy decides the Timing of next activations of the Data AgEntry Roles with respect to their Quality of Service requirements. Such Timing can be expressed in the form of Periodic Activation (P), Sporadic Activation (S), Aperiodic Activation (A), or Other Forms of Activation (O). The periodic activation discipline enforces certain properties on the interactions within the Data Pipeline, and is therefore the most suitable for Pipelines whose Delay Tolerance Level is Critical.

#### Equation 4: Control loops and feedback (agentic adaptation)

- State (what the pipeline “looks like now”):  $s_t$
  - Action (what agents decide):  $a_t$
  - Exogenous events (workload changes, failures):  $w_t$
- $$s_{t+1} = f(s_t, a_t, w_t)$$

A typical optimization policy chooses  $a_t$  to minimize expected future cost:

$$a_t = \pi(s_t) = \arg \min_a \mathbb{E}[J(s_t, a, w_t)]$$

#### Derivation:

1. Control loop observes state  $s_t$ .
2. Takes corrective action  $a_t$  (paper: “autonomously executing corrective actions”).
3. Environment evolves into  $s_{t+1}$ .
4. Choose actions that minimize a cost capturing QoS/QoI/compliance.

4.2. Quality of Service and Compliance Constraints  
Approaches to agent-based orchestration and optimization of healthcare data pipelines often consider Quality-of-Service (QoS) requirements, which address the performance and availability of the operations conducted by the various components, or the Quality-of-Information (QoI) requirements, which specify the properties of the data at a higher level of abstraction than the properties of the services. Agentic AI also enables the modeling of other forms of constraints and criteria, for example security compliance constraints, or high-level policies needed to manage information-sharing among participants with different roles involved in the pipeline execution.

QoS-aware scheduling handles the allocation of computing resources such as Central Processing Unit (CPU) and memory by optimizing the total travel cost for input data and the system response time. QoI-aware scheduling selects a subset of input datasets to minimize the total integration cost while preserving the minimum user-specified QoI-value for the output QoI-type. Resource-aware scheduling focuses on optimizing the execution schedule with respect to userspecified and agent-managed resource constraints and QoS requirements. For a healthcare scenario, the privacy of patients is also critical,

due to data-sharing regulations, and a privacy-preserving scheduling technique considers the cost of data-access granting when allocating operations to the agents that execute the operations and manage the exchanged data.

## 5. Data Governance, Privacy, and Security

The analysis, sharing, and publication of sensitive data demands distinct governance processes, which must be supported by AI-based methods and tools. These processes should cover both the orchestration and optimization of pipelines as well as the optimal trade-off between data accessibility and privacy protection. Privacy-preserving computations—e.g., Secure Multi-Party Computation (SMPC), Homomorphic Encryption, or Differential Privacy—enable the analysis of sensitive data without direct disclosure. Furthermore, agents in charge of decision making should minimize or avoid thus sensitive data travel.

Privacy-preserving techniques do not render distributed data sharing entirely safety-free, though. In fact, they can make sensitive data even more vulnerable, because users may share their sensitive data hinting at, e.g., Encrypted Database as a Service. For this reason, guarantees on the security of the agents and data providers involved in the orchestration and execution of the operation are of the utmost importance. Auditability and explainability methods and techniques help auditors and users understand the how and why of an AI-based decision (e.g., in the sharing or publication of sensitive data). Security considerations aim to convey the “securely” property of the proposed AI agents and solutions, ensuring that no unauthorized party is able to change or to sample the data of sensitive data holders.

### Equation 5: Privacy-preserving computation (equation-level foundations)

- True query:  $f(D)$
- Sensitivity:  $\Delta f = \max_{D, D'} \|f(D) - f(D')\|$  for neighboring datasets

Laplace mechanism:

$$\mathcal{M}(D) = f(D) + \eta, \eta \sim \text{Lap}\left(\frac{\Delta f}{\epsilon}\right)$$

### Step-by-step derivation idea:


1. Privacy goal: hide any single patient’s influence.
2. Worst-case change from adding/removing one record =  $\Delta f$ .
3. Add noise scaled to  $\Delta f/\epsilon$ : smaller  $\epsilon \rightarrow$  more noise  $\rightarrow$  stronger privacy.

### 5.1. Privacy-Preserving Computation

Privacy-

Preserving Computation represents a subdomain of Data Privacy that applies to specific operations in the Healthcare and Life Sciences sectors. Its primary purpose is to protect the confidentiality of a patient’s data while allowing it to be shared with the appropriate participants within the networking environment for analysis in accordance with the role-based access policies assigned to those specific participants. Such processes involve the usage of cryptographic techniques, secure hardware, and randomization, often referred to as third-party computation, or Trusted Third Party (TTP) privacy-preserving solutions. Solutions that do not involve the presence of a TTP rely on cryptographic primitives that allow calculations to occur without direct knowledge of the input data, while maintaining particular assurances in the final output. Common cryptographic primitives supporting these solutions include: Secure Multi-Party Computation (MPC) and Homomorphic Cryptography.

Secure Multiparty Computation, or Secure Multi-Party Computation, is a family of cryptographic techniques that expose a specific publicly verifiable function of all the inputs without revealing the individual inputs to any other entity. Entities share their data with a collection of other parties, where each party has access to a subset of the data and are not trusted with the data they collected from other participants. A trusted third party or an institution is chosen to derive the output from the inputs of the participants and the participants can only verify the result of the calculation. Secure Multi-Party Computation has an inherent solution to the problem of a TTP since the TTP has enough information to reconstruct the original input data. Homomorphic Cryptography permits the execution of a variety of calculations on cryptographic censorship without actually decrypting the ciphertext. A fundamental property of homomorphic encryption is that the level of a ciphertext does not affect the operation to be executed but only the size of the resulting ciphertext.




## 5.2. Auditability and Explainability

The inner workings of Agentic systems must be transparent to facilitate control by humans or AI-based regulators. For this purpose, an infrastructure of social networks of trust or social ALGs can be established, whose purpose is to provide trustworthy information to users, regulators, or third parties that inquire about the operations of an organization and about the AAIIs that have learned from them and are orchestrating their operations, and that need to prove compliance with GDPR rules. Such auditability is warranted through the use of process mining techniques. With respect to explainability, it has been shown that the execution of multi-agent systems can be semantically annotated such that explanations of opaque actions can be produced by revisiting the semantic annotation.

Nonetheless, enhancing transparency through audit reports, exposing the role allocation process, and providing guidance on how to compose more explainable Agentic systems and generate evaluations of their probative power, especially in critical settings, are still open research questions. Techniques belonging to XAI are expected to have a significant role in this respect, as will approaches based on current social standards of intelligence and ethical principles that consider not only the action outcomes, but also their motivations and intentions.

## 6. Conclusion

The development and deployment of agentic systems supporting the autonomous orchestration and optimization of healthcare data pipelines represent an important step towards overcoming current limitations of the field. Agents have been shown to address key challenges throughout the data lifecycle: facilitating interoperability between disparate data sources, ensuring data quality through suitable validation and cleaning procedures, making automated decisions about storage and processing resource provisioning, and adapting the pipeline's behaviour in response to changing operating conditions through proactive feedback loops. These capabilities empower healthcare data pipelines to operate independently, reducing the burden on human operators or system administrators and enabling continuous analysis of incoming data streams. Nevertheless, further research is needed to explore additional dimensions required by an agentic approach, such as the integration of optimization techniques for resource-aware scheduling, the ability to guarantee Quality of Service and compliance with security and privacy regulations, and mechanisms to ensure system auditability and explainability. In addition to such conditions for full autonomy, the use of multi-agent systems that collaborate towards a shared goal must also be extended to more advanced roles and patterns of interaction. Emerging trends in data engineering, such as the growing use of Data Mesh and Data Fabric concepts, are opening new avenues for agentic healthcare data pipeline architectures capable of seamlessly responding to the challenges posed by the massive increase in data volume, variety, and velocity and the proliferation of heterogeneous use cases across multiple stakeholders in the healthcare ecosystem.



### 6.1. Emerging Trends

On the far horizon, quantum computing will shift the physical limits of problem-solving computational systems from the speedup of exponentially large problems by a specific solver to the simultaneous solution of widely different types of problem for data analysis and pipeline orchestration. Massive quantum circuits will evolve, supporting completely new paradigms of quantum-enhanced cost function optimization, helping multi-agent pipelines jointly converge and trial-and-error convergence in quantum-enhanced training of task and domain-reinforced learning agent ai-processes. Beyond operation, a complex pipeline covering locations used by its agents will be feasible with sensor support. Unmanned agent-combot teams, capable of carrying their own networks and microbs, would provide the ultimate practical Maschine-AI candidate: a general purpose persistent agent-combot that could culturally reflect the population served—like war robots on a grand scale. In this bridge region, deception, self-dependence, and planning in agents than in their own staff will be remedies for current human-organizational and institutional toxicity. Quantum-and-sensor enhanced information middle agencies, white and black hat, will quickly emerge in the service of every agent-capitalist. In the service of crime detection, prevention, and prosecution preventive-and-blue hat-mids will counteract, probe, and by quire-the mandates, infrastructures, and behaviours of preventive-and-cybercrime machines.

| Discipline    | Latency predictability (0-10) | Best for delay-critical pipelines (0-10) | Operational flexibility (0-10) |
|---------------|-------------------------------|------------------------------------------|--------------------------------|
| Periodic (P)  | 9                             | 9                                        | 5                              |
| Sporadic (S)  | 6                             | 6                                        | 8                              |
| Aperiodic (A) | 3                             | 4                                        | 9                              |
| Other (O)     | 5                             | 5                                        | 7                              |

**Table: Activation disciplines (illustrative)**

## REFERENCES

- Amershi, S., Begel, A., Bird, C., et al. (2019). Software engineering for machine learning: A case study. *Proceedings of the International Conference on Software Engineering*, 291–300.
- Pamisetty, A., Paleti, S., Adusupalli, B., Singireddy, J., Inala, R., & Nagabhyru, K. C. (2025, September). Explainable AI Systems for Credit Scoring and Loan Risk Assessment in Digital Banking Platforms. In *2025 IEEE 13th International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS)* (pp. 1478-1483). IEEE.
- Armbrust, M., Zaharia, M., Xin, R. S., et al. (2015). *Apache Spark: A unified engine for big data processing*. *Communications of the ACM*, 59(11), 56–65.
- Garapati, R. S. (2025). An Intelligent IoT Security System: Cloud-Native Architecture with Real-Time AI Threat Detection and Web Visualization. *Journal homepage: <https://jmsronline.com>*, 2(06).
- Batini, C., & Scannapieco, M. (2016). *Data and information quality: Dimensions, principles and techniques*. Springer.
- Babaiah, C., Dobriyal, N., Shamila, M., Aitha, A. R., Patel, S. P., & Upodhyay, D. (2025, December). Intelligent Fault Detection and Recovery in Wireless Sensor Networks Using AI. In *2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG)* (pp. 1-6). IEEE.

7. Benjamens, S., Dhunoo, P., & Meskó, B. (2020). The state of artificial intelligence-based FDA-approved medical devices. *NPJ Digital Medicine*, 3, 118.
8. Nagabhyru, K. C. (2025). Beyond Automation: The 2025 Role of Agentic AI in Autonomous Data Engineering and Adaptive Enterprise Systems.
9. Bertsekas, D. P. (2012). *Dynamic programming and optimal control* (Vol. 1). Athena Scientific.
10. Vajpayee, A., Khan, S., Gottimukkala, V. R. R., Sharma, D., & Seshasai, S. J. (2025). Digital Financial Literacy 4.0: Consumer Readiness for AI-Driven Fintech and Blockchain Ecosystems. *International Insurance Law Review*, 33(S5), 963-973.
11. Brundage, M., Avin, S., Clark, J., et al. (2018). The malicious use of artificial intelligence. *arXiv*.
12. Nigam, N., Sireesha, B., Ediga, P., Segireddy, A. R., & Bokde, S. (2025, December). Comparative Evaluation of Cloud Security Algorithms Using Multiple Classifiers with an Optimized Intrusion Detection System. In *2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG)* (pp. 1-6). IEEE.
13. Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. *Mobile Networks and Applications*, 19, 171-209.
14. Pareyani, S., Goswami, S., Geetha, Y., Dimri, S. K., Niharika, D. S., & Amistapuram, K. (2025, December). Smart Resource Allocation in Wireless Sensor Networks Through AI Techniques. In *2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG)* (pp. 1-6). IEEE.
15. Dehghani, Z. (2022). *Data mesh*. O'Reilly Media.
16. Varri, D. B. S. V. (2025). Human-AI collaboration in healthcare security.
17. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4), 211-407.
18. Nagubandi, A. R. (2025). Cryptocurrency Market Spillovers: Risk Contagion Across Global Financial Systems.
19. European Parliament and Council of the European Union. (2016). *General Data Protection Regulation (GDPR)*. Official Journal of the European Union.
20. Yandamuri, U. S. AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology.
21. Gentry, C. (2009). A fully homomorphic encryption scheme. Stanford University.
22. Guntupalli, R. (2025). Federated Deep Learning for Predictive Healthcare: A Privacy-Preserving AI Framework on Cloud-Native Infrastructure. *Vascular and Endovascular Review*, 8(16s), 200-210.
23. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
24. Dutta, P., Mondal, A., Vadisetty, R., Polamarasetti, A., Guntupalli, R., & Rongali, S. K. (2025). A novel deep learning rule-based spike neural network (SNN) classification approach for diagnosis of intracranial tumors. *International Journal of Information Technology*, 17(9), 5705-5712.
25. He, J., Baxter, S., Xu, J., et al. (2019). The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*, 25, 30-36.
26. Enterprise-Scale Gen AI Orchestration Using Small LMs and LLM Agents for Intelligent ITSM and HRSD Automation in Enterprise Ecosystems. (2025). *MSW Management Journal*, 35(2), 1889-1897.
27. Holzinger, A. (2016). *Interactive machine learning for health informatics*. Springer.
28. FinOps Strategies for AI-Enabled Real-Time Compliance Platforms in Cloud Native Environments. (2025). *MSW Management Journal*, 35(2), 2080-2088.
29. IBM. (2023). *Data fabric architecture overview*. IBM Redbooks.
30. Velangani Divya Vardhan Kumar Bandi. (2024). Intelligent Data Platforms For Personalized Retail Analytics At Scale. *Metallurgical and Materials Engineering*, 30(4), 1011-1027. Retrieved from <https://metall-mater-eng.com/index.php/home/article/view/1011-1027>
31. Jennings, N. R., & Wooldridge, M. (1998). *Applications of intelligent agents*. Springer.
32. Sasi Kumar Kolla. (2023). Big Data-Driven Machine Learning Frameworks for Clinical Risk Prediction. *International Journal of Medical Toxicology and Legal Medicine*, 26(3 and 4), 44-59. Retrieved from <https://ijmtlm.org/index.php/journal/article/view/1456>
33. Kelly, C. J., Karthikesalingam, A., Suleyman, M., et al. (2019). Key challenges for delivering clinical impact with AI. *BMC Medicine*, 17, 195.
34. Kumar, K. M., Parasar, A., Walia, A., Inala, R., & Thulasimani, T. (2025, August). Enhancing Risk Management Strategies in Financial Institutions Using CNN and Support Vector Regression. In *2025 5th Asian Conference on Innovation in Technology (ASIANCON)* (pp. 1-6). IEEE.
35. Koller, D., & Friedman, N. (2009). *Probabilistic graphical models*. MIT Press.
36. Rao, A. N., Garapati, R. S., Suganya, R. T., Kaliappan, A., & Kamaleshwar, T. (2025, August). Smart Solar Harvesting and Power Management in IoT Nodes Through Deep Learning Models. In *2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS)* (pp. 1-6). IEEE.
37. Liu, F., et al. (2025). Foundational architecture for AI agents in healthcare. *Cell Reports Medicine*, 6(10), 102374.

38. Paleti, S., Baliyan, M., Aitha, A. R., Reddy, B. A., Bhadauria, G. S., & Sing, S. A. (2025, August). Graph—LSTM Hybrid Model for Improving Fraud Detection Accuracy in E-Commerce Financial Services. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.
39. Moreau, L., & Groth, P. (2013). Provenance: An introduction to PROV. Morgan & Claypool.
40. Nagabhyru, K. C., Rani, M., Reddy, D. S., & Krishnaraj, V. (2025, August). Machine Learning-Driven Fault Detection in Electric Vehicles via Hybrid Reinforcement Learning Model. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.
41. Obermeyer, Z., & Emanuel, E. (2016). Predicting the future—Big data and clinical medicine. *NEJM*, 375, 1216–1219.
42. Vijaya Rama Raju Gottimukkala. (2025). Agentic AI for Next-Generation Cross-Border Payments: Contextual Learning in Transaction Routing. *Journal of Informatics Education and Research*, 5(4). Retrieved from <https://jier.org/index.php/journal/article/view/3794>
43. Pearl, J. (2009). *Causality* (2nd ed.). Cambridge University Press.
44. Srikanth, T., Segireddy, A. R., & Elavarasi, S. A. (2025, October). STaSFormer-SGAD: Semantic Triplet-Aware Spatial Flow-Guided Spatio-Temporal Graph for Anomaly Detection in Surveillance Videos. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1-7). IEEE.
45. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *NEJM*, 380, 1347–1358.
46. Amistapuram, K. (2025). Agentic AI for Next-Generation Insurance Platforms: Autonomous Decision-Making in Claims and Policy Servicing. *Journal of Marketing & Social Research*, 2, 88-103.
47. Rieke, N., Hancox, J., Li, W., et al. (2020). Federated learning for digital health. *NPJ Digital Medicine*, 3, 119.
48. Varri, D. B. S. (2024). Adaptive and Autonomous Security Frameworks Using Generative AI for Cloud Ecosystems. Available at SSRN 5774785.
49. Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
50. Lebcir, I., Mageswari, S. U., Bhosale, Y. H., Nagubandi, A. R., & Mahabooba, M. M. Agile Strategic Management in the Age of Disruption: Leveraging AI and Data Analytics for Competitive Advantage.
51. Satyanarayanan, M. (2017). The emergence of edge computing. *Computer*, 50(1), 30–39.
52. 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.
53. Sheller, M. J., Reina, G. A., Edwards, B., et al. (2020). Multi-institutional deep learning without sharing patient data. *Brainlesion Workshop*.
54. GUNTUPALLI, R. (2025). EXPLAINABLE AI IN CLINICAL DECISION SUPPORT: INTERPRETABLE NEURAL MODELS FOR TRUSTWORTHY HEALTHCARE AUTOMATION. EXPLAINABLE AI IN CLINICAL DECISION SUPPORT: INTERPRETABLE NEURAL MODELS FOR TRUSTWORTHY HEALTHCARE AUTOMATION. *TPM—Testing, Psychometrics, Methodology in Applied Psychology*, 32(S9 (2025): Posted 15 December), 462-471.
55. Shortliffe, E. H., & Sepúlveda, M. J. (2018). Clinical decision support in the era of AI. *JAMA*, 320(21), 2199–2200.
56. Rongali, S. K. (2025, August). Deep Learning for Cybersecurity in Healthcare: A Mulesoft-Enabled Approach. In 2025 International Conference on Artificial Intelligence and Machine Vision (AIMV) (pp. 1-6). IEEE.
57. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning* (2nd ed.). MIT Press.
58. Siva Hemanth Kolla. (2023). Deep Learning–Driven Retrieval-Augmented Generation for Enterprise ITSM Automation: A Governance-Aligned Large Language Model Architecture . *Journal of Computational Analysis and Applications (JoCAAA)*, 31(4), 2489–2502. Retrieved from <https://www.eudoxuspress.com/index.php/pub/article/view/4774>
59. Tsamados, A., Aggarwal, N., Cowls, J., et al. (2022). The ethics of algorithms. *AI & Society*, 37, 215–230.
60. Davuluri, P. S. L. N. . (2024). AI-Driven Data Governance Frameworks for Automated Regulatory Reporting and Audit Readiness. *Metallurgical and Materials Engineering*, 30(4), 996–1010. Retrieved from <https://metall-mater-eng.com/index.php/home/article/view/1936>
61. Wooldridge, M. (2009). *An introduction to multiagent systems* (2nd ed.). Wiley.
62. Bandi, V. D. V. K. (2023). Production-Grade Machine Learning Pipelines For Healthcare Predictive Analytics. *South Eastern European Journal of Public Health*, 189–205. Retrieved from <https://www.seejph.com/index.php/seejph/article/view/7057>
63. Zhang, A., Xing, L., Zou, J., & Wu, J. C. (2022). Shifting ML for healthcare to deployment. *Nature Biomedical Engineering*, 6, 1330–1345.
64. Kolla, S. K. (2021). Architectural Frameworks for Large-Scale Electronic Health Record Data Platforms. *Current Research in Public Health*, 1(1), 1–19. Retrieved from <https://www.scipublications.com/journal/index.php/crph/article/view/1372>

65. Benford, S., et al. (2009). Emergent multi-agent architectures. *Autonomous Agents and Multi-Agent Systems*, 18, 15–45.
66. Inala, R. (2025). A Unified Framework for Agentic AI and Data Products: Enhancing Cloud, Big Data, and Machine Learning in Supply Chain, Insurance, Retail, and Manufacturing. *EKSPLORIUM-BULETIN PUSAT TEKNOLOGI BAHAN GALIAN NUKLIR*, 46(1), 1614-1628.
67. Ferber, J. (1999). Multi-agent systems: An introduction. Addison-Wesley.
68. Garapati, R. S., & Daram, D. S. B. (2025). AI-Enabled Predictive Maintenance Framework For Connected Vehicles Using Cloud-Based Web Interfaces. Available at SSRN 5524261.
69. Kephart, J. O., & Chess, D. M. (2003). The vision of autonomic computing. *Computer*, 36(1), 41–50.
70. Aitha, A. R., & Jyothi Babu, D. A. (2025). Agentic AI-Powered Claims Intelligence: A Deep Learning Framework for Automating Workers Compensation Claim Processing Using Generative AI. Available at SSRN 5505223.
71. Huhns, M. N., & Singh, M. P. (1998). Readings in agents. Morgan Kaufmann.
72. Nagabhyru, K. C., & Babu, A. J. Human In The Loop Generative AI: Redefining Collaborative Data Engineering For High Stakes Industries.
73. Erl, T. (2016). *Microservices design patterns*. Prentice Hall.
74. Gottimukkala, V. R. R. (2025). Generative AI for Exceptions and Investigations: Streamlining Resolution Across Global Payment Systems. *Journal of International Commercial Law and Technology*, 6(1), 969-972.
75. Fowler, M. (2018). *Refactoring* (2nd ed.). Addison-Wesley.
76. Segireddy, A. R. (2025). GENERATIVE AI FOR SECURE RELEASE ENGINEERING IN GLOBAL PAYMENT NETWORK. *Lex Localis: Journal of Local Self-Government*, 23.
77. Gamma, E., Helm, R., Johnson, R., & Vlissides, J. (1994). *Design patterns*. Addison-Wesley.
78. Amistapuram, K. (2025). GENERATIVE AI FOR CLAIMS EXCEPTIONS AND INVESTIGATIONS: ENHANCING RESOLUTION EFFICIENCY IN COMPLEX INSURANCE PROCESSES. Available at SSRN 5785482.
79. Zaharia, M., et al. (2010). Spark: Cluster computing with working sets. *HotCloud*.
80. Rongali, S. K., & Varri, D. B. S. (2025). AI in health care threat detection. *World Journal of Advanced Research and Reviews*, 25(3), 1784-1789.
81. Lakshman, A., & Malik, P. (2010). Cassandra. *ACM SIGOPS Operating Systems Review*, 44(2), 35–40.
82. Nagubandi, A. R. (2025). PIONEERING SELF-ADAPTIVE AI ORCHESTRATION ENGINES FOR REAL-TIME END-TO-END MULTI-COUNTERPARTY DERIVATIVES, COLLATERAL, AND ACCOUNTING AUTOMATION: INTELLIGENCE-DRIVEN WORKFLOW COORDINATION AT ENTERPRISE SCALE. *Lex Localis*, 23(S6), 8598-8610.
83. Stonebraker, M., & Çetintemel, U. (2005). One size fits all? *ICDE Proceedings*, 2–11.
84. Yandamuri, U. S. (2022). Big Data Pipelines for Cross-Domain Decision Support: A Cloud-Centric Approach. *International Journal of Scientific Research and Modern Technology*, 227.
85. Moreira, M. W. L., et al. (2018). IoT-based smart healthcare systems. *Sensors*, 18(4), 1155.
86. Guntupalli, R. (2025). Multi-Cloud vs. Hybrid Cloud Security: Key Challenges and Best Practices. *Hybrid Cloud Security: Key Challenges and Best Practices* (November 21, 2025).
87. Mell, P., & Grance, T. (2011). The NIST definition of cloud computing. NIST.
88. Rongali, S. K. (2025, August). AI-Powered Threat Detection in Healthcare Data. In *2025 International Conference on Artificial Intelligence and Machine Vision (AIMV)* (pp. 1-7). IEEE.
89. World Health Organization. (2021). *Ethics and governance of artificial intelligence for health*. WHO Press.
90. Kolla, S. H. (2024). RETRIEVAL-AUGMENTED GENERATION WITH SMALL LLMS FOR KNOWLEDGE-DRIVEN DECISION AUTOMATION IN ENTERPRISE SERVICE PLATFORMS. *Turkish Journal of Computer and Mathematics Education (TURCOMAT)*, 15(3), 476–486.  
<https://doi.org/10.61841/turcomat.v15i3.15497>
91. Moreau, L., et al. (2015). The W3C PROV family of specifications. *Future Generation Computer Systems*, 29(7), 161–165.
92. Davuluri, P. N. Integrating Artificial Intelligence into Event-Driven Financial Crime Compliance Platforms.
93. Van Roy, P. (2009). Self-management in distributed systems. *IEEE Computer*, 42(12), 40–47.
94. Vardhan Kumar Bandi, V. D. (2024). Automated Feature Engineering Systems in Large-Scale Healthcare Data Environments. *Journal of Neonatal Surgery*, 13(1), 2127–2141. Retrieved from <https://www.jneonatalurg.com/index.php/jns/article/view/10004>
95. Sutton, R. S. (2019). The bitter lesson. *Incomplete Ideas Blog*.
96. Floridi, L. (2013). *The ethics of information*. Oxford University Press.