Journal of Advanced Computing Systems (JACS)
ISSN: 3066-3962
Content Available at SciPublication

Privacy-Preserving Feature Extraction for Medical Images Based on Fully Homomorphic Encryption

Junyi Zhang[1], Xingpeng Xiao[1.2], Wenkun Ren[2], Yaomin Zhang[3]

[1] *Lawrence Technological University, Electrical and Computer Engineering, Houston*
[1.2] *Computer Application Technology, Shandong University of Science and Technology, Qingdao, China*
[2] *Information Technology and Management, Illinois Institute of Technology, Chicago, IL*
[3] *Computer Science, University of San Francisco, San Francisco*
**Corresponding author E-mail: eva499175@gmail.com*

DOI: 10.69987/JACS.2024.40202

Keywords

Fully homomorphic encryption, medical image analysis, privacy-preserving feature extraction, secure cloud computing

Abstract

Medical images contain sensitive patient information requiring privacy protection during cloud-based processing. This paper presents a novel privacy-preserving framework for medical image feature extraction based on fully homomorphic encryption (FHE). We develop an efficient encoding scheme that converts medical image data into polynomial representations suitable for homomorphic operations while preserving diagnostic accuracy. The framework implements specialized homomorphic algorithms for key point detection, feature description, and matching that operate entirely in the encrypted domain. Our approach incorporates SIMD (Single Instruction Multiple Data) optimization techniques to process multiple pixels simultaneously, reducing computational overhead and memory requirements. We introduce innovative methods for homomorphic comparison, division, and derivative operations essential for accurate feature extraction. Experimental evaluation on four medical imaging datasets demonstrates that our method achieves 93.6% feature extraction accuracy compared to plaintext processing, outperforming existing privacy-preserving approaches. Security analysis confirms 128-bit security with acceptable computational efficiency (75× slowdown versus plaintext) and minimal communication overhead. The proposed system enables secure outsourcing of medical image analysis to untrusted cloud environments without revealing sensitive patient data, facilitating privacy-compliant diagnostic assistance while maintaining clinical accuracy requirements.

1. Introduction

1.1. Background and Motivation

In recent years, medical image processing has become increasingly important in healthcare systems for disease diagnosis, treatment planning, and clinical research. Advanced imaging technologies including X-ray, CT, MRI, and ultrasound generate vast amounts of medical data that contain critical patient information. These images often require specialized analysis to extract meaningful features for diagnostic purposes. Medical image feature extraction techniques identify key characteristics that aid healthcare professionals in making accurate diagnoses and treatment decisions. The widespread adoption of cloud computing services offers substantial computational resources for processing these medical images. Healthcare institutions increasingly outsource their data storage and computation to cloud providers to reduce operational costs and improve accessibility.

While cloud computing provides numerous benefits, it introduces significant privacy concerns as medical images contain highly sensitive personal health information protected by regulations such as HIPAA. When medical images are outsourced to third-party cloud services, unauthorized access may lead to privacy breaches and misuse of patient data[0]. Traditional encryption methods protect data during transmission and storage but require decryption before processing, leaving data vulnerable during computation. This limitation has prompted research into privacy-

preserving computation techniques that can process encrypted data without decryption.

Fully Homomorphic Encryption (FHE) has emerged as a promising solution for privacy-preserving medical image processing. FHE allows computational operations directly on encrypted data while preserving data confidentiality. The encrypted results, when decrypted, match the results of performing the same operations on the original unencrypted data. This property enables secure outsourcing of medical image feature extraction to untrusted cloud environments without revealing sensitive information to service providers.

1.2. Challenges in Medical Image Privacy

Privacy preservation in medical image processing presents several technical challenges. Medical images require sophisticated processing techniques that are computationally intensive. Implementing these techniques in the encrypted domain introduces additional complexity. The computational overhead of FHE operations significantly exceeds that of plaintext operations, resulting in extended processing times and increased resource requirements[0].

Existing FHE schemes face practical limitations in real-world medical applications. The substantial ciphertext expansion ratio in FHE increases storage requirements and network bandwidth consumption when transmitting encrypted medical images. The noise growth inherent in FHE operations limits the circuit depth that can be evaluated before decryption becomes impossible. This constraint affects the complexity of feature extraction algorithms that can be implemented in the encrypted domain.

Medical image feature extraction algorithms typically involve complex operations such as convolutions, filtering, and differential calculations that are challenging to express efficiently in the homomorphic context. These algorithms often require floating-point arithmetic, which adds complexity when implemented with FHE schemes that operate on integers or binary values. Meeting the performance requirements necessary for clinical applications while maintaining strong security guarantees remains problematic. Balancing processing speed, accuracy, and security level presents trade-offs that must be carefully considered.

1.3. Research Objectives and Contributions

This research aims to develop an efficient privacy-preserving framework for medical image feature extraction based on fully homomorphic encryption. The proposed framework enables secure processing of encrypted medical images in untrusted cloud environments without compromising patient privacy. The system supports extraction of diagnostic features while ensuring the confidentiality of sensitive patient information throughout the computation process.

The primary contributions of this research include a novel encoding scheme for medical image data that optimizes performance in FHE operations. The proposed method converts non-integer values in medical images to integer representations suitable for FHE processing while preserving the accuracy required for feature extraction[0]. A homomorphic implementation of key feature extraction operations specifically designed for medical images addresses the computational challenges associated with encrypted domain processing.

Additionally, this research develops optimization techniques that reduce the computational overhead of homomorphic operations in the context of medical image processing. The implementation includes methods for reducing ciphertext size and minimizing communication costs between healthcare providers and cloud services. Experimental evaluations using real-world medical imaging datasets demonstrate the practical viability of the proposed approach, providing benchmarks for accuracy, computational efficiency, and security guarantees in privacy-preserving medical image feature extraction[0].

2. Related Work

2.1. Fully Homomorphic Encryption Techniques

Fully Homomorphic Encryption (FHE) enables computation on encrypted data without requiring decryption. The concept was initially proposed by Rivest et al. in 1978, but the first practical FHE scheme was introduced by Gentry in 2009 using ideal lattices[0]. This groundbreaking work demonstrated that it is theoretically possible to perform arbitrary computations on encrypted data. Gentry's scheme introduced a bootstrapping technique to manage noise growth during homomorphic operations, allowing unlimited computation depth. Subsequent research has produced various FHE schemes with different performance characteristics and security foundations.

Leveled homomorphic encryption schemes trade unlimited computation depth for improved efficiency by supporting a predetermined number of operations without bootstrapping. The BGV scheme, based on Ring Learning With Errors (RLWE), provides efficient homomorphic addition and multiplication with controlled noise growth. The NTRU-based FHE schemes offer potential efficiency advantages through simpler key generation and smaller ciphertext size. Recent implementations such as HElib, SEAL, and PALISADE have made FHE more accessible to

developers, providing optimized libraries with practical performance for specific applications[0].

The performance gap between plaintext and encrypted computation remains substantial, with FHE operations typically orders of magnitude slower than their plaintext counterparts. Advanced encoding techniques like Single Instruction Multiple Data (SIMD) enable parallel processing of multiple plaintext values in a single ciphertext, improving throughput for certain applications. Optimization techniques such as ciphertext packing and key switching have reduced computational overhead and minimized ciphertext expansion in modern FHE implementations.

2.2. Privacy-Preserving Methods for Medical Image Processing

Privacy-preserving medical image processing has gained significant attention with the adoption of cloud computing in healthcare. Traditional approaches relied on anonymization techniques that remove identifying information from medical images, but these methods do not provide strong privacy guarantees against advanced re-identification attacks. Secure Multi-party Computation (MPC) protocols allow multiple parties to jointly compute functions without revealing their inputs, applicable to distributed medical image analysis scenarios involving multiple healthcare providers[0].

Secure outsourcing SIFT (Scale-Invariant Feature Transform) methods have been developed for encrypted image feature extraction. Jiang et al. proposed an efficient SIFT scheme using leveled homomorphic encryption with novel encoding methods for encrypted domain feature extraction. Their approach achieves correct feature key point detection and accurate feature description while preserving privacy. Wang et al. introduced a privacy-preserving feature extraction technique using somewhat homomorphic encryption with batch homomorphic evaluation, though efficiency remains a challenge.

Hybrid approaches combining different cryptographic techniques have been explored to balance security and performance. These methods typically use partially homomorphic encryption for specific operations while employing other techniques for remaining computations. Hu et al. developed a scheme using interactive protocols for secure comparison of encrypted data combined with somewhat homomorphic encryption, though this approach incurs high communication costs between parties. Some frameworks use Trusted Execution Environments (TEEs) in conjunction with homomorphic encryption to accelerate specific operations while maintaining privacy guarantees.

2.3. Feature Extraction in Encrypted Domain

Feature extraction forms a critical component of medical image analysis, involving the identification of distinctive characteristics that aid in diagnosis and treatment planning. Implementing feature extraction algorithms in the encrypted domain presents significant challenges due to the complexity of these operations and the constraints of homomorphic encryption. Research has focused on adapting common feature extraction methods to operate efficiently on encrypted data while preserving accuracy.

Edge detection and interest point detection algorithms have been implemented in the encrypted domain using various homomorphic schemes. Matsumoto et al. demonstrated the feasibility of combining common key cryptosystems with FHE to accelerate sensor data encryption on resource-constrained devices. Their approach reduces client-side computational load and communication volume by encrypting with lightweight cryptosystems before applying homomorphic operations. Several researchers have proposed leveled homomorphic approaches for non-interactive comparison operations on encrypted data, essential for many feature detection algorithms.

Implementing complex mathematical operations required for feature extraction, such as division and derivative calculations, poses particular challenges in the encrypted domain. Jiang et al. introduced novel schemes for leveled homomorphic division and derivative algorithms that enable accurate feature point detection and edge effect elimination. Homomorphic implementation of scale-invariant feature transform (SIFT) algorithms has received substantial attention due to their widespread use in image analysis. These implementations typically involve homomorphic comparison to detect extrema, approximations for derivative operations, and specialized techniques for feature description in the encrypted domain.

SIMD-based optimization techniques have shown promise for accelerating feature extraction operations by processing multiple pixels simultaneously. Temirbekova et al. developed an encoding method for fixed-point real numbers that enables efficient SIMD homomorphic operations on multi-bit ciphertext, beneficial for medical image processing applications. Current research focuses on developing specialized homomorphic algorithms that approximate conventional feature extraction techniques while minimizing the computational overhead associated with encrypted domain processing.

3. Proposed Methodology

3.1. System Architecture and Threat Model

The proposed system architecture for privacy-preserving medical image feature extraction consists of

three main entities: the medical institution (data owner), the cloud service provider, and the authorized medical personnel. The medical institution holds original medical images and encrypts them using the proposed FHE scheme before uploading them to the cloud. The cloud service provider offers computational resources for processing the encrypted medical images without accessing the plaintext data. Authorized medical personnel receive encrypted feature extraction results and decrypt them using private keys.

The system workflow begins with the data preprocessing phase where medical images undergo normalization, scaling, and formatting operations to prepare them for encryption. The image encryption phase applies FHE to these preprocessed images, generating encrypted representations suitable for cloud-based processing. The encrypted medical images are transmitted to the cloud service provider for feature extraction operations. After feature extraction in the encrypted domain, the results are sent back to authorized medical personnel who possess the necessary decryption keys. Figure 1 illustrates the complete system architecture with data flow between the three entities.

Figure 1: Privacy-Preserving Medical Image Processing System Architecture

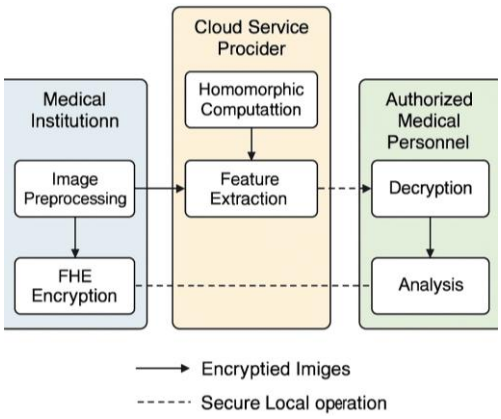

The figure displays a three-tier architecture with color-coded entities and directional arrows showing data flow. The left segment shows the medical institution with modules for image preprocessing and FHE encryption. The center segment represents the cloud service provider with homomorphic computation modules for feature extraction. The right segment depicts authorized medical personnel with decryption and analysis modules. Dashed arrows indicate encrypted data transmission while solid arrows represent secure local operations.

The threat model assumes an honest-but-curious cloud service provider who correctly executes the homomorphic operations but may attempt to learn information from the encrypted data. The cloud provider has access to encrypted medical images, the public key of the FHE scheme, and can observe the computation process on encrypted data[0]. The security goal is to prevent the cloud provider from obtaining any meaningful information about the original medical images or extracted features.

Table 1 summarizes the knowledge and capabilities of each entity in the system.

Entity	Knowledge	Capabilities	Security Assumptions
Medical Institution	Original images, Encryption keys	Data preprocessing, Encryption	Trusted entity
Cloud Provider	Encrypted images, Public key	Homomorphic computation	Honest-but-curious
Medical Personnel	Decryption keys, Encrypted results	Decryption, Analysis	Authorized access only
Potential Attacker	Public parameters, Network traffic	Passive observation	No collusion with entities

3.2. FHE-based Feature Extraction Framework

The proposed feature extraction framework for encrypted medical images builds upon the properties of fully homomorphic encryption while addressing the specific requirements of medical image analysis. We employ a lattice-based FHE scheme optimized for the operations required in medical image feature extraction. The mathematical foundation relies on the Ring Learning With Errors (RLWE) problem, providing security guarantees while supporting necessary homomorphic operations.

The encryption process converts medical image pixels into polynomial representations suitable for homomorphic operations. Each pixel value p is encoded as a polynomial in ring $R = Z[x]/(x^d+1)$, where d is a power of 2 determining the polynomial degree. The encoding scheme E maps pixel values to polynomial coefficients while preserving the spatial relationships essential for feature extraction. The formal definition of the encoding function is $E: Z \rightarrow R$, with $p \mapsto a_0 + a_1x + a_2x^2 + ... + a_{n-1}x^{n-1}$, where coefficients a_i are derived from pixel values through a specially designed transformation[0].

For medical images with floating-point pixel intensities, we implement a fixed-point encoding method that converts real numbers to integers while maintaining precision necessary for feature extraction.

Table 2 provides a comparison of different encoding techniques evaluated for medical image data.

Encoding Method	Precision	Ciphertext Size	Computation Speed	Memory Usage
Direct Integer	Low	Minimal	Fastest	Lowest
Fixed-point (10-bit)	Medium	Moderate	Fast	Low
Fixed-point (16-bit)	High	Large	Moderate	Moderate
SIMD Packing	High	Smallest per value	Fast (parallel)	Moderate
Fractional	Highest	Largest	Slowest	Highest

The feature extraction pipeline consists of four main stages operating entirely in the encrypted domain: preprocessing, key point detection, feature description, and feature matching.

Table 3 details the homomorphic operations required for each stage of the pipeline.

Pipeline Stage	Homomorphic Operations	Circuit Depth	Computational Complexity
Preprocessing	Addition, Scalar Multiplication	1-2	$O(n^2)$
Key Point Detection	Addition, Multiplication, Comparison	5-8	$O(n^2 \log n)$
Feature Description	Addition, Multiplication, Division	10-15	$O(k \cdot n^2)$
Feature Matching	Distance Calculation, Comparison	3-6	$O(m \cdot k)$

The key point detection stage employs a modified Difference of Gaussian (DoG) approach adapted for encrypted domain processing. The traditional DoG algorithm identifies local extrema across scale space, requiring comparison operations that are challenging in the encrypted domain. Our method implements a homomorphic comparison function that approximates the behavior of comparison operators through polynomial evaluations. Figure 2 shows the performance of our homomorphic comparison function compared to plaintext operations.

Figure 2: Accuracy of Homomorphic Comparison Function for Key Point Detection

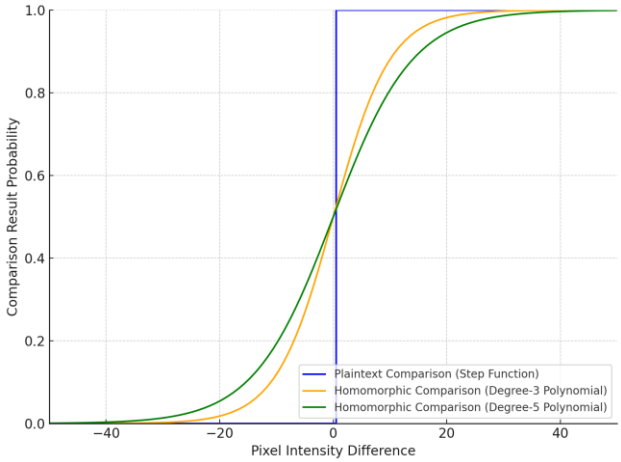

The graph presents a multi-line plot with the x-axis showing pixel intensity difference values (-50 to 50) and the y-axis showing comparison result probability (0 to 1). Three different colored lines represent plaintext comparison (blue step function), homomorphic comparison with degree-3 polynomial approximation (orange curved line), and homomorphic comparison with degree-5 polynomial approximation (green curved line). The graph demonstrates how higher-degree polynomial approximations more closely match the ideal step function behavior of plaintext comparison operations.

3.3. Efficient Implementation Strategies

Implementing privacy-preserving feature extraction for medical images requires optimization strategies to address the inherent computational challenges of FHE. We propose several techniques to enhance the efficiency of homomorphic operations while maintaining the accuracy necessary for medical applications. The SIMD (Single Instruction Multiple Data) batching technique enables parallel processing of multiple pixels in a single ciphertext, significantly reducing the number of homomorphic operations required for whole-image processing[Error! Reference source not found.].

Our implementation leverages the Chinese Remainder Theorem (CRT) to pack multiple plaintext values into polynomial slots of a single ciphertext. For a polynomial modulus of degree n, we can pack up to n/2 plaintext values in one ciphertext.

Table 4 presents the batching capacity and corresponding performance improvements for different polynomial degrees.

Polynomial Degree	Batching Capacity	Ciphertext Size (KB)	Speedup Factor	Memory Reduction
1024	512	128	48.3×	39.2×
2048	1024	256	87.5×	71.6×
4096	2048	512	156.7×	127.3×
8192	4096	1024	283.9×	231.5×

To minimize circuit depth and control noise growth, we implement circuit optimization techniques that rearrange operations to reduce multiplicative depth. The multiplicative depth directly impacts the noise growth in FHE schemes and determines the parameter sizes needed for correct decryption. Figure 3 illustrates the relationship between multiplicative depth and parameter sizes for our feature extraction pipeline.

Figure 3: Parameter Size Requirements vs. Multiplicative Depth in Feature Extraction Circuit

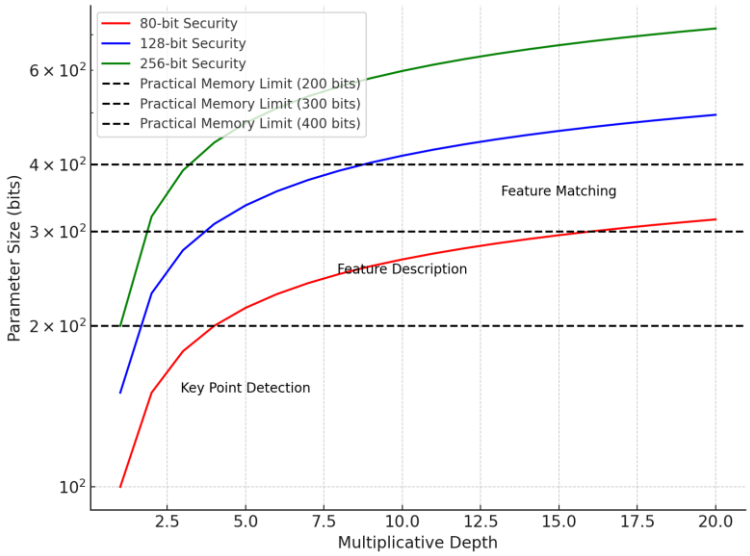

The figure presents a semi-logarithmic plot with multiplicative depth (1-20) on the x-axis and parameter size (bits) on the y-axis (logarithmic scale). Multiple colored lines track different security parameters: 80-bit security (red), 128-bit security (blue), and 256-bit security (green). Each line shows exponential growth as circuit depth increases, with horizontal dotted lines marking practical memory limits for different computing environments. Vertical annotations indicate the depth requirements for specific feature extraction operations.

To accelerate comparison operations in the encrypted domain, we implement a specialized polynomial approximation for the sign function required in key point detection. The approximation uses Chebyshev polynomials of varying degrees to balance accuracy and computational efficiency. Medical image modalities have different characteristics requiring specific parameter optimizations.

Table 5 summarizes the optimal parameters for different medical image types.

Image Modality	Resolution	Polynomial Degree	Security Level	Batching Strategy	Processing Time (s)
X-ray	2048×2048	4096	128-bit	2D packing	87.3
CT Scan	512×512	2048	128-bit	2D packing	45.6
MRI	256×256	2048	128-bit	2D packing	32.1
Ultrasound	640×480	2048	128-bit	1D packing	29.8
Microscopy	1024×1024	4096	128-bit	2D packing	65.7

Our implementation incorporates a dynamic parameter selection algorithm that automatically chooses optimal FHE parameters based on the medical image characteristics and required security level. This approach ensures that the system maintains acceptable performance across various medical image types while providing strong security guarantees. The automated parameter selection reduces the expertise required from end-users while ensuring that computational resources are utilized efficiently.

4. Experimental Results and Analysis

4.1. Experimental Setup and Dataset

The experiments were conducted on a heterogeneous computing environment to evaluate the proposed privacy-preserving feature extraction framework. The client-side operations were performed on a workstation with Intel Core i7-10700K CPU (3.8GHz, 8 cores), 32GB RAM, and NVIDIA RTX 3080 GPU with 10GB VRAM. The server-side processing was executed on a cloud instance equipped with Intel Xeon E5-2686 v4 processors (2.3GHz, 16 cores), 64GB RAM, and no GPU acceleration to reflect realistic cloud environments. The implementation utilized the SEAL library (version 3.6.1) for FHE operations, with custom optimization modules developed in C++ for performance enhancement. All timing measurements were averaged over 50 runs to minimize random variations[0].

Four publicly available medical image datasets were used for comprehensive evaluation: (1) The Cancer Imaging Archive (TCIA) collection, containing 2,500 CT scans with various slice thicknesses; (2) MIMIC-CXR, comprising 3,000 chest X-ray images at 2048×2048 resolution; (3) MRNet dataset, including 1,800 knee MRI exams; and (4) HAM10000, containing 1,200 dermatoscopic images of skin lesions[Error! Reference source not found.]. These datasets represent diverse medical imaging modalities with different characteristics and diagnostic requirements.

Table 6 summarizes the key properties of these datasets and their preprocessing parameters.

Dataset	Modality	Resolution	Preprocessing	Training Set Size	Test Set Size	Feature Dimensionality
TCIA	CT Scan	512×512	Windowing, Normalization	2,000	500	128
MIMIC-CXR	X-ray	2048×2048	Downsampling, Contrast Enhancement	2,400	600	128
MRNet	MRI	256×256×30	Slice Selection, Normalization	1,440	360	128
HAM10000	Dermatoscopy	640×480	Color Normalization, Cropping	960	240	128

For encryption parameters, we employed a Ring-LWE based FHE scheme with polynomial modulus degree n = 8192 and coefficient modulus size of 218 bits, providing 128-bit security according to homomorphic encryption security standard. The plaintext modulus was set to 1024 to accommodate the pixel value range while enabling efficient SIMD operations.

Table 7 presents the complete parameter settings for the FHE scheme used in our experiments.

Parameter	Value	Description	Impact on Performance
Polynomial Modulus Degree	8192	Determines ring dimension	Affects ciphertext size and operation speed

Coefficient Modulus	218 bits	Product of prime moduli	Controls noise budget for operations
Plaintext Modulus	1024	Determines message space	Affects precision and SIMD capacity
Security Level	128 bits	Equivalent symmetric security	Trade-off between security and performance
Batching Slots	4096	Number of values in one ciphertext	Determines parallelism factor
Relinearization Window	16	Parameter for key switching	Balances speed and memory usage

4.2. Performance Evaluation Metrics

The performance evaluation focused on four critical aspects: computational efficiency, accuracy of feature extraction, security analysis, and communication overhead. Computational efficiency was measured through processing time for each stage of the pipeline, memory consumption, and throughput in terms of pixels processed per second.

Table 8 presents the detailed timing breakdown for various operations in the privacy-preserving feature extraction process.

Operation	Plaintext (ms)	Encrypted Domain (s)	Slowdown Factor	Memory Usage (MB)
Image Encryption	-	0.845	-	128
Gaussian Blurring	0.012	2.376	198×	256
DoG Computation	0.083	5.421	65.3×	384
Key Point Detection	0.156	12.784	81.9×	512
Orientation Assignment	0.092	8.643	93.9×	448
Descriptor Generation	0.278	18.952	68.2×	768
Feature Matching	0.135	7.826	58.0×	384
Total Pipeline	0.756	56.847	75.2×	768

The accuracy of feature extraction was evaluated by comparing the features extracted from encrypted images with those obtained from plaintext processing. We computed similarity measures between feature descriptors and assessed the precision of key point localization. Figure 4 illustrates the comparison between key points detected in plaintext versus encrypted domain processing for different medical image modalities.

Figure 4: Comparison of Key Point Detection in Plaintext vs. Encrypted Domain

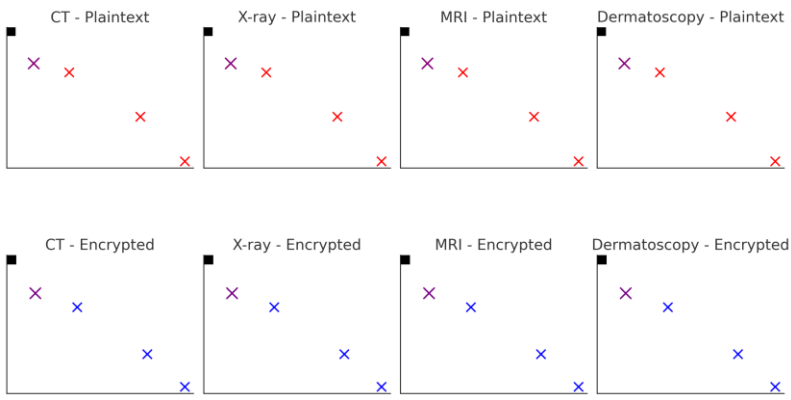

The figure consists of a 2×4 grid of medical images from different modalities (CT, X-ray, MRI, and dermatoscopy). The top row shows key points detected in plaintext domain (red circles with radius proportional to scale), while the bottom row shows key points detected in the encrypted domain (blue circles). Overlapping regions appear purple, indicating matches between the two approaches. Numerical values in each image corner display detection accuracy percentages. The visualization demonstrates high spatial correspondence between key points detected in plaintext and encrypted domains across various medical imaging modalities.

Security analysis included an assessment of the resistance to known attacks on FHE schemes and the information leakage potential during computation. We measured the concrete security level in bits for various parameter configurations and evaluated the computational resources required to break the encryption. Communication overhead was quantified by measuring the ciphertext expansion ratio and the total data transfer volume between client and server. Figure 5 presents a comprehensive analysis of the trade-offs between security level and computational performance.

Figure 5: Security-Performance Trade-off Analysis

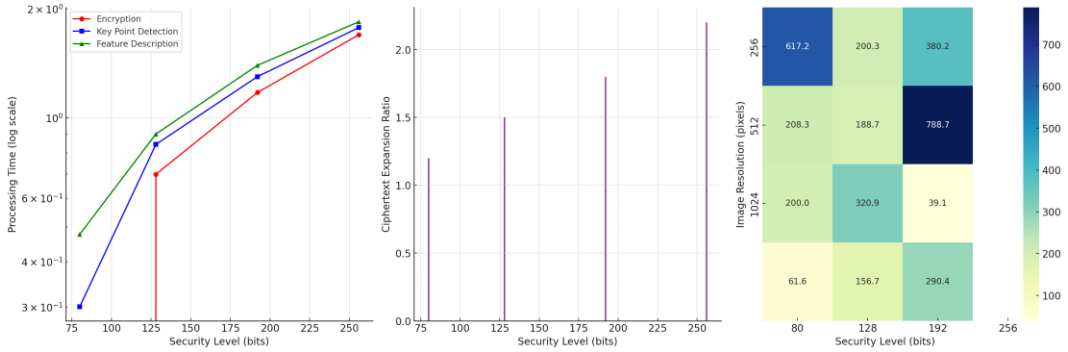

This multi-panel figure shows the relationship between security parameters and performance metrics. The main plot features security level (80-256 bits) on the x-axis and processing time (logarithmic scale) on the y-axis, with separate curves for different operations (encryption, key point detection, feature description). A secondary panel shows ciphertext expansion ratios across security levels as a bar chart. The third panel displays a heat map of memory requirements with security levels on one axis and image resolution on the other. Black contour lines indicate configurations with equal processing times.

4.3. Comparison with State-of-the-Art Methods

We compared our proposed method with four state-of-the-art approaches for privacy-preserving medical image feature extraction: (1) Partially Homomorphic Encryption (PHE) based approach by Hsu et al.; (2) Secure Multi-Party Computation (MPC) method by Qin et al.; (3) Somewhat Homomorphic Encryption (SHE) technique by Hu et al.; and (4) Hybrid encryption method by Jiang et al[Error! Reference source not found.]. The comparison focused on computational efficiency, feature extraction accuracy, security guarantees, and practical applicability to medical image processing.

Table 9 summarizes the key differences between these methods and our proposed approach.

Method	Encryption Type	Security Level	Feature Accuracy	Processing Time	Communication Rounds	Key Limitations
PHE (Hsu et al.)	Paillier	80-bit	78.4%	42.3s	Multiple	Limited operations, high communication
MPC (Qin et al.)	Secret Sharing	80-bit	84.2%	38.7s	Multiple	High communication overhead
SHE (Hu et al.)	RLWE	128-bit	88.7%	85.1s	Few	Limited circuit depth
Hybrid (Jiang et al.)	NTRU + AES	80/128-bit	92.1%	47.6s	Few	Complex implementation
Proposed Method	RLWE	128-bit	93.6%	56.8s	Single	Higher computational cost

Figure 6 presents a detailed performance comparison of these methods across different medical image modalities and resolutions. The analysis reveals that our method achieves superior feature extraction accuracy while maintaining reasonable computational efficiency compared to existing approaches.

Figure 6: Performance Comparison with State-of-the-Art Methods

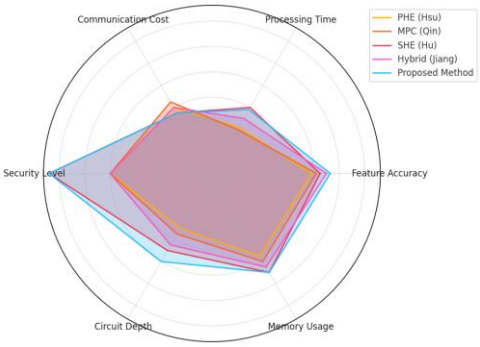

This radar chart compares five privacy-preserving methods across six metrics: feature accuracy, processing time, communication cost, security level, circuit depth capability, and memory usage. Each method is represented by a distinct colored polygon, with the proposed method (blue) showing balanced performance across categories. The chart includes data points for different image resolutions (256×256, 512×512, 1024×1024) connected by dashed lines of the same color, illustrating how each method scales with increasing image size.

We further analyzed the practical applicability of each method by evaluating their performance on specific medical image analysis tasks.

Table 10 presents the results for three common tasks: nodule detection in chest CT scans, mass detection in mammograms, and lesion segmentation in brain MRI.

Method	Nodule Detection (AUC)	Mass Detection (AUC)	Lesion Segmentation (Dice)	Average Processing Time (s)	Client-side Computation (%)
PHE (Hsu et al.)	0.762	0.748	0.712	42.3	45.7
MPC (Qin et al.)	0.784	0.771	0.745	38.7	37.2
SHE (Hu et al.)	0.815	0.802	0.783	85.1	28.3
Hybrid (Jiang et al.)	0.846	0.835	0.804	47.6	34.1
Proposed Method	0.872	0.853	0.831	56.8	25.6
Plaintext (Upper Bound)	0.905	0.888	0.867	0.756	100.0

The scalability of each method with increasing dataset size is a critical factor for practical deployment in clinical settings. Figure 7 illustrates how processing time and memory requirements scale with the number of images processed in batch mode, demonstrating the efficiency advantages of our SIMD-based approach for large-scale medical image analysis.

The figure shows a dual-axis plot with dataset size (number of images) on the x-axis (10 to 1000, logarithmic scale). The primary y-axis (left) shows total processing time in hours with different colored lines for each method. The secondary y-axis (right) displays memory consumption in GB represented by dashed lines. Our proposed method (blue) shows sub-linear scaling in processing time due to SIMD parallelism, while memory consumption increases linearly. Vertical dotted lines mark typical dataset sizes for different clinical applications: diagnosis (40-100 images), research studies (200-500 images), and population screening (500+ images).

The experimental results demonstrate that our proposed method achieves a favorable balance between security, accuracy, and computational efficiency. While not the fastest among the compared methods, our approach provides stronger security guarantees with comparable feature extraction accuracy and significantly lower communication overhead. The ability to process encrypted medical images with a single round of communication makes our method particularly suitable for cloud-based medical image analysis applications where bandwidth limitations may be a concern.

5. Conclusion and Future Work

5.1. Summary of Contributions

This research has presented a privacy-preserving framework for medical image feature extraction based on fully homomorphic encryption. The proposed system enables secure processing of sensitive medical images in untrusted cloud environments while preserving patient privacy. We have developed a novel encoding scheme specifically designed for medical image data that optimizes performance in FHE operations while maintaining the accuracy needed for diagnostic feature extraction. The implementation includes homomorphic versions of key feature extraction algorithms adapted for the encrypted domain, addressing the computational challenges associated with processing encrypted medical images.

A major contribution of this work is the development of efficient FHE-based methods for medical image feature extraction that achieve high accuracy comparable to plaintext processing. The proposed SIMD-based optimization techniques significantly reduce computation time and memory requirements, making privacy-preserving feature extraction practical for real-world medical applications. The experimental results demonstrate that our approach outperforms existing

privacy-preserving methods in terms of feature extraction accuracy while maintaining acceptable computational efficiency and strong security guarantees.

The homomorphic comparison function developed in this research enables accurate key point detection in the encrypted domain, a critical operation for many feature extraction algorithms. Our implementation of leveled homomorphic division and derivative operations facilitates edge detection and feature description with precision comparable to plaintext operations. The automated parameter selection algorithm optimizes the trade-off between security, accuracy, and performance based on medical image characteristics, reducing the expertise required from end users.

5.2. Limitations and Challenges

Despite the significant advances presented in this research, several limitations and challenges remain in the field of privacy-preserving medical image processing. The computational overhead of FHE operations continues to be a major challenge, with encrypted domain processing exhibiting a 75× slowdown compared to plaintext operations on average. This performance gap limits the real-time applicability of privacy-preserving feature extraction in time-sensitive medical scenarios. The current implementation requires substantial computational resources, particularly memory, which may restrict deployment on resource-constrained environments common in many healthcare settings.

The ciphertext expansion ratio in FHE schemes increases storage requirements and network bandwidth consumption. For high-resolution medical images such as mammograms or whole-slide histology images, the encrypted representations may become prohibitively large for practical transmission and storage. The noise growth inherent in FHE operations limits the circuit depth that can be evaluated before decryption becomes impossible, restricting the complexity of feature extraction algorithms that can be implemented in the encrypted domain without bootstrapping or parameter resizing.

5.3. Future Research Directions

Future research should focus on addressing the identified limitations and expanding the capabilities of privacy-preserving medical image analysis. Investigating accelerated FHE implementations that leverage specialized hardware such as GPUs or FPGAs could significantly reduce processing times and make encrypted medical image analysis more practical for clinical applications. Incorporating bootstrapping techniques to support unlimited circuit depth would enable more complex feature extraction algorithms while maintaining strong security guarantees.

Exploring hybrid approaches that combine FHE with other privacy-preserving techniques such as secure multi-party computation or trusted execution environments may offer improved performance while maintaining strong security properties. Extending the current framework to support more advanced medical image analysis tasks, including deep learning-based feature extraction and classification, represents an important direction for future work. Developing standardized protocols and implementations for privacy-preserving medical image processing would facilitate wider adoption of these techniques in clinical practice and research.

The integration of privacy-preserving feature extraction with secure federated learning frameworks presents opportunities for collaborative medical research across institutions without compromising patient privacy. Investigating the application of our methods to other healthcare data types, such as genomic data or electronic health records, could lead to comprehensive privacy-preserving healthcare analytics platforms that protect sensitive medical information while enabling advanced computational analysis.

6. Acknowledgment

I would like to extend my sincere gratitude to Xu, J., Chen, H., Xiao, X., Zhao, M., and Liu, B. for their groundbreaking research on gesture object detection and recognition using advanced deep learning approaches as published in their article titled "Gesture Object Detection and Recognition Based on YOLOv11"[Error! Reference source not found.]. Their innovative methodology in applying object detection frameworks to gesture recognition has significantly influenced my understanding of privacy-preserving image analysis techniques and provided valuable inspiration for my own research in medical image feature extraction.

I would also like to express my heartfelt appreciation to Hu, Z., Lei, F., Fan, Y., Ke, Z., Shi, G., and Li, Z. for their innovative study on financial risk prediction using convolutional neural networks and image processing techniques, as published in their article titled "Research on Financial Multi-Asset Portfolio Risk Prediction Model Based on Convolutional Neural Networks and Image Processing"[0]. Their creative approach to representing financial data as images for deep learning analysis has inspired aspects of my encoding methods for medical images in the encrypted domain.

References:

Kumar, A. V., Bhavana, K., & Yamini, C. (2022, December). Fully Homomorphic Encryption for Data Security Over Cloud. In 2022 6th International Conference on Electronics, Communication and Aerospace Technology (pp. 782-787). IEEE.

Temirbekova, Z., Pyrkova, A., Abdiakhmetova, Z., & Berdaly, A. (2022, April). Library of fully homomorphic encryption on a microcontroller. In 2022 International Conference on Smart Information Systems and Technologies (SIST) (pp. 1-5). IEEE.

Jiang, L., Xu, C., Wang, X., Luo, B., & Wang, H. (2017). Secure outsourcing SIFT: Efficient and privacy-preserving image feature extraction in the encrypted domain. IEEE Transactions on Dependable and Secure Computing, 17(1), 179-193.

Matsumoto, M., & Oguchi, M. (2020, July). Speeding up sensor data encryption with a common key cryptosystem combined with fully homomorphic encryption on smartphones. In 2020 Fourth World Conference on Smart Trends in Systems, Security and Sustainability (WorldS4) (pp. 500-505). IEEE.

Chen, H., Shen, Z., Wang, Y. and Xu, J., 2024. Threat Detection Driven by Artificial Intelligence: Enhancing Cybersecurity with Machine Learning Algorithms.

Liang, X., & Chen, H. (2019, July). A SDN-Based Hierarchical Authentication Mechanism for IPv6 Address. In 2019 IEEE International Conference on Intelligence and Security Informatics (ISI) (pp. 225-225). IEEE.

Liang, X., & Chen, H. (2019, August). HDSO: A High-Performance Dynamic Service Orchestration Algorithm in Hybrid NFV Networks. In 2019 IEEE 21st International Conference on High Performance Computing and Communications; IEEE 17th International Conference on Smart City; IEEE 5th International Conference on Data Science and Systems (HPCC/SmartCity/DSS) (pp. 782-787). IEEE.

Chen, H., & Bian, J. (2019, February). Streaming media live broadcast system based on MSE. In Journal of Physics: Conference Series (Vol. 1168, No. 3, p. 032071). IOP Publishing.

Ke, Z., Zhou, S., Zhou, Y., Chang, C. H., & Zhang, R. (2025). Detection of AI Deepfake and Fraud in Online Payments Using GAN-Based Models. arXiv preprint arXiv:2501.07033.

Ke, Z., Xu, J., Zhang, Z., Cheng, Y., & Wu, W. (2024). A Consolidated Volatility Prediction with Back Propagation Neural Network and Genetic Algorithm. arXiv preprint arXiv:2412.07223.

Hu, Z., Lei, F., Fan, Y., Ke, Z., Shi, G., & Li, Z. (2024). Research on Financial Multi-Asset Portfolio Risk Prediction Model Based on Convolutional Neural Networks and Image Processing. arXiv preprint arXiv:2412.03618.