

Diagnose Like a Radiologist: Hybrid Neuro-Probabilistic Reasoning for Attribute-Based Medical Image Diagnosis

Gangming Zhao, Quanlong Feng, Chaoqi Chen^{ID}, Zhen Zhou, and Yizhou Yu^{ID}, *Fellow, IEEE*

Abstract—During clinical practice, radiologists often use attributes, e.g., morphological and appearance characteristics of a lesion, to aid disease diagnosis. Effectively modeling attributes as well as all relationships involving attributes could boost the generalization ability and verifiability of medical image diagnosis algorithms. In this paper, we introduce a hybrid neuro-probabilistic reasoning algorithm for verifiable attribute-based medical image diagnosis. There are two parallel branches in our hybrid algorithm, a Bayesian network branch performing probabilistic causal relationship reasoning and a graph convolutional network branch performing more generic relational modeling and reasoning using a feature representation. Tight coupling between these two branches is achieved via a cross-network attention mechanism and the fusion of their classification results. We have successfully applied our hybrid reasoning algorithm to two challenging medical image diagnosis tasks. On the LIDC-IDRI benchmark dataset for benign-malignant classification of pulmonary nodules in CT images, our method achieves a new state-of-the-art accuracy of 95.36% and an AUC of 96.54%. Our method also achieves a 3.24% accuracy improvement on an in-house chest X-ray image dataset for tuberculosis diagnosis. Our ablation study indicates that our hybrid algorithm achieves a much better generalization performance than a pure neural network architecture under very limited training data.

Index Terms—Bayesian networks, deep neural networks, medical image analysis, neuro-probabilistic reasoning

1 INTRODUCTION

DUE to their rapid progress in the past ten years, deep neural networks have achieved tremendous success in boosting the performance of image recognition [18], [19], [27], [50] and other visual computing tasks [38], [48]. Such progress has also propelled forward many related research areas including medical image analysis. Nowadays, deep neural networks have been widely used for medical image analysis and diagnosis. They have demonstrated unprecedented accuracy on a variety of tasks, such as the detection of acute intracranial hemorrhage in head CT images [29], breast cancer diagnosis using mammography [39], the detection of diabetic retinopathy in retinal fundus photographs [14], and the classification of skin cancers [9].

Despite the high accuracy deep neural networks have achieved in medical image diagnosis, they still exhibit two

major limitations, insufficient generalization ability and verifiability. It is well known that deep neural networks require large amounts of training data. When high quality training data is scarce, as is typically the case in the medical image domain, the performance of trained deep neural networks on novel testing data may deteriorate significantly. In addition, deep neural networks typically produce a final result without showing how and why the result can be reached. This is undesired because very often computer-generated diagnostic results only serve as decision support for human doctors, who need to verify their credibility before accepting them.

To improve verifiability, a medical image diagnosis algorithm needs to learn the knowledge reasoning process followed by radiologists during clinical practice. They start from visual evidences in medical images and reach diagnostic conclusions by referring to causal relationships between diseases and the visual evidences. For example, radiologists use intuitive morphological and appearance characteristics of a pulmonary nodule, such as lobulation, spiculation, and texture, as visual evidences to assess whether the nodule is benign or malignant [21] (Fig. 1). In this paper, such characteristics of a lesion or, in general, radiological abnormalities in a medical image are defined as the attributes of the lesion or image. Thus, equipping a machine learning model with the capability of simultaneously predicting attributes as well as underlying diseases improves the causality and verifiability of the model now that radiologists can verify the predicted attributes as well as the knowledge reasoning process that goes from attributes to diagnostic conclusions.

In addition to causal relationships between diseases and attributes, there also exist relationships, such as co-occurrence,

- Gangming Zhao, Chaoqi Chen, and Yizhou Yu are with the Department of Computer Science, University of Hong Kong, Pokfulam, Hong Kong. E-mail: gnmzhao@connect.hku.hk, cqchen1994@gmail.com, yizhouy@acm.org.
- Quanlong Feng is with the Department of Geographical Information Engineering, China Agricultural University, Beijing 100107, China. E-mail: fengql@cau.edu.cn.
- Zhen Zhou is with AI Lab, Deepwise Healthcare, Beijing 100080, China. E-mail: zhouzhen@deepwise.com.

Manuscript received 23 Oct. 2020; revised 22 Sept. 2021; accepted 27 Sept. 2021. Date of publication 25 Nov. 2021; date of current version 3 Oct. 2022.

This work was supported in part by Beijing Municipal Science and Technology Planning Project under Grant Z201100005620008 and in part by the National Natural Science Foundation of China under Grant 81971616.

(Corresponding author: Yizhou Yu.)

Recommended for acceptance by C. Fowlkes.

Digital Object Identifier no. 10.1109/TPAMI.2021.3130759

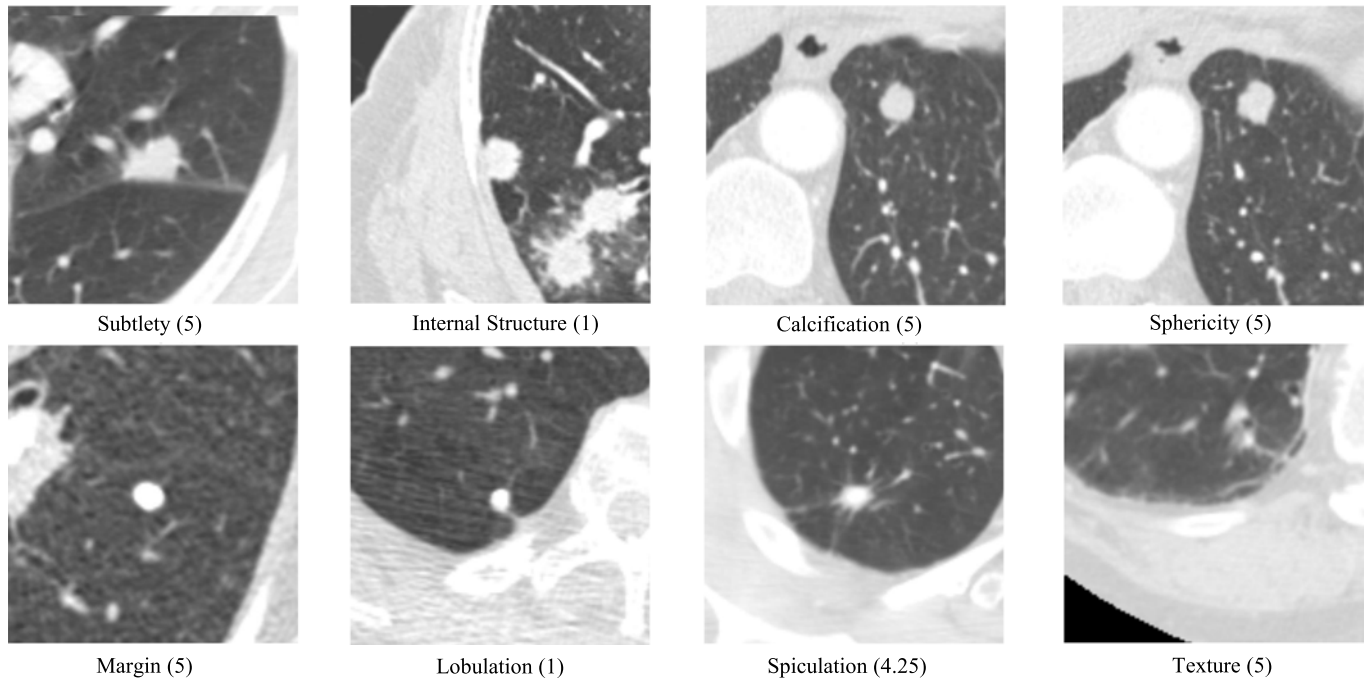


Fig. 1. Sample 2D slices of 3D pulmonary nodules from the LIDC-IDRI dataset, each with one of the following eight attributes: Subtlety, Internal Structure, Calcification, Sphericity, Margin, Lobulation, Spiculation, Texture. The annotated grade of each attribute is shown between the parentheses. For example, Subtlety (5) means its annotated grade is 5.

among the attributes themselves. Effectively modeling and exploiting the hidden relationships among these attributes as well as the relationships between attributes and the underlying diseases can boost the accuracy of attribute and disease prediction. Nonetheless, there have been few methods in medical image diagnosis exploiting such relationships.

To address generalization ability, we need to look at the strengths and limitations of both neural and probabilistic learning algorithms for relational modeling and reasoning. Examples of neural algorithms include graph neural networks [26] and relation networks [51] while examples of probabilistic learning algorithms include Bayesian networks [47] and factor graphs [11]. Neural algorithms are good at learning from large amounts of empirical data. The trained models can interpolate and approximate known information very well to reach high accuracy, but have relatively poor generalization ability, causality and explainability. On the other hand, probabilistic learning algorithms, which can be considered as a type of symbolic AI methods, can learn from smaller datasets. The trained models have better generalization ability, causality and explainability, but lower accuracy when facing real-world empirical data.

In this paper, to exploit the complementary strengths of neural and probabilistic learning algorithms as well as overcome the limitations of both ends, we introduce a hybrid neuro-probabilistic reasoning algorithm for verifiable attribute-based medical image diagnosis. Specifically, the proposed hybrid algorithm consists of multiple stages. It first employs a deep neural network backbone to extract features from an input image. Such features build the foundation of reasoning that happens in later stages. Next, a Bayesian network (BN) branch and a graph convolutional network (GCN) branch process the features extracted by the backbone and perform relational reasoning in parallel. The BN branch

first converts features into attributes as evidences, then performs probabilistic reasoning using the causal relationships between the disease and attributes, and finally produces their marginal posterior distributions. On the other hand, the GCN branch performs feature-based relational modeling. It represents attributes and generic relationships among them using a graph. The GCN branch exploits the relationships among the graph nodes to enhance the feature of every attribute. More accurate diagnoses and attribute classifications can be achieved using the enhanced features from the GCN. Although the BN and GCN belong to two parallel branches, they are not completely independent. We tightly couple the BN and GCN branches so that the BN is utilized to improve the causality and generalization ability of the GCN while the accuracy of the GCN is retained. Such tight coupling is achieved via a cross-network attention mechanism and the fusion of their classification results. By fusing the results from the two branches, causality enforced by the BN branch may be partially compromised. Therefore, a second Bayesian network module is appended at the end of the entire network to reinforce causality.

The entire network except the Bayesian network modules can be trained using gradient back-propagation as the inference phase of a Bayesian network is differentiable and gradients can be propagated through the two Bayesian networks in our algorithm. However, we notice that the compatibility between the structure of the Bayesian networks and the rest of our hybrid network significantly affects the performance of the entire network. Thus, once we learn the initial structure of the Bayesian networks using ground-truth annotations, the training procedure proceeds by alternating between two phases. The first phase updates the CNN and GCN weights using stochastic gradient descent while the second phase updates the structures and

parameters of the two Bayesian networks using the attribute classification results most recently predicted by their preceding stages in our hybrid network as training labels.

Note that verifiability is more stringent than explainability. For computer-generated diagnoses to be verifiable, the internal decision process of an algorithm not only should be explainable, but also needs to comply with the domain knowledge of human doctors. As mentioned earlier, in medical image diagnosis, radiologists start from visual evidences in medical images and reach diagnostic conclusions by referring to causal relationships between diseases and the visual evidences. Thus the algorithm needs to collect the same types of visual evidences as a doctor would do, and also learn the causal relationships between evidences and diseases as a doctor would do. The proposed algorithm in this paper models visual evidences as attributes and collects such visual evidences through attribute prediction. It further models causal and non-causal relationships using BN and GCN, which quantitatively encode such medical domain knowledge as “the presence of attribute A in this CT scan increases/decreases the likelihood that the patient has disease D” or “the occurrence of attribute A in this CT scan increases/decreases the chance that attribute B also occurs in the same scan”. Every node in these networks (BN and GCN) represents a high-level concept (a disease or an attribute), and a disease diagnosis exploits the relationships between the disease and a set of other high-level concepts. In comparison to previous medical image diagnosis methods, radiologists may find the knowledge reasoning process of the proposed algorithm more consistent with their own practice and, consequently, the disease diagnoses made by our method more credible.

We have successfully applied our hybrid reasoning algorithm to two challenging medical image diagnosis tasks. The first task is benign-malignant classification of pulmonary nodules in chest computed tomography (CT) images in the LIDC-IDRI benchmark dataset. As mentioned earlier, for pulmonary nodule classification, morphological and appearance attributes have been commonly used to assist clinical diagnosis [21]. The second task is tuberculosis (TB) diagnosis using chest X-ray images. For this task, we adopt an in-house dataset carefully annotated by senior medical experts. Every image in this TB dataset is not only annotated with a disease diagnosis, but also multiple types of radiological abnormalities, such as ‘Pulmonary Consolidation’, ‘Pulmonary Cavitation’, ‘Diffuse Nodules’, and ‘Fibrotic Appearance’, potentially caused by tuberculosis. According to the experimental results of pulmonary nodule classification on the LIDC-IDRI benchmark, our method achieves a new state-of-the-art accuracy of 95.36% and an AUC of 96.54%. On our in-house dataset for tuberculosis classification, our algorithm also achieves a 4.98% improvement in accuracy in comparison to previously best-performing algorithm for attribute-based medical image classification.

Our main contributions in this paper can be summarized as follows.

- We introduce a hybrid neuro-probabilistic reasoning algorithm for attribute-based medical image diagnosis. To support causal and verifiable relational modeling and reasoning, this algorithm tightly couples

Bayesian networks and a graph convolutional network. Such coupling is achieved via a cross-network attention mechanism and a classification result fusion scheme.

- We devise an effective training procedure for the proposed hybrid network. It alternates between two phases. The first phase updates the CNN backbone and GCN weights while the second phase updates the structures and parameters of the Bayesian networks.
- We have successfully applied the proposed hybrid reasoning algorithm to two representative medical image diagnosis tasks, benign-malignant classification of pulmonary nodules in chest CT images and tuberculosis diagnosis using chest X-ray images. The proposed algorithm achieves state-of-the-art performance on the LIDC-IDRI benchmark dataset for the first task and an in-house dataset for the second task.

2 RELATED WORK

2.1 Attribute Learning

Attributes, such as texture, color, and shape, are of great importance to describe objects. Attribute learning has been studied in computer vision for many years [1], [10], [28], [30], [33], [34], [42]. Ferrari *et al.* [10] proposed to use low-level semantic features for attribute representation and they presented a probabilistic generative model for visual attributes, together with an image likelihood learning algorithm. Human faces have many attributes, and remain a challenge for attribute learning. Kumar *et al.* [28] trained binary classifiers to recognize the presence or absence of describable aspects of facial visual appearance using traditional hand-crafted features. Liu *et al.* [37] proposed a CNN framework for face localization and attribute prediction, respectively. Attributes have also been exploited in tasks such as zero-shot learning [23], [30]. Effectively modeling the hidden relationships among attributes is useful for improving the accuracy of attribute prediction and causal association. Nonetheless, most of the early works in attribute learning did not model relationships among attributes and explore such relationships for attribute reasoning. The development of graph neural networks (GNN) [26] made it possible to learn relationships among attributes. For example, Meng *et al.* [41] used message passing to perform end-to-end learning of image representations, their relationships as well as the interplay among different attributes. They observed that relative attribute learning naturally benefits from exploiting the graph of dependencies among different image attributes. In this paper, we not only utilize a graph neural network to model the correlations among attributes, but also embed Bayesian networks into the whole framework to better model causality.

2.2 Bayesian Networks

Bayesian networks (BN), introduced by Judea Pearl [46], represent a natural approach to model causality and perform logical reasoning. Bayesian network learning includes two phases, structure learning and parameter learning. The most intuitive method for structure learning is that of ‘search and score,’ where one searches the space of directed acyclic graphs (DAGs) using dynamic programming and identifies the one that minimizes the objective function [58]. For parameter

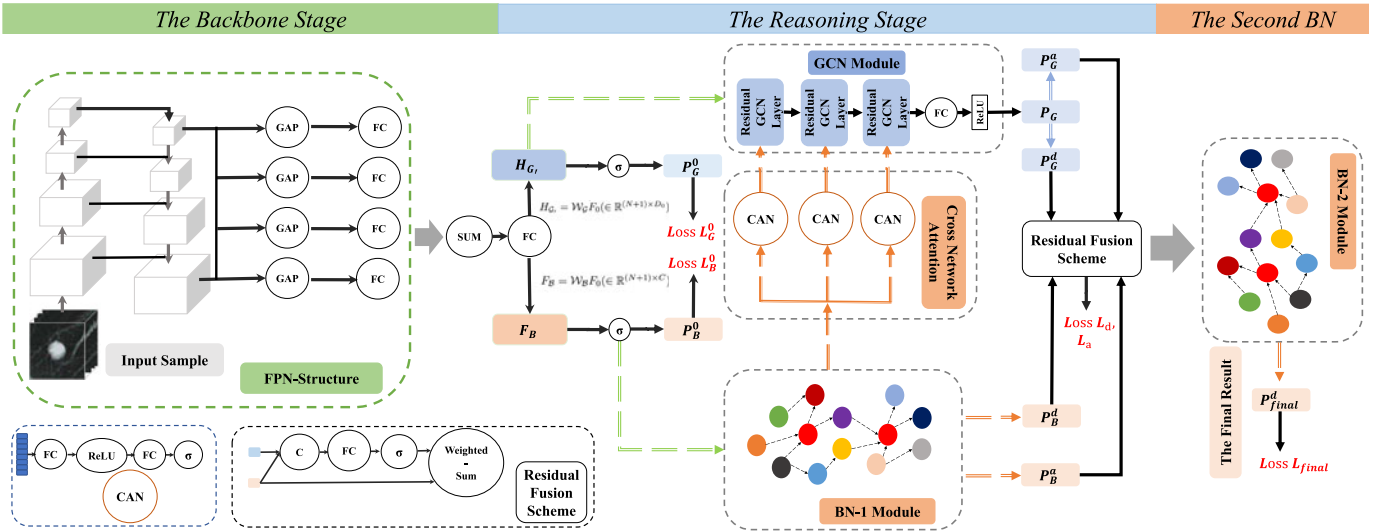


Fig. 2. Our proposed framework consists of three stages: 1) Deep feature extraction using an CNN-FPN backbone; 2) hybrid attribute relational modeling and reasoning using a pair of coupled Bayesian network and deep graph network; 3) causality reinforcement using a second Bayesian network. Note that BN-1 and BN-2 are distinct in both parameters and structures.

learning, the most frequently used method is maximum likelihood estimation (MLE) [46]. However, when given a dataset, BN cannot learn a feature representation, which limits its further development. Recently, researchers have started to focus on the integration between BN and deep learning. For example, Rohekar *et al.* [49] proposed to utilize BN models for learning better deep neural networks. Meanwhile, a few improved versions of BN have been proposed for new applications in computer vision. For example, Barik [3] improved BN using low-rank conditional probability tables. Elidan *et al.* [8] used the Copula Bayesian Network model for representing multivariate continuous distributions. The method in this paper differs from the above work in that it exploits the representation power of deep neural networks and the causality modeling capability of BN by embedding them into a unified framework.

2.3 Bayesian Neural Networks

In addition to poor interpretability, deep neural networks also suffer from overfitting on limited labeled data and the incapability of uncertainty analysis. To tackle these issues, researchers combined Bayesian methods with deep learning to construct Bayesian neural networks (BNN) for a probabilistic representation of uncertainty of the latter [12], [25], [56]. Specifically, a BNN estimates posterior distributions for the weights in a neural network. There exist a few estimation methods, such as variational inference [25] and dropout variational inference [12]. However, the proposed method in this paper differs significantly from Bayesian neural networks. Unlike BNNs, our proposed method aims to improve the causality and interpretability of deep learning via the fusion of deep learning and traditional graph-based probabilistic reasoning. It cannot perform uncertainty analysis. Nonetheless, due to the few trainable parameters in Bayesian networks, the proposed method can clearly mitigate overfitting on small training datasets, as demonstrated in one of our ablation studies.

2.4 Neuro-Symbolic Learning

The integration of connectionist (neural networks) and symbolic AI has long been a key issue in machine learning [32],

[53], [54], [65], [66]. Early works [53], [54] primarily focused on the injection of symbolic information as prior knowledge into a neural network. More recently, with the wide spread of deep learning, researchers [32], [65], [66] started to combine the merits of deep feature representation with symbolic reasoning to build strong, reliable and explainable neural-symbolic AI systems. Relevant studies have been conducted on visual question answering (VQA) [65], semantic parsing [66] and handwritten formula recognition [32]. One of the key issues in neural-symbolic learning is that most symbolic approaches are not differentiable, making end-to-end training difficult. In this study, we partially solved this issue by performing gradient back-propagation through both Bayesian networks, which serve as the symbolic reasoning module in our hybrid network, for a better feature representation and a more efficient training procedure.

3 ALGORITHM

3.1 Overview

The proposed hybrid neuro-probabilistic reasoning algorithm takes an image patch as the input and outputs a classification result, as shown in Fig. 2. The proposed algorithm consists of the following three stages.

The Backbone Stage. A backbone based on ResNet [18] or EfficientNet [57] is utilized for deep feature representation, which builds the foundation for reasoning in later stages. Moreover, a feature pyramid network (FPN) [36] is integrated into the backbone for multi-scale feature aggregation (Fig. 2). A distinct global average pooling (GAP) operator is applied to the feature map at each scale of the FPN to reduce the number of parameters. And for reducing information loss in low-level feature maps, we ensure that the dimension of the resulting feature from GAP is kept the same at different scales of the FPN. For instance, when the size of a feature map in the FPN is $64 \times 64 \times 256$, we first reshape it to $32 \times 16 \times 2048$ by concatenating the third dimensions of every 2×4 neighbors in the first two dimensions together, then use GAP to produce a 2048-dimensional feature vector. All GAP-reduced features go through parallel FC layers before they

are fused via a sum operator and a subsequent FC layer to obtain the final feature, F_0 , at the backbone stage.

The Reasoning Stage. Both a Bayesian network and a graph convolutional network are adopted for relational reasoning in parallel. The BN branch first converts deep features into attributes through attribute classification, then performs probabilistic causal reasoning using these attributes, and finally outputs marginal posterior distributions. The GCN branch first transforms the deep features from the backbone into a feature representation for individual attributes and represents the relationships among the attributes using an undirected graph, then enhances the feature of each attribute via graph convolutions. The BN and GCN are tightly coupled together through a cross-network attention mechanism and the fusion of their classification results, which will be described in Section 3.4.

The Second BN Module. To reinforce causality of the proposed hybrid algorithm, a second BN is appended at the end of the entire network.

3.2 Integration Between BN and Backbone

A Bayesian network is incorporated in the reasoning stage to model probabilistic dependencies among various attributes for disease diagnosis. Specifically, a Bayesian network, $\mathcal{B} = \langle \mathcal{V}_B, \mathcal{E}_B, \Theta \rangle$, is a directed acyclic graph (DAG) $\langle \mathcal{V}_B, \mathcal{E}_B \rangle$ with a conditional probability table (CPT) for each node, and Θ represents all parameters (CPTs), which encodes the joint probability distribution of the BN. Each node $v_i \in \mathcal{V}_B$ stands for a random variable, and a directed edge $e \in \mathcal{E}_B$ between two nodes (v_i, v_j) indicates v_i probabilistically depends on v_j . The training of a Bayesian network undergoes two phases, structure learning and parameter learning [7], [46], [47]. We adopt dynamic programming with the Bayesian information criterion (BIC) [58] for structure learning to determine the topology of the DAG and maximum likelihood estimation for parameter learning to determine Θ . Our BN module first learns its structure and then updates all parameters, i.e., the conditional probability tables at all nodes in the network. In this paper, each attribute has a corresponding node in the Bayesian network. Without loss of generality, assume there is only one disease we wish to diagnose, and that disease is an extra node in the network.

To integrate the Bayesian network with our backbone, the final feature from the backbone, F_0 , is fed to another FC-layer with weight matrix \mathcal{W}_B to obtain a feature map F_B . Thus, $F_B = \mathcal{W}_B F_0 (\in \mathbb{R}^{(N+1) \times C})$, where N is the number of attribute categories and C is the number of grades for each attribute or disease. Each attribute must have at least two grades to indicate whether that attribute is present or not. In practice, there is often a need to have more than two grades for each attribute or disease to indicate intermediate levels of certainty or severity. Note that different grades of the same attribute are mutually exclusive while different attributes are not because more than one attribute could be present at the same time. Thus, F_B can be interpreted as the concatenation of $N + 1$ C -dimensional vectors, one for each node in the Bayesian network. We further define $P_B^0 = \sigma(F_B) (\in [0, 1]^{(N+1) \times C})$, where the softmax operator σ is applied to each C -dimensional vector in F_B , and the result

represents a discrete distribution over the grades of an attribute or disease. This distribution becomes the input evidence at the corresponding node in the Bayesian network.

Once we perform inference in the Bayesian network using the conditional probability tables as well as the input evidences at all nodes, we obtain marginal posterior distributions at all nodes, including the extra node for disease diagnosis, as the output of the Bayesian network. Let the nodes in the Bayesian network be $v_i, i \in \{0, 1, \dots, n\}$. Without loss of generality, according to [46], the marginal posterior distribution $P_B(v_0)$ at node v_0 is formulated as

$$P_B(v_0) = \int \dots \int_V P(v_0, v_1, \dots, v_n) dv_1 \dots dv_n, \quad (1)$$

where

$$P(v_0, v_1, \dots, v_n) = \prod_{i=0}^n P(v_i | \text{Parents}(v_i)), \quad (2)$$

where $\text{Parents}(v_i)$ is NULL if v_i does not have any parent nodes. The equation in (2) is derived using the local Markov property of Bayesian networks. In practice, to evaluate (1), we use the belief propagation algorithm in [45].

3.3 GCN Module

In addition to the Bayesian network, we also adopt a graph convolutional network in the reasoning stage to perform feature-based relational modeling and reasoning among attributes to facilitate medical image diagnosis. Specifically, a GCN is based on an undirected graph $\mathcal{G} = (\mathcal{V}_G, \mathcal{E}_G)$, where \mathcal{V}_G and \mathcal{E}_G are the set of nodes and edges respectively. Here, there is a node corresponding to each attribute or disease, and the node holds a feature vector associated with that attribute or disease. Again, assume there is only one disease we wish to diagnose, and the total number of nodes is $N + 1$. If vertices v_i and v_j are connected in the graph, there might exist certain type of relations between them. We associate a learnable weight with every edge in the graph to indicate the strength of the connection. The GCN in our algorithm consists of L layers, all of which share the same number of nodes. But the edge weights in these layers may be different.

A general graph convolution operation \mathcal{F} at the l th layer can be formulated as follows:

$$\mathcal{F}(\mathcal{G}_l, \mathcal{W}_l) = \text{Update}(\text{Aggregate}(\mathcal{G}_l, \mathcal{W}_l^{\text{agg}}), \mathcal{W}_l^{\text{update}}), \quad (3)$$

where $\mathcal{G}_l = (\mathcal{V}_{\mathcal{G}_l}, \mathcal{E}_l)$ is the graph at the l th layer; $\mathcal{W}_l^{\text{agg}}$ and $\mathcal{W}_l^{\text{update}}$ are the learnable weights of the aggregation and update functions of the l th layer, respectively. We use a max-pooling node feature aggregator to pool the pairwise differences of features at node v_i and all of its neighbors. But the feature difference between v_i and one of its neighbors is modulated by the edge weight between them before the pooling operation. The node feature updater is a multi-layer perceptron (MLP) with batch normalization and ReLU as the activation function. The above MLP concatenates the original node feature with the aggregated feature as its input.

We adopt the residual graph convolution introduced in [31]. Let $H_{\mathcal{G}_l} \in \mathbb{R}^{N \times D_l}$ be the feature map holding the

features at all nodes of \mathcal{G}_l , where N is the number of attributes and D_l is the feature dimension in the l th layer. The residual graph convolution between the l th and $(l+1)$ th layers of the GCN is formulated as follows:

$$H_{\mathcal{G}_{l+1}} = \mathcal{F}(\mathcal{G}_l, \mathcal{W}_l) + H_{\mathcal{G}_l}. \quad (4)$$

To integrate the above GCN with our backbone, we need to obtain $N+1$ features from the final feature in the backbone, one for each node in the GCN. For this purpose, F_0 is fed to another FC-layer with weight matrix \mathcal{W}_G to obtain feature map H_G . Thus, $H_G = \mathcal{W}_G F_0 (\in \mathbb{R}^{(N+1) \times D_0})$, where N is the number of attribute categories and D_0 is the input feature dimension for every node in the GCN. H_G thus can be interpreted as the concatenation of $N+1$ D_0 -dimensional feature vectors, one for each node in the first layer of the GCN as input.

The final feature map of the GCN is $H_{G_L} \in \mathbb{R}^{(N+1) \times D_L}$. We further use a FC-layer followed by a softmax operator to obtain attribute and disease classification results in the following equation:

$$P_G = \sigma(\mathcal{W}_P H_{G_L}), \quad (5)$$

where $P_G \in [0, 1]^{(N+1) \times C}$ denotes the attribute and disease predictions.

3.4 Coupling Between BN and GCN

Although both the BN and GCN are utilized for relational reasoning, they are not completely independent. In this work, we tightly couple BN and GCN via a cross-network attention mechanism and the fusion of classification results. As for the former, BN provides attention for the node features in each layer of the GCN. Specifically, the marginal posterior distributions at all nodes of the BN are taken as the input of L node attention modules, one for each layer of the GCN to figure out the relative importance of attributes in that layer. Node attention values prescribed by the l th attention module are defined as follows:

$$M_{attn}^l = \sigma(\mathcal{W}_{l1} \text{ReLU}(\mathcal{W}_{l0} P_B(\mathcal{V}))), l \in 1, 2, \dots, L, \quad (6)$$

where $M_{attn}^l \in \mathbb{R}^{N+1}$ is the node attention vector for the l th layer of the GCN, \mathcal{V} includes all attribute and disease nodes of the BN, and $P_B(\mathcal{V}) \in \mathbb{R}^{(N+1) \times C}$ represents the concatenated marginal distributions at all nodes of the BN. Note that there is one attention value in M_{attn}^l for each node in the l th layer of the GCN. We name M_{attn}^l the spatial attention map for GCN nodes since nodes are spatial entities of a graph, similar to pixels in a CNN.

We also extend squeeze-and-excitation based channel-wise attention in SENet [20] to graph convolutional networks as follows to further enhance the features of every layer in the GCN

$$C_{attn}^l = \sigma(\mathcal{W}_{lex} \text{ReLU}(\mathcal{W}_{lsq} \text{GAP}(H_{\mathcal{G}_l}))), \quad (7)$$

where $C_{attn}^l \in \mathbb{R}^{D_l}$, $l \in 1, 2, \dots, L$ is the channel-wise attention vector for the l th layer of the GCN, \mathcal{W}_{lsq} and \mathcal{W}_{lex} are the weights for squeezing and excitation, respectively.

Once both spatial and channel-wise attention vectors have been calculated, first the same channel of all features

in \mathcal{G}_l is multiplied by its corresponding channel-wise attention value, then the entire feature at each node in \mathcal{G}_l is multiplied by its corresponding spatial attention value. Such spatially and channel-wise modulated features at the l th layer of the GCN serve as the input to the next graph convolution to produce the feature map $H_{\mathcal{G}_{l+1}}$ at the $(l+1)$ th layer.

As for the fusion of the disease classification results from BN and GCN, a residual fusion scheme is formulated as follows:

$$P_{fusion}^d = w_B \cdot P_B^d + (1 - w_B) \cdot \sigma(W_0 \text{Concat}(P_G^d, P_B^d)). \quad (8)$$

For attribute classification results, the same residual fusion scheme is adopted as follows:

$$P_{fusion}^a = w'_B \cdot P_B^a + (1 - w'_B) \cdot \sigma(W'_0 \text{Concat}(P_G^a, P_B^a)). \quad (9)$$

In these equations, P_B^d and P_B^a are the marginal posterior distributions at the disease and attribute nodes in the BN, P_G^d and P_G^a are the probabilistic disease and attribute classification results from the GCN, w_B and w'_B are two learnable trade-off coefficients, W_0 and W'_0 are the weight matrices of two fully connected layers, and $\text{Concat}()$ is the concatenation operator.

3.5 The Second BN Module

It should be noted that by fusing the results of both BN and GCN branches, causality enforced by the BN branch may be compromised to some extent. To tackle this issue, we append a second Bayesian network at the end of the entire network to reinforce causality. The fused attribute and disease soft classification results in (9) and (8) are used as the input evidences at the nodes in this second Bayesian network. The marginal posterior distribution at the disease node, denoted as P_{final}^d , can also be defined using equations similar to those in (1) and (2). In practice, it is also evaluated using the belief propagation algorithm. This marginal posterior distribution at the disease node of the second BN becomes the final disease prediction of the entire network. Note that BN-1 and BN-2 are distinct in both parameters and structures.

3.6 Training Scheme

Our entire network except the BN modules can be trained using stochastic gradient descent. Meanwhile, we observe that the inference phase of a BN is differentiable with respect to its inputs and parameters, but not its structure. That implies gradients can be backpropagated through a BN from its outputs (i.e., marginal posterior distributions) to its inputs (i.e., evidences). Parameter learning in a BN is tightly coupled with structure learning, and is driven by global statistics of the training set. Performing parameter learning separately from structure learning using stochastic gradient descent is unlikely to be very effective. Moreover, we discovered that the compatibility between the Bayesian network structures and the neural parts of our hybrid network significantly affects the overall performance. Therefore, we propose the following two alternating phases to train the entire hybrid network.

- The first phase updates the weights of the CNN backbone and the GCN branch using gradient back-propagation by fixing the structures and parameters of the two BNs. Gradients are simply backpropagated through these two BNs.
- The second phase updates the structures and conditional probability tables of the BN modules using the attribute and disease classification results most recently predicted by their preceding stages in the hybrid network as training labels. As for structure learning, dynamic programming is used with the same score function as in [58]. Once the structure of BN modules are determined, their conditional probability tables are updated immediately using maximum likelihood estimation.

It should be noted that we actually set a maximum value on the number of times the structures and parameters of the two BNs are updated (20 times in practice). Once the two phases have been alternated for such a number of times, we fix the structures and parameters of the two BNs and only update the weights in the neural parts of our hybrid network. In this way, we can guarantee the convergence of our training process.

The ground-truth annotations of the training samples are only used to train the initial structure and parameters of the BN modules. In subsequent iterations, for the first BN, the ground-truth attribute and disease annotations of a training sample are replaced with its attribute and disease classification results most recently produced by the backbone, i.e., P_B^0 ; for the second BN, they are replaced with its most recent fused attribute and disease classification results, P_{fusion}^d and P_{fusion}^a in (8) and (9). The effectiveness of this training strategy has been empirically verified. Potential reasons for the effectiveness are twofold: first, attribute classification results produced at intermediate stages of our hybrid network are soft labels, which are actually distributions, and have more complete information than the ground-truth attribute labels represented as one-hot vectors; second, using soft attribute labels generated within the hybrid network makes the trained Bayesian network structures more compatible with the neural parts of the hybrid network.

We rely on the belief propagation algorithm to perform inference in a Bayesian network and compute the marginal posterior distributions for its variables. Since we train the weights in the neural networks by propagating gradients through the Bayesian networks, we need to obtain the gradient of the belief propagation algorithm. A method for computing this gradient is described in the Appendix, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2021.3130759>.

3.6.1 Training Loss

During the first phase of the above alternating training procedure, we adopt the deep supervision strategy by distributing multiple supervision signals to various stages of our hybrid network. Most of the time, each supervision signal is primarily responsible for training a subset of the modules in the network. The overall loss function during the first phase of training is a weighted combination of five loss functions

$$\mathcal{L}_{all} = w_1 \mathcal{L}_G^0 + w_2 \mathcal{L}_B^0 + w_3 \mathcal{L}_a + w_4 \mathcal{L}_d + w_5 \mathcal{L}_{final}, \quad (10)$$

where $\sum_{i=1}^5 w_i = 1$. In our experiments, we always set $w_i = 0.2, i \in 1, 2, \dots, 5$.

The first two supervision signals corresponding to losses \mathcal{L}_B^0 and \mathcal{L}_G^0 are respectively located at the beginning of the BN and GCN branches. Two classifiers are involved. One of them is simply $P_B^0 = \sigma(F_B) \in [0, 1]^{(N+1) \times C}$, the input evidences at the nodes of the first Bayesian network. The cross-entropy loss for this classifier is written as follows:

$$\mathcal{L}_B^0 = - \sum_{i=1}^{N+1} \sum_{j=1}^C (y_{ij} \log(p_{ij}^{B,0}) + (1 - y_{ij}) \log(1 - p_{ij}^{B,0})). \quad (11)$$

The second classifier is defined as $P_G^0 = \sigma(\mathcal{W}_T H_G) \in [0, 1]^{(N+1) \times C}$, which takes H_G as the input and performs both attribute and disease classification. The cross-entropy loss for this classifier is written as follows:

$$\mathcal{L}_G^0 = - \sum_{i=1}^{N+1} \sum_{j=1}^C (y_{ij} \log(p_{ij}^{G,0}) + (1 - y_{ij}) \log(1 - p_{ij}^{G,0})). \quad (12)$$

In the above two loss functions, y_{ij} 's denote one-hot encoding of the ground-truth attribute or disease grades of training samples, $p_{ij}^{B,0}$ and $p_{ij}^{G,0}$ are components of P_B^0 and P_G^0 respectively. These two classification losses disentangle attributes from the deep features of the backbone under the supervision of ground-truth labels. These two supervision signals are primarily responsible for training the backbone as well as the connection layers between the backbone and the BN or GCN.

The second group of two supervision signals are applied to the fused results of the BN and GCN in (8) and (9). Two cross-entropy losses \mathcal{L}_a and \mathcal{L}_d are imposed again on the classification results of attributes and diseases, respectively. The first loss is imposed on the fused disease classification result, P_{fusion}^d in (8), and is formulated as follows:

$$\mathcal{L}_d = - \sum_{j=1}^C (y_j^d \log(p_j^d) + (1 - y_j^d) \log(1 - p_j^d)), \quad (13)$$

where y^d denotes one-hot encoding of the ground-truth disease grades of training samples, and p_j^d represents a component of P_{fusion}^d . The second loss \mathcal{L}_a is imposed on the fused attribute classification results, P_{fusion}^a in (9), and is formulated as follows:

$$\mathcal{L}_a = - \sum_{i=1}^N \sum_{j=1}^C (y_{ij}^a \log(p_{ij}^a) + (1 - y_{ij}^a) \log(1 - p_{ij}^a)), \quad (14)$$

where y^a denotes one-hot encoding of the ground-truth attribute grades of training samples, and p_{ij}^a represents a component of P_{fusion}^a . Since we do not update parameters in the BNs during the first phase of the training scheme, these two supervision signals are primarily responsible for training the backbone and the GCN.

The last supervision signal is located at the end of the entire hybrid network, and is responsible for training all modules in the entire network. A cross-entropy loss is imposed on the final disease classification result, P_{final}^d , defined in Section 3.5 and produced by the second Bayesian network. It is written as follows:

$$\mathcal{L}_{final} = - \sum_{j=1}^C (y_j^d \log(p_j^{final}) + (1 - y_j^d) \log(1 - p_j^{final})), \quad (15)$$

where p_j^{final} represents a component of P_{final}^d .

4 MEDICAL IMAGE DIAGNOSIS

In this section, we apply our attribute-based hybrid reasoning algorithm to two challenging medical image diagnosis tasks. The first task is benign-malignant classification of pulmonary nodules in chest computed tomography (CT) images. The second task is tuberculosis diagnosis using chest X-ray images. We define attributes and perform attribute relationship learning and reasoning for both tasks.

Following the common practice of previous attribute-based classification methods, not all attributes are used for training, i.e., we select attributes that have high causal relations with the diseases, and those unrelated attributes have been discarded.

4.1 Pulmonary Nodule Classification

Lung cancer gives rise to the most cancer-related deaths around the world [4]. Early diagnosis and treatment are of great importance to long-term survival of lung cancer patients. Benign-malignant classification of pulmonary nodules in chest CTs is vital for timely diagnosis of lung cancer [43]. In practice, radiologists often gain many clinical insights from the connections between domain knowledge and clinical symptoms as well as from the dependencies between diverse attributes and the underlying disease. This is applicable to pulmonary nodule diagnosis. For instance, when a nodule has an irregular shape and presents obvious spiculation, it has a high probability of being malignant.

4.1.1 LIDC-IDRI Dataset

The LIDC-IDRI dataset [2] from the Cancer Imaging Archive (TCIA) is one of the largest publicly available lung cancer datasets. It has 1018 clinical chest CT scans obtained from seven institutions. Each scan is associated with an XML file that details the locations of nodules on each 512×512 slice. The diameters of the nodules range from 3mm to 30mm. Each suspicious lesion is categorized as a non-nodule, a nodule < 3mm, or a nodule ≥ 3mm in diameter. Adopting the same setting as in [6], [15], [61], [62], [64], we only consider nodules ≥ 3mm in diameter since nodules < 3mm are not considered to be clinically relevant by current screening protocols [6], [15], [16], [22], [52], [55].

The malignancy of each nodule was evaluated with a 5-point scale, from benign to malignant, by up to four experienced thoracic radiologists. Following the same procedure used in previous studies [6], [15], [61], [62], [64], we select those nodules annotated by at least one radiologist and calculate the median malignancy level (MML) for each of them. Nodules with an MML less than or higher than 3 are respectively labeled as benign or malignant. To reduce the uncertainty of nodule malignancy evaluation, nodules with an MML of 3 (noted as ‘uncertain’) are excluded from our experiments as previous studies did. Thus there are a total of 1301 benign, 612 uncertain and 644 malignant nodules. The distribution of nodules over their MML is shown in Table 3. Apart from malignancy, eight attributes are also

graded for each nodule, i.e., subtlety, internal structure, calcification, sphericity, margin, lobulation, spiculation, and radiographic solidity. The grade of each attribute is between 1 and 5 except for internal structure (1-4) and calcification (1-6). We scale the grades of all attributes to 1-5 to maintain consistency, where 1 means the least obvious and 5 means the most obvious. Fig. 1 shows sample 2D slices with distinct attributes.

We also normalize all chest CTs to a unified voxel size of $1.0 \times 1.0 \times 1.0 \text{mm}^3$ using spline interpolation. In addition, following the same setting as in [61], [62], we assume the location of a nodule is the mean of the annotated centers of the nodule by radiologists. For each nodule, we crop a 64×64 patch centered on its location.

4.1.2 Experimental Setup

In all experiments, we perform 10-fold cross validation, and each model is independently trained and tested 5 times with randomly initialized weights on each fold. The overall performance of a model is assessed with several commonly used metrics, including the mean and standard deviation of accuracy, sensitivity/recall, specificity, precision with the cut-off value of 0.5, F-score and AUC (area under the receiver operator curve).

All models are trained for 160 epoches from scratch using PyTorch [44] on NVIDIA Titan X pascal GPUs while Adam [24] being the optimizer with the initial learning rate set to $1e-3$, which is reduced by a factor of 10 after every 30 epochs. The weight decay is set to $1e-4$. Both ResNet and EfficientNet are adopted as the backbone models. We have implemented their 3D versions for the LIDC dataset by revising the original 2D versions. The size of the input image patches is $64 \times 64 \times 64$, and the batch size is 6 on a single GPU. Separate models are trained with and without data augmentation, which includes standard operations including flipping, rotation, and random cropping.

4.1.3 Comparison With the State of the Art

We have compared our hybrid algorithm with several state-of-the-art models for pulmonary nodule classification on the LIDC-IDRI dataset and the results are shown in Table 1, where O2 and O2* represent our models trained without and with standard data augmentation. The proposed method achieves the best performance under every evaluation metric when data augmentation is applied while it is still ranked first under every evaluation metric except for specificity when data augmentation is not used. Such a performance demonstrates the effectiveness of hybrid neural and probabilistic reasoning proposed in this paper, as well as the important role of attribute learning in medical image diagnosis. Specifically, method A adopts a multi-crop convolutional neural network trained a subset of 825 nodules from LIDC and achieves an accuracy of 87.14%. Method B trains a 3D CNN with a subset of 1144 nodules and achieves a higher accuracy (91.26%). All methods from C to H train models using the same number of nodules, which is 1945. However, method H introduced additional 1839 unlabeled nodules for semi-supervised learning and achieves the highest accuracy (92.53%) among methods from C to H. Method I proposes a multi-scale 3D ResNet with a cost-

TABLE 1
Performance Comparison of Lung Nodule Classification Models on the LIDC-IDRI Dataset

Methods		Number		Results(%) (mean±standard deviation)					
		B	M	Accuracy	Sensitivity / Recall	Specificity	AUC	Precision	F-score
A	Shen <i>et al.</i> , 2017 [55] (Multi-crop CNN)	528	297	87.14	77.00	93.00	93.00	Not given	Not given
B	Hussein <i>et al.</i> , 2017 [21] (3D CNN)	635	509	91.26	Not given	Not given	Not given	Not given	Not given
C	Han <i>et al.</i> , 2015 [15] (3D GLCM feature+SVM)	1301	644	85.38±0.10	70.20±0.15	92.80±0.20	88.19±0.16	82.85±0.38	75.99±0.10
D	Dhara <i>et al.</i> , 2016 [6] (Multi-visual features)	1301	644	87.90±0.17	84.50±0.19	89.09±0.25	93.77±0.15	79.31±0.37	81.82±0.21
E	Xie <i>et al.</i> , 2018 [63] (Deep + visual features)	1301	644	88.73±0.15	84.40±0.20	90.88±0.13	94.02±0.20	82.09±0.24	83.23±0.21
F	Xie <i>et al.</i> , 2017 [60] (TMME with Resnet-50)	1301	644	91.01±0.10	83.83±0.15	94.56±0.13	95.35±0.15	88.40±0.24	86.07±0.15
G	Xie <i>et al.</i> , 2019a [61] (Knowledge-based)	1301	644	91.60±0.15	86.52±0.25	94.00±0.30	95.70±0.24	87.75±0.52	87.13±0.16
H	Xie <i>et al.</i> , 2019b [62] (Semi-Supervised)	1301	644	92.53±0.05	84.94±0.17	96.28±0.08	95.81± 0.19	Not given	Not given
I	Xu <i>et al.</i> , 2020 [64] (Multi-Scale Cost-Sensitive)	1156	556	92.64± 0.12	85.58 ± 0.44	95.87±1.26	94.00±0.26	90.39±0.48	87.91±0.11
T1	Low-Level-Feature [10]	1301	644	88.73±0.15	84.40±0.20	90.88±0.13	94.02±0.20	82.09±0.24	83.23±0.21
T2	Basic-visual-Feature [28]	1301	644	91.01±0.10	83.83±0.15	94.56±0.13	95.35±0.15	88.40±0.24	84.07±0.15
M1	ResNet-50	1301	644	88.14±0.23	82.17±0.14	90.77±0.15	91.21±0.14	82.18±0.11	82.36±0.14
M2	Efficient-B4	1301	644	89.21±0.12	83.83±0.24	91.18±0.22	92.05±0.27	86.88±0.19	83.04±0.23
M3	ResNet-50-FPN	1301	644	90.01±0.13	84.23±0.21	91.54±0.21	92.81±0.31	84.26±0.12	84.71±0.21
M4	Efficient-B4-FPN	1301	644	90.91±0.22	85.74±0.13	92.27±0.15	93.23±0.10	87.10±0.24	86.97±0.14
M31	ResNet-50-FPN-GCN-Relation [51]	1301	644	91.60±0.15	86.52±0.25	92.32±0.15	93.70±0.24	86.75±0.52	85.13±0.16
M41	Efficient-B4-FPN-GCN-Relation [51]	1301	644	92.15±0.12	86.97±0.23	93.13±0.11	93.89±0.23	87.14±0.21	85.57±0.24
M32	ResNet-50-FPN-GRU [41]	1301	644	91.41±0.11	86.12±0.14	92.92±0.19	93.61±0.23	87.44±0.21	85.22±0.15
M42	Efficient-B4-FPN-GRU [41]	1301	644	92.21±0.10	86.94±0.11	93.89±0.24	93.72±0.12	88.24±0.13	86.63±0.21
O1	Our-ResNet-50-FPN	1301	644	93.74±0.17	89.23±0.21	95.76±0.24	96.12±0.12	94.21±0.14	88.27±0.31
O2	Our-Efficient-B4-FPN	1301	644	95.31±0.15	90.51±0.15	96.15±0.22	96.47±0.31	95.95±0.24	88.83±0.45
O2*	Our-Efficient-B4-FPN*	1301	644	95.36±0.10	91.01±0.16	96.47±0.12	96.54±0.32	95.96±0.21	89.13±0.15

'B' and 'M' are the Numbers of Benign and Malignant Lung Nodules.

sensitive loss and delivers a higher performance (92.64%) than method H using a subset of 1712 nodules. In comparison, when adopting the same backbone (3D ResNet) as method I, our method (O1) achieves a better performance (93.74%), indicating the importance of attribute relationship modeling in enhancing the performance of pulmonary nodule classification. Most importantly, Table 1 shows that when using an EfficientB4-FPN backbone, our model with data augmentation (O2*) achieves the highest accuracy of 95.36%, and our model without data augmentation (O2) gains 4.41% over the baseline using the same backbone (M4).

In addition, Table 2 shows that under the training setting and testing protocol adopted in [59], our algorithm also achieves the highest accuracy, further demonstrating the effectiveness of our method. In this comparison, we adopt the same backbone (DenseNet) as in [59], and do not apply any data augmentation during training. During testing, we adopt the same "off-by-one" accuracy as in [59], which considers attribute/malignancy grading within ± 1 of the ground truth as correct results. Note that although we do

not exploit the additional supervision signal provided by the nodule segmentation masks, that are used in [59], our method still achieves the highest accuracy of 98.12%.

We have further compared with attribute relationship modeling methods including M31, M41, M32, and M42. For a fair comparison, methods in the same comparison always adopt the same backbone, which is either 3D ResNet-50 or EfficientNet-B4. The relational network module proposed in [51] and the GRU module proposed in [41] respectively achieve an improvement of 1.59% and 1.30% in accuracy when compared with the 3D ResNet-50 baseline while our method improves the accuracy of the same baseline by 3.63%. The same conclusion can be drawn for the EfficientNet-B4 backbone. We have also compared with two classic attribute learning methods T1 and T2, and they deliver a lower accuracy of 88.73% and 91.01%, respectively. In summary, all the above well-designed experiments demonstrate that the proposed hybrid neuro-probabilistic reasoning algorithm achieves clearly better performance than existing methods on the LIDC-IDRI dataset.

Sample input images and their associated attribute and disease classification results from our method and a few other baseline methods are shown in Fig. 3. In addition, we provide the importance of individual attributes for pulmonary nodule classification, i.e., quantitatively evaluating how important individual attributes are in making the prediction. To be specific, we measure how much the final disease classification probability is affected by deactivating

TABLE 2
Comparison With Existing Lung Nodule Classification Models Under a Second Training and Testing Protocol

Methods	Accuracy %	# Nodules
TumorNet [22]	82.47	1145
TumorNet-attribute [22]	92.31	1145
SHC-DCNN [5]	82.40	1432
MCNN [55]	86.84	1100
CNN-MTL [21]	91.26	1340
PN-SAMP-S1 [59]	92.03	1404
PN-SAMP-S2 [59]	95.30	1404
PN-SAMP-M [59]	97.58	1404
Ours	98.12	1404

TABLE 3
Distribution of Median Malignancy Levels (MML) in the LIDC-IDRI Dataset for Lung Nodule Classification

Dataset	Benign		Uncertain		Malignant
MML	1	2	3	4	5
# Nodules	358	943	612	474	170

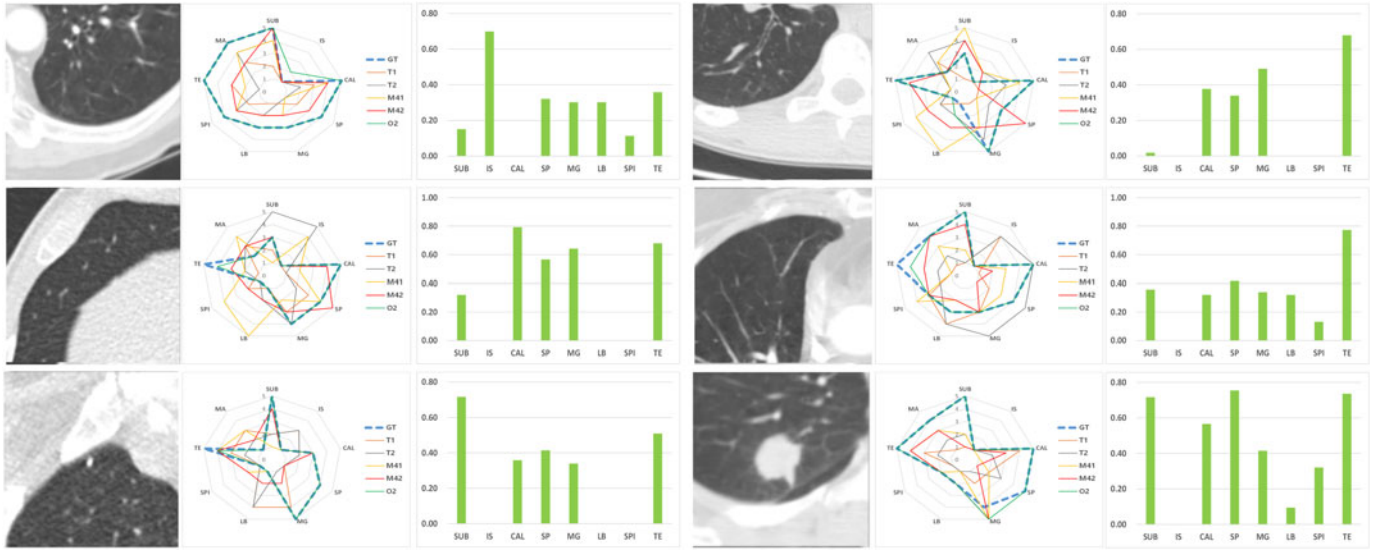


Fig. 3. Sample results of our method (O2) and four baseline methods (T1, T2, M41, M42) on the LIDC-IDRI dataset. Short names are defined as follows: SUB (Subtlety), IS (Internal Structure), CAL (Calcification), SP (Sphericity), MG (Margin), LB (Lobulation), SL (Spiculation), TE (Texture), and MA (Malignancy). The radar charts visualize the actual grades of attributes and disease, and there exist five possible grades for each attribute or disease. GT stands for the ground truth. The Histograms denote the numerical importance of different attributes for diagnosing malignancy (MA).

one attribute at a time in the input signals to BN-1 and GCN. For the input signals to GCN, an attribute can be deactivated by setting its corresponding feature vector to an average feature for that attribute over the samples where that attribute attains the lowest possible grade. For the input signals to BN-1, since each attribute in the LIDC-IDRI dataset is associated with 5 grades, deactivation means assigning the lowest grade (1) to the attribute.

4.1.4 Ablation Study

We have performed a systematic ablation study on the LIDC dataset to verify the effectiveness of the components in our framework. The results are shown in Table 4. Specifically, we have explored the effects of both the first and second BN (BN-1 and BN-2), cross-network attention and residual fusion (CNA-RES), SE channel-wise attention (SEatt) [20], gradient backpropagation through BN (GradBN), BN and DNN alternating training (AlterTrain) and the Relational Network Module (Relation) [51] on the EfficientNet-B4 backbone. Table 4 shows that all proposed components or strategies contribute

to the final performance of our algorithm. More specifically, we have found out that ‘CNA-RES’ is the most important component, the performance would drop by 1.67% if it was removed, which verifies the effectiveness of our proposed network coupling strategy. When both ‘BN-2’ and ‘CNA-RES’ are removed and the final result of the network is computed by a simple average of the results from BN-1 and GCN, the network becomes an ensemble model integrating the BN-1 and GCN branches. The performance drops significantly by 3.2% under this ensemble model setting. This result indicates that various components in our hybrid framework are highly coupled, and altogether they deliver a performance much better than a simple ensemble model. ‘GradBN’ and ‘AlterTrain’ also have a significant impact, and the performance drops by 1.5% and 1.0% respectively if either of them is removed. In addition, completely removing BN-1 and its coupling with GCN would decrease the performance by 2.90% while removing BN-2 alone would decrease the performance by 0.57%. Moreover, the performance would drop by 3.16% when both BN modules are removed and a relational network is appended after the GCN module. This result demonstrates

TABLE 4
Ablation Study of Our Classification Algorithm on the LIDC-IDRI Dataset

GCN	BN-2	CNA-RES	SEatt	BN-1	GradBN	AlterTrain	Relation	Accuracy	Sensitivity / Recall	Specificity	AUC	Precision	F-score
✓	✓	✓	✓	✓	✓	✓	✗	95.31±0.15	90.51±0.15	96.15±0.22	96.47±0.31	95.95±0.24	88.83±0.45
✓	✗	✓	✓	✓	✓	✓	✗	94.74±0.09 ^{10.57}	88.92±0.13	95.86±0.41	95.83±0.09	92.54±0.06	88.04±0.04
✓	✗	✗	✓	✓	✓	✓	✗	92.11±0.14 ^{13.20}	88.14±0.22	93.74±0.15	94.91±0.26	88.11±0.09	86.88±0.12
✗	✗	✗	✗	✓	✓	✓	✗	91.21±0.22 ^{4.10}	87.26±0.14	93.21±0.13	93.47±0.54	88.01±0.09	86.07±0.56
✓	✓	✗	✓	✓	✓	✓	✗	93.64±0.11 ^{11.67}	88.01±0.15	95.03±0.13	95.02±0.13	91.42±0.13	87.92±0.04
✓	✓	✓	✗	✓	✓	✓	✗	94.01±0.02 ^{11.30}	88.64±0.08	95.23±0.07	95.32±0.11	91.11±0.04	87.13±0.09
✓	✓	✗	✓	✗	✓	✓	✗	92.41±0.14 ^{12.90}	88.14±0.22	93.74±0.15	94.91±0.26	92.11±0.09	86.88±0.12
✓	✓	✓	✓	✓	✗	✓	✗	93.81±0.10 ^{11.50}	88.04±0.32	93.74±0.15	94.91±0.26	90.21±0.07	86.96±0.10
✓	✓	✓	✓	✓	✓	✗	✗	94.31±0.10 ^{11.00}	88.28±0.41	94.23±0.12	94.87±0.76	89.11±0.22	87.03±0.20
✓	✗	✗	✗	✗	✗	✗	✓	92.15±0.12 ^{13.16}	86.97±0.23	93.13±0.11	93.89±0.23	87.14±0.21	85.57±0.24
✓	✗	✗	✗	✗	✗	✗	✗	91.31±0.12 ^{14.00}	87.83±0.24	93.15±0.13	94.07±0.12	86.65±0.12	87.04±0.13

‘BN-1’ and ‘BN-2’ Denote the First and Second BN Modules, ‘CNA-RES’ Denotes Cross-Network Attention and Residual Fusion, ‘SEatt’ Means Adopting SE Channel-Wise Attention in the GCN, ‘GradBN’ Means Gradient Backpropagation through the BN Modules, ‘AlterTrain’ Denotes the Proposed Alternating Training Strategy. ‘Relation’ Means Removing Both BN Modules and Appending the Relational Network from [51] after the GCN Module.




Fig. 4. The Bayesian network structure learned from the LIDC-IDRI dataset. Notations are defined as follows. SPI: spiculation; MAL: malignancy; SUB: subtlety; CAL: calcification; MAR: margin; LOB: lobulation.

the strong dependency modeling ability of BN modules. Fig. 4 shows the structure of the first BN module (BN-1) learned from the LIDC-IDRI dataset. This learned BN structure indicates that the node for malignancy has one parent node (i.e., spiculation) and four children nodes (i.e., subtlety, calcification, margin and lobulation).

We can draw the following conclusions from the results given in Table 4,

- ‘BN-1’, ‘CNA-RES’, ‘GradBN’, ‘AlterTrain’ and ‘SENet’ are important to the classification performance on the LIDC-IDRI dataset. Note that removing ‘AlterTrain’ actually means using ground-truth attribute labels

instead of soft attribute labels for training the BN modules because the BN structures will not change if we always use ground-truth attribute labels during alternating training.

- BN modules are vital for attribute relational modeling. Their performance surpasses the performance of the advanced relational network.
- Although a simple ensemble model consisting of BN and GCN can already improve the performance, our proposed cross-network attention mechanism takes full advantage of BN and GCN, thus performs better.

4.2 Tuberculosis Diagnosis

Chest radiographs are a type of medical images that can be conveniently acquired for disease diagnosis. Radiologists can find out many diseases quickly via observing chest X-ray images, which are useful for early diagnosis and intervention. Tuberculosis [35] has the highest mortality around the globe among infectious diseases. However, if it was detected at an early stage from chest radiographs, the death rate could decrease by 70% [17]. Meanwhile, tuberculosis is a type of bacterial infection that can give rise to multiple types of radiological abnormalities in chest radiographs such as diffuse nodules and fibrotic appearance. Just like morphological characteristics of

TABLE 5
Distribution of Disease and Radiological Abnormalities in the TB-Xatt Dataset for Tuberculosis Diagnosis

Abnormality/ Disease	Fibrotic Streaks	Pulmonary Consolidation	Diffuse Nodules	Pulmonary Cavitation	Atelectasis	Multiple Nodules	Pleural Effusion	Pulmonary Tuberculosis
#Images	1,640	1,200	1,500	1,400	800	1,200	900	1,972



Fibrotic Streaks



Pulmonary Consolidation



Diffuse Nodules



Pulmonary Cavitation



Atelectasis



Multiple Nodules



Pleural Effusion



Pulmonary Tuberculosis

Fig. 5. Sample chest X-ray images used for building the TB-Xatt dataset, each with one of the following seven attributes except the last one: Fibrotic Streaks, Pulmonary Consolidation, Diffuse Nodules, Pulmonary Cavitation, Atelectasis, Multiple Nodules, and Pleural Effusion. The last image shows a case of the disease, Pulmonary Tuberculosis.




Fig. 6. Construction of image samples in the TB-Xatt dataset. First, lung segmentation in chest X-ray images using a trained segmentation model. Second, resizing the left and right lung patches to the same resolution and placing them side by side.

lung nodules, these radiological abnormalities can be treated as attributes to facilitate the diagnosis of tuberculosis by medical experts. Therefore, we choose tuberculosis diagnosis as an additional task to further evaluate the proposed hybrid neuro-probabilistic reasoning algorithm, which is able to model the dependencies between tuberculosis and radiological abnormalities to improve the reasoning capability of deep learning methods.

4.2.1 TB-Xatt Dataset

As for the task of tuberculosis diagnosis, there are a total of 14200 frontal chest X-ray images in an in-house dataset named TB-Xatt. The retrospective study on this dataset has been approved by the institutional review board of West China Hospital. Every image in this dataset has been carefully

annotated by certificated and experienced radiologists. Image annotations include 7 types of radiological abnormalities, 'Pulmonary Consolidation', 'Pulmonary Cavitation', 'Diffuse Nodules', 'Fibrotic Streaks', 'Atelectasis', 'Multiple Nodules', 'Pleural Effusion' and one disease, 'Pulmonary Tuberculosis'. On average, there are 3 occurrences of disease or radiological abnormalities per image and 1326 annotated images for each type of disease or radiological abnormalities. The distribution of images over disease and radiological abnormalities is presented in Table 5. In this study, we aim to improve the classification performance of 'Pulmonary Tuberculosis' with the help of the seven types of radiological abnormalities, which are treated as the attributes of a chest X-ray image in our hybrid algorithm since these seven types of radiological abnormalities are the main expressions of active tuberculosis in the lung. Fig. 5 shows sample images with these attributes and disease from the TB-Xatt dataset. Note that the image samples in the dataset are pre-processed versions of the original X-ray images. We first perform lung segmentation in the original X-ray images and extract bounding boxes of the left and right lungs as shown in Fig. 6, then resize these bounding boxes to 512×256 and place each pair of left and right bounding boxes side by side as a single image sample. In all experiments, we perform 10-fold cross validation, and each model is independently trained and tested 5 times with randomly initialized weights on each fold, the same as we did on LIDC. The performance of a model is also assessed with the same set of metrics as LIDC, including accuracy, sensitivity/recall, specificity, precision, F-score and AUC.

4.2.2 Experimental Setup

Different from LIDC, the ResNet and EfficientNet backbones are pre-trained on ImageNet. Adam is also the optimizer with the initial learning rate set to $1e-3$, and the batch size is 32. All models are trained 60 epochs and the learning rate is reduced by a factor of 10 after 20 and 40 epochs.

4.2.3 Comparison With the State of the Art

On our in-house TB-Xatt dataset for tuberculosis diagnosis, we have also compared our hybrid algorithm with existing

TABLE 6
Performance Comparison of Tuberculosis Diagnosis Models on the TB-Xatt Dataset

Methods		Results(%) (mean±standard deviation)					
		Accuracy	Sensitivity / Recall	Specificity	AUC	Precision	F-score
A1	Low-Level-Feature [10]	86.14±0.14	83.11±0.26	89.17±0.10	90.11±0.76	81.10±0.12	81.71±0.11
A2	Basic-visual-Feature [28]	87.23±0.12	84.16±0.27	92.00±0.09	91.27±0.63	83.26±0.71	83.44±0.17
S1	Attention-Guide [13]	93.12±0.10	86.17±0.16	91.22±0.34	93.67±0.43	87.21±0.53	85.64±0.29
S2	ADINet [40]	93.43±0.19	87.24±0.25	91.87±0.11	94.11±0.29	87.44±0.54	84.21±0.13
M1	ResNet-50	90.17±0.11	90.44±0.12	89.16±0.23	90.03±0.21	82.09±0.12	84.15±0.23
M2	Efficient-B4	93.21±0.25	91.54±0.22	92.55±0.11	92.67±0.19	84.27±0.98	87.23±0.09
M3	ResNet-50-FPN	91.23±0.28	91.96±0.42	91.22±0.33	91.54±0.07	83.27±0.08	85.26±0.44
M4	Efficient-B4-FPN	94.19±0.19	92.78±0.09	92.97±0.01	94.01±0.10	86.11±0.02	89.17±0.08
M31	ResNet-50-FPN-GCN-Relation [51]	91.74±0.29	92.16±0.32	93.23±0.25	92.26±0.12	86.14±0.11	86.54±0.13
M41	Efficient-B4-FPN-GCN-Relation [51]	95.17±0.11	92.29±0.18	94.01±0.12	94.83±0.21	86.77±0.14	90.21±0.19
M32	ResNet-50-FPN-GRU [41]	91.66±0.13	92.06±0.34	94.26±0.26	93.13±0.12	86.01±0.21	87.11±0.98
M42	Efficient-B4-FPN-GRU [41]	94.65±0.32	92.94±0.18	95.19±0.13	94.88±0.23	86.98±0.12	91.03±0.24
O1	Our-ResNet-50-FPN	94.67±0.11	94.22±0.33	96.25±0.23	97.22±0.43	90.29±0.22	90.01±0.10
O2	Our-Efficient-B4-FPN	97.13±0.19	95.51±0.10	97.11±0.16	98.10±0.34	92.44±0.31	91.78±0.27
O2*	Our-Efficient-B4-FPN*	98.41±0.10	96.32±0.11	97.91±0.21	98.91±0.12	93.24±0.10	92.77±0.14

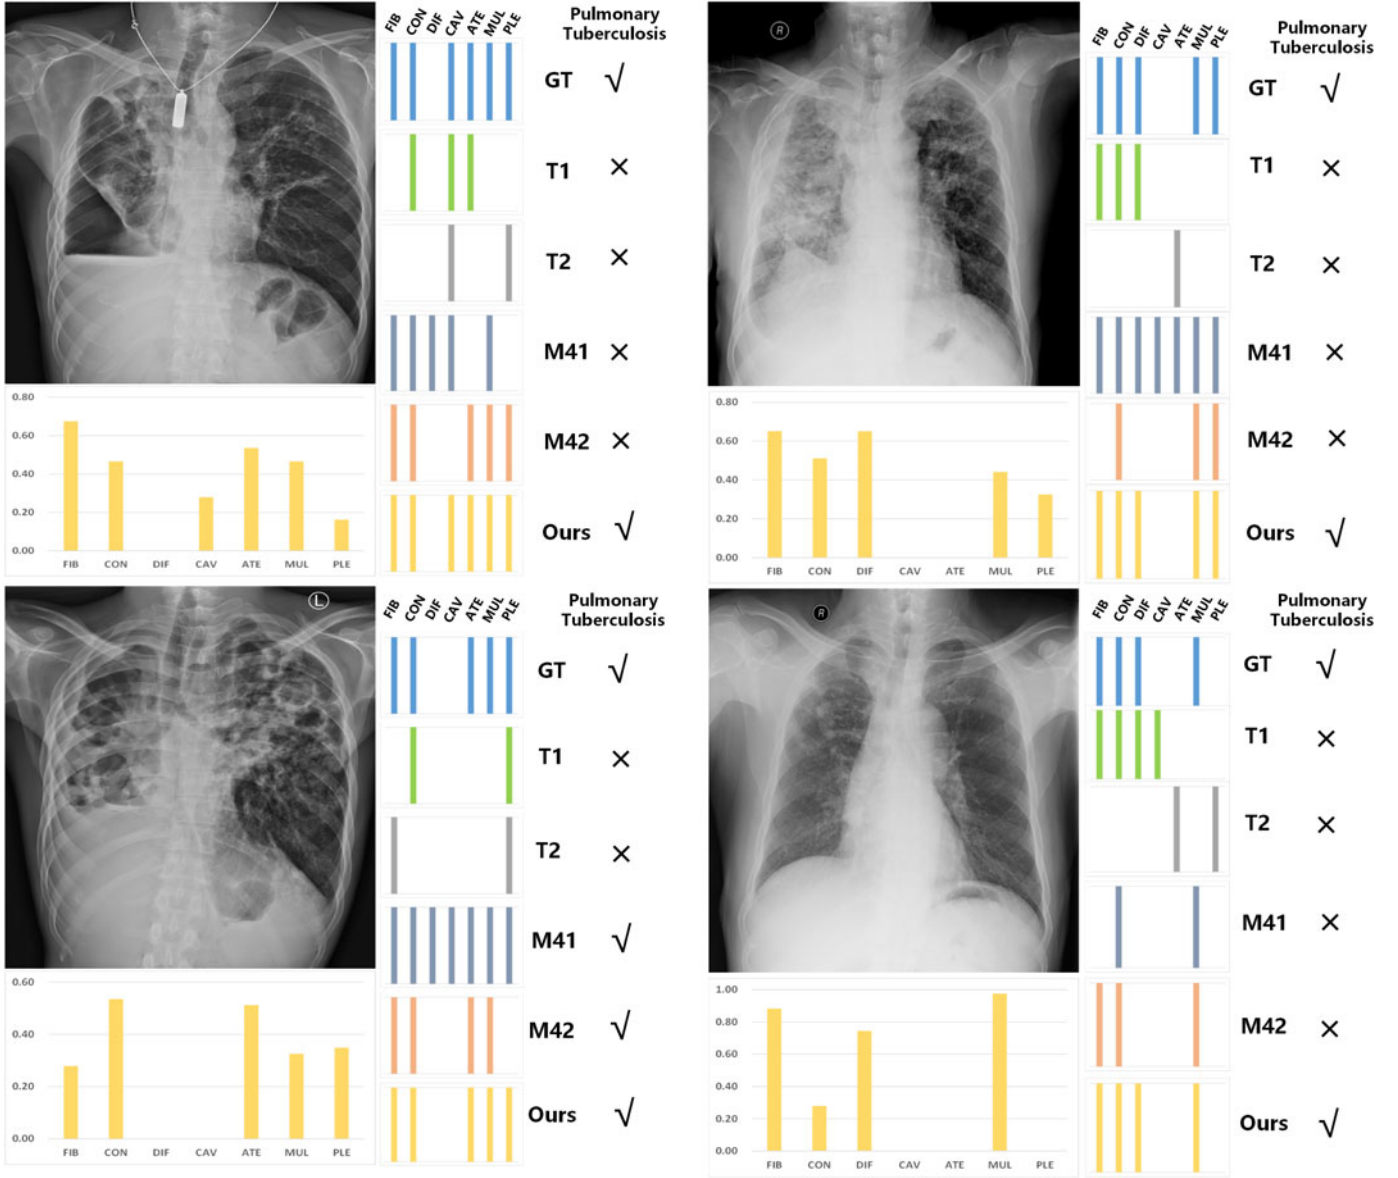


Fig. 7. Sample results of our method (O2) and four baseline methods (T1, T2, M41, M42) on the in-house TB-Xatt dataset. Short names are defined as follows: FIB (Fibrotic Streaks), CON (Pulmonary Consolidation), DIF (Diffuse Nodules), CAV (Pulmonary Cavitation), ATE (Atelectasis), MUL (Multiple Nodules), and PLE (Pleural Effusion). The colored bars visualize attribute classification results from individual methods as well as the ground truth. GT stands for the ground truth. ✓ and ✗ visualize the diagnoses of pulmonary tuberculosis from individual methods as well as the ground truth. The histograms under the X-ray images visualize the importance of individual attributes for diagnosing pulmonary tuberculosis.

methods, including two state-of-the-art models (S1 and S2), two attribute learning methods (A1 and A2) and two relationship modeling methods (M31, M41 and M32, M42). The results are shown in Table 6, where O2 and O2* represent our models trained without and with standard data augmentation, respectively. No matter whether data augmentation is used, the proposed algorithm achieves the best performance under all evaluation metrics, including accuracy, sensitivity/recall, specificity, precision with the cut-off value of 0.5, F-score and AUC. Specifically, both A1 and A2 are classic attribute learning methods that use hand-crafted features, and achieve an accuracy of 86.14% and 87.23%, respectively. When compared with state-of-the-art disease classification models S1 and S2, the proposed hybrid algorithm demonstrates better performance with 4.01% and 3.70% improvement in accuracy, respectively. Furthermore, as in the

pulmonary nodule classification task, we implemented two attribute relationship modeling methods (M31, M41 and M32, M42) based on the ResNet-50 and EfficientNet-B4 backbones. Results indicate that our method achieves better performance due to its strong causality modeling and relationship reasoning capabilities. The classification performance of our method with ResNet-50-FPN and EfficientNet-B4-FPN backbones are 94.67% and 97.13%, which respectively gain 3.44% and 2.94% over the two baselines using the same backbones. The results in Table 6 demonstrate the effectiveness of our proposed algorithm on the tuberculosis diagnosis task. Sample input images and their associated attribute and disease classification results from our method and a few other baseline methods are shown in Fig. 7, where we have binary attribute and disease classifications. We also evaluate the importance of individual

TABLE 7
Ablation Study of Our Classification Algorithm on the TB-Xatt Dataset

GCN	BN-2	CNA-RES	SEatt	BN-1	GradBN	AlterTrain	Relation	Accuracy	Sensitivity / Recall	Specificity	AUC	Precision	F-score
✓	✓	✓	✓	✓	✓	✓	✗	97.13±0.19	95.51±0.10	97.11±0.16	98.10±0.34	92.44±0.31	91.78±0.27
✓	✗	✓	✓	✓	✓	✓	✗	96.78±0.11 ^{10.35}	94.26±0.13	96.78±0.23	97.12±0.11	90.93±0.17	90.03±0.07
✓	✓	✗	✓	✓	✓	✓	✗	96.21±0.44 ^{10.57}	94.12±0.31	96.22±0.41	97.17±0.41	90.61±0.20	90.01±0.20
✗	✗	✗	✗	✓	✓	✓	✗	95.69±0.12 ^{11.09}	93.58±0.17	95.44±0.32	96.71±0.21	90.01±0.10	89.96±0.24
✓	✓	✓	✗	✓	✓	✓	✗	96.89±0.32 ^{10.24}	94.47±0.55	96.79±0.42	97.87±0.31	91.02±0.33	90.64±0.25
✓	✓	✗	✓	✗	✓	✓	✗	95.69±0.14 ^{11.44}	94.55±0.43	94.78±0.11	95.94±0.12	87.11±0.23	91.21±0.54
✓	✓	✓	✓	✓	✗	✓	✗	96.12±0.24 ^{10.66}	94.66±0.31	94.51±0.23	95.17±0.21	90.14±0.32	90.11±0.23
✓	✓	✓	✓	✓	✓	✗	✗	96.53±0.13 ^{10.60}	94.71±0.24	94.62±0.33	95.19±0.47	90.28±0.34	90.51±0.33
✓	✗	✗	✗	✗	✗	✗	✓	94.97±0.21 ^{12.81}	93.09±0.08	93.46±0.29	94.93±0.21	86.77±0.14	90.21±0.19

'BN-1' and 'BN-2' Denote the First and Second BN Modules, 'CNA-RES' Denotes Cross-Network Attention and Residual Fusion, 'SEatt' Means Adopting SE Channel-Wise Attention in the GCN, 'GradBN' Means Gradient Backpropagation through the BN Modules, 'AlterTrain' Denotes the Proposed Alternating Training Strategy. 'Relation' Means Removing Both BN Modules and Appending the Relational Network from [51] After the GCN Module.

TABLE 8
Performance Comparison of Our Tuberculosis Diagnosis Models Trained With Increasingly Smaller Datasets

Methods		Results(%) (mean±standard deviation)					
		Accuracy	Sensitivity / Recall	Specificity	AUC	Precision	F score
P1	Low-Level-Feature [10]	86.14±0.14	83.11±0.26	89.17±0.10	90.11±0.76	81.10±0.12	81.71±0.11
	Basic-visual-Feature [28]	87.23±0.12	84.16±0.27	92.00±0.09	91.27±0.63	83.26±0.71	83.44±0.17
	Efficient-B4-FPN-GCN-Relation [51]	94.97±0.21	93.09±0.08	93.46±0.29	94.93±0.21	86.77±0.14	90.21±0.19
	Efficient-B4-FPN	94.19±0.19	92.78±0.09	92.97±0.01	94.01±0.10	86.11±0.02	89.17±0.08
	Our-Efficient-B4-FPN	97.13±0.19^{12.94}	95.51±0.10	97.11±0.16	98.10±0.34	92.44±0.31	91.78±0.27
P2	Low-Level-Feature [10]	84.23±0.12	80.07±0.04	86.22±0.20	86.78±0.15	78.24±0.09	78.16±0.27
	Basic-visual-Feature [28]	83.45±0.22	81.07±0.54	89.21±0.12	87.13±0.01	80.55±0.44	79.21±0.03
	Efficient-B4-FPN-GCN-Relation [51]	91.26±0.34	91.11±0.25	90.24±0.27	91.45±0.67	84.12±0.03	87.64±0.11
	Efficient-B4-FPN	91.21±0.14	90.16±0.22	89.98±0.34	92.12±0.22	82.56±0.13	86.26±0.17
	Our-Efficient-B4-FPN	93.21±0.44 ^{12.11}	91.02±0.33	91.20±0.21	93.56±0.17	87.65±0.12	89.43±0.14
P3	Low-Level-Feature [10]	80.11±0.09	75.67±0.12	80.32±0.55	81.26±0.64	72.13±0.07	72.51±0.62
	Basic-visual-Feature [28]	79.14±0.24	78.27±0.63	82.11±0.10	80.65±0.44	76.17±0.23	74.56±0.35
	Efficient-B4-FPN-GCN-Relation [51]	87.24±0.12	86.10±0.65	87.21±0.46	87.13±0.55	81.21±0.46	83.27±0.35
	Efficient-B4-FPN	86.25±0.22	84.23±0.27	84.22±0.54	82.45±0.32	79.15±0.21	80.24±0.55
	Our-Efficient-B4-FPN	92.44±0.65 ^{16.19}	90.01±0.21	89.26±0.17	89.22±0.10	86.14±0.25	85.23±0.17
P4	Low-Level-Feature [10]	76.21±0.09	79.27±0.31	74.32±0.45	77.26±0.13	66.45±0.16	67.84±0.17
	Basic-visual-Feature [28]	75.67±0.54	77.15±0.13	76.55±0.43	76.21±0.22	72.56±0.34	71.27±0.21
	Efficient-B4-FPN-GCN-Relation [51]	80.11±0.01	78.20±0.15	81.31±0.11	80.13±0.24	75.56±0.31	76.27±0.45
	Efficient-B4-FPN	78.11±0.12	79.26±0.36	78.23±0.34	80.15±0.24	76.27±0.15	73.62±0.14
	Our-Efficient-B-FPN	86.98±0.27 ^{18.87}	87.23±0.19	87.63±0.24	88.11±0.12	85.23±0.11	80.16±0.47

The Amount of Training Data in P1-P4 Has the Following Ratio, 1: 0.75 : 0.5: 0.25.

attributes for pulmonary tuberculosis diagnosis. Similarly, we measure how much the final disease classification probability is affected by deactivating one attribute at a time in the input signals to BN-1 and GCN. In this case, for the input signals to BN-1, deactivation means setting the classification probability of a specific attribute to zero. The deactivation scheme for the input signals to GCN is the same as that for the LIDC-IDRI dataset.

4.2.4 Ablation Study

We have also conducted an ablation study on the TB-Xatt dataset to verify the effectiveness of the components in our proposed framework. The results are shown in Table 7. As in the pulmonary nodule classification task, EfficientNet-B4 is the chosen backbone, and performance is measured when individual components are removed or replaced. Similar patterns in performance have been found here as well, and

all the considered components contribute to the final performance of our hybrid algorithm.

Moreover, we have also tested the performance of our algorithm under the condition of a limited training set by sampling an increasingly smaller subset of the TB-Xatt dataset as the training set while keeping the size of the validation and testing sets unchanged. The sampling results are shown in Table 9, where P1 means all original training samples, and P2, P3 and P4 mean 75%, 50% and 25% of the original training samples, respectively. For a fair comparison, we perform stratified sampling and ensure the same ratio between the number of positive and negative samples under all settings. This ratio is always 1/8 for the training sets, and 1/3 for the validation and testing sets. All results are shown in Table 8, which indicates that the proposed algorithm can still perform better than the baselines when given a decreasing number of training samples. Note that all baselines adopt EfficientNet-B4 as the backbone since we empirically find

TABLE 9
Specification of Datasets Used for the Performance Comparison in Table 8

Dataset	P1	P2	P3	P4
#Training Samples	11360	8530	5680	2840
#Validation Samples	1420	1420	1420	1420
#Testing Samples	1420	1420	1420	1420
P/N ratio in Train	1/8	1/8	1/8	1/8
P/N ratio in Val	1/3	1/3	1/3	1/3
P/N ratio in Test	1/3	1/3	1/3	1/3

TABLE 10
Classification Accuracy of Input and Output Signals of BN-1 and GCN on TB-Xatt Dataset

Method	P1	P4
BN-1 input	91.04	78.22
BN-1 output	93.10	85.44
GCN input	92.12	78.23
GCN output	96.36	80.28
Residual Fusion (GCN and BN-1)	97.07	86.10

out it is the best backbone under different settings. Furthermore, the performance gap between our algorithm and its corresponding baseline widens when the number of training samples decreases. When only 25% of the original training samples are used, this performance gap reaches 8.87% for the EfficientNet-B4 backbone. More surprisingly, although the number of training samples used in P4 is only half of that in P3, the performance of our models trained in the P4 setting is comparable and sometimes even better than the performance of corresponding baseline models in the P3 setting. These results demonstrate the strong generalization ability of our models trained using a small dataset.

We further investigate the classification performance of the individual output signals from the BN-1 and GCN modules to show how the two mechanisms work collaboratively when the size of training data varies. As shown in Table 10, (1) when the training dataset is small, BN-1 achieves better performance than GCN; (2) when the training dataset is large, GCN achieves better performance than BN-1. The justification is that (1) when the training dataset is small, BN-1 probabilistically models the causal relationships between attributes and diseases to mitigate the problem of insufficient data; (2) when the training dataset is large, GCN has the capability to learn expressive and powerful feature representations from the large-scale annotated dataset. Moreover, we notice that our residual fusion scheme always achieves higher accuracy than both individual networks (BN-1 and GCN). It has become clear that the strong generalization ability of our trained models is made possible by the complementarity between our neural and probabilistic reasoning algorithms. The neural branch in our network realizes its full power when the training set is large while the probabilistic branch prevents overfitting from happening when the training set is small as probabilistic models are better at learning from smaller datasets.

5 CONCLUSION AND DISCUSSION


In this paper, we have introduced a hybrid neuro-probabilistic reasoning algorithm for verifiable attribute-based medical image diagnosis. There are two parallel branches in our hybrid algorithm, a Bayesian network branch performing probabilistic causal relationship reasoning and a graph convolutional network branch performing more generic relational modeling and reasoning using a feature representation. Tight coupling between these two branches is achieved via a cross-network attention mechanism and residual fusion of classification results. We train the hybrid network by alternatively updating neural network weights and the structure/parameters of the Bayesian networks. We have successfully applied our hybrid reasoning algorithm to two challenging medical image diagnosis tasks, benign-malignant classification of pulmonary nodules in chest computed tomography images and tuberculosis diagnosis using chest X-ray images. On the LIDC-IDRI benchmark dataset for pulmonary nodule classification, our method achieves a new state-of-the-art accuracy of 95.36% and an AUC of 96.54%. Our method also achieves a 2.94% accuracy improvement on the in-house tuberculosis diagnosis dataset. Our ablation study indicates that our hybrid algorithm achieves a much more robust performance than a pure neural network architecture under very limited training data.

REFERENCES

- [1] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid, "Label-embedding for attribute-based classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2013, pp. 819–826.
- [2] S. G. Armato 3rd et al., "The lung image database consortium (LIDC) and image database resource initiative (IDRI): A completed reference database of lung nodules on CT scans," *Med. Phys.*, vol. 38, no. 2, pp. 915–931, 2011.
- [3] A. Barik and J. Honorio, "Learning Bayesian networks with low rank conditional probability tables," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 8964–8973.
- [4] F. Bray, J. Ferlay, I. Soerjomataram, R. L. Siegel, L. A. Torre, and A. Jemal, "Global cancer statistics 2018: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries," *Cancer J. Clin.*, vol. 68, no. 6, pp. 394–424, 2018.
- [5] M. Buty, Z. Xu, M. Gao, U. Bagci, A. Wu, and D. J. Mollura, "Characterization of lung nodule malignancy using hybrid shape and appearance features," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Interv.*, 2016, pp. 662–670.
- [6] A. K. Dhara, S. Mukhopadhyay, A. Dutta, M. Garg, and N. Khandelwal, "A combination of shape and texture features for classification of pulmonary nodules in lung CT images," *J. Digit. Imag.*, vol. 29, no. 4, pp. 466–475, 2016.
- [7] D. Eaton and K. Murphy, "Bayesian structure learning using dynamic programming and MCMC," in *Proc. 23rd Conf. Uncertainty Artif. Intell.*, 2007, pp. 101–108.
- [8] G. Elidan, "Copula Bayesian networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2010, pp. 559–567.
- [9] A. Esteva et al., "Correction: Corrigendum: Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 546, no. 7660, pp. 686–686, 2017.
- [10] V. Ferrari and A. Zisserman, "Learning visual attributes," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2008, pp. 433–440.
- [11] B. J. Frey, F. R. Kschischang, H.-A. Loeliger, and N. Wiberg, "Factor graphs and algorithms," in *Proc. Annu. Allerton Conf. Commun. Control Comput.*, 1997, pp. 666–680.
- [12] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1050–1059.
- [13] Q. Guan, Y. Huang, Z. Zhong, Z. Zheng, L. Zheng, and Y. Yang, "Thorax disease classification with attention guided convolutional neural network," vol. 131, pp. 38–45, 2020.

- [14] V. Gulshan *et al.*, "Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs," *JAMA*, vol. 316, no. 22, pp. 2402–2410, 2016.
- [15] F. Han *et al.*, "Texture feature analysis for computer-aided diagnosis on pulmonary nodules," *J. Digit. Imag.*, vol. 28, no. 1, pp. 99–115, 2015.
- [16] F. Han *et al.*, "A texture feature analysis for diagnosis of pulmonary nodules using LIDC-IDRI database," in *Proc. IEEE Int. Conf. Med. Imag. Phys. Eng.*, 2013, pp. 14–18.
- [17] P. Hart and I. Sutherland, "BCG and vole bacillus vaccines in the prevention of tuberculosis in adolescence and early adult life," *Brit. Med. J.*, vol. 2, no. 6082, pp. 293–295, 1977.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [19] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 630–645.
- [20] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7132–7141.
- [21] S. Hussein, K. Cao, Q. Song, and U. Bagci, "Risk stratification of lung nodules using 3D CNN-based multi-task learning," in *Proc. Int. Conf. Inf. Process. Med. Imag.*, 2017, pp. 249–260.
- [22] S. Hussein, R. Gillies, K. Cao, Q. Song, and U. Bagci, "TumorNet: Lung nodule characterization using multi-view convolutional neural network with Gaussian process," in *Proc. IEEE 14th Int. Symp. Biomed. Imag.*, 2017, pp. 1007–1010.
- [23] H. Jiang, R. Wang, S. Shan, Y. Yang, and X. Chen, "Learning discriminative latent attributes for zero-shot classification," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 4233–4242.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Representations*, 2015.
- [25] D. P. Kingma, T. Salimans, and M. Welling, "Variational dropout and the local reparameterization trick," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2015, pp. 2575–2583.
- [26] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. 5th Int. Conf. Learn. Representations*, 2017.
- [27] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [28] N. Kumar, A. C. Berg, P. N. Belhumeur, and S. K. Nayar, "Attribute and simile classifiers for face verification," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, 2009, pp. 365–372.
- [29] W. Kuo, C. Hane, P. Mukherjee, J. Malik, and E. L. Yuh, "Expert-level detection of acute intracranial hemorrhage on head computed tomography using deep learning," *Proc. Nat. Acad. Sci. USA*, vol. 116, no. 45, pp. 22 737–22 745, 2019.
- [30] C. H. Lampert, H. Nickisch, and S. Harmeling, "Attribute-based classification for zero-shot visual object categorization," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 3, pp. 453–465, Mar. 2014.
- [31] G. Li, M. Müller, A. Thabet, and B. Ghanem, "DeepGCNs: Can GCNs go as deep as CNNs?," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 9266–9275.
- [32] Q. Li, S. Huang, Y. Hong, Y. Chen, Y. N. Wu, and S.-C. Zhu, "Closed loop neural-symbolic learning via integrating neural perception, grammar parsing, and symbolic reasoning," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 5884–5894.
- [33] K. Liang, H. Chang, B. Ma, S. Shan, and X. Chen, "Unifying visual attribute learning with object recognition in a multiplicative framework," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 7, pp. 1747–1760, Jul. 2019.
- [34] K. Liang, Y. Guo, H. Chang, and X. Chen, "Incomplete attribute learning with auxiliary labels," in *Proc. 26th Int. Joint Conf. Artif. Intell.*, 2017, pp. 2252–2258.
- [35] C.-H. Lin *et al.*, "Tuberculosis mortality: Patient characteristics and causes," *BMC Infect. Dis.*, vol. 14, no. 1, 2014, Art. no. 5.
- [36] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 936–944.
- [37] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 3730–3738.
- [38] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 3431–3440.
- [39] S. M. McKinney *et al.*, "International evaluation of an AI system for breast cancer screening," *Nature*, vol. 577, no. 7788, pp. 89–94, 2020.
- [40] Q. Meng and S. Shin'ichi, "ADINet: Attribute driven incremental network for retinal image classification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 4032–4041.
- [41] Z. Meng, N. Adluru, H. J. Kim, G. Fung, and V. Singh, "Efficient relative attribute learning using graph neural networks," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 552–567.
- [42] W. Min, S. Mei, L. Liu, Y. Wang, and S. Jiang, "Multi-task deep relative attribute learning for visual urban perception," *IEEE Trans. Image Process.*, vol. 29, pp. 657–669, Aug. 2019.
- [43] D. E. Ost and M. K. Gould, "Decision making in patients with pulmonary nodules," *Amer. J. Respir. Crit. Care Med.*, vol. 185, no. 4, pp. 363–372, 2012.
- [44] A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 8024–8035.
- [45] J. Pearl, "Reverend bayes on inference engines: A distributed hierarchical approach," in *Proc. 2nd AAAI Conf. Artif. Intell.*, 1982, pp. 133–136.
- [46] J. Pearl, "Bayesian networks," in *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA, USA: MIT Press, 1998, pp. 149–153.
- [47] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Amsterdam, The Netherlands: Elsevier, 2014.
- [48] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [49] R. Y. Rohekar, S. Nisimov, Y. Gurwicz, G. Koren, and G. Novik, "Constructing deep neural networks by Bayesian network structure learning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2018, pp. 3047–3058.
- [50] O. Russakovsky *et al.*, "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, 2015.
- [51] A. Santoro *et al.*, "A simple neural network module for relational reasoning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 4967–4976.
- [52] A. A. A. Setio *et al.*, "Pulmonary nodule detection in CT images: False positive reduction using multi-view convolutional networks," *IEEE Trans. Med. Imag.*, vol. 35, no. 5, pp. 1160–1169, May 2016.
- [53] J. W. Shavlik, "A framework for combining symbolic and neural learning," *Mach. Learn.*, vol. 14, pp. 321–331, Mar. 1994.
- [54] J. W. Shavlik, R. J. Mooney, and G. G. Towell, "Symbolic and neural learning algorithms: An experimental comparison," *Mach. Learn.*, vol. 6, pp. 111–143, 1991.
- [55] W. Shen, M. Zhou, F. Yang, C. Yang, and J. Tian, "Multi-scale convolutional neural networks for lung nodule classification," in *Proc. Int. Conf. Inf. Process. Med. Imag.*, 2015, pp. 588–599.
- [56] Z. Shi *et al.*, "Bayesian VoxDRN: A probabilistic deep voxelwise dilated residual network for whole heart segmentation from 3D MR images," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Interv.*, 2018, pp. 569–577.
- [57] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019.
- [58] E. Wit, E. V. D. Heuvel, and J.-W. Romeijn, "All models are wrong...: An introduction to model uncertainty," *Statistica Neerlandica*, vol. 66, no. 3, pp. 217–236, 2012.
- [59] B. Wu, Z. Zhou, J. Wang, and Y. Wang, "Joint learning for pulmonary nodule segmentation, attributes and malignancy prediction," in *Proc. IEEE 15th Int. Symp. Biomed. Imag.*, 2018, pp. 1109–1113.
- [60] Y. Xie, Y. Xia, J. Zhang, D. D. Feng, M. Fulham, and W. Cai, "Transferable multi-model ensemble for benign-malignant lung nodule classification on chest CT," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Interv.*, 2017, pp. 656–664.
- [61] Y. Xie *et al.*, "Knowledge-based collaborative deep learning for benign-malignant lung nodule classification on chest CT," *IEEE Trans. Med. Imag.*, vol. 38, no. 4, pp. 991–1004, Apr. 2019.
- [62] Y. Xie, J. Zhang, and Y. Xia, "Semi-supervised adversarial model for benign-malignant lung nodule classification on chest CT," *Med. Image Anal.*, vol. 57, pp. 237–248, 2019.
- [63] Y. Xie, J. Zhang, Y. Xia, M. Fulham, and Y. Zhang, "Fusing texture, shape and deep model-learned information at decision level for automated classification of lung nodules on chest CT," *Inf. Fusion*, vol. 42, pp. 102–110, 2018.
- [64] X. Xu *et al.*, "MSCS-DeepLN: Evaluating lung nodule malignancy using multi-scale cost-sensitive neural networks," *Med. Image Anal.*, vol. 65, 2020, Art. no. 101772.

- [65] K. Yi, J. Wu, C. Gan, A. Torralba, P. Kohli, and J. B. Tenenbaum, "Neural-symbolic VQA: Disentangling reasoning from vision and language understanding," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2018, pp. 1039–1050.
- [66] P. Yin, C. Zhou, J. He, and G. Neubig, "StructVAE: Tree-structured latent variable models for semi-supervised semantic parsing," in *Proc. 56th Annu. Meeting Assoc. Comput. Linguistics*, 2018, pp. 754–765.



Gangming Zhao received the BS degree from the School of Mathematical Sciences, Anhui University, Hefei, China, in 2016, and the MS degree from the Center for Research on Intelligent Perception and Computing (CRIPAC), University of Chinese Academy of Sciences (UCAS), Beijing, China in 2019. He is currently working toward the PhD degree in the Department of Computer Science, University of Hong Kong, Hong Kong. Before that, he was a machine learning researcher with the AI Lab of Deepwise Healthcare.

His current research interests include computer vision, medical image analysis, and deep learning.



Quanlong Feng received the BS and MS degrees in measuring and control technology and instruments from the University of Science and Technology Beijing, Beijing, China, in 2010 and 2012, respectively, and the PhD degree in cartography and geography information system from the University of Chinese Academy of Sciences (UCAS), Beijing, China in 2016. He is currently an associate professor with the College of Land Science and Technology, China Agricultural University. His current research interests include using machine and

deep learning methods for scene understanding, semantic segmentation and object detection from multi-source and multi-sensor images.



Chaoqi Chen received the BS and ME degrees from Xiamen University, Xiamen, China in 2017 and 2020, respectively. He is currently working toward the PhD degree in the Department of Computer Science, University of Hong Kong, Hong Kong. Before that, he was a machine learning researcher with the AI Lab of Deepwise Healthcare. His research spans computer vision and machine learning, with special interests in deep learning, transfer learning, and their visual applications.



Zhen Zhou received the BS degree from North China Electric Power University, Beijing, China in 2012, and the PhD degree in pattern recognition and intelligent systems from the Institute of Automation, Chinese Academy of Sciences, Beijing, China in 2017. Afterwards, he was a post-doctoral researcher working on medical image analysis with Peking University. He is currently a principle researcher with the AI Lab, Deepwise Healthcare. His research interests include the theory and applications of deep learning, and intelligent medical image analysis.



Yizhou Yu (Fellow, IEEE) received the PhD degree from the University of California at Berkeley, Berkeley, California, in 2000. He is a professor with the University of Hong Kong, and was a faculty member with the University of Illinois at Urbana-Champaign for 12 years. He is a recipient of 2002 U.S. National Science Foundation CAREER Award and ACCV 2018 Best Application Paper Award. He has served on the editorial board of the *IET Computer Vision*, *The Visual Computer*, and *IEEE Transactions on Visualization and Computer Graphics*. He has also

served on the program committee of many leading international conferences, including CVPR, ICCV, and SIGGRAPH. His current research interests include computer vision, deep learning, AI for medicine, and geometric computing.

► **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.**