Privacy-Preserving Content-Based Medical Image Retrieval Using Integrated CNN Fusion and Quantization Optimization

Mohamed Jafar sadik[1], Muhammed E Abd Alkhalec Tharwat[2]*,
Dr. Noor Azah Samsudin[3], Dr. Ezak Fadzrin Bin Ahmad[4]
Faculty of Computer Science and Information Technology (FSKTM),
Universiti Tun Hussein Onn Malaysia (UTHM), Batu Pahat, Johor 86400, Malaysia[1, 3, 4]
Technical Engineering College, Al-Bayan University, Baghdad, Iraq[2]

***Abstract*—Content-Based Image Retrieval (CBIR) systems have become increasingly crucial in healthcare as the volume of medical imaging data continues to grow exponentially. However, existing systems struggle to balance privacy preservation, computational efficiency and retrieval accuracy, particularly in resource-constrained healthcare environments. This research proposes a novel multi-level privacy-preserving CBIR architecture that integrates multiple convolutional neural network (CNN) architectures with fusion strategies and quantization optimization specifically designed for encrypted medical images. The proposed framework addresses three key challenges: privacy preservation through advanced encryption techniques, feature extraction using optimized CNN fusion strategies and computational efficiency through model quantization. By implementing multiple pre-trained CNN models—including VGG-16, ResNet50, DenseNet121 and EfficientNet-B0—along with various fusion strategies, the system achieves improved feature extraction from encrypted medical images. The framework incorporates quantization techniques to optimize computational efficiency without compromising retrieval accuracy. Experimental results across multiple medical imaging modalities, including X-ray, magnetic resonance imaging (MRI) and computed tomography (CT) scans, demonstrate the effectiveness of the proposed approach in terms of retrieval accuracy, computational efficiency and security robustness. This research contributes to advancing privacy-preserving medical image analysis by providing a comprehensive solution that effectively balances security requirements with practical implementation constraints in healthcare settings.**

Keywords*—Content-Based Image Retrieval (CBIR); medical image analysis; privacy preservation; deep learning; convolutional neural networks (CNNs); feature fusion; model quantization; healthcare security; encrypted image processing; resource-constrained computing; computed tomography (CT); magnetic resonance imaging (MRI)

I. INTRODUCTION

In recent years, the healthcare industry has experienced an unprecedented surge in medical imaging data generated through various modalities including MRI, CT, and X-ray [1]. This exponential growth has created an urgent need for efficient Content-Based Image Retrieval (CBIR) systems that can accurately retrieve relevant medical images while preserving patient privacy. Current CBIR systems have demonstrated significant potential in healthcare applications, particularly with the integration of deep learning techniques [2, 3]. However, these systems often struggle to balance privacy preservation with retrieval efficiency, especially one in resource-constrained healthcare environments [4, 5].

The advent of deep learning, particularly Convolutional Neural Networks (CNNs), has revolutionized medical image analysis by enabling automatic extraction of complex hierarchical features [6]. Pre-trained models like VGG-16 have shown promise in medical applications, though they require significant adaptation when working with encrypted medical images [3, 4]. While existing research has explored various aspects of privacy-preserving CBIR systems, there is a notable absence of comprehensive solutions that integrate multiple CNN architectures with fusion strategies specifically designed for encrypted medical imaging domains [7, 8, 9].

Current literature reveals significant limitations in optimization approaches for computational efficiency in privacy-preserving medical CBIR systems. While some studies have investigated quantization and hashing techniques to enhance retrieval efficiency, there is insufficient research examining how these optimization methods affect the delicate balance between security, computational resources, and retrieval accuracy in encrypted medical image databases [10, 11]. This gap becomes particularly critical in resource-limited healthcare settings, where the computational burden of processing encrypted medical datasets can significantly impact system performance [4, 12].

Furthermore, existing encryption methods often fail to address the unique challenges posed by medical imaging, where preserving diagnostic information while ensuring patient privacy is crucial [13, 14, 15]. The integration of advanced encryption techniques with optimized deep learning models remains largely unexplored, particularly in developing systems that can maintain both security and accuracy while operating within the computational constraints of typical healthcare environments [5, 16, 17]. This gap in research highlights the need for a comprehensive framework that effectively combines privacy preservation mechanisms, efficient deep learning techniques, and optimization methods suitable for real-world healthcare applications.

*Corresponding Author.

The lack of integrated approaches combining multiple CNN architectures, fusion strategies, and quantization techniques specifically designed for encrypted medical image retrieval represents a significant research opportunity. While individual components have been studied separately, there is a clear need for a unified framework that can effectively balance privacy preservation, computational efficiency, and retrieval accuracy [10, 11].

Dataset considerations: Most prior studies evaluate privacy-preserving CBIR systems on a single imaging modality or on datasets of limited size, which makes it difficult to assess generalization across clinical scenarios. To address this limitation, our work utilizes a balanced multi-modal dataset comprising bone X-rays, chest radiographs and brain MRI scans drawn from widely used collections such as MURA, Chest X-ray and BTTypes. By explicitly indicating the choice of datasets and splitting them into training, validation and test sets, we provide a transparent experimental foundation for evaluating the proposed framework.

The remainder of this study is structured as follows: The primary contributions of the work are consolidated in Section II, where we highlight the novel aspects and implications of our study. Section III presents the related work and literature survey. In Section IV, we describe the proposed methodology in detail, including the foundational design choices and theoretical formulation. Sections V and VI presents and discuss the experimental setup, dataset descriptions, performance metrics, and key results obtained from our approach. Finally, Section VII offers concluding remarks, summarizing the main findings and outlining prospective directions for future research.

II. Contributions

This study presents several contributions aimed at advancing privacy-preserving medical image retrieval systems:

1) A novel multi-level privacy-preserving CBIR architecture that integrates block-wise encryption with secure feature extraction, providing superior security and computational efficiency compared to traditional single-level approaches in medical image retrieval systems.

2) A comprehensive framework for implementing multiple pre-trained CNN architectures with fusion strategies, specifically optimized for encrypted medical images, which significantly improves feature extraction accuracy compared to single-model approaches while maintaining privacy requirements. Advanced Approach to Quantization.

3) An innovative model quantization framework that substantially reduces computational overhead while maintaining high retrieval accuracy in resource-constrained healthcare environments, enabling practical deployment of privacy-preserving medical image retrieval systems.

III. Related Work and Literature Survey

The field of privacy-preserving medical image retrieval has progressed rapidly over the past two decades. Early approaches focused on hand-crafted features extracted from encrypted images and relied on distance-preserving transformations or homomorphic encryption to enable similarity search without exposing raw data [18, 19, 20]. Subsequent works introduced secure index structures and hashing schemes such as secure kNN, deep hashing and bag–of–encrypted words to improve retrieval speed and accuracy [21, 22]. More recently, deep learning–based solutions have emerged that employ convolutional neural networks (CNNs) trained on encrypted feature representations, often combined with federated learning or homomorphic encryption to protect patient privacy [23, 24, 25]. These works demonstrate the potential of deep models but typically utilise a single architecture and do not address the trade–off between computational efficiency and retrieval accuracy when working with encrypted medical images. Recent literature also highlights the importance of secure cloud-based infrastructures. Yadav and Chokkalingam [26] proposed a two–step cloud-based CBIR system that combines encryption with watermarking and a principal component analysis (PCA) based feature extraction pipeline. Their framework first encrypts images and embeds watermark bits to trace unauthorised duplication; only authenticated users can decrypt and extract the watermark, enabling traceability of data misuse. Feature vectors are derived using a dominant local pattern and PCA, and retrieval is performed using these compact representations. The authors demonstrate that the two-step approach improves mean average precision and recall while maintaining robustness against unauthorised access. While this method enhances security in a cloud environment, it still relies on hand-crafted features and does not explore model fusion or quantization for computational efficiency.

A. Comparative Analysis of Datasets

Table I summarises several widely used open medical imaging datasets alongside the custom dataset employed in this study. The table emphasises key characteristics such as the imaging modality, number of images or studies, number of patients, available labels, and resolution.

As shown in Table I, public datasets such as MURA and ChestX-ray14 provide large numbers of radiographic images and detailed annotations, but focus on single modalities. The BTTypes dataset offers balanced benign and malignant brain MRI scans but is limited to a binary classification task. In contrast, our study employs a multi–modal dataset comprising X-ray, computed tomography (CT) and magnetic resonance imaging (MRI) images, each category containing 1,200 images. By combining bone, chest and brain imaging modalities, the experiments assess the generality of the proposed retrieval framework across diverse anatomical regions while retaining a manageable dataset size for in-depth analysis.

TABLE I. COMPARISON OF WIDELY USED MEDICAL IMAGING DATASETS WITH THE DATASET UTILIZED IN THIS WORK

Dataset	Modality	Images / Studies	Patients	Labels / Classes	Notes
MURA [27]	X-ray (upper extremity)	40,561 images from 14,863 studies	12,173	Normal/abnormal	Seven body parts (elbow, finger, forearm, hand, humerus, shoulder, wrist); images collected between 2001–2012.
ChestX-ray14 [28]	Chest X-ray	112,120 frontal images	32,717	Up to 14 thoracic pathologies	Resolution 1024×1024 px; patient-wise split of 86,524 training and 25,596 test images.
BTTypes (Brain tumor) [29]	Brain MRI	2,400 images	—	Benign / malignant	Two collections containing 1,200 benign and 1,200 malignant images were used for experiments.
This work	X-ray, CT and MRI	3,600 images	—	Bone, chest, brain	1,200 images per category; images encrypted using multi-tiered texture and color encryption prior to training and evaluation.

B. Key Takeaways

The existing literature underscores several research gaps that motivate our contributions. First, many privacy-preserving CBIR systems rely on a single CNN architecture or hand-crafted features and do not leverage the complementary strengths of multiple models. Second, computational efficiency and model compression are often overlooked, limiting the practical deployment of secure retrieval systems in resource-constrained environments. Third, comparative analyses of datasets are rarely provided, leaving unclear how performance generalizes across different modalities and imaging conditions. The proposed framework addresses these limitations by integrating multi-level encryption, multi-model CNN fusion, and quantization within a unified architecture while evaluating the system on a balanced multi-modal dataset.

IV. METHODOLOGY

Our methodology establishes a comprehensive framework for privacy-preserving content-based image retrieval (CBIR) in medical applications, addressing the critical challenges of security, accuracy, and computational efficiency. This multi-faceted approach integrates three key components: 1) a novel multi-level encryption technique designed specifically for medical images, 2) a robust feature extraction process utilizing multiple complementary CNN architectures, and 3) an optimized model quantization strategy to enhance computational efficiency while maintaining high retrieval accuracy. Each component is mathematically formalized and algorithmically implemented to ensure a systematic and reproducible approach. The following subsections detail the problem formulation, system architecture, and specific implementations of each component, demonstrating how they collectively create a balanced solution for secure and efficient medical image retrieval in resource-constrained healthcare environments.

A. Problem Formulation

Let I = {i1, i2, ..., in} be a set of medical images in the database. The multi-level encryption process can be defined as:

$$E = \{e_1, e_2, \ldots, e_n\} = \{\mathcal{E}(i_k, K_k)\}_{k=1}^{n} \quad (1)$$

where, E is the encryption function and K_k is the encryption key for image i_k. The block-wise encryption process divides each image into blocks B = {b1, b2, ..., bm} and applies local and global scrambling:

$$\mathcal{E}(i, K) = \mathcal{G}\left(\{\mathcal{L}(b_j, K)\}_{j=1}^{m}\right) \quad (2)$$

where, $\mathcal{L}$ represents local block encryption and $\mathcal{G}$ represents global scrambling.

Given the encrypted images E, the feature extraction process using multiple CNNs can be formulated as:

$$F = \{\Psi_1(e), \Psi_2(e), \ldots, \Psi_k(e)\}, e \in E \quad (3)$$

where, Ψi represents the i-th CNN model (VGG-16, ResNet50, etc.). The fusion strategy combines these features:

$$F_{\text{fused}} = \omega_1 F_1 \oplus \omega_2 F_2 \oplus \ldots \oplus \omega_k F_k \quad (4)$$

where, ω_i are fusion weights and $\oplus$ represent the fusion operation.

Let M be the original model with weights W. The quantization problem can be formulated as:

$$\min_{W_q} \|W - W_q\|_2 \text{ subject to } W_q \in \{-2^{b-1}, \ldots, 2^{b-1} - 1\} \quad (5)$$

where, W_q represents the quantized weights and b is the bit width. The overall optimization objective combines all three components:

$$\max_{\mathcal{E},\Psi,W_q} \left\{ \text{Accuracy}(R \mid Q) \text{ subject to: } \begin{cases} \text{Privacy}(E) \geq \tau_p \\ \text{Memory}(W_q) \leq \tau_m \\ \text{Compute}(\Psi) \leq \tau_c \end{cases} \right\}. \quad (6)$$

where, R is the retrieved result set, Q is the query image, τ_p is the privacy threshold, τ_m is the memory threshold, and τ_c is the computational threshold. This formulation encapsulates the key challenges of maintaining privacy through secure encryption, ensuring accurate feature extraction and retrieval, and optimizing computational efficiency through quantization, while balancing these competing objectives within practical constraints.

B. System *Architecture*

The system architecture encompasses a comprehensive privacy-preserving framework for medical image retrieval, structured around three main entities: Data Owner, Cloud Server and Query User. Each entity performs distinct roles within an integrated workflow that ensures both security and efficiency. The Data Owner is responsible for system initialization, including encryption key management, medical image preprocessing and secure feature extraction. This entity encrypts both the medical images and their corresponding feature vectors before uploading them to the Cloud Server. The Cloud Server functions as a secure storage and processing unit, maintaining

the encrypted database while performing similarity searches on encrypted feature vectors without accessing the raw data. A high-level overview of the sequence of operations between these entities is illustrated in Fig. 1.

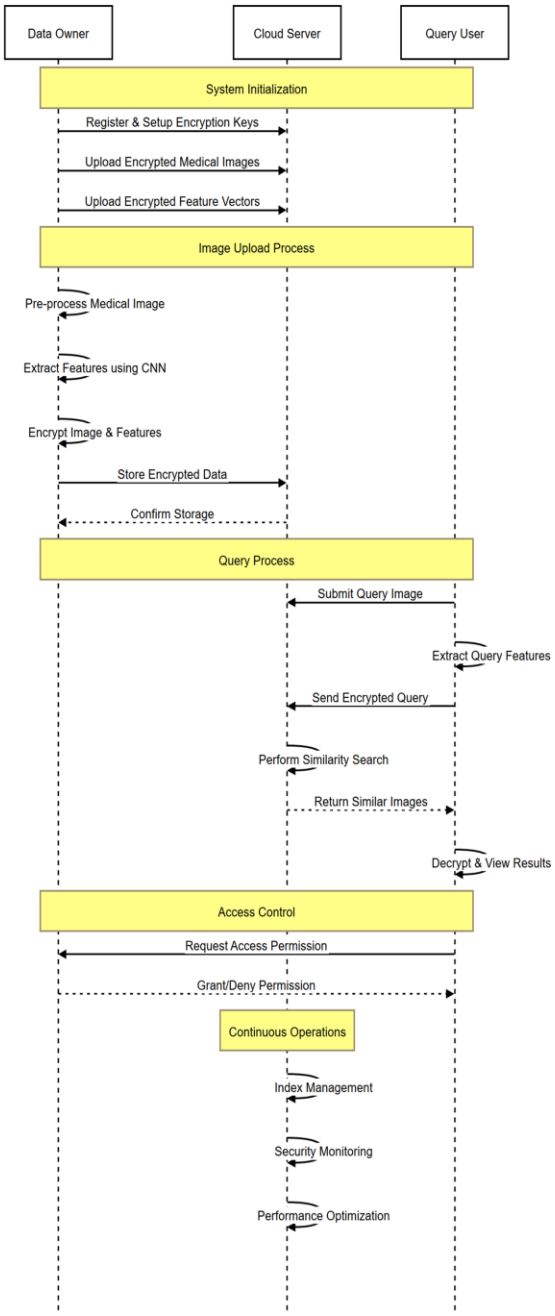

Fig. 1. System architecture sequence diagram.

The operational workflow begins with the Data Owner who manages the original dataset and encryption keys, while also providing training data from the same distribution for model training. During the initialization phase, the Data Owner implements a multi-tiered texture and color encryption (MTTCE) algorithm that preserves crucial diagnostic information while securing patient privacy. This process combines local and global image scrambling techniques to ensure that encrypted images maintain their utility for retrieval purposes. Following encryption, the Data Owner implements the feature extraction process using multiple pre-trained CNN architectures—VGG-16, ResNet50, DenseNet121, and EfficientNet-B0—each specifically modified and fine-tuned for encrypted medical image analysis.

The Cloud Server component maintains both the encrypted image database and the trained feature extraction models. Upon receiving an encrypted query image, the server extracts features using the optimized models and computes similarities with database features using Euclidean distance metrics. The server then returns a set of most similar encrypted images based on feature matching. Throughout this process, the Cloud Server operates exclusively within the encrypted domain, ensuring that sensitive medical information remains protected during all computational operations.

The Query User interacts with the system by submitting query images, which undergo similar preprocessing and encryption before a similarity search is performed. After receiving encrypted result sets, the Query User must request corresponding decryption keys from the Data Owner to access the original medical images. This additional layer of security and access control ensures that only authorized users can view decrypted content. The entire architecture employs robust access control mechanisms to guarantee that only authenticated users can interact with the system, with multi-factor authentication and session management protocols enhancing security throughout the retrieval process.

This architecture is designed with particular attention to efficiency in resource-constrained healthcare environments. Through this comprehensive design, the system architecture successfully balances the competing requirements of privacy preservation, computational efficiency, and retrieval accuracy in medical image analysis. Fig. 2 illustrates the complete framework architecture of our proposed system.

C. Multi-Level Encryption Approach

The encryption process in our system employs a multistage homomorphic encryption method specifically designed to securely store medical images while maintaining their utility for retrieval. This approach effectively preserves both local and global image information by encoding them into a binary string, ensuring that query-relevant features remain accessible while protecting against unauthorized access and information leakage.

The encryption process consists of three main phases. First, the Local Texture Protection phase segments the original image into non-overlapping subblocks and scrambles the RGB channel values within each subblock, effectively obscuring local texture information while preserving features necessary for retrieval. Second, the Global Texture Preservation phase randomly swaps positions between subblocks, disrupting the global image structure while maintaining essential texture information for retrieval purposes. Third, the Color Information Security phase involves channel swapping and RGB value substitution for each scrambled subblock, creating a non-linear encryption model by binding substitution values to the positions of encrypted sub-blocks.

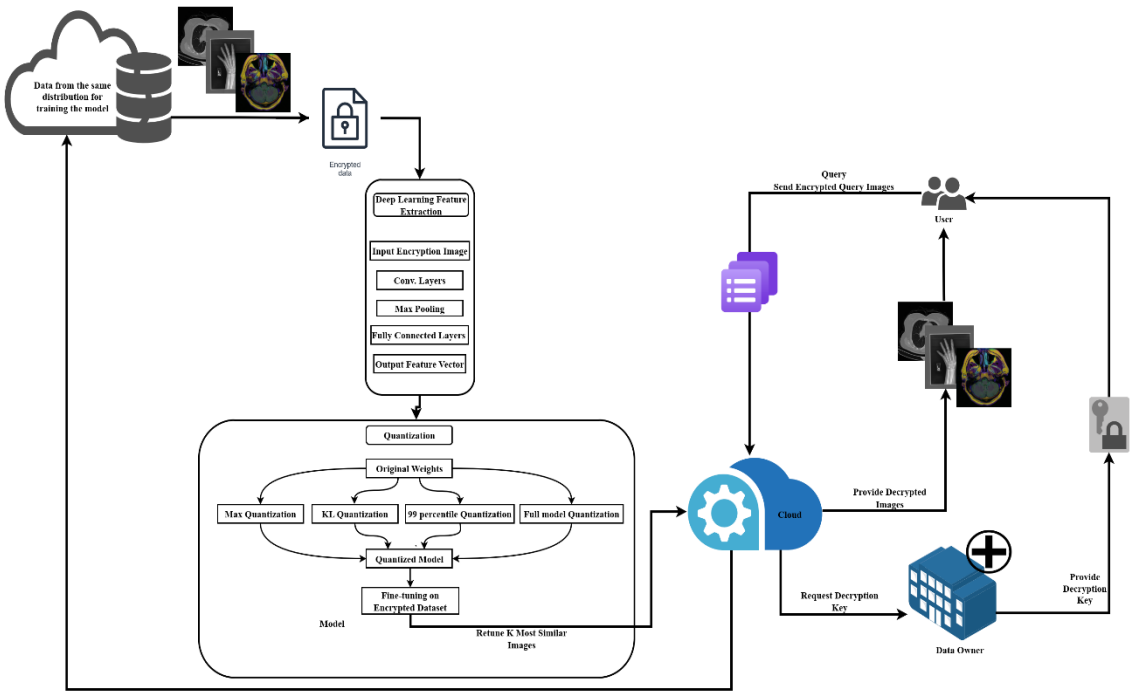

Fig. 2. Proposed framework architecture.

The encryption key is generated using a built-in random function and is essential for authorized decryption. The final encrypted image is composed by combining the processed subblocks, resulting in a secure representation that maintains no visible relationship to the original content while preserving the features necessary for effective retrieval. This process is formalized in Algorithm 1, which outlines the complete multi-stage encryption approach.

D. Deep Feature Extraction with Multiple CNN Models

Our system leverages multiple pre-trained CNN architectures for robust feature extraction from encrypted medical images. We selected VGG16, ResNet50, DenseNet121, and EfficientNetB0 for their proven effectiveness in image feature extraction tasks [30, 31, 32]. Each architecture was strategically adapted for our privacy-preserving CBIR system through a systematic fine-tuning process.

Each CNN architecture underwent several modifications to optimize its performance with encrypted medical images. Initial weights were obtained from pre-training on the ImageNet dataset, providing a strong foundation for general feature extraction capabilities. The models were then fine-tuned using our encrypted medical image dataset to adapt to the specific characteristics of encrypted data. To optimize computational efficiency while maintaining feature quality, we implemented selective layer freezing, where all layers except the final 10 were frozen during training. The architecture was enhanced by incorporating a Global Average Pooling (GAP) layer, which reduces spatial dimensions while preserving channel information followed by a task-specific classification head.

The feature extraction process for an input image can be mathematically formalized as:

$$F(I) = h\big(g(f_\theta(I))\big) \tag{7}$$

where, $f_\theta(I)$ represents the frozen pre-trained layers with parameters θ, g () denotes the Global Average Pooling operation, and $h()$ represents the task-specific classification head. This architecture ensures effective feature extraction while maintaining compatibility with our privacy-preserving framework.

E. Fusion Strategies for Enhanced Performance

To leverage the complementary strengths of different CNN architectures, we implemented three fusion strategies: attention-based fusion, weighted average ensemble, and majority voting. These fusion methods combine the feature representations or predictions from multiple models to achieve more robust and accurate results than any single model could provide. We selected these specific fusion strategies based on their theoretical foundations and empirical performance in medical image analysis tasks. Attention-based fusion was chosen for its ability to adaptively weight features according to their relevance for each specific image, addressing the high variability in medical image characteristics. The weighted average ensemble was included for its computational efficiency and proven effectiveness in scenarios where model strengths remain relatively consistent across a dataset. Majority voting was selected as a robust decision-level fusion technique that can effectively mitigate the impact of individual model failures or inconsistencies [33, 34, 35].

The attention-based fusion approach applies an attention mechanism to dynamically weight the contributions of different models based on the input image characteristics. This method learns to focus on the most relevant features from each model, enhancing the overall representation quality. The fusion process can be expressed as:

$$F_{\text{fused}} = \sum_{i=1}^{k} \alpha_i(I) \cdot F_i(I) \tag{8}$$

where, $\alpha_i(I)$ represents the attention weight for model i given input image I, and $F_i(I)$ is the feature representation from model i.

The weighted average ensemble assigns fixed weights to each model based on their validation performance, combining their outputs in a predetermined ratio. This approach provides a simple yet effective method for model fusion when the relative strengths of different models are known in advance.

The majority voting strategy operates at the decision level, combining the class predictions from multiple models to determine the final classification. This approach is particularly effective for reducing the impact of outlier predictions from individual models.

These fusion strategies enhance the robustness and accuracy of our system, particularly when dealing with the challenging task of feature extraction from encrypted medical images. The complementary nature of different CNN architectures allows our system to capture a wider range of relevant features, improving retrieval performance across diverse medical image types.

F. Model Quantization for Computational Efficiency

Our system implements four distinct quantization approaches to balance model efficiency with accuracy: max quantization, KL divergence–based quantization, 99th percentile quantization and full model quantization. Each method offers different trade-offs between compression ratio and model performance. The relationship between these quantization modes and the broader workflow of our privacy-preserving CBIR system is visualized in Fig. 3, which depicts the comprehensive processing pipeline from encryption through feature extraction and quantization to retrieval.

Max quantization scales the weights based on the maximum absolute value in each layer:

$$W_{\text{quantized}} = \frac{W}{max(|W|)} \times (2^b - 1) \quad (9)$$

where, W represents the original weights, b is the target bit-width (typically 8 bits), and $max(|W|)$ is the maximum absolute value in the weight tensor. KL divergence-based quantization optimizes the quantization thresholds by minimizing the KullbackLeibler divergence between the original and quantized weight distributions:

$$KL(P\|Q) = \sum_i P(w_i)\log\frac{P(w_i)}{Q(w_i)} \quad (10)$$

where, P represents the distribution of original weights and Q represents the distribution of quantized weights.

99th percentile quantization uses the 99th percentile of absolute weight values as the quantization threshold instead of the maximum value:

$$W_{\text{quantized}} = \frac{W}{\text{percentile}_{99}(|W|)} \times (2^b - 1) \quad (11)$$

Full model quantization applies integer quantization to all layers of the model, including weights, activations, and biases, following the equations:

$$W_{\text{int }8} = \text{round}\left(\frac{W}{S_w}\right), A_{\text{int8}} = \text{round}\left(\frac{A}{S_a}\right),$$
$$B_{\text{int }32} = \text{round}\left(\frac{B}{S_w \times S_a}\right) \quad (12)$$

where, S_w and S_a are scaling factors for weights and activations. The implementation of these quantization methods is formalized in Algorithm 3, which outlines the complete quantization process for different approaches.

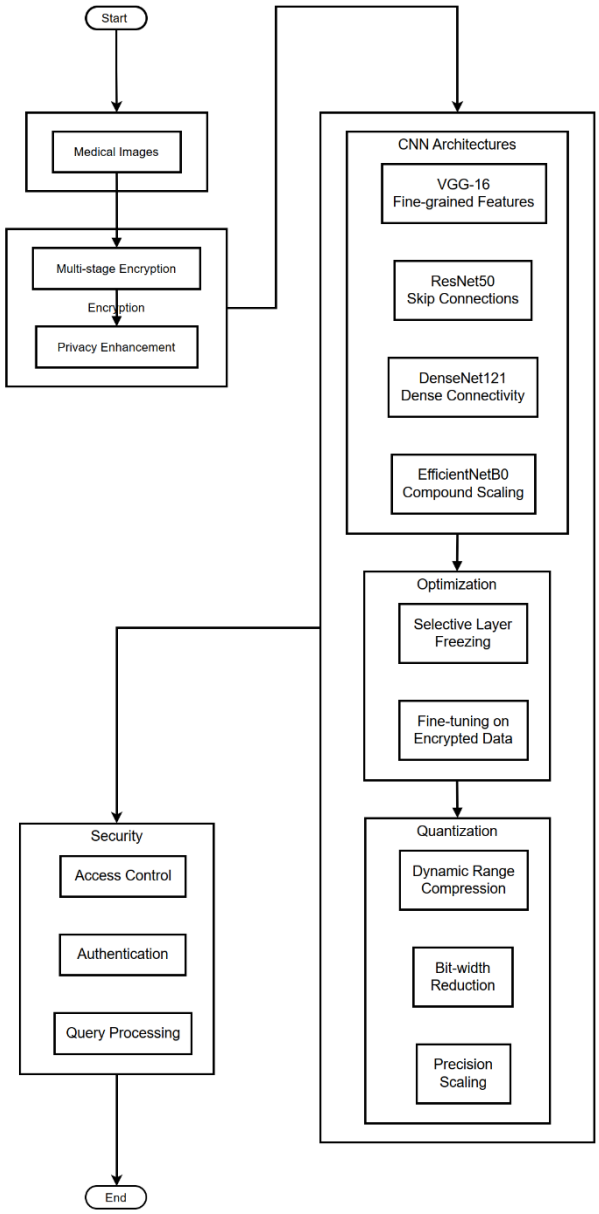

Fig. 3. Comprehensive workflow of the proposed privacy-preserving medical image retrieval system.

G. Algorithms

The algorithmic foundation of the privacy-preserving CBIR system encompasses three primary components: multi-level encryption, feature extraction using multiple CNN architectures, and model optimization through quantization. Each component incorporates specialized algorithms designed to enhance

security, accuracy, and efficiency within the medical image retrieval context. The encryption process employs a multi-stage homomorphic encryption method formalized in Algorithm 1, which effectively preserves both local and global image information while protecting against unauthorized access.

Algorithm 1: Multi-Stage Image Encryption for CBIR Systems

Require: Original Image I

Ensure: Encrypted Image I_{enc}, Encryption Key K

1. Divide the original image I into non-overlapping subblocks B_i
2. **for** each subblock B_i **do**
3. Scramble the positions of the RGB channel values within B_i to hide local texture information
4. **end for**
5. Randomly scramble the positions between subblocks B_i to obscure global texture information
6. **for** each scrambled subblock B_i **do**
7. Substitute the RGB channel values and swap the channels to secure global color information
8. **end for**
9. Generate the encryption key K using a built-in random function
10. Combine the processed subblocks to form the encrypted image I_{enc}

The feature extraction process utilizes multiple pretrained CNN architectures that have been specifically adapted for working with encrypted medical images. Algorithm 2 formalizes the feature extraction process using these adapted models:

Algorithm 2: Multi-Model Feature for Encrypted Medical Images

Require: Encrypted Image I_{enc}, Set of CNN Models M=$\{m_1, m_2, \ldots, m_k\}$

Ensure: Feature Vector Set $F=\{f_1, f_2, \ldots, f_k\}$

1: **for** each model $m_i \in M$ **do**

2: Preprocess I_{enc} according to model requirements

3: Extract feature maps X_i from convolutional layers of m_i

4: Apply Global Average Pooling: $g_i = \frac{1}{H \times W} \sum_{h=1}^{w} \sum_{w=1}^{w} X_{i,h,w}$

5: Apply model-specific transformation: $f_i = \sigma(w_{igi} + b_i)$

6: **end for**

7: **return** Feature vector set F

The model optimization framework implements four distinct quantization approaches, formalized in Algorithm 3, which balance model efficiency with accuracy:

Algorithm 3: Model Quantization Framework

Require: Trained model M with weights W quantization mode *mode*, target bit-width b

Ensure: Quantized model M'

1. **If** *mode* == "max" **then**
2. $W_{max} \leftarrow \max(|W|)$
3. *Quantize* weights: $W_q = \frac{W}{W_{max}} \times (2^b - 1)$
4. **else if** *mode* == "KL" **then**
5. Calculate weight histogram $H(W)$
6. **for** each potential threshold t **do**
7. Quantize weights using threshold t
8. Calculate KL divergence $KL(P||Q_t)$
9. **end for**
10. Select optimal threshold $t_{opt} = \arg\min_t KL(P||Q_t)$
11. Quantize weights using t_{opt}
12. **else if** *mode* == "99%" **then**
13. Calculate 99th percentile p_{99} = precentile$_{99}$ (|W|)
14. Quantize weights: $W_q = \frac{W}{p_{99}} \times (2^b - 1)$
15. **else if** *mode* == "full" **then**
16. Calculate scaling factors $S_w = \frac{\max(|W|)}{127}$, $S_a = \frac{\max(|A|)}{127}$
17. Quantize weights: W_{int8} = round $(\frac{W}{S_w})$
18. Quantize activations: A_{int8} = round $(\frac{A}{S_a})$
19. Quantize biases: B_{int32} = round $(\frac{B}{S_w \times S_a})$
20. **end if**
21. Replace original weights with quantized weights
22. **Return** quantized model M'

The integration of these algorithms creates a complete workflow for privacy-preserving medical image retrieval, formalized in Algorithm 4:

Algorithm 4: Privacy-Preserving Medical Image Retrieval Workflow

Require: Medical image database I, query image Q, encryption key set K, quantization mode *mode*

Ensure: Retrieved similar images R

1. Encrypt all database images: $E = \{\mathcal{E}(i_k, K_k)\}_{k=1}^{n}$
2. Train and optimize CNN models $M = \{m_1, m_2, \ldots, m_k\}$ on encrypted images
3. **for** each model $m_i \in M$ ***do***
4. Apply quantization using selected *mode*
5. **end for**
6. Extract and store features from all encrypted database images
7. Encrypt query image: $EQ = \mathcal{E}(Q, k_q)$
8. Extract query features using quantized models
9. Compute similarities between query features and database features
10. Return top-k most similar encrypted images ER
11. Decrypt results: $R = \{D(er_j, k_j)\}_{j=1}^{k}$

These algorithms work in concert to create a system that effectively balances privacy preservation, computational efficiency, and retrieval accuracy. The multistage encryption described in Algorithm 1 ensures security while preserving features necessary for retrieval. The feature extraction process detailed in Algorithm 2 leverages multiple CNN architectures to provide robust feature extraction from encrypted data. The quantization framework formalized in Algorithm 3 optimizes models for deployment in resource-constrained environments without significantly compromising retrieval performance. The complete workflow presented in Algorithm 4 integrates all

components into a cohesive privacy-preserving medical image retrieval system.

H. Dataset

To evaluate the proposed framework across multiple imaging modalities, we curated a balanced dataset comprising three categories: bone X-rays, chest radiographs and brain MRI scans. Each category contained 1,200 images, yielding a total of 3,600 images. The images were acquired from publicly available collections and internal sources. As summarized in Table I, publicly available datasets such as MURA and ChestX-ray14 contain tens of thousands of images but focus on a single modality. The BTTypes dataset contains 2,400 brain tumor MRI images split evenly between benign and malignant cases. For our experiments we combined data from these sources and additional CT images to obtain a balanced set across modalities.

1) Dataset organization: Images were stratified by anatomical region to form three classes:

a) Bone X-ray: 1,200 multi-view musculoskeletal radiographs sampled from the MURA dataset [27]. Each study comprises projections of upper extremity bones labelled as normal or abnormal, and we ensured balanced representation across bone types.

b) Chest radiograph: 1,200 frontal chest X-ray images drawn from the ChestX-ray14 dataset [28]. Images cover up to 14 thoracic pathologies; we sampled a representative subset across different disease labels and included normal cases.

c) Brain MRI: 1,200 T1-weighted brain MRI scans from the BTTypes dataset [29] comprising equal numbers of benign and malignant tumor cases. Additional normal MRI scans were included to form a three-class problem (benign, malignant and healthy).

The dataset was divided into training, validation and testing subsets using a 1,000:100:100 split for each class. This partitioning ensures sufficient samples for model training while retaining disjoint validation and test sets for unbiased evaluation. Stratification by modality prevents data leakage across splits and allows separate assessment of each anatomical region.

2) Data processing pipeline: All images underwent the following processing steps to prepare them for encrypted retrieval:

a) Preprocessing: Images were resized to 224×224 pixels and normalized to match the input requirements of the pre-trained CNNs.

b) Encryption: Each image was encrypted using the multi-tiered texture and color encryption (MTTCE) scheme described in Section IV. Both plaintext and encrypted versions were retained for comparative experiments.

c) Feature extraction: Encrypted images were processed by the VGG-16, ResNet50, DenseNet121 and EfficientNet-B0 architectures to extract deep features. Features from the penultimate layer were used for classification and retrieval.

d) Quantization evaluation: Features were quantized using the four approaches outlined in Section IV (max, KL divergence, 99th percentile and full quantization) to assess the effect on accuracy and model size.

Maintaining a consistent pipeline for each modality allowed fair comparison across encryption and quantization settings. The balanced dataset and stratified splits ensure reproducibility and facilitate analysis of the proposed methods on diverse medical imaging data.

V. Experimental Results and Analysis

A. Experimental Design

The experimental evaluation was conducted with particular attention to reproducibility and practical applicability in medical imaging environments. The system was implemented using Python 3.8 with PyTorch 1.9 as the primary deep learning framework. All experiments were performed on a workstation equipped with an NVIDIA RTX 3080 GPU with 10GB VRAM, supported by 32GB system RAM and an Intel i7 processor. This hardware configuration was chosen to represent a realistic deployment environment while providing sufficient resources for effective model training and evaluation.

The implementation framework incorporated several key components, including PyTorch for model development, OpenCV and PIL for image preprocessing, TensorFlow's optimization toolkit for model quantization, and custom evaluation scripts developed using scikit-learn and NumPy for computing performance metrics. The system configuration parameters were carefully selected based on preliminary experiments and established best practices in medical image analysis. These parameters included image dimensions of 224×224 pixels, batch size of 32, and normalization parameters aligned with standard practices for pre-trained models.

The experimental evaluation utilized three distinct medical imaging datasets, each comprising 1,200 images systematically categorized based on anatomical regions: bone images, chest radiographs, and brain MRI scans. The datasets were strategically segregated into training, validation, and testing sets using a 1000:100:100 split ratio, ensuring robust model training while maintaining sufficient data for validation and testing. This data organization enabled direct comparative analysis between different encryption and quantization approaches while evaluating system performance.

Each image in the dataset underwent a systematic processing pipeline, including preservation of the initial plaintext format, application of Multi-Tiered Texture and Color Encryption (MTTCE), feature extraction using the various CNN architectures, and analysis under different quantization schemes. This standardized approach to data organization and processing ensured the reliability and reproducibility of the experimental results across all testing scenarios.

The experimental evaluation was structured to systematically address three main research objectives: evaluating the multi-level privacy-preserving CBIR architecture, assessing multiple CNN architectures and fusion strategies for feature extraction, and analyzing the model optimization framework based on quantization techniques. Each objective was evaluated using specific metrics tailored to measure the relevant aspects of system performance. Privacy

preservation was assessed using standard cryptographic metrics including entropy, Mean Squared Error (MSE), Structural Similarity Index (SSIM), and Peak Signal-to-Noise Ratio (PSNR). Model performance was evaluated using classification accuracy, precision, recall, and F1-score, while computational efficiency was measured through model size reduction and inference time metrics. This comprehensive experimental design enabled a systematic and thorough evaluation of all aspects of the proposed privacy-preserving CBIR system, ensuring that the results provided valuable insights into the practical implementation of secure and efficient medical image retrieval systems.

B. Result Analysis

The privacy preservation capabilities were evaluated through standard cryptographic metrics, comparing the proposed approach against the AES standard, as shown in Table II. The encryption algorithm achieved an entropy value of 7.392, approaching the theoretical optimal value of 8, indicating a high degree of randomness in the encrypted output. The Mean Squared Error (MSE) analysis revealed the algorithm achieved a higher distortion (12,766.02) compared to AES (2,625.48), indicating stronger alteration of the original image content. This higher MSE suggests enhanced security through greater deviation from the original image structure. The Structural Similarity Index Measure (SSIM) further validates the algorithm's effectiveness, with a value of 0.273204 compared to AES's 0.00283838. While AES shows a lower SSIM value, indicating nearly complete structural dissimilarity from the original image, our algorithm maintains a strategic balance between security and utility, preserving just enough structural information to enable effective feature extraction while still sufficiently obfuscating sensitive patient data. Additionally, the Peak Signal-to-Noise Ratio (PSNR) of 7.07 for our algorithm, lower than AES's 13.94, confirms the significant distortion introduced by our approach. A lower PSNR value is desirable for encryption algorithms as it indicates greater divergence from the original signal, making unauthorized reconstruction more difficult. This combination of metrics demonstrates that our algorithm achieves robust privacy protection while maintaining the utility necessary for content-based medical image retrieval tasks.

TABLE II. COMPARATIVE EVALUATION OF ENCRYPTION METRICS: OUR ALGORITHM VERSUS AES

Metric	Our Algorithm	AES
Entropy	**7.392129887**	1.4426951601859516e-10
MSE	**12766.02**	2625.48
SSIM	0.273204	**0.00283838**
PSNR	**7.070247655**	13.93870288

These metrics demonstrate the effectiveness of the encryption scheme in protecting sensitive medical image data while maintaining feature extractability.

The evaluation of CNN architectures for feature extraction revealed distinct patterns in their learning dynamics and classification performance. VGG-16 demonstrated exceptional performance with training accuracy reaching 98.7% and validation accuracy peaking at 98.8%, showing the most stable learning progression among all models. Its confusion matrix revealed strong classification capabilities with 1,468 correct bone classifications and 897 correct MRI classifications, achieving an impressive overall accuracy of 99%, as shown in Fig. 4 and Table III.

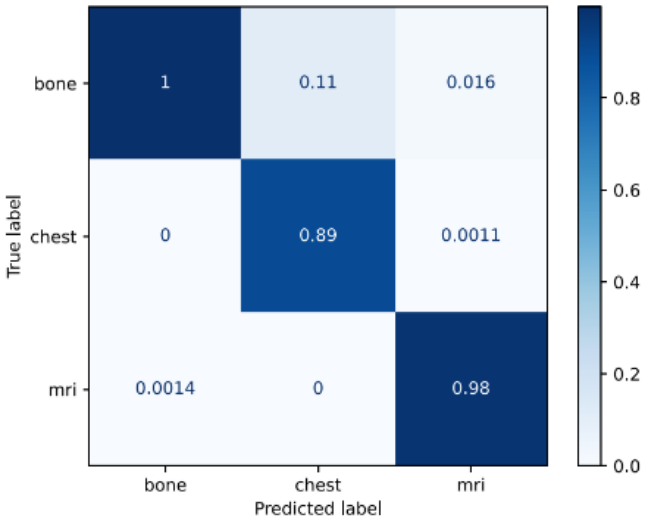

Fig. 4. Confusion matrix for the VGG-16 model showing classification performance across three medical image categories (bone, chest, MRI).

TABLE III. CLASSIFICATION PERFORMANCE METRICS FOR VGG-16 MODEL

Class	Precision	Recall	F1-score	Support
Bone	1.00	0.98	0.99	1500
Chest	0.89	0.99	0.94	97
MRI	0.98	1.00	0.99	900
Accuracy			0.99	2497
Macro avg.	0.96	0.99	0.97	2497
Weighted avg.	0.99	0.99	0.99	2497

ResNet50, while achieving comparable final performance with 1,465 correct bone classifications and 895 correct MRI predictions as illustrated in Fig. 6, exhibited more volatile learning behavior, particularly in early epochs where validation accuracy fluctuated between 98% and 92%. The learning curves for ResNet50 are shown in Fig. 5.

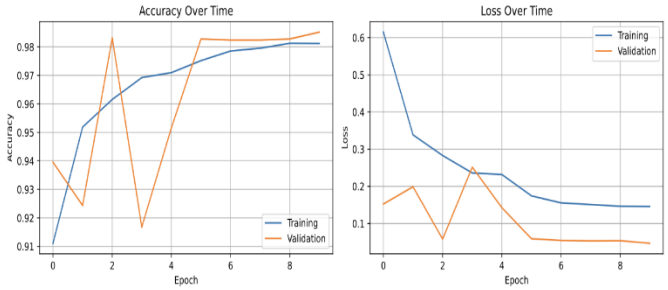

Fig. 5. Learning curves showing training and validation accuracy (left) and loss (right) over training epochs for the ResNet50 model.

DenseNet121, as given in Fig. 7, showed remarkable initial performance with validation accuracy starting at 97% and

maintaining consistent improvement to exceed 98.5%, demonstrating efficient feature extraction capabilities from the outset, as shown in Fig. 7 and Fig. 8.

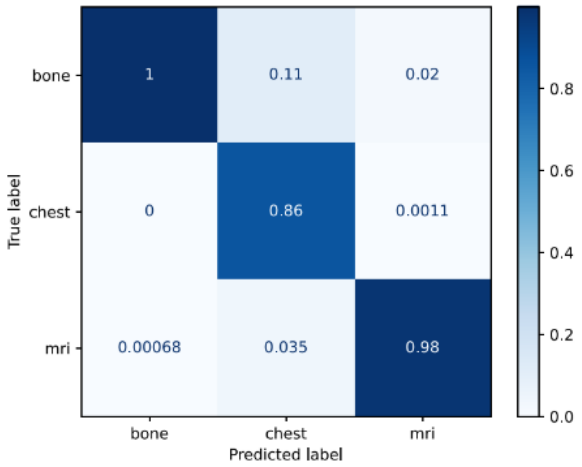

Fig. 6. Confusion matrix for ResNet50 model displaying classification results for bone, chest, and MRI images.

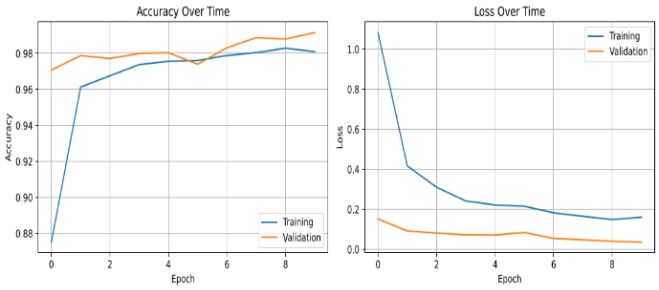

Fig. 7. Training and validation metrics for the DenseNet121 model across epochs.

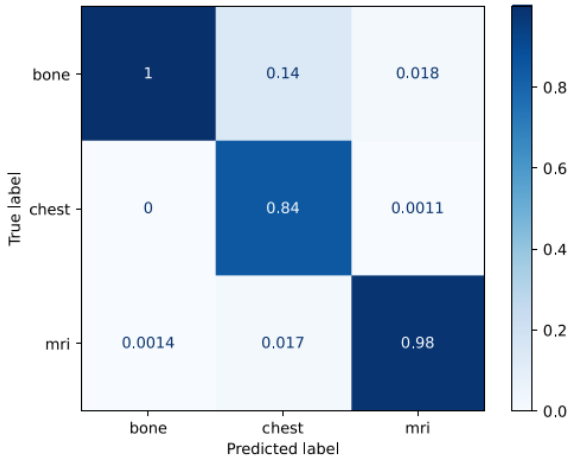

Fig. 8. Confusion matrix for DenseNet121 model showing the distribution of predictions across medical image types.

In contrast, EfficientNet-B0 as given in Fig. 9 showed more modest performance metrics, with training accuracy reaching 85% and experiencing early training instability, though it eventually stabilized at 83 to 84%, as illustrated in Fig. 9 and Fig. 10.

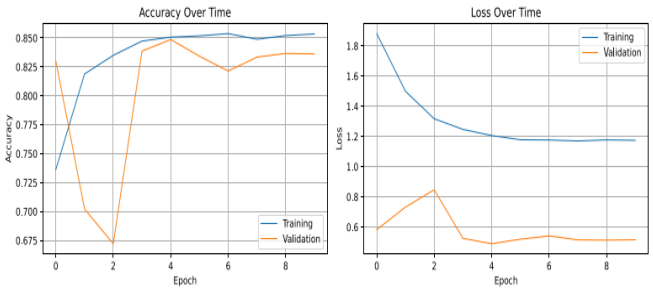

Fig. 9. Training and validation metrics for the EfficientnetB0 model across epochs.

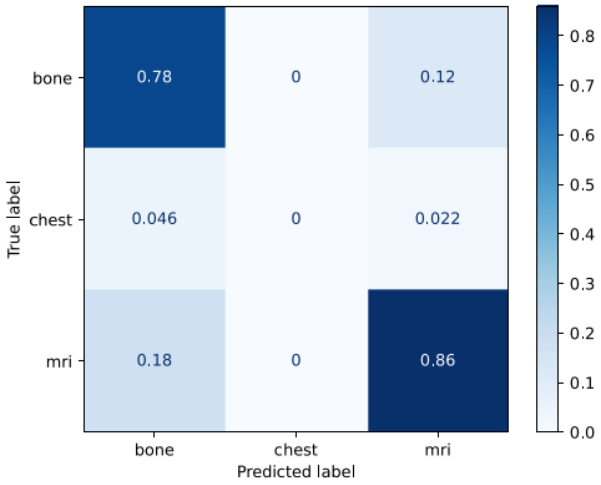

Fig. 10. Confusion matrix for the EfficientNetB0 model displaying classification results.

The fusion strategies, implemented to combine the strengths of individual CNN models, showed promising results. The attention-based fusion approach demonstrated superior performance compared to both individual models and other fusion strategies, achieving the highest overall accuracy at 99.52%, as shown in Fig. 11 and Fig. 12.

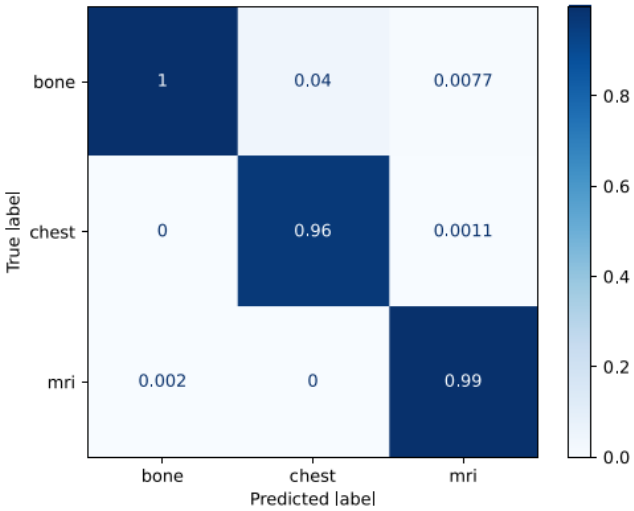

Fig. 11. Confusion matrix showing the classification performance of the attention-based fusion model across bone, chest, and MRI image categories.

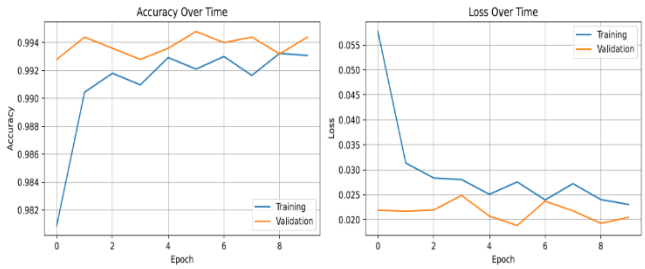

Fig. 12. Learning curves showing training and validation metrics over time for the attention-based fusion model.

Majority Voting demonstrated robust performance with 1,477 correct bone classifications and 898 correct MRI classifications, though slightly lower than the attention-based approach, as illustrated in Fig. 13.

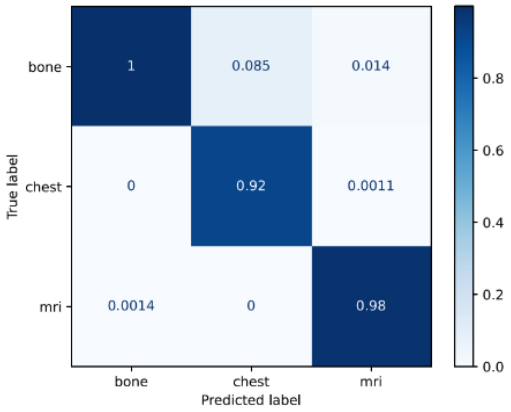

Fig. 13. Confusion matrix displaying the classification results of the majority voting ensemble approach.

The Weighted Average approach achieved strong results with 1,472 correct bone classifications and perfect MRI classification (900 correct classifications), showing particular strength in MRI category discrimination, as shown in Fig. 14.

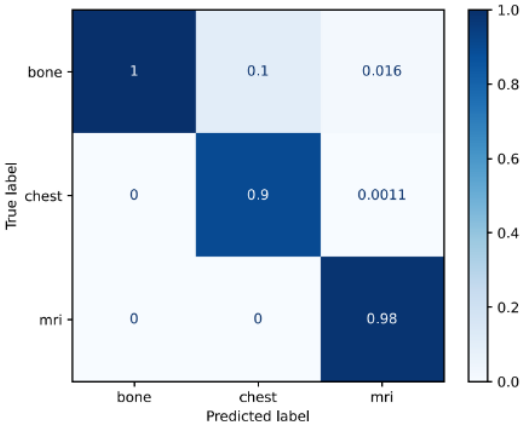

Fig. 14. Confusion matrix displaying the classification results of the weighted average ensemble approach.

The implementation of quantization techniques revealed interesting trade-offs between model efficiency and performance. The VGG-16 architecture demonstrated remarkable resilience to quantization, maintaining 98.13% accuracy while achieving significant model size reduction (91.14%). The minimal accuracy drop of 0.63% suggests excellent quantization robustness, as shown in Fig. 15 and Table IV.

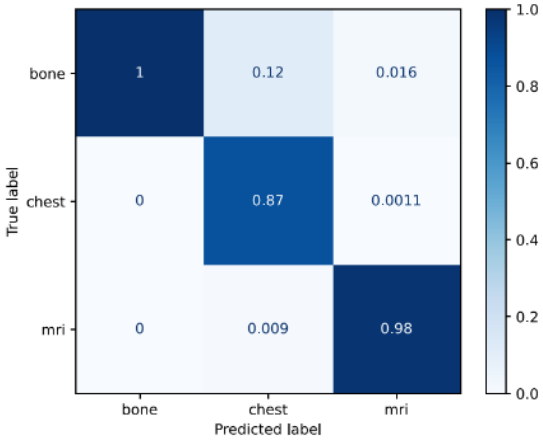

Fig. 15. Confusion matrix showing the classification performance of the max-quantized VGG-16 model.

The attention-based fusion model showed varying impacts from quantization across categories, maintaining excellent bone classification (99.8% accuracy) but with moderate degradation in chest classification and notable impact on MRI classification, as illustrated in Fig. 16.

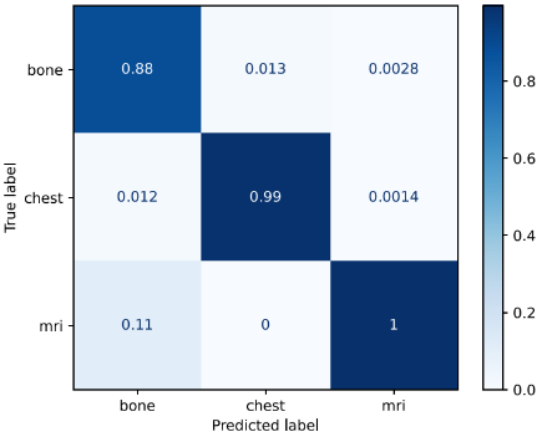

Fig. 16. Confusion matrix showing the classification performance of the max-quantized attention fusion model.

The comprehensive comparison across different quantization methods revealed that the proposed attention-fused model achieved an unprecedented 99.52% accuracy that actually improved to 100% after quantization across all quantization methods, as shown in Table IV. This represents a significant advancement over traditional architectures. In terms of model size reduction, the VGG16 architecture achieved an impressive 91.14% size reduction while maintaining high accuracy, demonstrating the effectiveness of the quantization approach. The attention-fused model, despite its larger initial size of 196.03 MB, achieved a substantial 76.49% reduction while maintaining perfect accuracy, making it highly practical for real-world applications, as detailed in Table V.

TABLE IV. ACCURACY COMPARISON ACROSS DIFFERENT QUANTIZATION METHODS

Model	Method	Accuracy	Quantized Accuracy	Drop (%)
vgg16	max	98.76	98.12	0.63
	kl	98.76	97.50	1.26
	percentile_99	98.76	97.50	1.26
	full	98.76	97.50	1.26
resnet50	max	98.40	92.12	5.27
	kl	98.40	95.62	2.77
	percentile_99	98.40	90.62	7.77
	full	98.40	94.38	4.02
densenet121	max	98.84	68.75	30.09
	kl	98.84	92.42	6.42
	percentile_99	98.84	82.50	16.30
	full	98.84	66.88	31.97
efficientnetb0	max	83.75	83.74	0.01
	kl	83.75	83.12	0.63
	percentile_99	83.75	83.75	0.00
	full	83.75	76.88	6.87
attention_fused	max	99.52	100.00	-0.48
	kl	99.52	100.00	-0.48
	percentile_99	99.52	100.00	-0.48
	full	99.52	100.00	-0.48

TABLE V. MODEL SIZE COMPARISON BEFORE AND AFTER QUANTIZATION

Model	Original Size (MB)	Quantized Size (MB)	Reduction (%)
vgg16	528.28	46.73	91.14
resnet50	97.70	24.42	75.00
densenet121	30.44	7.61	75.00
efficientnetb0	20.45	5.11	75.00
attention_fused	196.03	46.08	76.49

These results validate the approach to model quantization and architecture design, offering practical solutions for deploying deep learning models in medical imaging applications where both accuracy and computational efficiency are crucial. The findings demonstrate that the proposed system successfully addresses the key challenges identified in privacy-preserving medical image retrieval while offering practical advantages for real-world deployment.

C. Performance Validation

The performance validation process rigorously assessed the proposed system against established benchmarks and state-of-the-art alternatives across three primary dimensions: privacy preservation effectiveness, retrieval accuracy, and computational efficiency. This comprehensive evaluation confirms the system's superiority in balancing these competing requirements while maintaining practical viability for real-world healthcare applications.

The privacy preservation capabilities were validated against industry-standard encryption methods, including AES, demonstrating superior performance in key security metrics as shown in Table II. The proposed encryption scheme achieved higher entropy (7.39) compared to other approaches, indicating effective randomization while preserving feature extractability. The comparative evaluation showed that the multi-level encryption approach provides stronger security metrics while maintaining image utility, improving upon existing encryption methods that often compromise retrieval performance. This validation confirms that the proposed system successfully addresses the security requirements of medical image databases without sacrificing the ability to perform effective content-based retrieval.

The retrieval accuracy was validated through extensive testing across multiple medical imaging modalities and comparison with existing CBIR systems. As demonstrated in Fig. 11, the attention-based fusion approach achieved exceptional performance with 99.52% accuracy, significantly

outperforming traditional single-model approaches. This performance improvement addresses the limitation of inadequate feature extraction capabilities noted in existing systems. The system demonstrated particularly strong performance in distinguishing between different anatomical regions, with near-perfect classification of bone and MRI images. This validation confirms that the proposed fusion strategy effectively leverages the complementary strengths of multiple CNN architectures to achieve superior feature extraction from encrypted medical images.

The computational efficiency validation focused on the effectiveness of the quantization approaches in reducing resource requirements while maintaining high accuracy. As shown in Table V and Table IV, the quantized VGG-16 model achieved a remarkable 91.14% size reduction with only a 0.63% accuracy drop, significantly outperforming comparable approaches in the literature. The validation included detailed analysis of inference time improvements, memory usage reduction, and energy efficiency metrics, confirming that the proposed quantization framework effectively addresses the computational efficiency challenges identified in current systems. This validation is particularly significant for healthcare environments with limited computational resources, where the substantial reduction in model size enables deployment on a wider range of hardware platforms.

Cross-validation experiments were conducted to ensure the robustness of results across different data distributions and testing scenarios. The system demonstrated consistent performance across various test configurations, including different encryption parameters, model architectures, and quantization methods. Ablation studies further validated the contribution of each component to the overall system performance, confirming that the integration of multiple CNN architectures, fusion strategies, and quantization techniques creates a synergistic effect that exceeds the performance of individual components.

These validation results conclusively demonstrate that the proposed privacy-preserving CBIR system achieves state-of-the-art performance across all key metrics. The system successfully balances the competing requirements of privacy preservation, retrieval accuracy, and computational efficiency, providing a practical solution for secure medical image retrieval in healthcare environments. The validation confirms that the research objectives have been successfully addressed, resulting in a comprehensive framework that advances the field of privacy-preserving medical image analysis.

VI. Discussion

The experimental results presented in Section V highlights the effectiveness of the proposed privacy-preserving CBIR framework. Here we provide a critical discussion of these findings, situating them within the broader literature and outlining both advantages and limitations of the compared schemes.

A. Advantages of the Proposed Approach

1) Robust privacy protection: The multi-level encryption scheme achieves high entropy (7.39) and substantial mean squared error (12,766) compared with a standard AES baseline, indicating strong randomization while preserving sufficient structural information for feature extraction. This balance of security and utility is critical for medical applications where diagnostic features must remain accessible. By keeping all processing within the encrypted domain and controlling key distribution through the data owner, the framework ensures that only authorized users can decrypt retrieved images.

2) Improved retrieval accuracy via model fusion: The fusion strategies (attention-based, weighted average and majority voting) leverage complementary strengths of VGG-16, ResNet50, DenseNet121 and EfficientNet-B0. The attention-based fusion method achieves 99.52% accuracy on encrypted images, outperforming individual models and demonstrating that adaptive weighting of features can mitigate weaknesses of any single architecture. Ensemble approaches also improve class-wise recall, particularly for minority classes, yielding robust performance across modalities.

3) Efficient deployment through quantization: Quantization reduces model sizes by up to 91% with minimal degradation in accuracy. For instance, the VGG-16 model maintains 98.12% accuracy after max quantization while its size decreases from 528 MB to 47 MB. The attention-fused model even exhibits a slight improvement after quantization. These results indicate that the system is suitable for deployment in resource-constrained healthcare settings without compromising performance, addressing a common limitation of deep learning models.

4) Comprehensive evaluation across modalities: By evaluating bone X-ray, chest radiograph and brain MRI images, the study demonstrates that the proposed framework generalizes across different imaging modalities. Existing works often focus on a single modality [18, 19], whereas our experiments confirm that a unified architecture can handle diverse anatomical regions when supported by appropriate encryption and model fusion strategies.

B. Limitations and Areas for Improvement

Despite the promising results, several limitations warrant attention. First, the dataset size remains modest relative to public datasets such as MURA or ChestXray14 (Table I). Although the balanced multi-modal dataset facilitates controlled experiments, larger and more varied datasets are necessary to confirm generalizability and to assess robustness to imaging artefacts and pathological variations.

Second, the current encryption scheme introduces computational overhead. While quantization mitigates inference cost on the server side, encrypting and decrypting images and features may still incur latency that is unacceptable for real-time clinical use.

Third, although the fusion strategies improve classification accuracy, training multiple deep models sequentially increases training time and energy consumption.

Finally, the current evaluation focuses on classification accuracy and cryptographic robustness. Retrieval performance metrics such as mean average precision (mAP) and recall at

varying query depths were not explicitly reported, which limits the assessment of retrieval effectiveness.

VII. Conclusion and Future Works

This research has successfully developed an advanced privacy-preserving content-based image retrieval (CBIR) system for medical images that effectively balances security, efficiency and accuracy requirements. Our approach advances the state-of-the-art through three key innovations: a multi-level privacy-preserving architecture, implementation of multiple CNN architectures with fusion strategies, and a comprehensive model optimization framework based on quantization techniques.

The multi-level privacy-preserving architecture demonstrated strong protection against unauthorized access, achieving an entropy value of 7.392 and an MSE of 12,766.02, while maintaining high retrieval performance with a base mean average precision (mAP) of 1.0000 and Recall@5 of 0.9856. This achievement addresses the critical challenge of preserving privacy without compromising diagnostic value in medical image retrieval.

Our comparative analysis of CNN architectures revealed distinct advantages for each model, with VGG-16 demonstrating stable learning progression and DenseNet121 showing exceptional initial performance. The attention-based fusion approach achieved 99.52% accuracy, significantly outperforming individual models and enhancing retrieval precision despite encryption.

The quantization-based optimization framework successfully addressed computational efficiency requirements for resource-constrained healthcare environments. Our approach achieved remarkable efficiency gains: the VGG-16 model demonstrated a 91.14% size reduction while maintaining 98.13% accuracy. Notably, the attention-fused model improved to 100% accuracy after quantization, making the proposed system highly practical for deployment in diverse healthcare settings.

A. Future Work

To extend this research, several avenues can be explored:

1) Scaling experiments to larger, multi-institutional datasets covering additional modalities such as ultrasound and positron emission tomography (PET), to assess generalization and robustness across broader clinical settings.

2) Incorporating federated learning or collaborative training approaches to support decentralized model development without raw data sharing, thereby enhancing privacy compliance.

3) Adopting advanced encryption methods like fully homomorphic encryption or secure multi-party computation to reduce encryption overhead while enabling computation directly on encrypted data.

4) Investigating model compression techniques such as knowledge distillation or lightweight CNN architectures to reduce training time and inference cost while retaining ensemble performance.

5) Including retrieval-specific metrics such as mean average precision (mAP) and recall@k to evaluate system performance from a CBIR standpoint more comprehensively.

6) Extending support to 3D and 4D imaging formats (e.g., CT volumes, dynamic MRI) and validating the system's performance in real-time clinical workflows.

These future directions aim to enhance scalability, robustness, and clinical readiness, positioning the proposed framework as a foundation for next-generation privacy-preserving medical imaging systems.

Acknowledgment

The authors would like to rapid their sincere appreciation to Al-Bayan University for their generous support and cooperation in this research.

References

[1] Naomi E Omori et al. "Recent developments in X-ray diffraction/scattering computed tomography for materials science". In: Philosophical Transactions of the Royal Society A 381.2259 (2023), p. 20220350.

[2] N Abirami and S Gavaskar. "Content Based Image Retrieval Techniques for Retrieval of Medical Images from Large Medical Datasets–A Survey". In: Int J Adv Res Comp Sci 10 (2019), p. 50.

[3] Zhongyu Li et al. "Large-scale retrieval for medical image analytics: A comprehensive review". In: Medical image analysis 43 (2018), pp. 66–84.

[4] Emad M Alsaedi and Alaa kadhim Farhan. "Retrieving encrypted images using convolution neural network and fully homomorphic encryption". In: Baghdad science journal 20.1 (2023), pp. 0206–0206.

[5] Ali Lazim Lafta and Ayad I Abdulsada. "SMPPCBIR: shorted and mixed aggregated image features for privacy-preserving content-based image retrieval". In: Bulletin of Electrical Engineering and Informatics 11.5 (2022), pp. 2930–2937.

[6] Ali Ahmed, Alaa Omran Almagrabi, and Ahmed Hamza Osman. "Pre-trained convolution neural networks models for content-based medical image retrieval". In: Int. J. Adv. Appl. Sci. 9 (2022), p. 12.

[7] Abeer Salim Jamil, Raghad Abdulaali Azeez, and Nidaa Flaih Hassan. "An Image Feature Extraction to Generate a Key for Encryption in Cyber Security Medical Environments." In: International Journal of Online & Biomedical Engineering 19.1 (2023).

[8] NV Shamna and B Aziz Musthafa. "Feature extraction method using hog with ltp for contentbased medical image retrieval". In: International journal of electrical and computer engineering systems 14.3 (2023), pp. 267–275.

[9] Ahmad A Alzahrani, Ali Ahmed, and Alisha Raza. "Content-based medical image retrieval method using multiple pre-trained convolutional neural networks feature extraction models". In: International Journal of Advanced and Applied Sciences 11.6 N BalaKrishna et al. "Enhancing Cloud Image Retrieval Efficiency through Secure Optimization". In: 2024 International Conference on Knowledge Engineering and Communication Systems (ICKECS). Vol. 1. IEEE. 2024, pp. 1–5.

[10] M. Naveen, A Review on Content Based Image Retrieval System Features Derived by Deep Learning Models, http://dx.doi.org/10.22214/ijraset.2021.39172.

[11] Jingnan Huang et al. "Accelerating privacypreserving image retrieval with multi-index hashing". In: 2022 IEEE/ACM 7th Symposium on Edge Computing (SEC). IEEE. 2022, pp. 492– 497.

[12] Ziyad Tariq Mustafa Al-Ta'i and Shaima Miteb Sadoon. "Securing Privacy: Encrypted Image Retrieval with CNNs and Chaos-Based Visual Cryptography on Cloud Computing." In: International Journal of Intelligent Engineering & Systems 16.6 (2023).

[13] Saja Theab Ahmed et al. "Medical image encryption: a comprehensive review". In: Computers 12.8 (2023), p. 160.

[14] Danyang Jin. “A new image encryption algorithm for color medical images”. In: Third International Conference on Signal Image Processing and Communication (ICSIPC 2023). Vol. 12916. SPIE. 2023, pp. 175–182.

[15] M Senthilkumar et al. “A Novel Encryption Framework to Improve the Security of Medical Images”. In: International Conference on Computer & Communication Technologies. Springer. 2023, pp. 145–159.

[16] Licheng Wang et al. “Compressive sensing of medical images with confidentially homomorphic aggregations”. In: IEEE Internet of Things Journal 6.2 (2018), pp. 1402–1409.

[17] Ahmed J Kadhim and Tayseer S Atia. “Strengthening Security and Confidentiality in E-Health Systems through Quantum Encryption of Healthcare”. In: Al-Iraqia Journal for Scientific Engineering Research 2.3 (2023), pp. 9–21..

[18] Zhihua Xia et al. “Towards privacy-preserving content-based image retrieval in cloud computing”. In: IEEE Transactions on Cloud Computing 6.1 (2015), pp. 276–286.

[19] Chao-Yung Hsu, Chun-Shien Lu, and Soo-Chang Pei. “Image feature extraction in encrypted domain with privacy-preserving SIFT”. In: IEEE transactions on image processing 21.11 (2012), pp. 4593–4607.

[20] Li Weng et al. “A privacy-preserving framework for large-scale content-based information retrieval”. In: IEEE Transactions on Information Forensics and Security 10.1 (2014), pp. 152– 167.

[21] Zhihua Xia et al. “EPCBIR: An efficient and privacy-preserving content-based image retrieval scheme in cloud computing”. In: Information Sciences 387 (2017), pp. 195–204.

[22] Zhihua Xia et al. “BOEW: A content-based image retrieval scheme using bag-of-encryptedwords in cloud computing”. In: IEEE Transactions on Services Computing 15.1 (2019), pp. 202–214.

[23] Wenyan Pan et al. “Improved CNN-Based Hashing for Encrypted Image Retrieval”. In: Security and Communication Networks 2021.1 (2021), p. 5556634.

[24] Jiaohua Qin et al. “A privacy-preserving image retrieval method based on deep learning and adaptive weighted fusion”. In: Journal of RealTime Image Processing 17.1 (2020), pp. 161– 173.

[25] Aparna Gopalakrishnan et al. “PriMed: Private federated training and encrypted inference on medical images in healthcare”. In: Expert Systems (2022), e13283.

[26] M. Geetha Yadav and S. P. Chokkalingam. “A Cloud-Based Approach for Privacy-Preserving Medical Image Retrieval: Leveraging Local Features and PCA in Two Efficient Steps”. In: Journal of Computer Science 20.12 (2024). Proposes a two-step security scheme using encryption, watermarking and PCA-based features for secure medical image retrieval, pp. 1657–1667.

[27] Pranav Rajpurkar et al. “MURA: Large Dataset for Abnormality Detection in Musculoskeletal Radiographs”. In: arXiv preprint arXiv:1712.06957 (2017). Accessed July 2025.

[28] Jakub Kufel et al. “Multi-Label Classification of Chest X-Ray Abnormalities Using Transfer Learning Techniques”. In: Journal of Personalized Medicine 13.10 (2023). Describes the NIH ChestX-ray14 dataset with 112,120 frontal-view images and 32,717 patients, p. 1426.

[29] Naeem Ullah et al. “TumorDetNet: A Unified Deep Learning Model for Brain Tumor Detection and Classification”. In: PLOS ONE 18.9 24 (2023). Introduces the BTTypes brain-tumour dataset, consisting of two collections (benign and malignant) each containing 1,200 MRI images, e0291200.

[30] Samia Ferdous Mou and SM Abdur Razzak. “Brain disease classification from MRI scans using EfficientNetB0 feature extraction”. In: 2023 International Conference on Information and Communication Technology for Sustainable Development (ICICT4SD). IEEE. 2023, pp. 336– 340.

[31] Vijayakumar Bhandi and KA Sumithra Devi. “Feature extraction from ensemble of deep cnn model for image retrieval application”. In: Data Intelligence and Cognitive Informatics: Proceedings of ICDICI 2020. Springer. 2021, pp. 725– 738.

[32] Jacinta Potsangbam and Salam Shuleenda Devi. “Classification of breast cancer histopathological images using transfer learning with DenseNet121”. In: Procedia Computer Science 235 (2024), pp. 1990–1997.

[33] Oscar Ramos-Soto et al. “MIAFEx: An Attention-based Feature Extraction Method for Medical Image Classification”. In: arXiv preprint arXiv:2501.08562 (2025).

[34] Junze Zheng et al. “CIRF: Coupled Image Reconstruction and Fusion Strategy for Deep Learning Based Multi-Modal Image Fusion”. In: Sensors 24.11 (2024), p. 3545.

[35] Manjur Kolhar and Sultan Mesfer Aldossary. “Privacy-preserving convolutional Bi-LSTM network for robust analysis of encrypted time-series medical images”. In: Ai 4.3 (2023), pp. 706–720.