Beyond COVID-19 Diagnosis: Prognosis with Hierarchical Graph Representation Learning

Chen Liu, Jinze Cui, Dailin Gan, and Guosheng Yin

The University of Hong Kong, Pokfulam, Hong Kong
{liuchen, cjz1206, davidgan, gyin}@hku.hk

Abstract. Coronavirus disease 2019 (COVID-19), the pandemic that is spreading fast globally, has caused over 181 million confirmed cases. Apart from the reverse transcription polymerase chain reaction (RT-PCR), the chest computed tomography (CT) is viewed as a standard and effective tool for disease diagnosis and progression monitoring. We propose a diagnosis and prognosis model based on graph convolutional networks (GCNs). The chest CT scan of a patient, typically involving hundreds of sectional images in a sequential order, is formulated as a densely connected weighted graph. A novel distance aware pooling is proposed to abstract the node information hierarchically, which is robust and efficient for such densely connected graphs. Our method, combining GCNs and distance aware pooling, can integrate the information from all slices in the chest CT scans for optimal decision making, which leads to the state-of-the-art accuracy in the COVID-19 diagnosis and prognosis. With less than 1% of the total number of parameters in the baseline 3D ResNet model, our method achieves 94.8% accuracy for diagnosis, which represents a 2.4% improvement over the baseline on the same dataset. In addition, we can localize the most informative slices with disease lesions for COVID-19 within a large sequence of chest CT images. The proposed model can produce visual explanations for the diagnosis and prognosis, making the decision more transparent and explainable, while RT-PCR only leads to the test result with no prognosis information. The prognosis analysis can help hospitals or clinical centers designate medical resources more efficiently and better support clinicians to determine the proper clinical treatment.

Keywords: COVID-19 Diagnosis · Lesion Localization · GCN · Prognosis.

1 Introduction

Coronavirus disease 2019 (COVID-19) has resulted in an ongoing pandemic in the world. To control the sources of infection and cut off the channels of transmission, rapid testing and detection are of vital importance. The reverse transcription polymerase chain reaction (RT-PCR) is a widely-used screening technology and viewed as the standard diagnostic method for suspected cases. However, this method highly relies upon the required lab facilities and the diagnostic kits.

In addition, the sensitivity of RT-PCR is not high enough for early diagnosis [1, 5]. To mitigate the limitations of RT-PCR, the computed tomography (CT) has been widely used as an effective complementary tool by providing medical images of the lung to reveal the details of the disease and its prognosis [7, 3], for which RT-PCR cannot. Additionally, CT is also useful in monitoring the COVID-19 disease progression and therapeutic efficacy evaluation [12, 9].

The chest CT slices of a patient have a sequential and hierarchical data structure. The relationship between slices possesses more information than the order of the slices. The adjacent ones with the same abnormality could be considered as one lesion. The slices containing the same type of lesions may not be continuous as the lesions are distributed in various parts of the lung. We propose a diagnosis and prognosis system that combines graph convolutional networks (GCNs) and a distance aware pooling procedure, which integrates the information from all slices in the chest CT scans for optimal decision making. Our major contributions are three-fold: (1) Owing to the sequential structure of CT images, this is the first work to utilize GCNs to extract node information hierarchically, and conduct both diagnosis and prognosis for COVID-19. The prognosis can help facilitate medical resources, e.g., use of ventilators or admission to Intensive Care Units (ICUs), more efficiently by triaging mild or severe patients. (2) A novel pooling method, called distance aware pooling, is proposed to aggregate the graph, i.e., the patient's CT scan, effectively. The new pooling method integrated with GCNs can aggregate a densely connected graph efficiently. (3) The proposed model can localize the most informative slices within a chest CT scan, which significantly reduces the amount of work for radiologists.

2 Methodology

We propose a GCN-based diagnosis and prognosis method that models the sequential slices of CT scans hierarchically.

To downsample and learn graph-level representation from the input node features, a novel distance aware pooling method is proposed. The node features refer to the slices in a CT scan. The model gradually extracts information from the slice level to the patient level by graph convolution and pooling. Eventually, a higher-level representation is learned and further used for diagnosis, prognosis, and lesion localization. The schema of our model is illustrated in Figure 1, which is composed of GCNs, pooling modules, a multilayer perceptron (MLP) classifier, and a one-drop localization module. The graph convolution-based method can integrate all slices in the chest CT scans for optimal decision making.

Furthermore, we propose the one-drop localization to localize the most informative slices and reduce redundancies, so that radiologists may focus on those recommended slices with the most suspected lesion areas. Consequently, the proposed model can produce visual explanations for the diagnosis and prognosis. The localized slices may help uncover how a given classification method arrives at the conclusion. Specifically, it reveals if a method is merely overfitting on the training data by examining the slices it attends to. Conversely, it may also be

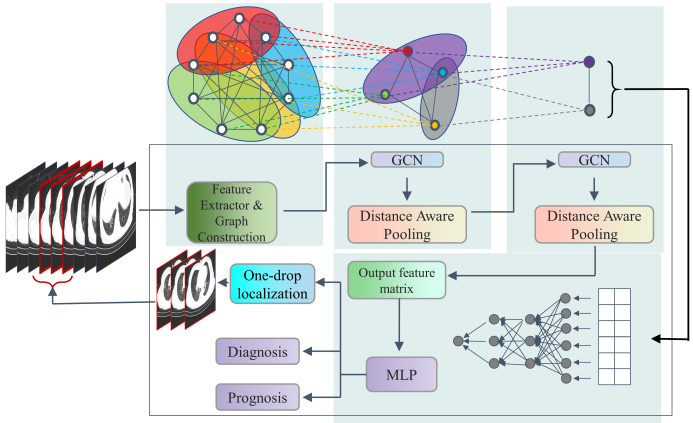

Fig. 1: Schema of our model structure. The CT scan of one patient is converted to a densely connected graph. The GCN and the Distance Aware Pooling (DAP) are integrated to learn a graph-level representation. At each of the two layers, the node embedding is learned by GCN, and the cluster membership is calculated by DAP. The aggregated graph, which is on the top right corner, is passed to an MLP. Meanwhile, the one-drop localization can localize the most informative slices in the CT scan in a weakly supervised manner.

used to identify the intrinsic dataset bias, notably the data acquisition bias [2], and thus guide the data collection process.

2.1 Problem statement

Let $\mathcal{G}$ ($\mathbb{V}$, $\mathbb{E}$) be a patient's CT scan graph, with $|\mathbb{V}| = N$ nodes and $|\mathbb{E}|$ edges, where $|\cdot|$ represents the cardinality of a set. For each $v_i \in \mathbb{V}$, $\boldsymbol{x}_i$ is the corresponding d-dimensional vector. Let $\boldsymbol{X} \in \mathbb{R}^{N\times d}$ be the node feature matrix, and $\boldsymbol{A}^{adj} \in \mathbb{R}^{N\times N}$ be the weighted adjacency matrix. Each entry in $\boldsymbol{A}^{adj}$ is defined using cosine similarity, which is $\boldsymbol{A}^{adj}_{i,j} = <\boldsymbol{x}_i, \boldsymbol{x}_j>/(\|\boldsymbol{x}_i\| \cdot \|\boldsymbol{x}_j\|)$.

Each graph $\mathcal{G}$ has a label y. For diagnosis, the label represents its class from COVID-19 positive, common pneumonia, or normal individuals. For prognosis, the class indicates whether a COVID-19 positive patient develops into severe/critical illness status. Thus, the diagnosis and prognosis of COVID-19 is a task of graph classification. Given a training dataset $\mathbb{T} = \{(\mathcal{G}_1, y_1), \ldots, (\mathcal{G}_M, y_M)\}$, the goal is to learn a mapping $f : \mathcal{G} \to y$, which classifies a graph $\mathcal{G}$ into the corresponding class y. Our model is composed of two modules: f_1 includes node convolution and feature pooling, and f_2 involves an MLP classifier. At each layer of f_1, the node embeddings and cluster membership are learned iteratively. The first module can be written as $f_1 : \mathcal{G} \to \mathcal{G}^p$, where $\mathcal{G}^p$ is the pooled graph with fewer nodes and a hierarchical feature representation. The second module is $f_2 : \mathcal{G}^p \to y$, which utilizes the graph-level representation learned for patient diagnosis and prognosis. The two modules are integrated in an end-to-end fashion.

Thus, the cluster assignment is learnt merely based on the graph classification objective.

2.2 Node convolution and feature pooling

Node convolution Node convolution applies the graph convolutional network to obtain a high-level node feature representation in the feature matrix $\boldsymbol{X}$. Although several methods exist to construct the convolutional network, the method recommended by [8] is particularly effective for our case, which is given by $\boldsymbol{X}^{(l+1)} = \sigma(\sqrt{\boldsymbol{D}^{(l)}}\boldsymbol{A}^{adj(l)}\sqrt{\boldsymbol{D}^{(l)}}\boldsymbol{X}^{(l)}\boldsymbol{W}^{(l)})$, where $\boldsymbol{D}^{(l)}$ is the diagonal degree matrix of $\boldsymbol{A}^{adj(l)} - \boldsymbol{I}$, and $\boldsymbol{W}^{(l)} \in \mathbb{R}^{d\times k}$ is a learnable weight matrix at the l-th layer. Due to the application of feature pooling, the topology of the graph changes at each layer, and thus the dimensions of matrices involved are reduced accordingly.

Distance aware pooling We propose an innovative pooling method, which includes graph-based clustering and feature pooling. Below, we outline the pooling module and illustrate how it is integrated into an end-to-end GCN-based model. Empirically, it is shown to be more robust for densely connected graphs. The overall structure of the pooling method is shown in Figure 2.

• **Improved receptive field** The concept of the receptive field, RF, used in the convolutional neural network (CNN) was extended to GNNs [11]. In particular, RF^{node} is defined as the number of hops required to cover the neighborhood of a given node, such that given a chosen node, a cluster can be obtained based on a fixed receptive field h. However, this design may not be applicable to densely connected graphs, because one node may be connected to most of the nodes in the graph even for a small value of h. Hence, we define an improved receptive field for densely connected graphs, RF^d, denoted by h^d. It is a radius centered at a given node and retains the edge weight information in the clustering process. The value of h^d is not restricted to integers. Define $\mathcal{N}(v_i)$ as the local neighborhood of the node v_i, and $\mathcal{N}_{h^d}(v_i)$ as the RF^d neighborhood of the node v_i with a radius h^d, and $\forall v_j \in \mathcal{N}(v_i)$, v_i and v_j are connected by the edge (v_i, v_j)

• **Node clustering** Inspired by the clustering and ranking ideas in [11] and [6], we propose a local node clustering and score ranking method. Each node is considered as a center of a cluster for a given h^d. Then, we score all the clusters and choose the top k proportion of them to represent the next layer's nodes with pooled feature values, where k is a hyperparameter.

• **Cluster ranking** Given a node v_i and a radius h^d, $\mathcal{N}_{h^d}(v_i)$ is the corresponding RF^d neighborhood. Let $\mathbb{I}_1(v_i)$ be the index set of the nodes in $\mathcal{N}_{h^d}(v_i)$. Define the distance matrix $\boldsymbol{A}^{dis}$ as $\boldsymbol{A}^{dis} = \mathbb{1} - \boldsymbol{A}^{adj}$, where $\mathbb{1}$ is the matrix with all entries of 1's. The cluster score is defined as $\alpha_i = \sum_{m,n\in\mathbb{I}_1(v_i)} \boldsymbol{A}^{dis}_{m,n}/|\mathcal{N}_{h^d}(v_i)|$, where $m \neq n$. If α_i is small, nodes in $\mathcal{N}_{h^d}(v_i)$ are close to each other. The top k proportion of clusters form the next layer's nodes.

• **Selecting cluster centers** We define $\mathbb{V}_j = \mathcal{N}[v_j] \cap \mathcal{N}_{h^d}(v_i)$, $\forall v_j \in \mathcal{N}_{h^d}(v_i)$ as the set of nodes connected to v_j in $\mathcal{N}_{h^d}(v_i)$, where $\mathcal{N}[v_j]$ is the closed neighborhood of node v_j. Let $\mathbb{I}_2(v_j)$ be the index set of the nodes in $\mathbb{V}_j$. The node

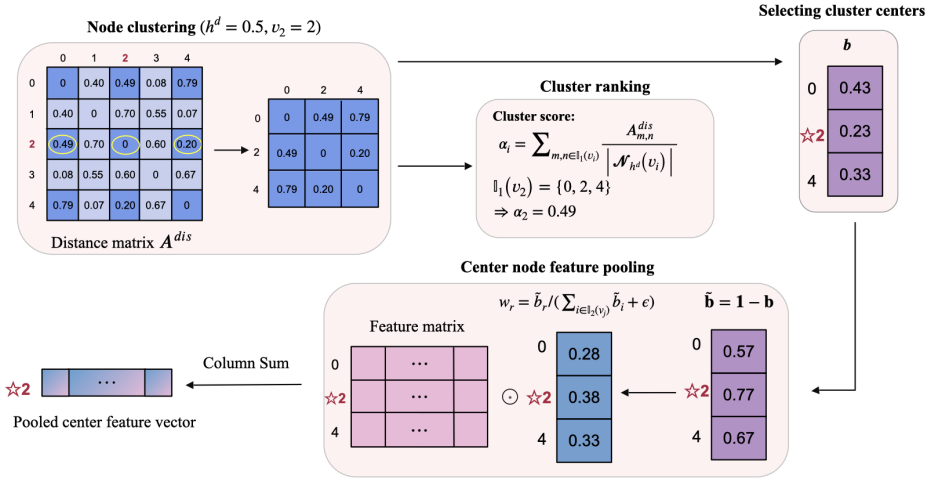

Fig. 2: The structure of the Distance Aware Pooling method with a numeric example. Four steps in the figure include node clustering, cluster ranking, selecting cluster centers, and center node feature pooling, which correspond to those detailed in Section 2.2. In this example, the proposed center in the node clustering step is node $v_2 = 2$, and the corresponding RF^d neighborhood with $h^d = 0.5$ includes nodes 0, 2, and 4. It should be noted that the proposed center may not be the final center for this cluster. The cluster score α_i of each cluster with the proposed center $v_i = i$ can be calculated. With α_i, the top k proportion of clusters can be obtained. In the selecting cluster centers step, $\boldsymbol{b}$ is the average of each row sum, and node v_2 is still chosen as the center with the least average distance value. In the center node feature pooling step, w_r is the weight of nodes derived from $\boldsymbol{b}$. In this way, we can derive the weighted pooled center feature vector.

score is $b_j = \sum_{c \in \mathbb{I}_2(v_i)} \boldsymbol{A}^{dis}_{j,c} / |\mathbb{V}_j|$. The node with the smallest value of b_j is chosen as the cluster center, and represents a new node in the next layer.

• **Center node feature pooling** Based on the node scores, we rank the nodes in $\mathbb{V}_j$ and assign a weight w_r to $\boldsymbol{x}_r$, where $r \in \mathbb{I}_2(v_j)$. Define $\boldsymbol{b}$ as a vector containing all the weights of nodes in $\mathbb{V}_j$, and let $\tilde{\boldsymbol{b}} = \mathbf{1} - \boldsymbol{b}$, where $\mathbf{1}$ is a vector with all 1's. The weight vector $\boldsymbol{w}$ is defined as $w_r = \tilde{b}_r / (\sum_{i \in \mathbb{I}_2(v_j)} \tilde{b}_i + \epsilon)$, where ϵ is an extremely small positive value to avoid 0 in the denominator. The value of the pooled center feature is defined as $\boldsymbol{x}^p = \boldsymbol{X}_r \boldsymbol{w}$, and then $\boldsymbol{x}^p$ is used as the center feature representing $\mathcal{N}_{h^d}(v_i)$.

• **Next layer node connectivity** Following the idea in [15], the connectivity of the nodes in the next layer is preserved as follows. According to the above ranking and pooling methods, a pooled graph $\mathcal{G}^p$ with the node set $\mathbb{V}^p$ is obtained. The next step is to determine the pooled adjacency matrix $\boldsymbol{A}^{adj,p}$. Define matrix $\boldsymbol{S}$ such that the columns of $\boldsymbol{S}$ are the top k clusters' weight vectors $\boldsymbol{w}$. Hence, the pooled adjacency matrix is defined as $\boldsymbol{A}^{adj,p} = \boldsymbol{S}^\top \boldsymbol{A}^{adj} \boldsymbol{S}$.

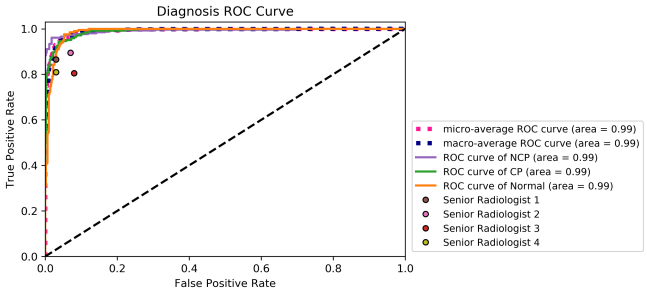

Fig. 3: The ROC curve and AUC (area under the curve) for diagnosis. 'NCP', 'CP', and 'Normal' indicate COVID-19 positive patients, common pneumonia patients, and healthy individuals respectively. The filled dots represent the performance on 'NCP' diagnosis of senior radiologists with 15 to 25 years of clinical experience [16]. It shows that our method is comparable to the senior radiologists.

2.3 Classification and Localization Based on Pooled Graphs

Graph classification A hierarchical representation for each patient is abstracted by the above GCNs with distance aware pooling. The representation is an $N' \times D'$ feature matrix, where N' is the number of clusters and D' is the number of features in each cluster. For each feature, the mean and the maximum over the N' clusters are calculated. Subsequently, we obtain a $2D'$ feature vector, where the first and second D' elements are the mean and maximal values of each feature respectively. Then, an MLP classifies the $2D'$ feature vector into one of the three classes: COVID-19 positive, common pneumonia, or normal.

Weakly supervised informative CT slices localization Localizing the CT slices with lesions is also vital for the diagnosis. Therefore, we propose the one-drop localization to select the most informative CT slices for the model to make a decision. This method does not require mask annotation for lesions, which can be learned in a weakly supervised manner. We are inspired by the backward elimination of stepwise regression used for knowledge discovery [4], an automatic procedure for variable selection. For each patient with the $N' \times D'$ feature matrix, we first predict the target class using the above MLP classifier. Then, we occlude one cluster each time to obtain N' new feature matrices with an equal size of $(N' - 1) \times D'$. For each of the new feature matrices, the score of the target class is calculated by the same MLP classifier. The occluded cluster with the lowest score from the N' results is chosen. Because ignoring this cluster leads to the lowest score of the target class, the cluster should contain the most crucial information. Finally, the cluster center is determined and the top k^s CT slices with the highest similarity to the center are further localized, where k^s is the number of selected slices.

3 Experiments and Results

For COVID-19 diagnosis and prognosis, we compare our model with the baseline, a 3D ResNet-18 classification network [16], and the state-of-the-art graph classification methods. Moreover, we appraise whether the proposed method can deliver meaningful and interpretable clusters on the input chest CT scans by comparing the localization results with the slices containing lesions (ground truth).

3.1 Dataset and Implementation Details

Dataset We utilize the CT dataset from the 2019 Novel Coronavirus Resource (2019nCoVR) [16]. The dataset includes the complete chest CT scans of 929 COVID-19 positive patients, 964 common pneumonia patients, and 849 healthy individuals. The dataset also provides these patients' clinical prognosis, whether the patients developed into severe/critical illness status, referring to admission to ICUs, mechanical ventilation, or death. The prognosis analysis could support the hospitals to designate medical resources more efficiently. In addition, the dataset summarizes the slices with lesions for COVID-19 positive and common pneumonia patients, which can be used to evaluate the one-drop localization method. Each CT slice is normalized to the dimension 256×256.

Data preprocessing We apply the following two methods for chest CT scan image feature extraction.

• **CNN feature extraction** We utilize Inception V3 [14] pretrained on ImageNet [13]. The feature map of the bottleneck layer, i.e., the last layer before the flatten operation, is regarded as the node representation in the graph.

• **Wavelet decomposition extraction** Considering each slice in CT scans as a 2-dimensional signal, it can be viewed as a function with two variables, which can be reconstructed as a summation of wavelet functions multiplying their coefficients for a given resolution [10]. We choose the Haar wavelet function with resolution 3. The flattened approximation matrix of the image signal, which is of dimension 1024, is used as the feature embedding of a slice in a CT scan.

Implementation Details We use systematic sampling to ensure that 48 slices for each CT scan are chosen. CT scans of 60% individuals are randomly chosen as the training set, 25% as the test set, and the remaining 15% for the validation. To avoid information leakage, the dataset is split according to individuals instead of the CT slices. Similar to [17] and [15], we repeat the aforementioned data splitting on 20 random seeds. For each random seed, the model is trained from scratch. The maximum, average and standard deviation of test accuracies are reported. We use five GCN layers and a two-layer MLP classifier. The pooling proportion k is set as 0.8. The negative log-likelihood loss is used for graph classification. The Adam optimizer is applied with an initial learning rate 0.0002 and a linear decay schedule. For prognosis, the parameters of GCN and pooling are initialized using those pretrained on the diagnosis task. All models are trained for 128 epochs with possible early stopping.

(a) **Performance Evaluation of COVID-19 Diagnosis**

Method	Feature Extractor	Average Accuracy (SD)	Best Accuracy	Time (s/epoch)
GCN-DAP	Inception V3	**93.93% (0.41%)**	**94.80%**	22.70
GCN-DAP	Wavelet	83.65% (1.01%)	85.16%	20.90
GCN-ASAP	Inception V3	75.20% (18.70%)	93.74%	30.00
GCN-ASAP	Wavelet	51.43% (11.90%)	81.50%	27.25
GCN-DiffPool	Inception V3	71.22% (23.73%)	94.31%	18.35
GCN-HGP-SL	Inception V3	93.89% (0.39%)	94.22%	45.60

(b) **Performance Evaluation of COVID-19 Prognosis**

Method	Feature Extractor	Average Accuracy (SD)	Best Accuracy	Time (s/epoch)
GCN-DAP	Inception V3	**82.70% (3.90%)**	**91.39%**	7.67
GCN-DAP	Wavelet	78.98% (3.38%)	84.95%	1.83
GCN-ASAP	Inception V3	67.90% (11.09%)	82.80%	9.67
GCN-ASAP	Wavelet	60.22% (5.63%)	72.04%	2.30

Table 1: Performance evaluation of COVID-19 diagnosis and prognosis, where 'GCN-DAP' represents the proposed GCN-based method integrated with the distance aware pooling. 'ASAP', 'DiffPool', and 'HGP-SL' refer to the state-of-the-art hierarchical pooling methods. The last column is the average time in seconds to complete one training epoch for each model using a single NVIDIA V100 GPU.

3.2 Quantitative Results

Diagnosis and prognosis performance According to the ROC (receiver operating characteristic) curve for diagnosis in Figure 3, our method is comparable to the senior radiologists with 15 to 25 years of clinical experience. We also compare our method with the clinically applicable AI system based on 3D ResNet-18 [16], which reached 92.49% diagnosis accuracy. With less than 1% of the total number of parameters, our method achieves an improvement of 2.4% over this CNN-based state-of-the-art model.

In Table 1, we compared the performance of our model and GCN with the state-of-the-art hierarchical pooling methods, including adaptive structure aware pooling (ASAP) [11], differentiable pooling (DiffPool) [15], and hierarchical graph pooling with structure learning (HGP-SL) [17]. We observed that the Inception V3 feature extraction method constantly outperforms the wavelet decomposition method under the same model configuration. The gradient explosion occurs in around 50% of the runs under the ASAP and DiffPool, resulting in optimization failures, while this issue has not been witnessed during the training of our method and HGP-SL. Thus, the standard deviations of ASAP and DiffPool are much higher. Additionally, our method outperforms HGP-SL marginally but the training of our method is about 2 times faster. The training curves using DAP versus the aforementioned hierarchical pooling methods over 20 runs are presented in Figure 4, which shows that our model improves the training convergence, and DAP consistently outperforms ASAP and DiffPool across almost all runs.

Weakly supervised lesion localization results In addition to graph classification, our model can also localize the most informative CT slices for each CT scan in the test set using the procedure in Section 2.3 with $k^s = 10$. Because the

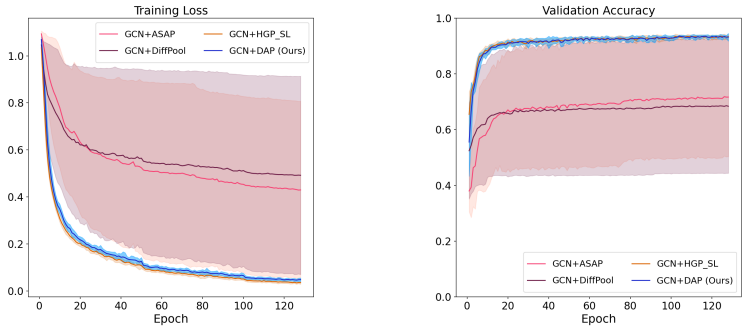

Fig. 4: Training curves of the GCN diagnosis model using DAP versus three hierarchical pooling methods over 20 runs, with varied random seeds and train-validation-test split. The solid lines represent the mean training loss and validation accuracy, and the shaded areas exhibit the intervals of one standard deviation. It shows that DAP consistently outperforms ASAP and DiffPool across almost all runs, and converges much faster.

chosen CT slices are the most decisive ones for diagnosis, they contain lesions related to COVID-19 or common pneumonia. We compare our selected slices to the CT slices with labeled lesions for each CT scan. Considering the slices in the same CT scan are ordered sequentially, we compare the slices between the first and last slices localized by our model with those labeled with lesions in the dataset. Among patients in the test set with 20 random seeds, the average precision and recall are 57.39% and 79.89% with standard deviations 3.32% and 3.94%, respectively. The localization results of four patients are presented in the Supplementary Material for illustration.

4 Conclusion

This paper introduces an efficient and robust GCN-based diagnosis and prognosis system with a distance aware pooling method. It can cluster nodes and learn the patient-level representation hierarchically. Unlike previous diagnosis methods based on CT scans, our model can produce coarse localization to highlight the potential slices with lesions, making the clinical decision more interpretable and reliable. To strengthen confidence in this framework and move towards clinical use, we should ensure that the model can provide explanations for the prediction instead of merely outputting the result. The localization can help analyze prediction failure, and improve researchers' understanding of the effect of adversarial attacks in the medical imaging domain.

Bibliography

[1] Ai, T., Yang, Z., Hou, H., Zhan, C., Chen, C., Lv, W., Tao, Q., Sun, Z., Xia, L.: Correlation of chest CT and RT-PCR testing for coronavirus disease 2019 (COVID-19) in china: A report of 1014 cases. Radiology **296**(2), E32–E40 (2020)

[2] Biondetti, G.P., Gauriau, R., Bridge, C.P., Lu, C., Andriole, K.P.: "Name that manufacturer". Relating image acquisition bias with task complexity when training deep learning models: experiments on head CT. arXiv preprint arXiv:2008.08525 (2020)

[3] Chung, M., Bernheim, A., Mei, X., Zhang, N., Huang, M., Zeng, X., Cui, J., Xu, W., Yang, Y., Fayad, Z.A., Jacobi, A., Li, K., Li, S., Shan, H.: CT imaging features of 2019 novel coronavirus (2019-nCoV). Radiology **295**(1), 202–207 (2020)

[4] Cios, K.J., Pedrycz, W., Swiniarski, R.W.: Data mining and knowledge discovery. In: Data Mining Methods for Knowledge Discovery, pp. 1–26. Springer (1998)

[5] Fang, Y., Zhang, H., Xie, J., Lin, M., Ying, L., Pang, P., Ji, W.: Sensitivity of chest CT for COVID-19: Comparison to RT-PCR. Radiology **296**(2), E115–E117 (2020)

[6] Gao, H., Ji, S.: Graph U-Nets. In: International Conference on Machine Learning. pp. 2083–2092 (2019)

[7] Huang, C., Wang, Y., Li, X., Ren, L., Zhao, J., Hu, Y., Zhang, L., Fan, G., Xu, J., Gu, X., et al.: Clinical features of patients infected with 2019 novel coronavirus in Wuhan, China. The Lancet **395**(10223), 497–506 (2020)

[8] Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: International Conference on Learning Representations (2017)

[9] Liechti, M.R., Muehlematter, U.J., Schneider, A.F., Eberli, D., Rupp, N.J., Hötker, A.M., Donati, O.F., Becker, A.S.: Manual prostate cancer segmentation in MRI: Interreader agreement and volumetric correlation with transperineal template core needle biopsy. European Radiology **30**(9), 4806–4815 (2020)

[10] Mallat, S.G.: A theory for multiresolution signal decomposition: The wavelet representation. IEEE Transactions on Pattern Analysis and Machine Intelligence **11**(7), 674–693 (1989)

[11] Ranjan, E., Sanyal, S., Talukdar, P.P.: ASAP: Adaptive structure aware pooling for learning hierarchical graph representations. Proceedings of the AAAI Conference on Artificial Intelligence **34**(04), 5470–5477 (2020)

[12] Rodriguez-Morales, A.J., Cardona-Ospina, J.A., Gutiérrez-Ocampo, E., Villamizar-Peña, R., Holguin-Rivera, Y., Escalera-Antezana, J.P., Alvarado-Arnez, L.E., Bonilla-Aldana, D.K., Franco-Paredes, C., Henao-Martinez, A.F., et al.: Clinical, laboratory and imaging features of COVID-

19: A systematic review and meta-analysis. Travel Medicine and Infectious Disease **34**, 101623 (2020)
[13] Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. International Journal of Computer Vision **115**(3), 211–252 (2015)
[14] Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2818–2826 (2016)
[15] Ying, Z., You, J., Morris, C., Ren, X., Hamilton, W., Leskovec, J.: Hierarchical graph representation learning with differentiable pooling. In: Advances in Neural Information Processing Systems. pp. 4800–4810 (2018)
[16] Zhang, K., Liu, X., Shen, J., Li, Z., Sang, Y., Wu, X., Zha, Y., Liang, W., Wang, C., Wang, K., et al.: Clinically applicable AI system for accurate diagnosis, quantitative measurements, and prognosis of COVID-19 pneumonia using computed tomography. Cell **181**(6), 1423–1433 (2020)
[17] Zhang, Z., Bu, J., Ester, M., Zhang, J., Yao, C., Yu, Z., Wang, C.: Hierarchical graph pooling with structure learning. arXiv preprint arXiv:1911.05954 (2019)