Connecting the Dots: Graph Neural Network Powered Ensemble and Classification of Medical Images

1st Aryan Singh
Dept. of Electronic and Computer Engineering
University of Limerick
Limerick, Ireland
aryan.singh@ul.ie

2nd Pepijn Van de Ven
Dept. of Electronic and Computer Engineering
University of Limerick
Limerick, Ireland
pepijn.vandeven@ul.ie

3rd Ciarán Eising
Dept. of Electronic and Computer Engineering
University of Limerick
Limerick, Country
ciaran.eising@ul.ie

4th Patrick Denny
Dept. of Electronic and Computer Engineering
University of Limerick
Limerick, Ireland
patrick.denny@ul.ie

Abstract—Deep learning models have demonstrated remarkable results for various computer vision tasks, including the realm of medical imaging. However, their application in the medical domain is limited due to the requirement for large amounts of training data, which can be both challenging and expensive to obtain. To mitigate this, pre-trained models have been fine-tuned on domain-specific data, but such an approach can suffer from inductive biases. Furthermore, deep learning models struggle to learn the relationship between spatially distant features and their importance, as convolution operations treat all pixels equally. Pioneering a novel solution to this challenge, we employ the Image Foresting Transform to optimally segment images into superpixels. These superpixels are subsequently transformed into graph-structured data, enabling the proficient extraction of features and modeling of relationships using Graph Neural Networks (GNNs). Our method harnesses an ensemble of three distinct GNN architectures to boost its robustness. In our evaluations targeting pneumonia classification, our methodology surpassed prevailing Deep Neural Networks (DNNs) in performance, all while drastically cutting down on the parameter count. This not only trims down the expenses tied to data but also accelerates training and minimizes bias. Consequently, our proposition offers a sturdy, economically viable, and scalable strategy for medical image classification, significantly diminishing dependency on extensive training data sets. Our code is available at Github.

Index Terms—Graph Neural Networks, Medical imaging, Computer vision, Classification

I. INTRODUCTION

Deep Neural Networks (DNNs) have been proven to be effective for computer vision tasks and are increasingly being utilized in medical imaging research. These models have evinced state-of-the-art performance in an array of tasks including object detection, image classification [28], semantic segmentation, and instance segmentation [29]. Acquiring labeled data in the medical domain can be an arduous and costly undertaking, and data quality may vary considerably across different sources. Consequently, an ideal model would require less training data and fewer parameters, leading to greater efficiency and reduced computational resources. Moreover, such models would offer superior generalizability [35], a crucial feature that is often limited in DNNs. While DNNs can perform well on the data they were trained on, they may struggle to generalize to new and unseen data, which highlights a critical issue with the use of DNNs in the medical domain [11] where labeled data of high quality can be scarce.

Notably, while DNNs employ convolutional kernels for fixed local pixel grid connectivity and pooling for global feature extraction, they overlook the inherent topological structure of medical images. This omission, crucial for optimal medical image understanding and representation, has been highlighted [42], [43].

Graph-based neural networks (GNNs) adeptly handle variable-sized heterogeneous graph input [45], allowing for adaptability across diverse data and tasks. These models can discern intricate geometric interrelationships in image datasets [44], enhancing predictive performance [10]. GNNs, capable of analyzing complex interconnected phenomena, are increasingly used in medical imaging [41]. Their application in medical imaging, including classification [40] and segmentation [39], has seen significant progress. GNNs have proven effective in multi-modal data-based medical image analysis [33], [36] and in human-object interaction detection using deep neural network (DNN) features [34]. Motivated by these advancements, we aim to explore the synergistic use of DNN

features and GNNs in transducing medical images into graphs.

In this study, GNNs were applied to classify pneumonia from medical images of the lungs and have demonstrated effectiveness in this task. Our approach utilizes a method that involves transforming X-ray images into graphs, taking into account the topological connections between the various features present. This graph structure not only enhances the global level representation captured by DNNs, but also provides a more comprehensive understanding of the image's overall structure. By leveraging this graph structure, our approach can effectively capture both locally and globally relevant features, resulting in improved performance in X-ray image analysis tasks. Our study presents a compelling comparison between the performance of GNNs and three prominent DNNs, which serves to highlight the significant advantages of incorporating GNNs alongside our method for graph creation. By showcasing the improved performance achieved through the use of GNNs, we provide strong evidence for the effectiveness of our approach and its potential for application in the medical imaging domain. We made efforts to enhance performance by ensembling multiple GNN architectures trained on varying graph sizes, reaching a sensitivity score as high as 99%, a critical factor in the medical domain. By ensembling different GNNs, we can leverage the strengths of each model while mitigating their individual weaknesses, ultimately yielding a more robust and accurate system.

II. METHODS

Our method is divided into two distinct stages.

The first stage, as illustrated in Figure 1, involves dividing the input image into smaller, homogeneous regions called superpixels by grouping adjacent pixels based on aggregating criteria such as color, intensity, or texture similarity. We then create a graph with these superpixels as nodes and establish edges between nodes based on region adjacency and homogeneity explained in section II-B. The features for each superpixel are obtained using a DNN. The extracted features from the superpixel segments of the image. including basic features like edges, corners, and textures from the early stages of the network. As the image gets passed through the deeper layers of the network, it extracts more complex and abstract features such as objects, scenes, and different parts of objects. These features are then used as node features in the graph.

The second stage, as depicted in Figure 1, classifies the entire graph for pneumonia. Utilizing a GNN for graph classification generates distinct embeddings for pneumonia and non-pneumonia graphs, which incorporate the features extracted using a DNN as node features. The aggregation of neighbouring node features contributes to better embedding, enhancing graph-level classification performance.

A. Superpixel Segmentation

We have used superpixel segmentation to divide an image into smaller regions, that are more homogeneous in terms of colour, texture, and other features. Superpixel segmentation overcomes the limitations of primitive segmentation methods, which tend to produce irregular and fragmented regions. Using superpixel segmentation, we effectively isolate regions of interest and obtain a more detailed understanding of the underlying image structure. To generate superpixels, we have implemented two algorithms: Simple Linear Iterative Clustering (SLIC) [5] and Dynamic and Iterative Spanning Forest (DISF) [6].

1) SLIC: The SLIC algorithm begins with dividing the image into a grid of small, rectangular cells of size $S \times S$ where $S = \sqrt{\frac{N}{k}}$ where N is the number of pixels and k is the desired number of superpixels. The input to the SLIC algorithm is the number k. Each pixel is then assigned to one of the initial clusters which are sampled on a regular grid spaced S pixels apart. In the case of images in the CIELAB colour space, each pixel is allocated to a cluster based on its distance to the cluster centre, represented by the tuple $C_i = [l_i, a_i, b_i, x_i, y_i]$, where l_i is the lightness, a_i is the colour axis ranging from green to red, b_i is the colour axis ranging from blue to yellow, and x_i, y_i are the cluster centre coordinates. Each pixel is associated with its nearest cluster. Once all pixels have been assigned a cluster, the cluster centre is adjusted to the mean of the $[l, a, b, x, y]^T$ vector applied to all pixels belonging to the cluster. One notable advantage of SLIC over other clustering methods is that it only considers pixels in a small, local region (of size $2S \times 2S$) when forming clusters, which makes the algorithm faster to run.

2) DISF: The DISF algorithm takes two inputs: initial seed value s and the number of superpixels k. An image (I) is defined as a set of pixels D_I, and a function $I(p)$ that assigns local attributes to each pixel p in D_I. Seeds are placed at equal distances in a grid format, with distance $d = \sqrt{\frac{N}{k}}$, where N is a set of the neighbourhood nodes of pixel p_i. $N \subseteq D_I$ and k is the desired number of superpixels. Thus one can define an image as a graph $\mathcal{G}$ as (N, A, I) where A is the adjacency relation $A \subset N \times N$.

DISF leverages the Image Foresting Transform (IFT) [46] to achieve better superpixel segmentation.

At the core of the DISF algorithm, the IFT method involves the iterative construction of minimum spanning trees using the image as a graph with seeds as nodes. In each iteration, the superpixel assignments are refined by adjusting the pixel similarity measures. This is accomplished by minimizing the cost map across all paths in the tree. The iterations continue, progressively enhancing superpixel assignments until convergence is reached or a predefined maximum number of iterations is completed.

One of the most salient benefits of the DISF algorithm lies in its ability to surmount the limitations of conventional superpixel methods, particularly in terms of initial seed availability. By employing an oversampling strategy, DISF increases the initial seed count, improving the likelihood of including all relevant seeds and resulting in more representative superpixels. This, in turn, enhances the overall segmentation quality [7]. Additionally, DISF removes seeds that produce the smallest trees in homogeneous regions of the image, based on vectors
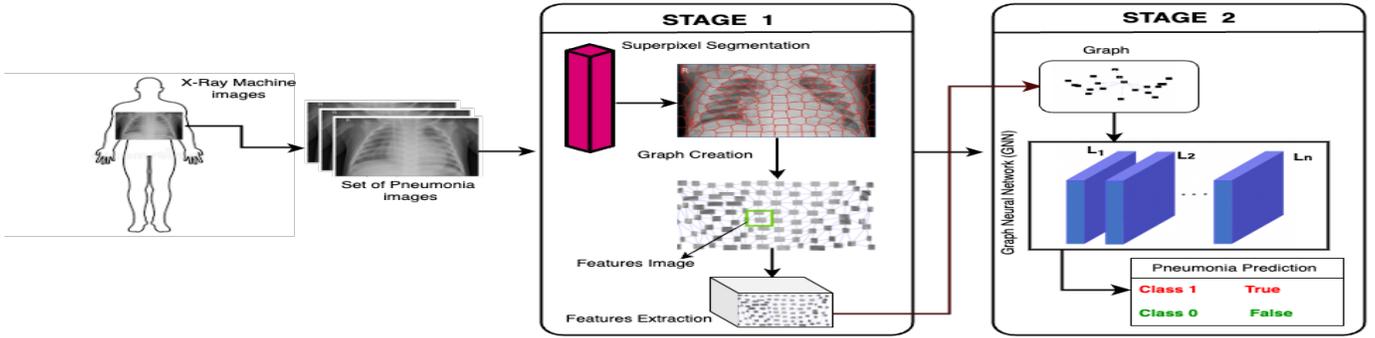
Fig. 1. Image to graph pipeline stage 1 and stage 2

representing each superpixel's properties and its neighboring superpixels. This ensures that the ultimate superpixels are more homogeneous and accurately represent the image's structure, culminating in superior segmentation outcomes.

B. Graph Creation

After segmenting an image into superpixels, we proceed to construct a Region Adjacency Graph (RAG). In this graph, each superpixel is represented as a node, and the nodes are connected to their neighboring nodes based on their spatial proximity within the image. These connections between the nodes are called edges.

To assign weights to these edges, we calculate the mean colour value of the pixels within each superpixel and then determine the differences in these mean values between adjacent superpixels. In order to augment the graph with additional information, we can incorporate features extracted from a DNN as node attributes II-C. These features can help to improve the graph's utility in subsequent image classification tasks.

Once the RAG has been generated and enhanced with the DNN extracted features, it is then saved as a pickle file, for easy retrieval and use in future classification tasks.

In the next sub-section, we will delve into the methods used for feature extraction.

C. Superpixel Feature Extraction

Once we have generated the superpixels of an image, we use pre-trained DNNs for feature extraction. We have chosen ResNet18, DenseNet121, and EfficientNetB0 to extract features from the superpixel segments. To achieve this, we first generate a new image by merging the segmented mask region with a bounding rectangle, which is constructed using the minimum and maximum pixel coordinates of the respective segment. We then resize the derived image to a standardized dimension, ensuring compatibility with the input requirements of the DNNs during the feature extraction phase.

We provide a detailed description of the employed DNNs for feature extraction below:

1) ResNet18: ResNet [17] Deep Convolutional Neural Networks make use of stacked layers (depth) to capture low to high-level features. The depth of the network is important

because it produces good results, but also introduces several challenges. One main problem that arises from stacking multiple layers is that of vanishing/exploding gradients [20], which is partially solved by normalized initialization. However, as the network gets deeper, the accuracy starts to decline and training loss starts to increase. These problems arise primarily because, of the identity mapping of the features being propagated from one layer to another. It becomes difficult to optimize performance as the model becomes deeper. ResNet adds a skip connection that uses an identity function to bypass the non-linear activation. The advantage of this architecture is the flow of gradient through the identity function from the later layer to earlier layers, but the summation operation between the identity function and output from a layer impedes the information flow in the network. ResNet mitigates this issue by using residue blocks in DNNs, which do not add additional parameters to the model but achieve better accuracy. ResNets have been successfully applied in the medical imaging domain and have produced state-of-the-art results. For example, ResNet has been applied for medical image retrieval [22] and to breast cancer classification from histopathological images [21].

2) DenseNet121: DenseNet, is a DNN architecture characterized by its unique connectivity pattern, which establishes connections between each layer and all other layers in a feed-forward manner. This dense connectivity ensures that the feature map of every preceding layer serves as input for the current layer, while its own feature map contributes as input to all subsequent layers. The distinctive structure of DenseNet enables each layer to receive direct inputs from all preceding layers. These inputs undergo concatenation, followed by a series of three consecutive operations. The process begins with batch normalization (BN), which is designed to improve the stability and convergence of the network. Next, the rectified linear unit (ReLU) is applied as an activation function to introduce non-linearity and enhance the model's capacity to learn complex patterns. Finally, a 3×3 convolution (Conv) operation is performed to extract features from the input data, as described in the DenseNet literature [18]. This combination of operations within DenseNet allows for efficient feature extraction and contributes to the overall effectiveness of our proposed framework.

3) EfficientNetB0: EfficientNet [19] is an innovative approach to scaling deep neural network architectures, tailored to specific use cases depending on data or resource constraints. The key advantage of EfficientNet lies in its ability to enhance model performance without increasing the number of floating-point operations per second (FLOPS), making it a more efficient choice compared to other DNNs.

The primary concept behind EfficientNet is the introduction of a method to uniformly scale all dimensions of a neural network, including depth, width, and resolution, by employing efficient compound coefficients. By scaling a model's computational resources by a factor of 2^N, EfficientNet adjusts the depth of the network by α^N, the width of the network by β^N, and the image size by γ^N. These coefficients, α, β, and γ, are determined through a grid search performed on the base model.

The compound scaling method employed by EfficientNet facilitates a well-balanced trade-off between the accuracy and efficiency of the networks, enabling users to optimize their models based on available resources. This ingenious approach is tailored such that the network architecture is adapted effectively to meet the specific needs of various use cases while maintaining computational efficiency.

D. Graph Neural Networks

For this study, we examined various types of spectral and spatial GNNs, including Graph Convolutional Neural Networks (GCNN), Graph Attention Networks, and Graph Isomorphism Networks. To address the limitations of individual models, we also tested combining all of these GNNs through ensemble methods. We provide a brief overview of each GNN in this section.

1) Graph Convolutional Neural Networks (GCNN): GC-NNs are a class of deep-learning models specifically designed for graph-structured data. They extend the concept of convolutional operations from the realm of grid-like data, such as images, to irregular domains represented by graphs. GCNNs can be broadly categorized into two approaches: spectral and spatial.

The spectral approach leverages the principles of graph signal processing and transforms the graph to the spectral domain by utilizing eigenvectors of the graph's Laplacian matrix. To simplify filters, a learned a learnable diagonal matrix, g_w is used, and various designs for the filter g_w have been proposed. GCNNs employ techniques to avoid computing Laplacian eigenvectors for efficiency [2].

In contrast, the spatial approach directly defines convolution on the graph based on its topology, using feature aggregation through message passing. The convolution is defined according to the number of hops allowed between a node and its neighbours [27], aiming to maintain local invariance similar to traditional CNNs, despite varying neighbour sizes.

2) Graph Attention Networks (GAT): The issue of over-smoothing has long plagued traditional GCNNs as the addition of multiple layers leads to simple graph embeddings for different class graphs. GAT addresses this problem by utilizing an attention mechanism to minimize over-smoothing.

GAT [13] makes use of attention-based convolution operation, which assigns different weights for neighbours to make the learning process more robust and stable, thus alleviating the noise effects. The attention-based mechanism allows for dealing with variable-sized input and making use of the most relevant part of the input to make inferences. GAT makes use of a self-attention layer and multi-head attention mechanisms [12].

The input to the attention layer is $\mathbf{h} = \left\{ \vec{h}_1, \vec{h}_2, \ldots, \vec{h}_N \right\}, \vec{h}_i \in \mathbb{R}^F$, where N is the number of nodes, and F is the number of features in each node which produces the output $\mathbf{h}' = \left\{ \vec{h}'_1, \vec{h}'_2, \ldots, \vec{h}'_N \right\}, \vec{h}'_i \in \mathbb{R}^{F'}$. The self-attention mechanism is applied to every node, where the attention coefficient is computed as

$$e_{ij} = a \left(\mathbf{W} \vec{h}_i, \mathbf{W} \vec{h}_j \right)$$

which absorbs the importance of node j 's features to node i, this operation is executed across all nodes, it results in the loss of the inherent graph structure. This issue is overcome by performing masked attention over nodes, e_{ij} for nodes $j \in \mathcal{N}_i$, where $\mathcal{N}_i$ is some neighbourhood of node i in the graph. For reason that first-order nodes $N(v) = \{u \in V \mid (u, v) \in E\}$ of a node can be of different sizes, and the coefficient is normalized across all choices of j using the softmax function

$$\alpha_{ij} = \text{softmax}_j \left(e_{ij} \right) = \frac{\exp \left(e_{ij} \right)}{\sum_{k \in \mathcal{N}_i} \exp \left(e_{ik} \right)}.$$

The normalized attention coefficients are used to compute the linear combination of features of all nodes. The aggregated features from each head are concatenated or averaged to obtain $\vec{h}'_1$, and finally a non-linear activation is applied:

$$\vec{h}'_i = \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \mathbf{W} \vec{h}_j \right).$$

3) GRAPH ISOMORPHISM NETWORK (GIN): GIN [14] was developed to test the power of GNNs, in comparison to the Weisfeiler-Lehman (WL) test of graph isomorphism [15] which is an algorithm used to determine if two graphs are isomorphic, meaning they have the same structure and patterns, but with different labels. The WL algorithm uses a recursive procedure to generate hashes which are used to find the structural similarity between two graphs. GIN makes use of the WL algorithm in continuous feature space. GIN aggregates node features using a Learnable Neighborhood Aggregation operator, which is an epsilon-multilayer perceptron (E-MLP) function. E-MLP is a unique variation of the traditional multi-layer perceptron (MLP) that incorporates a learnable epsilon (ϵ) parameter. This function accepts the node features and the features of its neighbouring nodes as input and adjusts the aggregation of these features using the learnable ϵ parameter in the MLP process. The E-MLP enables the model to learn

the optimal balance between a node's own features and the features of its neighbours for improved graph representation learning. This parameter controls the level of expressiveness of the MLP, allowing GIN to learn more complex features. The E-MLP adjusts the contribution of each feature based on the learned ϵ value. This enables the model to adaptively decide the importance of a node's own features relative to its neighbours' features, leading to better embeddings for graph classification.

$$h_v^{(k)} = \text{MLP}^{(k)} \left(\left(1 + \epsilon^{(k)} \right) \cdot h_v^{(k-1)} + \sum_{u \in \mathcal{N}(v)} h_u^{(k-1)} \right).$$

The output of MLP can be further used for various tasks like node classification, link prediction, or graph-level classification.

E. Ensemble Graph Neural Networks

Various trade-offs exist between network architecture and model performance in graph-based methods. For instance, GCNN performance may decline with added layers due to the over-smoothing of node representations, causing indistinguishable inter-class nodes. To address this issue, the jumping knowledge network is a proposed solution that selects features from more representative nodes.

GATs apply attention to all nodes, primarily focusing on immediate neighbours rather than the entire graph structure, which can result in suboptimal performance for tasks needing comprehensive graph understanding.

GINs have aggregation power equivalent to the Weisfeiler-Lehman test for distinguishing graphs. However, GIN embeddings may vary in quality for non-isomorphic graphs with similar structures and limited features.

Ensemble methods offer a solution to the individual limitations of various GNNs by utilizing a group of GNNs. Each GNN in the ensemble focuses on capturing distinct aspects of the graph's structure and information. By combining the outputs of these GNNs, a more comprehensive feature representation of the graph can be generated, which captures a wider range of information.

Ensemble methods can be a useful technique for combining the predictions of multiple GNNs to improve classification performance [16]. One common way to ensemble GNNs for classification is to concatenate the output of each GNN and use the concatenated output as the input to a final classifier. We have ensembled GCNN, GAT and GIN to get new graph embedding for classification.

We aim to leverage the flexibility of GNNs, which can be trained on graph data with varying numbers of nodes, allowing them to make accurate predictions on graphs with different node counts. This adaptability is particularly advantageous when dealing with graphs of arbitrary sizes generated through diverse methods, catering to a wide range of use cases.

In the following section, we will provide a detailed explanation of the experiments and results obtained from our analysis of the pneumonia image dataset. In addition, we will also conduct an ablation study to examine the performance of the different models used in this research.

III. EXPERIMENTS

In this section, we introduce the dataset that was used for training the models. We then describe the various experiments that were conducted for training different GNNs. Hyperparameter tuning was conducted using Optuna, resulting in a learning rate of 0.001, weight decay of 0.001, and a dropout rate of 0.5%

A. Datasets

The dataset used in our study [3] is organized into three distinct subsets: training, testing, and validation, each containing subfolders for the two image categories: Pneumonia and Normal. The dataset comprises a total of 5,856 X-Ray images in JPEG format, with 1,583 images in the Normal category and 4,273 images in the Pneumonia category. Out of the 5,856 images, 1,341 Normal images and 3,875 Pneumonia images are used for training and 234 Normal and 390 Pneumonia images were used for testing. Images were converted into graph representations, maintaining the original folder structure, and saved in pickle format for GNN training. Each experiment generated a different graph dataset.

B. Models

- GCN: a GCN architecture with 4 graph convolution layers was trained to utilize graph representationsII-B generated from images as input. Node feature extraction was performed using ResNet18, DenseNet121, and EfficientNet-B0 models.
- GAT: we trained a GAT model with a multi-headed attention mechanism. The architecture consisted of two layers of GAT convolution, with a dropout rate of 0.3% applied after each convolutional layer. The learning rate was set to 0.001. 8 attention heads were utilized in each layer, allowing the model to focus on different aspects of the input graph.
- GIN: we trained a GIN model with 3 GIN convolution layers to generate graph embeddings. The GIN convolution layers were designed to capture the underlying structure and relationships of the graph. Additionally, a global add pooling operation was used to aggregate neighborhood features and extract higher-level representations. The learning rate was set to 0.001, and a dropout rate of 0.5% was applied as a regularization technique to prevent overfitting. The resulting embeddings were concatenated and used for performing graph-level predictions.
- Ensemble Model: we combined the best performing pretrained GNNs from the above 3 models and ensemble their results to make the classification. The learning rate was set to 0.001 with a dropout rate of 0.5%.

The models were trained using graph datasets generated with superpixel values of 5, 10, 50, 100, 150, and 300. The features were obtained by utilizing Resnet18, EfficientNet-B0, and

DenseNet121, which generated features with sizes of 512, 1280, and 1024, respectively. As a result, we have established 18 benchmark datasets for model training and classification.

C. Classification

Our primary goal was to accurately predict labels. To evaluate the performance of various models, we utilized the accuracy metric. The performance of various models on different graph datasets can be seen in Table I, where we present the accuracies for each conducted experiment. Table II shows the performance of Ensemble models. We also analyzed the model sensitivity performance across the datasets, which were generated using varying numbers of superpixels and feature extraction methods.

We achieved an optimal performance of 93.108% in classifying pneumonia, and using DNNs, we achieved an accuracy of 92.81% IV . Our model parameters are 100 times fewer than state-of-the-art dense neural network models as shown in Table V, and training and inference are also faster in comparison. The 3 DNNs had an average train time of 155.702 seconds, while the 3 GNNs had an average train time of 42.759. Our results show that the model performance reached saturation after 20 epochs, and we recorded a higher sensitivity score with our pipeline compared to DNN models.

IV. ANALYSIS

In this section, we evaluate our graph dataset preparation method by examining alternative design options. We present the outcomes of alterations in different stages of the experiments, such as adding layers to GNN models, evaluating the quality and time consumption of two superpixel generation techniques, and comparing the feature extraction quality across three DNNs.

A. Effects of superpixel segment number on Model Performance

Our findings indicate that smaller image segments (a higher number of them) limit relational features, mainly due to the diffuse nature of pneumonia features in the dataset. As a result, DNNsII-C may miss out on fully capturing the image's context, adversely affecting classification accuracy, as reflected in Table I. Conversely, when segments are larger but fewer, their interconnectedness is better captured. This enables the CNN to capture the full context and relationships between the objects in the image, thereby enhancing classification performance.

B. Effects of Superpixel methods on Model Performance

The impact of selecting a method for superpixel creation on the performance of a GNN model is significant. All reported performance is based on the DISF method. Another significant contrast observed between the two methods was that SLIC had a shorter image segmentation time compared to DISF. This observation can be explained by the fact that DISF demands a higher initial seed value to construct superpixels. The time

taken for segmentation is a crucial factor as it is one of the bottlenecks in the process of converting an image to a graph. Table III contains a comparison of the time taken by each algorithm.

C. Effects of Feature extraction methods on Model Performance

DNN models like ResNet18, EfficientNet-b0, and DenseNet extract features of sizes 512, 1280, and 1024 respectively. Larger feature sizes increase sensitivity but demand more computation. While performance often correlates with feature size, this isn't consistent with graphs having edge weights. This suggests analyzing images as graphs can help models learn interconnected feature relationships.

Fine-tuning DNNs (ResNet18, Densenet121, EfficientNet-B0) for X-ray images improves GNN model performance. Table IV shows the DNN model's performance. Using the top model from training as a feature extractor has resulted in an average improvement of 2-3% for all GNN models compared to using an untrained pre-trained DNN model on the pneumonia dataset.

D. Effects of model complexity on Model Performance

For GCNN, performance plateaus after 4 convolutional layers, and adding more causes instability due to over-smoothing, resulting in poor graph embeddings.

For GAT, while performance doesn't notably improve beyond 8 attention heads, it does increase model complexity, leading to longer training and inference times.

For GIN, performance doesn't notably improve beyond 4 convolutional layers. The choice of graph-level pooling significantly influences performance, with global add pooling outperforming mean, minimum, and maximum aggregation methods.

V. CONCLUSION

In this study, we explored various techniques for converting medical images into graphs for classification purposes. We evaluated several GNNs and compared their performance with DNNs. The following steps outline our process:

1) Evaluating SLIC and DISF methodologies for image segmentation.
2) Constructing a graph from the image, with each superpixel segment representing a graph node.
3) Employing various DNNs to extract features and assign them as node attributes.
4) Training and testing three distinct GNN architectures.
5) Combining the top-performing GNNs from different architectures through ensembling.

Our results showed that GNNs performed well with 100 times fewer parameters than DNNs (refer to Table V and Table IV for comparison).

GNN models' independence from graph size enabled us to ensemble various models trained on different graph sizes, thereby improving sensitivity scores, a vital aspect in the medical field. The novelty of our work lies in the use of IFT

TABLE I
ACCURACIES OF GRAPH NEURAL NETWORK MODELS FOR EACH EXPERIMENT

Superpixels	GCNN			GAT			GIN		
	ResNet	EfficientNet	DenseNet	ResNet	EfficientNet	DenseNet	ResNet	EfficientNet	DenseNet
5	0.911	0.883	0.900	0.892	0.892	0.897	0.910	0.900	0.895
10	**0.927**	0.900	0.900	0.900	0.907	0.916	**0.931**	0.921	0.905
50	0.924	0.908	0.892	**0.929**	0.897	0.875	0.900	0.911	0.905
100	0.899	0.886	0.883	0.833	0.860	0.841	0.892	0.889	0.891
150	0.878	0.878	0.870	0.830	0.842	0.833	0.895	0.897	0.908
300	0.863	0.865	0.818	0.834	0.854	0.746	0.887	0.891	0.878

TABLE II
ENSEMBLE RESULT OBTAINED FROM THE GRAPH DATASET CREATED WITH 10 SUPERPIXELS

Feature Extraction method	Accuracy	AUC	Sensitivity
DenseNet121	0.899	0.872	0.979
EffiecientNet-b0	0.852	0.810	0.992
ResNet18	0.895	0.865	0.987

TABLE III
TIME TAKEN (IN SECONDS) BY SLIC AND DISF FOR SEGMENTATION

Superpixels	SLIC	DISF
5	0.0214	0.1043
10	0.0226	0.1007
50	0.0298	0.2280
100	0.0320	0.1975
150	0.0353	0.3717
300	0.0360	0.4833

TABLE IV
DNN PERFORMANCE ON PNEUMONIA DATASET

DNN	Accuracy	Parameters
densenet121	0.9281	7,978,856
effiecientnet-b0	0.8187	5,288,548
renset18	0.9125	11,689,512

TABLE V
TRAINABLE PARAMETERS IN GNNS FOR LARGEST FEATURE SIZE 1280

GNN	Parameters
GCNN	699426
GAT	67938
GIN	99266

to group like pixels together, creating nodes in a graph format. This method effectively spots spread-out irregularities, such as pneumonia, in the chest X-ray images.

In our upcoming work, we aim to test our process on a wide range of medical datasets. Given that our method achieved a 99.13% accuracy on the MNIST dataset, we believe it has the potential for success on other datasets as well.

ACKNOWLEDGMENT

REFERENCES

[1] Johnson, Alistair E. W. and Pollard, Tom J. and Mark, Roger G., MIMIC-III Clinical Database (version 1.4), 2016, https://doi.org/10.13026/C2XW26
[2] Kipf, Thomas N. and Welling, Max, Semi-Supervised Classification with Graph Convolutional Networks, https://arxiv.org/abs/1609.02907, DOI: 10.48550/ARXIV.1609.02907
[3] Kermany, Daniel Zhang, Kang Goldbaum, Michael, Labeled Optical Coherence Tomography (OCT) and Chest X-Ray Images for Classification, 2018, https://data.mendeley.com/datasets/rscbjbr9sj/2, DOI: 10.17632/rscbjbr9sj.2
[4] Chengsheng Mao, Liang Yao, Yuan Luo, ImageGCN: Multi-Relational Image Graph Convolutional Networks for Disease Identification with Chest X-rays, https://arxiv.org/pdf/1904.00325v1.pdf
[5] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine S¨usstrunk, SLIC Superpixels, EPFL Technical, Report, 2010
[6] Belem, F.C., Guimaraes, S.J.F., Falcao, A.X., Superpixel Segmentation Using Dynamic and Iterative Spanning Forest, IEEE Signal Processing Letters, 2020, 27, 1440-1444, https://doi.org/10.1109/lsp.2020.3015433

[7] Vargas-Munoz, J.E., Chowdhury, A.S., Alexandre, E.B., Galvao, F.L., Miranda, P.A.V., Falcao, A.X., An Iterative Spanning Forest Framework for Superpixel Segmentation, IEEE Transactions on Image Processing, 2019, https://doi.org/10.1109/tip.2019.2897941

[8] Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C., Shpanskaya, K., Lungren, M.P., Ng, A.Y., CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning, 2017, https://arxiv.org/abs/1711.05225, DOI: 10.48550/ARXIV.1711.05225

[9] Bruna, J., Zaremba, W., Szlam, A., LeCun, Y., Spectral Networks and Locally Connected Networks on Graphs, 2013, https://arxiv.org/abs/1312.6203, DOI: 10.48550/ARXIV.1312.6203

[10] Dong, Y., Liu, Q., Du, B., Zhang, L., Weighted Feature Fusion of CNN and GAN for Hyperspectral Image Classification, IEEE Transactions on Image Processing, 2022, 31, 1559-1572, DOI: 10.1109/TIP.2022.3144017

[11] S.K. Zhou, H. Greenspan, C. Davatzikos, J.S. Duncan, B.V. Ginneken, A. Madabhushi, J.L. Prince, D. Rueckert, R.M. Summers, A Review of Deep Learning in Medical Imaging: Imaging Traits, Tech Trends, Case Studies, and Future Promises, Proceedings of the IEEE, 2021, https://doi.org/10.1109/jproc.2021.3054390

[12] Vaswani, Ashish, Shazeer, Noam, Parmar, Niki, Uszkoreit, Jakob, Jones, Llion, Gomez, Aidan N., Kaiser, Lukasz, Polosukhin, Illia, Attention Is All You Need, arXiv, 2017, https://arxiv.org/abs/1706.03762, DOI: 10.48550/ARXIV.1706.03762

[13] Veličković, Petar, Cucurull, Guillem, Casanova, Arantxa, Romero, Adriana, Liò, Pietro, Bengio, Yoshua, Graph Attention Networks, arXiv, 2017, https://arxiv.org/abs/1710.10903, DOI: 10.48550/ARXIV.1710.10903

[14] Keyulu Xu, Weihua Hu, Jure Leskovec, Stefanie Jegelka, How Powerful are Graph Neural Networks?, arXiv:1810.00826, 2019

[15] B.YU. WEISFEILER AND A.A. LEMAN, THE REDUCTION OF A GRAPH TO CANONICAL FORM AND THE ALGEBRA WHICH APPEARS THEREIN

[16] Kosasih, E.E., Cabezas, J., Sumba, X., Bielak, P., Tagowski, K., Idanwekhai, K., Tjandra, B.A., Jamasb, A.R., On Graph Neural Network Ensembles for Large-Scale Molecular Property Prediction, arXiv, 2021, https://arxiv.org/abs/2106.15529, DOI: 10.48550/ARXIV.2106.15529

[17] He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, Sun, Jian, Deep Residual Learning for Image Recognition, arXiv, 2015, https://arxiv.org/abs/1512.03385, DOI: 10.48550/ARXIV.1512.03385

[18] Huang, Gao, Liu, Zhuang, van der Maaten, Laurens, Weinberger, Kilian Q., Densely Connected Convolutional Networks, arXiv, 2016, https://arxiv.org/abs/1608.06993, DOI: 10.48550/ARXIV.1608.06993

[19] Tan, Mingxing, Le, Quoc V., EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks, arXiv, 2019, https://arxiv.org/abs/1905.11946, DOI: 10.48550/ARXIV.1905.11946

[20] Philipp, G., Song, D., Carbonell, J.G., The exploding gradient problem demystified, arXiv, 2017, https://arxiv.org/abs/1712.05577, DOI: 10.48550/ARXIV.1712.05577

[21] Alom, M.Z., Yakopcic, C., Taha, T.M., Asari, V.K., Breast Cancer Classification with Inception Recurrent Residual CNN, arXiv, 2018, https://arxiv.org/abs/1811.04241, DOI: 10.48550/ARXIV.1811.04241

[22] Ayyachamy, S., Alex, V., Khened, M., Krishnamurthi, G., Medical image retrieval using Resnet-18, Medical imaging 2019: imaging informatics for healthcare, SPIE

[23] Lee, H., Park, H., Yoon, K., Better Generalization in Multi-Module GNNs, arXiv, 2022, https://arxiv.org/abs/2209.06589, DOI: 10.48550/ARXIV.2209.06589

[24] Shuman, David I. and Narang, Sunil K. and Frossard, Pascal and Ortega, Antonio and Vandergheynst, Pierre, The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains, 2013, https://doi.org/10.1109/MSP.2012.2235192

[25] Defferrard, Michaël and Bresson, Xavier and Vandergheynst, Pierre, Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering, 2016, https://arxiv.org/abs/1606.09375

[26] Bronstein, Michael M. and Bruna, Joan and LeCun, Yann and Szlam, Arthur and Vandergheynst, Pierre, Geometric Deep Learning: Going beyond Euclidean data, 2017, https://doi.org/10.1109/MSP.2017.2693418

[27] Hamilton, William L. and Ying, Rex and Leskovec, Jure, Inductive Representation Learning on Large Graphs, 2017, https://arxiv.org/abs/1706.02216

[28] Lecun, Y. and Bottou, L. and Bengio, Y. and Haffner, P., Gradient-based learning applied to document recognition, 1998, https://doi.org/10.1109/5.726791

[29] Sultana, Farhana and Sufian, Abu and Dutta, Paramartha, Evolution of Image Segmentation using Deep Convolutional Neural Network: A Survey, 2020, https://doi.org/10.1016/j.knosys.2020.106062

[30] Gilmer, Justin and Schoenholz, Samuel S. and Riley, Patrick F. and Vinyals, Oriol and Dahl, George E., Neural Message Passing for Quantum Chemistry, 2017, https://arxiv.org/abs/1704.01212

[31] Bodnar, Cristian and Di Giovanni, Francesco and Chamberlain, Benjamin Paul and Liò, Pietro and Bronstein, Michael M., Neural Sheaf Diffusion: A Topological Perspective on Heterophily and Oversmoothing in GNNs, 2022, https://arxiv.org/abs/2202.04579

[32] Sze, Vivienne and Chen, Yu-Hsin and Yang, Tien-Ju and Emer, Joel, Efficient Processing of Deep Neural Networks: A Tutorial and Survey, 2017, https://arxiv.org/abs/1703.09039

[33] Ding, Kexin and Zhou, Mu and Wang, Zichen and Liu, Qiao and Arnold, Corey W. and Zhang, Shaoting and Metaxas, Dimitri N., Graph Convolutional Networks for Multi-modality Medical Imaging: Methods, Architectures, and Clinical Applications, 2022, https://arxiv.org/abs/2202.08916

[34] Liang, Zhijun and Rojas, Juan and Liu, Junfa and Guan, Yisheng, Visual-Semantic Graph Attention Networks for Human-Object Interaction Detection, 2021, https://arxiv.org/abs/2001.02302

[35] Ba, Lei Jimmy and Caruana, Rich, Do Deep Nets Really Need to be Deep?, 2014, https://arxiv.org/abs/1312.6184

[36] Keicher, M., Burwinkel, H., Bani-Harouni, D., Paschali, M., Czempiel, T., Burian, E., Makowski, M.R., Braren, R., Navab, N., Wendler, T., U-GAT: Multimodal Graph Attention Network for COVID-19 Outcome Prediction, 2021, https://arxiv.org/abs/2108.00860

[37] Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M., ChestX-Ray8: Hospital-Scale Chest X-Ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases, 2017, https://doi.org/10.1109/CVPR.2017.369

[38] Johnson, Alistair E. W. and Pollard, Tom J. and Greenbaum, Nathaniel R. and Lungren, Matthew P. and Deng, Chih-ying and Peng, Yifan and Lu, Zhiyong and Mark, Roger G. and Berkowitz, Seth J. and Horng, Steven, MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs, 2019, https://arxiv.org/abs/1901.07042

[39] Ehteshami Bejnordi B, Balkenhol M, Litjens G, Holland R, Bult P, Karssemeijer N, van der Laak JA, Automated Detection of DCIS in Whole-Slide H&E Stained Breast Histopathology Images, 2016, https://doi.org/10.1109/TMI.2016.2550620

[40] Ahmedt-Aristizabal, David and Armin, Mohammad Ali and Denman, Simon and Fookes, Clinton and Petersson, Lars, Graph-Based Deep Learning for Medical Diagnosis and Analysis: Past, Present and Future, Sensors, 2021, https://doi.org/10.3390/s21144758

[41] Shen, Yiqing and Zhou, Bingxin and Xiong, Xinye and Gao, Ruitian and Wang, Yu Guang, How GNNs Facilitate CNNs in Mining Geometric Information from Large-Scale Medical Images, 2022, https://arxiv.org/abs/2206.07599

[42] Galon J, Costes A, Sanchez-Cabo F, Kirilovsky A, Mlecnik B, Lagorce-Pagès C, Tosolini M, Camus M, Berger A, Wind P, Zinzindohoué F, Bruneval P, Cugnenc PH, Trajanoski Z, Fridman WH, Pagès F, Type, density, and location of immune cells within human colorectal tumors predict clinical outcome, 2006, https://doi.org/10.1126/science.1129139

[43] Feichtenbeiner A, Haas M, Büttner M, Grabenbauer GG, Fietkau R, Distel LV., Critical role of spatial interaction between CD8 and Foxp3 cells in human gastric cancer: the distance matters. Cancer Immunol Immunother, 2014, https://doi.org/10.1007/s00262-013-1491-x

[44] Bronstein, Michael M. and Bruna, Joan and LeCun, Yann and Szlam, Arthur and Vandergheynst, Pierre, Geometric Deep Learning: Going beyond Euclidean data, IEEE Signal Processing Magazine, 2017, https://doi.org/10.1109/msp.2017.2693418

[45] Chen, Bingzhi and Li, Jinxing and Lu, Guangming and Yu, Hongbing and Zhang, David, Label Co-Occurrence Learning With Graph Convolutional Networks for Multi-Label Chest X-Ray Image Classification, IEEE Journal of Biomedical and Health Informatics, https://doi.org/10.1109/JBHI.2020.2967084

[46] Falcão AX, Stolfi J, de Alencar Lotufo R, The image foresting transform: theory, algorithms, and applications, 2004, https://doi.org/10.1109/tpami.2004.1261076