

GIIM: Graph-based Learning of Inter- and Intra-view Dependencies for Multi-view Medical Image Diagnosis

Tran Bao Sam, Hung Vu, Dao Trung Kien, Tran Dat Dang, Van Ha Tang, Steven Truong

NVIDIA
Ho Chi Minh City, Vietnam
sbtran@nvidia.com

Abstract

Computer-aided diagnosis (CADx) has become vital in medical imaging, but automated systems often struggle to replicate the nuanced process of clinical interpretation. Expert diagnosis requires a comprehensive analysis of how abnormalities relate to each other across various views and time points, but current multi-view CADx methods frequently overlook these complex dependencies. Specifically, they fail to model the crucial relationships within a single view and the dynamic changes lesions exhibit across different views. This limitation, combined with the common challenge of incomplete data, greatly reduces their predictive reliability. To address these gaps, we reframe the diagnostic task as one of relationship modeling and propose GIIM, a novel graph-based approach. Our framework is uniquely designed to simultaneously capture both critical intra-view dependencies between abnormalities and inter-view dynamics. Furthermore, it ensures diagnostic robustness by incorporating specific techniques to effectively handle missing data, a common clinical issue. We demonstrate the generality of this approach through extensive evaluations on diverse imaging modalities, including CT, MRI, and mammography. The results confirm that our GIIM model significantly enhances diagnostic accuracy and robustness over existing methods, establishing a more effective framework for future CADx systems.

Code — <https://github.com/dtrgkien/giim>

Introduction

Medical image classification is a fundamental component of medical image analysis, crucial for supporting accurate diagnoses and effective treatment planning across a variety of imaging modalities like computed tomography (CT), magnetic resonance imaging (MRI), X-ray and mammography. Initially, the field relied on traditional classification methods using predefined and handcrafted feature extraction. While useful, these techniques were often limited in their flexibility and adaptability. An advanced example of this feature-based approach is radiomics (Lambin et al. 2012), which extracts detailed statistical and gray-level information from regions of interest (ROIs) and has shown promise when paired with machine learning predictors (Wu et al. 2019), (Liu et al. 2021), (Cameron et al. 2022), (Tang et al.

2023). A significant transformation came with the arrival of deep learning techniques, especially convolutional neural networks (CNNs). This paradigm shift introduced automatic feature extraction, leading to substantial performance improvements in a wide range of diagnostic tasks. CNNs have demonstrated exceptional efficacy in classifying complex patterns, such as identifying interstitial lung disease in chest CT scans (Marios et al. 2016), (Mikhael et al. 2023) or distinguishing between different types of brain tumors (Abiwinanda et al. 2019), (Ayadi et al. 2021), (Ullah et al. 2024) and early detecting abnormalities in mammography (Luo et al. 2024), (Gudhe et al. 2024).

Alongside CNNs, transformer-based architectures like attention modules and Vision Transformers (ViTs) have shown strong potential in medical imaging (Weibin et al. 2022), (Manigrasso et al. 2025). For example, (Weibin et al. 2022) introduced a model with intra-phase and inter-phase attention, enabling focus on key spatial locations within a view and tracking changes across phases to capture temporal dynamics. Recently, Graph Neural Networks (GNNs) (Scarselli et al. 2008) and their heterogeneous variants (Xiao et al. 2019) have become powerful tools for modeling complex dependencies in perceptual data. GNNs excel at representing multifaceted relationships and diverse entity types, leading to improved performance. Several recent works have applied GNNs to medical diagnosis, such as SSL-MNGCN and MMGCN (Ibrahim et al. 2022), which classify mammography images by modeling texture and spatial features as graphs. GNNs have also been used to contextualize connections between different body parts in lung CT scans, enhancing the understanding of complex anatomical relationships (Sun et al. 2022).

Within the realm of medical image analysis, multi-view images have become indispensable, offering complementary perspectives that enrich the information available to radiologists. The utilization of multiphase CT volumes (Sam et al. 2024), (Tran et al. 2025) enables the monitoring of changes in organ intensity over time, providing critical insights into patient health. Similarly, multi-view mammography (TX et al. 2022) improves tumor detection by capturing images from different angles, reducing the chances of missing lesions hidden by overlying tissues. Some methods leverage these multiple views by first aligning the images and then combining them into a single, fused image (Sharif

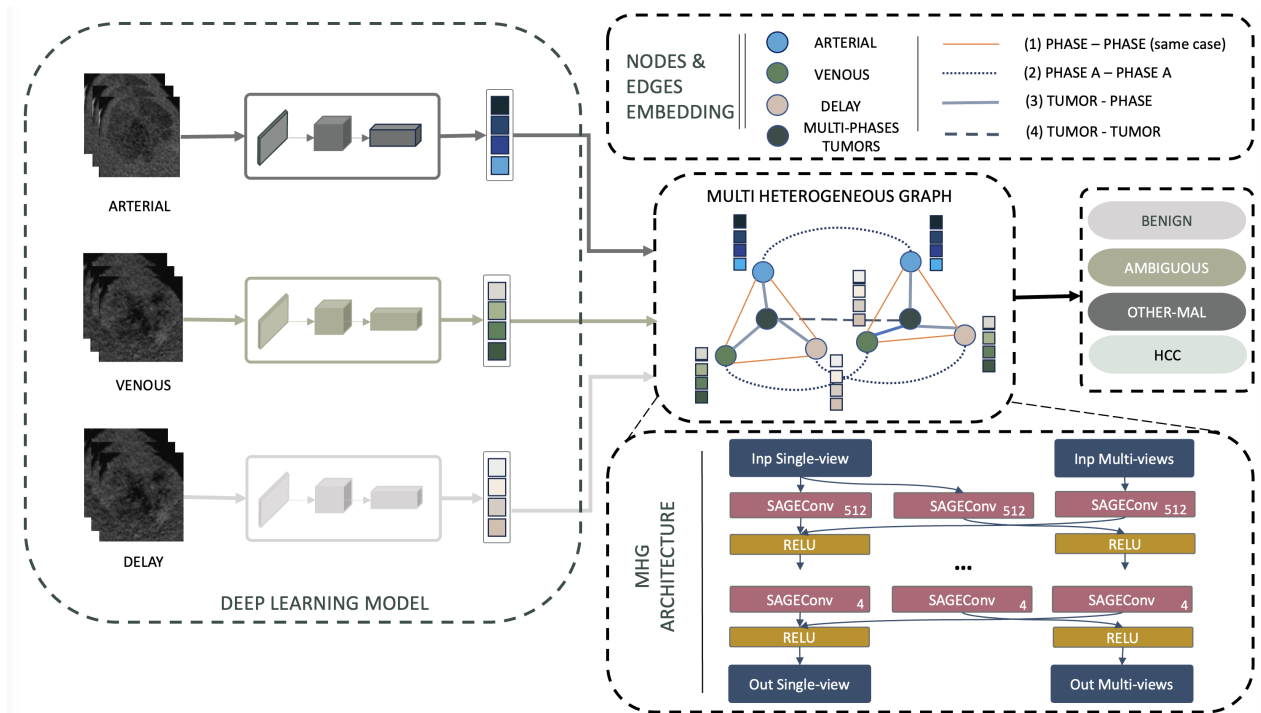


Figure 1: Illustration of our approach for focal liver tumor classification. Two main components consist of (i) the extraction of single-view features and (ii) GIIM.

et al. 2021), (Ouhmich et al. 2019). These techniques often use separate CNN encoders for each view to extract features, which are then fused through strategies like addition, multiplication, or averaging for classification tasks. Additionally, pretrained models, such as multiphase convolutional dense networks, have been developed to analyze four-phase mammograms, leading to improved lesion detection from multiphase data (Su-E et al. 2020), (Weibin et al. 2022).

A comprehensive clinical diagnosis requires a holistic assessment of how abnormalities relate to each other. Factors such as tumor size, location, and spatial relationships between multiple tumors are critical in cancer diagnosis (Lee et al. 2025), (Wu et al. 2018), (Zhang et al. 2020). Despite advances in CADx, many existing methods, including those based on ML, CNNs, transformers, and GNNs, typically analyze lesions independently, overlooking crucial diagnostic clues from interdependencies among abnormalities. They fail to model relationships within a single view (intra-view) and changes across different views or time points (inter-view dependencies). This limitation is particularly problematic for diagnosing small or ambiguous tumors. To address these gaps, we introduce our proposed method, **GIIM**, as it is a Graph-based approach that models both Inter- and Intra-view dependencies for Multi-view medical image classification. Our novel architecture, based on Multi-Heterogeneous Graphs (MHGs), captures complex relationships and correlations between all lesions of a patient. This graph-based paradigm offers advantages over methods like CNNs and Transformers, which require fixed-size inputs. GNNs are inherently flexible, modeling a variable number of

lesions and their intricate connections, making GIIM framework suitable for clinical diagnosis. GIIM also mitigates issues associated with incomplete multi-view data, which can negatively affect diagnostic performance. Multi-view imaging provides comprehensive information but often encounters missing data issues, hindering image-based analysis and raising uncertainties in data classification (Patel et al. 2020), (Zhang et al. 2022), (Mengyao et al. 2023). Addressing these limitations significantly enhances medical image classification performance and potentially aids in clinical practice.

The main contributions can be highlighted as follows.

1. **A Novel GIIM Architecture:** We propose GIIM architecture based on MHGs that is uniquely designed to integrate both intra-view and inter-view structural dependencies. This breakthrough enhances diagnostic accuracy across diverse imaging modalities, including CT, MRI, and mammography.
2. **Robustness to Incomplete Data:** Our approach effectively addresses the common challenge of incomplete multi-view medical images. We introduce four techniques for missing data representation to ensure robust performance in missing-view scenarios.
3. **Extensive Experimental Validation:** The proposed model is evaluated on multiple real medical datasets under both complete and incomplete data settings. Through extensive experiments, our GIIM significantly improves diagnostic outcomes and outperforms existing methods in accuracy and area under the receiver operating characteristic curve (AUC).

Algorithm 1: GIIM Training Pipeline

Input: Dataset D of all tumor images, grouped by patient.
Output: V trained feature extractors $\{f_v\}_{v=1}^V$ and a trained MHG model \mathcal{M}_{MHG} .

```
1: Stage 1: Train Single-View Feature Extractors
2: for  $v = 1 \rightarrow V$  do
3:   Let  $D_v$  be the dataset of all tumor images for view  $v$ .
4:   Initialize single-view model  $N_v$  (e.g., ConvNeXt).
5:   for each training epoch do
6:     for each tumor sample  $(x_{tumor}, y_{tumor})$  in  $D_v$  do
7:       Predict  $\hat{y} \leftarrow N_v(x_{tumor})$ .
8:       Calculate loss  $\mathcal{L}(\hat{y}, y_{tumor})$ .
9:       Update parameters of  $N_v$ .
10:    end for
11:  end for
12:  Freeze the trained backbone of  $N_v$  and denote it as
    feature extractor  $f_v$ .
13: end for
14:
15: Stage 2: Train the MHGs
16: Group dataset by patient:  $D_{patient} = \{(P_i, Y_i)\}_{i=1}^{NumPatients}$ ,
    where  $Y_i$  is the set of all tumor-level labels for patient  $i$ .
17: Initialize the MHG model  $\mathcal{M}_{MHG}$ .
18: for each training epoch do
19:   for each patient case  $(P_i, Y_i)$  in  $D_{patient}$  do
20:     Let PatientFeatures  $\leftarrow []$ .
21:     for each lesion  $l$  in patient  $P_i$  do
22:       Extract features for all its views:  $N_{single} =$ 
          $\{f_v(x_{l,v}) \mid v = 1 \dots V\}$ .
23:       Append  $N_{single}$  to PatientFeatures.
24:     end for
25:     {Construct a HG for patient  $i$  using Eq 1,2,3,4.}
26:      $G_i \leftarrow \text{ConstructMHG}(\text{PatientFeatures})$ .
27:     Let  $Z_i \leftarrow \mathcal{M}_{MHG}(G_i)$ . { $Z_i$  contains output logits
        for each tumor node.} using Eq.5,6,7.
28:     Calculate loss  $\mathcal{L}(Z_i, Y_i)$  over all tumor nodes.
29:     Update parameters of model  $\mathcal{M}_{MHG}$ .
30:   end for
31: end for
32: return  $\{f_v\}_{v=1}^V$  and  $\mathcal{M}_{MHG}$ 
```

Proposed Method

This section presents the proposed method that comprises two main parts: (i) the extraction of single-view features and (ii) GIIM as a comprehensive multi-view approach. The visual representation of our approach is illustrated in Figure 1. The pseudo code of the proposed method is also presented in Algorithm 1.

Single View Feature Extraction

The single-view network architecture is designed to automatically extract features from images, which are then mapped to predictions via two fully connected layers. For this task, we adopt the ConvNeXt (Liu et al. 2022) architecture as our feature extractor. ConvNeXt (Liu et al. 2022) is particularly well-suited for medical image analysis because its modern convolutional design effectively captures

both the large-scale context and fine-grained details crucial for accurate diagnosis. Its use of large-kernel spatial convolutions (7×7) provides a larger receptive field to better analyze the morphology of sizable tumors and understand the spatial relationship between a lesion and surrounding tissue. Additionally, its efficient depthwise separable convolutions are adept at learning detailed spatial patterns, which helps maintain lesion boundaries and sensitivity to changes in tumor intensity. Our methodology involves individually training distinct models on their corresponding single-view datasets. For a dataset with V views, this process generates V separate models, each tailored to a specific view to produce meaningful features. These extracted features are then harnessed as inputs for the subsequent multi-view model, enabling a comprehensive analysis of the data’s underlying correlations.

GIIM: Graph-based Model of Inter and Intra-view Dependencies for Multi-view

This section details how we represent multi-view medical data as a heterogeneous graph and the message-passing mechanism used for learning. We denote the total number of samples in the dataset as S . Each sample, which contains V different views, is represented as a heterogeneous graph G_s from the collection $\{G_s\}_{s=1}^S$. Each graph G_s is composed of a set of vertices (nodes) N and a collection of edges E , and is associated with a class label y_s .

Node Representation Two primary types of nodes exist within each graph:

- **Single-view node** (N_{single}): A set of nodes $\{N_{single_v}\}_{v=1}^V$, where each node represents one of the V distinct views.
- **Multi-view node** (M_{multi}): A single node created by concatenating the features of all single-view nodes: $M_{multi} = \parallel_{v=1}^V (N_{single_v})$.

Edge Representation The set of edges E models the relationships between nodes. Let $N_{v,j}$ be the node for single-view v of lesion L_j , and M_j be the multi-view summary node for lesion L_j . The total edge set E is the union of four distinct subsets:

- **Intra-tumor, Inter-view** (E_{intra}): Connects different views of the same lesion to capture temporal changes.

$$E_{intra} = \{(N_{v,j}, N_{v',j}) \mid j \in L, v \neq v'\} \quad (1)$$

- **Single-to-Multi-view** (E_{s-m}): Connects a single-view node to its corresponding multi-view summary node.

$$E_{s-m} = \{(N_{v,j}, M_j) \mid j \in L, v \in V\} \quad (2)$$

- **Inter-tumor, Single-view** ($E_{inter-s}$): Connects different lesions observed in the same view.

$$E_{inter-s} = \{(N_{v,j}, N_{v,j'}) \mid v \in V, j \neq j'\} \quad (3)$$

- **Inter-tumor, Multi-view** ($E_{inter-m}$): Models the high-level contextual relationship between all lesions in the case, which is useful for identifying small lesions by their proximity to larger ones.

$$E_{inter-m} = \{(M_j, M_{j'}) \mid j \neq j'\} \quad (4)$$

Heterogeneous Message Passing To learn from this structure, we employ a heterogeneous message-passing scheme where the feature update for any node n at layer k separately processes its neighbors based on their type.

1. Neighbor Aggregation We compute two distinct aggregated feature vectors for node n : one from its single-view neighbors and one from its multi-view neighbors.

- Aggregation from Single-view neighbors ($N_{single}(n)$):

$$h_{N_{single}(n)}^k \leftarrow \frac{1}{|N_{single}(n)|} \sum_{u \in N_{single}(n)} \mathbf{W}_{single}^k h_u^{k-1} \quad (5)$$

- Aggregation from Multi-view neighbors ($M_{multi}(n)$):

$$h_{M_{multi}(n)}^k \leftarrow \frac{1}{|M_{multi}(n)|} \sum_{u \in M_{multi}(n)} \mathbf{W}_{multi}^k h_u^{k-1} \quad (6)$$

Here, \mathbf{W}_{single}^k and \mathbf{W}_{multi}^k are distinct, trainable weight matrices for layer k that learn different transformations for each neighbor type.

2. Final Update The final feature vector h_n^k is computed by concatenating the node’s own previous state, h_n^{k-1} , with the two aggregated neighborhood vectors. This combined vector is then passed through a final dense layer with a non-linear activation function σ .

$$h_n^k \leftarrow \sigma \left(\mathbf{W}^k \cdot \text{CONCAT}(h_n^{k-1}, h_{N_{single}(n)}^k, h_{M_{multi}(n)}^k) \right) \quad (7)$$

Where \mathbf{W}^k is the main weight matrix for the update step in layer k .

MHG Architecture In terms of architecture design, the proposed MHG-based model contains a sequence of five SAGEConv layers with varying units of 512, 256, 128, 64, and 32. The final SAGEConv layer has a number of units equivalent to the output classes, used to compute the prediction probability, thus supplanting the need for a fully connected layer. A detailed illustration of the model architecture is provided in Figure 1. In the following, we present the proposed GIIM for typical multi-phase and multi-view medical imaging diagnosis problems.

GIIM in Medical Context

Liver Tumor Classification on Multi-phase CT Accurate classification and characterization of focal liver lesions (FLLs) are vital in guiding effective patient treatment strategies. Liver examination is commonly conducted using various medical imaging techniques, including CT, MRI, ultrasound, and X-ray. Among these, multi-phase contrast-enhanced CT emerges as an effective approach for detecting FLLs, particularly those related to hepatocellular carcinoma (HCC). Such FLLs can be categorized by processing the multi-phase CT scans comprising arterial, venous, and delay phases.

The proposed GIIM model is capable of harnessing the distinctive features presented across various phases for classifying the FLLs. Let S denote the number of patients

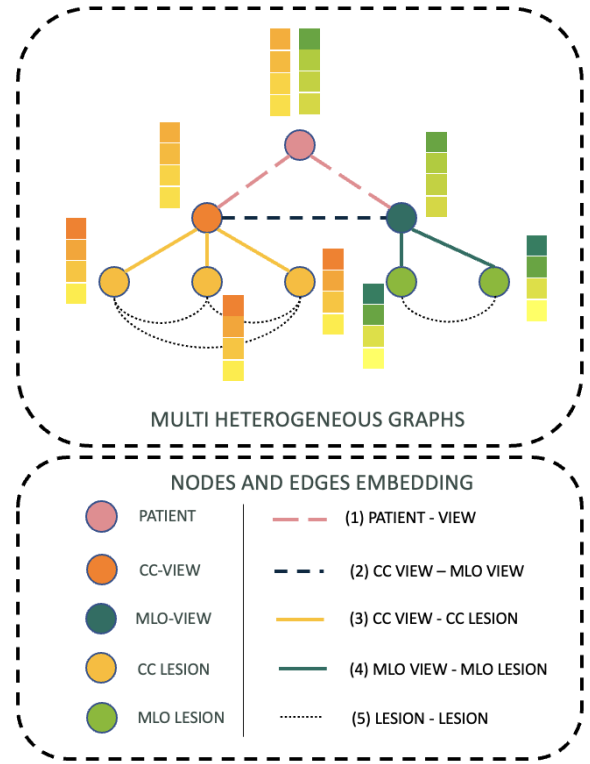


Figure 2: Our proposed GIIM for BI-RADS classification.

present in the dataset. Each patient is associated with L lesions, and every lesion is observed across three distinct phases denoted by $V = \{Arterial, Venous, Delay\}$. Our GIIM, denoted as $G_{MHG_s} = \{G(N, E)\}_{s=1}^S$, comprises node embeddings $\{N = \{N_{single}, M_{multi}\}_{l=1}^L$.

For FLL classification, the edge set E models four key relationships: E_{intra} capture temporal changes within a single lesion, while E_{s-m} integrate these phase-specific details into a holistic summary. To model context between different lesions, $E_{inter-s}$ correlate them at specific time points, and $E_{inter-m}$ connect their overall summary nodes. This inter-lesion modeling is particularly beneficial as it provides contextual clues, acknowledging that certain tumor types often appear together.

BIRADS Classification on Mammograms Breast cancer has now taken the lead as the most prevalent cancer among women, as confirmed by statistics from the International Agency for Research on Cancer (IARC) in December 2020 (Organization Dec. 15, 2020 [Online]). Mammography, a widely adopted X-ray imaging technique for breasts, plays a crucial role in CADx systems aimed at supporting radiologists in evaluating breast cancer risk. For a patient, there are two distinct breast sides: left and right, each with two views - CC (top-down) and MLO (side). Within these views, L lesions are present. CADx in mammograms faces challenges when a lesion in the CC view might not be visible in the MLO view. Mapping lesion A from the CC view to lesion B in the MLO view becomes complicated due to

the lack of necessary information and differing observation angles. Therefore, GIIM designed for BI-RADS classification slightly differ from GIIM for FLLs classification. The single-view node embedding $N_v = \frac{1}{L} \sum_{l=1}^L f_v(l_v)$ represents the mean of all lesion features within that specific view, while the multi-view node embedding $M_m = \parallel_{v=1}^V M_v$ signifies the concatenation of different single-view features. The detailed architecture of GIIM for BI-RADS classification is illustrated in Figure 2.

Breast Lesions Classification on MRI In this scenario, $V = \{Pre-contrast, Post-contrast, Subtraction\}$ was defined as a set representing the three distinct MRI sequences acquired for each case study. A function $f_v(\cdot)$ for $v \in V$ represents the feature extractor for each sequence. This function extracts relevant information from the input data l_v associated with sequence v . The proposed GIIM is denoted as $G_{MHGs} = \{G(N, E)\}$. Here, $G(N, E)$ represents a graph for each case study, where N represents the set of nodes and E represents the edges connecting them. The GIIM model utilizes two types of node embeddings for each patient: (1) Single-Sequence Features $N_{single} = f_v(l_v)$ represents the features extracted from sequence v using the corresponding feature extractor f_v . (2) Multi-Sequence Features $M_{multi} = \parallel_{v \in V} f_v(l_v)$ represents the concatenated features from all sequences for a particular lesion. This captures the combined information across all sequences. By incorporating both single-sequence and multi-sequence features, the model aims to capture the individual characteristics of each sequence and the relationships between them, potentially leading to improved breast lesion classification on MRI data.

GIIM in the Missing-view Context

To address the challenge of missing views, our approach first generates missing features and then uses them within multi-view methodologies. This exploration aims to identify suitable multi-view techniques and establish a meaningful representation of features for the missing views. Let $\mathcal{F}_v = \{\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_n\}$ denote the feature map extracted from the single-view model, where $\mathcal{F}_v^a \in \mathcal{F}_v$ denotes the feature maps representing available views. The goal is to generate missing view feature maps \mathcal{F}_v^m from \mathcal{F}_v^a with $\mathcal{F}_v^a \cup \mathcal{F}_v^m = \mathcal{F}_v$.

Missing view: Constant defines the feature vector for the missing view as a constant zero vector $F_v^{1 \times c} = [0.0]^{1 \times c}$

Missing view: Learnable treats the feature vector as a set of learnable parameters that are optimized during training. The final vector is produced by pooling a randomly initialized learnable tensor and then normalizing it using the Frobenius norm $\hat{F}_v^{1 \times c} = \frac{\hat{F}_v}{\|\hat{F}_v\|_F}$

Missing view: RAG-based We additionally applied a retrieval-augmented strategy focused on imputation rather than synthesis. If the feature was absent, we merged the available features, computed similarity across the dataset, identified the entry with the highest similarity, and used its corresponding missing features to populate the missing field.

To calculate the missing feature \mathcal{F}_v^m , we can use the following formula:

$$\mathbf{x} = \begin{bmatrix} \mathcal{F}_v^{a,1} \\ \dots \\ \mathcal{F}_v^{a,k} \end{bmatrix}, \quad \mathbf{x}_i = \begin{bmatrix} \mathcal{F}_{v,db}^{a,1} \\ \dots \\ \mathcal{F}_{v,db}^{a,k} \end{bmatrix}, \quad s_i = \frac{\mathbf{x}^T \mathbf{x}_i}{\|\mathbf{x}\|_2 \|\mathbf{x}_i\|_2}$$

$$i^* = \arg \max_i s_i, \quad \mathcal{F}_v^m = \mathcal{F}_{v_{j^*}}^m \quad (8)$$

Available features in \mathcal{F}_v^a are stacked column vectors \mathbf{x} (query) and \mathbf{x}_i (database). Calculate the similarity s_i for each pair $(\mathbf{x}, \mathbf{x}_i)$, identify the index i^* with the highest similarity, and copy the corresponding missing value $\mathcal{F}_{v_{j^*}}^m$ into \mathcal{F}_v^m .

Missing view: Covariance-based Besides, we propose a method to impute missing features by leveraging a covariance-based similarity metric. The method begins by constructing a feature space from a database of k complete available views. This space is defined by a set of vectors $\tilde{\mathcal{F}}_v^a = \{\Delta_1, \Delta_2, \dots, \Delta_k\}$, where each vector $\Delta_i = \mathcal{F}_v^{a,1} - \sum_{j=2}^k \mathcal{F}_v^{a,j} \in R^c$ quantifies the difference among available features for a given sample $i = 1, \dots, n$. A covariance matrix $\Sigma \in R^{c \times c}$ is then computed over the feature space \mathcal{F} to capture the statistical relationships between the feature components:

$$\Sigma = \frac{1}{n-1} \sum_{i=1}^n (\Delta_i - \mu)(\Delta_i - \mu)^T \quad (9)$$

where $\mu \in R^c$ is the mean vector of the feature space \mathcal{F} . To generate the missing feature \mathcal{F}_v^m for a query sample, its corresponding available feature difference vector, Δ^q , is first computed. A covariance similarity score, s_j , is then calculated between this query vector and each vector Δ_j in the database. The index j^* of the database sample with the highest similarity is identified, and its delay feature is adopted, as follows:

$$s_j = (\Delta^q)^T \Sigma \Delta_j$$

$$j^* = \arg \max_j s_j, \quad \text{then } \mathcal{F}_v^m = \mathcal{F}_{v_{j^*}}^m \quad (10)$$

This process imputes missing value by borrowing from the most statistically similar complete sample in feature space.

Experiments

This section first presents the datasets used for model evaluations and then describes experimental results and analysis.

Datasets

Liver Tumor Dataset Our method's performance is evaluated on a private dataset comprising 920 abdominal CT examinations conducted between 2016 and 2022. These exams originate from 877 patients, ranging in age from 5 to 96 years, and contain four distinct 3D CT images corresponding to different phases of contrast enhancement: non-contrast, arterial, venous, and delayed phases. We stratify the

	Group	Data		Number of Tumors				
		# Patients	# Samples	Benign	Ambiguous	Malignant	HCC	Total
Train	625	659	561	195	147	491	1394	
Test	252	261	176	43	56	165	440	
Total	877	920	737	238	203	656	1834	

Table 1: The statistics of the liver tumor dataset.

Type	Liver			VinDr-Mammo			BreastDM		
	Methods	Acc (%)	AUC (%)	Methods	Acc (%)	AUC (%)	Methods	Acc (%)	AUC (%)
Single	Arterial	65.91	82.78	CC	62.56	77.81	Pre-contrast	82.97	82.94
	Venous	68.64	84.91	MLO	61.57	78.01	Post-contrast	72.34	75.29
	Delay	66.82	82.83				Subtraction	76.59	77.84
Multi	NN-based	75.45	89.09	NN-based	67.48	82.21	NN-based	80.85	87.35
	ML-based	73.63	88.00	ML-based	66.87	80.86	ML-based	82.98	79.41
	Attention-based	73.41	88.53	Attention-based	68.09	81.00	Attention-based	85.1	76.37
	GIIM (ours)	78.20	91.05	GIIM (ours)	71.17	82.54	GIIM (ours)	87.23	89.02

Table 2: Performance metrics in terms of accuracy and AUC of different single- and multi-phase methods evaluated on the liver tumor test set, the VinDr-Mammo dataset and the BreastDM dataset.

dataset into two groups using label stratification: a training set (659 cases) and a test set (261 cases). Our classification categorizes lesions into four distinct classes based on their pathological characteristics: benign, ambiguous, malignant other than HCC, and HCC tumors. Benign lesions encompass a variety of liver tumors such as focal nodular hyperplasia (FNH), cyst, hemangioma, abscess, differential diagnosis (DD), and adenoma. The ambiguous label contains two types of liver tumors, LIRADS 3 and ambiguous DD. The malignant group contains three types of tumors: intrahepatic cholangiocarcinoma (ICC), metastases, and malignant DD. Detailed dataset statistics are listed in Table 1.

VinDr-Mammo Dataset The VinDr-Mammo dataset (Nguyen et al. 2023) is a comprehensive benchmark for full-field digital mammography, containing 5,000 four-view examinations with 20,000 scans. Each scan is annotated with breast-level BI-RADS and findings information, making it the most extensive public resource for BI-RADS assessment. The mammography images were acquired from H108 and HMUH sources. In our approach, we extract ROIs representing lesions from mammograms for the BI-RADS classification task. Each lesion is annotated with an internal BI-RADS label (3-5). A sample consists of CC and MLO views, each containing multiple lesion ROIs. The global label for the sample is the highest pathology, denoted as $y_G = \max([y_{l=1}^{L_{CC}} || y_{l=1}^{L_{MLO}}])$. Model performance is evaluated at the sample level. The dataset comprises 783 samples (1,930 lesions), split into a **training set** of 620 samples (1,532 lesions) and a **test set** of 163 samples (398 lesions). The label distribution for BI-RADS $\{3, 4, 5\}$ is $N_{train} = \{264, 271, 85\}$ and $N_{test} = \{71, 71, 21\}$ for the training and test sets, respectively.

BreastDM Dataset BreastDM (Zhao et al. 2023) is a dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) dataset curated for advancing research in breast lesion classification, particularly focusing on malignant and

benign differentiations. The dataset encompasses a total of 232 cases, comprising 147 malignant and 85 benign lesions. Each case within the dataset includes three distinct types of sequences: pre-contrast, post-contrast, and subtraction.

Experiment Results and Analysis

Multi-view Experiments In modern medical image analysis, multi-view classification techniques can be broadly categorized by their core methodology. These include traditional ML methods that excel with structured data, CNN based approaches that focus on local spatial patterns, Attention/Transformer-based models designed to capture global dependencies, and GNN based techniques that explicitly model relationships between different entities. To comprehensively evaluate GIIM, we compare its performance against established techniques from these other domains. To ensure a **fair comparison**, all methods share the same core feature extraction backbone, which was originally designed for single-view models. This isolates the comparison to the effectiveness of the multi-view learning technique itself. The classifiers we evaluate against are:

- **NN-based:** A classifier that utilizes two fully connected layers to process the extracted features.
- **ML-based:** The LightGBM classifier, a powerful gradient-boosting framework for structured feature data, as implemented by (TX et al. 2022).
- **Attention-based:** Phase attention modules introduced by (Weibin et al. 2022), which use an attention mechanism to weigh the importance of features from different views.

The performance of these approaches on the private liver dataset, the VinDr-Mammo dataset and BreastDM dataset is presented in Table 2. It can be observed from Table 2 that multi-view approaches achieved significant higher performances, showcasing large margin of consistent improvements of approximately 12% in accuracy and 8.3% in AUC on the in-house dataset, 7.8% in accuracy and 4.8% in AUC

Method		Test on 100% missing-view					Test on full view				
		$\eta = 0.0$	$\eta = 0.2$	$\eta = 0.5$	$\eta = 0.7$	$\eta = 1.0$	$\eta = 0.0$	$\eta = 0.2$	$\eta = 0.5$	$\eta = 0.7$	$\eta = 1.0$
Liver Dataset	NN-based	70.00	72.50	70.23	71.82	72.50	75.45	73.41	71.36	71.59	72.50
	ML-based	71.59	72.95	72.05	72.27	72.73	73.63	73.86	72.73	73.18	72.73
	Attention-based	67.50	73.18	71.36	67.95	72.73	73.41	71.82	72.95	64.09	67.50
	GIIM (constant)	72.27	73.86	72.73	74.09	73.41	78.20	75.68	73.41	72.50	73.86
	GIIM (learnable)	73.64	73.86	73.86	72.95	70.91	78.20	74.09	76.82	76.82	65.45
	GIIM (RAG-based)	74.31	72.73	72.95	71.81	72.50	78.20	74.55	75.91	76.14	76.82
	GIIM (Covariance)	70.91	74.32	72.95	73.18	72.73	78.20	78.41	75.91	77.95	74.31
VinDr-Mammo	NN-based	63.19	63.19	63.19	64.42	63.80	67.48	63.19	63.80	65.03	62.58
	ML-based	58.28	61.96	63.80	63.19	62.58	66.87	65.03	63.80	63.80	62.58
	Attention-based	63.19	65.03	65.64	65.64	66.26	68.09	63.19	63.80	65.03	63.04
	GIIM (constant)	61.35	65.03	66.26	66.26	66.26	71.17	66.87	65.03	65.03	66.26
	GIIM (learnable)	63.19	65.03	64.42	66.26	65.64	71.17	64.42	63.19	61.35	60.74
	GIIM (RAG-based)	59.51	60.74	62.58	61.96	61.96	71.17	65.64	66.87	68.71	66.26
	GIIM (Covariance)	61.35	62.58	59.51	63.80	63.80	71.17	67.48	65.64	66.26	66.26

Table 3: Comparison in terms of classification accuracy with missing view rate $\eta = [0.0, 0.2, 0.5, 0.7, 1.0]$ on Liver and VinDr-Mammo datasets.

on VinDr-Mammo, as well as 5% in accuracy and 7% in AUC on BreastDM when compared to single-view methods. This result can justify the fact that multi-view medical imaging provides a complete perspective of lesions, capturing insights from mammogram observation angles to different observation timeframes in CT scans and breast MRI.

Among the multi-view approaches, GIIM model consistently outperforms other multi-view methods, with an 3% improvement in accuracy and 2% increment in AUC on the liver tumor dataset. For the mammogram dataset, our GIIM achieves the highest accuracy among the tested methods. The same observation can be drawn from results obtained on the BreastDM dataset, in which the proposed GIIM obtained highest accuracy of 87.23% and AUC of 89.02%.

Missing-view Experiments In medical analysis, missing views present a significant challenge for accurate diagnosis. This issue occurs for various reasons, including routine clinical protocols that may not require all views (e.g., using only arterial and venous phases for FLL classification), as well as technical errors or patient refusal. Such incomplete data can severely degrade or even deactivate the predictive performance of multi-view models. To assess the influence of incomplete view data on model performance, we trained models on datasets with varying rates of missing views (η) and then evaluated on two distinct test sets: **Full-view Test**, where all samples have complete views, and **Miss-view Test**, where all samples are missing one view. For FLLs classification, we consider the scenario where the Delay phase is missing, whereas for BI-RADS classification, we assume the CC view is missing.

First, we compared our basic GIIM model (using a constant vector for missing views) against other multi-view methods. As shown in Table 3, **GIIM (constant)** model consistently and significantly outperforms the NN-based, ML-based and Attention-based methods in both the full view and miss view test sets across two datasets. This demonstrates the fundamental advantage of explicitly modeling inter- and intra-lesion relationships for multi-view diagnosis.

We then investigated four different methods for handling missing data within GIIM framework: constant, learnable, RAG-based, and Covariance-based. A clear overall trend is that all methods perform better on the **Full-view Test** than the **Miss-view Test** around 4% in Liver dataset and 5% in VinDr-Mammo. This confirms the intuitive idea that having complete data at the time of diagnosis leads to better accuracy. Interestingly, there is a trade-off between performance on the two test sets. On **Full-view Test**, RAG-based and Covariance methods often achieve the best results. This suggests that these methods learn to generate a feature vector with a similar distribution to a real view’s features, which is beneficial when a full set of views is available. On **Miss-view Test**, however, the simple constant vector method is often more effective. A likely explanation is that using a unique, constant vector makes the missing node obvious to the graph, causing the model to learn to rely more heavily on the available nodes.

Conclusion

This paper introduces GIIM, a novel MHGs-based approach for medical image classification that leverages both intra- and inter-view information inherent in several imaging modalities. Through extensive experimental validation on both internal and public datasets, our proposed method yields significantly enhanced diagnostic precision compared to several existing single-phase and other multi-view techniques. Furthermore, for the cases of incomplete viewing data, our GIIM demonstrates remarkable performance and robustness compared to existing multi-view approaches. Our study opens up a promising approach to constructing and developing an efficient CADx, that can work well even when lacking input information and thus potentially benefit practical clinics. Our ongoing work involves exploring novel MHGs architectures for medical imaging analysis, specifically focusing on exploring lesion relationships and integrating expert knowledge to construct powerful MHGs.

References

- Abiwinanda, N.; Hanif, M.; Hesaputra, S. T.; Handayani, A.; and Mengko, T. R. 2019. Brain tumor classification using convolutional neural network. In *World Congress on Medical Physics and Biomedical Engineering 2018: June 3-8, 2018, Prague, Czech Republic (Vol. 1)*, 183–189. Springer.
- Ayadi, W.; Elhamzi, W.; Charfi, I.; and Atri, M. 2021. Deep CNN for brain tumor classification. *Neural processing letters*, 53: 671–700.
- Cameron, S.; Krithika, S.; Carsten; Seong, C. Y.; Rajan, J.; and Debashis, G. 2022. A Pipeline for the implementation and visualization of explainable machine learning for medical imaging using radiomics features. *Sensors*, 22(14): 5205.
- Gudhe, N. R.; Mazen, S.; Sund, R.; Kosma, V.-M.; Behravan, H.; and Mannermaa, A. 2024. A Multi-View Deep Evidential Learning Approach for Mammogram Density Classification. *IEEE Access*.
- Ibrahim, M.; Shagufta, H.; Gary, C.; et al. 2022. Multi-Graph Convolutional Neural Network for Breast Cancer Multi-task Classification. *Irish Conference on Artificial Intelligence and Cognitive Science*.
- Lambin, P.; Rios-Velazquez, E.; Leijenaar, R.; Carvalho, S.; Van Stiphout, R. G.; Granton, P.; Zegers, C. M.; Gillies, R.; Boellard, R.; Dekker, A.; et al. 2012. Radiomics: extracting more information from medical images using advanced feature analysis. *European journal of cancer*, 48(4): 441–446.
- Lee, C. U.; Kim, J. H.; Kim, G.; Jeon, J.; Kim, J. W.; Lee, Y. S.; Min, K.; Tae, J. H.; Choi, S. Y.; Chang, I. H.; and Choi, J. 2025. Impact of tumor size and sex on diagnostic discrepancies in renal mass imaging. *Nature Scientific Report*.
- Liu, X.; Khalvati, F.; Namdar, K.; Fischer, S.; Lewis, S.; Taouli, B.; Haider, M. A.; and Jhaveri, K. S. 2021. Can machine learning radiomics provide pre-operative differentiation of combined hepatocellular cholangiocarcinoma from hepatocellular carcinoma and cholangiocarcinoma to inform optimal treatment planning? *European Radiology*, 31(1): 244–255.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A ConvNet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Luo, L.; Wang, X.; Lin, Y.; Ma, X.; Tan, A.; and Chan, R. 2024. Deep Learning in Breast Cancer Imaging: A Decade of Progress and Future Directions. *IEEE Reviews in Biomedical Engineering*.
- Manigrasso, F.; Milazzo, R.; Russo, A. S.; Lamberti, F.; Strand, F.; Pagnani, A.; and Morra, L. 2025. Mammography classification with multi-view deep learning techniques: Investigating graph and transformer-based architectures. *Medical Image Analysis*.
- Marios, A.; Stergios, C.; Lukas, E.; Andreas, C.; and Stavroula, M. 2016. Lung pattern classification for interstitial lung diseases using a deep convolutional neural network. *IEEE transactions on medical imaging*, 35(5): 1207–1216.
- Mengyao, X.; Zongbo, H.; Changqing, Z.; Yichen, B.; and Qinghua, H. 2023. Exploring and Exploiting Uncertainty for Incomplete Multi-View Classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19873–19882.
- Mikhael, P.; Wohlgend, J.; Karstens, L.; Xiang, J.; Takigami, A.; and Bourgouin, P. 2023. Sybil: A Validated Deep Learning Model to Predict Future Lung Cancer Risk From a Single Low-Dose Chest Computed Tomography. *Journal of Clinical Oncology*.
- Nguyen, H. T.; Nguyen, H. Q.; Pham, H. H.; Lam, K.; Le, L. T.; Dao, M.; and Vu, V. 2023. VinDr-Mammo: A large-scale benchmark dataset for computer-aided diagnosis in full-field digital mammography. In *Scientific Data (Nature)*.
- Organization, W. H. Dec. 15, 2020 [Online]. Latest global cancer data: Cancer burden rises to 19.3 million new cases and 10.0 million cancer deaths in 2020.
- Ouhmich, F.; Agnus, V.; Noblet, V.; Heitz, F.; and Pessaux, P. 2019. Liver tissue segmentation in multiphase CT scans using cascaded convolutional neural networks. *International journal of computer assisted radiology and surgery*, 14: 1275–1284.
- Patel, D. V.; Hsu, C.-C.; Varghese, B.; Cen, S.; Hwang, D.; Gill, I.; Duddalwar, V.; and Oberai, A. A. 2020. Probabilistic Recovery of Missing Images in Contrast-Enhanced CT. In *Proceedings of the NeurIPS Workshop on Medical Imaging*.
- Sam, T. B.; Huy, T. D.; Dao, C. T.; Lam, T. T.; Tang, V. H.; and Truong, S. Q. 2024. A multi-phase multi-graph approach for focal liver lesion classification on ct scans. *Proceedings of the Asian Conference on Computer Vision*.
- Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2008. The graph neural network model. *IEEE transactions on neural networks*, 20(1): 61–80.
- Sharif, M. I.; Khan, M. A.; Alhoussein, M.; Aurangzeb, K.; and Raza, M. 2021. A decision support system for multimodal brain tumor classification using deep learning. *Complex & Intelligent Systems*, 1–14.
- Su-E, C.; Lin-Qi, Z.; Si-Chi, K.; Wen-Qi, S.; Bing, H.; Si-Dong, X.; Yi-Nan, C.; Hui, L.; Si-Min, C.; Ting, J.; et al. 2020. Multiphase convolutional dense network for the classification of focal liver lesions on dynamic contrast-enhanced computed tomography. *World journal of gastroenterology*, 26(25): 3660.
- Sun, L.; Yu, K.; Batmanghelich, K.; et al. 2022. Context-aware Self-supervised Learning for Medical Images Using Graph Neural Network. *arXiv preprint arXiv:2207.02957*.
- Tang, H.; Duong, S.; Nguyen, C.; et al. 2023. Wavelet radiomics features from multiphase CT images for screening hepatocellular carcinoma: analysis and comparison. *Nature journal*.
- Tran, Q. H.; Sam, T. B.; Tang, V. H.; and Truong, S. Q. 2025. Feature-Driven LI-RADS Classification in Liver Cancer: A CNN-Based Approach for CT Imaging Under Data Scarcity. *IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*.
- TX, N. H.; Sam, T.; B, N. D.; H, P. H.; and Q, N. H. 2022. A novel multi-view deep learning approach for BI-RADS

and density assessment of mammograms. In *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2144–2148. IEEE.

Ullah, M. S.; Khan, M. A.; Albarakati, H. M.; Damaševičius, R.; and Alsenan, S. 2024. Multimodal brain tumor segmentation and classification from MRI scans based on optimized DeepLabV3+ and interpreted networks information fusion empowered with explainable AI. *Computers in Biology and Medicine*.

Weibin, W.; Fang, W.; Qingqing, C.; Shuyi, O.; Yutaro, I.; Xianhua, H.; Lanfen, L.; Hongjie, H.; Ruofeng, T.; and Yen-Wei, C. 2022. Phase Attention Model for Prediction of Early Recurrence of Hepatocellular Carcinoma with Multi-phase CT Images and Clinical Data. *Frontiers in Radiology*, 2: 856460.

Wu, G.; Wu, J.; Wang, B.; Zhu, X.; Shi, X.; and Ding, Y. 2018. Importance of tumor size at diagnosis as a prognostic factor for hepatocellular carcinoma survival: a population-based study. *Cancer Management and Research*.

Wu, J.; Liu, A.; Cui, J.; Chen, A.; Song, Q.; and Xie, L. 2019. Radiomics-based classification of hepatocellular carcinoma and hepatic haemangioma on precontrast magnetic resonance images. *BMC medical imaging*, 19(1): 1–11.

Xiao, W.; Houye, J.; Chuan, S.; Bai, W.; Yanfang, Y.; Peng, C.; and Philip, Y. 2019. Heterogeneous graph attention network. In *The world wide web conference, 2022–2032*.

Zhang, N.; Liang, R.; Gensheimer, M. F.; Guo, M.; Zhu, H.; inming Yu; Diehn, M.; Jr, B. W. L.; Li, R.; and Wu, J. 2020. Early response evaluation using primary tumor and nodal imaging features to predict progression-free survival of locally advanced non-small cell lung cancer. *Theranostics journal*.

Zhang, Y.; He, N.; Yang, J.; Li, Y.; Wei, D.; Huang, Y.; Zhang, Y.; He, Z.; and Zheng, Y. 2022. mmformer: Multimodal medical transformer for incomplete multimodal learning of brain tumor segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 107–117. Springer.

Zhao, X.; Liao, Y.; Xie, J.; and et al, X. H. 2023. BreastDM: A DCE-MRI dataset for breast tumor image segmentation and classification. *Computers in Biology and Medicine*.