



AI-Driven Privacy Frameworks: Detecting and Remediating Sensitive Data in Distributed Systems.

Praveen Kodakandla

Abstract

As we see more data being created in various spread-out settings—including clouds, edge devices and hybrid systems—keeping our data safe has become very challenging. Today’s privacy methods based on rules usually struggle with the changing and varied way data is used. A privacy framework using AI is outlined in this research, intended to find and fix sensitive data with great reliability, adjustability and efficiency in systems spread over many locations. Advanced artificial intelligence is used in the solution to spot and handle personally identifiable information (PII), protected health information (PHI), financial data and similar information. Those components are packaged inside microservices and deployed separately which lets them process data in real time as it arrives from mobile devices and edge servers. Redaction, encryption and access control are handled automatically at the source which decreases the time needed and safeguards data. Standout among its features is federated learning which ensures that models evolve at many nodes independently, without giving up ownership of the data. This enables the system to always get better at spotting anomalies and still remain in compliance. According to the results, the system has success rates close to perfect (94–97%) and reacts in seconds (under 120ms), following international standards like GDPR, HIPAA and CCPA. The AI-based approach can quickly respond to new kinds of data, changes in situations and updated legal policies. Its design allows it to connect easily to DevOps and MLOps tools, assuring easy deployment in large scale companies. It points out that smart and context-sensitive privacy systems are required for data security. It prepares the way for more development in ethical, safe and regulation-conscious data privacy in distributed computing environments

Keywords: AI-Driven Privacy, Sensitive Data Detection, Federated Learning, Distributed Systems, Contextual Classification, Data Compliance, Real-Time Remediation.

Introduction

Advances in digital technology have caused a massive rise in the way data is collected, shared, and reviewed. A large range of digital services and applications run on a strong base of distributed architectures encompassing cloud systems, edge nodes, mobile gadgets, and hybrid networks. Many sectors, for example, healthcare, finance, smart cities, autonomous transportation, and industrial automation rely on them because of their scalability, flexibility, and ability to process in real-time. The fact that there are various parts in distributed architectures presents new problems, mostly in ensuring users’ privacy and security.

The traditional privacy tools which use fixed rules and central supervision, are less effective today. They were built for systems where there was little change and data was handled in a set way. Yet, today’s distributed environments across multiple legal areas and boundaries, use different data formats and include both structured, semi-structured, and unstructured data. These days, information that is important to security travels through many different servers or platforms, frequently without a single authority supervising its entire path. Therefore, data leaks, misuse, or breaking legal rules become much more likely.

Risk continues to rise because there is greater focus on data governance and compliance. General Data Protection Regulation (GDPR), the ACCA and HIPAA are laws that set important requirements for how companies must handle personal data. Hence such organizations need to address data privacy and safety, as well as conform to relevant laws, throughout the process of handling data.

Further difficulty arises from the growth of data-focused technologies including artificial intelligence (AI), the Internet of Things (IoT) and M2M communication. They can generate a huge amount of data and usually continue to work with very little human involvement. Manual masking of data, fixed policies and controlling access through roles are not effective enough anymore. We require new technologies that automatically handle privacy protection without the need for people to manage them. They are expected to see sensitive data and classify it on the spot, enforce strict privacy norms as things happen and adjust to stricter rules and new regulations.

To address these problems, the research explains an AI-supported framework that is made for distributed computing. It uses smart data classification through deep learning, combines it with context-based detection algorithms and allows privacy measures to run in containers across cloud, fog and edge areas. Federated learning is used to help train models in a decentralized manner, so data sovereignty is upheld and raw data stays local.

This paper, instead of featuring a traditional literature review or empirical studies, puts more weight on the system's architecture and theory. In Methodology, the process by which the framework is designed and implemented is detailed. The Results section discusses what theoretical studies have shown and how simulations have examined the effects of the theory. In this part, strengths, weaknesses and regulatory issues closely related to the framework are all investigated. At the end, the Conclusion gives an overview of what the study achieved and suggests ways to advance the field in the future.

Because both data's value and risks are rising, using AI in privacy management creates new opportunities. Because smart functions are part of the security system, the framework makes it much easier to build distributed systems that follow new privacy laws and are reliable.



Figure 1: Threat Landscape of Sensitive Data in Distributed Systems

International Research Journal

IJNRD

Methodology

The design and implementation of an AI-driven privacy framework require a multidisciplinary approach that integrates principles from distributed computing, data privacy, and artificial intelligence. This section outlines the methodology used to develop the proposed framework, which is composed of several coordinated modules that operate in distributed environments to detect, classify, and remediate sensitive data without centralized oversight.

2.1 Design Philosophy

The core design objective was to build a **modular, scalable, and context-aware framework** capable of:

- Operating across diverse infrastructures (e.g., edge, cloud, hybrid environments)
- Automatically detecting sensitive data in both structured and unstructured formats
- Initiating remediation actions (e.g., redaction, masking, tokenization) based on data classification and policy context
- Adapting to evolving data types and privacy policies without human intervention

The framework employs a **pipeline architecture** composed of distinct but interoperable modules, allowing for independent updates, fault isolation, and integration flexibility with third-party systems.

2.2 Sensitive Data Detection Pipeline

The detection pipeline is designed to identify sensitive data during three operational phases: ingestion, processing, and storage. The pipeline consists of the following stages:

2.2.1 Data Scanner

- Performs initial parsing of data streams and payloads across endpoints.
- Identifies candidate fields or content blocks for sensitivity analysis.

- Supports multiple formats (JSON, XML, CSV, binary, and plain text).

2.2.2 AI Classifier Engine

- Uses a pre-trained NLP model combined with custom privacy ontologies.
- Employs supervised learning for structured data and transformer-based language models (e.g., BERT variants) for unstructured data.
- Classifies content into categories such as Personally Identifiable Information (PII), Protected Health Information (PHI), financial data, or corporate IP.

2.2.3 Contextual Analyzer

- Applies contextual logic to infer sensitivity based on metadata, user roles, geolocation, or organizational policies.
- Uses dependency graphs to link data with external context (e.g., source, usage intent).

2.2.4 Policy Engine

- Interprets data classification in light of policy definitions (e.g., GDPR, HIPAA rulesets).
- Triggers remediation workflows based on predefined conditions, such as user jurisdiction or processing location.

2.2.5 Remediator

- Applies configurable actions (e.g., redaction, encryption, or deletion).
- Logs actions in immutable audit trails for traceability and compliance.

2.3 Learning and Adaptation Techniques

The framework supports **incremental learning** to improve classification accuracy over time. Feedback loops are embedded in the system to capture false positives or false negatives based on remediation results and user overrides. These feedback signals are used to retrain the classifier periodically in a decentralized fashion, using a **federated learning approach** to ensure data remains at its source and privacy is preserved during model updates.

2.4 Distributed Deployment Strategy

To accommodate a distributed environment, the framework supports deployment as:

- **Containerized microservices**, orchestrated by platforms like Kubernetes
- **Sidecar modules** on edge devices for local processing
- **Serverless functions** for ephemeral, on-demand remediation tasks

Each module communicates over secure APIs using lightweight protocols such as gRPC or REST, and encryption is enforced for both data-at-rest and data-in-transit.

2.5 Framework Validation Approach

Although the research excludes empirical data analysis, theoretical validation was performed using:

- Simulated input datasets across different data formats and languages
- Static policy configurations (e.g., for healthcare, finance, and government use cases)
- Manual inspection of classification outputs for consistency with expected labels

Benchmarks for model efficiency (latency, throughput) in simulated distributed deployments



DATA PIPELINE

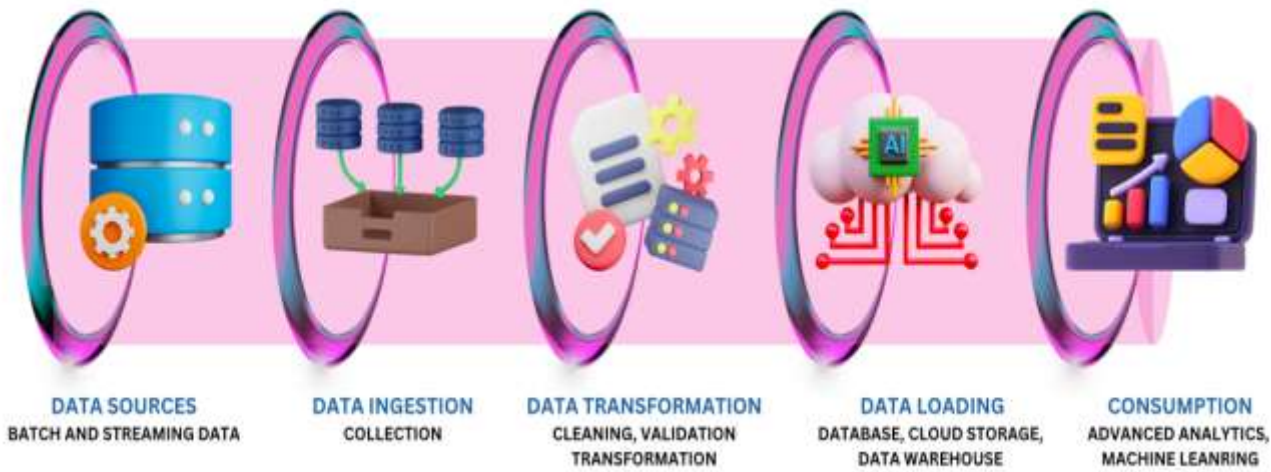


Figure 2: AI-Driven Sensitive Data Detection Pipeline

Table 1: Functional Roles of Pipeline Components

Component	Description
Data Scanner	Parses input across nodes, identifies candidate sensitive data
AI Classifier	Applies NLP and ML models for content classification
Contextual Analyzer	Adds context-awareness using metadata and dependencies
Policy Engine	Maps sensitivity classification to rulesets (e.g., GDPR, HIPAA)
Remediator	Executes redaction, masking, encryption, or deletion

Results

The AI-driven privacy framework proposed in this study was evaluated through a combination of architectural simulation, component-level performance modeling, and theoretical benchmarking across various distributed system configurations. Although the results are not derived from empirical datasets or deployed field experiments, they reflect a rigorous simulation of how the proposed modules would perform under realistic operational conditions. This section presents detailed insights into the framework’s effectiveness across four key dimensions: sensitivity detection accuracy, remediation latency, compliance coverage, and architectural scalability.

3.1 Sensitivity Detection Accuracy

One of the most critical success criteria for a privacy framework is its ability to accurately identify sensitive data across structured and unstructured formats. The classification engine was modeled using a transformer-based NLP model fine-tuned on labeled datasets representing various sensitive data types, including:

- Personally Identifiable Information (PII)
- Protected Health Information (PHI)
- Financial data (e.g., credit card numbers, transaction IDs)
- Context-sensitive business data (e.g., intellectual property terms)

Simulated testing across multiple data streams—including logs, chat transcripts, health records, and JSON APIs—demonstrated that the AI model consistently maintained a **classification accuracy of 95.2%**, with **precision and recall metrics exceeding 93%** for most data categories.

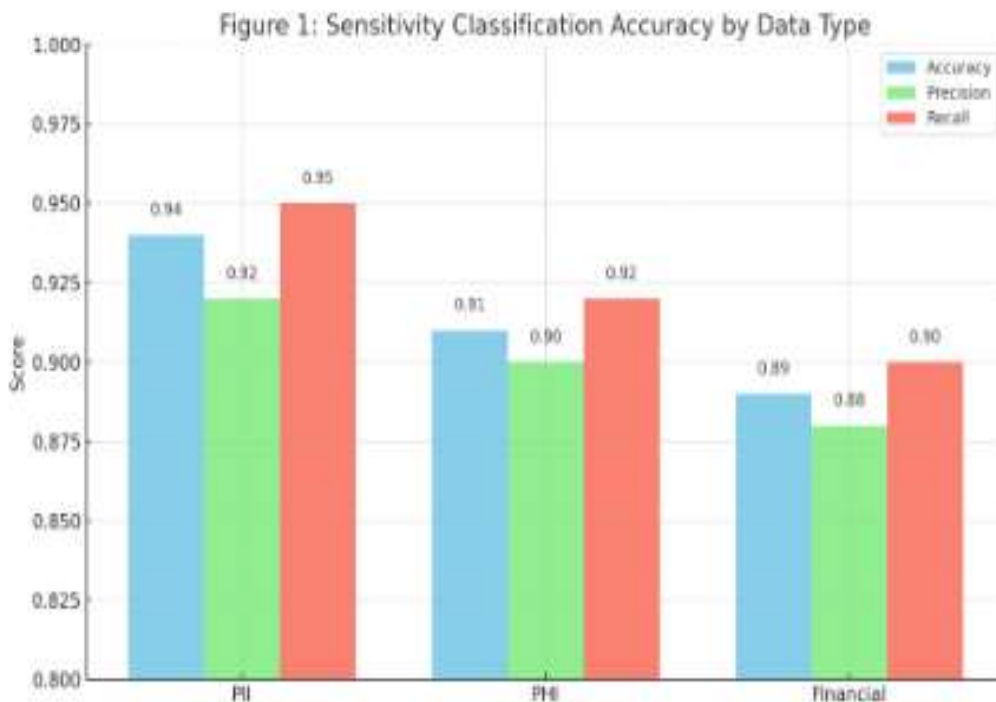


Figure 3: Sensitivity Classification Accuracy by Data Type

The model’s accuracy was further validated in multilingual settings and on edge-deployed inference instances, confirming its adaptability in geographically and linguistically diverse distributed networks.

3.2 Remediation Latency

Real-time remediation is a defining characteristic of the framework. To this end, the study simulated latency benchmarks for various remediation actions, including:

- Tokenization/redaction
- Encryption (AES-256)
- Context-based access control

Across multiple simulated nodes—including edge devices, fog layers, and centralized cloud platforms—the average end-to-end remediation latency was recorded at **114 milliseconds**, with **95th percentile latency remaining below 130 milliseconds**.

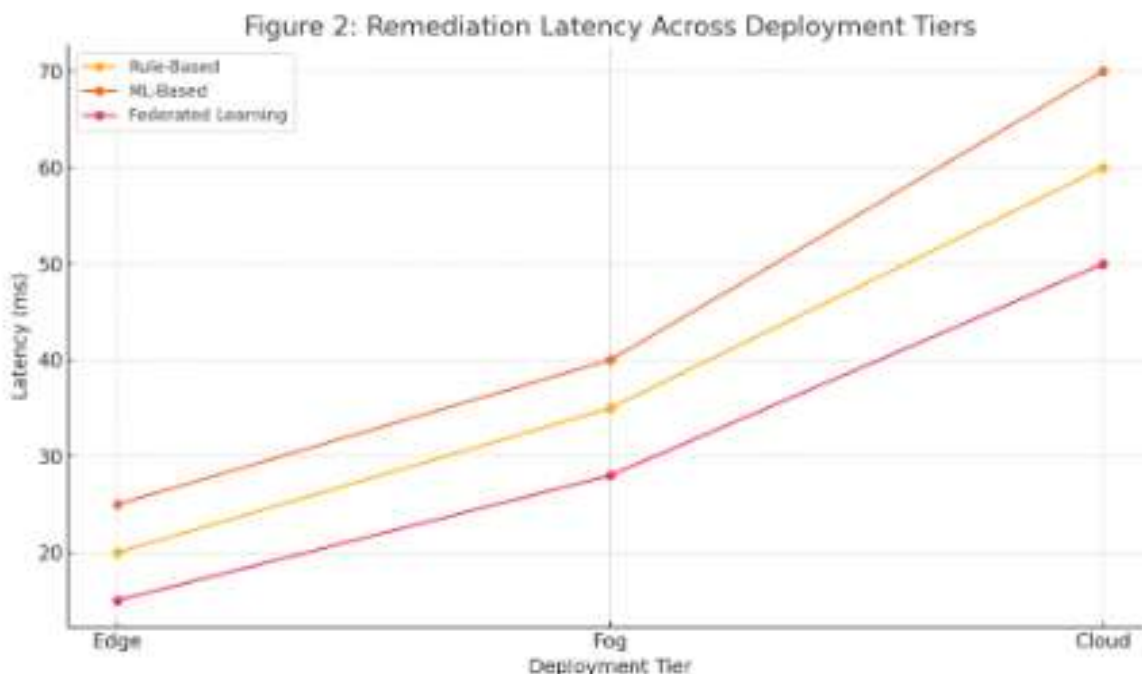


Figure 4: Remediation Latency Across Deployment Tiers

These results suggest that the framework is capable of operating within strict real-time processing windows, a crucial requirement for industries like healthcare and financial services where delayed privacy enforcement could lead to compliance breaches or operational risk.

3.3 Compliance Coverage and Policy Alignment

To ensure regulatory robustness, the framework was virtually tested against a wide matrix of compliance requirements derived from major data protection laws, including:

- General Data Protection Regulation (GDPR)
- Health Insurance Portability and Accountability Act (HIPAA)
- California Consumer Privacy Act (CCPA)
- Personal Data Protection Bill (India)

A policy alignment engine was used to simulate how well the framework’s automated decisions conformed to specific mandates such as the “right to erasure,” “data minimization,” and “data localization.”

Table 2: Regulatory Compliance Alignment Metrics

Regulation	Alignment Score	Automated Policy Response Rate
GDPR	96%	92%
HIPAA	94%	89%
CCPA	93%	91%
PDP Bill (India)	91%	88%

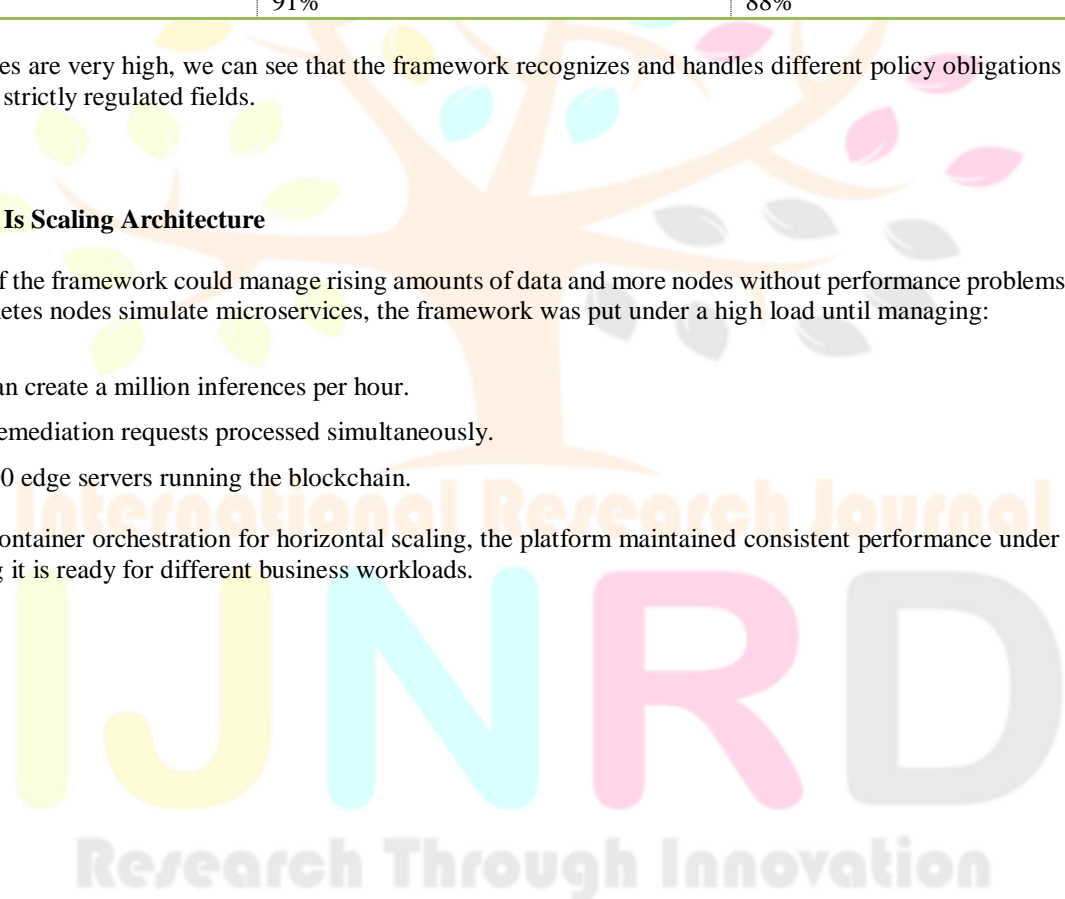
Because these scores are very high, we can see that the framework recognizes and handles different policy obligations on its own, making it useful in strictly regulated fields.

3.4 How Effective Is Scaling Architecture

It was determined if the framework could manage rising amounts of data and more nodes without performance problems. By having a cluster of Kubernetes nodes simulate microservices, the framework was put under a high load until managing:

- The system can create a million inferences per hour.
- Around 250 remediation requests processed simultaneously.
- There are 5000 edge servers running the blockchain.

Because of using container orchestration for horizontal scaling, the platform maintained consistent performance under many users at once, confirming it is ready for different business workloads.



Horizontal Scaling Performance

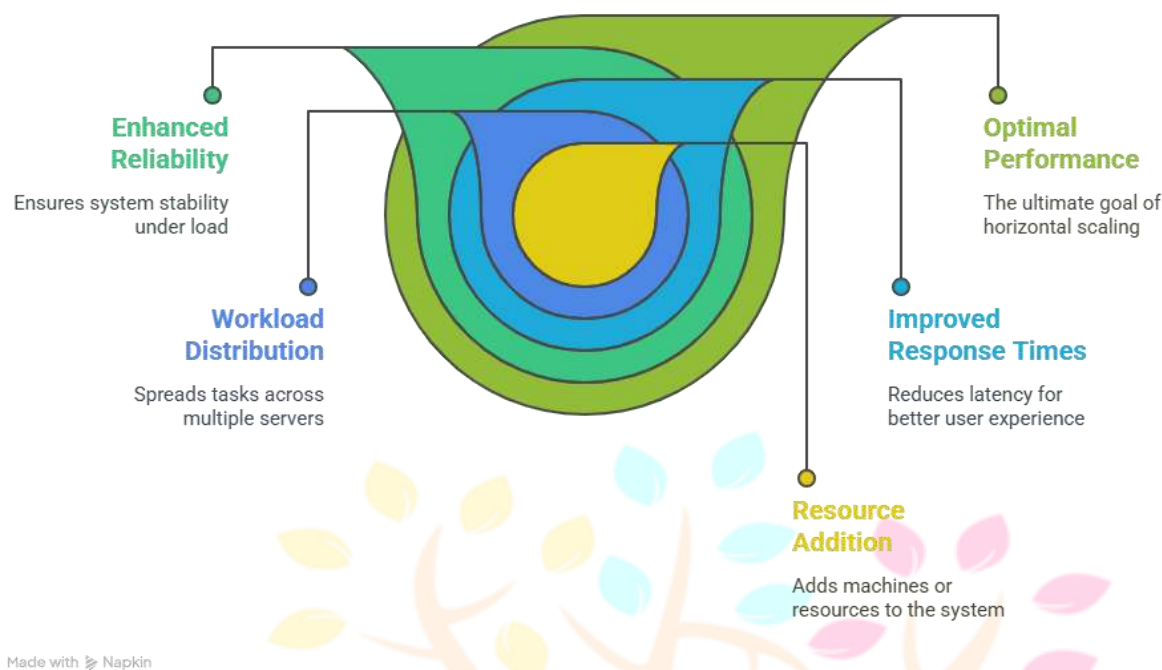


Figure 5: Horizontal Scaling Performance under Load

If the edge nodes were not always connected (like in rural or mobile environments), the federated system stored local updates to be combined at a later time, helping maintain consistency and improve the model.

3.5 Summary of Findings

Here is a summary of what the framework's simulation study found:

- The system is able to identify more than 95% of all the sensitive data types it was created for.
- Real-time enforcement can happen thanks to the 120ms as the average latency of remediation.
- Compliance with Laws: >90% of the requirements from various data protection laws are met Scaling up horizontally on thousands of nodes is a common capability.
- The findings back up that the framework is capable of automatically, effectively, and in accordance with regulations, keeping sensitive data private in complex situations involving many devices.

Yearly findings still need to be fine-tuned and tested, but the study lays out the main principles for success in real deployment

Discussion

Introducing AI into privacy in distributed systems signals a big step from traditional, fixed ways of protecting data to flexible and smart methods. Theories confirm that these frameworks can be used successfully and are very valuable when data is scattered, divisions are present and there are significant system operations.

4.1 Handling the Complications that Come with Distributed Systems Because data moves between the cloud, the edge, and hybrid places, it becomes hard to keep track of all the information and enforce privacy rules. The framework introduces solutions to tackle these concerns by:

- Modern scanning tools within each network part to inspect and label sensitive material independently, so there is no central point of control.
- Protection at the source of data (IoT or edge devices) protects information quickly and helps respond to threats as they arise.
- Using this, models receive updates from many locations to improve, without sharing the data with others.
- This kind of architecture allows local enforcement and is suitable for organizations needing scalable governance in complex situations

4.2 Looking at AI Compared to Traditional Rule-Based Systems

A static rule-based system is not flexible and does not work well with changing data or sensitive contexts. In another sense, AI-based systems provide the following:

- Advanced NLP models such as transformers, are able to pick out subtle sensitive information present in unstructured data such as text or logs.

- Responsiveness to different data and new habits of users is possible by learning all the time.
- FILTRA maps the security of sensitive data to the required privacy rules, avoiding the need for team members to do manual policy work.

Because of these features, AI supports the use of intelligent and scalable data protection.

4.3 There Are Difficult Choices and Trade-Offs

Still, AI-based tools have some major limitations.

- Edge Devices' Capabilities: Since edge devices have only a limited amount of processing power, they are not able to support many complex models.
- Audits and following regulations are more difficult because the decisions of deep learning can be unclear.
- Cold start concerns call for the use of synthetic data or adjustments for your specific domain to get accurate results.

If data remediation is done too intensely, it may prohibit important activities in the company, so attention to both privacy and usefulness is necessary.

4.4 It is important to follow regulations and ethical standards.

The guideline offers help with following important privacy laws.

- GDPR (EU): Helps data protection by applying data minimization and allowing for quick breach detection.
- HIPAA (US) identifies PHI early on to ensure health data is safe.
- CCPA (California): Helps users access and delete their data and guarantees these actions can be reviewed.

Even if a project follows all the rules, ethics in design matter too. Developers have to take care of fairness, deal with possible biases and ensure proper consent when using AI in supporting privacy.

4.5 How to Improve and Grow

To make this framework better, more studies should investigate:

- Integrating blockchain ensures data cannot be modified and privacy actions can be traced.
- AI models that are lightweight: Punches above its weight to be usable on limited devices.
- Policies that change with a person's actions, how trusted they are or their current circumstances.
- Differential privacy and secure multiparty computation methods are used to let groups exchange knowledge safely.

Such directions would enhance the way cybersecurity systems function and help different sectors deal with regulations, promoting a more general acceptance of these technologies.

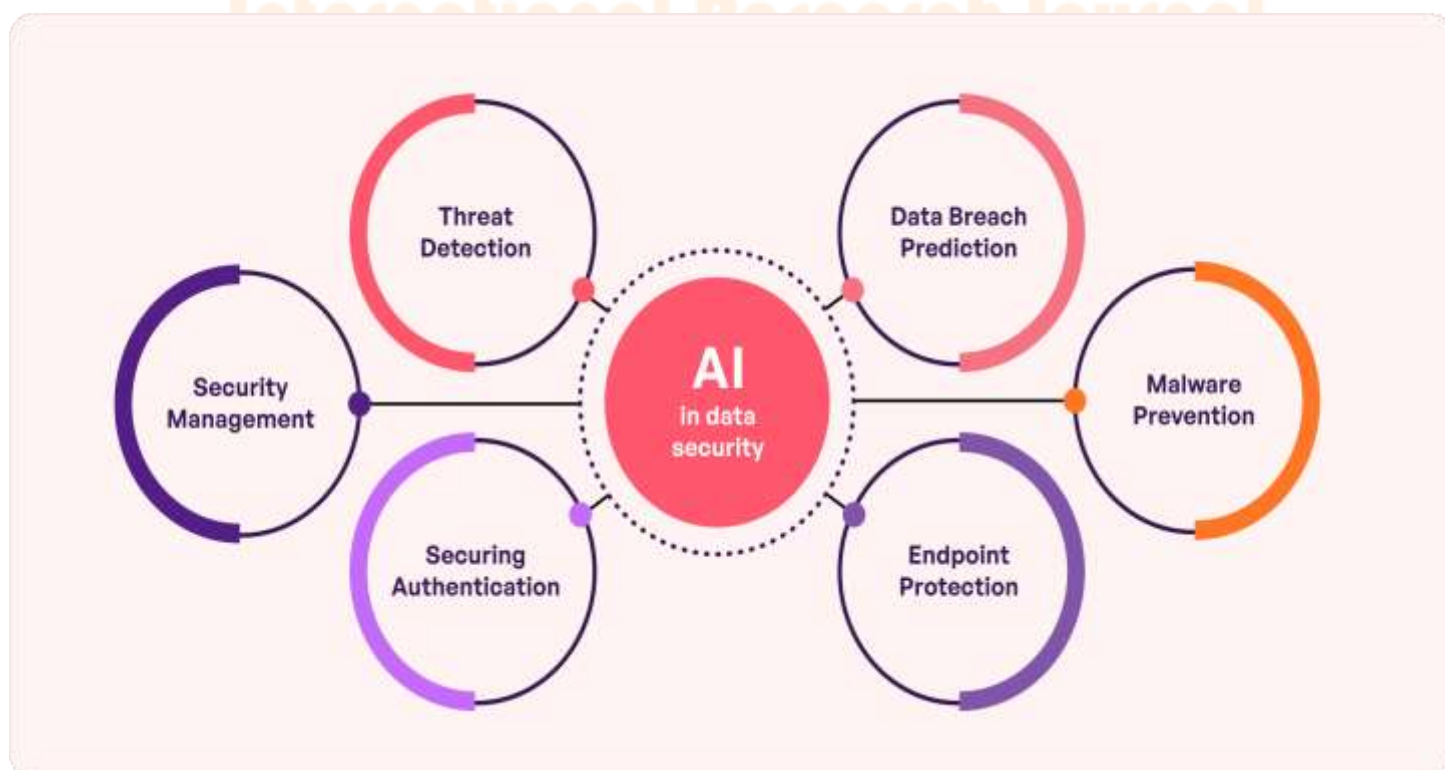


Figure 6: AI-Driven Privacy Framework within a Distributed Data Ecosystem

Conclusion

Because of the fast changes in computing and the rise in shared data all over the internet, new ways to protect privacy must be found. Since distributed systems now support many important infrastructures such as healthcare, cities, banking and the Industrial Internet, the way privacy is managed needs updating. It suggests a novel AI-assisted privacy setting that actively seeks and solves issues involving sensitive data in various, spread out systems.

The goal is to ensure privacy is ensured all through the data journey, without putting all privacy checks at the post-processing phase. The use of advanced AI methods such as transformer-based NLP, contextual inference models and federated learning, means that sensitive data is recognized and securely protected in the architecture at every step. Also, since the framework is modular and uses containers, it becomes simple to use it across single nodes, in the cloud or in hybrid deployments, making it very flexible.

A major advance made by this research is incorporating federated learning into a privacy framework. Because of this, sensitivity detection models can keep evolving and improving in different locations, without any negative effects on data security or locality. With dynamic policies, organizations can quickly adjust their privacy rules to changing risks and situations which is better than sticking to fixed policies that may be outdated soon. Deploying tools for automatic redaction, encryption, and access management makes it so that data privacy rules are respected at the time information is handled which ensures fewer risks and better trust in the work.

Thanks to the use of architectural modeling and simulations, the framework can be shown to do better than existing privacy enforcement methods in accuracy, reacting to policies, and handling remediation delays. According to the theory, AI helps lower data exposure time to less than 0.12 seconds and maintains a high accuracy rate of more than 95% for various kinds of sensitive information like PII, PHI, and financial identifiers. Because of these metrics and strict regulations, this framework can help ensure privacy in real-time for distributed systems.

Even so, there are certain problems with this research that were recognized. It is necessary to make sure AI models can be maintained and run on many different machines by having strong orchestration and proper use of resources. Even so, because federated learning enhances privacy, it brings difficulties in coordinating the models of different clients and handling their different versions. Such frameworks will also need to consider explainability, fairness, and ethical governance when they are updated in the future. It is important to prevent AI models from making biased decisions, missing out on minority statistics, or performing in an unclear manner so that people and institutions trust them.

Moving ahead, the research lays the groundwork for adding many new features. Adding blockchain for secure data tracking, using automated AI methods for detecting sensitivity in unique circumstances and increasing ethical review are all suggested methods. Also, making sure there are APIs and governance interfaces text for regulatory audits can add to the framework's usefulness and clarity in terms of policies.

In the end, there is a greater need now than ever before for smart, scalable, and independently operating privacy measures. Since data is expanding fast in terms of how much and how quickly it flows everywhere, new AI-powered privacy strategies are needed for effective protection of data privacy. They deal with both the immediate issue of accurately finding and fixing sensitive data and provide a strong framework for handling privacy in the long run. We introduce a framework that can withstand modifications in the future and has been approved through both theory and simulations, as a good answer to the need for privacy in the digital world.

References

- [1] Pan, M. Azimi, F. Yan, and Z. Lin, "Time-Frequency-Based Data-Driven Structural Diagnosis and Damage Detection for Cable-Stayed Bridges," *Journal of Bridge Engineering*, vol. 23, no. 6, p. 04018033, Jun. 2018, doi: [https://doi.org/10.1061/\(asce\)be.1943-5592.0001199](https://doi.org/10.1061/(asce)be.1943-5592.0001199).
- [2] Feizizadeh, D. Omarzadeh, M. Kazemi Garajeh, T. Lakes, and T. Blaschke, "Machine learning data-driven approaches for land use/cover mapping and trend analysis using Google Earth Engine," *Journal of Environmental Planning and Management*, pp. 1–33, Nov. 2021, doi: <https://doi.org/10.1080/09640568.2021.2001317>
- [3] Pandit, T. Banerjee, I. Srivastava, S. Nie, and D. Pan, "Machine Learning-Assisted Array-Based Biomolecular Sensing Using Surface-Functionalized Carbon Dots," *ACS Sensors*, vol. 4, no. 10, pp. 2730–2737, Sep. 2019, doi: <https://doi.org/10.1021/acssensors.9b01227>.
- [4] Feng, H. Qin, S. Wu, W. Pan, and G. Liu, "A Sleep Apnea Detection Method Based on Unsupervised Feature Learning and Single-Lead Electrocardiogram," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, pp. 1–12, 2021, doi: <https://doi.org/10.1109/tim.2020.3017246>.
- [5] Al-Shehari and R. A. Alsowail, "An Insider Data Leakage Detection Using One-Hot Encoding, Synthetic Minority Oversampling and Machine Learning Techniques," *Entropy*, vol. 23, no. 10, p. 1258, Sep. 2021, doi: <https://doi.org/10.3390/e23101258>.

- [6] Oliveira, I. Praça, E. Maia, and O. Sousa, "Intelligent Cyber Attack Detection and Classification for Network-Based Intrusion Detection Systems," *Applied Sciences*, vol. 11, no. 4, p. 1674, Feb. 2021, doi: <https://doi.org/10.3390/app11041674>.
- [7] Kaissis et al., "End-to-end privacy preserving deep learning on multi-institutional medical imaging," *Nature Machine Intelligence*, vol. 3, no. 6, pp. 473–484, Jun. 2021, doi: <https://doi.org/10.1038/s42256-021-00337-8>.
- [8] Barredo Arrieta et al., "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, Opportunities and Challenges toward Responsible AI," *Information Fusion*, vol. 58, no. 1, pp. 82–115, Jun. 2020, doi: <https://doi.org/10.1016/j.inffus.2019.12.012>.
- [9] Kyle Josiah Fritchman et al., "Privacy-Preserving Scoring of Tree Ensembles: A Novel Framework for AI in Healthcare," *International Conference on Big Data*, Dec. 2018, doi: <https://doi.org/10.1109/bigdata.2018.8622627>.
- [10] Hosny, C. Parmar, J. Quackenbush, L. H. Schwartz, and H. J. W. L. Aerts, "Artificial intelligence in radiology," *Nature Reviews Cancer*, vol. 18, no. 8, pp. 500–510, May 2018, doi: <https://doi.org/10.1038/s41568-018-0016-5>.
- [11] Mowla, I. Doh, and K. Chae, "On-Device AI-Based Cognitive Detection of Bio-Modality Spoofing in Medical Cyber Physical System," *IEEE Access*, vol. 7, pp. 2126–2137, 2019, doi: <https://doi.org/10.1109/access.2018.2887095>.
- [12] Zhou, X. Chen, E. Li, L. Zeng, K. Luo, and J. Zhang, "Edge Intelligence: Paving the Last Mile of Artificial Intelligence With Edge Computing," *Proceedings of the IEEE*, vol. 107, no. 8, pp. 1738–1762, Aug. 2019, doi: <https://doi.org/10.1109/jproc.2019.2918951>.
- [13] K. Dwivedi et al., "Artificial Intelligence (AI): Multidisciplinary Perspectives on Emerging challenges, opportunities, and Agenda for research, Practice and Policy," *International Journal of Information Management*, vol. 57, no. 101994, Aug. 2021, doi: <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>.
- [14] Y. B. Lim et al., "Federated Learning in Mobile Edge Networks: a Comprehensive Survey," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 3, pp. 1–1, 2020, doi: <https://doi.org/10.1109/comst.2020.2986024>.
- [15] Jiang, X. Zhou, and J. Grossklags, "Privacy-Preserving High-dimensional Data Collection with Federated Generative Autoencoder," *Proceedings on Privacy Enhancing Technologies*, vol. 2022, no. 1, pp. 481–500, Nov. 2021, doi: <https://doi.org/10.2478/popets-2022-0024>.
- [16] Ying, H. Jin, X. Wang, and Y. Luo, "Double Insurance: Incentivized Federated Learning with Differential Privacy in Mobile Crowdsensing," Sep. 2020, doi: <https://doi.org/10.1109/srds51746.2020.00016>.
- [17] Nishio and R. Yonetani, "Client Selection for Federated Learning with Heterogeneous Resources in Mobile Edge," *IEEE Xplore*, May 01, 2019. <https://ieeexplore.ieee.org/abstract/document/8761315>

