Contents lists available at ScienceDirect



Computerized Medical Imaging and Graphics



journal homepage: http://www.elsevier.com/locate/compmedimag

Deep Reinforcement Learning Framework for Thoracic Diseases Classification via Prior Knowledge Guidance

Weizhi Nie^a, Chen Zhang^a, Dan Song^{a,*}, Lina Zhao^b, Yunpeng Bai^{e,f}, Keliang Xie^{b,c,d}, Anan Liu^a

^aSchool of Electrical and Information Engineering, Tianjin University, Tianjin 300072, China

^{Vb}Department of Critical Care Medicine, Tianjin Medical University General Hospital, Tianjin 300052, China

C Department of Anesthesiology, Tianjin Medical University General Hospital, Tianjin 300052, China

🗖 Tianjin Institute of Anesthesiology, Tianjin 300052, China

^eDepartment of Cardiac Surgery, Chest Hospital, Tianjin University, Tianjin 300222, China

Clinical school of Thoracic, Tianjin Medical University, Tianjin 300052, China

ARTICLE INFO

Article history: Received 1 May 2013 Received 1 May 2013 Accepted 13 May 2013 Available online 15 May 2013 Communicated by S. Sarkar 2000 MSC: 41A05, 41A10, 65D05, 65D17 Keywords: Deep Reinforcement Learning, Medical Image Processing, Chest X-ray Images, Thoracic Diseases Classification

ABSTRACT

The chest X-ray is often utilized for diagnosing common thoracic diseases. In recent years, many approaches have been proposed to handle the problem of automatic diagnosis based on chest X-rays. However, the scarcity of labeled data for related diseases still poses a huge challenge to an accurate diagnosis. In this paper, we focus on the thorax disease diagnostic problem and propose a novel deep reinforcement learning framework, which introduces prior knowledge to direct the learning of diagnostic agents and the model parameters can also be continuously updated as the data increases, like a person's learning process. Especially, 1) prior knowledge can be learned from the pre-trained model based on old data or other domains' similar data, which can effectively reduce the dependence on target domain data, and 2) the framework of reinforcement learning can make the diagnostic agent as exploratory as a human being and improve the accuracy of diagnosis through continuous exploration. The method can also effectively solve the model learning problem in the case of few-shot data and improve the generalization ability of the model. Finally, our approach's performance was demonstrated using the well-known NIH ChestX-ray 14 and CheXpert datasets, and we achieved competitive results. The source code can be found here: https://github.com/NeaseZ/MARL.

© 2023 Elsevier B. V. All rights reserved.

1. Introduction

The chest X-ray (CXR) is one of the most commonly used clinical examinations in clinical scenarios. Early screening of many diseases relies on CXR data, such as heart disease, lung infection, throat examination, etc. This is because CXR data is easier to obtain and costs less. Traditional diagnosis and treat-

*Corresponding author

e-mail: dan.song@tju.edu.cn (Dan Song),

ment rely heavily on the experience of full-time doctors. Effective automated diagnosis and treatment technology can greatly improve diagnosis and treatment efficiency while also assisting doctors in improving the accuracy of early patient diagnosis and treatment. However, the sparsity of medical data and the diversity of patients' clinical outcomes have been major challenges in the development of automated diagnosis and treatment technology.

In recent years, many proposed approaches, like Fu et al. (2021); Abbas et al. (2021) have handled the CXR image diagnosis problem as a multi-label classification problem. Their

oliverwhite@126.com (Yunpeng Bai)



Fig. 1. The motivation of our approach: Each diagnosis can be based on the doctor's experience which can be seen as the utilization of the doctor's prior knowledge. Meanwhile, the diagnosis can also increase or update the experience of the doctor, which then helps to make a more accurate diagnosis in the future.

work relies on the large scale of training data and parameter selection. These approaches are difficult to improve further because of the high inner similaritySarvamangala and Kulkarni (2021) between classes in CXR, their traditional convolutional network-based methods cannot distinguish well. Furthermore, many medical image datasets are difficult and expensive to achieve on a large scale, so improving data use efficiency is critical. In recent years, some researchers have focused on how to achieve high classification accuracy based on the fewshot learning technique. Cherti.et al.Cherti and Jitsev (2021) combined pre-training and transfer learning to large-scale natural and medical images datasets and verify the effectiveness of their method by some experiments. Singh.et al.Singh et al. (2021) formulated a few-shot learning problem and presented a meta-learning-based model, which can adapt to rare disease classes with the availability of a few images. Furthermore, some methods consider utilizing external knowledge as auxiliary information to improve the performance of the model. Chenet al. Chen et al. (2022) proposed a multi-label annotation framework based on a external medical knowledge graph and model relevant lesion labels more comprehensively, but the professional medical knowledge graph is always hard to obtain and difficult to generalize to other tasks.

1.1. Motivation

Based on the above analysis, we hope the automatic diagnosis model could consider the context of a medical image and also utilize some additional prior knowledge like that of a professor doctor. Meanwhile, we also hope the automatic diagnosis model can be updated like the learning process of a doctor based on the increase in clinic data, like Fig.1. Thus, the learning process can be divided into two parts: 1) Obtain experience and update: the experience can be obtained from the other domain or the target domain data. The experience can be seen as the pre-trained module's output. 2) Make diagnosis: The module makes the final diagnosis based on the current medical image and also according to the prior knowledge learned from old data or other domain's data.

Here, each diagnosis is influenced by the previous experience (previous module's parameter), and each diagnosis can also increase or update the experience, which is very similar to the human's knowledge growing process. Obviously, this process can be seen as a Markov decision process (MDP). Thus, it is natural that the reinforcement learning (RL) framework can be used to represent or estimate this process. However, we also need to solve two questions below:

- How to learn and save the prior knowledge based on limited image data. The prior knowledge is a very abstract theory, so we need to define the format of the prior knowledge, which should be represented and utilized for helping to make the diagnosis, then ease the problem of catastrophic forgetting.
- How to utilize the prior knowledge to guide the training of the diagnosis model and inject the knowledge to the model. We should consider fusing the prior knowledge into the diagnosis model and make the right diagnosis.

In this paper, we propose a novel multi-agent reinforcement learning (MARL) framework to solve the common thorax medical image classification problem. We design multiple prior knowledge agents to obtain the prior knowledge and the diagnostic agent to make the final diagnosis based on the actions of prior knowledge agents. The incremental data in the framework will also be used to update the prior knowledge learning and extraction module, which can be learned from source or target domain data and output in the same visual and embedding formats as the prior knowledge. The diagnostic agent can be seen as the main agent in our proposed framework. It will be used to handle the diagnostic problem of thoracic diseases based on prior knowledge and input medical images.

We introduce three agents in this work. 1) Semantic agent: This agent is used to provide an initial embedding for classification and to guide the learning of parameters. 2) Visual agent: This agent is used to get the coarse region of interest information in the image. We hope that it will be able to provide visual information as a prior knowledge format. 3) Diagnostic agent: This agent is used to make the final diagnosis by combining the visual information and the semantic embedding. Meanwhile, the reinforcement learning framework can make the action selection process more exploratory, so when our proposed framework faces a tough case, i.e., an easily-confused case, it could make a bolder diagnostic choice and may perform better. The performance of our technique is demonstrated using the famous NIH ChestX-ray14 and CheXpert datasets.

1.2. Contribution

The key contributions of our work are followed as:

- We propose a novel multi-agent reinforcement learning framework that can address the abnormality classification problem. This universal framework can introduce the prior knowledge to guide the diagnostic agent to improve the final performance of diagnosis;
- We propose a uniform multi-information fusion module based on a transformer to solve the correlation among agents. We also propose a exploration-added model to handle the multi-label classification problem and high inter-class similarity problem;

• We validated the effectiveness of our method based on some popular datasets. Several current efficient methods are used for comparison, the final experimental findings show that our method is superior;

The remainder of our paper is outlined below. In Section 2, we present the related work. The proposed solution is presented in Section 3. In Section 4, we present key experiments. This section also includes a summary of the experimental results, and we will demonstrate the effectiveness of our approach by using it to solve a variety of thorax diseases classification problems. Implementation details are in Section 4.3. Finally, in Section 5, we draw a conclusion for this work and outline future work possibilities. The supplementary materials are attached.

2. Related Works

2.1. CXR Image Classification

The release of some large-scale CXR datasets with more than one hundred thousand images, like NIH ChestX-ray14Wang et al. (2017) drives the development of data hungry deep learning for the task of CXR image analysis. Besides the CXR dataset we mentioned above, CheXpertIrvin et al. (2019) is another large-scale chest radiograph dataset that is widely used in CXR image processing. Rocha et al. Rocha et al. (2022) trained a spatial transformer network on CheXpert, without the need for localization labels. Liu et al.Liu et al. (2019) argued that traditional deep learning approaches dealing with the presence of potentially misaligned or unrelated objects throughout the CXR image may lead to unnecessary noise, and the restriction of image resolution may lead to the loss of image details so it is difficult to detect the pathology with small lesion area, so they proposed a framework trained with high resolution images and utilizing domain knowledge at the same time. Saleem et al.Saleem et al. (2021) used a transfer learning technique to detect tuberculosis and achieved a decent result. Zhu et al. Zhu et al. (2022) argued that existing deep networks typically use the global mean pooling layer to generate features for classifiers, but the relative size, absolute size, and location of occurrence may limit classification performance. So they proposed a pixelwise classification and attention network, which can ameliorate the above problems. Minaee et al. Minaee et al. (2020) used different types of CNN-based backbones to detecte lung lesions in COVID-19 patients, while Park et al. Park et al. (2022) utilized transformer-based method to deal with similar problems. Ke et al.Ke et al. (2021) investigated different ImageNetDeng et al. (2009) pretrained architectures' performance and parameter efficiency on CheXpert. Paul et al. Paul et al. (2021) proposed a few-shot CXR diagnostic method and introduced a saliencybased classifier to extract features from the output of the CNN and then classify. In this paper, to solve the CXR image classification task, we hope to propose a model as simple as possible without degrading the classification performance.

2.2. Multi-label Classification

Multi-label classification is a relatively more challenging task compared with single-label classification, in which each sample may have more than one associated label. Recently, multi-label classification has attracted much research attention. Zhang *et al.*Zhang et al. (2020) proposed a unified deep learning framework for the scenario of multi-label unknown image classification, whose results are comparable to most relevant methods. Class imbalance problem is the main challenge of multi-label classification, which means partly categories occur more frequently in the data space than othersTarekegn et al. (2021), Jain *et al.*Jain et al. (2017) noticed that category imbalance is common in clinic diagnosis and classifier tends to neglect the important impact of the minority category, therefore they proposed an algorithm to help biased classifier to perform well on minority class.

How to model relations between labels is another challenge in the multi-label classification task, as natural images are intrinsically and frequently relevant and co-occurrence. The author, inSun and Lee (2017) proposed a novel framework to explore the unseen relationship between labels and address the class imbalance problem at the same time. Wu *et al.* Wu *et al.* (2018) mentioned a cost-sensitive model to tackle the issue that traditional multi-label classification tasks usually ignore the label correlations by using binary relevanceBoutell *et al.* (2004) to find associations between labels. The cost-sensitive loss in the model deals with the label imbalance problem at the same time. Chen *et al.*Chen *et al.* (2019b) considered the cooccurrence of different objects in different images and proposed a framework to model the label relationship based on Graph Convolution Network (GCN)Kipf and Welling (2016).

2.3. Transformer In Vision Tasks

TransformerVaswani et al. (2017) was initially applied in natural language processing(NLP) tasks, Devlin et al. Devlin et al. (2018) proposed a model named BERT based on Transformer and obtained a prominent progress on 11 NLP tasks. Kant et al.Kant et al. (2018) proposed a model to classify the multiemotion sentiment by training a Transformer encoder-decoder architecture, Vila et al. Vila et al. (2018) proposed an attentionbased end-to-end Transformer model to translate Spanish to English and performed well. Recently, the huge application prospect of Transformer has been developed for computer vision tasks. Dosovitskiy et al. Dosovitskiy et al. (2020) proposed the Vision Transformer (ViT), which is directly applied to image patches to classify the images. Liu et al. Liu et al. (2021b) proposed a novel hierarchical vision Transformer model named Swin Transformer, which could be used as an alternative option when choosing a computer vision task's backbone. Carion et al.Carion et al. (2020) proposed an end-to-end object detection model named DETR based on Transformer and achieve a prominent detection result.

2.4. Reinforcement Learning for Medical Image Processing

Reinforcement Learning simply means that agents choose the appropriate behavior in order to maximize the reward in a specific dynamic environment which can naturally be applied in games such as atariMnih et al. (2013). Besides, reinforcement learning for medical image processing has developed rapidly in recent years, Dou *et al*.Dou et al. (2019) proposed a RL

based framework to deal with the fetal brain tissue location task. Liao et al.Liao et al. (2020) proposed a relative model with cross-entropy gain based reward to realise 3D medical image segmentation task by treating each voxel as an agent, making it a MARL problem naturally. Considered that the processing of medical images might require large number of labeled images, Stember et al. Stember and Shalu (2020) proposed a RL based method by training a deep Q network(DQN)Mnih et al. (2013) to detect brain lesions on MRI. Ghesu et al. Ghesu et al. (2016) proposed an agent learning model based on RL for anatomical landmark detection in medical images by considering image parsing as a policy encoding problem. Hou et al. Hou et al. (2021) proposed a CXR report-generation framework using cascaded encoders, decoders, and a reward module. The proposed method combined adversarial training with RL and considered both results and literature fluency. In order to overcome the obstacles of traditional machine learning methods in dealing with anatomy detection, such as the use of computationally sub-optimal search schemes, Ghesu et al. Ghesu et al. (2017) proposed a RL-based framework combined with multiscale analysis that considered the 3D detection task as a behavior learning task, training a well-designed agent to detect the abnormal area in CT images.

3. Our Approach

In this section, we first introduce the problem definition and illustrate the details of our approach in the next subsections.

3.1. Problem Definition

In this study, the diagnostic process can be seen as MDP then handled using the RL framework. In each stage of diagnosis, the doctor should make an accurate diagnosis based on the medical image and the doctor's prior knowledge like Fig.1. Based on this analysis, we apply the multi-agent reinforcement learning framework, which includes multiple prior agents and one main agent. The framework is shown in Fig.2. The prior agents are used to extract the prior knowledge and the main agent make the final diagnosis based on prior agents' actions. The details of definitions are followed as:

- Agent: We propose a prior knowledge agent set, Q_p = {Q_p¹, Q_p², ..., Q_pⁿ}. They can provide different formats of prior knowledge to help the main agent make an accurate diagnosis. The main agent is used to judge the disease categories; Q_d is defined as the diagnostic agent (main agent). In this work, we define two prior knowledge agents, Q_p^s and Q_p^v, which represent semantic agent and visual agent, respectively. We will detail these agents in the next few subsections.
- Action: According to prior agents Q_p^v and Q_p^s , we define the a_p^v and a_p^s as their action respectively. For the main agent, we define a_d as the final diagnostic action selection. The visual agent Q_p^v is utilized to locate the area of disease in the medical image *I*. The action a_p^v is the coarse semantic segmentation image. The actions a_p^s is the preliminary classification representation, which can be obtained by pre-trained modal.

- Environment: The environment \mathcal{E} provides the influence based on the action of agents, the influence will also be used to update the visual agent and the semantic agent. The prior knowledge agents Q_p utilize the input query CXR images I and its associated label as its environment, the label embeddings inside Q_d are initialized by the prior knowledge of Q_p . Moreover, when we get the TD target, it also interacts with the prior agents and diagnostic agents, so the environment is dynamically interactive in our setting.
- **State**: The state is the feature vector *s* of input CXR image in each stage. Here, we apply the classic CNN model to extract *s*, which is trained by old data or other domain's similar data.
- **Reward**: We define the r^d as the reward of main agent Q_d and r_i^p as the reward of prior agent Q_p^i . Discount return is discount accumulative reward. The discount return of main agent is $U_t = \sum_{j=0}^{\infty} \gamma^j \cdot r_{t+j}^d$, where $\gamma \in [0, 1)$ is the discount factor, t is the index of training episode. The discount return of prior agent is $U_t^{p_i} = \sum_{j=0}^{\infty} \gamma^j \cdot r_{t+j}^{p_i}$, For Q_d and Q_p , we define that when a specific pathology is correctly classified, r^d and r_i^q equal to 1 else -1.

Based on these definitions, our goal is to seek a target $Q_d^*(s_t, a_t^v, a_t^s)$ function, which can provide the final accurate diagnosis, where a_t^v and a_t^s are the action of prior agents in time *t*. We also need to optimize the agent $Q_v^*(s_t)$ and $Q_s^*(s_t)$. Thus, the goal of RL is to maximize the expected cumulative reward:

$$max_{w_{s},w_{v},w_{d}}\mathbb{E}[U_{t}|s_{t},a_{t}^{v},a_{t}^{s}] + \sum_{i=0}^{n}\mathbb{E}[U_{t}^{p_{i}}|s_{t}],$$
(1)

where w_d denotes parameters of Q_d function, w_s and w_v denote the parameters of Q_p , t denotes the index of learning episode, U denotes the expectation of reward. n denotes the number of prior agents. In this paper, n = 2.

In our framework, we extract the feature as the state s_t when we input the medical image I. Then, the prior agents Q_p can select a set of action a^v and a^s . The main agent Q_d outputs the action a_t based on s_t , a^v and a^s . Here, we use vanilla temporal difference(TD) algorithmWatkins (1989) to train the DQN and define the loss \mathcal{L}_{TD} for the RL part based on the one-step TD error.

Based on feedback of the environment, we can expect achieving the related true label. These label can also be used to finetune the prior agents via the classic cross-entropy loss. The optimization function 1 can be rewritten as:

$$max_{w_{s},w_{v},w_{d}}\mathbb{E}[U_{t}|s_{t},a_{t}^{v},a_{t}^{s}] + \sum_{i=0}^{n}\mathbb{E}[U_{t}^{p_{i}}|s_{t}] + \mathbb{E}[y_{t}|s_{t},a_{t}^{v},a_{t}^{s}] + \sum_{i=0}^{n}\mathbb{E}[y_{t}|s_{t}],$$
(2)

where y_t is the true label of input image *I*. Here, $\mathbb{E}[U_t|s_t, a_t^v, a_t^s]$ and $\mathbb{E}[U_t^{p_i}|s_t]$ can apply the TD loss to optimization the model parameters. $\mathbb{E}[y_t|s_t, a_t^v, a_t^s]$ and $\mathbb{E}[y_t|s_t]$ can apply the classic



Fig. 2. The framework of our approach. The input is medical image *I*. This framework includes two prior knowledge agents Q_{prior}^v , Q_{prior}^s , and a diagnostic agent Q_d . The related actions are a_{prior}^v , a_{prior}^s , and a_d . The environment provides the ground truth y_t . Given a quadruple notation (s_t , a_t , r_t , s_{t+1}), the prediction value of DQN is \hat{q}_t , TD target is \hat{y}_t , then update the parameters in DQN via TD algorithm and loss function.

cross-entropy loss to solve the optimization problem. Thus, the final object can be written as:

$$\mathcal{L} = \sum_{i=0}^{n} \mathcal{L}_{p} + \sum_{i=0}^{n} \mathcal{L}_{TD}^{p} + \mathcal{L}_{TD} + \mathcal{L}_{d}, \qquad (3)$$

where \mathcal{L}_p and \mathcal{L}_d denotes the cross-entropy losses of prior agents and main agent, \mathcal{L}_{TD}^p denotes the TD loss of prior agent and \mathcal{L}_{TD} denotes the TD loss of main agents. *n* denotes the number of prior agents. We will details these losses in the later implementation stage.

Until now, we design the process of model learning based on the RL framework. However, in each diagnostic case, the model only selects the action with the highest value as the final diagnosis. This condition loses the human ability to explore in some tough cases. To handle this problem and make the diagnostic model smarter and more daring, we applied the ϵ -greedyWatkins (1989) algorithm to give our model the ability to explore. The key idea of this policy is with the probability ϵ to explore while with the probability (1- ϵ) to use the learned knowledge, which ensures that non-optimal cases can be chosen. Well in this setting, we want to explore the nonoptimal actions as our choice and make full use of our learned prior knowledge. In order to deal with the problem of the exploration-exploitation dilemma, we choose ϵ -greedy policy as our behavior policy to control the selection of actions:

- 1 ε: When the probability is 1 ε, we chose the action with the highest score as the diagnostic result. Note that our target problem is a multi-label classification problem, so on the premise that the classification threshold is satisfied, each agent may choose multiple actions.
- ϵ: When the probability is ε, we expect the agent can explore the different diagnosis like human. However, we do not allow this agent to randomly select the actions as the diagnosis results, which also should follow some prior knowledge. Thus, we delete the action with the highest score, and then select the remainder of the highest scores from the remaining actions as diagnostic results.

Apparently in the whole training process, ϵ is decreasing gradually, exploration at first and exploitation after learning

some useful information. In the training stage, we choose $\epsilon = \max\{\epsilon_{min}, 1 - \frac{1 - \epsilon_{min} * step}{total}\}$. ϵ equals to 0 when testing the model. We will detail some key elements in our proposed framework in the next few subsections.

3.2. Prior Agent

3.2.1. Semantic Agent

The semantic agent (Q_p^s) is used to extract and provide prior knowledge as the semantic format from the query CXR image *I*. We design the semantic information extraction network based on the ResNet structureHe et al. (2016). Here, we construct the 2D network with ResNet-50 and 3 additional convolutional layers with different kernel sizes, which can effectively extract the detail information from different scale of medical image. The action of this agent can be seen as the preliminary delineation of some pathologies. The initialized parameter of this agent can be learned by the training data, which can also be learned by other cross-domain data or other large-scale data as the prior knowledge (experience). We will discuss its different performance in experiment section.

By the semantic agent, we can obtain the output features $f_s^i \in \mathbb{R}^{d \times C}$, which represents the extracted feature and mainly contains information about i^{th} category from medical image. We employ the classifier to process each f_s^i into probability value p_s^i , which represents the classification result for f_s^i . We take p_s^i as the final confidence of i^{th} category. The output is composed of all categories confidences, which is denoted as $p_t \in \mathbb{R}^C$. The loss function \mathcal{L}_p can be written:

$$\mathcal{L}_p = -\sum_{i=1}^C t_j^i \log(p_i^i),\tag{4}$$

where $\mathbf{T}_j = \{t_j^1, t_j^2, ..., t_j^C\}$ represents the label of medical image I_j , t_j^i is a binary parameter that if the current CXR image I_j belongs to category *i*, $t_j^i = 1$, otherwise $t_j^i = 0$. In the test phase, we compare each value in p_t , and take the category with highest value as the classification result.

3.2.2. Visual Agent

This agent is used to extract visual attention area of the query medical image I as the visual prior knowledge. The action of



Fig. 3. Foreground Attention Block (FAB). We use the cascaded structure of channel- and position-wise attention to produce the foreground attention.



Fig. 4. The structure of the diagnostic agent.

 Q_p^{ν} is the foreground image feature a_p^{ν} . The structure of this agent is shown in Fig.3, visual agent (Q_p^{ν}) is used to provide prior knowledge as visual format. This structure is inspired by the previous work Ouyang et al. (2020). The foreground attention block(FAB) applies a self-attention mechanism to learn and emphasize the coarse foreground attention maps. The channelwise attention module is adapted from squeeze-and-excitation (SE) blockHu et al. (2018) which can focus on helpful channels, and then the position-wise attention module can use channelweighted vectors to give a preliminary and sketchy estimation of the areas. Note that the visual agent is also trained via the loss function (4) and the action of the visual agent is the coarse ROI of an input CXR image.

3.3. Diagnostic Agent

This agent can be seen as the final classification agent. The classic transformer structure is applied in this agent. The goal of transformer is used to fuse the output feature of Q_p^s and Q_p^v . Then, update the fuse feature based on the inner label embeddingLiu et al. (2021a) to achieve the final image feature representation. The detail of architecture is shown in Fig.4, which is similar to the standard Transformer decoder structure without self-attention module. Each Transformer decoder layer l updates the fused feature F_{l-1} from the output of its previous layers based on the label embedding L_l and spacial feature as follows:

$$q_{l} = L_{l} + p_{L_{l}}, k_{l} = f + p_{f}, v_{l} = f,$$

$$\dot{F}_{l} = MultiHead(q_{l}, k_{l}, v_{l}) + L_{l-1},$$

$$F_{l} = FFN(LN(\dot{F}_{l})) + \dot{F}_{l},$$
(5)

where *l* is the index of transformer layer, *f* is the image feature from prior knowledge agent, $p \in \mathbb{R}^{C \times d}$ denotes the learnable position encoding. In the first layer, the input label embedding is the initialized $Q_0 \in \mathbb{R}^{C \times d}$, where *C* is number of categories and *d* is the dimension of decoder. The output is the fused feature $F_l \in \mathbb{R}^{C \times d}$, which is actually a_d and represents the extracted feature for *C* categories at the last layer. Transformer decoder have a built-in cross-attention mechanism, we find common multi-label classification losses work well in dealing with CXR images classification tasks. But as we mentioned above about the data distribution imbalance problem, we want to deal with the label imbalance problem further via choosing a proper loss, then we choose an efficient version of focal lossLiu et al. (2021a), where we can choose different γ values for positive and negative values. Given an input CXR images *I*, we can predict its category probabilities $p = [p_1, \ldots, p_C]^T \in \mathbb{R}^C$ using our framework. We leverage the following asymmetric focal lossRidnik et al. (2021) to implement $\mathcal{L}_{diagnostic}(y_t | f_{cls}(a_d))$, the second part need to minimize in the objective function (3) to calculate the loss for each training CXR image *I*:

$$\mathcal{L}_{d} = \frac{1}{C} \sum_{c=1}^{C} \begin{cases} (1-p_{c})^{\gamma+} \log(p_{c}), & y_{c} = 1, \\ (p_{c})^{\gamma-} \log(1-p_{c}), & y_{c} = 0, \end{cases}$$
(6)

where y_c equals to 0 when the processed CXR image has label c and equals to 1 otherwise. γ + is set to 0 and γ - is set to 1 in our experiment. When completely training all the CXR images in our chosen training data set, we average (6) to calculate the loss to optimize the module. Furthermore, as we define the Q_d as the main agent and the DQN in our framework, we also optimize its parameters via the TD algorithm:

$$\mathcal{L}_{TD} = \frac{1}{2} [Q_d(s_t, a_s, a_v, w_t^d) - \hat{y}_t]^2,$$
(7)

$$\mathcal{L}_{TD}^{p} = \frac{1}{2} [Q_{p}(s_{t}, \boldsymbol{w}_{t}^{p}) - \hat{y}_{t}^{p}]^{2}, \qquad (8)$$

where, w_d and w_p are the network parameter of main agent and prior agent. $\hat{y}_t = r_t + \gamma \cdot max Q_d(s_{t+1}, a; w_{t-1})$ is the TD target. \hat{y}_t^p has the same equation. At this point, the details of each part of the total loss function (3) have been described. Similar to the work in Mnih et al. (2013), we store the experiences $e_t = (s_t, a_t, r_t, s_{t+1}) \in \mathcal{D}$ that agents learned at each time-step by experience reply, then sample the experiences from reply buffer by Q-learning updates during the inner loop of, after that agents can choose an action based on the ϵ -greedy policy. The full algorithm is presented in Algorithm 1.

4. Experiment

4.1. Data Sets

We evaluate our method on the publicly released NIH ChestX-ray14 and CheXpert datasets. NIH ChestX-ray14 consists of 112,120 frontal-view X-ray images of 30,805 unique patients, and we report the 14 thorax abnormal diseases classification performance on the testing set. CheXpert is another large-scale dataset for chest X-rays released by Stanford University. It contains 224,316 chest radio graphs of 65,240 unique patients. 14 observations are labeled in radiology reports, capturing uncertainties inherent in radiography interpretation. Part of the CXR images is a frontal view, and the rest of the data set is a lateral view.

Algorithm 1 MARL with DQN

Require: Image data *I* and corresponding ground truth y_t , initialize learnable parameters w_s of semantic agent, w_v of visual agent, w_d of diagnostic agent with random weights

Ensure: Initialize reply memory \mathcal{D} to capacity \mathcal{N}

Prior knowledge agents select actions and update parameters w_s and w_v by the image data *I* and its related ground truth y_t based on the loss function (4) and (8)

2:	tor	episode	=1, 1	M do
----	-----	---------	-------	------

- 3: Initialise s
- 4: **for** t = 1, T **do**

```
5: Diagnostic agent select an action a_d via \epsilon-greedy
```

- 6: Execute action a_d and observe reward r_d
- 7: Set $s_{t+1} = s_{t-d}$
- 8: Store $(s_{t,d}, a_{t,d}, r_{t,d}, s_{t+1,d})$ as a transition in \mathcal{D}

```
9: Sample random minibatch of transitions
```

- $(s_{j_d}, a_{j_d}, r_{j_d}, s_{j+1_d})$ from \mathcal{D}
- 10: Get \hat{y}_{j-d} base on TD target
- 11: Update w_d by (7) and (6)
- 12: Fine-tune w_s and w_v

end for

end for

Return: Updated DQN parameters $w_{\{s,v,q\}}$ via (3)

4.2. Evaluation Metrics

In our study, we chose the area under the receiver operating characteristic (ROC) curve, namely the AUC to evaluate the performance of the abnormal thorax disease classification. We conduct some experiments to test our proposed model's performance compared with the previous state-of-the-art (SOTA) framework. The AUC score of each pathology and the average AUC score over all pathologies are reported, respectively.

4.3. Implementation Details

We use PyTorch for implementation. We adopt vanilla Resnet-50He et al. (2016) as our backbone. Our experiment is operated by using NVIDIA GeForce RTX 3090 with 24GB memory. For training, we perform data augmentation by using RandAugmentCubuk et al. (2020) and normalizing CXR images with zero mean and one std rather than mostly used ImageNetDeng et al. (2009) pretrained images' mean and std according to the nature of CXR images. We use the AdamKingma and Ba (2014) optimizer with True-weight-decayLoshchilov and Hutter (2017)1*e*-2 and maximal learning rate is 1*e*-3, and we choose initial ϵ_{min} equals to 0.2 during the training step. At the same time, we resize the original images to and 512 × 512 and randomly horizontal flipping.

4.4. Comparison with State-of-the-art Methods

4.4.1. NIH ChestX-ray14

We first compared with some SOTA methods on the NIH ChestX-ray14 dataset. The AUC score of each pathology is summarized in Table.1. The average AUC score of our baseline arrives at 0.866 across the 14 thoracic diseases. It is competitive with or better than the previous works. In this experiment, we do not apply additional data to train the prior agents in order to make sure the comparison experiment is fair.

From these results, we have the following observations:

- The results obtained by our method are consistent with the optimal algorithms for "Cardiomegaly", "Infiltration", "Emphysema", "Pneumonia" and "Hernia". The last two pathologies are the least frequent of all pathologies, they account for about only 1.1% percent of the total data. All of the compared methods perform poorly in these two pathologies, which can be attributed to the long tail distribution problem. In order to solve above problem, our method guides the multi-information for the final pathology representation with the help of prior agents. The goal of our approach is the pursuit of optimal overall agents. It demonstrates the effectiveness of multi-agent framework indirectly.
- The four diseases, "Cardiomegaly", "Emphysema", "Edema", and "Fibrosis" account for only 6.6% of total pathologies, but we achieve SOTA AUC compared with other previous work. The reason is that we utilize the exploratory part of our MARL framework to give our model the opportunity to try some more reasonable choices, then solve the long tail distribution problem effectively. Meanwhile, the model of prior knowledge can provide effective auxiliary information and additional experience. Other comparable methods mostly rely on the quality of training data, but when they train with fewer samples, they are often powerless.

4.4.2. CheXpert

We also report the performance of our proposed MARL framework on the CheXpert data set in this section. We need to note that the data structure in CheXpert is different from NIH ChestX-ray14 data set, so we based on the baseline workIrvin et al. (2019), map the uncertain labels to 1 as it works well and make the proposed model outputs the maximum probability of the observations across the views. Meanwhile, based on the baseline, we also choose to evaluate 5 classes which called competition tasks, which selected according to clinical importance and prevalence. Our results are shown in Table.2 and we adopt different backbones. Same as the experiment on NIH ChestX-ray14, when Resnet is used as the backbone we got the better results. We present the performance obtained by a single model rather than the performance of Irvin et al. Irvin et al. (2019), which is from the ensemble of 30 models. In our proposed framework, the average AUC scores over five pathologies achieve 0.925.

In this data set, when using our framework and the uncertain labels are set to "ones", we get the mean AUC of 0.925, while when using our framework and the uncertain labels are set to "zeros", the AUC of each pathology is 0.831, 0.877, 0.933, 0.902, 0.938, respectively, the average AUC score is 0.8962, which surpasses the corresponding baseline "U-Zeros". Here, we notice that our model shows its superiority for some pathologies except "Edema", the performance of "Cardiomegaly" and "Atelectasis" are significantly improved (0.832 vs 0.890, 0.858 vs 0.922), which improved 6.97% and 7.46%, respectively.

Abnormality	Wang <i>et al</i> . Wang et al. (2017)	Xi <i>et al.</i> Ouyang et al. (2020)	ImageGCN Mao et al. (2022)	DGFN Gong et al. (2021)	Ours
Atelectasis	0.70	0.77	0.80	0.82	0.83 (0.82, 0.84)
Cardiomegaly	0.81	0.87	0.89	0.93	0.93 (0.91, 0.94)
Effusion	0.76	0.83	0.87	0.88	0.90 (0.89, 0.92)
Infiltration	0.66	0.71	0.70	0.75	0.75 (0.74, 0.76)
Mass	0.69	0.83	0.84	0.88	0.89 (0.88, 0.91)
Nodule	0.67	0.79	0.77	0.79	0.81 (0.80, 0.82)
Pneumonia	0.66	0.82	0.72	0.78	0.82 (0.81, 0.83)
Pneumothorax	0.80	0.88	0.90	0.89	0.91 (0.90, 0.93)
Consolidation	0.70	0.74	0.80	0.81	0.82 (0.81, 0.85)
Edema	0.81	0.84	0.88	0.89	0.91 (0.90, 0.92)
Emphysema	0.83	0.94	0.92	0.94	0.94 (0.94, 0.95)
Fibrosis	0.79	0.83	0.83	0.82	0.85 (0.84, 0.86)
Pleural_Thickening	0.68	0.79	0.79	0.81	0.83 (0.81, 0.84)
Hernia	0.87	0.91	0.94	0.92	0.94 (0.93, 0.95)
Mean AUC	0.745	0.819	0.832	0.850	0.866 (0.859, 0.873)

Table 1. Comparation Of AUC Scores with previous SOTA Works on the official split of NIH ChestX-ray14. We report the AUC with 95 % confidence interval (CI) of our method. The best performance of each pathology is shown in bold.

Table 2. Comparison of AUC scores with different models on CheXpert validation set. Our approach shows the results under different backbone, DenseNet and ResNet.

Abnormality	U-Zeros Irvin et al. (2	019) Ir	U-Ones tvin et al. (2019)	Pham et	+LSR al. (2021)	Xi <i>et al.</i> Ouyang et al.	$(2020) \begin{vmatrix} Ou \\ (Dens) \end{vmatrix}$	urs eNet) Ours (ResNet)
Atelectasis	0.811		0.858	0.	.825	0.920	0.9	016 0.922
Cardiomegaly	0.840		0.832	0.	.855	0.886	0.8	0.890
Consolidation	0.932		0.899	0.	.937	0.907	0.9	0.942
Edema	0.929		0.941	0.	.930	0.937	0.9	034 0.932
Pleural Effusion	0.931		0.934	0.	.923	0.933	0.9	035 0.938
Mean AUC	0.889		0.893	0.	.894	0.917	0.9	020 0.925

4.5. Few-shot and Data-efficiency Problem

We conduct an experiment on the NIH ChestX-ray14 data set to demonstrate that our proposed MARL framework can handle the data-efficient and few-shot learning problem.

In terms of the few-shot learning problem, we split the NIH dataset, in which 9 types of pathologies with about 1,100 samples are chosen as auxiliary set and the remaining 5 types of samples are taken as the test set. We perform 5-way 1-shot, 5-way 5-shot, and 5-way 10-shot classification tasks on the NIH data set. Take 5-way 5-shot as an example: there are 5 different pathology categories with 5 support images and 10 query images per category, *i.e.* $5 \times 5 + 10 \times 5 = 75$ images in each episode. We pre-trained the prior agent using the CheXpert dataset. The two datasets do not have the same category. As a result, we can only guarantee dimensionality consistency in the initial guided learning. Naturally, we then apply fine-tuning at the next time. To divide the query images into the closest category, we use a Euclidean and Cosine distance metric, as well as a softmax

function.

Compared with other methods using few-shot learning Jin et al. (2021); Jiang et al. (2022); Finn et al. (2017); Gordon et al. (2018) as shown in Table.3, we achieve the acc of 44.26, 49.63, and 52.64 in the 1, 5, and 10-shot scenarios. Experimental results prove that the prior knowledge learning module and the reinforcement learning setting in this framework can help to solve the few-shot learning problem, and the exploration ability of our method can fit the nature of humans, which helps to improve model performance. At the same time, when we train the model with the same scale of data as other few-shot learning methods, we can also achieve better classification performance.

We also modify the number of training data to further observe the impact of the number of training data on model performance. The related experimental results are shown in Fig.5, at the same time, we reproduce two other works Xiao et al. (2023); Seyyed-Kalantari et al. (2020) to compare with our pro-



Fig. 5. The process of prior knowledge accumulation and illustration of data-efficient. The ordinate is AUC scores while the abscissa means the percentage of total data volume of CXR images a doctor has learned. We denote CheXNet and M-MAE as the method in Seyyed-Kalantari et al. (2020); Xiao et al. (2023) respectively.

 Table 3. The 3-way, 1, 5, 10-shot classification accuracy(%). The best results are highlighted.

FSL Methods	1-shot	5-shot	10-shot
MAMLFinn et al. (2017)	39.96	46.83	48.69
VersaGordon et al. (2018)	37.94	45.81	47.13
M-LearnerJiang et al. (2022)	43.37	47.95	50.58
Ours	44.26	49.63	52.64

posed method on the matter of data efficient. From these results, we can find that we train the diagnostic agent with only about 80% of CXR images, the mean AUC of the framework is comparative with previous SOTA work. Besides, we can find that when selecting the same amount of data for training, our results always outperform than the other two methods. We think the reason is that the RL framework allows each training to have more exploration capabilities, allowing the model to perform reasoning like a person, thereby amplifying the role of training data.

4.6. Ablation Study

In this section, we conduct the ablation study to prove the superiority of our proposed MARL thorax diseases classification framework. We run the entire ablation experiment on the NIH ChestX-ray14 data set to demonstrate the significance of each agent and the MARL setting. The related experiments are shown in Table.4. Note that only model 8 remove the RL-related settings. "+" represents the framework using that element, while "-" means the framework removing the element. Model 9 which contains three "+" is our proposed integrated MARL framework, while Model 8 is a model based on Model 9 but removes all the MARL-related settings. We can conclude that the prior knowledge extraction module of our proposed framework is important to help the diagnostic agent make a decision. We can also confirm the diagnostic agent is good at decoupling and fusing prior knowledge from the visual and semantic agents. Furthermore, the RL-related setting is key to making a more accurate diagnosis.

 Table 4. Ablation study. Note that model 8 removes the reinforcement learning settings, compared with model 9.

Model		Visual Agent	Semantic Agent	Diagnostic Agent	AUC
	1	-	+	+	0.862
	2	-	-	+	0.856
	3	-	-	-	0.839
	4	-	+	-	0.850
	5	+	+	-	0.838
	6	+	-	+	0.861
	7	+	-	-	0.841
	8	+	+	+	0.837
	9	+	+	+	0.866

Based on this setting, we can achieve the following observations:

- The visual agent is the only used part in model 7, and we can find the result is just a little higher than the worst in model 3, which is easy to understand. If none of the components are used, the whole structure is totally not integrated into our proposed model 9. In the case of model 7, we only focus on the foreground information of the input CXR images and ignore the detailed information under the CXR images, thus the results are not good. But we can find that the result is still better than part of the previous work, so the result still demonstrates the performance of the visual agent.
- The semantic agent brings some performance improvements according to model 5, 7 and 1, 2. The results show that a strong spatial feature extractor is important and can get detailed information about the features of CXR images. According to the results of models 4, 5, and 7, we can find that when we use both the visual agent and the semantic agent, the model's performance is worse than the model using only the semantic agent. It means the foreground attention block introduces some redundant information and then impairs the final classification results. Furthermore, we can see that the model with only a semantic agent outperforms some previous SOTA works because it validated the semantic agent's validity.
- According to the results of models 6 and 7, the diagnostic (main) agent takes advantage of the transformer decoder structure and achieves a significant improvement over the visual and semantic modules. Compared with the other models, model 9 achieves a new state-of-the-art AUC of 0.866, which represents a 2% improvement and is better than almost all of the previous works. In this case, the prior agent, trained with all training data, is crucial to the final diagnostic result. That seems a little unfair. In order to reduce this interference, we randomly initialized the prior

agent and updated it with the new data. It also achieves the mean AUC of 0.856 and also better than all of the other baseline models.

• The proposed MARL framework is superior to the traditional deep learning framework in Model 8. In order to demonstrate the performance of the MARL framework. we introduce Model 8, which is the model in which we remove the multi-agent reinforcement learning related settings such as action, the temporal difference algorithm, and so on. It means that we only applied the L_p and L_d losses to optimize the model's parameters. It achieves an AUC of 0.837, which is about 3% lower than the integrated proposed model 9. This result demonstrates the effectiveness of the RL mechanism in our approach; when we remove it, the model 8 cannot benefit from the interaction and exploration mechanisms in reinforcement learning. Besides, the results in model 8 is still competitively compared to previous methods, it also proves the rationality of our proposed model structure.

In general, our proposed framework includes all modules, which gives us the best mean AUC of 0.866. The corresponding experiments and analyses demonstrate the performance of these modules.

4.7. Prior Knowledge from Other Domain

In this paper, we introduce the prior knowledge information to guide the final diagnosis. The prior agent can be seen as the Large-scale pre-trained models (PTMs). If this prior knowledge comes from data in other domains, will it have new effects on the main agent? This question is somewhat similar. If we find an orthopedic doctor to see a chest disease, will the orthopedic experience play a positive role? In order to explore this problem, we applied the other data set to initialize the parameters of the visual agent and the semantic agent. Then, we trained the main agent and fine-tuned the prior agents based on the NIH ChestX-ray14 data set.

The related experimental results are shown in Fig.6. "NIH" means that we only use the NIH training data to initialize the parameters of agents; "NIH-" means that we remove the prior agents and only apply the main agent to make the final diagnosis. "PASCAL+NIH" means that we applied PASCAL data to train prior agents and NIH data to fine-tune all agents. "CheXpert+NIH" means that we applied CheXpert data to train prior agents. We can achieve the following observations:

• The prior agents can speed up the training of the main agent compared with that of "NIH-". "NIH-" achieved the worst performance. The other three methods introduced the prior agent. We also find that "CheXpert+NIH" has the fastest training speed, followed by "PASCAL+NIH", and "NIH" is the worst. "CheXpert+NIH" and "PAS-CAL+NIH" applied the CHeXpert and Pascal data to intilize the parameters of prior agents. This means that the main agent uses more data, regardless of whether this data is consistent with the target domain data. It still allows the diagnostic agent to see more data. More data can provide



Fig. 6. Different experiment setting's mean AUC with experiment epoch.

better training results, and many deep learning algorithms have proved this conclusion. In these results, More data leads to faster convergence and better results. This also proves the necessity and effectiveness of the prior agents.

Similar data with target domain can improve the final performance. We can find that "CheXpert+NIH" has the best results, followed by "NIH", and "PASCAL+NIH" is the worst. The reason rely on the quality of prior agent's parameter. CheXert data has a similar distribution to target domain data in NIH. PASCAL data has obviously different with target data. Thus, "CheXpert+NIH" has the best result. "NIH" has a better result than that of "PASCAL+NIH". The reason is that the large distribution difference with the target domain data has had a certain negative impact. This is also in line with people's common sense. After all, an orthopedic doctor needs a certain amount of study before he can become a chest doctor.

4.8. The Performance of ϵ -greedy

In the training step, we utilize the ϵ -greedy algorithm to add some exploratory and speed up the convergence of the model. This is like a human being making some tentative decisions on certain issues without experience. In this section, we adjust the value of the parameter ϵ to control the scale of exploration, and then observe its impact on the final result. The related experiment is shown in Fig.7.

The diagnostic result increases with the increase of ϵ value, and when ϵ =0.2, the best result is achieved. But when ϵ continued to increase, the diagnostic results showed a rapid decline. When ϵ is equal to 0.2, it means that in one-fifth of the cases, the highest-scoring result was not chosen as the recognition result. If we continue to increase the value of ϵ , it means that we discard the experience and increase the uncertainty of the diagnosis. This condition will lead to that the model introduces more wrong results in the training, which also leads to a rapid decrease in the results. From these results, we can find that



Fig. 7. The correlation of mean AUC and exploration.

appropriate exploration is effective and can speed up the convergence and training speed of the model. However, too much exploration will lose its original meaning, abandon the guiding significance of the real results, and introduce more uncertainties. In general, this experiment demonstrated the effectiveness of training mechanism. This idea can be applied to many similar problems.

4.9. Parameter Selection

The parameters of transformer in diagnostic agents are worth further discussion. Here, we try to figure out how many decoder layers and multi-heads are in the inner structure of the Transformer in order to test its most important information-splitting and -combining functions in this vision task. At first, we set up 2 layers of transformer decoders, each with 4 heads, and got an average AUC of 0.866. Under the condition of 4 heads, we achieve a mean AUC of 0.843, 0.866, 0.853, 0.836, 0.829, and 0.825 when we set 1-6 layers, respectively. Under the condition of 2 layers, we achieve a mean AUC of 0.837, 0.849, 0.866, and 0.826 when we set 1, 2, 4, and 8 heads, respectively.

To further verify the impact of layers and multi-headed transformers, we chose to visualize some attention maps of transformer structure. We are also interested in finding out the role of multi-head attention in this task. We plot the mean of each head's cross-attention maps, which represent the similarities of a given query and extracted spatial features. From Fig.8, we can see that when the number of heads equals to 4, the attention maps better identify the lesion of thorax diseases than others, it testify in this specific task, it is good to use 4 heads. Besides, we can see that when the multi-heads are not enough, the attention weights are scattered and not very accurate, and when there are 8 multi-heads, the attention maps are chaotic, which means the multi-head attention mechanism introduces some redundant information and adds some interference to the classification task, which indicates that the redundant heads are not utilized as the other four heads already collect sufficient information for classification. Similar to the layer attention map visualization experiment shown in Fig.9, we can find that when the number of



Fig. 8. Visualize results of average the maps of all heads. The leftmost column is the input CXR image, and disease areas are marked according to official documents by us, diseases from the top to bottom are "Mass", "Pneumonia", "Mass", the right 4 columns average the maps of all heads, 1, 2, 4, and 8 heads, respectively. Best view in colors.



Fig. 9. Visualize results of layer attention maps, all pathologies are "Cardiomegaly". Best view in colors.

layers of the transformer decoder equals 2, the visualization result shows better learned information than others, proving that the layer setting is important in this task.

4.10. Algorithm Complexity

The computational consumption is also a factor that should be considered for clinical usage. In general, our computational resource consumption is moderate compared to other methods.

We mostly focus on the model parameters and the size of the input image that affect how much GPU power our proposed MARL framework needs. For a certain input CXR image resolution, e.g. 256×256 , we use Resnet-18 and Resnet-50 in our framework, and the scale of model parameters in our proposed framework increases from 39.87M to 175.89M.

Taking a 256×256 input image as an example, as shown in Table.5, the FLOPs (floating point operations) for our proposed framework are 63.57G MACs (Multiplication and Accumulation) to 42.78G MACs from Resnet-50 to VGGNet, which the range of change is not small. In practice, training our proposed framework with 256×256 images and Resnet-18 or Resnet-50 on an NVIDIA GeForce RTX 3090 GPU with 24 GB memory costs nearly 0.033 or 0.145 seconds per image. Large-resolution images would cost much more time, e.g., when we

Methods	Backbone	Resolution	FLOPs(GMACs)
Yao <i>et al</i> . Yao et al. (2018)	DenseNet	224	13.57
Wang et al. Wang et al. (2017)	AlexNet	224	14.54
Wang et al. Wang et al. (2017)	VGGNet	224	30.95
LLAGNet Chen et al. (2019a)	DenseNet	256	34.96
Ours	ResNet	256	63.57
Ours	VGGNet	256	42.78
Ours	DenseNet	256	52.31
Ours	ResNet	512	159.96

Table 5. Floating point operations with different backbone and resolution

use 512×512 as the input CXR image resolution with Resnet-18 or Resnet-50 training on the same GPU, it costs nearly 0.069 or 0.251 seconds per CXR image. By the way, loading data into memory would take a lot of time compared to a smaller input resolution. The FLOPs increase about 96.39G (63.57G vs. 159.96G) MACs for 512×512 with Resnet-18 as our backbone compared with input resolution 256×256 , respectively.

Besides, the GPU consumption is extremely different when loading the training data with different input image sizes or batch sizes. In our experiment, when we adopt Resnet-18 as our backbone, training our proposed framework with 16 images (256×256) in a mini-batch costs about 2.4 GB of GPU memory. When the input is set to 32 images (256×256) in a mini-batch, 3.1 GB of GPU memory is required. While the input is set to 32 images (512×512) in a mini-batch, nearly 5.0 GB of GPU memory is required. In the above experiments, training with images of a higher resolution leads to better performance, but the amount of GPU memory used also goes up.

5. Conclusion

In this paper, we propose a new multi-agent reinforcement learning framework to solve the multi-label CXR image classification problem. This framework uses diagnostic agents' previous knowledge to guide their learning, just like how a person learns. Prior knowledge is learned from the pre-trained model based on old data or similar data from other domains, which can effectively reduce the dependence on target domain data and speed up convergence. 2) The framework of reinforcement learning can make the diagnostic agent as exploratory as a human and improve the accuracy of diagnosis through continuous exploration. This design makes the whole model more intelligent and more in line with human learning rules. Meanwhile, the method can effectively solve the few-shot model learning problem and improve the model's generalization ability. We evaluated our proposed method on two public data sets, NIH ChestX-ray14 and CheXpert. The experimental results demonstrate the performance of our approach.

We next briefly discuss the limitations of our proposed method and future work possibilities. Firstly, our proposed MARL framework adopts the simplest form of multi-agent reinforcement learning, and more complicated multi-agent relationships are worth further study. Secondly, as we introduce reinforcement learning into the traditional classification framework, we will consume more computing resources and train much more slowly than in the traditional classification format, we will try to make our models lightweight in the future. We believe that further research will help us solve these issues.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (62272337) and the Natural Science Foundation of Tianjin (16JCZDJC31100, 16JCZDJC31100).

References

- Abbas, A., Abdelsamea, M.M., Gaber, M.M., 2021. Classification of covid-19 in chest x-ray images using detrac deep convolutional neural network. Applied Intelligence 51, 854–864.
- Boutell, M.R., Luo, J., Shen, X., Brown, C.M., 2004. Learning multi-label scene classification. Pattern recognition 37, 1757–1771.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S., 2020. End-to-end object detection with transformers, in: European Conference on Computer Vision, Springer. pp. 213–229.
- Chen, B., Li, J., Lu, G., Zhang, D., 2019a. Lesion location attention guided network for multi-label thoracic disease classification in chest x-rays. IEEE journal of biomedical and health informatics 24, 2016–2027.
- Chen, Y., Zhao, X., Tang, B., 2022. Boosting lesion annotation via aggregating explicit relations in external medical knowledge graph. Artificial Intelligence in Medicine 132, 102376.
- Chen, Z.M., Wei, X.S., Wang, P., Guo, Y., 2019b. Multi-label image recognition with graph convolutional networks, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 5177–5186.
- Cherti, M., Jitsev, J., 2021. Effect of large-scale pre-training on full and few-shot transfer learning for natural and medical images. arXiv preprint arXiv:2106.00116.
- Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V., 2020. Randaugment: Practical automated data augmentation with a reduced search space, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 702–703.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database, in: 2009 IEEE conference on computer vision and pattern recognition, Ieee. pp. 248–255.
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K., 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al., 2020. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.
- Dou, H., Yang, X., Qian, J., Xue, W., Qin, H., Wang, X., Yu, L., Wang, S., Xiong, Y., Heng, P.A., et al., 2019. Agent with warm start and active

termination for plane localization in 3d ultrasound, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 290–298.

- Finn, C., Abbeel, P., Levine, S., 2017. Model-agnostic meta-learning for fast adaptation of deep networks, in: International conference on machine learning, PMLR, pp. 1126–1135.
- Fu, J., Thomas, H.R., Li, C., 2021. Tortuosity of porous media: Image analysis and physical simulation. Earth-Science Reviews 212, 103439.
- Ghesu, F.C., Georgescu, B., Mansi, T., Neumann, D., Hornegger, J., Comaniciu, D., 2016. An artificial agent for anatomical landmark detection in medical images, in: International conference on medical image computing and computer-assisted intervention, Springer. pp. 229–237.
- Ghesu, F.C., Georgescu, B., Zheng, Y., Grbic, S., Maier, A., Hornegger, J., Comaniciu, D., 2017. Multi-scale deep reinforcement learning for real-time 3d-landmark detection in ct scans. IEEE transactions on pattern analysis and machine intelligence 41, 176–189.
- Gong, X., Xia, X., Zhu, W., Zhang, B., Doermann, D., Zhuo, L., 2021. Deformable gabor feature networks for biomedical image classification, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 4004–4012.
- Gordon, J., Bronskill, J., Bauer, M., Nowozin, S., Turner, R.E., 2018. Versa: Versatile and efficient few-shot learning, in: Third workshop on Bayesian Deep Learning.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778.
- Hou, D., Zhao, Z., Liu, Y., Chang, F., Hu, S., 2021. Automatic report generation for chest x-ray images via adversarial reinforcement learning. IEEE Access 9, 21236–21250.
- Hu, J., Shen, L., Sun, G., 2018. Squeeze-and-excitation networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 7132–7141.
- Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghgoo, B., Ball, R., Shpanskaya, K., et al., 2019. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison, in: Proceedings of the AAAI conference on artificial intelligence, pp. 590–597.
- Jain, A., Ratnoo, S., Kumar, D., 2017. Addressing class imbalance problem in medical diagnosis: A genetic algorithm approach, in: 2017 international conference on information, communication, instrumentation and control (ICICIC), IEEE. pp. 1–8.
- Jiang, H., Gao, M., Li, H., Jin, R., Miao, H., Liu, J., 2022. Multi-learner based deep meta-learning for few-shot medical image classification. IEEE Journal of Biomedical and Health Informatics.
- Jin, Y., Lu, H., Zhu, W., Yan, K., Gao, Z., Li, Z., 2021. Ctfc: A convolution and visual transformer based classifier for few-shot chest x-ray images, in: 2021 2nd International Conference on Artificial Intelligence and Computer Engineering (ICAICE), IEEE. pp. 616–622.
- Kant, N., Puri, R., Yakovenko, N., Catanzaro, B., 2018. Practical text classification with large pre-trained language models. arXiv preprint arXiv:1812.01207.
- Ke, A., Ellsworth, W., Banerjee, O., Ng, A.Y., Rajpurkar, P., 2021. Chextransfer: performance and parameter efficiency of imagenet models for chest xray interpretation, in: Proceedings of the Conference on Health, Inference, and Learning, pp. 116–124.
- Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- Kipf, T.N., Welling, M., 2016. Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907.
- Liao, X., Li, W., Xu, Q., Wang, X., Jin, B., Zhang, X., Wang, Y., Zhang, Y., 2020. Iteratively-refined interactive 3d medical image segmentation with multi-agent reinforcement learning, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9394–9402.
- Liu, H., Wang, L., Nan, Y., Jin, F., Wang, Q., Pu, J., 2019. Sdfn: Segmentationbased deep fusion network for thoracic disease classification in chest x-ray images. Computerized Medical Imaging and Graphics 75, 66–73.
- Liu, S., Zhang, L., Yang, X., Su, H., Zhu, J., 2021a. Query2label: A simple transformer way to multi-label classification. arXiv preprint arXiv:2107.10834.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B., 2021b. Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030.
- Loshchilov, I., Hutter, F., 2017. Decoupled weight decay regularization. arXiv

preprint arXiv:1711.05101.

- Mao, C., Yao, L., Luo, Y., 2022. Imagegcn: Multi-relational image graph convolutional networks for disease identification with chest x-rays. IEEE Transactions on Medical Imaging .
- Minaee, S., Kafieh, R., Sonka, M., Yazdani, S., Soufi, G.J., 2020. Deep-covid: Predicting covid-19 from chest x-ray images using deep transfer learning. Medical image analysis 65, 101794.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M., 2013. Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602.
- Ouyang, X., Karanam, S., Wu, Z., Chen, T., Huo, J., Zhou, X.S., Wang, Q., Cheng, J.Z., 2020. Learning hierarchical attention for weakly-supervised chest x-ray abnormality localization and diagnosis. IEEE Transactions on Medical Imaging.
- Park, S., Kim, G., Oh, Y., Seo, J.B., Lee, S.M., Kim, J.H., Moon, S., Lim, J.K., Ye, J.C., 2022. Multi-task vision transformer using low-level chest x-ray feature corpus for covid-19 diagnosis and severity quantification. Medical Image Analysis 75, 102299.
- Paul, A., Tang, Y.X., Shen, T.C., Summers, R.M., 2021. Discriminative ensemble learning for few-shot chest x-ray diagnosis. Medical image analysis 68, 101911.
- Pham, H.H., Le, T.T., Tran, D.Q., Ngo, D.T., Nguyen, H.Q., 2021. Interpreting chest x-rays via cnns that exploit hierarchical disease dependencies and uncertainty labels. Neurocomputing 437, 186–194.
- Ridnik, T., Ben-Baruch, E., Zamir, N., Noy, A., Friedman, I., Protter, M., Zelnik-Manor, L., 2021. Asymmetric loss for multi-label classification, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 82–91.
- Rocha, J., Pereira, S.C., Pedrosa, J., Campilho, A., Mendonça, A.M., 2022. Attention-driven spatial transformer network for abnormality detection in chest x-ray images, in: 2022 IEEE 35th International Symposium on Computer-Based Medical Systems (CBMS), IEEE. pp. 252–257.
- Saleem, H.N., Sheikh, U.U., Khalid, S.A., 2021. Classification of chest diseases from x-ray images on the chexpert dataset, in: Innovations in Electrical and Electronic Engineering. Springer, pp. 837–850.
- Sarvamangala, D., Kulkarni, R.V., 2021. Convolutional neural networks in medical image understanding: a survey. Evolutionary intelligence, 1–22.
- Seyyed-Kalantari, L., Liu, G., McDermott, M., Chen, I.Y., Ghassemi, M., 2020. Chexclusion: Fairness gaps in deep chest x-ray classifiers, in: BIOCOM-PUTING 2021: proceedings of the Pacific symposium, World Scientific. pp. 232–243.
- Singh, R., Bharti, V., Purohit, V., Kumar, A., Singh, A.K., Singh, S.K., 2021. Metamed: Few-shot medical image classification using gradientbased meta-learning. Pattern Recognition, 108111.
- Stember, J., Shalu, H., 2020. Deep reinforcement learning to detect brain lesions on mri: a proof-of-concept application of reinforcement learning to medical images. arXiv preprint arXiv:2008.02708.
- Sun, K.W., Lee, C.H., 2017. Addressing class-imbalance in multi-label learning via two-stage multi-label hypernetwork. Neurocomputing 266, 375–389.
- Tarekegn, A., Giacobini, M., Michalak, K., 2021. A review of methods for imbalanced multi-label classification. Pattern Recognition, 107965.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need, in: Advances in neural information processing systems, pp. 5998–6008.
- Vila, L.C., Escolano, C., Fonollosa, J.A., Costa-Jussa, M.R., 2018. End-to-end speech translation with the transformer., in: IberSPEECH, pp. 60–63.
- Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M., 2017. Chestxray8: Hospital-scale chest x-ray database and benchmarks on weaklysupervised classification and localization of common thorax diseases, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2097–2106.
- Watkins, C.J.C.H., 1989. Learning from delayed rewards .
- Wu, G., Tian, Y., Liu, D., 2018. Cost-sensitive multi-label learning with positive and negative label pairwise correlations. Neural Networks 108, 411–423.
- Xiao, J., Bai, Y., Yuille, A., Zhou, Z., 2023. Delving into masked autoencoders for multi-label thorax disease classification, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 3588–3600.
- Yao, L., Prosky, J., Poblenz, E., Covington, B., Lyman, K., 2018. Weakly supervised medical diagnosis and localization from multiple resolutions. arXiv preprint arXiv:1803.07703.
- Zhang, Y., Wang, Y., Liu, X.Y., Mi, S., Zhang, M.L., 2020. Large-scale multi-

label classification using unknown streaming images. Pattern Recognition 99, 107100.

Zhu, X., Pang, S., Zhang, X., Huang, J., Zhao, L., Tang, K., Feng, Q., 2022. Pcan: Pixel-wise classification and attention network for thoracic disease classification and weakly supervised localization. Computerized Medical Imaging and Graphics 102, 102137.