Cureus Journal of Computer Science
https://doi.org/10.7759/s44389-026-00110-1

Review Article

Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques

B Lavanya [1, ✉], S Janani [1]

1. *Department of Computer Science, University of Madras, Chennai, IND*

Received: April 24, 2026 | Review began: May 11, 2026 | Review ended: June 02, 2026 | Published: June 09, 2026

© **Copyright** 2026
This is an open access article distributed under the terms of the Creative Commons Attribution License CC-BY 4.0., which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract

Data privacy concerns have become more critical than ever as machine learning and applied intelligence systems permeate sensitive industries such as healthcare, finance, national security, and personal services. This necessitates the development of privacy-preserving strategies for protecting private information while retaining the utility of intelligent models. This survey provides a comprehensive overview of privacy-preserving machine learning, with an emphasis on the cryptographic and statistical methods that are transforming how safe learning systems are built. The study starts by examining the most important components of the machine learning model and figuring out which of these may be protected to solve important privacy problems. The article then explores modern cryptographic techniques, including homomorphic encryption, zero-knowledge proofs, secure multiparty computations, and a statistical approach called differential privacy, that support contemporary privacy-preserving machine learning solutions. The study then explores how these strategies are applied independently and in hybrid systems to achieve accuracy, efficiency, and balance of privacy. This survey provides promising direction for protecting sensitive information during real-world model training and inference, offering insights into the design of trustworthy applied intelligence systems.

Categories: Security and Privacy, Cryptography, Machine Learning (ML)

Keywords: privacy-preserving machine learning, cryptographic techniques, homomorphic encryption, secure multiparty computation, differential privacy

Introduction And Background

The potential of machine learning (ML) lies in its ability to identify patterns within extensive datasets. However, this reliance on data also creates vulnerabilities that nefarious individuals may exploit. Data security is even more important in Industry 5.0 and Healthcare 5.0, where increased personalisation, collaboration between humans and AI, and real-time decision-making all depend on private and sensitive information. When used without privacy-protecting techniques, standard models can accidentally reveal private information. This issue becomes increasingly critical in cloud or dispersed systems because data often traverses unreliable networks. Conventional security methods protect data during storage and transmission, but innovative cryptographic techniques guarantee data protection during processing, thus reconciling usability with privacy. Implementing Privacy-Preserving Machine Learning (PPML) involves maintaining the security of sensitive data while the model is being trained, which protects the model's internal parameters and outputs [1]. Figure *1* illustrates a schematic depiction of the procedures employed to protect data during the various phases of the ML pipeline, based on the concepts discussed by Rechberger and Walch [1] and Syed et al. [2].

How to cite this article:
Lavanya B, Janani S (June 09, 2026) Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques. Cureus J Comput Sci 3 : es44389-026-00110-1. DOI https://doi.org/10.7759/s44389-026-00110-1

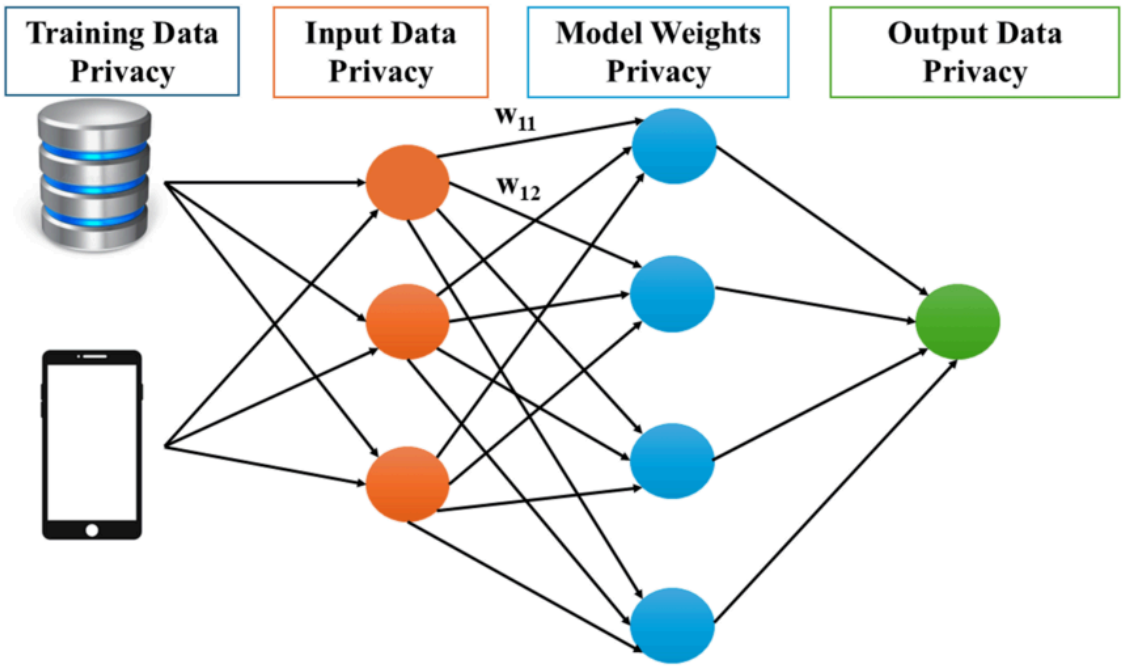

FIGURE 1: Schematic Depiction of Privacy-Sensitive Components in a Machine Learning Workflow

Developed based on the concepts discussed by Rechberger and Walch [1] and Syed et al. [2].

Syed et al. [2] specify that strong protections are needed to keep training data private. If these steps are not taken, there is a real risk that individuals outside the organisation could rebuild the original dataset just by deploying the model. Keeping raw input data private is a key part of input privacy. Model weight privacy is just as important. It means hiding the parameters the model has learned, which protects both proprietary knowledge and the model's internal logic. Finally, output privacy safeguards the model's predictions from being shared with the data owner so they do not accidentally get shared with people outside the company.

Review methodology

In order to provide a comprehensive overview of PPML, this survey examines the evolution of privacy-preserving techniques from foundational principles to state-of-the-art contributions in secure applied intelligence systems and emerging hybrid PPML frameworks. The literature search was conducted using combinations of terms such as “Privacy-Preserving Machine Learning,” “Homomorphic Encryption,” “Secure Multi-Party Computation,” “Zero-Knowledge Proof,” “Differential Privacy,” “Applied Intelligence,” and “Secure AI.” The inclusion criteria considered articles published between 2019 and 2025, along with foundational contributions, collected from major scientific databases and digital libraries such as IEEE Xplore, ACM Digital Library, SpringerLink, ScienceDirect, and arXiv.

The initial search yielded 119 publications. After removing duplicate records, non-English articles, workshop summaries, inaccessible full-text articles, studies lacking sufficient technical details, and publications not directly related to PPML, the remaining studies were subjected to title, abstract, and full-text screening. Following the eligibility assessment, 46 studies were selected for detailed review and analysis. The selected literature particularly included studies reporting implementation details, scalability analysis, performance evaluation, privacy guarantees, or practical deployment considerations related to cryptographic and statistical privacy-preserving approaches in ML.

How to cite this article:
Lavanya B, Janani S (June 09, 2026) Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques. Cureus J Comput Sci 3 : es44389-026-00110-1. DOI https://doi.org/10.7759/s44389-026-00110-1

Comparative analysis was then performed based on privacy guarantees, computational complexity, scalability, model utility, and applicability to real-world ML systems. Putting these points of view together shows that the survey's contributions are based on three main ideas:

a) Crypto-stat Fusion: This study provides an extensive analysis of cryptographic techniques and statistical frameworks, which have generally been regarded in isolation in prior surveys.

b) PPML Benchmarking: It systematically examines different PPML approaches across a range of practical factors.

c) Hybrid Frameworks: It highlights emerging hybrid PPML frameworks, points out open research challenges, and offers guidance for building scalable, verifiable, and ready-to-deploy PPML systems.

The subsequent sections are structured to move from cryptographic paradigms to statistical methods. Figure *2*, developed for this study, presents a compact visual abstraction of the survey's structural progression, facilitating a clearer understanding of how cryptographic and statistical techniques jointly contribute to next-generation privacy in applied intelligence systems.

How to cite this article:
Lavanya B, Janani S (June 09, 2026) Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques. Cureus J Comput Sci 3 : es44389-026-00110-1. DOI https://doi.org/10.7759/s44389-026-00110-1

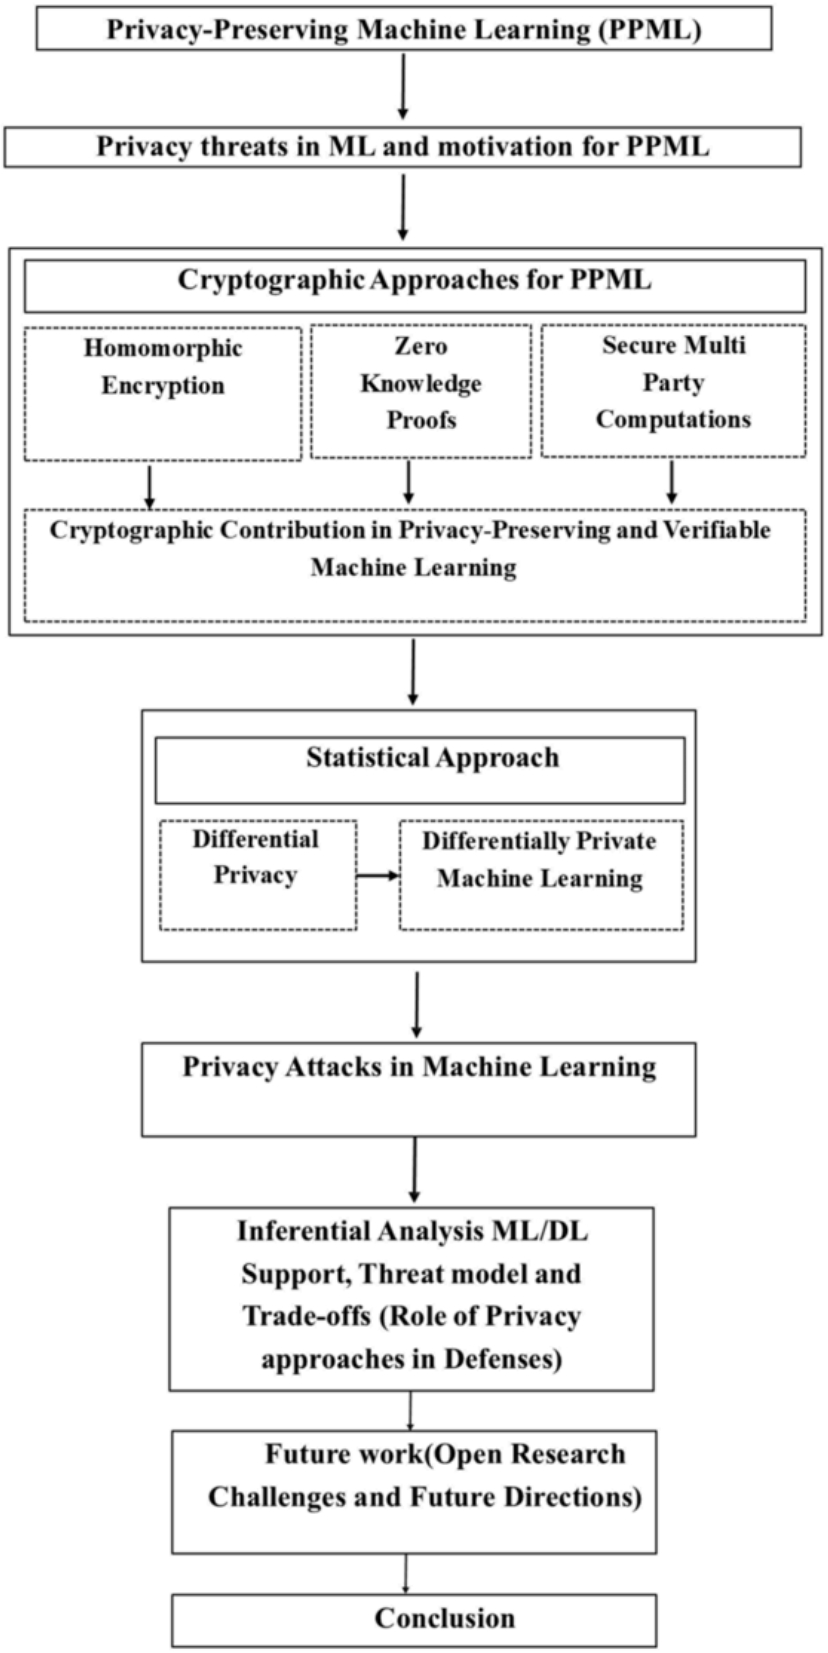

How to cite this article:
Lavanya B, Janani S (June 09, 2026) Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques. Cureus J Comput Sci 3 : es44389-026-00110-1. DOI https://doi.org/10.7759/s44389-026-00110-1

FIGURE 2: Architectural Flow of the Survey Framework

Conceptual representation developed for the present study.

Review

Homomorphic encryption

Homomorphic encryption (HE) allows computations on encrypted data without exposing the actual values [3]. It enables processing directly within the ciphertext domain, keeping sensitive information hidden throughout the workflow. The encrypted result of these computations corresponds exactly to what would be obtained if performed on the plain text. This capability makes HE particularly suitable for cloud-based and PPML applications [4], where sensitive data can be analyzed, trained, or inferred upon without revealing the underlying information to external servers or third parties.

Partial HE

One approach to group HE systems is based on their ability to apply mathematical operations on secret information while maintaining the structural links between the plaintext and ciphertext domains. These systems can be divided into two groups based on their operational capacity. Additive homomorphism makes it possible to add plaintext values using ciphertext processes, as shown by Yan et al. [5]. The Paillier cryptosystem [6] is a semantically secure additive HE technique that is based on the composite residuosity problem, which is well-suited for aggregating data.

$$D\left(E(m_1) \cdot E(m_2) \bmod n^2\right) = (m_1 + m_2) \bmod n$$

E = Encryption Function

D = Decryption Function

m_1, m_2 = Plaintext

n = Public modules, product of two large prime numbers

Two multiplicative homomorphic techniques are RSA [7] and ElGamal [8]. The interaction between the modulus-encryption operation and the RSA encryption system exemplifies the multiplicative homomorphism property.

$$C \equiv E(M) \equiv M^e \pmod{n}$$

$$(m_1^e \cdot m_2^e) \bmod n = (m_1 \cdot m_2)^e \bmod n$$

Fully Homomorphic Encryption on Privacy-Preserving Approaches

Fully Homomorphic Encryption (FHE), often regarded as the ultimate goal of encryption, enables an untrusted third party to process encrypted data without exposing sensitive information, ensuring secure computation in diverse networking and distributed systems [9]. Falcetta and Roveri [10] outline a detailed process for tailoring deep learning models to meet the strict computational and algorithmic limits of HE schemes, specifically the Brakerski-Fan-Vercauteren (BFV) scheme. Instead of BFV, they used the RNS-CKKS (Residue Number System-Cheon-Kim-Kim-Song) FHE scheme for classifying images on encrypted data using Resnet-20 (Residual Network). This gave them accuracy similar to that of models trained on plaintext data while keeping the data private [11]. Bootstrapping was a major step forward that got around the limits of classical FHE. It allows for any level of computation depth without needing large parameters that are too big to compute. A full framework was built on this, including the CKKS approach, to make it possible for AI services that protect privacy to work on edge devices with limited resources [12]. We tested this method by using encrypted calculations on four main AI models: Linear Regression, Fisherface (for recognising faces), Convolutional Neural Network (CNN), and Density-Based Spatial Clustering of Applications with Noise (DBSCAN, for clustering). Improvements in hardware acceleration and better software libraries are making the performance limits of traditional FHE less and less important. In

How to cite this article:
Lavanya B, Janani S (June 09, 2026) Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques. Cureus J Comput Sci 3 : es44389-026-00110-1. DOI https://doi.org/10.7759/s44389-026-00110-1

this context, a functional privacy-preserving system for tree-based ensemble models such as XGBoost has been created for sensitive healthcare information, employing HE grounded in ring learning with errors secured by the CKKS S(IP) framework [13]. This keeps patient information private.

Even though FHE lets you compute on encrypted data, it still has significant issues, such as being 10^3-10^5 times slower than normal operations and requiring frequent bootstrapping. As explored in [14-18], hybrid approaches that combine FHE with other cryptographic primitives or frameworks are becoming increasingly popular for fixing these challenges. Table *1* compares key HE algorithms to show how far they have come and how useful they are in the real world. It shows when they were discovered, how hard they are to use, and their usability.

HE Algorithm	HE Type	Year	Complexity	Usability
Privacy Homomorphism	Foundation	1970	Less	No
RSA	Partial (Multiplicative)	1977	Less	No
ElGamal	Partial (Multiplicative)	1985	Less	No
Paillier	Partial (Additive)	1985	Less	No
Gentry	Fully (Lattice-Based)	2009	Very High	No
BGV	Fully (Ring-LWE Based)	2011	High (but optimized for depth)	Yes - IoT, Big Data
GSW	Fully (LWE without linearization)	2013	Moderate (less than BGV)	Yes - Edge Computing
CKKS	Approximate HE	2017	Moderate (optimized for ML workloads)	Yes - ML

TABLE 1: Comparison of HE Algorithms

BGV, Brakerski–Gentry–Vaikuntanathan; CKKS, Cheon–Kim–Kim–Song; GSW, Gentry–Sahai–Waters; HE, Homomorphic Encryption; IoT, Internet of Things; LWE, Learning with Errors; RSA, Rivest–Shamir–Adleman

Table *1* highlights the progression of HE from computationally efficient but functionally restricted partial schemes toward more expressive fully homomorphic frameworks. While RSA and ElGamal provide lower computational complexity, their limited operational capabilities constrain their use in PPML. Conversely, CKKS has emerged as a prominent scheme for ML workloads owing to its support for efficient approximate arithmetic. Nevertheless, FHE schemes continue to exhibit substantial computational and memory requirements, underscoring the ongoing challenge of balancing privacy guarantees with practical scalability and performance.

Several well-known open-source libraries exist to help people use HE schemes effectively on encrypted data while still protecting their privacy [19]. Some well-known libraries are as follows:

• Microsoft SEAL (2022): Recognized for its easy-to-use API and high performance. It has a lot of well-documented examples to help with implementation.

• PALISADE (2019): This framework provides efficient lattice cryptography primitives and advanced HE techniques, enabling modular development of privacy-centric applications.

• Helib (2020): This library, made by IBM, speeds up the Smart-Vercauteren ciphertext packing method, making it easier to process encrypted data.

• HEAAN (2018): Enables approximate arithmetic on encrypted fixed-point numbers, emulating floating-point precision with tolerable error margins.

Zero-knowledge proof

Zero-knowledge proofs (ZKPs) are a kind of cryptography that enables one side, the prover, to convince the other, the verifier, of the validity of the statement through interactive proofs without revealing any additional knowledge [20]. An interactive proof system consists of two interactive Turing machines.

1. Prover (A): Has unlimited computational power.

2. Verifier (B): Runs in probabilistic polynomial time.

3. For a language L, the interactive proof system (A, B) for L meets the following requirements:

4. Completeness: The verifier accepts with high probability if x $\in$ Lx \ in L.

5. Soundness: If x $\notin$ Lx \ not in L, no prover can convince the verifier with non-negligible probability.

6. Zero knowledge: The prover does not give away any information during the protocol except for the fact that the statement is true.

ZKPs in ML allow the confirmation of things like model accuracy or forecasts without revealing sensitive model specifics. This makes it possible to safely validate calculations while maintaining the privacy of the underlying data [21].

Non-Interactive Proof System

Since interactive protocols require multiple cycles between the proving party and the verification expert, while non-interactive zero-knowledge (NIZK) proofs require only one communication, in 1988, NIZK proofs [22] were constructed using shared random strings under the quadratic residuosity assumption. NIZKs are extensively utilized in fields such as distributed ML because these proofs provide independent verification with effectiveness, scalability, and security.

zk-SNARKs and zk-STARK

ZKPs allow one person to demonstrate to another the truthfulness of a given claim without disclosing any underlying information by enabling quick and private verification. Two well-known ZKPs systems are zk-SNARKs (Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge) and zk-STARKs (Zero-Knowledge Scalable Transparent Arguments of Knowledge). zk-SNARKs let you prove that you know something without giving away the actual information. Bitansky et al. [23] made zk-STARKs by using publicly verifiable randomness and hash-based encryption, which provide scalability, transparency, and post-quantum security. These properties enable verifiable computation, secure authentication, and confidential transactions in privacy-preserving applications.

Hybrid Approach

Nassar et al. [24] used zk-SNARKs and zk-STARKs together to make it easier to quickly and safely check decentralised applications. This protects against quantum attacks and keeps the integrity of the computations. The method uses advanced polynomial commitment schemes like KZG (Kate-Zaverucha-Goldberg) commitments, DARK commitments (for post-quantum security), and the Poseidon hash function to make sure that both privacy and verifiability are possible.

ZKP for Verified ML

zkCNN is a ZKP method [25] that validates CNN-based ML predictions while preserving model privacy. It addresses the challenge of ensuring prediction integrity by enabling users to verify that results come from genuine CNN models without revealing any details about the model parameters or architecture. Blockchain designated verifier proof, a type of ZKP, protects privacy in blockchain apps by letting verifiers make fake proofs [26]. It has a designated-verifier mechanism and a Chameleon Hash function to make non-transferability better.

$$r \leftarrow \mathrm{CH}(pk, g^k, p)$$

g - makes a cyclic group G

k - random number from a big prime field

p - random text that holds the hash function

KAIZEN (Sumcheck-based proof for verifying each gradient descent iteration) makes iterative training through gradient descent possible by giving proofs that do not depend on how many iterations there are [27]. This proposed method may effectively be employed in copyright verification, distributed training, and proof of model ownership, since it enables a prover to demonstrate successful model training on a committed dataset without disclosing further information. KAIZEN makes it possible for deep neural networks to use zkPoT (Zero Knowledge Proof of Training). The Sum-check Protocol is used to prove that steps in gradient descent are correct. A cryptographic system called zkLLM [28] was built to safely check LLM calculations. Non-arithmetic operations, such as activation functions, are taken care of, but demonstrations of attention processes are improved using tlookup, a tensor-based lookup argument made for deep learning operations. This method makes sure that verification is scalable and useful for real-world applications by making sure that it is comprehensive with a very minimal chance of mistakes. In 15 minutes, zkLLM makes proofs, checks them in less than 3 seconds, and performs them quickly in memory (using just 23GB of GPU RAM). To enable secure and verifiable ML, Aziz et al. [29] suggest ZkVML (Zero-Knowledge Verifiable Machine Learning), a privacy-preserving system that integrates zk-SNARKs, zk-STARKs, arithmetic circuits, and neural network inference verification. Without disclosing sensitive input data, model parameters, or intermediate calculations, the framework guarantees that ML predictions are computed accurately. Major issues with outsourced AI systems, such as inference tampering, model manipulation, and privacy leaks in insecure cloud environments, are resolved by ZkVML. For small and medium-sized neural network models, experimental assessment shows effective proof production, minimal verification latency, robust privacy guarantees, and consistent inference accuracy. By combining ZKPs with deep learning techniques, a framework was developed to address important issues in decentralised finance, such as transaction privacy leaks, fraudulent transaction activity, and lack of trust in decentralised verification systems [30]. The work focuses on making it possible to verify transactions in a safe and private manner without disclosing private financial data during blockchain validation. To improve transaction integrity and intelligent anomaly analysis, the system uses deep learning-based fraud detection, decentralised verification methods, and cryptographic proof generation. For secure decentralised finance environments, experimental evaluation shows enhanced verification accuracy, effective fraud detection capability, decreased privacy threats, and acceptable computing overhead and verification latency.

Secure multi-party computation

Secure multi-party computation, or SMPC, is a sort of private computing that keeps data sources and systems safe and secret while also keeping the data itself safe [31]. It also lets many people work together to compute a function using their own data without having to share that information with each other, which enables the utilization of data without sacrificing privacy.

Two-Party Computation

In the early 1980s, Andrew Yao first proposed the idea of SMPC (Through Famous Yao's Millionaires' Problem) [32]. This work applies the Millionaires' Problem as a prime illustration of secure two-party computation. In this problem, two individuals compare their wealth without giving exact numbers. Yao's cryptographic protocols let Alice and Bob safely compute a function f(i,j) by turning it into a circuit and encrypting the gates, by preserving their privacy. This allows the parties to assess the circuit without being aware of each other's inputs. In advance of controlling knowledge transfer in

the protocol M = (MA, MB), Andrew Yao formalized the garbled circuits [33] technique in 1986 by introducing a new cryptographic tool that generates an integer N = p·q such that its prime factors p and q remain hidden from both parties individually but recoverable jointly.

Multi-Party Computations

The GMW protocol [34], developed by Goldreich, Micali, and Wigderson in 1987, constitutes one of the fundamental protocols in SMPC. A Boolean circuit comprising AND and XOR gates is employed to represent the function $f(x_1, x_2, ..., x_n)$. SMPC is a privacy-preserving approach that allows multiple participants to jointly compute a function $(y_1, y_2, ..., y_i) \rightarrow f(x_1, x_2, ..., x_i)$, where each party Pi provides input xi and computes yi. In this setting, each party contributes its own input and receives the corresponding output. However, no party learns anything about the other's inputs beyond what can be inferred from the final result. As a result, each party's data is kept safe and confidential while still enabling a practical group calculation that can be widely used in decentralized applications and untrusted environments. Figure *3* illustrates a multi-party computation model where multiple participants submit inputs to a function *f*, which computes corresponding outputs without revealing individual inputs, based on the concepts discussed by Goldreich et al. [34].

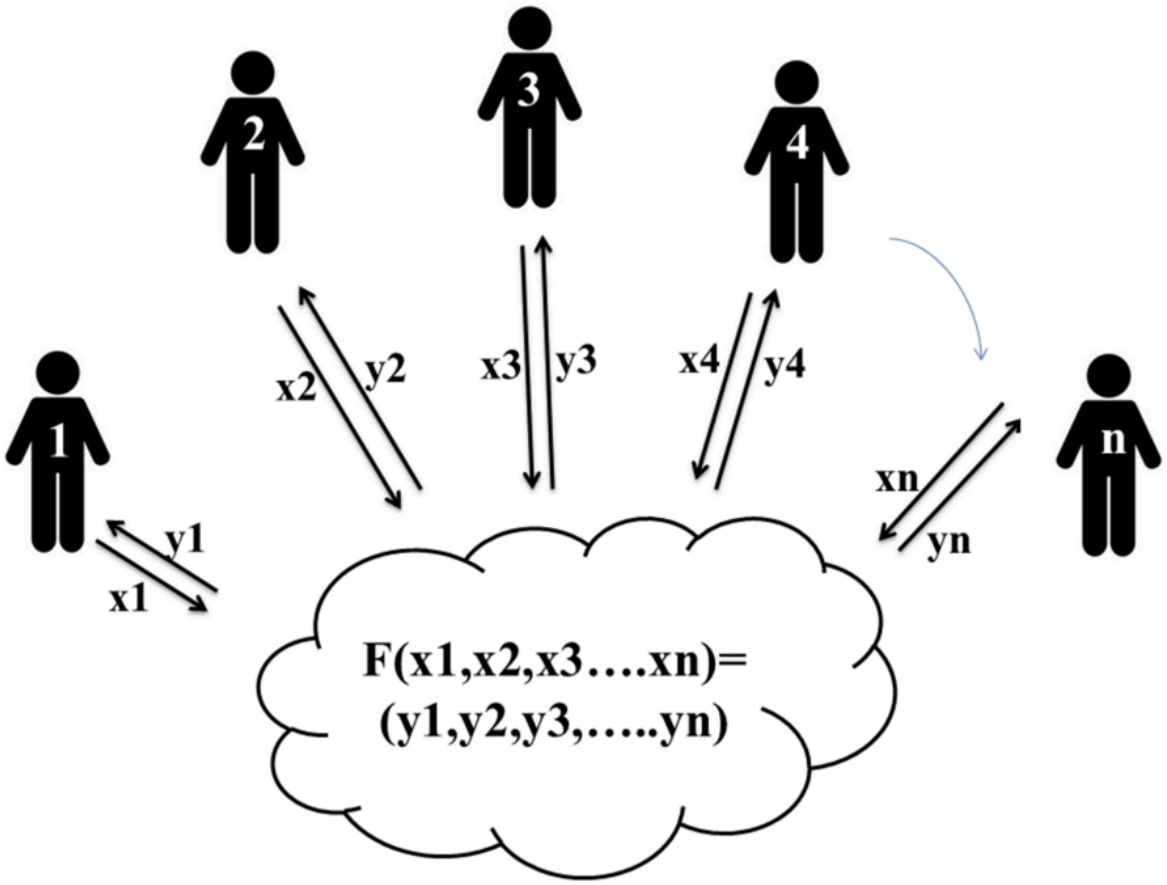

FIGURE 3: Diagram of Multi-Party Computations

Schematic representation developed based on the concepts discussed by Goldreich et al. [34].

SMPC on a Privacy-Preserving Approach

To improve PPML performance in cloud contexts, Hernandez et al. [35] proposed an optimization framework for pipelined SMPC. The framework employs a multi-objective optimization approach to minimize execution time and lower cloud costs. Their research advances effective and scalable PPML by shedding light on trade-offs between latency, bandwidth, and security. Although this research predominantly depends on CPU-based implementations, recent studies have commenced examining GPU-assisted computation. An example is CRYPTGPU [36], a GPU-accelerated MPC framework using PyTorch and CrypTen, which shows that GPU-enabled SMPC scales well for deep learning, allowing private training and inference on models like ResNet-152, with a 36-fold speedup over CPU-based frameworks like FALCON.

Expanding on the search for deep learning based on scalable SMPC, a collection of three-party protocols has been suggested by authors for the safe training of deep neural networks utilizing SMPC [37] without altering algorithms like Adam or SoftMax. The creation of 23-bit fixed-point arithmetic protocols that adhere to single-precision accuracy requirements utilized in frameworks such as TensorFlow and PyTorch is part of their technique. On networks like VGG16, significant performance improvements - up to 51× speedups - are observed during deep learning model training. To address real-world challenges in healthcare data sharing, authors proposed SMPC-based framework [38] that mitigates privacy leaks, unauthorized access, inference attacks, and regulatory compliance issues in sensitive medical environments. The study incorporates cutting-edge privacy-preserving technologies like SMPC, Garbled Circuits, Additive and Multiplicative Secret Sharing, HE, Differential Privacy (DP), and ZKPs. The framework tackles issues including data breaches, eavesdropping, malevolent participant attacks, and computational trust constraints that are related to centralised healthcare systems. While adhering to GDPR and HIPAA regulations, the authors further enhance security with error-detection techniques, secure communication channels, and cryptographic primitives. Computational efficiency, communication overhead, scalability, privacy protection, data confidentiality, integrity, and availability are the main concerns of performance evaluation. Experimental case studies in EHR analysis, secure health monitoring, and genomic data sharing show effective privacy-preserving cooperation without disclosing private patient data. A novel framework [39] tackles the high costs associated with fully encrypted training by keeping only labels private and secret-shared, which is useful for applications such as ad conversion prediction. It introduces two SMPC-DP hybrid protocols: one employing Randomised Response for label perturbation and another offering a label-private version of DP-SGD, enabling efficient training without full encryption SMPC. This contribution aligns with hybrid approaches that combine hardware-based Trusted Execution Environments, HE, or SMPC with DP.

The following section introduces DP, a statistical approach, and outlines its key noise-addition mechanisms as well as recent advances in differentially private ML.

Differential privacy

DP is a formal method that provides useful aggregate data while safeguarding individual privacy by adding carefully calibrated noise to outputs, inputs, or gradients [40]. It strikes the right balance between privacy and usefulness, making sure that no one can deduce the data of one individual. The Laplace, Gaussian, and exponential mechanisms are the most common ways to get DP. Each one operates optimally with specific sorts of data queries and levels of privacy.

Laplace mechanism: The Laplace mechanism achieves DP by adding noise that follows a Laplace distribution to the results of numerical queries (counts, sums, averages). Noise magnitude depends on global sensitivity and privacy budget ε, with the distribution centered at 0 and scale parameter *b*.

$$\mathbf{Lap}(x \mid b) = \frac{1}{2b}\exp\left(-\frac{|x|}{b}\right)$$

Despite its effectiveness, the Laplace mechanism ignores range constraints. Croft et al. [41] suggest a technique for normalizing and truncating the Laplace probability density function to solve this problem. This change scales the probability density function so that its total probability remains 1 and limits the output to a valid range. This approach increases usefulness by raising the probability density around the genuine query value and minimizing predicted distortion.

How to cite this article:
Lavanya B, Janani S (June 09, 2026) Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques. Cureus J Comput Sci 3 : es44389-026-00110-1. DOI https://doi.org/10.7759/s44389-026-00110-1

Gaussian mechanism: The Gaussian mechanism also uses additive noise, but it defines privacy guarantees based on the hypothesis testing trade-off between two shifted Gaussian distributions [42]. It is helpful when the extra noise needs to meet the standards of (ε, δ)-DP, which is a less strict version of pure DP. This mechanism is often used when Laplace noise could cause too much instability in high-dimensional data contexts or large-scale ML systems.

Exponential mechanism: The Exponential mechanism is a key part of DP [43]. It is used when the output is categorical or non-numeric, such as when identifying a feature or class label. The exponential mechanism works best for tasks such as classification, hyperparameter selection, or feature selection, where adding numerical noise is ineffective or unproductive.

Differential Private ML With Enhanced Data Confidentiality

One method for training ML models with DP is PATE (Private Aggregation of Teacher Ensembles) [44]. This is very crucial for deep learning. PATE is the process of using many "teacher" models that have been trained on subsets of the data to create "noisy labels" for the training data. The disruptive labels are then used to train the final "student" model. This helps maintain DP. The PATE-ADLM (Adaptive Deep Learning Mechanism) [45] architecture was suggested to make deep learning more effective and flexible. Through this work, the author integrated PATE with ADLM, which was made just for CNNs. ADLM changes the learning settings in real time based on how complex the input data is and what the instructor models say.

ADLM modifies the learning parameters in real time. Additionally, the framework emphasizes that its adaptive learning approach simplifies models and reduces training time. In 2024, Ju et al. [46] addressed the key limitations of using the Laplace mechanism with the PATE method. The integration of exponential processes instead of laplace mechanism picks outcomes at random based on a scoring function, which works better for outputs that can be grouped into categories, like class labels, therefore making it perfect for applications like image classification where it is important to be sure that you choose the right class else results in high sensitivity to utility-score design, making the privacy-utility trade-off difficult to control and potentially causing inconsistent performance across datasets.

Privacy attacks in ML

ML models trained on sensitive data face multiple privacy attacks that attempt to infer information about the data, model parameters, or training process. This section summarizes the main attack classes and maps each to the corresponding PPML defences, thereby highlighting how different cryptographic and statistical protections help reduce specific leakage vectors.

Model Inversion Attack/Reconstruction Attack

The goal of a model inversion attack is to get sensitive information from a trained model. GuardML [17] protects against this by using a hybrid HE framework to encrypt data and parameters. The service provider never sees plaintext inputs or decrypted model parameters, so an attacker cannot get to useful gradients, weights, or activations. This stops the usual leakage pathways that are used in model inversion attacks. On the other hand, zkCNN [25] checks CNN predictions without revealing parameters or intermediate data, thereby restricting the reconstruction of the information needed.

Gradient Leakage/Deep Leakage From Gradients

Gradient-based attacks show a central privacy hole in distributed ML. They let attackers flip shared gradients and get back sensitive training data. The SMPC-based Adam framework [37] stops this from happening by doing all calculations through three-party secret sharing, which makes sure that gradients are never shown and cannot be inverted.

Membership Inference Attack

Membership inference attacks aim to find whether a given data item was a subset of the model's learning set under adversarial probing. DP-based models implicitly address membership inference attacks by imposing bounds on the influence of any individual training sample, particularly through a noise mechanism [41] and knowledge transfer via

noisy aggregation (PATE) [44]. where multiple teacher models are trained to supervise a student model, ensuring that membership information from any single training record is obscured while still preserving the overall utility of the learned model.

Inferential analysis

Table *2* compares cryptographic and statistical techniques in terms of scalability and trade-offs, while indicating the ML pipeline components protected by each approach.

How to cite this article:
Lavanya B, Janani S (June 09, 2026) Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques. Cureus J Comput Sci 3 : es44389-026-00110-1. DOI https://doi.org/10.7759/s44389-026-00110-1

Cryptographic Approach	ML Pipeline Security	ML/DL Paradigm It Supports	Reported Results	Scalability	Limitation and Trade-Offs
HE [11]	Input data privacy	Optimized for CNNs (LeNet-1/5)	Accuracy: 90.67% on Cifar-10 dataset	Enables deeper networks but increases memory/computation costs quadratically	1.5×10^5 slower ~$10^2\times$ more memory
HE [17]	Privacy preserving inference	Shallow CNN, Compatible with PocketNN	Accuracy: 87.06% on ECG (MIT-BIH dataset)	Supports multi-client scenarios by centralizing key management at the analyst	Quantization constraints, CSP burden
ZKP [25]	Verifiable inference	CNNs (VGG16, LeNet, ReLU)	0.44 s to generate a proof on the MNIST dataset	Supports VGG16 (15M parameters) and 16-layer CNNs	Faster prover but recursion adds memory overhead
ZKP [28]	Verifiable inference	Transformer-based LLMs	15 m to generate proof on 13 billion parameter model on the C4 dataset	Supports 13B-parameter models (OPT, LLaMa-2) with 2048-token sequences	Limited to 16-bit fixed-point arithmetic
SMPC [39]	Secure collaborative computation	feedforward networks	MPC Runtime: 1.77 s in randomized response and 4.02 s on Label DP-SGD	Extendable to larger networks like AlexNet and VGG16	High communication and computation overhead, especially in WAN settings
DP [44]	Training stage of ML pipeline	CNN and ADLM	93.18% for MNIST dataset and 87.79% for the SVHN dataset with $\varepsilon = 0.05$	Adaptive to Fewer Teacher models	Trade-off between model accuracy and privacy
DP [46]	Privacy preserving training	DL and ML models especially CNN architectures (LeNet-5)	98.27% on MNIST dataset and 64.79% on the Cifar-10 dataset	Supports a large-scale query model with fixed privacy budget	Human-in-the-loop introduces dependency

TABLE 2: Inferential Analysis of Cryptographic Approaches in PPML

ADLM, Adaptive Deep Learning Mechanism; CNN, Convolutional Neural Network; CSP, Cloud Service Provider; DL, Deep Learning; HE, Homomorphic Encryption; LLM, Large Language Model; ML, Machine Learning; MNIST, Modified National Institute of Standards and Technology; MPC, Multi-Party Computation; PPML, Privacy-Preserving Machine Learning; ReLU, Rectified Linear Unit; SMPC, Secure Multi-Party Computation; WAN, Wide Area Network; ZKP, Zero-Knowledge Proof

The comparative analysis reveals that the practical adoption of PPML techniques is largely governed by trade-offs between privacy guarantees, computational efficiency, scalability, and model utility. HE provides strong confidentiality by enabling computation on encrypted data, but its computational and memory overhead remain significant barriers to large-scale deployment. ZKPs offer verifiable ML and enhanced trustworthiness; however, proof generation can introduce considerable computational costs for complex models. SMPC facilitates privacy-preserving collaborative learning without exposing sensitive information, although communication overhead may limit its efficiency in distributed environments. DP provides formal privacy guarantees and strong scalability but may reduce model utility due to noise injection. These observations suggest that hybrid PPML frameworks combining cryptographic and statistical approaches represent a promising direction for balancing privacy, efficiency, and practical deployment requirements.

Future work

The surveyed literature consistently shows a privacy-utility trade-off. However, current research rarely explores beyond this binary perspective to include new or combined trade-off metrics, such as robustness to distribution shifts, resilience to data impairments, or model interpretability. This results in a significant evaluation gap in the current PPML research.

Concurrently, advancements in lightweight cryptographic accelerators and adaptive hybrid protocols that effectively combine statistical privacy features with cryptographic techniques are needed for successful deployment. These adaptive methods minimize costs while preserving strong privacy guarantees by tailoring protection levels based on data sensitivity and system constraints. In order to transform PPML from theoretically sound models into dependable, scalable, and implementable systems, these criteria must be followed collectively.

Conclusions

This survey provided a comprehensive overview of PPML by examining cryptographic techniques, including HE, ZKPs, and SMPC, alongside DP as a statistical privacy-preserving approach. The comparative analysis demonstrates that no single technique can independently satisfy the diverse requirements of privacy, utility, scalability, and computational efficiency across all ML applications. Rather, each approach offers distinct advantages and limitations, resulting in important trade-offs that must be carefully considered based on the deployment scenario and threat model. The findings further indicate a growing trend toward hybrid PPML frameworks that combine cryptographic and statistical methods to leverage complementary strengths while mitigating individual limitations. Although substantial progress has been made in enabling privacy-aware ML, challenges related to computational overhead, scalability, communication complexity, and privacy-utility trade-offs continue to hinder widespread adoption. Consequently, PPML should be viewed as a rapidly evolving research area rather than a fully mature solution space. Continued advances in efficient hybrid architectures, standardized evaluation methodologies, and real-world deployment studies will be essential for developing trustworthy and practical privacy-preserving applied intelligence systems.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Acquisition, analysis, or interpretation of data: B Lavanya

How to cite this article:
Lavanya B, Janani S (June 09, 2026) Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques. Cureus J Comput Sci 3 : es44389-026-00110-1. DOI https://doi.org/10.7759/s44389-026-00110-1

Critical review of the manuscript for important intellectual content: B Lavanya

Supervision: B Lavanya

Concept and design: S Janani

Drafting of the manuscript: S Janani

Disclosures

Conflicts of interest: In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

Data Availability Statements

Data sharing is not applicable. This work is theoretical; no datasets were generated or analyzed.

Acknowledgements

This paper is submitted through the International Conference on Applied Intelligence and Innovative Computational Solutions (IC-AIICS 2026), organized by the Gandhigram Rural Institute, India. This research is supported by the University Grants Commission (UGC) of India under the Junior Research Fellowship (JRF) scheme.

References

1. Rechberger C, Walch R: Privacy-preserving machine learning using cryptography. Security and Artificial Intelligence. Batina L, Bäck T, Buhan I, Picek S (ed): Springer, Cham, Switzerland; 2022. 13049:109-129. 10.1007/978-3-030-98795-4_6
2. Syed D, Refaat SS, Bouhali O: Privacy preservation of data-driven models in smart grids using homomorphic encryption. Information. 2020, 11:357. 10.3390/info11070357
3. Gilbert C, Gilbert M: The effectiveness of homomorphic encryption in protecting data privacy. International Journal of Research Publication and Reviews. 2024, 5:3235-3256. 10.2139/ssrn.5259722
4. Ameur Y, Bouzefrane S, Audigier V: Application of homomorphic encryption in machine learning. Emerging Trends in Cybersecurity Applications. Daimi K, Alsadoon A, Peoples C, El Madhoun N (ed): Springer, Cham, Switzerland; 2023. 391-410. 10.1007/978-3-031-09640-2_18
5. Yan X, Zhou G, Huang Y, Meng W, Nguyen AT, Huang H: Secure estimation using partially homomorphic encryption for unmanned aerial systems in the presence of eavesdroppers. IEEE Transactions on Intelligent Vehicles. 2025, 10:2508-2518. 10.1109/tiv.2024.3378288
6. Paillier P: Public-key cryptosystems based on composite degree residuosity classes. Advances in Cryptology — EUROCRYPT '99. Stern J (ed): Springer, Berlin, Heidelberg; 1999. 1592:223-238. 10.1007/3-540-48910-X_16
7. Rivest RL, Shamir A, Adleman L: A method for obtaining digital signatures and public-key cryptosystems. Communications of the ACM. 1978, 21:120-126. 10.1145/359340.359342
8. Elgamal T: A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Transactions on Information Theory. 1985, 31:469-472. 10.1109/tit.1985.1057074
9. Pulido-Gaytan LB, Tchernykh A, Cortés-Mendoza JM, Babenko M, Radchenko G: A survey on privacy-preserving machine learning with fully homomorphic encryption. High Performance Computing. CARLA 2020. Nesmachnow S, Castro H, Tchernykh A (ed): Springer, Cham, Switzerland; 2021. 1327:115-129. 10.1007/978-3-030-68035-0_9
10. Falcetta A, Roveri M: Privacy-preserving deep learning with homomorphic encryption: an introduction. IEEE Computational Intelligence Magazine. 2022, 17:14-25. 10.1109/mci.2022.3180883

11. Lee JW, Kang H, Lee Y, et al.: Privacy-preserving machine learning with fully homomorphic encryption for deep neural network. IEEE Access. 2022, 10:30039-30054. 10.1109/access.2022.3159694

12. Khan MJ, Fang B, Cimino G, Cirillo S, Yang L, Zhao D: Privacy-preserving artificial intelligence on edge devices: a homomorphic encryption approach. 2024 IEEE International Conference on Web Services (ICWS). IEEE, Shenzhen, China; 2024. 395-405. 10.1109/ICWS62655.2024.00061

13. Balaban B, Magara SS, Yilgor C, et al.: Privacy-preserving machine learning (PPML) inference for clinically actionable models. IEEE Access. 2025, 13:37431-37456. 10.1109/access.2025.3540261

14. Sivaraman HK, Rangaiah L: Optimizing data survivability in unattended wireless sensor networks: a machine learning approach to cluster head selection and hybrid homomorphic encryption. Indonesian Journal of Electrical Engineering and Informatics (IJEEI). 2025, 13:137-155. 10.52549/.v13i1.5998

15. Cho J, Ha J, Kim S, et al.: Transciphering framework for approximate homomorphic encryption. Advances in Cryptology - ASIACRYPT 2021. Tibouchi M, Wang H (ed): Springer, Cham, Switzerland; 2021. 13092:640-669. 10.1007/978-3-030-92078-4_22

16. Li B, Micciancio D, Schultz-Wu M, Sorrell J: Securing approximate homomorphic encryption using differential privacy. Advances in Cryptology - CRYPTO 2022. Dodis Y, Shrimpton T (ed): Springer, Cham, Switzerland; 2022. 13507:560-589. 10.1007/978-3-031-15802-5_20

17. Frimpong E, Nguyen K, Budzys M, Khan T, Michalas A: GuardML: efficient privacy-preserving machine learning services through hybrid homomorphic encryption. SAC '24: Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing. 2024, 953-962. 10.1145/3605098.3635983

18. Correia P, Costa I, Amorim I, Maia E, Praça I: Federated learning: an approach with hybrid homomorphic encryption. Computer Security. ESORICS 2025 International Workshops. Laborde R, Garcia-Alfaro J, Yazdinejad A, et al. (ed): Springer, Cham, Switzerland; 2026. 16233:254-273. 10.1007/978-3-032-16165-9_15

19. Al Badawi A, Bates J, Bergamaschi F, et al.: OpenFHE: open-source fully homomorphic encryption library. WAHC'22: Proceedings of the 10th Workshop on Encrypted Computing & Applied Homomorphic Cryptography. 2022, 53-63. 10.1145/3560827.3563379

20. Goldwasser S, Micali S, Rackoff C: The knowledge complexity of interactive proof systems. SIAM Journal on Computing. 1989, 18:186-208. 10.1137/0218012

21. Zhang Y, Fan Z: Research on Zero knowledge with machine learning. Journal of Computing and Electronic Information Management. 2024, 12:105-108. 10.54097/6awase9w

22. De Santis A, Micali S, Persiano G: Non-interactive zero-knowledge proof systems. Advances in Cryptology — CRYPTO '87. Pomerance C (ed): Springer, Berlin, Heidelberg; 1988. 293:52-72. 10.1007/3-540-48184-2_5

23. Bitansky N, Canetti R, Chiesa A, Tromer E: From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. ITCS '12: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. 2012, 326-349. 10.1145/2090236.2090263

24. Nassar ST, Hamdy A, Nagaty K: Hybrid zk-STARK and zk-SNARK framework for privacy-preserving smart contract data feeds. 2024 International Conference on Computer and Applications (ICCA). IEEE, Cairo, Egypt; 2024. 1-7. 10.1109/ICCA62237.2024.10928100

25. Liu T, Xie X, Zhang Y: zkCNN: zero knowledge proofs for convolutional neural network predictions and accuracy. CCS '21: Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security. 2021, 2968-2985. 10.1145/3460120.3485379

26. Chi PW, Lu YH, Guan A: A privacy-preserving zero-knowledge proof for blockchain. IEEE Access. 2023, 11:85108-85117. 10.1109/access.2023.3302691

27. Abbaszadeh K, Pappas C, Katz J, Papadopoulos D: Zero-knowledge proofs of training for deep neural networks. CCS '24: Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security. 2024, 4316-4330. 10.1145/3658644.3670316

28. Sun H, Li J, Zhang H: zkLLM: zero knowledge proofs for large language models. CCS '24: Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security. 2024, 4405-4419. 10.1145/3658644.3670334

29. Aziz MB, Naushad AS, Siddiqui M, Shamsi JA: ZkVML: zero-knowledge verifiable machine learning. Asia Pacific Advanced Network. APANConf 2024. Herath D, Jayasinghe U, Ul Islam I, Ahmad J, Ragel R, Bandaranayake A (ed): Springer, Cham, Switzerland; 2025. 2412:220-239. 10.1007/978-3-031-89813-6_14

How to cite this article:
Lavanya B, Janani S (June 09, 2026) Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques. Cureus J Comput Sci 3 : es44389-026-00110-1. DOI https://doi.org/10.7759/s44389-026-00110-1

30. Kavipriya G, Prasanth L: Privacy-preserving transaction verification in decentralized finance using zero-knowledge proofs and deep learning. Information Security, Privacy and Digital Forensics. ICISPD 2024. Chaudhary NK, Iyengar S, Modi C, Patel SJ (ed): Springer, Singapore; 2026. 1456:139-149. 10.1007/978-981-95-2196-8_8

31. Khan S: Secure multi-party computation for privacy preservation in big data analytics. Journal of Big Data Privacy Management. 2024, 2:168-179. 10.71146/jbdpm42

32. Yao AC: Protocols for secure computations. 23rd Annual Symposium on Foundations of Computer Science (SFCS 1982). IEEE, Chicago, IL; 1982. 160-164. 10.1109/SFCS.1982.38

33. Yao AC-C: How to generate and exchange secrets. 27th Annual Symposium on Foundations of Computer Science (SFCS 1986). IEEE, Toronto, ON; 1986. 162-167. 10.1109/SFCS.1986.25

34. Goldreich O, Micali S, Wigderson A: How to play ANY mental game. STOC '87: Proceedings of the nineteenth annual ACM symposium on Theory of computing. 1987, 218-229. 10.1145/28395.28420

35. Hernandez R, Bautista OG, Akkaya K: Optimizing the parameters of pipelined multi-party computation for privacy-preserving machine learning applications. ICC 2024 - IEEE International Conference on Communications. IEEE, Denver, CO; 2024. 938-943. 10.1109/ICC51166.2024.10623002

36. Tan S, Knott B, Tian Y, Wu DJ: CryptGPU: fast privacy-preserving machine learning on the GPU [PREPRINT]. arXiv. 2021, 10.48550/arXiv.2104.10949

37. Attrapadung N, Hamada K, Ikarashi D, et al.: Adam in private: secure and fast training of deep neural networks with adaptive moment estimation. Proceedings on Privacy Enhancing Technologies. 2022, 2022:746-767. 10.56553/popets-2022-0131

38. Ahammed MF, Labu MR: Privacy-preserving data sharing in healthcare: advances in secure multiparty computation. Journal of Medical and Health Studies. 2024, 5:37-47. 10.32996/jmhs.2024.5.2.4

39. Yuan S, Shen M, Mironov I, Nascimento ACA: Practical, label private deep learning training based on secure multiparty computation and differential privacy [PREPRINT]. Cryptology ePrint Archive. 2021,

40. Selvi A, Liu H, Wiesemann W: Differential privacy via distributionally robust optimization. Operations Research. 2025, 74:356-376. 10.1287/opre.2023.0218

41. Croft W, Sack J-R, Shi W: Differential privacy via a truncated and normalized Laplace mechanism. Journal of Computer Science and Technology. 2022, 37:369-388. 10.1007/s11390-020-0193-z

42. Dong J, Roth A, Su WJ: Gaussian differential privacy. Journal of the Royal Statistical Society Series B: Statistical Methodology. 2022, 84:3-37. 10.1111/rssb.12454

43. Zhang S, Peng T, Sun P, Zhang J: Differential privacy work set selection method based on exponential mechanism for linear support vector machines. 2025 10th International Conference on Computer and Communication System (ICCCS). IEEE, Chengdu, China; 2025. 1-9. 10.1109/ICCCS65393.2025.11069575

44. Papernot N, Abadi M, Erlingsson U, Goodfellow I, Talwar K: Semi-supervised knowledge transfer for deep learning from private training data [PREPRINT]. arXiv. 2017, 10.48550/arXiv.1610.05755

45. Tang S, Wang W: The differential privacy framework for convolutional neural network using PATE, ADLM and SGD models. 2023 International Conference on Networking and Network Applications (NaNA). IEEE, Qingdao, China; 2023. 479-483. 10.1109/NaNA60121.2023.00085

46. Ju Q, Xia R, Li S, Zhang X: Privacy-preserving classification on deep learning with exponential mechanism. International Journal of Computational Intelligence Systems. 2024, 17:39. 10.1007/s44196-024-00422-x

How to cite this article:
Lavanya B, Janani S (June 09, 2026) Next-Generation Privacy in Applied Intelligence: A Unified Examination of Cryptographic and Statistical Privacy-Enhancing Techniques. Cureus J Comput Sci 3 : es44389-026-00110-1. DOI https://doi.org/10.7759/s44389-026-00110-1