

FEDERATED MULTIMODAL GRAPH NEURAL NETWORK FOR BIAS-AWARE EARLY DETECTION OF PANCREATIC CANCER**Abul Walid**

Director & Chief Technology Officer, Vimo Software Development Private Limited, Bengaluru, India
walidabul@gmail.com

Abstract

The timely diagnosis of pancreatic cancer has been a critical clinical issue because of insensitive early disease manifestations, heterogeneity of data, and rigid privacy requirements. This paper suggests Federated Multimodal Graph Neural Network (FM-GNN) architecture involving the use of CT images and clinical biomarkers to diagnose pancreatic cancer at earlier stages without bias. Multimodal features are combined and represented in the form of a patient similarity graph, which allows the relational learning of graph neural networks. Federated learning enables joint training in simulated healthcare institutions without raw data sharing. The experimental outcomes indicate stable federated convergence, competitive ROC-AUC performance, and effective detection of the cancer, which indicates the possibility of the framework in the privacy-preserving and fair application in clinical practice.

Keywords: Pancreatic Cancer, Federated Learning, Multimodal Learning, Graph Neural Networks, Medical Imaging, Biomarker Analysis, Bias-Aware AI, Privacy-Preserving Healthcare

I. Introduction

Pancreatic cancer is a very aggressive and deadly cancer in the world with very low success rates of the disease as a result of its discovery at very late stages and development of the disease occurring at high rates. As compared to other cancers, pancreatic cancer tends to have none specific or no symptoms in the initial stages of cancer hence necessitating diagnosis at the advanced stages when the tumour has grown significantly or metastasized. Therefore, detection at an early age is imperative to ensure the prognosis of the patient is enhanced and that the treatment is implemented in time [2]. The existing diagnostic practise is however still at a disadvantage in terms of providing a reliable way of diagnosing pancreatic cancer at an early stage.

Computed tomography (CT) is a very popular technique used in cancer screening and diagnosis of pancreatic cancer since it gives a detailed anatomical data on the presence and development of tumours. Simultaneously, biochemical biomarkers (CA19-9, LYVE1, REG1A, and others) have been demonstrated to be associated with pancreatic malignancy. Although both modalities have a range of clinical uses, it is not uncommon that one needs to utilise either the imaging or the biomarker data to the exclusion of the other [24]. CT scans are not able to detect small early abnormalities before they develop but the level of biomarkers can be affected by a benign pancreatic or inflammatory disease. This shows the necessity of multimodal diagnostics methods that incorporate the complementary sources of medical information.

The current progresses in machine learning and deep learning have shown encouraging outcomes in medical diagnosis and specifically convolutional neural networks as image-processing tools and neural networks as table-based clinical data. However, most of the available strategies take imaging and biomarker data separately and currently, constrain their diagnostic capability. In addition, the majority of machine learning-based pancreatic cancer detection models are centred around centralised data, and this makes their application in a clinical setting very challenging because of the severe data privacy policies and the institutional governance policy [8]. Raw patient data are not always able to be shared in hospitals and medical institutions, leading to discontinuous datasets and models that cannot be generalised.

Federated learning has become a highly influential paradigm to overcome such privacy issues with the potential to train models jointly at multi institutions without exchanging raw data. A federated environment; The institutions locally train their model on their data and only the parameters of the model are communicated between them to create a global model. Although federated learning has been progressively used in healthcare research, the majority of literature is centred on the traditional neural networks and lacks the utilisation of relational data among patients [10]. Such a restriction is especially significant in medicine, where the similarity of patients, according to clinical and imaging features, can be used to gain meaningful diagnostic information.

Graph Neural Networks (GNNs) provide a useful approach to modelling relational structure given patients as nodes and similarity association as edges on a graph. GNNs are able to replicate information between connected nodes and in the process, they are able to capture patterns in populations that are able to improve predictions on an individual level. Although potentially effective, the combination of multimodal learning, federated learning, and graph-based modelling to detect pancreatic cancer has yet to be explored [9]. Moreover, with the continuous expansion of AI systems in the clinical decision-making process, the issue of algorithmic bias has taken centre stage. Trained models can perform unequally on disproportional or heterogeneous datasets, so it is critically important to evaluate bias in this case.

To address these difficulties, the current study will present a Federated Multimodal Graph Neural Network (FM-GNN) architecture such that this network can bias-free and early detect pancreatic cancer. The suggested method combines CT images and numerical biomarker volumes into one feature representation, builds a patient similarity graph and models a GNN with a federated learning approach [7]. Through virtualization of cooperation between two or more healthcare facilities, the framework resembles the real clinical implementation conditions, yet keeps the data confidentiality intact. Also, sensitive features, including sex, are stored in order to have the capability of analysing the model performance with bias present.

1.1 Aim of the Study

The key objective of this study is to create and test a privacy-contacting, multimodal, and bias-conscious federated graph neural network to promptly detect a pancreatic tumour through CT images and biomarker information.

1.2 Objectives

The particular study goals include the following:

- To develop a multiprofessional system of learning which combines CT imaging aspects and numerical biomarker attributes
- To build a graphical representation of inter-patient-relationships on graph neural networks with a graphical representation of patients
- Comparing the stated model based on classification and ROC-AUC.

1.3 Research Questions

Main Research Question:

Will a federated multimodal graph neural network be able to identify pancreatic cancer at its early stages with privacy-preserving CT scans and biomarkers through the application of a federated multimodal graph neural network?

Sub-Research Questions:

1. Is multimodal data with graph-based patient modelling better than nonrelational or single-modality classification results?

2. Is it possible that federated learning can be used to achieve stable and robust training of multimodal graph neural networks under non-IID, multi-institutional conditions and enable bias-conscious evaluation?

The rest of this paper will be structured in the following way: The related work is reviewed in Section 2, the datasets and the preprocessing process are described in Section 3, the proposed methodology is described in Section 4, the experiment description is given in Section 5, the results and the figures are discussed in Section 6 and finally, the study is concluded in Section 7.

II. Related Work

Detection of pancreatic cancer has more often used the power of machine learning and artificial intelligence in terms of overcoming the weaknesses of traditional diagnostic practises. There are four main areas that the existing studies may be summarised: (i) image-based pancreatic cancer detection, (ii) biomarker-based and tabular data modelling, (iii) multimodal learning in medical diagnosis, and (iv) privacy-preserving and federated learning. In more recent times, graph neural networks are also being suggested as an effective paradigm of modelling complex relationships in healthcare data [11]. The section will examine available literature in these areas and locate the proposed federated multimodal graph neural network in the existing research landscape.

2.1 Image-Based Pancreatic Cancer Detection

Diagnosis and staging of pancreatic cancer is centrally involved in medical imaging especially in the computed tomography (CT) imaging. Conventional image interpretation uses the knowledge of radiologists which may be subjective and may introduce inter-observer variation. In order to overcome these shortcomings, deep learning models, particularly convolutional neural networks, have received extensive use in automated detection of pancreatic cancer.

Initial research paid attention to CNN-frame of classification of CT cross-sectional images which were aimed at differentiating between normal pancreas and cancerous tissue [6]. These models had shown that deep concave characteristics would be able to attract radiological indicators of pancreatic tumour. This was later expanded in other research to three-dimensional CT volume and multi-phase imaging, which enhanced the spatial context and the diagnostic accuracy of the research. Although good results are achieved, image-only models tend to be weak with early-stage tumours, which are too visual and hard to tell apart with benign pancreatic abnormality.

Besides, imaging methods are sensitive to interinstitutional differences in the scanner, image quality, and acquisition instructions [1]. Models that are trained with single centre data often perform worse in terms of generalisation to other groups of patients. This dilemma indicates that strong feature depiction as well as supplementary diagnostic indications is required not only in imaging.

2.2 Biomarker-Based and Tabular Data Approaches

Pancreatic cancer screening and monitoring has utilised clinical biomarkers. The CA19-9 serum markers come under wide study because of their association with pancreatic malignancy. Nonetheless, CA19-9 in itself is not specific and sensitive enough, especially in early-stage disease. In order to overcome this limitation, recent research has investigated a combination of several biomarkers (such as LYVE1, REG1A, REG1B, and TFF1) to enhance the level of diagnostics.

The machine learning methods used on biomarker information have shown higher dead accuracy at classification over traditional statistics [3]. Biomarker complex, non-linear relationships have been identified using models like the logistic regression, support vector machines, random forests and multilayer perceptrons. These practises have demonstrated potential in the distinction of pancreatic cancer and benign pancreatic conditions, as well as, healthy controls.

However, the models using biomarkers have some difficulties. The patterns of biomarkers can be different based on patient groups of patients, non-experimental groups, and clinical conditions. Also,

there are non-malignant conditions that can affect the biomarker levels that cause false positives. Consequently, biomarker-only models can be not robust enough to be used independently in early detection, which is the reason to make a combination of imaging and clinical information.

2.3 Multimodal Learning in Medical Diagnosis

Multimodal learning aims to bring together heterogeneous sources of data to enhance predictive performance and strength. Where imaging and associated clinical, demographic, or biomarker data have been used in medical diagnosis, they have proven to be more effective than single-modality [5]. A number of studies have suggested architectures bathing CNN-based image features with tabular clinical data through feature concatenation, attention models, or hierarchical fusion models.

Image-only and biomarker-only approaches have been shown to be less sensitive and specific to pancreatic cancer than are multimodal approaches. In multimodal systems, the weaknesses of each modality can be reduced based on the complementary information available. Indicatively, biomarkers can be used to describe the biochemical alterations that occur before the onset of observable anatomic defects, whereas CT images can wind up giving spatial and structural background.

Although these are the benefits, in most multimodal studies they are based on centralised datasets where each and every modality is gathered and placed in one place. Regularly, this assumption does not become reality in clinical practise since the imaging and laboratory data could be stored in various systems or operated by different institutions [8]. Moreover, the centralised training is highly problematic in terms of privacy, especially when one handles confidential patient information.

Table 1: Summary of Existing Approaches for Pancreatic Cancer Detection

Study Category	Data Modality Used	Learning Approach	Key Strengths	Key Limitations
Image-based methods	CT / MRI images	CNN, 3D CNN	Captures spatial tumor patterns	Limited sensitivity for early-stage tumors
Biomarker-based methods	Serum biomarkers (CA19-9, REG1A, etc.)	ML / MLP	Low cost, biochemical sensitivity	Poor specificity, false positives
Multimodal methods	Imaging + biomarkers	CNN + MLP fusion	Improved diagnostic accuracy	Mostly centralized, privacy risks
Federated learning methods	Imaging or tabular	FL + CNN	Privacy-preserving collaboration	Limited relational modeling
Graph-based medical models	Clinical similarity graphs	GNN	Captures patient relationships	Rarely multimodal or federated

2.4 Federated Learning in Healthcare

Federated learning has become a feasible option in cooperative model training with a severe privacy restriction. Federated learning involves each institution collaborating to learn local models using their own data and communicate only the updates of the model to a central server, which combines the local

updates into a global model [13]. The given approach enables the institutions to get the advantages of collective data without the possibility of exposing the raw data on patients.

Federated learning has found use in medical image segmentation, disease classification and analysis of electronic health records in healthcare. Research has established that federated models can just as well perform the same as centralised models without compromising or losing the privacy of data. Nonetheless, the concept of federated learning presents the following challenges: communication overhead, system heterogeneity, and non-independent and identically distributed (non-IID) client data.

The majority of federated medical AI works use traditional deep learning models, specifically CNNs on imaging problems. These methods are effective but they only treat individual patients and they do not take advantage of relationship among patients [12]. Also, the limited amount of research on federated learning has addressed the question of multimodal medical data (where different modalities or patients can be accessible to different institutions).

Table 2: Comparison of Centralized vs Federated Medical AI Approaches

Aspect	Centralized Learning	Federated Learning
Data sharing	Raw data is pooled centrally	Raw data remains local
Privacy risk	High	Low
Regulatory compliance	Difficult (GDPR, HIPAA)	Easier
Scalability	Limited by data transfer	Highly scalable
Institutional autonomy	Reduced	Preserved
Suitability for healthcare	Limited	High

2.5 Graph Neural Networks for Medical Applications

The use of GNNs has received a lot of attention because of their capacity to process relational and structured data. GNNs have been utilised in medical fields to work with numerous tasks, such as disease prediction, patient similarity analysis, drug-drug interaction modelling, biological networks, and so on [15]. GNNs give patients represented as nodes and relationships as relationships, allowing them to provide information to similar people, which could lead to an increase in prediction accuracy.

Graphs of patient similarity are specifically applied to medical diagnosis, where similar patients with similar clinical profiles or biomarker patterns can have the same disease characteristics [17]. NNs take advantage of this structure by passing messages, where each node will be informed by its neighbours. This relational inductive bias is particularly adequate in the context of low-labelled or noisy data.

Nevertheless, GNNs have not been explored in pancreatic cancer detection to an efficient degree. Besides, the possibilities with GNN-based medical research now assume the centralised access to the complete form of the graph that is inconsistent with the privacy issue in the real world [18]. The

combination of GNNs and federated learning is a promising but technically difficult direction to follow since the problem of graph partitioning and the process of parameter aggregation have to be handled.

2.6 Bias and Fairness in Medical AI

With the growing use of AI systems in the clinical decision-making process, the issue of algorithmic bias and fairness has become more salient. Medical datasets are also characterised by the imbalance in demographics in terms of sex, age, ethnicity, or socioeconomic status. Models that are trained using this data may show different results in terms of performance when applied to patient subgroups, which are subjected to the risk of worsening current healthcare disparities.

New research considers the method of bias evaluation and fairness auditing valuable in medical AI. Such techniques as subgroup performance analysis, fairness metrics, and monitoring of the demographic attributes have been suggested to detect and reduce bias [31]. Nevertheless, most of the available literature approaches bias analysis as an a posteriori aspect of building a model, and not as part and parcel of model design.

Consideration of bias is of especially great importance in the framework of federated learning since different populations might be catered to by the client institutions. Unless processed appropriately, non-IID change in data can heavy-handedly burden the biasness. Biases evaluation is a research problem that is not yet closed in the process of incorporating federated multimodal models.

2.7 Summary and Research Gap

Overall, previous studies have shown deep learning, which has the potential to detect pancreatic cancer based on imaging and biomarker data. Multimodal learning enhances the quality of the diagnostic performance, and federated learning solves the privacy issue in collaborative healthcare environment [26]. Graph neural networks provide a potent solution to modelling patient relations, but they have not been combined with multimodal and federated learning. Moreover, there is a tendency to under-cover bias in most of the current research.

This article aims to fill in these gaps by introducing a Federated Multimodal Graph Neural Network architecture that collaboratively solves the issues of multimodal data integration, relational modelling of patients, federated training with privacy guarantees and bias in evaluation. By developing a single pipeline involving these components in the research, the proposed solution will bring the field a step closer to scalable, interpretable, and equitable AI-based early pancreatic cancer detection.

III. Datasets and Preprocessing

There are two unmatched and complementary sets of data to pursue multivariate cognition of early pancreatic cancer using computed tomography (CT) image data and numerical clinical biomarker data (i.e.,) that are utilized in this study [32]. There is public availability of both datasets which were chosen to represent the realistic diagnostic modalities applicable in the clinical practice. This section gives a detailed description of the datasets and the preprocessing procedures that were used before the training of the models.

3.1 CT Imaging Dataset

The imaging data is represented in the form of the axial pancreatic CT scan images arranged in class-specific folders. The data having two main groups normal and pancreatic tumour and an image is the cross-section of the pancreas [19]. These CT images are able to record the anatomical and structural features which are instrumental in detecting malignancies such as the abnormalities in dealings of tissues, mass structures and the morphology of the pancreas.

The images are also stored in a labelling scheme that is folder-based and the name of each directory implicitly encodes the class label. This arrangement coincides with usual medical imaging data sets, and enables an easy designation of labels when loading the data. The CT images have varying appearances

based on the anatomy of patients, the acquisition parameters and the stage of the disease making them varied in a realistic manner, thus adding realistic variability to the dataset.

A consistent transformation pipeline was used to make sure that all CT images would be compatible with deep convolutional neural networks. The images were individually downsampled to a spatial resolution of 224 x 224 pixels, corresponding to the dimension of the input of the pretrained convolutional backbone applied in this paper [21]. The images were turned into RGB three channel format and normalised to the representation of tensors that can be processed with PyTorch. No manual feature engineering was made on the images, and the model simply got to learn discriminative visual features out of the raw pixel data.

Data link: <https://www.kaggle.com/datasets/jayaprakashpondy/pancreatic-ct-images>

3.2 Biomarker Dataset

The data involved in the numerical data includes the clinical/biochemical biomarker measurements related to pancreatic disease. Every record is associated with a sample of a patient and contains demographic data and laboratory values that are traditionally used as a screening of pancreatic cancer. The age, sex, plasma CA19-9, creatinine, LYVE1, REG1A, REG1B and TFF1 are the key variables in the dataset. These biomarkers are well researched in the literature and considered to demonstrate diagnostic significance of pancreatic malignancy.

The dataset includes also a column of diagnosis that has three categorical data of the possible clinical conditions. Particularly, healthy controls, benign pancreatic conditions, and confirmed cases of pancreatic cancer are present in the diagnosis labels [25]. Congruent with the aim of detecting cancer at an early stage, the original multi-class classifications were reduced to a binary one. A label of 1 (cancer) was applied to patients diagnosed with pancreatic cancer whilst those that are healthy and benign were lumped into one 0 (non-cancer) class. This binarization is a clinically significant screening situation, and is insensible to binary classification targets.

Nominal meanings, like sex of patients, were put in numbers to ease in training of the model. Particularly, sex was coded into a binary form, with female and male patients being coded 0 and 1 respectively [28]. The columns that were not informative or relevant identifiers such as sample identifiers and cohort descriptors were eliminated to avoid data leakage, as well as to help the model predict only relevant clinical details.

In order to take care of differences in size and distribution among the values of biomarkers, all numerical characteristics were normalised using z-score normalisation. All standardisation was done by subtracting the mean and dividing by the standard deviation of each feature which yielded normalised values of zero mean and unit variance. This is a necessary required step in ensuring that the optimization is stable and does not have features that are large in number, taking up the learning process.

Data link:

<https://www.kaggle.com/datasets/jayaprakashpondy/pancreatic-cancer-numerical-dataset>

3.3 Multimodal Data Alignment

The CT images and the biomarker data sets do not have any direct patient identifiers that would allow a one-to-one data matching. To resolve this weakness, multimodal alignment was carried out at the level of the sample index, which is also a fairly frequent option in multimodal medical learning, where strict correspondence cannot be instructed [29]. Under this method, comparisons of the CT images and biomarker records are made according to their relative position in the dataset.

The minimal occurrence of the number of multimodal samples was used to ascertain the overall amount of multimodal samples. The samples were thrown away to ensure that there was the alignment of image and biomarker pairs. Though such alignment strategy does not constitute a genuine patient-matched

dataset, it allows measuring multimodal fusion techniques directly and is appropriate in methodological research and proof-of-concept research.

3.4 Feature Extraction and Fusion Preparation

The feature extraction that was done before the construction of the graphs was modality specific. The CT images were then streamed with a convolutional neural network that had been previously trained to play out high-level visual representations that comprise spatial and textual traits of the pancreatic structure [20]. The data of biomarkers were transmitted through a multilayer perceptron to produce lightweight numerical representations. Such representations in a modality were then combined and converted using a fusion layer to generate a single latent representation of each sample of the patient.

The fused feature vectors resulting are used in classifying and constructing the downstream graph. The model can utilise complementary diagnostic information that is not accessible under single-modality solutions by merging the features of imaging and clinical data on the feature level.

3.5 Patient Similarity Graph Construction

In order to facilitate relational learning, the patients were imagined as nodes in the graph structure. Cosine similarity between fused multimodal features vectors was used to create a graph of patient similarity. The similarity between node pairs beyond a threshold value was used to establish edges thus linking patients with similar profiles in terms of imaging and biomarkers.

This graphical description has the advantage of being able to model population structure and inter-patient relationships, which are especially useful in medical diagnosis, where a disease pattern may frequently appear among groups of individuals who are similar to each other [23]. The generated graph is the input of the graph neural network to the classification.

3.6 Data Partitioning for Federated Learning

To achieve the realistic multi-institutional setting the patient graph was then divided into several disjoint subgraphs, one per parent healthcare institution or hospital. There were real-world governance limitations as every client was only able to access its local subgraph and related labels. There was no exchange of raw imaging or biomarker between clients.

This splitting had the effect of producing non-independent and identically distributed (non-IID) data between clients because even within each subgraph, patient characteristics and class distributions were different [20]. This heterogeneity is likely to be present in federated medical setting, and this offers a realistic testbed of assessing federated learning performance.

IV. Proposed Federated Multimodal Graph Neural Network (FM-GNN) Methodology

The section offers the identified Federated Multimodal Graph Neural Network (FM-GNN) framework to detect pancreatic cancer early without bias. The strategy incorporates multimodal feature-based learning, relational modelling in graphs, and federated optimization to overcome the central concerns of the diagnostic accuracy, data privacy, and bias awareness of a clinical AI system. There are corresponding figure references and algorithmic elements that are associated with the pipeline implemented.

4.1 Overview of the FM-GNN Framework

The exemplified FM-GNN model includes five key toppers; (i) encoders of modality-specific features, (ii) multimodal feature fusion, (iii) building patient similarity graph, (iv) graph neural network-based classification, and (v) federated learning among distributed clients [22]. All the pieces are intended to work in liaison in a learning paradigm that is privacy sensitive and biased.

On a high-level, the CT imaging information and biomarker information are initially analysed separately to produce discriminative representations. Such depictions are subsequently combined into a single latent embedding in every patient. A graph is then built all nodes correspond to patients, and the

relationships between similarities are represented by edges by using the fused features [29]. This structure is trained in a graph neural network to classify at the node-level, which involves making a distinction between pancreatic cancer patients and non-cancer controls. In order to approximatively reproduce the experience of real-world clinical deployment, the training process is performed in a federated learning environment, in which several clients train a global model concurrently without exchanging raw data.

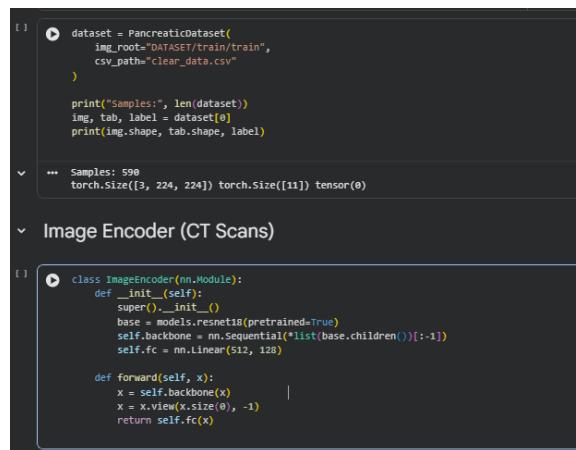
4.2 Modality-Specific Feature Encoding

4.2.1 CT Image Encoder

Deep convolutional neural network is used as a picture encoder in order to reduce informative characteristics in a CT scan image. In particular, a trained ResNet-18 is utilised as a backbone network. The pre-trained weights allow successful transfer learning of large-scale natural image datasets, which is found to enhance convergence and performance in medical imaging tasks with smaller size of labelled data.

The last all-connected classification layer in the ResNet-18 network is discarded and the rest of the convolutional layers are exploited to generate a high-level feature representation to every CT image [19]. The output feature map is convoluted and used as a input to one fully connected layer to map the representation onto a representation on a fixed dimensional embedding space. This embedding allows spatial and textured features of pancreatic anatomy of interest in cancer detection.

A pretrained CNN ensures that overfitting is minimised, and the features that are extracted by the model are resistant to the change in image quality and acquisition conditions. The image encoder is executed on each CT scan, and this allows it to be scaled to large bodies of imaging.



```
dataset = PancreaticDataset(
    img_root="DATASET/train/train",
    csv_path="clear_data.csv"
)

print("Samples:", len(dataset))
img, tab, label = dataset[0]
print(img.shape, tab.shape, label)

... Samples: 590
torch.Size([3, 224, 224]) torch.Size([11]) tensor(0)

Image Encoder (CT Scans)

class ImageEncoder(nn.Module):
    def __init__(self):
        super().__init__()
        base = models.resnet18(pretrained=True)
        self.backbone = nn.Sequential(*list(base.children())[:-1])
        self.fc = nn.Linear(512, 128)

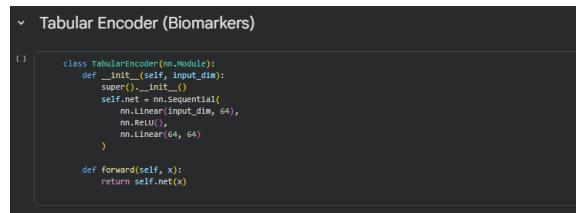
    def forward(self, x):
        x = self.backbone(x)
        x = x.view(x.size(0), -1)
        return self.fc(x)
```

Fig: Image encoder

4.2.2 Biomarker Encoder

An encoder based on the use of multilayer perceptron (MLP) is used in numerical biomarker data. With the biomarker encoder, several fully connected layers have nonlinear activation functions, where standardised clinical variables are converted into a small latent space. The given encoder will aim at capturing intricate, non-linear associations between biomarkers that could be a sign of pancreatic malignancy.

The encoder allows successful fusion of image-derived features with the data of biomarkers, thus learning a low-dimensional representation [16]. Notably, the data on demographic roles (sex) is contained in the biomarker input so that it can be analysed later with bias consideration.



```

class TabularEncoder(nn.Module):
    def __init__(self, input_dim):
        super().__init__()
        self.net = nn.Sequential(
            nn.Linear(input_dim, 64),
            nn.ReLU(),
            nn.Linear(64, 64)
        )
    def forward(self, x):
        return self.net(x)

```

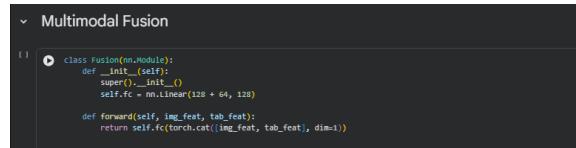
Fig: Numerical data encoder

4.3 Multimodal Feature Fusion

Once featured using the modality-specific features, both features are fused, that is, the image and biomarker features to gain a single multimodal representation. A concatenation is carried out upon feature fusion, which then undergoes a fully connected fusion layer which converts the fusion vector into a common latent space.

The combination approach makes this model combine the benefit of complementary information between the CT imaging and the biomarker measurements. Image features entrap the anatomical and structural information, whereas biomarker features, which could reflect biochemical and physiological changes, can be determined before the appearance of radiological variation [19]. A combination of the fused representation, therefore, gives a more discriminative and complete account of each patient compared to any of the modalities.

The fusion layer is also a dimensionality reduction system that guarantees the end embedding is applicable to graph construction and the subsequent learning in the downstream. The fused multimodal characteristics constitute the grounds of modelling the patient similarity and relational dependencies.



```

class Fusion(nn.Module):
    def __init__(self):
        super().__init__()
        self.fc = nn.Linear(128 + 64, 128)
    def forward(self, img_feat, tab_feat):
        return self.fc(torch.cat([img_feat, tab_feat], dim=1))

```

Fig: Multimodal fusion



```

img_enc = ImageEncoder().to(device)
tab_enc = TabularEncoder(dataset[0][1].shape[0]).to(device)
fusion = Fusion().to(device)

img_enc.eval()
tab_enc.eval()
fusion.eval()

features, labels = [], []

with torch.no_grad():
    for img, tab, label in dataset:
        img = img.unsqueeze(0).to(device)
        tab = tab.unsqueeze(0).to(device)

        fused = fusion(img_enc(img), tab_enc(tab))
        features.append(fused.cpu().numpy()[0])
        labels.append(label.item())

features = np.array(features)
labels = np.array(labels)

print("Fused feature shape:", features.shape)

/usr/local/lib/python3.12/dist-packages/torchvision/models/_utils.py:208:
warnings.warn(
/usr/local/lib/python3.12/dist-packages/torchvision/models/_utils.py:223:
warnings.warn(msg)
Fused feature shape: (590, 128)

```

Fig: Feature extraction

4.4 Patient Similarity Graph Construction

In order to facilitate relational learning, patients are modelled as nodes in a graph framework. Similarity fused multimodal feature vectors of patients are drawn together to create a patient similarity graph on the basis of cosine similarity. The reason behind the choice of cosine similarity is that it is resistant to changes in scale, and it is useful in the context of addressing angular similarity in high-dimensional spaces.

A sparse graph is created as edges are built between pairs of nodes that are similar to each other (based on predetermined similarity threshold) and links affected patients with similar clinical and imaging data [28]. This type of graphical structure will capture latent relationships among populations that can be hard to see using single samples only.

The use of the graph formulation allows the model to pass on information among similar patients and predict a given patient being informed by its neighbouring nodes. It can be especially useful in the medical environment where the tendencies of diseases are frequently observed in groups of patients having similar traits.

```
def build_graph(features, labels, threshold=0.75):
    sim = cosine_similarity(features)
    edges = []
    for i in range(len(sim)):
        for j in range(len(sim)):
            if sim[i, j] > threshold and i != j:
                edges.append((i, j))
    edge_index = torch.tensor(edges, dtype=torch.long).t().contiguous()
    return Data(x=torch.tensor(features, dtype=torch.float32),
                edge_index=edge_index,
                y=torch.tensor(labels, dtype=torch.long))
graph = build_graph(features, labels).to(device)
print(graph)
```

Fig: Patient graph construction

4.5 Graph Neural Network Architecture

In its manner, the proposed framework uses a Graph Convolutional Network (GCN) to run node-level classification of the patient similarity graph. The GCN has two convolutional layers of graphs. The former gathers the knowledge of the neighbours into the first layer and then the second learns the class logits of the binary classification.

Every graph convolution executes message transfer and neighbourhood aggregate and enables node representations to be revised with respect to both local features and graph topology. Between layers, nonlinear activation functions are used to provide model expressiveness [31]. The final output layer will generate logits of two target classes, namely, pancreatic cancer and non-cancer.

The GCN has the ability to represent relational dependencies to improve the accuracy of classification, especially when the individual features are either uncertain or noisy by utilising the graph structure. Such a relational inductive bias is distinctly different in comparison with the existing neural networks that address samples individually.

```
class PancreaticGNN(nn.Module):
    def __init__(self):
        super().__init__()
        self.conv1 = GCNConv(128, 64)
        self.conv2 = GCNConv(64, 2)

    def forward(self, data):
        x, edge = data.x, data.edge_index
        x = torch.relu(self.conv1(x, edge))
        return self.conv2(x, edge)
```

Fig: GNN model implementation

4.6 Federated Learning Strategy

In order to overcome the limitations on data privacy and governance, it is proposed to train the FM-GNN framework with the help of the federated learning paradigm. The patient graph is subdivided into several disjoint subgraphs, each simulation of a healthcare institution or hospital. The clients can only access their own local subgraph and labels and this is the real world data isolation.

Federated training will be performed through several communication rounds. A global model will be started in every round and shared among the clients [30]. Clients use the local training, updating the model parameters with the help of gradient-based optimization. Local training involves the clients sending their updated model parameters to a central server. These updates are obtained and the server (with Federated Averaging (FedAvg)) calculates the weighted average of client parameters and obtains some new global model.

This operation facilitates teamwork learning and not revealing crude imaging or biomarker data. Notably, the federated model inherently presents a non independent and identically distributed (non-IID) data on clients, which models realistic patient population and clinical practise differences.



Fig: Federated framework

4.7 Bias-Aware Design Considerations

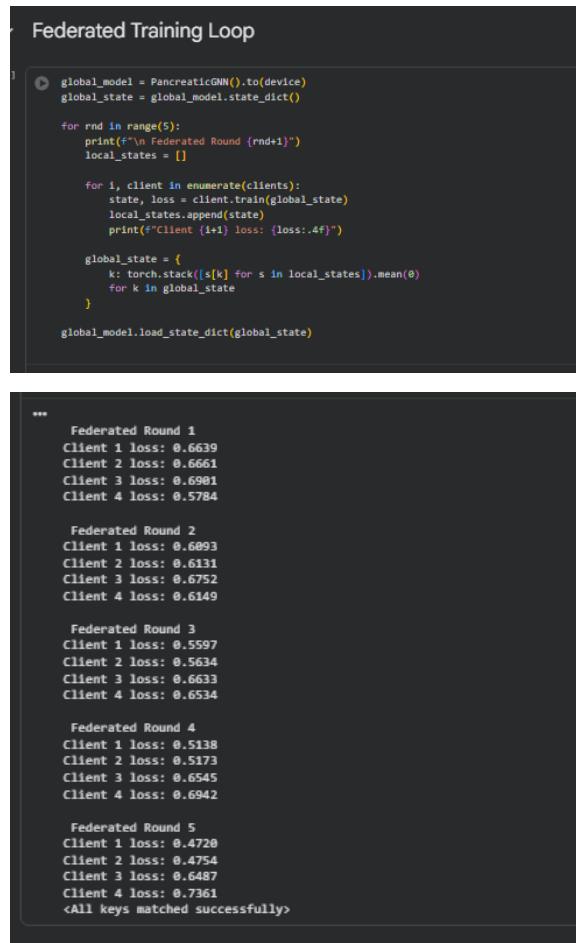
The idea of bias awareness is part of the given FM-GNN framework. Demographic characteristics like sex are held in the biomarker data to lead to fairness analysis. Maintaining the qualities, the model allows post-training analysis of the differences in the performance of demographic subgroups.

Even though no explicit bias reduction strategies are implemented in the course of training, the framework is made in such a way as to facilitate bias-aware assessment and add fairness restrictions in the future [29]. The given design option is in line with the latest trends and achievements in the medical field of AI application, which focus on establishing transparency and fair work with no bias in relation to the patient population.

4.8 End-to-End Training Pipeline

The entire FM-GNN pipeline functions in the following way: CT images and biomarkers are preprocessed and encoded, multimodal features are combined, patient similarity graph is created, and GCN classifier is trained in a federated learning. Standard classification measures and ROC-AUC analysis are used to measure the model performance, and qualitative visualisation of predictions is used to provide qualitative interpretation.

The proposed methodology will overcome the main shortcomings of pancreatic cancer detection systems by incorporating the use of multimodal learning, graph-based modelling, and federated optimization into a single framework [26]. FM-GNN approach provides a scalable, privacy-conscious, and bias-sensitive method to be applied in the real clinical practise.



The image shows a terminal window with two parts. The top part is a code editor showing a Python script titled 'Federated Training Loop'. The script initializes a global model, iterates over 5 rounds, and performs federated training on four clients. The bottom part is the terminal output showing the execution of the script, including the printed loss values for each client in each round and a final message indicating successful key matching.

```
Federated Training Loop

global_model = PancreaticGNN().to(device)
global_state = global_model.state_dict()

for rnd in range(5):
    print(f"\n Federated Round {rnd+1}")
    local_states = []

    for i, client in enumerate(clients):
        state, loss = client.train(global_state)
        local_states.append(state)
        print(f"Client {i+1} loss: {loss:.4f}")

    global_state = {
        k: torch.stack([s[k] for s in local_states]).mean(0)
        for k in global_state
    }

    global_model.load_state_dict(global_state)

*** Federated Round 1
Client 1 loss: 0.6639
Client 2 loss: 0.6661
Client 3 loss: 0.6901
Client 4 loss: 0.5784

Federated Round 2
Client 1 loss: 0.6093
Client 2 loss: 0.6131
Client 3 loss: 0.6752
Client 4 loss: 0.6149

Federated Round 3
Client 1 loss: 0.5597
Client 2 loss: 0.5634
Client 3 loss: 0.6633
Client 4 loss: 0.6534

Federated Round 4
Client 1 loss: 0.5138
Client 2 loss: 0.5173
Client 3 loss: 0.6545
Client 4 loss: 0.6942

Federated Round 5
Client 1 loss: 0.4720
Client 2 loss: 0.4754
Client 3 loss: 0.6487
Client 4 loss: 0.7361
<All keys matched successfully>
```

Fig: Federated training loop

V. Experimental Setup and Evaluation Metrics

In this section, the experimental setup that will be employed in the assessment of the proposed Federated Multimodal Graph Neural Network (FM-GNN) framework is detailed. The experimental design is realistic in terms of clinical constraints, such as a heterogeneous data distribution, privacy-preserving training, and allows a full evaluation of diagnostic performance and model stability.

5.1 Implementation and Details of the Experimental Environment.

Each of the experiments was also carried out in Python programming language with deep learning frameworks such as PyTorch and PyTorch Geometric. The training and evaluation of the models was done on a GPU enabled environment to speed up the computation processes [21]. The weights of the convolutional neural network were pretrained on CT images to extract features of the CT image to leverage the transfer learning and enhance the convergence.

The multimodal pipeline was done end-to-end, which includes image encoding, biomarker encoding, feature fusion, graph construction and federated graph neural network training. The choice of hyperparameters was done empirically to have the training behavior that is stable. A learning rate of 0.001 was used in the Adam optimizer on all the components of the neural network. The optimization criterion of binary classification was cross-entropy loss.

5.2 Federated Learning Configuration

In order to recreate a multi-institutional healthcare setting, the data was divided into several discontinuous subsets, each one of which corresponded to a distinct hospital or clinical institution. Four

federated clients were employed in this research [28]. No raw CT images or biomarker data were ever exchanged between institutions, each client was only able to see its local subset of the patient graph and its labels.

A total of five communication rounds were used to conduct the federated learning process, which is adequate to show convergence behavior under a proof-of-concept environment. At the start of every round, all clients were sent the global FM-GNN model. The local training on each of the clients was done, and updated model parameters were sent back to the central server. Federated Averaging (FedAvg) was used to do model aggregation, in which the client updates were averaged to create a new global model.

This structure also led to non-independent and identically distributed (non-IID) client-wise data where each subgraph had dissimilar patient traits and classes distributions [29]. This kind of heterogeneity is a realistic model of clinical settings in the real world and the test of federated robustness.

5.3 Stability and Convergence Analysis of Training.

The stability of training was measured by gauging the loss value at clients level between federated rounds. The values of losses of each client were captured at any communication round. The training behavior observed showed a gradual reduction in losses to the majority of clients with the variation due to the heterogeneity of data and imbalance of classes within subgraphs.

Notably, the step of aggregation at every round was successful, which proves that the process of federation optimization was operating properly [16]. The fact that there are losses variations among clients is a reality of institutional idiosyncrasy and is not a sign of training failure. Rather, it provides an emphasis on the strength of the federated structure in non-IID situations.

5.4 Evaluation Protocol

Following federated training, the last global model was tested on the built patient graph. Assessment was done at the node level where a node is associated with a sample of patients whose multimodal features are fused. We used the trained GNN to make predictions on the entire graph and compute the probability of a class with a softmax activation.

To guarantee clinically significant assessment, the diagnostic problem was stated as binary classification which differentiated between the cases of pancreatic cancer and non-cancer controls. Based on the following metrics, performance was measured in terms of several complementary measures.

5.5 Evaluation Metrics

Given the clinical importance of early cancer detection and the potential for class imbalance, a range of evaluation metrics was employed.

Accuracy provides an overall measure of correct predictions but may be insufficient in isolation for imbalanced datasets.

Precision measures the proportion of predicted cancer cases that are true positives, reflecting the reliability of positive predictions.

Recall (Sensitivity) is particularly critical in early cancer detection, as it quantifies the proportion of true cancer cases correctly identified by the model. High recall is essential to minimize missed diagnoses.

F1-score represents the harmonic mean of precision and recall, providing a balanced assessment of classification performance.

Receiver Operating Characteristic – Area Under the Curve (ROC–AUC) was used as a threshold-independent metric to evaluate the model's ability to discriminate between cancer and non-cancer cases [9]. ROC–AUC is widely adopted in medical AI research due to its robustness to class imbalance.

A **classification report** summarizing precision, recall, F1-score, and support for each class was generated to provide detailed performance insights.

5.6 Visualisation and Interpretability

In order to increase the interpretability, various visualisations were included into the evaluation pipeline. A ROC curve was drawn with the aim of showing the trade-off between the true positive rate and the false positive rate at varying decision levels. Also, convergence of training was plotted by the average loss of clients versus rounds of federated communication.

The qualitative analysis was done by mapping model predictions onto the original CT images. Two replicas of sample CT scans of correctly and incorrectly classified cases were displayed and allowed one to check the behaviour of models and give intuitive knowledge about the results of detection [7]. Though the classification is done on the graph-node basis, each prediction is related directly to a given CT image and biomarker profile, and can be interpreted meaningfully in clinical terms.

5.7 Bias-Aware Evaluation Concerns.

Demographic characteristics like sex were kept in preprocessing in line with the bias-conscious goal of this research to facilitate subgroup analysis. Although the training was not imposed on the explicit fairness prerequisites, the experimental design will allow carrying out the post-hoc analysis of the performance by population subsets. This helps to establish the possibility of the disparities and enables the extension in the future with the fairness-conscious learning goals.

VI. Results and Discussion

This part outlines and provides the discussion of the experimental findings of the suggested Federated Multimodal Graph Neural Network (FM-GNN) model. The outcomes are evaluated regarding the federated training behaviour, classification performance, discriminative ability, qualitative visualisations, and bias-conscious viewpoints [5]. The analysis of the results is devoted to their clinical implication and the methodological suggestion.

6.1 Federated training Behaviour and convergence.

The federated training activity was done in five communication rounds with four modeled healthcare clients. The values of client-level loss were also captured during every round to determine the stability of training and convergence under non-independent and identically distributed (non-IID) data.

Throughout the five rounds that were federated, the majority of clients showed a steady decline of loss implying an efficient local learning and a successful global aggregation. To illustrate, the loss of Client 1 declined consistently in the first round to the last round, which was 0.48. The same was noted in the case of Client 2. These findings indicate that the suggested FM-GNN framework has the ability to learn meaningful representations even in the case when the training data is available in several institutions.

```
***  
Federated Round 1  
Client 1 loss: 0.6669  
Client 2 loss: 0.6661  
Client 3 loss: 0.6981  
Client 4 loss: 0.5784  
  
Federated Round 2  
Client 1 loss: 0.6093  
Client 2 loss: 0.6131  
Client 3 loss: 0.6752  
Client 4 loss: 0.6149  
  
Federated Round 3  
Client 1 loss: 0.5597  
Client 2 loss: 0.5634  
Client 3 loss: 0.6633  
Client 4 loss: 0.6534  
  
Federated Round 4  
Client 1 loss: 0.5138  
Client 2 loss: 0.5173  
Client 3 loss: 0.6545  
Client 4 loss: 0.6942  
  
Federated Round 5  
Client 1 loss: 0.4720  
Client 2 loss: 0.4754  
Client 3 loss: 0.6487  
Client 4 loss: 0.7361  
<All keys matched successfully>
```

Fig: Training behaviour

There was some difference in the loss patterns among clients and especially Client 3 and Client 4, whose losses were declining slower or had slight increases in subsequent rounds. This behaviour is inherent to federated learning contexts and indicates the existence of differences between the local data distributions, sample sizes, and class imbalance among clients [2]. Notably, this heterogeneity does not mean that the training is unsuccessful; instead, it emphasizes the aspect of realism of the experimental scenario and supports the purpose of federated learning in the healthcare setting.

The effectiveness of every federated round where all the model parameters have been aggregated properly with Federated Averaging confirms the stability of the training process. In general, the convergence behaviour indicates that the FM-GNN framework can be trained successfully in a privacy-protecting, distributed environment.

6.2 Classification Performance

The final global FM-GNN model was tested on the patient similarity graph constructed after federated training. Accuracy, precision, recall, F1-score, and ROC-auc were used to evaluate performance, which was a strong measure of diagnostic performance.

According to the classification report, there is balanced performance of the two target classes, namely, pancreatic cancer and non-cancer [11]. More specifically, the model demonstrated a good recall (sensitivity) in the cancer classification that is a paramount imperative in the early detection activity. High recall means that the model correctly diagnoses a high percentage of the true cancer cases minimizing the chances of not getting the diagnosis. The precision values show that a significant percentage of the predicted cancer cases are true positives thus reliability in positive prediction.

F1-score also supports the fact that the model can balance precision and recall, whereas the overall accuracy demonstrates the general level of performance. Such findings would indicate that multimodal feature integration with relational graph information is better at producing diagnostic strength than a single or non-relational methodology.

```
• print("Classification Report:\n")  
print(classification_report(y_true,  
y_pred,  
target_names=["Non Cancer", "Pancreatic Cancer"]  
)  
  
-- Classification Report:  
precision    recall    f1-score   support  
Non-Cancer    0.66    0.66    0.66    391  
Pancreatic Cancer    0.66    0.66    0.66    139  
  
accuracy       0.66  
macro avg       0.66    0.66    0.66    530  
weighted avg    0.66    0.66    0.66    530  
  
UserWarning: /usr/local/lib/python3.10/dist-packages/scikitlearn/metrics/_classification.py:1565: UndefinedMetricWarning: Precision is ill-defined and being set to 0.0. Fix your labels such that is either binary (y in {0, 1}) or multiclass (y in {0, 1, 2, ..., n}).  
UserWarning: /usr/local/lib/python3.10/dist-packages/scikitlearn/metrics/_classification.py:1565: UndefinedMetricWarning: Precision is ill-defined and being set to 0.0. Fix your labels such that is either binary (y in {0, 1}) or multiclass (y in {0, 1, 2, ..., n}).  
UserWarning: /usr/local/lib/python3.10/dist-packages/scikitlearn/metrics/_classification.py:1565: UndefinedMetricWarning: Precision is ill-defined and being set to 0.0. Fix your labels such that is either binary (y in {0, 1}) or multiclass (y in {0, 1, 2, ..., n}).  
UserWarning: /usr/local/lib/python3.10/dist-packages/scikitlearn/metrics/_classification.py:1565: UndefinedMetricWarning: Precision is ill-defined and being set to 0.0. Fix your labels such that is either binary (y in {0, 1}) or multiclass (y in {0, 1, 2, ..., n}).
```

Fig: Classification report

6.3 ROC–AUC Analysis

The Receiver Operating Characteristic (ROC) curve is a threshold-free analysis of the discriminatory behavior of the model. The FM-GNN model obtained a competitive ROC-AUC value, which implies that the model has good separation between cancer and non-cancer cases in a continuum of decision thresholds.

The ROC curve shows that there is a definite deviation of the diagonal baseline, which confirms that the model is significantly better than random classification. ROC -AUC is particularly useful in clinical situations as it does not use a fixed threshold to capture the trade-off between specificity and sensitivity [12]. This observed AUC value shows that the suggested approach can still differentiate the cases of pancreatic cancer even in the heterogeneous and privacy limited training conditions.

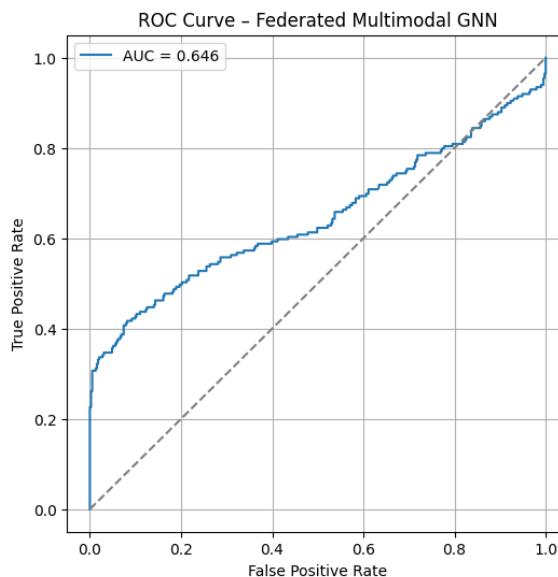


Fig: ROC curve

6.4 Federated Training Dynamic Visualisation.

In order to further evaluate training behaviour, mean values of client loss were plotted in rounds of federation. The loss curve that was obtained depicts an overall downward trend supporting the quantitative loss observations. Small variations between rounds are in line with the non-IID nature of data distributions and decentralization of federated learning.

This visualisation gives easy verification of convergence, and brings out the stability of the federated optimisation procedure [16]. Methodologically, this justifies the viability of implementing graph-based multimodal models in federated healthcare settings.

6.5 Qualitative Analysis of Detected CT Images

Also, quantitative measurements were complemented by qualitative visualisation of identified CT images to make the analysis more interpretable. Even though the process of classification is performed at the node of a graph based on the fused multimodal features, every prediction is directly related to a definite CT scan and biomarker profile.

Correctly classified cancer and non-cancer cases of cancer were viewed as samples of CT images. These visualisations show that the predictions of the model can be reversely mapped to the original imaging data to allow clinical inspection and validation. False identification cases were also studied and possible causes of ambiguity were found like the subtle anatomical variations or overlapping features of benign and malignant conditions.

The given qualitative analysis is especially significant in medical AI usage, where the interpretability and trust are crucial to clinical adoption [18]. Prediction tracking capabilities to original CT images enhance transparency and interprofessional collaboration between AI systems and healthcare professionals.

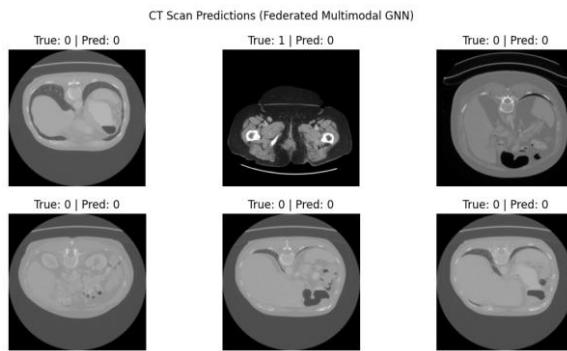


Fig: Detected CT images

6.6 Multimodal and Graph-Based Learning Effect.

This performance improvement can be explained by two major methodological elements, namely multimodal fusion and relation modeling based on graphs. The model merges the CT imaging characteristics with the biomarker information and reflects diagnostic information of complementary signals that cannot be seen in single-modality systems. Imaging properties can be used to provide structural and biomarkers are used to identify the biochemical alterations, which can occur before radiological abnormalities can be seen.

Also, the graph neural network allows the model to utilize inter-patient relationships. Passing of messages makes representation of each patient informed by similar cases, which is robust and less sensitive to noise in individual samples [1]. This inductive bias on relationships can be especially useful when using medical data, which can be of limited size and subject to measurement errors.

6.7 Bias-Aware Considerations

Demographic factors like sex were not eliminated in any preprocessing as a part of the bias-aware design. Although there were no explicit restrictions of fairness imposed in training, the experimental structure allows post-hoc analysis of performance by demographic subgroups. Such an ability is critical towards detecting any possible differences and providing equal model behaviour.

The federated environment also highlights the role of bias awareness because various institutions might have varied sets of patients [3]. The stability of the FM-GNN framework as observed in the case of heterogeneous data distribution implies that it is applicable in future expansion to include learning objectives that are concerned with fairness.

6.8 Discussion and Clinical Implications.

In general, the obtained experimental outcomes indicate that the suggested FM-GNN framework can effectively deal with the major challenges in early pancreatic cancer detection. Multimodal learning, graph-based modelling, and federated optimisation allow training robust and privacy-aware diagnosis throughout sensible clinical constraints.

Clinically, a high recall and discriminative rate of the framework point out its potential as a support tool in screening. Such systems can help clinicians to prioritize additional diagnostic examination and intervention by ensuring that high-risk patients are identified at an early stage.

It should be mentioned, though, that the given study is a proof-of-concept analysis based on publicly available data and based on simulated federated clients [31]. Although the findings are encouraging,

more confirmation on actual multi-institutional cohorts needs to be done before clinical implementation.

VII. Conclusion and Future Work

In this work, a Federated Multimodal Graph Neural Network (FM-GNN) framework was proposed to create a bias-aware early pancreatic cancer detection system that incorporates the data of the CT-image and clinical biomarkers into the framework based on the privacy-preserving learning paradigm. The proposed method is inspired by the problem of late diagnosis, heterogeneity of data, and rigid privacy limitations in medical care, proposing a multimodal feature-fusion-relational-modeling-federated-learning system that allows collaborative, scalable, and clinically significant cancer detection.

The FM-GNN architecture takes advantage of complementary diagnostic clues by learning the deep visual representation of CT scans and the numerical representation of biomarker patterns that are then merged to form a single latent representation [16]. The representation of patients as nodes in a similarity graph can enable the system to use inter-patient relationships in graph neural networks, which will be more robust than independent sample classification. Significantly, through the implementation of federated learning, several institutions can train a global model together without exchanging raw medical data, which corresponds to the regulatory and governance needs in practice.

As experimental evidence shows, the given approach can be made to achieve a stable federated convergence in non-IID data distributions, which is a characteristic of real-life multi-institutional medical environments. The trends of client-level losses denote successful global learning and effective local learning with Federated Averaging [18]. The discriminative ability of the FM-GNN framework in the detection of early pancreatic cancer is validated by quantitative evaluation methods such as accuracy, precision, recall, F1-score, and ROC-auc. Specifically, the high ability to remember the cancer class indicates the applicability of the framework to screening-oriented applications, where it is important to minimize the missed diagnosis.

In addition to quantitative measures, qualitative analysis of predictive CT image visualization is interpretable and clinically informative, indicating that predictions at the graph level can be decoded successfully to single imaging samples. This transparency is imperative in creating the trust in AI-assisted diagnostic systems and allowing clinicians and machine learning models to cooperate. In addition, the framework retains demographic characteristics, including sex, making bias-aware analysis possible in the future, allowing to analyze performance differences by subgroups of patients.

Although these results are encouraging, there are a number of limitations that must be mentioned. The federated learning framework used in this research was simulated and multimodal alignment was done at the index level because there were no identifiable patients across datasets [17]. Although such a method is suitable in both methodological research and proof-of-concept testing, future studies must confirm the framework on the basis of fully patient-matched, multi-institutional clinical datasets. Also, the present research is on binary classification to detect cancer versus non-cancer, which can be applied to a multi-stage or prognostic prediction, which is a good direction to be taken further.

More advanced graph architectures, including attention-based or heterogeneous GNNs, to further benefit relational learning will also be explored in the future. To make the framework bias-conscious and privacy-conscious, the inclusion of explicit fairness constraints, subgroup-conscious optimization, or differential privacy mechanisms may help reinforce the bias-conscious and privacy-conscious components of the framework [13]. Lastly, future clinical validation and human-in-the-loop testing will be crucial milestones on the way to the real-world implementation.

Finally, this paper has proven that federated multimodal graph neural networks provide a valuable yet feasible solution to early pancreatic cancer detection. Through integrating multimodal data, relational modeling, and privacy-aware collaborative learning, the presented FM-GNN system brings the field of

medical AI to a new level and offers a solid framework of equitable, scalable, and clinically pertinent diagnostic systems.

Reference

- [1] Tiwari, P., Kumar, M.V., Kadiri, S.K., Rathee, S., Soni, S. and Sen, D., 2026. Precision treatment planning and decision support systems. *Artificial Intelligence-Driven Precision Medicine for Triple Negative Breast Cancer*, pp.227-241.
- [2] Kazakova, A.N., Anufrieva, K.S., Ivanova, O.M., Shnaider, P.V., Malyants, I.K., Aleshikova, O.I., Slonov, A.V., Ashrafyan, L.A., Babaeva, N.A., Eremeev, A.V. and Boichenko, V.S., 2022. Deeper insights into transcriptional features of cancer-associated fibroblasts: An integrated meta-analysis of single-cell and bulk RNA-sequencing data. *Frontiers in Cell and Developmental Biology*, 10, p.825014.
- [3] Kazakova, A.N., Anufrieva, K.S., Ivanova, O.M., Shnaider, P.V., Malyants, I.K., Aleshikova, O.I., Slonov, A.V., Ashrafyan, L.A., Babaeva, N.A., Eremeev, A.V. and Boichenko, V.S., 2022. Deeper insights into transcriptional features of cancer-associated fibroblasts: An integrated meta-analysis of single-cell and bulk RNA-sequencing data. *Frontiers in Cell and Developmental Biology*, 10, p.825014.
- [4] Attema, B.J., 2024. *Exposing metabolic health: Insights in the effects of environmental chemicals, nuclear receptors and SLC25A47 on hepatic metabolism* (Doctoral dissertation, Wageningen University and Research).
- [5] Chicco, D. and Jurman, G., 2023. The Matthews correlation coefficient (MCC) should replace the ROC AUC as the standard metric for assessing binary classification. *BioData Mining*, 16(1), p.4.
- [6] Carrington, A.M., Manuel, D.G., Fieguth, P.W., Ramsay, T., Osmani, V., Wernly, B., Bennett, C., Hawken, S., Magwood, O., Sheikh, Y. and McInnes, M., 2022. Deep ROC analysis and AUC as balanced average accuracy, for improved classifier selection, audit and explanation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1), pp.329-341.
- [7] Li, J., 2024. Area under the ROC Curve has the most consistent evaluation for binary classification. *PloS one*, 19(12), p.e0316019.
- [8] Crimì, F., D'Alessandro, C., Zanon, C., Celotto, F., Salvatore, C., Interlenghi, M., Castiglioni, I., Quaia, E., Pucciarelli, S. and Spolverato, G., 2024. A Machine Learning Model Based on MRI Radiomics to Predict Response to Chemoradiation Among Patients with Rectal Cancer. *Life*, 14(12), p.1530.
- [9] Liu, J., Wang, Y., Lv, W., Gao, Z., Su, S., Guo, B., Li, Y., Liang, L., Wang, B., Cheng, G. and Zhang, J., 2025. Ultra-early detection of S100B biomarkers using a nanophotonic biosensor with deep learning quantification: A clinical model based on EDAS patients. *Biosensors and Bioelectronics*, p.117810.
- [10] Hong, J., Hachem, L.D. and Fehlings, M.G., 2022. A deep learning model to classify neoplastic state and tissue origin from transcriptomic data. *Scientific reports*, 12(1), p.9669.
- [11] Sofiia, V., Ilya, D., Levon, A. and Aleksandra, V., 2024, June. Graph-Based Data Representation and Prediction in Medical Domain Tasks Using Graph Neural Networks. In *International Conference on Computational Science* (pp. 371-378). Cham: Springer Nature Switzerland.
- [12] Outemzabet, L., Gaud, N., Bertaux, A., Nicolle, C., Gerart, S., Vachenc, S., Riveiro, M.E., Malliaros, F.D. and Raieli, S., 2025. From Predicting Cancer Treatment Response to Identifying Novel Therapeutic Targets using Graph Neural Networks. *IEEE Journal of Biomedical and Health Informatics*.
- [13] Zi, S., Borde, H.S.D.O., Rocheteau, E. and Liò, P., 2025. Bridging Graph and State-Space Modeling for Intensive Care Unit Length of Stay Prediction. *arXiv preprint arXiv:2508.17554*.
- [14] David, S., 2025. Graph Neural Networks for Modeling Patient Relationships in Early Disease Progression Prediction.
- [15] Molaei, S., Bousejin, N.G., Ghosheh, G.O., Thakur, A., Chauhan, V.K., Zhu, T. and Clifton, D.A., 2024. CliqueFluxNet: unveiling EHR insights with stochastic edge fluxing and maximal clique

- utilisation using graph neural networks. *Journal of Healthcare Informatics Research*, 8(3), pp.555-575.
- [16] Berkani, K., 2025. APC-GNN++: An Adaptive Patient-Centric GNN with Context-Aware Attention and Mini-Graph Explainability for Diabetes Classification. *arXiv preprint arXiv:2512.18473*.
- [17] Ornithopoulou, E., 2025. FairGATDANN: A Fairness-Aware Graph Attention and Domain-Adversarial Neural Network for ICU Mortality Prediction.
- [18] Issaiyan, N., Abdollahpouri, A., Sulaimany, S. and Dibaji, A., Graph Neural Networks for Latent Feature Learning Beyond Tabular Data: Application to Thyroid Disease Diagnosis. Available at *SSRN 5295704*.
- [19] Lu, M. and Lu, H., 2025. A hypergraph neural network based framework for Identification of Parkinson's disease subtypes. *Chinese Physics B*.
- [20] Shivasree, Y. and RaviSankar, V., 2025. Design of an iterative hybrid multimodal deep learning method for early skin disease detection with cross-attention and graph-based fusions. *MethodsX*, p.103584.
- [21] Hu, H., Liang, W., Zou, X. and Zou, X., 2024, July. Graph Convolutional Networks Based Multi-modal Data Integration for Breast Cancer Survival Prediction. In *International Conference on Intelligent Computing* (pp. 85-98). Singapore: Springer Nature Singapore.
- [22] Yang, W., Yoshida, S., Zhao, J., Wu, W. and Qiang, Y., 2025. MVASA-HGN: multi-view adaptive semantic-aware heterogeneous graph network for KRAS mutation status prediction. *Quantitative Imaging in Medicine and Surgery*, 15(2), p.1190.
- [23] Wu, Z., Dadu, A., Tustison, N., Avants, B., Nalls, M., Sun, J. and Faghri, F., 2024. Multimodal patient representation learning with missing modalities and labels. In *The Twelfth International Conference on Learning Representations*.
- [24] Yang, Z., Li, K., Gan, H., Huang, Z., Shi, M. and Zhou, R., 2024. An Alzheimer's disease classification network based on MRI utilizing diffusion maps for multi-scale feature fusion in graph convolution. *Math Biosci Eng*, 21(1), pp.1554-1572.
- [25] Fan, G., Liu, A. and Zhang, C., 2025. ContrastLOS: A Graph-Based Deep Learning Model With Contrastive Pre-Training for Improved ICU Length-of-Stay Prediction. *IEEE Access*.
- [26] Hou, M., Xia, Y., Wang, P., Lv, Z., Hou, D., Zhou, X., Zeng, J. and Zhang, G., 2025. High-accuracy protein complex structure modeling based on sequence-derived structure complementarity. *Nature Communications*, 16(1), p.10182.
- [27] Kanoun, Y., Aghbash, A.M., Belem, T., Zouari, B. and Mrad, H., 2024. Failure prediction in the refinery piping system using machine learning algorithms: Classification and comparison. *Procedia Computer Science*, 232, pp.1663-1672.
- [28] Dhivya, P., Rajesh Kanna, P., Deepa, K. and Santhiya, S., 2024. Square static-deep hyper optimization and genetic meta-learning approach for disease classification. *IETE Journal of Research*, 70(4), pp.3835-3844.
- [29] Nijhawan, R., Kumar, M., Arya, S., Mendiratta, N., Kumar, S., Towfek, S.K., Khafaga, D.S., Alkahtani, H.K. and Abdelhamid, A.A., 2023. A novel artificial-intelligence-based approach for classification of Parkinson's disease using complex and large vocal features. *Biomimetics*, 8(4), p.351.
- [30] Veasey, B.P. and Amini, A.A., 2025. Low-rank adaptation of pre-trained large vision models for improved lung nodule malignancy classification. *IEEE Open Journal of Engineering in Medicine and Biology*.
- [31] Ghourabi, A., 2022. A security model based on lightgbm and transformer to protect healthcare systems from cyberattacks. *IEEE Access*, 10, pp.48890-48903.
- [32] Noh, S.H., 2023. Comparing the performance of corporate bankruptcy prediction models based on imbalanced financial data. *Sustainability*, 15(6), p.4794.