
Extraction of Airways using Graph Neural Networks

Raghavendra Selvan
University of Copenhagen
raghav@di.ku.dk

Thomas Kipf
University of Amsterdam
t.n.kipf@uva.nl

Max Welling
University of Amsterdam
CIFAR*
m.welling@uva.nl

Jesper H. Pedersen
University Hospital of Copenhagen
jped0106@regionh.dk

Jens Petersen
University of Copenhagen
phup@di.ku.dk

Marleen de Bruijne
University of Copenhagen
Erasmus MC, Rotterdam
marleen@di.ku.dk

Abstract

We present extraction of tree structures, such as airways, from image data as a graph refinement task. To this end, we propose a graph auto-encoder model that uses an encoder based on graph neural networks (GNNs) to learn embeddings from input node features and a decoder to predict connections between nodes. Performance of the GNN model is compared with mean-field networks in their ability to extract airways from 3D chest CT scans.

1 Introduction

Extraction of trees, like airways, from image data is useful in many image analysis applications. In this work, we pose tree extraction as a graph refinement task, of recovering underlying connectivity that corresponds to structures of interest from an over-complete graph. In previous work from the authors [1], graph refinement for extraction of tree structures was performed in a probabilistic inference setting using mean-field networks (MFNs) yielding competitive results.

Graph neural networks (GNNs) [2, 3] are a new class of models that are interpreted as generalisation of the message passing algorithms such as MFNs [4]. These models can be used as trainable end-to-end systems allowing inference using message passing even in cases where closed form algorithms do not exist [5]. The lack of prior work on using GNNs for graph refinement in medical imaging motivated us to examine the usefulness of GNNs for such tasks and present a preliminary investigation of using a variant of GNN, introduced in [6] known as Graph Auto-Encoders (GAEs), to extract airways.

2 Method

Given an over-complete input graph with N nodes, an adjacency matrix, $\mathbf{A}_I \in \{0, 1\}^{N \times N}$, describing connectivity between nodes, and a node feature matrix, $\mathbf{X} \in \mathbb{R}^{N \times F}$, where F is dimensionality of the features, we seek a model that can recover the ground truth adjacency matrix, $\mathbf{A} \in \{0, 1\}^{N \times N}$, corresponding to structures of interest from the input graph, i.e. $f(\mathbf{A}_I, \mathbf{X}) \rightarrow \mathbf{A}$. For airway data, we pre-process the CT scans into graph-structured data, following the method in [7], such that each node is associated with a 7-dimensional Gaussian density, with local radius, position and orientation in 3D, and their variances, as the node feature vector: $\mathbf{x}_i = [\mathbf{x}_\mu^i, \mathbf{x}_{\sigma^2}^i]$, comprising of mean, $\mathbf{x}_\mu^i \in \mathbb{R}^{7 \times 1}$, and variance for each feature, $\mathbf{x}_{\sigma^2}^i \in \mathbb{R}^{7 \times 1}$.

GAEs are comprised of an encoder that learns an embedding, $\mathbf{Z} \in \mathbb{R}^{N \times E}$, based on an input adjacency matrix and node features: $\mathbf{Z} = f(\mathbf{A}_I, \mathbf{X})$, where E is the dimensionality of the learnt

*Canadian Institute for Advanced Research

1st Conference on Medical Imaging with Deep Learning (MIDL 2018), Amsterdam.

embedding space. The encoder $f(\cdot)$ comprises of a graph convolution operation and a non-linearity $\sigma(\cdot)$. The encoder used in this work uses a formulation similar to [8], with weights $\mathbf{W}_0^{(l)} \in \mathbb{R}^{E \times E}$ and $\mathbf{W}_1^{(l)} \in \mathbb{R}^{E \times E}$ and hidden activations $\mathbf{H}^{(l)} \in \mathbb{R}^{N \times E}$ at layer l , and is given as:

$$\mathbf{H}^{(l+1)} = \sigma\left(\mathbf{H}^{(l)}\mathbf{W}_0^{(l)} + \mathbf{D}^{-1}\mathbf{A}_l\mathbf{H}^{(l)}\mathbf{W}_1^{(l)}\right), \quad (1)$$

where \mathbf{D} is the degree matrix obtained from the adjacency matrix with diagonal entries, $D_{ii} = \sum_{j=1}^N A_{ij}$. Note that for the first layer, we set $\mathbf{H}^{(0)} = \mathbf{X}$ and for the last layer $\mathbf{H}^{(L)} = \mathbf{Z}$. Also, notice from (1) that the encoder updates each node with a weighted transformation of node features of its neighbours. A non-linear decoder is used to predict the adjacency matrix from thus encoded embedding: $\alpha = g(\mathbf{Z})$. We use a radial basis function kernel as the decoder operating on each pair of nodes i.e.,

$$\alpha_{ij} = \exp\left(-\frac{1}{2}(\mathbf{z}_i - \mathbf{z}_j)^2\right), \quad \forall(i, j). \quad (2)$$

Position and radius information of adjacent nodes, according to α , are used to generate rough binary segmentations by drawing 3D spheres along each edge.

Dice loss: Weights of the GNN model, $\mathbf{W}^{(l)}$, are learned by minimising dice loss between predicted adjacency at the final layer, α , and ground truth labels, \mathbf{A} , of training data. Dice loss is useful to account for class imbalance (between edge and no-edge classes) and is given as:

$$\mathcal{L}(\mathbf{A}, \alpha) = 1 - \frac{2 \sum_{i,j=1}^N A_{ij} \alpha_{ij}}{\sum_{i,j=1}^N A_{ij}^2 + \sum_{i,j=1}^N \alpha_{ij}^2}. \quad (3)$$

3 Experiments and Results

We use two disjoint subsets of CT scans of size 32 and 100 subjects from the Danish lung cancer screening trial [9]. The first subset, with 32 scans, has expert-verified reference segmentations and is used for comparing the methods, whereas the second subset, of 100 scans, has reference segmentations obtained automatically using the airway extraction method in [10] and it is used to perform pre-training and tune the hyperparameters.


Tuning of GNN model is performed in three stages: hyperparameter tuning, pre-training and final model training. Using the dice loss in (3) for tuning, we set the number of GNN layers to 3 and number of hidden units per layer to 32 (see Figure 1a). We use Adam optimiser with learning rate 0.005 and batch size of 4 for training the GNN model. Further, we use all 100 scans with automatic reference segmentations to pre-train the GNN model. For the comparing mean-field network model, we use the hyperparameter setting reported in [1].

Centerlines are extracted from the generated binary segmentations to compute centerline distance, which is used to evaluate the performance of the methods. It is defined as $d_{err} = (d_{FP} + d_{FN})/2$, where d_{FP} is the average minimum Euclidean distance from extracted centerline points to reference centerline points, and d_{FN} is the average minimum Euclidean distance from reference centerline points to extracted centerline points.


Results Experiments comparing GNN and MFN models were conducted on the 32 reference scans split into 24 for training and 8 for test purposes. Table 1 summarises centerline distance error, d_{err} , on the test set. We show two versions of the GNN model: a stand-alone model (GNN) and with merged predictions from MFN (GNN+MFN), obtained from union of the predictions from both methods. Visualisation of MFN and GNN predictions on one test set image along with the reference segmentations is shown in Figure 1b.

Table 1: Performance comparison using centerline distance

Method	d_{FN} (mm)	d_{FP} (mm)	d_{err} (mm)
MFN	2.571	0.835	1.703 ± 0.186
GNN	2.890	3.913	3.402 ± 0.386
GNN+MFN	2.014	3.345	2.679 ± 0.264



(a) Plot showing the influence of increasing training set size on average dice score on the validation set for different hidden units per GNN layer.



(b) Airway tree centerlines for one of the test cases obtained from MFN predictions (blue) overlaid with the reference segmentations (pink surface) and the centerlines from GNN model (yellow).

4 Discussion

The error measures in Table 1 indicate that as a stand-alone method, GNN model does not compare favourably to MFN model. For the combined case of GNN+MFN we see that there is reduction in the d_{FN} measure indicating that the GNN model is able to add missing branches that are not detected by the MFN model. There are two main reasons, we believe, for the observed performance of GNN model. Firstly, a limited amount of labelled data, which is reflected in Figure 1a, wherein it appears that improvement in accuracy may follow from further increases in training set size. Secondly, the investigated GNN model models only nodes and does not take pairwise node interactions into account. For airway extraction tasks, modelling pairwise interactions can be beneficial [1]. A variation of GNN which explicitly learns representations for pairs of connected nodes may improve performance.

As a final remark, with our evaluation we observe that use of hand-crafted models such as MFNs can be useful in cases where domain knowledge is available. However, end-to-end trainable models using GNNs can be useful when there is limited insight into the exact nature of problems and may have greater potential with future growth in dataset sizes.

Acknowledgements This work was funded by the Independent Research Fund Denmark (DFR) and SAP SE Berlin.

References

- [1] Raghavendra Selvan, Max Welling, Jens Petersen, Jesper H Pedersen, and Marleen de Bruijne. Mean field network based graph refinement with application to airway tree extraction. *arXiv:1804.03348*, 2018.
- [2] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- [3] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.
- [4] Yujia Li and Richard Zemel. Mean-field networks. *arXiv preprint arXiv:1410.5884*, 2014.
- [5] KiJung Yoon, Renjie Liao, Yuwen Xiong, Lisa Zhang, and et.al. Inference in probabilistic graphical models by graph neural networks. *arXiv preprint arXiv:1803.07710*, 2018.
- [6] Thomas N Kipf and Max Welling. Variational graph auto-encoders. In *NIPS Bayesian Deep Learning Workshop*, 2016.
- [7] Raghavendra Selvan, Jens Petersen, Jesper H Pedersen, and et.al. Extraction of airways with probabilistic state-space models and Bayesian smoothing. In *Graphs in Biomedical Image Analysis*. 2017.
- [8] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- [9] Jesper H Pedersen, Haseem Ashraf, Dirksen, and et.al. The danish randomized lung cancer CT screening trial—overall design and results of the prevalence round. *Journal of Thoracic Oncology*, 4(5), 2009.
- [10] Pechin Lo, Jon Sparring, Jesper Pedersen, and et.al. Airway tree extraction with locally optimal paths. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2009.