

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Deep Learning for Automated Medical Image Analysis

Permalink

<https://escholarship.org/uc/item/78w60726>

Author

Zhu, Wentao

Publication Date

2019

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Deep Learning for Automated Medical Image Analysis

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Computer Science

by

Wentao Zhu

Dissertation Committee:
Professor Xiaohui Xie, Chair
Professor Charless C. Fowlkes
Professor Sameer Singh

2019

DEDICATION

To my wife, little son and my parents

TABLE OF CONTENTS

	Page
LIST OF FIGURES	vi
LIST OF TABLES	x
LIST OF ALGORITHMS	xi
ACKNOWLEDGMENTS	xii
CURRICULUM VITAE	xiii
ABSTRACT OF THE DISSERTATION	xvi
1 Introduction	1
1.1 Dissertation Outline and Contributions	3
2 Adversarial Deep Structured Nets for Mass Segmentation from Mammograms	7
2.1 Introduction	7
2.2 FCN-CRF Network	9
2.3 Adversarial FCN-CRF Nets	10
2.4 Experiments	12
2.5 Conclusion	16
3 Deep Multi-instance Networks with Sparse Label Assignment for Whole Mammogram Classification	17
3.1 Introduction	17
3.2 Deep MIL for Whole Mammogram Mass Classification	19
3.2.1 Max Pooling-Based Multi-Instance Learning	20
3.2.2 Label Assignment-Based Multi-Instance Learning	22
3.2.3 Sparse Multi-Instance Learning	23
3.3 Experiments	24
3.4 Conclusion	27

4	DeepLung: Deep 3D Dual Path Nets for Automated Pulmonary Nodule Detection and Classification	29
4.1	Introduction	29
4.2	Related Work	32
4.3	DeepLung Framework	33
4.3.1	3D Faster R-CNN with Deep 3D Dual Path Net for Nodule Detection	33
4.3.2	Gradient Boosting Machine with 3D Dual Path Net Feature for Nodule Classification	37
4.3.3	DeepLung System: Fully Automated Lung CT Cancer Diagnosis	38
4.4	Experiments	39
4.4.1	Datasets	39
4.4.2	Preprocessing	40
4.4.3	DeepLung for Nodule Detection	40
4.4.4	DeepLung for Nodule Classification	43
4.4.5	DeepLung for Fully Automated Lung CT Cancer Diagnosis	45
4.5	Discussion	47
4.5.1	Nodule Detection	47
4.5.2	Nodule Classification	48
4.6	Conclusion	50
5	DeepEM: Deep 3D ConvNets With EM For Weakly Supervised Pulmonary Nodule Detection	51
5.1	Introduction	51
5.2	DeepEM for Weakly Supervised Detection	54
5.3	Experiments	57
5.4	Conclusion	60
6	AnatomyNet: Deep Learning for Fast and Fully Automated Whole-volume Segmentation of Head and Neck Anatomy	62
6.1	Introduction	62
6.2	Materials and Methods	66
6.2.1	Data	66
6.2.2	Network architecture	67
6.2.3	Loss function	69
6.2.4	Handling missing annotations	71
6.2.5	Implementation details and performance evaluation	73
6.3	Results	73
6.3.1	Determining down-sampling scheme	74
6.3.2	Choosing network structures	75
6.3.3	Choosing loss functions	77
6.3.4	Comparing to state-of-the-art methods	79
6.3.5	Visualizations on MICCAI 2015 test	80
6.3.6	Visualizations on independent samples	81

6.4	Discussion	81
6.4.1	Impacts of training datasets	81
6.4.2	Limitations	83
6.5	Conclusion	84
7	Conclusion	90
	Bibliography	92

LIST OF FIGURES

	Page
2.1 The proposed adversarial deep FCN-CRF network with four convolutional layers followed by CRF for structured learning.	8
2.2 The empirical estimation of a priori on INbreast (left) and DDSM-BCRP (right) training datasets (a). Trimap visualizations on the DDSM-BCRP dataset, segmentation groundtruth (first column), trimap of width 2 (second column), trimaps of width 3 (third column) (b).	11
2.3 Visualization of segmentation results using the FCN (first row), Adversarial FCN (second row), Joint FCN-CRF (third row), Adversarial FCN-CRF (fourth row) on the test sets of INbreast dataset. Each column denotes a test sample. Red lines denote the ground truth. Green lines or points denote the segmentation results. Adversarial training provides sharper and more accurate segmentation boundaries than methods without adversarial training. . .	15
2.4 Accuracy comparisons among FCN, Adversarial FCN, Joint FCN-CRF and Adversarial FCN-CRF in trimaps with pixel width 1, 2, 3, 4, 5 on the INbreast dataset (a) and the DDSM-BCRP dataset (b). The adversarial training improves segmentation accuracy around boundaries.	16
3.1 The framework of whole mammogram classification. First, we use Otsu’s segmentation to remove the background and resize the mammogram to 227×227 . Second, the deep MIL accepts the resized mammogram as input to the convolutional layers. Here we use the convolutional layers in AlexNet [59]. Third, logistic regression with weight sharing over different patches is employed for the malignant probability of each position from the convolutional neural network (CNN) feature maps of high channel dimensions. Then the responses of the instances/patches are ranked. Lastly, the learning loss is calculated using max pooling loss, label assignment, or sparsity loss for the three different schemes.	19
3.2 Histograms of mass width (a) and height (b), mammogram width (c) and height (d). Compared to the size of whole mammogram ($1,474 \times 3,086$ on average after cropping), the mass of average size (329×325) is tiny, and takes about 2% of a whole mammogram.	20

3.3	The visualization of predicted malignant probabilities for instances/patches in four resized mammograms. The first row is the resized mammogram. The red rectangle boxes are mass regions from the annotations on the dataset. The color images from the second row to the last row are the predicted malignant probability from logistic regression layer for (a) to (d) respectively, which are the malignant probabilities of patches/instances. Max pooling-based, label assignment-based, sparse deep MIL are in the second row, third row, fourth row respectively.	26
4.1	The framework of DeepLung. DeepLung first employs 3D Faster R-CNN to generate candidate nodules. Then it uses deep 3D DPN to extract deep features from the detected and cropped nodules. Lastly, GBM with deep features, detected nodule size, and raw pixels is employed for classification. Patient-level diagnosis can be achieved by fusing the classification results of detected nodules in the CT.	30
4.2	Illustration of dual path connection [12], which benefits both from the advantage of residual learning [43] and that of dense connection [47] from network structure design intrinsically.	34
4.3	The 3D Faster R-CNN framework contains 3D dual path blocks and a U-net-like encoder-decoder structure. We design 26 layers 3D dual path network for the encoder subnetwork. The model employs 3 anchors and multi-task learning loss, including coordinates (x, y, z) and diameter d regression, and candidate box classification. The numbers in boxes are feature map sizes in the format (#slices*#rows*#cols*#maps). The numbers above the connections are in the format (#filters #slices*#rows*#cols).	36
4.4	The deep 3D dual path network framework in the nodule classification subnetwork, which contains 30 3D dual path connection blocks. After the training, the deep 3D dual path network feature is extracted for gradient boosting machine to do nodule diagnosis. The numbers are of the same formats as Fig. 4.3.	37
4.5	The 3D Faster R-CNN network with 3D residual blocks. It contains several 3D residual blocks. We employ a deep 3D residual network of 18 layers as the encoder subnetwork, which is an extension from 2D Res18 net [43].	41
4.6	Sensitivity (Recall) rate with respect to false positives per scan. The FROC (average recall rate at the false positives as 0.125, 0.25, 0.5, 1, 2, 4, 8) of 3D Res18 Faster R-CNN is 83.4%, while the FROC of 3D DPN26 Faster R-CNN is 84.2% with only $\frac{1}{4}$ of the parameters as 3D Res18 Faster R-CNN. The 3D Res18 Faster R-CNN has a total recall rate 94.6% for all the detected nodules, while 3D DPN26 Faster R-CNN has a recall rate 95.8%	42
4.7	Visualization of central slices for nodule ground truths and detection results. We randomly choose nodules (red circle boxes in the first row) from test fold 1. Detection results are shown in the blue circles of second row. The center slice numbers are shown below the images. The last row shows detection probability. The DeepLung performs well for nodule detection.	48

4.8	Visualization of central slices for nodule classification results on test fold 1. We choose nodules that are predicted right by the DeepLung, but annotated wrong by some doctors. The numbers below the nodules are model predicted malignant probabilities followed by which annotation of doctors is wrong. The first seven nodules are benign nodules. The rest nodules are malignant nodules. The DeepLung performs well for nodule classification.	49
5.1	Illustration of DeepEM framework. Faster R-CNN is employed for nodule proposal generation. Half-Gaussian model and logistic regression are employed for central slice and lobe location respectively. In the E-step, we utilize all the observations, CT slices, and weak label to infer the latent variable, nodule proposals, by maximum a posteriori (MAP) or sampling. In the M-step, we employ the estimated proposals to update parameters in the Faster R-CNN and logistic regression.	52
5.2	(a) Empirical estimation of half-Gaussian model for $P(z_m \mathbf{H}_m; \boldsymbol{\theta})$ on LUNA16. (b) FROC (%) comparison among Faster R-CNN, DeepEM with MAP, DeepEM with Sampling on LUNA16.	58
5.3	Detection visual comparison among Faster R-CNN, DeepEM with MAP and DeepEM with Sampling on nodules randomly sampled from Tianchi. DeepEM provides more accurate detection (central slice, center and diameter) than Faster R-CNN.	60
6.1	The AnatomyNet is a variant of U-Net with only one down-sampling and squeeze-and-excitation (SE) residual building blocks. The number before symbol @ denotes the number of output channels, while the number after the symbol denotes the size of feature map relative to the input. In the decoder, we use concatenated features. Hybrid loss with dice loss and focal loss is employed to force the model to learn not-well-classified voxels. Masked and weighted loss function is used for ground truth with missing annotations and balanced gradient descent respectively. The decoder layers are symmetric with the encoder layers. The SE residual block is illustrated in the upper right corner.	69
6.2	The frequency of voxels for each class on MICCAI 2015 challenge dataset. Background takes up 98.18% of all the voxels. Chiasm takes only 0.35% of the foreground which means it only takes about 1/100,000 of the whole-volume CT image. The huge imbalance of voxels in small-volumed organs causes difficulty for small-volumed organ segmentation.	70
6.3	Visualizations on four test CT images. Rows from top to bottom represent brain stem, brain stem 3D, mandibular, mandibular 3D, parotid left and right, and parotid left and right 3D. Each column represents one CT image. Green represents ground truths, and red denotes predictions. Yellow is the overlap.	86

6.4	Visualizations on four test CT images. Rows from top to bottom represent submandibular left and right, submandibular left and right 3D, optic nerve left and right, and chiasm. Each column represents one CT image. Green represents ground truths, and red denotes predictions. Yellow is the overlap. The AnatomyNet performs well on small-volumed anatomies.	87
6.5	Visualizations for the first four anatomy on the first four holdout CT images. There is no ground truth for mandible and submandibular glands. Because this is a different source from MICCAI 2015, the annotations of brain stem and chiasm are inconsistent with those from MICCAI 2015. The AnatomyNet generalizes well for hold out test set.	88
6.6	Visualizations for the rest five anatomies on the first four holdout CT images.	89

LIST OF TABLES

	Page
2.1 Kernel sizes of sub-nets ($\#kernel \times \#width \times \#height$).	14
2.2 Dices (%) on INbreast and DDSM-BCRP datasets.	14
3.1 Accuracy Comparisons of the proposed deep MILs and related methods on test sets.	25
4.1 Nodule classification comparisons on LIDC-IDRI dataset.	43
4.2 Nodule-level diagnosis accuracy (%) between nodule classification subnetwork in DeepLung and experienced doctors on doctor’s individually confident nodules.	44
4.3 Statistical property of predicted malignant probability for borderline nodules (%)	44
4.4 Comparison between DeepLung’s nodule classification on all detected nodules and doctors on all nodules.	46
4.5 Patient-level diagnosis accuracy(%) between DeepLung and experienced doctors on doctor’s individually confident CTs.	46
5.1 FROC (%) comparisons among Faster R-CNN with 3D ResNet18, DeepEM with MAP, DeepEM with Sampling on Tianchi.	60
6.1 The numbers of the nine annotated anatomies from 261 head and neck training CT images.	72
6.2 Sizes of encoder blocks in U-Nets with different numbers of down-samplings. The number before symbol @ denotes the number of output channels, while the number after the symbol denotes the size of feature map relative to the input.	74
6.3 Performances of U-Net models with different numbers of down-sampling layers, measured with Dice coefficients.	75
6.4 Performance comparison on different network structures	77
6.5 Comparisons of test performances of models trained with different loss functions, evaluated with Dice coefficients.	78
6.6 Performance comparisons with state-of-the-art methods, showing the average DSC on the test set.	79
6.7 Performance comparisons of models trained with different datasets.	82
6.8 Average 95th percentile Hausdorff distance (unit: mm) comparisons with state-of-the-art methods	83

LIST OF ALGORITHMS

	Page
1 DeepEM for Weakly Supervised Detection	57

ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisor, Professor Xiaohui Xie. Your guidance, expertise, insights, wisdom and optimism have been instrumental to my PhD career. Next, thank you to my dear wife, Dr. Yufang Huang. She gives me a family with my lovely little son. It is the greatest happiness being a husband and a father. And it is the kind of happiness that accompanies me during my PhD study. Thank you to my mother-in-law helping me take care of my little son. Thank you to my parents supporting me all the time. Thank you to Professor Charless C. Fowlkes, Professor Sameer Singh, Professor Jack Xin and Professor Ramesh Jain for your expertise and insights in my thesis and advancement. Finally, I want to thank all the friends I made and mentors during my PhD study: Yeeleng Scott Vang, Elmira Forouzmand, Yifei Chen, Cuiling Lan and Wenjun Zeng in Microsoft Research, Yi Wang and Xiaodi Hou in TuSimple, Chaochun Liu (now at JD.com) in Baidu Research, Wei Fan and Zhen Qian in Tencent Medical AI Labs, Le Lu in PingAn Tech, and many many other friends. It is your help and the friendship that support me in the fruitful and exciting PhD life.

I came into UCI through the Computer Science PhD program. Thank you to everyone involved in the program a tremendous amount of gratitude providing such a high quality and standard program.

I gratefully acknowledge Breast Research Group, INESC Porto, Portugal, LUng Nodule Analysis 2016 (LUNA16), National Institute of Health (NIH), National Cancer Institute (NCI), The National Lung Screening Trial (NLST), The Cancer Imaging Archive (TCIA), Alibaba Cloud TianChi providing datasets for my PhD research. I gratefully acknowledge UCI Associated Graduate students (AGS), AAAI 2016, NIPS 2018, Tencent giving me travel grants and awards to attend conferences.

CURRICULUM VITAE

Wentao Zhu

EDUCATION

Doctor of Philosophy in Computer Science University of California, Irvine	2019 <i>Irvine, California</i>
Master of Science in Computer Sciences Institute of Computing Technology, Chinese Academy of Sciences	2015 <i>Beijing, China</i>
Bachelor of Science in Computational Mathematics Shandong University	2012 <i>Shandong, China</i>

RESEARCH EXPERIENCE

Graduate Research Assistant University of California, Irvine	2015–2019 <i>Irvine, California</i>
--	---

TEACHING EXPERIENCE

Reader	2015–2017
Reader University of California, Irvine	2019 <i>Irvine, California</i>

REFEREED JOURNAL PUBLICATIONS

AnatomyNet: Deep learning for fast and fully automated whole-volume segmentation of head and neck anatomy 2018
Medical Physics

REFEREED CONFERENCE PUBLICATIONS

DeepEM: Deep 3D ConvNets with EM for weakly supervised pulmonary nodule detection September 2018
MICCAI

Deeplung: Deep 3d dual path nets for automated pulmonary nodule detection and classification March 2018
WACV

Adversarial deep structured nets for mass segmentation from mammograms April 2018
ISBI

Deep multi-instance networks with sparse label assignment for whole mammogram classification September 2017
MICCAI

Co-Occurrence feature learning for skeleton based action recognition using regularized deep LSTM networks February 2016
AAAI

SOFTWARE

AnatomyNet

<https://github.com/wentaozhu/AnatomyNet-for-anatomical-segmentation>

Deep learning for fast and fully automated whole-volume segmentation of head and neck anatomy.

DeepEM

<https://github.com/wentaozhu/DeepEM-for-Weakly-Supervised-Detection>

Deep 3D ConvNets with EM for weakly supervised pulmonary nodule detection.

DeepLung

<https://github.com/wentaozhu/DeepLung>

Deep 3d dual path nets for automated pulmonary nodule detection and classification.

Adversarial DSN

<https://github.com/wentaozhu/adversarial-deep-structural-networks>

Adversarial Deep Structural Networks for Mammographic Mass Segmentation.

Deep MIL

<https://github.com/wentaozhu/deep-mil-for-whole-mammogram-classification>

Deep multi-instance networks with sparse label assignment for whole mammogram classification.

Regularized Deep LSTM

<http://www.escience.cn/system/file?fileId=87579>

Co-Occurrence feature learning for skeleton based action recognition using regularized deep LSTM networks.

ABSTRACT OF THE DISSERTATION

Deep Learning for Automated Medical Image Analysis

By

Wentao Zhu

Doctor of Philosophy in Computer Science

University of California, Irvine, 2019

Professor Xiaohui Xie, Chair

Medical imaging is an essential tool in many areas of medical applications, used for both diagnosis and treatment. However, reading medical images and making diagnosis or treatment recommendations require specially trained medical specialists. The current practice of reading medical images is labor-intensive, time-consuming, costly, and error-prone. It would be more desirable to have a computer-aided system that can automatically make diagnosis and treatment recommendations.

Recent advances in deep learning enable us to rethink the ways of clinician diagnosis based on medical images. Early detection has proven to be critical to give patients the best chance of recovery and survival. Advanced computer-aided diagnosis systems are expected to have high sensitivities and small low positive rates. How to provide accurate diagnosis results and explore different types of clinical data is an important topic in the current computer-aided diagnosis research.

In this thesis, we will introduce 1) mammograms for detecting breast cancers, the most frequently diagnosed solid cancer for U.S. women, 2) lung Computed Tomography (CT) images for detecting lung cancers, the most frequently diagnosed malignant cancer, and 3) head and neck CT images for automated delineation of organs at risk in radiotherapy. First, we will show how to employ the adversarial concept to generate the hard examples

improving mammogram mass segmentation. Second, we will demonstrate how to use the weakly labelled data for the mammogram breast cancer diagnosis by efficiently design deep learning for multi-instance learning. Third, the thesis will walk through DeepLung system which combines deep 3D ConvNets and Gradient Boosting Machine (GBM) for automated lung nodule detection and classification. Fourth, we will show how to use weakly labelled data to improve existing lung nodule detection system by integrating deep learning with probabilistic graphic model. Lastly, we will demonstrate the AnatomyNet which is thousands faster and more accurate than previous methods on automated anatomy segmentation.

Chapter 1

Introduction

Deep learning has been a powerful and successful tool to lead the era of artificial intelligence (AI) in recent a few years. It has achieved surprising or even over human-level performance on image classification [43], speech recognition [109], reading comprehension [17] etc. On the other hand, autonomous driving, personal assistant devices (Google home, Alexa etc), AI for health-care are emerging topics with the aforementioned breakthroughs from deep learning. In this thesis, we focus on several topics, including mammograms, pulmonary computed tomography (CT), head and neck CT images, and deep learning for automated medical image analysis which potentially assists radiologists to improve the diagnosis.

Breast cancer is the second common cancer for women in United States, and mammogram screening is shown to be an effective way to reduce the death caused by breast cancer [71]. However, traditional computer-aided diagnosis system is exposed to be of high false positive rate. The main difficulty for deep learning based medical image computing research is that the scarce of data leads to easily over-fitting [19]. The thesis tries to alleviate the challenge from two directions. Firstly, we try to improve the generalization ability for deep learning based segmentation [121], which is an essential problem for medical image analysis because

the mass segmentation produces morphological features which provides effective evidence for mammogram diagnosis. Secondly, compared with detection or segmentation ground-truth, image-level labels can be easily obtained from electronic medical report (EMR). We design a novel framework for breast cancer diagnosis using image-level label which is different from traditional methods relying on regions of interest (ROIs) which require great efforts to annotate [21, 119].

Lung cancer is the cancer causing most death. Computed tomography (CT) screening is an effective way for accurate and early diagnosis which significantly improves the survival rate. Recently deep learning has been witnessing widespread adoption in various medical image applications. Considering CT images are spatial 3D data, we firstly construct 3D convolutional neural networks for lung CT nodule analysis system which includes nodule detection and classification [118]. However, training complicated deep neural nets requires large-scale datasets labeled with ground truth, which are often unavailable or costly and difficult to scale in many medical image domains. On the other hand, electronic medical records (EMR) consists of plenty of partial information on the content of each medical image. We explore how to tap this vast weakly labeled data from electronic medical records (EMR), but currently unexplored data source to improve pulmonary nodule detection [120].

Radiation therapy (RT) is a common treatment option for head and neck (HaN) cancer. An important step involved in RT planning is the delineation of organs-at-risks (OARs) based on HaN computed tomography (CT). However, manually delineating OARs is time-consuming as each slice of CT images needs to be individually examined and a typical CT consists of hundreds of slices. Automating OARs segmentation has the benefit of both reducing the time and improving the quality of RT planning. Existing anatomy auto-segmentation algorithms use primarily atlas-based methods, which are time-consuming and require sophisticated atlas creation and cannot adequately account for anatomy variations among patients. Designing a fast automated OARs delineation system with fewer radiologists correction can reduce the

delay time for treatment and potentially improve the performance of radiotherapy.

1.1 Dissertation Outline and Contributions

The thesis is outlined as follows:

In chapter 2, we propose a novel end-to-end network for mammographic mass segmentation which employs a fully convolutional network (FCN) to model a potential function, followed by a conditional random field (CRF) to perform structured learning [121]. Because the mass distribution varies greatly with pixel position in the ROIs, the FCN is combined with a position priori. Further, we employ adversarial training to eliminate over-fitting due to the small sizes of mammogram datasets. Multi-scale FCN is employed to improve the segmentation performance. Experimental results on two public datasets, INbreast and DDSM-BCRP, demonstrate that our end-to-end network achieves better performance than state-of-the-art approaches. These contributions are released as an open-source software package called adversarial-deep-structural-networks, which are publicly available¹. Portions of this chapter were published as part of [121].

In chapter 3, we propose end-to-end trained deep multi-instance networks for mass classification based on whole mammogram without the aforementioned ROIs [119], inspired by the success of using deep convolutional features for natural image analysis and multi-instance learning (MIL) for labeling a set of instances/patches. We explore three different schemes to construct deep multi-instance networks for whole mammogram classification. Experimental results on the INbreast dataset demonstrate the robustness of proposed networks compared to previous work using segmentation and detection annotations. These contributions are released as an open-source software package called deep-mil-for-whole-mammogram-

¹<https://github.com/wentaozhu/adversarial-deep-structural-networks>

classification, which are publicly available². Portions of this chapter were published as part of [119].

In chapter 4, we present a fully automated lung computed tomography (CT) cancer diagnosis system, DeepLung [118]. DeepLung consists of two components, nodule detection (identifying the locations of candidate nodules) and classification (classifying candidate nodules into benign or malignant). Considering the 3D nature of lung CT data and the compactness of dual path networks (DPN), two deep 3D DPN are designed for nodule detection and classification respectively. Specifically, a 3D Faster Regions with Convolutional Neural Net (R-CNN) is designed for nodule detection with 3D dual path blocks and a U-net-like encoder-decoder structure to effectively learn nodule features. For nodule classification, gradient boosting machine (GBM) with 3D dual path network features is proposed. The nodule classification subnetwork is validated on a public dataset from LIDC-IDRI, on which it achieves better performance than the state-of-the-art approaches and surpasses the performance of experienced doctors based on image modality. Within the DeepLung system, candidate nodules are detected first by the nodule detection subnetwork, and nodule diagnosis is conducted by the classification subnetwork. Extensive experimental results demonstrate that DeepLung has performance comparable to the experienced doctors both for the nodule-level and patient-level diagnosis on the LIDC-IDRI dataset. These contributions are released as an open-source software package called DeepLung, which are publicly available³. Portions of this chapter were published as part of [118].

In chapter 5, we propose DeepEM, a novel deep 3D ConvNet framework augmented with expectation-maximization (EM), to mine weakly supervised labels in EMRs for pulmonary nodule detection [120]. Experimental results show that DeepEM can lead to 1.5% and 3.9% average improvement in free-response receiver operating characteristic (FROC) scores on LUNA16 and Tianchi datasets, respectively, demonstrating the utility of incomplete infor-

²<https://github.com/wentaozhu/deep-mil-for-whole-mammogram-classification>

³<https://github.com/wentaozhu/DeepLung>

mation in EMRs for improving deep learning algorithms. These contributions are released as an open-source software package called DeepEM, which are publicly available⁴. Portions of this chapter were published as part of [120].

In chapter 6, we propose an end-to-end, atlas-free 3D convolutional deep learning framework for fast and fully automated whole-volume HaN anatomy segmentation [116]. Our deep learning model, called AnatomyNet, segments OARs from head and neck CT images in an end-to-end fashion, receiving whole-volume HaN CT images as input and generating masks of all OARs of interest in one shot. AnatomyNet is built upon the popular 3D U-net architecture, but extends it in three important ways: 1) a new encoding scheme to allow auto-segmentation on whole-volume CT images instead of local patches or subsets of slices, 2) incorporating 3D squeeze-and-excitation residual blocks in encoding layers for better feature representation, and 3) a new loss function combining Dice scores and focal loss to facilitate the training of the neural model. These features are designed to address two main challenges in deep-learning-based HaN segmentation: a) segmenting small anatomies (i.e., optic chiasm and optic nerves) occupying only a few slices, and b) training with inconsistent data annotations with missing ground truth for some anatomical structures. We collect 261 HaN CT images to train AnatomyNet, and use MICCAI Head and Neck Auto Segmentation Challenge 2015 as a benchmark dataset to evaluate the performance of AnatomyNet. The objective is to segment nine anatomies: brain stem, chiasm, mandible, optic nerve left, optic nerve right, parotid gland left, parotid gland right, submandibular gland left, and submandibular gland right. Compared to previous state-of-the-art results from the MICCAI 2015 competition, AnatomyNet increases Dice similarity coefficient by 3.3% on average. AnatomyNet takes about 0.12 seconds to fully segment a head and neck CT image of dimension $178 \times 302 \times 225$, significantly faster than previous methods. In addition, the model is able to process whole-volume CT images and delineate all OARs in one pass, requiring little pre- or post-processing. Deep learning models offer a feasible solution to the problem of

⁴<https://github.com/wentaozhu/DeepEM-for-Weakly-Supervised-Detection>

delineating OARs from CT images. We demonstrate that our proposed model can improve segmentation accuracy and simplify the auto-segmentation pipeline. With this method, it is possible to delineate OARs of a head and neck CT within a fraction of a second. These contributions are released as an open-source software package called AnatomyNet, which are publicly available⁵. Portions of this chapter were published as part of [116].

⁵<https://github.com/wentaozhu/AnatomyNet-for-anatomical-segmentation>

Chapter 2

Adversarial Deep Structured Nets for Mass Segmentation from Mammograms

2.1 Introduction

According to the American Cancer Society, breast cancer is the most frequently diagnosed solid cancer and the second leading cause of cancer death among U.S. women. Mammogram screening has been proven to be an effective way for early detection and diagnosis, which significantly decrease breast cancer mortality. Mass segmentation provides morphological features, which play crucial roles for diagnosis.

Traditional studies on mass segmentation rely heavily on hand-crafted features. Model-based methods build classifiers and learn features from masses [6, 9]. There are few works using deep networks for mammogram [30]. Dhungel et al. employed multiple deep belief networks (DBNs), Gaussian mixture model (GMM) classifier and a priori as potential functions, and

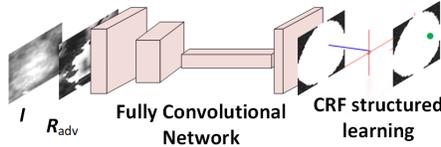


Figure 2.1: The proposed adversarial deep FCN-CRF network with four convolutional layers followed by CRF for structured learning.

structured support vector machine (SVM) to perform segmentation [19]. They further used CRF with tree re-weighted belief propagation to boost the segmentation performance [20]. A recent work used the output from a convolutional network (CNN) as a complimentary potential function, yielding the state-of-the-art performance [18]. However, the two-stage training used in these methods produces potential functions that easily over-fit the training data.

In this work, we propose an end-to-end trained adversarial deep structured network to perform mass segmentation (Fig. 2.1). The proposed network is designed to robustly learn from a small dataset with poor contrast mammographic images. Specifically, an end-to-end trained FCN with CRF is applied. Adversarial training is introduced into the network to learn robustly from scarce mammographic images. Different from DI2IN-AN using a generative framework [28], we directly optimize pixel-wise labeling loss. To further explore statistical property of mass regions, a spatial priori is integrated into FCN. We validate the adversarial deep structured network on two public mammographic mass segmentation datasets. The proposed network is demonstrated to outperform other algorithms for mass segmentation consistently.

Our main contributions in this work are: (1) We propose an unified end-to-end training framework integrating FCN+CRF and adversarial training. (2) We employ an end-to-end network to do mass segmentation while previous works require a lot of hand-designed features or multi-stage training. (3) Our model achieves the best results on two most commonly used mammographic mass segmentation datasets.

2.2 FCN-CRF Network

Fully convolutional network (FCN) is a commonly used model for image segmentation, which consists of convolution, transpose convolution, or pooling [65]. For training, the FCN optimizes maximum likelihood loss function

$$\mathcal{L}_{FCN} = -\frac{1}{N \times N_i} \sum_{n=1}^N \sum_{i=1}^{N_i} \log P_{fcn}(y_{n,i} | \mathbf{I}_n; \boldsymbol{\theta}), \quad (2.1)$$

where $y_{n,i}$ is the label of i th pixel in the n th image \mathbf{I}_n , N is the number of training mammograms, N_i is the number of pixels in the image, and $\boldsymbol{\theta}$ is the parameter of FCN. Here the size of images is fixed to 40×40 and N_i is 1,600.

CRF is a classical model for structured learning, well suited for image segmentation. It models pixel labels as random variables in a Markov random field conditioned on an observed input image. To make the annotation consistent, we use $\mathbf{y} = (y_1, y_2, \dots, y_i, \dots, y_{1,600})^T$ to denote the random variables of pixel labels in an image, where $y_i \in \{0, 1\}$. The zero denotes pixel belonging to background, and one denotes it belonging to mass region. The Gibbs energy of fully connected pairwise CRF is [58]

$$E(\mathbf{y}) = \sum_i \psi_u(y_i) + \sum_{i < j} \psi_p(y_i, y_j), \quad (2.2)$$

where unary potential function $\psi_u(y_i)$ is the loss of FCN in our case, pairwise potential function $\psi_p(y_i, y_j)$ defines the cost of labeling pair (y_i, y_j) ,

$$\psi_p(y_i, y_j) = \mu(y_i, y_j) \sum_m w^{(m)} k_G^{(m)}(\mathbf{f}_i, \mathbf{f}_j), \quad (2.3)$$

where label compatibility function μ is given by the Potts model in our case, $w^{(m)}$ is the learned weight, pixel values I_i and positions p_i can be used as the feature vector \mathbf{f}_i , $k_G^{(m)}$ is

the Gaussian kernel applied to feature vectors [58],

$$k_G(\mathbf{f}_i, \mathbf{f}_j) = [\exp(-|I_i - I_j|^2/2), \exp(-|p_i - p_j|^2/2)]^T. \quad (2.4)$$

Efficient inference algorithm can be obtained by mean field approximation $Q(\mathbf{y}) = \prod_i Q_i(y_i)$ [58]. The update rule is

$$\begin{aligned} \tilde{Q}_i^{(m)}(l) &\leftarrow \sum_{i \neq j} k_G^{(m)}(\mathbf{f}_i, \mathbf{f}_j) Q_j(l) \text{ for all } m, \\ \check{Q}_i(l) &\leftarrow \sum_m w^{(m)} \tilde{Q}_i^{(m)}(l), \quad \hat{Q}_i(l) \leftarrow \sum_{l' \in \mathcal{L}} \mu(l, l') \check{Q}_i(l'), \\ \check{\check{Q}}_i(l) &\leftarrow \exp(-\psi_u(y_i = l)) - \hat{Q}_i(l), \quad Q_i \leftarrow \frac{1}{Z_i} \exp(\check{\check{Q}}_i(l)), \end{aligned} \quad (2.5)$$

where the first equation is the message passing from label of pixel i to label of pixel j , the second equation is re-weighting with the learned weights $w^{(m)}$, the third equation is compatibility transformation, the fourth equation is adding unary potentials, and the last step is normalization. Here $\mathcal{L} = \{0, 1\}$ denotes background or mass. The initialization of inference employs unary potential function as $Q_i(y_i) = \frac{1}{Z_i} \exp(-\psi_u(y_i))$. The mean field approximation can be interpreted as a recurrent neural network (RNN) [115].

2.3 Adversarial FCN-CRF Nets

The shape and appearance priori play important roles in mammogram mass segmentation [18, 29]. The distribution of labels varies greatly with position in the mammographic mass segmentation. From observation, most of the masses are located in the center of region of interest (ROI), and the boundary areas of ROI are more likely to be background (Fig. 2.2(a)).

The conventional FCN provides independent pixel-wise predictions. It considers global class



Figure 2.2: The empirical estimation of a priori on INbreast (left) and DDSM-BCRP (right) training datasets (a). Trimap visualizations on the DDSM-BCRP dataset, segmentation groundtruth (first column), trimap of width 2 (second column), trimaps of width 3 (third column) (b).

distribution difference corresponding to bias in the last layer. Here we employ a priori for position into consideration

$$P(y_i|\mathbf{I};\boldsymbol{\theta}) \propto w_i P_{fcn}(y_i|\mathbf{I};\boldsymbol{\theta}), \quad (2.6)$$

where w_i is the empirical estimation of mass varied with the pixel position i , and $P_{fcn}(y_i|\mathbf{I};\boldsymbol{\theta})$ is the predicted mass probability of conventional FCN. In the implementation, we added an image sized bias in the softmax layer as the empirical estimation of mass for FCN to train network. The $-\log P(y_i|\mathbf{I};\boldsymbol{\theta})$ is used as the unary potential function for $\psi_u(y_i)$ in the CRF as RNN. For multi-scale FCN as potential functions, the potential function is defined as $\psi_u(y_i) = \sum_{u'} w_{(u')} \psi_{u'}(y_i)$, where $w_{(u')}$ is the learned weight for unary potential function, $\psi_{u'}(y_i)$ is the potential function provided by FCN of each scale.

Adversarial training provides strong regularization for deep networks. The idea of adversarial training is that if the model is robust enough, it should be invariant to small perturbations of training examples that yield the largest increase in the loss (adversarial examples [91]). The perturbation \mathbf{R} can be obtained as $\min_{\mathbf{R}, \|\mathbf{R}\| \leq \epsilon} \log P(\mathbf{y}|\mathbf{I} + \mathbf{R}; \boldsymbol{\theta})$. In general, the calculation of exact \mathbf{R} is intractable especially for complicated models such as deep networks. The linear approximation and L_2 norm box constraint can be used for the calculation of perturbation as $\mathbf{R}_{adv} = -\frac{\epsilon \mathbf{g}}{\|\mathbf{g}\|_2}$, where $\mathbf{g} = \nabla_{\mathbf{I}} \log P(\mathbf{y}|\mathbf{I}; \boldsymbol{\theta})$. For adversarial FCN, the network predicts label of each pixel independently as $p(\mathbf{y}|\mathbf{I}; \boldsymbol{\theta}) = \prod_i P(y_i|\mathbf{I}; \boldsymbol{\theta})$. For adversarial CRF as RNN, the prediction of network relies on mean field approximation inference as $p(\mathbf{y}|\mathbf{I}; \boldsymbol{\theta}) =$

$\prod_i Q(y_i|\mathbf{I};\boldsymbol{\theta})$.

The adversarial training forces the model to fit examples with the worst perturbation direction. The adversarial loss is

$$\mathcal{L}_{adv}(\boldsymbol{\theta}) = -\frac{1}{N} \sum_{n=1}^N \log P(\mathbf{y}_n|\mathbf{I}_n + \mathbf{R}_{adv,n};\boldsymbol{\theta}). \quad (2.7)$$

In training, the total loss is defined as the sum of adversarial loss and the empirical loss based on training samples as

$$\mathcal{L}(\boldsymbol{\theta}) = \mathcal{L}_{adv}(\boldsymbol{\theta}) - \frac{1}{N} \sum_{n=1}^N \log P(\mathbf{y}_n|\mathbf{I}_n;\boldsymbol{\theta}) + \frac{\lambda}{2} \|\boldsymbol{\theta}\|^2, \quad (2.8)$$

where λ is the l_2 regularization factor for $\boldsymbol{\theta}$, $P(\mathbf{y}_n|\mathbf{I}_n;\boldsymbol{\theta})$ is either mass probability prediction in the FCN or a posteriori approximated by mean field inference in the CRF as RNN for the n th image \mathbf{I}_n .

2.4 Experiments

We validate the proposed model on two most commonly used public mammographic mass segmentation datasets: INbreast [68] and DDSM-BCRP dataset [44]. We use the same ROI extraction and resize principle as [18, 19, 20]. Due to the low contrast of mammograms, image enhancement technique is used on the extracted ROI images as the first 9 steps in [5], followed by pixel position dependent normalization. The preprocessing makes training converge quickly. We further augment each training set by flipping horizontally, flipping vertically, flipping horizontally and vertically, which makes the training set 4 times larger than the original training set.

For consistent comparison, the Dice index metric is used to evaluate segmentation perfor-

mance and is defined as $\frac{2 \times TP}{2 \times TP + FP + FN}$. For a fair comparison, we re-implement a two-stage model [18], and obtain similar result (Dice index 0.9010) on the INbreast dataset.

- FCN is the network integrating a position priori into FCN (denoted as FCN 1 in Table 2.1).
- Adversarial FCN is FCN with adversarial training.
- Joint FCN-CRF is the FCN followed by CRF as RNN with an end-to-end training scheme.
- Adversarial FCN-CRF is the Jointly FCN-CRF with end-to-end adversarial training.
- Multi-FCN, Adversarial multi-FCN, Joint multi-FCN-CRF, Adversarial multi-FCN-CRF employ 4 FCNs with multi-scale kernels, which can be trained in an end-to-end way using the last prediction.

The prediction of Multi-FCN, Adversarial multi-FCN is the average prediction of the 4 FCNs. The configurations of FCNs are in Table 2.1. Each convolutional layer is followed by 2×2 max pooling. The last layers of the four FCNs are all two 40×40 transpose convolution kernels with soft-max activation function. We use hyperbolic tangent activation function in middle layers. The parameters of FCNs are set such that the number of each layer’s parameters is almost the same as that of CNN used in the work [18]. We use Adam with learning rate 0.003. The λ is 0.5 in the two datasets. The ϵ used in adversarial training are 0.1 and 0.5 for INbreast and DDSM-BCRP datasets respectively. Because the boundaries of masses on the DDSM-BCRP dataset are smoother than those on the INbreast dataset, we use larger perturbation ϵ . For the CRF as RNN, we use 5 time steps in the training and 10 time steps in the test phase empirically.

The INbreast dataset is a recently released mammographic mass analysis dataset, which provides more accurate contours of lesion region and the mammograms are of high quality.

Table 2.1: Kernel sizes of sub-nets ($\#kernel \times \#width \times \#height$).

Net.	First layer	Second layer	Third layer
FCN 1	$6 \times 5 \times 5$	$12 \times 5 \times 5$ conv.	$588 \times 7 \times 7$
FCN 2	$9 \times 4 \times 4$	$12 \times 4 \times 4$ conv.	$588 \times 7 \times 7$
FCN 3	$16 \times 3 \times 3$	$13 \times 3 \times 3$ conv.	$415 \times 8 \times 8$
FCN 4	$37 \times 2 \times 2$	$12 \times 2 \times 2$ conv.	$355 \times 9 \times 9$

Table 2.2: Dices (%) on INbreast and DDSM-BCRP datasets.

Methodology	INbreast	DDSM-BCRP
Cardoso et al. [9]	88	N/A
Beller et al. [6]	N/A	70
Deep Structure Learning [19]	88	87
TRW Deep Structure Learning [20]	89	89
Deep Structure Learning + CNN [18]	90	90
FCN	89.48	90.21
Adversarial FCN	89.71	90.78
Joint FCN-CRF	89.78	90.97
Adversarial FCN-CRF	90.07	91.03
Multi-FCN	90.47	91.17
Adversarial multi-FCN	90.71	91.20
Joint multi-FCN-CRF	90.76	91.26
Adversarial multi-FCN-CRF	90.97	91.30

For mass segmentation, the dataset contains 116 mass regions. We use the first 58 masses for training and the rest for test, which is of the same protocol as [18, 19, 20]. The DDSM-BCRP dataset contains 39 cases (156 images) for training and 40 cases (160 images) for testing [44]. After ROI extraction, there are 84 ROIs for training, and 87 ROIs for test. We compare schemes with other recently published mammographic mass segmentation methods in Table 2.2.

Table 2.2 shows the CNN features provide superior performance on mass segmentation, outperforming hand-crafted feature based methods [6, 9]. Our enhanced FCN achieves 0.25% Dice index improvement than the traditional FCN on the INbreast dataset. The adversarial training yields 0.4% improvement on average. Incorporating the spatially structured learning further produces 0.3% improvement. Using multi-scale model contributes the most to segmentation results, which shows multi-scale features are effective for pixel-wise classification in mass segmentation. Combining all the components together achieves the best performance with 0.97%, 1.3% improvement on INbreast, DDSM-BCRP datasets respectively. The possible reason for the improvement is adversarial scheme eliminates the over-fitting. We calculate

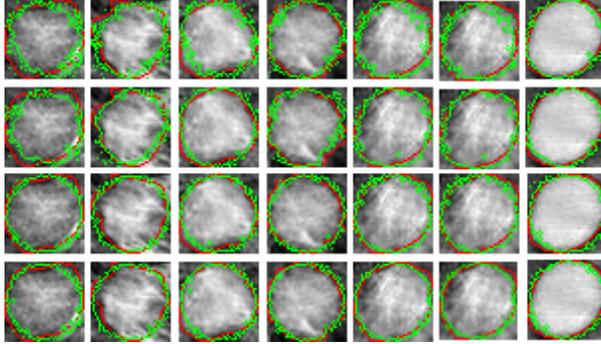


Figure 2.3: Visualization of segmentation results using the FCN (first row), Adversarial FCN (second row), Joint FCN-CRF (third row), Adversarial FCN-CRF (fourth row) on the test sets of INbreast dataset. Each column denotes a test sample. Red lines denote the ground truth. Green lines or points denote the segmentation results. Adversarial training provides sharper and more accurate segmentation boundaries than methods without adversarial training.

the p-value of McNemars Chi-Square Test to compare our model with [18] on the INbreast dataset. We obtain p-value < 0.001 , which shows our model is significantly better than model [18].

To better understand the adversarial training, we visualize segmentation results in Fig. 2.3. We observe that the segmentations in the second and fourth rows have more accurate boundaries than those of the first and third rows. It demonstrates the adversarial training improves FCN and FCN-CRF.

We further employ the prediction accuracy based on trimap to specifically evaluate segmentation accuracy in boundaries [56]. We calculate the accuracies within trimap surrounding the actual mass boundaries (groundtruth) in Fig. 2.4. Trimaps on the DDSM-BCRP dataset is visualized in Fig. 2.2(b). From the figure, accuracies of Adversarial FCN-CRF are 2-3 % higher than those of Joint FCN-CRF on average and the accuracies of Adversarial FCN are better than those of FCN. The above results demonstrate that the adversarial training improves the FCN and Joint FCN-CRF both for whole image and boundary region segmentation.



Figure 2.4: Accuracy comparisons among FCN, Adversarial FCN, Joint FCN-CRF and Adversarial FCN-CRF in trimaps with pixel width 1, 2, 3, 4, 5 on the INbreast dataset (a) and the DDSM-BCRP dataset (b). The adversarial training improves segmentation accuracy around boundaries.

2.5 Conclusion

In this work, we propose an end-to-end adversarial FCN-CRF network for mammographic mass segmentation. To integrate the priori distribution of masses and fully explore the power of FCN, a position priori is added to the network. Furthermore, adversarial training is used to handle the small size of training data by reducing over-fitting and increasing robustness. Experimental results demonstrate the superior performance of adversarial FCN-CRF on two commonly used public datasets.

Chapter 3

Deep Multi-instance Networks with Sparse Label Assignment for Whole Mammogram Classification

3.1 Introduction

According to the American Cancer Society, breast cancer is the most frequently diagnosed solid cancer and the second leading cause of cancer death among U.S. women [1]. Mammogram screening has been demonstrated to be an effective way for early detection and diagnosis, which can significantly decrease breast cancer mortality [71]. Traditional mammogram classification requires extra annotations such as bounding box for detection or mask ground truth for segmentation [10, 55, 98]. Other work have employed different deep networks to detect ROIs and obtain mass boundaries in different stages [21]. However, these methods require hand-crafted features to complement the system [57], and training data to be annotated with bounding boxes and segmentation ground truths which require expert

domain knowledge and costly effort to obtain. In addition, multi-stage training cannot fully explore the power of deep networks.

Due to the high cost of annotation, we intend to perform classification based on a raw whole mammogram. Each patch of a mammogram can be treated as an instance and a whole mammogram is treated as a bag of instances. The whole mammogram classification problem can then be thought of as a standard MIL problem. Due to the great representation power of deep features [39, 117, 122], combining MIL with deep neural networks is an emerging topic. Yan et al. used a deep MIL to find discriminative patches for body part recognition [112]. Patch based CNN added a new layer after the last layer of deep MIL to learn the fusion model for multi-instance predictions [45]. Shen et al. employed two stage training to learn the deep multi-instance networks for pre-detected lung nodule classification [84]. The above approaches used max pooling to model the general multi-instance assumption which only considers the patch of max probability. In this paper, more effective task-related deep multi-instance models with end-to-end training are explored for whole mammogram classification. We investigate three different schemes, i.e., max pooling, label assignment, and sparsity, to perform deep MIL for the whole mammogram classification task.

The framework for our proposed end-to-end trained deep MIL for mammogram classification is shown in Fig. 3.1. To fully explore the power of deep MIL, we convert the traditional MIL assumption into a label assignment problem. As a mass typically composes only 2% of a whole mammogram (see Fig. 3.2), we further propose sparse deep MIL. The proposed deep multi-instance networks are shown to provide robust performance for whole mammogram classification on the INbreast dataset [68].

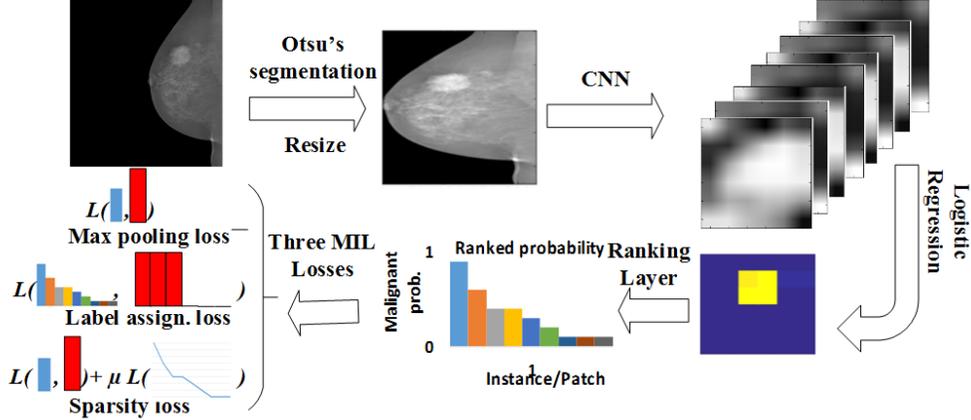


Figure 3.1: The framework of whole mammogram classification. First, we use Otsu’s segmentation to remove the background and resize the mammogram to 227×227 . Second, the deep MIL accepts the resized mammogram as input to the convolutional layers. Here we use the convolutional layers in AlexNet [59]. Third, logistic regression with weight sharing over different patches is employed for the malignant probability of each position from the convolutional neural network (CNN) feature maps of high channel dimensions. Then the responses of the instances/patches are ranked. Lastly, the learning loss is calculated using max pooling loss, label assignment, or sparsity loss for the three different schemes.

3.2 Deep MIL for Whole Mammogram Mass Classification

Unlike other deep multi-instance networks [45, 112], we use a CNN to efficiently obtain features of all patches (instances) at the same time. Given an image \mathbf{I} , we obtain a much smaller feature map \mathbf{F} of multi-channels N_c after multiple convolutional layers and max pooling layers. The $(\mathbf{F})_{i,j,:}$ represents deep CNN features for a patch $\mathbf{Q}_{i,j}$ in \mathbf{I} , where i, j represents the pixel row and column index respectively, and $:$ denotes the channel dimension.

The goal of our work is to predict whether a whole mammogram contains a malignant mass (BI-RADS $\in \{4, 5, 6\}$ as positive) or not, which is a standard binary classification problem. We add a logistic regression with weights shared across all the pixel positions following \mathbf{F} , and an element-wise sigmoid activation function is applied to the output. To clarify it, the

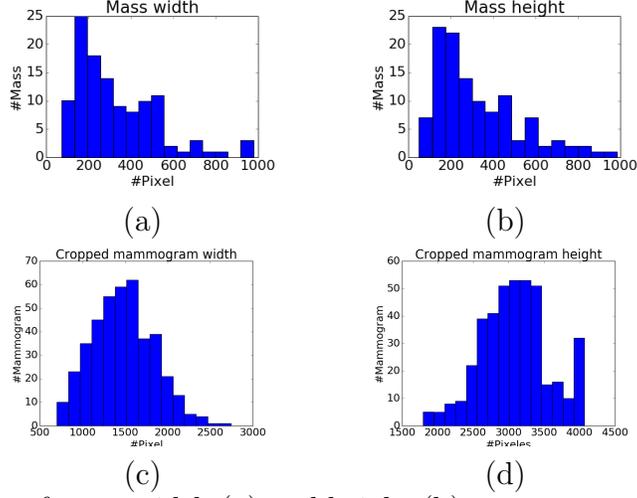


Figure 3.2: Histograms of mass width (a) and height (b), mammogram width (c) and height (d). Compared to the size of whole mammogram ($1,474 \times 3,086$ on average after cropping), the mass of average size (329×325) is tiny, and takes about 2% of a whole mammogram.

malignant probability of feature space’s pixel (i, j) is

$$r_{i,j} = \text{sigmoid}(\mathbf{a} \cdot \mathbf{F}_{i,j,:} + b), \quad (3.1)$$

where \mathbf{a} is the weights in logistic regression, b is the bias, and \cdot is the inner product of the two vectors \mathbf{a} and $\mathbf{F}_{i,j,:}$. The \mathbf{a} and b are shared for different pixel positions i, j . We can combine $r_{i,j}$ into a matrix $\mathbf{r} = (r_{i,j})$ of range $[0, 1]$ denoting the probabilities of patches being malignant masses. The \mathbf{r} can be flattened into a one-dimensional vector as $\mathbf{r} = (r_1, r_2, \dots, r_m)$ corresponding to flattened patches $(\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_m)$, where m is the number of patches.

3.2.1 Max Pooling-Based Multi-Instance Learning

The general multi-instance assumption is that if there exists an instance that is positive, the bag is positive [22]. The bag is negative if and only if all instances are negative. For whole mammogram classification, the equivalent scenario is that if there exists a malignant mass, the mammogram \mathbf{I} should be classified as positive. Likewise, negative mammogram

\mathbf{I} should not have any malignant masses. If we treat each patch \mathbf{Q}_i of \mathbf{I} as an instance, the whole mammogram classification is a standard multi-instance task.

For negative mammograms, we expect all the r_i to be close to 0. For positive mammograms, at least one r_i should be close to 1. Thus, it is natural to use the maximum component of \mathbf{r} as the malignant probability of the mammogram \mathbf{I}

$$p(y = 1|\mathbf{I}, \boldsymbol{\theta}) = \max\{r_1, r_2, \dots, r_m\}, \quad (3.2)$$

where $\boldsymbol{\theta}$ is the weights in deep networks.

If we sort \mathbf{r} first in descending order as illustrated in Fig. 3.1, the malignant probability of the whole mammogram \mathbf{I} is the first element of ranked \mathbf{r} as

$$\begin{aligned} \{r'_1, r'_2, \dots, r'_m\} &= \text{sort}(\{r_1, r_2, \dots, r_m\}), \\ p(y = 1|\mathbf{I}, \boldsymbol{\theta}) &= r'_1, \quad \text{and} \quad p(y = 0|\mathbf{I}, \boldsymbol{\theta}) = 1 - r'_1, \end{aligned} \quad (3.3)$$

where $\mathbf{r}' = (r'_1, r'_2, \dots, r'_m)$ is descending ranked \mathbf{r} . The cross entropy-based cost function can be defined as

$$\mathcal{L}_{maxpooling} = -\frac{1}{N} \sum_{n=1}^N \log(p(y_n|\mathbf{I}_n, \boldsymbol{\theta})) + \frac{\lambda}{2} \|\boldsymbol{\theta}\|^2 \quad (3.4)$$

where N is the total number of mammograms, $y_n \in \{0, 1\}$ is the true label of malignancy for mammogram \mathbf{I}_n , and λ is the regularizer that controls model complexity.

One disadvantage of max pooling-based MIL is that it only considers the patch \mathbf{Q}'_1 (patch of the max malignant probability), and does not exploit information from other patches. A more powerful framework should add task-related priori, such as sparsity of mass in whole mammogram, into the general multi-instance assumption and explore more patches for training.

3.2.2 Label Assignment-Based Multi-Instance Learning

For the conventional classification tasks, we assign a label to each data point. In the MIL scheme, if we consider each instance (patch) \mathbf{Q}_i as a data point for classification, we can convert the multi-instance learning problem into a label assignment problem.

After we rank the malignant probabilities $\mathbf{r} = (r_1, r_2, \dots, r_m)$ for all the instances (patches) in a whole mammogram \mathbf{I} using the first equation in Eq. 3.3, the first few r'_i should be consistent with the label of whole mammogram as previously mentioned, while the remaining patches (instances) should be negative. Instead of adopting the general MIL assumption that only considers the \mathbf{Q}'_1 (patch of malignant probability r'_1), we assume that 1) patches of the first k largest malignant probabilities $\{r'_1, r'_2, \dots, r'_k\}$ should be assigned with the same class label as that of whole mammogram, and 2) the rest patches should be labeled as negative in the label assignment-based MIL.

After the ranking/sorting layer using the first equation in Eq. 3.3, we can obtain the malignant probability for each patch

$$p(y = 1|\mathbf{Q}'_i, \boldsymbol{\theta}) = r'_i, \quad \text{and} \quad p(y = 0|\mathbf{Q}'_i, \boldsymbol{\theta}) = 1 - r'_i. \quad (3.5)$$

The cross entropy loss function of the label assignment-based MIL can be defined

$$\begin{aligned} \mathcal{L}_{labelassign.} = & -\frac{1}{mN} \sum_{n=1}^N \left(\sum_{j=1}^k \log(p(y_n|\mathbf{Q}'_j, \boldsymbol{\theta})) + \right. \\ & \left. \sum_{j=k+1}^m \log(p(y = 0|\mathbf{Q}'_j, \boldsymbol{\theta})) \right) + \frac{\lambda}{2} \|\boldsymbol{\theta}\|^2. \end{aligned} \quad (3.6)$$

One advantage of the label assignment-based MIL is that it explores all the patches to train the model. Essentially it acts a kind of data augmentation which is an effective technique

to train deep networks when the training data is scarce. From the sparsity perspective, the optimization problem of label assignment-based MIL is exactly a k -sparse problem for the positive data points, where we expect $\{r'_1, r'_2, \dots, r'_k\}$ being 1 and $\{r'_{k+1}, r'_{k+2}, \dots, r'_m\}$ being 0. The disadvantage of label assignment-based MIL is that it is hard to estimate the hyper-parameter k . Thus, a relaxed assumption for the MIL or an adaptive way to estimate the hyper-parameter k is preferred.

3.2.3 Sparse Multi-Instance Learning

From the mass distribution, the mass typically comprises about 2% of the whole mammogram on average (Fig. 3.2), which means the mass region is quite sparse in the whole mammogram. It is straightforward to convert the mass sparsity to the malignant mass sparsity, which implies that $\{r'_1, r'_2, \dots, r'_m\}$ is sparse in the whole mammogram classification problem. The sparsity constraint means we expect the malignant probability of part patches r'_i being 0 or close to 0, which is equivalent to the second assumption in the label assignment-based MIL. Analogously, we expect r'_1 to be indicative of the true label of mammogram \mathbf{I} .

After the above discussion, the loss function of sparse MIL problem can be defined

$$\mathcal{L}_{sparse} = \frac{1}{N} \sum_{n=1}^N \left(-\log(p(y_n|\mathbf{I}_n, \boldsymbol{\theta})) + \mu \|\mathbf{r}'_n\|_1 \right) + \frac{\lambda}{2} \|\boldsymbol{\theta}\|^2, \quad (3.7)$$

where $p(y_n|\mathbf{I}_n, \boldsymbol{\theta})$ can be calculated in Eq. 3.3, $\mathbf{r}'_n = (r'_1, r'_2, \dots, r'_m)$ for mammogram \mathbf{I}_n , $\|\cdot\|_1$ denotes the \mathcal{L}_1 norm, μ is the sparsity factor, which is a trade-off between the sparsity assumption and the importance of patch \mathbf{Q}'_1 .

From the discussion of label assignment-based MIL, this learning is a kind of exact k -sparse problem which can be converted to \mathcal{L}_1 constrain. One advantage of sparse MIL over label assignment-based MIL is that it does not require assign label for each patch which is hard

to do for patches where probabilities are not too large or small. The sparse MIL considers the overall statistical property of \mathbf{r} .

Another advantage of sparse MIL is that, it has different weights for general MIL assumption (the first part loss) and label distribution within mammogram (the second part loss), which can be considered as a trade-off between max pooling-based MIL (slack assumption) and label assignment-based MIL (hard assumption).

3.3 Experiments

We validate the proposed models on the most frequently used mammographic mass classification dataset, INbreast dataset [68], as the mammograms in other datasets, such as DDSM dataset [8], are of low quality. The INbreast dataset contains 410 mammograms of which 100 containing malignant masses. These 100 mammograms with malignant masses are defined as positive. For fair comparison, we also use 5-fold cross validation to evaluate model performance as [21]. For each testing fold, we use three folds for training, and one fold for validation to tune hyper-parameters. The performance is reported as the average of five testing results obtained from cross-validation.

We employ techniques to augment our data. For each training epoch, we randomly flip the mammograms horizontally, shift within 0.1 proportion of mammograms horizontally and vertically, rotate within 45 degree, and set 50×50 square box as 0. In experiments, the data augmentation is essential for us to train the deep networks.

For the CNN network structure, we use AlexNet and remove the fully connected layers [59]. Through CNN, the mammogram of size 227×227 becomes $256 \times 6 \times 6$ feature maps. Then we use steps in Sec. 3.2 to do MIL. Here we employ weights pretrained on the ImageNet due to the scarce of data. We use Adam optimization with learning rate 5×10^{-5} for training

Table 3.1: Accuracy Comparisons of the proposed deep MILs and related methods on test sets.

Methodology	Dataset	Set-up	Accu.	AUC
Ball et al. [5]	DDSM	Manual+feat.	0.87	N/A
Varela et al. [98]	DDSM	Manual+feat.	0.81	N/A
Domingues et al. [24]	INbr.	Manual+feat.	0.89	N/A
Pretrained CNN [21]	INbr.	Auto.+feat.	0.84±0.04	0.69±0.10
Pretrained CNN+Random Forest [21]	INbr.	Auto.+feat.	0.91 ± 0.02	0.76±0.23
AlexNet	INbr.	Auto.	0.81±0.02	0.79±0.03
AlexNet+Max Pooling MIL	INbr.	Auto.	0.85±0.03	0.83±0.05
AlexNet+Label Assign. MIL	INbr.	Auto.	0.86±0.02	0.84±0.04
AlexNet+Sparse MIL	INbr.	Auto.	0.90±0.02	0.89 ± 0.04

models [4]. The λ for max pooling-based and label assignment-based MIL are 1×10^{-5} . The λ and μ for sparse MIL are 5×10^{-6} and 1×10^{-5} respectively. For the label assignment-based MIL, we select k from $\{1, 2, 4, 6, 8\}$ based on the validation set.

We firstly compare our methods to previous models validated on DDSM dataset and INbreast dataset in Table 3.1. Previous hand-crafted feature-based methods require manually annotated detection bounding box or segmentation ground truth even in test denoting as manual [5, 24, 98]. The feat. denotes requiring hand-crafted features. Pretrained CNN uses two CNNs to detect the mass region and segment the mass, followed by a third CNN to do mass classification on the detected ROI region, which requires hand-crafted features to pretrain the network and needs multi-stages training[21]. Pretrained CNN+Random Forest further employs random forest and obtained 7% improvement. These methods are either manually or need hand-crafted features or multi-stages training, while our methods are totally automated, do not require hand-crafted features or extra annotations even on training set, and can be trained in an end-to-end manner.

The max pooling-based deep MIL obtains better performance than the pretrained CNN using 3 different CNNs and detection/segmentation annotation in the training set. This shows the superiority of our end-to-end trained deep MIL for whole mammogram classification. According to the accuracy metric, the sparse deep MIL is better than the label assignment-based MIL, which is better than the max pooling-based MIL. This result is consistent with

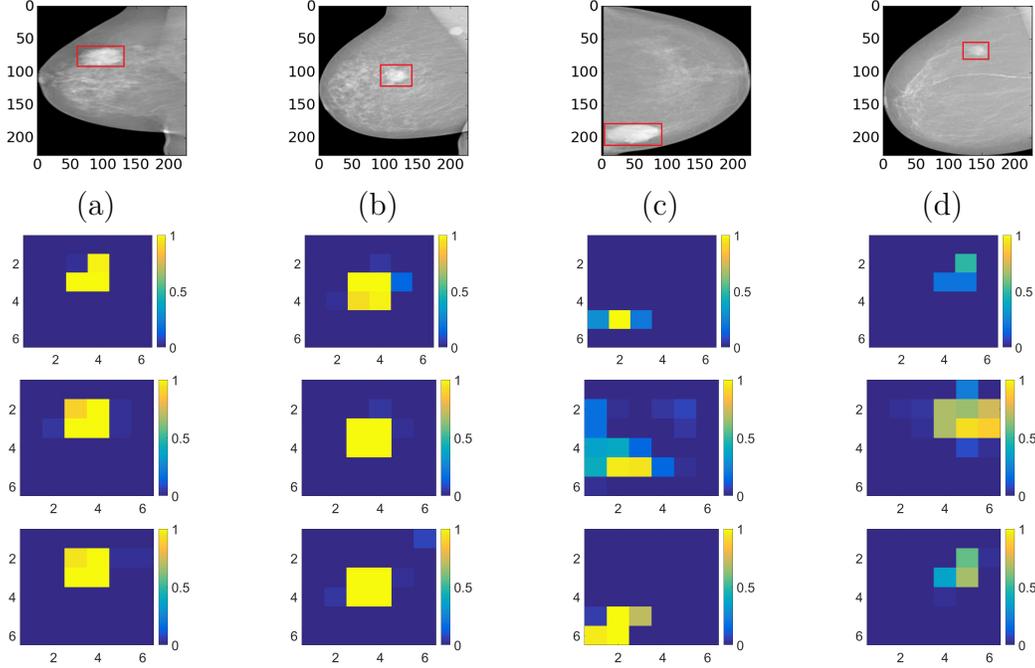


Figure 3.3: The visualization of predicted malignant probabilities for instances/patches in four resized mammograms. The first row is the resized mammogram. The red rectangle boxes are mass regions from the annotations on the dataset. The color images from the second row to the last row are the predicted malignant probability from logistic regression layer for (a) to (d) respectively, which are the malignant probabilities of patches/instances. Max pooling-based, label assignment-based, sparse deep MIL are in the second row, third row, fourth row respectively.

previous discussion that the sparsity assumption benefited from not having hard constraints of the label assignment assumption, which employs all the patches and is more efficient than max pooling assumption. Our sparse deep MIL achieves competitive accuracy to random forest-based pretrained CNN, while much higher AUC than previous work, which shows our method is more robust. The main reasons for the robust results using our models are as follows. Firstly, data augmentation is an important technique to increase scarce training datasets and proves useful here. Secondly, the transfer learning that employs the pretrained weights from ImageNet is effective for the INBreast dataset. Thirdly, our models fully explore all the patches to train our deep networks thereby eliminating any possibility of overlooking malignant patches by only considering a subset of patches. This is a distinct advantage over previous networks that employ several stages MIL consisting of detection and segmentation.

To further understand our deep MIL, we visualize the responses of logistic regression layer for four mammograms on test set, which represents the malignant probability of each patch, in Fig. 3.3. We can see the deep MIL learns not only the prediction of whole mammogram, but also the prediction of malignant patches within the whole mammogram. Our models are able to learn the mass region of the whole mammogram without any explicit bounding box or segmentation ground truth annotation of training data. The max pooling-based deep multi-instance network misses some malignant patches in (a), (c) and (d). The possible reason is that it only considers the patch of max malignant probability in training and the model is not well learned for all patches. The label assignment-based deep MIL misclassifies some patches in (d). The possible reason is that the model sets a constant k for all the mammograms, which causes some mis-classifications for small masses. One of the potential applications of our work is that these deep MIL networks could be used to do weak mass annotation automatically, which provides evidence for the diagnosis.

3.4 Conclusion

In this paper, we propose end-to-end trained deep MIL for whole mammogram classification. Different from previous work using segmentation or detection annotations, we conduct mass classification based on whole mammogram directly. We convert the general MIL assumption to label assignment problem after ranking. Due to the sparsity of masses, sparse MIL is used for whole mammogram classification. Experimental results demonstrate more robust performance than previous work even without detection or segmentation annotation in the training.

In future work, we plan to extend the current work by: 1) incorporating multi-scale modeling such as spatial pyramid to further improve whole mammogram classification, 2) employing the deep MIL to do annotation or provide potential malignant patches to assist diagnoses,

and 3) applying to large datasets and expected to have improvement if the big dataset is available.

Chapter 4

DeepLung: Deep 3D Dual Path Nets for Automated Pulmonary Nodule Detection and Classification

4.1 Introduction

Lung cancer is the most common cause of cancer-related death in men. Low-dose lung CT screening provides an effective way for early diagnosis, which can sharply reduce the lung cancer mortality rate. Advanced computer-aided diagnosis systems (CADs) are expected to have high sensitivities while at the same time maintaining low false positive rates. Recent advances in deep learning enable us to rethink the ways of clinician lung cancer diagnosis.

Current lung CT analysis research mainly includes nodule detection [23, 25], and nodule classification [49, 85, 86, 111]. There is few work on building a complete lung CT cancer diagnosis system for fully automated lung CT cancer diagnosis using deep learning, integrating both nodule detection and nodule classification. It is worth exploring a whole lung CT

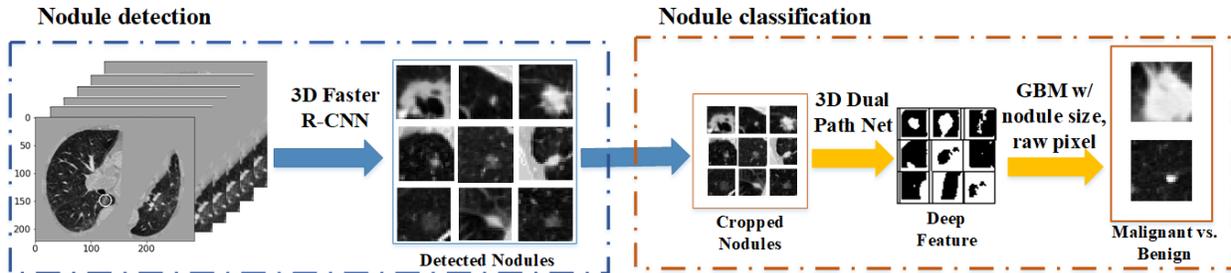


Figure 4.1: The framework of DeepLung. DeepLung first employs 3D Faster R-CNN to generate candidate nodules. Then it uses deep 3D DPN to extract deep features from the detected and cropped nodules. Lastly, GBM with deep features, detected nodule size, and raw pixels is employed for classification. Patient-level diagnosis can be achieved by fusing the classification results of detected nodules in the CT.

cancer diagnosis system and understanding how far the performance of current deep learning technology differs from that of experienced doctors. To our best knowledge, this is the first work for a fully automated and complete lung CT cancer diagnosis system using deep nets.

The emergence of large-scale dataset, LUNA16 [81], accelerates the nodule detection related research. Typically, nodule detection consists of two stages, region proposal generation and false positive reduction. Traditional approaches generally require manually designed features such as morphological features, voxel clustering and pixel thresholding [53, 69]. Recently, deep ConvNets, such as Faster R-CNN [62, 76] and fully ConvNets [65, 103, 104, 105, 121], are employed to generate candidate bounding boxes [23, 25]. In the second stage, more advanced methods or complex features, such as carefully designed texture features, are used to remove false positive nodules. Because of the 3D nature of CT data and the effectiveness of Faster R-CNN for object detection in 2D natural images [48], we design a 3D Faster R-CNN for nodule detection with 3D convolutional kernels and a U-net-like encoder-decoder structure to effectively learn latent features [78]. The U-Net structure is basically a convolutional autoencoder, augmented with skip connections between encoder and decoder layers [78]. Although it has been widely used in the context of semantic segmentation, being able to capture both contextual and local information should be very helpful for nodule detections

as well. Because 3D ConvNet has too many parameters and is hard to train on public lung CT datasets of relatively small sizes, 3D dual path network is employed as the building block since deep dual path network is more compact and provides better performance than deep residual network at the same time [12].

Before the era of deep learning, feature engineering followed by classifiers is a general pipeline for nodule classification [40]. After the public large-scale dataset, LIDC-IDRI [3], becomes available, deep learning based methods have become dominant for nodule classification research [85, 117]. Multi-scale deep ConvNet with shared weights on different scales has been proposed for the nodule classification [86]. The weight sharing scheme reduces the number of parameters and forces the multi-scale deep ConvNet to learn scale-invariant features. Inspired by the recent success of dual path network (DPN) on ImageNet [12, 16], we propose a novel framework for CT nodule classification. First, we design a deep 3D dual path network to extract features. Considering the excellent power of gradient boosting machines (GBM) given effective features, we use GBM with deep 3D dual path features, nodule size and cropped raw nodule CT pixels for the nodule classification [34].

At last, we build a fully automated lung CT cancer diagnosis system, DeepLung, by combining the nodule detection network and nodule classification network together, as illustrated in Fig. 4.1. For a CT image, we first use the detection subnetwork to detect candidate nodules. Next, we employ the classification subnetwork to classify the detected nodules into either malignant or benign. Finally, the patient-level diagnosis result can be achieved for the whole CT by fusing the diagnosis result of each nodule.

Our main contributions are as follows: 1) To fully exploit the 3D CT images, two deep 3D ConvNets are designed for nodule detection and classification respectively. Because 3D ConvNet contains too many parameters and is hard to train on relatively small public lung CT datasets, we employ 3D dual path networks as the components since DPN uses less parameters and obtains better performance than residual network [12]. Specifically, inspired

by the effectiveness of Faster R-CNN for object detection [48], we propose 3D Faster R-CNN for nodule detection based on 3D dual path network and U-net-like encoder-decoder structure, and deep 3D dual path network for nodule classification. 2) Our classification framework achieves better performance compared with state-of-the-art approaches, and the performance surpasses the performance of experienced doctors on the largest public dataset, LIDC-IDRI dataset. 3) The fully automated DeepLung system, nodule classification based on detection, is comparable to the performance of experienced doctors both on nodule-level and patient-level diagnosis.

4.2 Related Work

Traditional nodule detection requires manually designed features or descriptors [66]. Recently, several works have been proposed to use deep ConvNets for nodule detection to automatically learn features, which is proven to be much more effective than hand-crafted features. Setio et al. proposes multi-view ConvNet for false positive nodule reduction [80]. Due to the 3D nature of CT scans, some work propose 3D ConvNets to handle the challenge. The 3D fully ConvNet (FCN) is proposed to generate region candidates, and deep ConvNet with weighted sampling is used in the false positive candidates reduction stage [25]. Ding et al. and Liao et al. use the Faster R-CNN to generate candidate nodules, followed by 3D ConvNets to remove false positive nodules [23, 62]. Due to the effective performance of Faster R-CNN [48, 76], we design a novel network, 3D Faster R-CNN with 3D dual path blocks, for the nodule detection. Further, a U-net-like encoder-decoder scheme is employed for 3D Faster R-CNN to effectively learn the features [78].

Nodule classification has traditionally been based on segmentation [27] and manual feature design [2]. Several works designed 3D contour feature, shape feature and texture feature for CT nodule diagnosis [27, 40, 107]. Recently, deep networks have been shown to be effective

for medical images. Artificial neural network was implemented for CT nodule diagnosis [90]. More computationally effective network, multi-scale ConvNet with shared weights for different scales to learn scale-invariant features, is proposed for nodule classification [86]. Deep transfer learning and multi-instance learning is used for patient-level lung CT diagnosis [85, 119]. A comparison on 2D and 3D ConvNets is conducted and shown that 3D ConvNet is better than 2D ConvNet for 3D CT data [111]. Further, a multi-task learning and transfer learning framework is proposed for nodule diagnosis [49]. Different from their approaches, we propose a novel classification framework for CT nodule diagnosis. Inspired by the recent success of deep dual path network (DPN) on ImageNet [12], we design a novel totally 3D DPN to extract features from raw CT nodules. Due to the superior power of gradient boost machine (GBM) with complete features, we employ GBM with different levels of granularity ranging from raw pixels, DPN features, to global features such as nodule size for the nodule diagnosis. Patient-level diagnosis can be achieved by fusing the nodule-level diagnosis.

4.3 DeepLung Framework

The fully automated lung CT cancer diagnosis system, DeepLung, consists of two parts, nodule detection and classification. We design a 3D Faster R-CNN for nodule detection, and propose GBM with deep 3D DPN features, raw nodule CT pixels and nodule size for nodule classification.

4.3.1 3D Faster R-CNN with Deep 3D Dual Path Net for Nodule Detection

Inspired by the success of dual path network on the ImageNet [12, 16], we design a deep 3D DPN framework for lung CT nodule detection and classification in Fig. 4.3 and Fig.

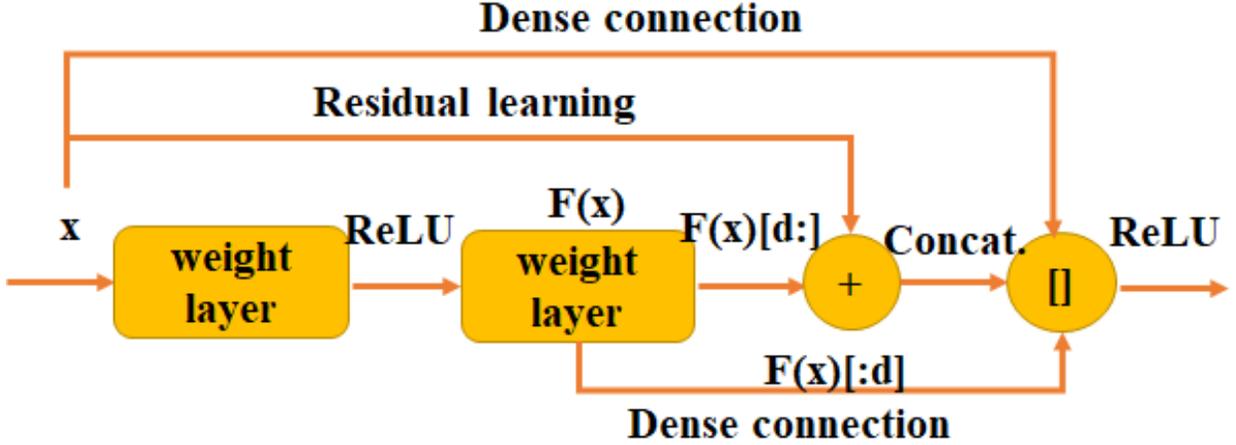


Figure 4.2: Illustration of dual path connection [12], which benefits both from the advantage of residual learning [43] and that of dense connection [47] from network structure design intrinsically.

4.4. Dual path connection benefits both from the advantage of residual learning and that of dense connection [43, 47]. The shortcut connection in residual learning is an effective way to eliminate gradient vanishing phenomenon in very deep networks. From a learned feature sharing perspective, residual learning enables feature reuse, while dense connection has an advantage of keeping exploiting new features [12]. And densely connected network has fewer parameters than residual learning because there is no need to relearn redundant feature maps. The assumption of dual path connection is that there might exist some redundancy in the exploited features. And dual path connection uses part of feature maps for dense connection and part of them for residual learning. In implementation, the dual path connection splits its feature maps into two parts. One part, $\mathbf{F}(\mathbf{x})[d:]$, is used for residual learning, the other part, $\mathbf{F}(\mathbf{x})[:d]$, is used for dense connection as shown in Fig. 4.2. Here d is a hyper-parameter for deciding how many new features to be exploited. The dual path connection can be formulated as

$$\mathbf{y} = \mathbf{G}([\mathbf{x}[:d], \mathbf{F}(\mathbf{x})[:d], \mathbf{F}(\mathbf{x})[d:] + \mathbf{x}[d:]]), \quad (4.1)$$

where \mathbf{y} is the feature map for dual path connection, \mathbf{G} is used as ReLU activation function, \mathbf{F} is convolutional layer functions, and \mathbf{x} is the input of dual path connection block. Dual path connection integrates the advantages of the two advanced frameworks, residual learning for feature reuse and dense connection for keeping exploiting new features, into a unified structure, which obtains success on the ImageNet dataset[16]. We design deep 3D neural nets based on 3D DPN because of its compactness and effectiveness.

The 3D Faster R-CNN with a U-net-like encoder-decoder structure and 3D dual path blocks is illustrated in Fig. 4.3. Due to the GPU memory limitation, the input of 3D Faster R-CNN is cropped from 3D reconstructed CT images with pixel size $96 \times 96 \times 96$. The encoder network is derived from 2D DPN [12]. Before the first max-pooling, two convolutional layers are used to generate features. After that, eight dual path blocks are employed in the encoder subnetwork. We integrate the U-net-like encoder-decoder design concept in the detection to learn the deep nets efficiently [78]. In fact, for the region proposal generation, the 3D Faster R-CNN conducts pixel-wise multi-scale learning and the U-net is validated as an effective way for pixel-wise labeling. This integration makes candidate nodule generation more effective. In the decoder network, the feature maps are processed by deconvolution layers and dual path blocks, and are subsequently concatenated with the corresponding layers in the encoder network [113]. Then a convolutional layer with dropout (dropout probability 0.5) is used for the second last layer. In the last layer, we design 3 anchors, 5, 10, 20, for scale references which are designed based on the distribution of nodule sizes. For each anchor, there are 5 parts in the loss function, classification loss L_{cls} for whether the current box is a nodule or not, regression loss L_{reg} for nodule coordinates x, y, z and nodule size d .

If an anchor overlaps a ground truth bounding box with the intersection over union (IoU) higher than 0.5, we consider it as a positive anchor ($p^* = 1$). On the other hand, if an anchor has IoU with all ground truth boxes less than 0.02, we consider it as a negative

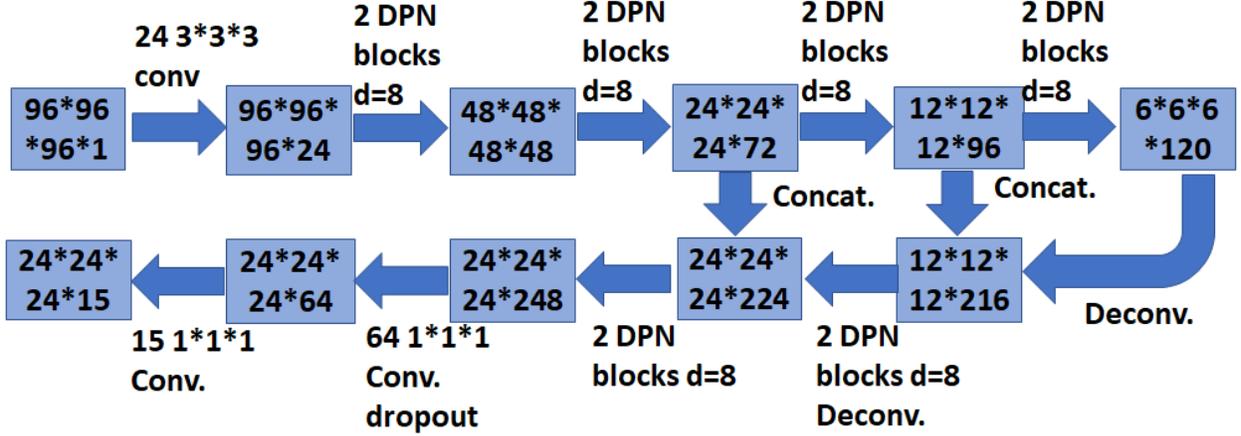


Figure 4.3: The 3D Faster R-CNN framework contains 3D dual path blocks and a U-net-like encoder-decoder structure. We design 26 layers 3D dual path network for the encoder subnetwork. The model employs 3 anchors and multi-task learning loss, including coordinates (x, y, z) and diameter d regression, and candidate box classification. The numbers in boxes are feature map sizes in the format ($\#slices \times \#rows \times \#cols \times \#maps$). The numbers above the connections are in the format ($\#filters \quad \#slices \times \#rows \times \#cols$).

anchor ($p^* = 0$). The multi-task loss function for the anchor i is defined as

$$L(p_i, \mathbf{t}_i) = \lambda L_{cls}(p_i, p_i^*) + p_i^* L_{reg}(\mathbf{t}_i, \mathbf{t}_i^*), \quad (4.2)$$

where p_i is the predicted probability for current anchor i being a nodule, \mathbf{t}_i is the predicted relative coordinates for nodule position, which is defined as

$$\mathbf{t}_i = \left(\frac{x - x_a}{d_a}, \frac{y - y_a}{d_a}, \frac{z - z_a}{d_a}, \log\left(\frac{d}{d_a}\right) \right), \quad (4.3)$$

where (x, y, z, d) are the predicted nodule coordinates and diameter in the original space, (x_a, y_a, z_a, d_a) are the coordinates and scale for the anchor i . For ground truth nodule position, it is defined as

$$\mathbf{t}_i^* = \left(\frac{x^* - x_a}{d_a}, \frac{y^* - y_a}{d_a}, \frac{z^* - z_a}{d_a}, \log\left(\frac{d^*}{d_a}\right) \right), \quad (4.4)$$

where (x^*, y^*, z^*, d^*) are nodule ground truth coordinates and diameter. The λ is set as 0.5.

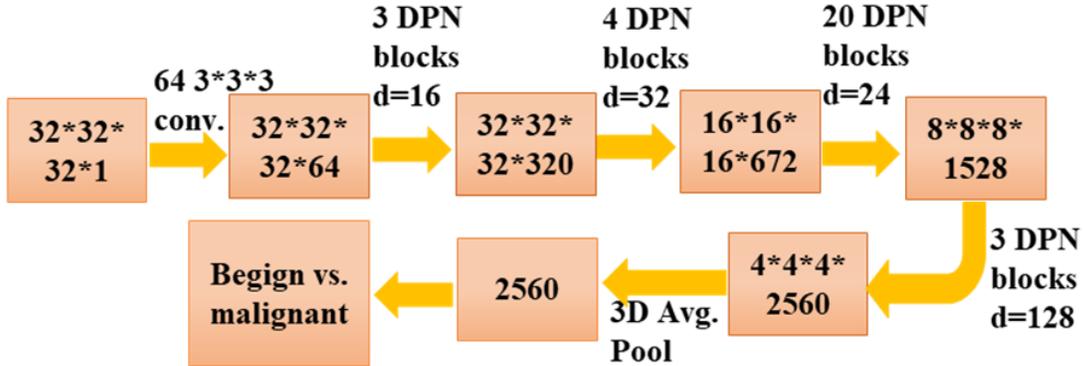


Figure 4.4: The deep 3D dual path network framework in the nodule classification subnetwork, which contains 30 3D dual path connection blocks. After the training, the deep 3D dual path network feature is extracted for gradient boosting machine to do nodule diagnosis. The numbers are of the same formats as Fig. 4.3.

For L_{cls} , we used binary cross entropy loss function. For L_{reg} , we used smooth l_1 regression loss function [38].

4.3.2 Gradient Boosting Machine with 3D Dual Path Net Feature for Nodule Classification

For CT data, advanced method should be effective to extract 3D volume feature [111]. We design a 3D deep dual path network for the 3D CT lung nodule classification in Fig. 4.4. The main reason we employ dual modules for detection and classification is that classifying nodules into benign and malignant requires the system to learn finer-level features, which can be achieved by focusing only on nodules. In addition, it allows to introduce extra features in the final classification. We firstly crop CT data centered at predicted nodule locations with size $32 \times 32 \times 32$. After that, a convolutional layer is used to extract features. Then 30 3D dual path blocks are employed to learn higher level features. Lastly, the 3D average pooling and binary logistic regression layer are used for benign or malignant diagnosis.

The deep 3D dual path network can be used as a classifier for nodule diagnosis directly. And

it can also be employed to learn effective features. We construct feature by concatenating the learned deep 3D DPN features (the second last layer, 2,560 dimension), nodule size, and raw 3D cropped nodule pixels. Given complete and effective features, GBM is an effective method to build an advanced classifier from these features [34]. We validate the feature combining nodule size with raw 3D cropped nodule pixels, employ GBM as a classifier, and obtain 86.12% test accuracy averagely. Lastly, we employ GBM with the constructed feature and achieve the best diagnosis performance.

4.3.3 DeepLung System: Fully Automated Lung CT Cancer Diagnosis

The DeepLung system includes the nodule detection using the 3D Faster R-CNN, and nodule classification using GBM with constructed feature (deep 3D dual path features, nodule size and raw nodule CT pixels) in Fig. 4.1.

Due to the GPU memory limitation, we first split the whole CT into several $96 \times 96 \times 96$ patches, process them through the detector, and combine the detected results together. We only keep the detected boxes of detection probabilities larger than 0.12 (threshold as -2 before sigmoid function). After that, non-maximum suppression (NMS) is adopted based on detection probability with the intersection over union (IoU) threshold as 0.1. Here we expect to not miss too many ground truth nodules.

After we get the detected nodules, we crop the nodule with the center as the detected center and size as $32 \times 32 \times 32$. The detected nodule size is kept for the classification model as a part of features. The deep 3D DPN is employed to extract features. We use the GBM and construct features to conduct diagnosis for the detected nodules. For pixel feature, we use the cropped size as $16 \times 16 \times 16$ and center as the detected nodule center in the experiments. For patient-level diagnosis, if one of the detected nodules is positive (cancer), the patient is a

cancer patient, and if all the detected nodules are negative, the patient is a negative patient.

4.4 Experiments

We conduct extensive experiments to validate the DeepLung system. We perform 10-fold cross validation using the detector on LUNA16 dataset. For nodule classification, we use the LIDC-IDRI annotation, and employ the LUNA16’s patient-level dataset split. Finally, we also validate the whole system based on the detected nodules both on patient-level diagnosis and nodule-level diagnosis.

In the training, for each model, we use 150 epochs in total with stochastic gradient descent optimization and momentum as 0.9. The used batch size is relied on the GPU memory. We use weight decay as 1×10^{-4} . The initial learning rate is 0.01, 0.001 at the half of training, and 0.0001 after the epoch 120.

4.4.1 Datasets

LUNA16 dataset is a subset of the largest public dataset for pulmonary nodules, the LIDC-IDRI dataset [3, 81]. LUNA16 dataset only has the detection annotations, while LIDC-IDRI contains almost all the related information for low-dose lung CTs including several doctors’ annotations on nodule sizes, locations, diagnosis results, nodule texture, nodule margin and other informations. LUNA16 dataset removes CTs with slice thickness greater than 3mm, slice spacing inconsistent or missing slices from LIDC-IDRI dataset, and explicitly gives the patient-level 10-fold cross validation split of the dataset. LUNA16 dataset contains 888 low-dose lung CTs, and LIDC-IDRI contains 1,018 low-dose lung CTs. Note that LUNA16 dataset removes the annotated nodules of size smaller than 3mm.

For nodule classification, we extract nodule annotations from LIDC-IDRI dataset, find the mapping of different doctors’ nodule annotations with the LUNA16’s nodule annotations, and get the ground truth of nodule diagnosis by taking different doctors’ diagnosis averagely (Do not count the 0 score for diagnosis, which means N/A.). If the final average score is equal to 3 (uncertain about malignant or benign), we remove the nodule. For the nodules with score greater than 3, we label them as positive. Otherwise, we label them as negative. Because CT slides were annotated by anonymous doctors, the identities of doctors (referred to as Drs 1-4 as the 1st-4th annotations) are not strictly consistent. As such, we refer them as “simulated” doctors. To make our results reproducible, we only keep the CTs within LUNA16 dataset, and use the same cross validation split as LUNA16 for classification.

4.4.2 Preprocessing

Three automated preprocessing steps are employed for the input CT images. We firstly clip the raw data into $[-1200, 600]$. Secondly, we transform the range linearly into $[0, 1]$. Thirdly, we use the LUNA16’s given segmentation ground truth and remove the useless background.

4.4.3 DeepLung for Nodule Detection

We train and evaluate the detector on LUNA16 dataset following 10-fold cross validation with given patient-level split. In the training, we use the flipping, randomly scale from 0.75 to 1.25 for the cropped patches to augment the data. The evaluation metric, FROC, is the average recall rate at the average number of false positives as 0.125, 0.25, 0.5, 1, 2, 4, 8 per scan, which is the official evaluation metric for LUNA16 dataset [81]. In the test phase, we use detection probability threshold as -2 (before sigmoid function), followed by NMS with IoU threshold as 0.1.

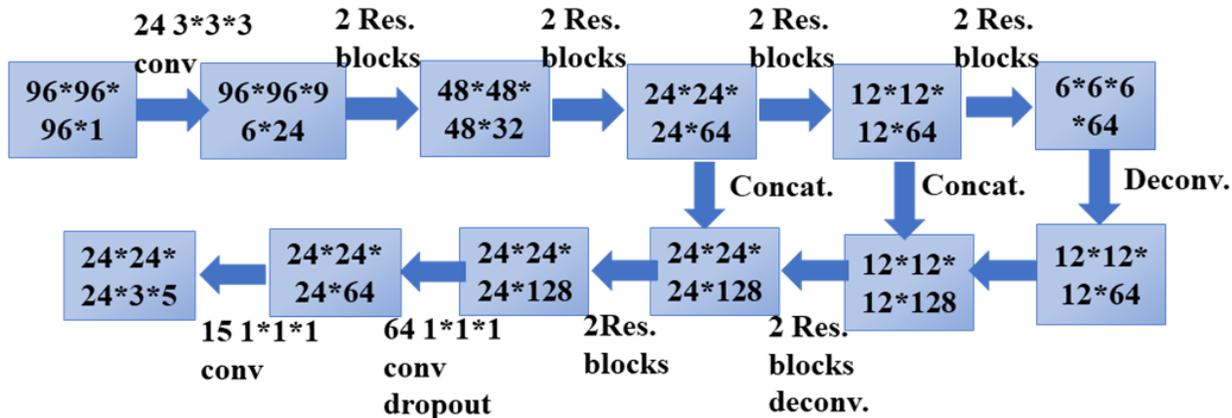


Figure 4.5: The 3D Faster R-CNN network with 3D residual blocks. It contains several 3D residual blocks. We employ a deep 3D residual network of 18 layers as the encoder subnetwork, which is an extension from 2D Res18 net [43].

To validate the superior performance of proposed deep 3D dual path network for detection, we employ a deep 3D residual network as a comparison in Fig. 4.5. The encoder part of compared network is a deep 3D residual network of 18 layers, which is an extension from 2D Res18 net [43]. Note that the 3D Res18 Faster R-CNN contains 5.4M trainable parameters, while the 3D DPN26 Faster R-CNN employs 1.4M trainable parameters, which is only $\frac{1}{4}$ of that in 3D Res18 Faster R-CNN.

The FROC performance on LUNA16 is visualized in Fig. 4.6. The solid line is interpolated FROC based on true prediction. The 3D DPN26 Faster R-CNN achieves **84.2%** FROC without any false positive nodule reduction stage, which is better than 83.9% using two-stage training [25]. The 3D DPN26 Faster R-CNN using only $\frac{1}{4}$ of the parameters performs better than the 3D Res18 Faster R-CNN, which demonstrates the superior suitability of the 3D DPN for detection. Ding et al. obtains 89.1% FROC using 2D Faster R-CNN followed by extra false positive reduction classifier [23], while we only employ enhanced Faster R-CNN with deep 3D dual path for detection. We have recently applied the 3D model to Alibaba Tianchi Medical AI on nodule detection challenge and were able to achieve top accuracy on a hold-out dataset.

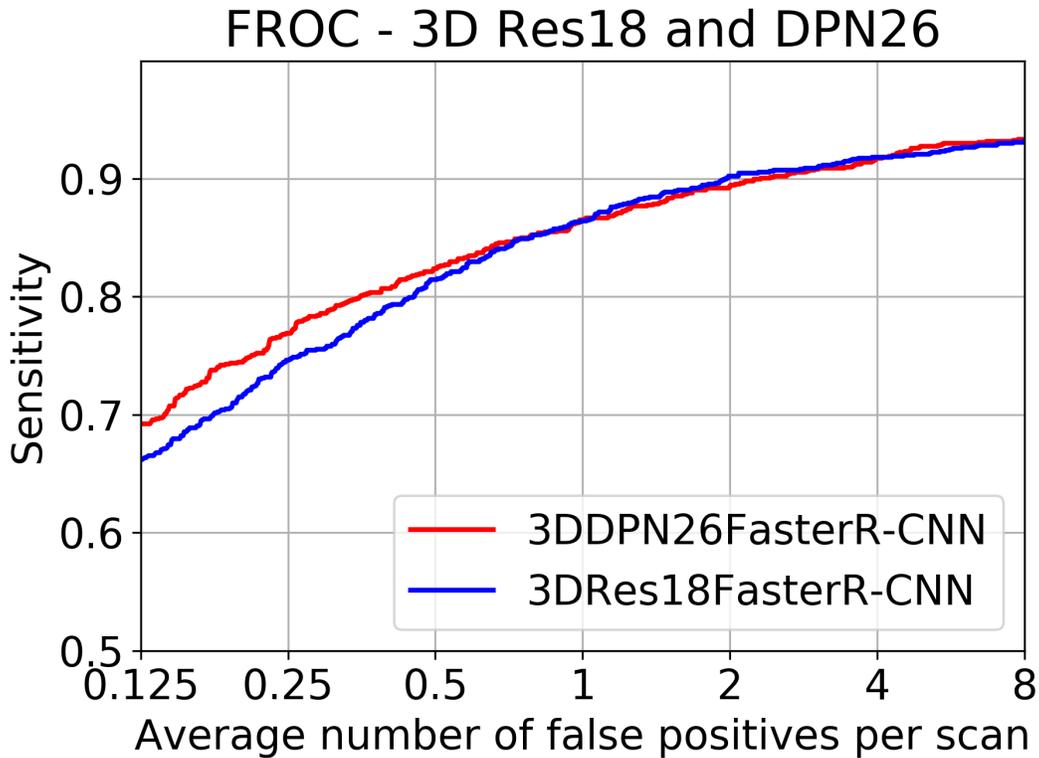


Figure 4.6: Sensitivity (Recall) rate with respect to false positives per scan. The FROC (average recall rate at the false positives as 0.125, 0.25, 0.5, 1, 2, 4, 8) of 3D Res18 Faster R-CNN is 83.4%, while the FROC of 3D DPN26 Faster R-CNN is **84.2%** with only $\frac{1}{4}$ of the parameters as 3D Res18 Faster R-CNN. The 3D Res18 Faster R-CNN has a total recall rate 94.6% for all the detected nodules, while 3D DPN26 Faster R-CNN has a recall rate **95.8%**.

Table 4.1: Nodule classification comparisons on LIDC-IDRI dataset.

Models	Accuracy (%)	Year
Multi-scale CNN [86]	86.84	2015
Slice-level 2D CNN [111]	86.70	2016
Nodule-level 2D CNN [111]	87.30	2016
Vanilla 3D CNN [111]	87.40	2016
Multi-crop CNN [87]	87.14	2017
Deep 3D DPN	88.74	2017
Nodule Size+Pixel+GBM	86.12	2017
All feat.+GBM	90.44	2017

4.4.4 DeepLung for Nodule Classification

We validate the nodule classification performance of the DeepLung system on the LIDC-IDRI dataset with the LUNA16’s split principle, 10-fold patient-level cross validation. There are 1,004 nodules left and 450 nodules are positive. In the training, we firstly pad the nodules of size $32 \times 32 \times 32$ into $36 \times 36 \times 36$, randomly crop $32 \times 32 \times 32$ from the padded data, horizontal flip, vertical flip, z-axis flip the data for augmentation, randomly set $4 \times 4 \times 4$ patch as 0, and normalize the data with the mean and standard deviation obtained from training data. The total number of epochs is 1,050. The learning rate is 0.01 at first, then became 0.001 after epoch 525, and turned into 0.0001 after epoch 840. Due to time and resource limitation for training, we use the fold 1, 2, 3, 4, 5 for test, and the final performance is the average performance on the five test folds. The nodule classification performance is concluded in Table 4.1.

From the table 4.1, our deep 3D dual path network (DPN) achieves better performance than those of Multi-scale CNN [86], Vanilla 3D CNN [111] and Multi-crop CNN [87], because of the strong power of 3D structure and deep dual path network. GBM with nodule size and raw nodule pixels with crop size as $16 \times 16 \times 16$ achieves comparable performance as multi-scale CNN [86] because of the superior classification performance of gradient boosting machine (GBM). Finally, we construct feature with deep 3D dual path network features, 3D

Table 4.2: Nodule-level diagnosis accuracy (%) between nodule classification subnetwork in DeepLung and experienced doctors on doctor’s individually confident nodules.

	Dr 1	Dr 2	Dr 3	Dr 4	Average
Doctors	93.44	93.69	91.82	86.03	91.25
DeepLung	93.55	93.30	93.19	90.89	92.74

Table 4.3: Statistical property of predicted malignant probability for borderline nodules (%)

Prediction	< 0.1 or > 0.9	< 0.2 or > 0.8	< 0.3 or > 0.7	< 0.4 or > 0.6
Frequency	64.98	80.14	89.75	94.80

Faster R-CNN detected nodule size and raw nodule pixels, and obtain **90.44%** accuracy, which shows the effectiveness of deep 3D dual path network features.

Compared with Experienced Doctors on Their Individually Confident Nodules

We compare our predictions with those of four “simulated” experienced doctors on their individually confident nodules (with individual score not 3). Note that about 1/3 annotations are 3. Comparison results are concluded in Table 4.2.

From Table 4.2, these doctors’ confident nodules are easy to be diagnosed nodules from the performance comparison between our model’s performances in Table 4.1 and Table 4.2. To our surprise, the average performance of our model is **1.5%** better than that of experienced doctors even on their individually confident diagnosed nodules. In fact, our model’s performance is better than 3 out of 4 doctors (doctor 1, 3, 4) on the confident nodule diagnosis task. The result validates deep network surpasses human-level performance for image classification [43], and the DeepLung is better suited for nodule diagnosis than experienced doctors.

We also employ Kappa coefficient, which is a common approach to evaluate the agreement between two raters, to test the agreement between DeepLung and the ground truth [60]. The kappa coefficient of DeepLung is 85.07%, which is significantly better than the average

kappa coefficient of doctors (81.58%). To evaluate the performance for all nodules including borderline nodules (labeled as 3, uncertain between malignant and benign), we compute the log likelihood (LL) scores of DeepLung and doctors’ diagnosis. We randomly sample 100 times from the experienced doctors’ annotations as 100 “simulated” doctors. The mean LL of doctors is -2.563 with a standard deviation of 0.23. By contrast, the LL of DeepLung is -1.515, showing that the performance of DeepLung is 4.48 standard deviation better than the average performance of doctors, which is highly statistically significant. It is important to analysis the statistical property of predictions for borderline nodules that cannot be conclusively classified by doctors. Interestingly, 64.98% of the borderline nodules are classified to be either malignant (with probability > 0.9) or benign (with probability < 0.1) in Table 4.3. DeepLung classified most of the borderline nodules of malignant probabilities closer to zero or closer to one, showing its potential as a tool for assisted diagnosis.

4.4.5 DeepLung for Fully Automated Lung CT Cancer Diagnosis

We also validate the DeepLung for fully automated lung CT cancer diagnosis on the LIDC-IDRI dataset with the same protocol as LUNA16’s patient-level split. Firstly, we employ our 3D Faster R-CNN to detect suspicious nodules. Then we retrain the model from nodule classification model on the detected nodules dataset. If the center of detected nodule is within the ground truth positive nodule, it is a positive nodule. Otherwise, it is a negative nodule. Through this mapping from the detected nodule and ground truth nodule, we can evaluate the performance and compare it with the performance of experienced doctors. We adopt the test fold 1, 2, 3, 4, 5 to validate the performance the same as that for nodule classification.

Different from pure nodule classification, the fully automated lung CT nodule diagnosis relies on nodule detection. We evaluate the performance of DeepLung on the detection true

Table 4.4: Comparison between DeepLung’s nodule classification on all detected nodules and doctors on all nodules.

Method	TP Set	FP Set	Doctors
Acc. (%)	81.42	97.02	74.05-82.67

Table 4.5: Patient-level diagnosis accuracy(%) between DeepLung and experienced doctors on doctor’s individually confident CTs.

	Dr 1	Dr 2	Dr 3	Dr 4	Average
Doctors	83.03	85.65	82.75	77.80	82.31
DeepLung	81.82	80.69	78.86	84.28	81.41

positive (TP) set and detection false positive (FP) set individually in Table 4.4. If the detected nodule of center within one of ground truth nodule regions, it is in the TP set. If the detected nodule of center out of any ground truth nodule regions, it is in FP set. From Table 4.4, the DeepLung system using detected nodule region obtains **81.42%** accuracy for all the detected TP nodules. Note that the experienced doctors obtain 78.36% accuracy for all the nodule diagnosis on average. The DeepLung system with fully automated lung CT nodule diagnosis still achieves above average performance of experienced doctors. On the FP set, our nodule classification subnetwork in the DeepLung can reduce 97.02% FP detected nodules, which guarantees that our fully automated system is effective for the lung CT cancer diagnosis.

Compared with Experienced Doctors on Their Individually Confident CTs

We employ the DeepLung for patient-level diagnosis further. If the current CT has one nodule that is classified as positive, the diagnosis of the CT is positive. If all the nodules are classified as negative for the CT, the diagnosis of the CT is negative. We evaluate the DeepLung on the doctors’ individually confident CTs for benchmark comparison in Table 4.5.

From Table 4.5, DeepLung achieves **81.41%** patient-level diagnosis accuracy. The perfor-

mance is **99%** of the average performance of four experienced doctors, and the performance of DeepLung is better than that of doctor 4. Thus DeepLung can be used to help improve some doctors' performance, like that of doctor 4, which is the goal for computer aided diagnosis system. For comparison, we valid those of 4 individual doctors on their individual confident CTs. The Kappa coefficient of DeepLung is 63.02%, while the average Kappa coefficient of doctors is 64.46%. It shows the predictions of DeepLung are of good agreement with ground truths for patient-level diagnosis, and are comparable with those of experienced doctors.

4.5 Discussion

In this section, we are trying to explain the DeepLung by visualizing the nodule detection and classification results.

4.5.1 Nodule Detection

We randomly pick nodules from test fold 1 and visualize them in red circles of the first row in Fig. 4.7. Detected nodules are visualized in blue circles of the second row. Because CT is 3D voxel data, we can only plot the central slice for visualization. The third row shows the detection probabilities for the detected nodules. The central slice number is shown below each slice. The diameter of the circle is relative to the nodule size.

From the central slice visualizations in Fig. 4.7, we observe the detected nodule positions including central slice numbers are consistent with those of ground truth nodules. The circle sizes are similar between the nodules in the first row and the second row. The detection probability is also very high for these nodules in the third row. It shows 3D Faster R-CNN works well to detect the nodules from test fold 1.

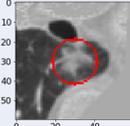
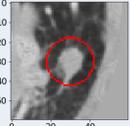
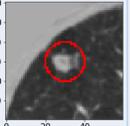
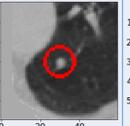
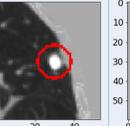
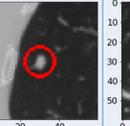
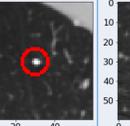
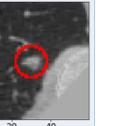
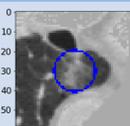
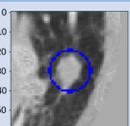
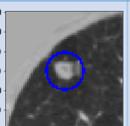
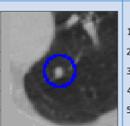
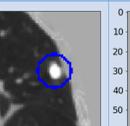
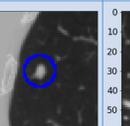
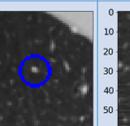
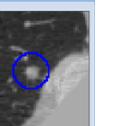
Ground truth center slice	 Slice 120	 Slice 150	 Slice 127	 Slice 191	 Slice 86	 Slice 108	 Slice 171	 Slice 271
Detection center slice	 Slice 118	 Slice 149	 Slice 127	 Slice 191	 Slice 85	 Slice 108	 Slice 170	 Slice 272
Detection prob. (%)	99.46	99.93	100.00	99.86	99.95	99.90	99.61	99.52

Figure 4.7: Visualization of central slices for nodule ground truths and detection results. We randomly choose nodules (red circle boxes in the first row) from test fold 1. Detection results are shown in the blue circles of second row. The center slice numbers are shown below the images. The last row shows detection probability. The DeepLung performs well for nodule detection.

4.5.2 Nodule Classification

We also visualize the nodule classification results from test fold 1 in Fig. 4.8. We choose nodules that is predicted right, but annotated wrong by some doctors. The first seven nodules are benign nodules, and the rest nodules are malignant nodules. The numbers below the figures are the DeepLung predicted malignant probabilities, followed by which annotation of doctors is wrong. For the DeepLung, if the probability is large than 0.5, it predicts malignant. Otherwise, it predicts benign. For an experienced doctor, if a nodule is big and has irregular shape, it has a high probability to be a malignant nodule.

From Fig. 4.8, we can observe that doctors mis-diagnose some nodules. The reason is that, humans are not good at processing 3D CT data, which is of low signal to noise ratio. Maybe the doctor cannot find some weak irregular boundaries or consider some tissues as nodule boundaries, which is the possible reason why there are false negatives or false positives for doctors' annotations. In fact, even for high quality 2D natural image, the performance of deep network surpasses that of humans [43]. From the doctors' individual experience point

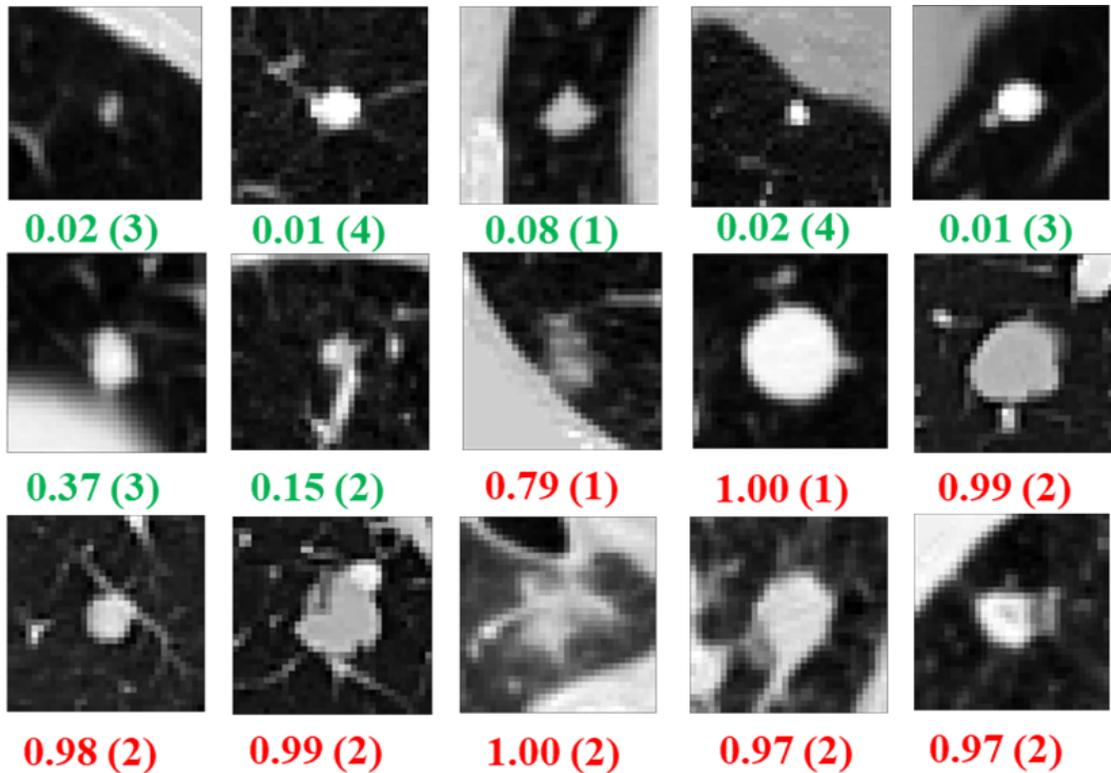


Figure 4.8: Visualization of central slices for nodule classification results on test fold 1. We choose nodules that are predicted right by the DeepLung, but annotated wrong by some doctors. The numbers below the nodules are model predicted malignant probabilities followed by which annotation of doctors is wrong. The first seven nodules are benign nodules. The rest nodules are malignant nodules. The DeepLung performs well for nodule classification.

of view, some doctors intend to be optimistic to some nodules, and some doctors intend to be pessimistic to them. And they can just observe one slice each time. Some irregular boundaries are vague. The machine learning based methods can learn these complicated rules and high dimensional features from these doctors' annotations, and avoid such bias. From the above analysis, the DeepLung can be considered as a tool to assist the diagnosis for doctors. Combing the DeepLung and doctor's own diagnosis is an effective way to improve diagnosis accuracy.

4.6 Conclusion

In this work, we propose a fully automated lung CT cancer diagnosis system, DeepLung, based on deep learning. DeepLung consists of two parts, nodule detection and classification. To fully exploit 3D CT images, we propose two deep 3D convolutional networks based on 3D dual path networks, which is more compact and can yield better performance than residual networks. For nodule detection, we design a 3D Faster R-CNN with 3D dual path blocks and a U-net-like encoder-decoder structure to detect candidate nodules. The detected nodules are subsequently fed to nodule classification network. We use a deep 3D dual path network to extract classification features. Finally, gradient boosting machine with combined features are trained to classify candidate nodules into benign or malignant. Extensive experimental results on public available large-scale datasets, LUNA16 and LIDC-IDRI datasets, demonstrate the superior performance of the DeepLung system.

Chapter 5

DeepEM: Deep 3D ConvNets With EM For Weakly Supervised Pulmonary Nodule Detection

5.1 Introduction

Lung cancer is the most common cause of cancer-related death in men. Low-dose lung computed tomography (CT) screening provides an effective way for early diagnosis and can sharply reduce the lung cancer mortality rate. Advanced computer-aided diagnosis (CAD) systems are expected to have high sensitivities while maintaining low false positive rates to be truly useful. Recent advance in deep learning provides new opportunities to design more effective CAD systems to help facilitate doctors in their effort to catch lung cancer in their early stages.

The emergence of large-scale datasets such as the LUNA16 [82] has helped to accelerate research in nodule detection. Typically, nodule detection consists of two stages: nodule

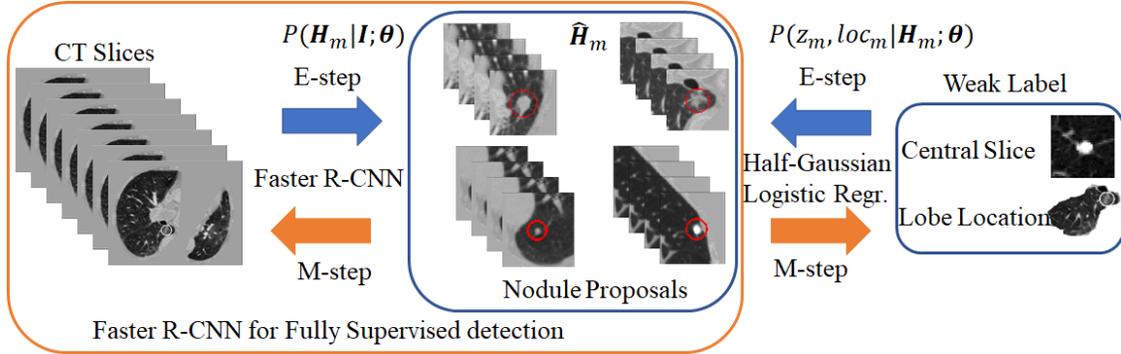


Figure 5.1: Illustration of DeepEM framework. Faster R-CNN is employed for nodule proposal generation. Half-Gaussian model and logistic regression are employed for central slice and lobe location respectively. In the E-step, we utilize all the observations, CT slices, and weak label to infer the latent variable, nodule proposals, by maximum a posteriori (MAP) or sampling. In the M-step, we employ the estimated proposals to update parameters in the Faster R-CNN and logistic regression.

proposal generation and false positive reduction. Traditional approaches generally require hand-designed features such as morphological features, voxel clustering and pixel thresholding [53, 66, 69]. More recently, deep convolutional architectures were employed to generate the candidate bounding boxes. Setio et al. proposed multi-view convolutional network for false positive nodule reduction [80]. Several work employed 3D convolutional networks to handle the challenge due to the 3D nature of CT scans. The 3D fully convolutional network (FCN) was proposed to generate region candidates and deep convolutional network with weighted sampling was used in the false positive reduction stage [25, 62, 93, 118]. CASED proposed curriculum adaptive sampling for 3D U-net training in nodule detection [54, 78]. Ding et al. used Faster R-CNN to generate candidate nodules, followed by 3D convolutional networks to remove false positive nodules [23]. Due to the effective performance of Faster R-CNN [76], Faster R-CNN with a U-net-like encoder-decoder scheme was proposed for nodule detection [118].

A prerequisite to utilization of deep learning models is the existence of an abundance of labeled data. However, labels are especially difficult to obtain in the medical image analysis domain. There are multiple contributing factors: a) labeling medical data typically requires

specially trained doctors; b) marking lesion boundaries can be hard even for experts because of low signal-to-noise ratio in many medical images; and c) for CT and magnetic resonance imaging (MRI) images, the annotators need to label the entire 3D volumetric data, which can be costly and time-consuming. Due to these limitations, CT medical image datasets are usually small, which can lead to over-fitting on the training set and, by extension, poor generalization performance on test sets [121].

By contrast, medical institutions have large amount of weakly labeled medical images. In these databases, each medical image is typically associated with an electronic medical report (EMR). Although these reports may not contain explicit information on detection bounding box or segmentation ground truth, it often includes the results of diagnosis, rough locations and summary descriptions of lesions if they exist. We hypothesize that these extra sources of weakly labeled data may be used to enhance the performance of existing detector and improve its generalization capability.

There are previous attempts to utilize weakly supervised labels to help train machine learning models. Deep multi-instance learning was proposed for lesion localization and whole mammogram classification [119]. The two-stream spatio-temporal ConvNet was proposed to recognize heart frames and localize the heart using only weak labels for whole ultrasound image of fetal heartbeat [37]. Different pooling strategies were proposed for weakly supervised localization and segmentation respectively [7, 31, 102]. Papandreou et al. proposed an iterative approach to infer pixel-wise label using image classification label for segmentation [72]. Self-transfer learning co-optimized both classification and localization networks for weakly supervised lesion localization [50]. Different from these works, we consider nodule proposal as latent variable and propose DeepEM, a new deep 3D convolutional nets with Expectation-Maximization optimization, to mine the big data source of weakly supervised label in EMR as illustrated in Fig. 5.1. Specifically, we infer the posterior probabilities of the proposed nodules being true nodules, and utilize the posterior probabilities to train nodule

detection models.

5.2 DeepEM for Weakly Supervised Detection

Notation We denote by $\mathbf{I} \in \mathbb{R}^{h \times w \times s}$ the CT image, where h , w , and s are image height, width, and number of slices respectively. The nodule bounding boxes for \mathbf{I} are denoted as $\mathbf{H} = \{\mathbf{H}_1, \mathbf{H}_2, \dots, \mathbf{H}_M\}$, where $\mathbf{H}_m = \{x_m, y_m, z_m, d_m\}$, the (x_m, y_m, z_m) represents the center of nodule proposal, d_m is the diameter of the nodule proposal, and M is the number of nodules in the image \mathbf{I} . In the weakly supervised scenario, the nodule proposal \mathbf{H} is a latent variable, and each image \mathbf{I} is associated with weak label $\mathbf{X} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_M\}$, where $\mathbf{X}_m = \{loc_m, z_m\}$, $loc_m \in \{1, 2, 3, 4, 5, 6\}$ is the location (right upper lobe, right middle lobe, right lower lobe, left upper lobe, lingula, left lower lobe) of nodule \mathbf{H}_m in the lung, and z_m is the central slice of the nodule.

For fully supervised detection, the objective function is to maximize the log-likelihood function for observed nodule ground truth \mathbf{H} given image \mathbf{I} as

$$\mathcal{L}(\boldsymbol{\theta}) = \log P(\mathbf{H} \cup \bar{\mathbf{H}} | \mathbf{I}; \boldsymbol{\theta}) = \frac{1}{M} \sum_{m=1}^M \log P(\mathbf{H}_m | \mathbf{I}; \boldsymbol{\theta}) + \frac{1}{N} \sum_{n=1}^N \log P(\bar{\mathbf{H}}_n | \mathbf{I}; \boldsymbol{\theta}), \quad (5.1)$$

where $\bar{\mathbf{H}} = \{\bar{\mathbf{H}}_1, \bar{\mathbf{H}}_2, \dots, \bar{\mathbf{H}}_N\}$ are hard negative nodule proposals [76], $\boldsymbol{\theta}$ is the weights of deep 3D ConvNet. We employ Faster R-CNN with 3D Res18 for the fully supervised detection because of its superior performance.

For weakly supervised detection, nodule proposal \mathbf{H} can be considered as a latent variable. Using this framework, image \mathbf{I} and weak label $\mathbf{X} = \{(loc_1, z_1), (loc_2,$

$z_2), \dots, (loc_M, z_M)\}$ can be considered as observations. The joint distribution is

$$\begin{aligned} P(\mathbf{I}, \mathbf{H}, \mathbf{X}; \boldsymbol{\theta}) &= P(\mathbf{I}) \prod_{m=1}^M (P(\mathbf{H}_m | \mathbf{I}; \boldsymbol{\theta}) P(\mathbf{X}_m | \mathbf{H}_m; \boldsymbol{\theta})) \\ &= P(\mathbf{I}) \prod_{m=1}^M (P(\mathbf{H}_m | \mathbf{I}; \boldsymbol{\theta}) P(loc_m | \mathbf{H}_m; \boldsymbol{\theta}) P(z_m | \mathbf{H}_m; \boldsymbol{\theta})). \end{aligned} \quad (5.2)$$

To model $P(z_m | \mathbf{H}_m; \boldsymbol{\theta})$, we propose using a half-Gaussian distribution based on nodule size distribution because z_m is correct if it is within the nodule area (center slice of \mathbf{H}_m as $z_{\mathbf{H}_m}$, and nodule size σ can be empirically estimated based on existing data) for nodule detection in Fig. 5.2(a). For lung lobe prediction $P(loc_m | \mathbf{H}_m; \boldsymbol{\theta})$, a logistic regression model is used based on relative value of nodule center $(x_{\mathbf{H}_m}, y_{\mathbf{H}_m}, z_{\mathbf{H}_m})$ after lung segmentation. That is

$$P(z_m, loc_m | \mathbf{H}_m; \boldsymbol{\theta}) = \frac{2}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{|z_m - z_{\mathbf{H}_m}|^2}{2\sigma^2}\right) \frac{\exp(\mathbf{f}(\mathbf{H}_m)\boldsymbol{\theta}_{loc_m})}{\sum_{loc_m=1}^6 \exp(\mathbf{f}(\mathbf{H}_m)\boldsymbol{\theta}_{loc_m})}, \quad (5.3)$$

where $\boldsymbol{\theta}_{loc_m}$ is the associated weights with lobe location loc_m for logistic regression, feature $\mathbf{f}(\mathbf{H}_m) = (\frac{x_{\mathbf{H}_m}}{x_{\mathbf{I}}}, \frac{y_{\mathbf{H}_m}}{y_{\mathbf{I}}}, \frac{z_{\mathbf{H}_m}}{z_{\mathbf{I}}})$, and $(x_{\mathbf{I}}, y_{\mathbf{I}}, z_{\mathbf{I}})$ is the total size of image \mathbf{I} after lung segmentation. In the experiments, we found the logistic regression converges quickly and is stable.

The expectation-maximization (EM) is a commonly used approach to optimize the maximum log-likelihood function when there are latent variables in the model. We employ the EM algorithm to optimize deep weakly supervised detection model in equation 5.2. The expected complete-data log-likelihood function given previous estimated parameter $\boldsymbol{\theta}'$ in deep 3D Faster R-CNN is

$$\begin{aligned} Q(\boldsymbol{\theta}; \boldsymbol{\theta}') &= \frac{1}{M} \sum_{m=1}^M \mathbb{E}_{P(\mathbf{H}_m | \mathbf{I}, z_m, loc_m; \boldsymbol{\theta}')} [\log P(\mathbf{H}_m | \mathbf{I}; \boldsymbol{\theta}) \\ &\quad + \log P(z_m, loc_m | \mathbf{H}_m; \boldsymbol{\theta})] + \mathbb{E}_{Q(\bar{\mathbf{H}}_n | \mathbf{z})} [\log P(\bar{\mathbf{H}}_n | \mathbf{I}; \boldsymbol{\theta})], \end{aligned} \quad (5.4)$$

where $\mathbf{z} = \{z_1, z_2, \dots, z_m\}$. In the implementation, we only keep hard negative proposals far away from weak annotation \mathbf{z} to simplify $Q(\bar{\mathbf{H}}_n | \mathbf{z})$. The posterior distribution of latent

variable \mathbf{H}_m can be calculated by

$$P(\mathbf{H}_m|\mathbf{I}, z_m, loc_m; \boldsymbol{\theta}') \propto P(\mathbf{H}_m|\mathbf{I}; \boldsymbol{\theta}')P(z_m, loc_m|\mathbf{H}_m; \boldsymbol{\theta}'). \quad (5.5)$$

Because Faster R-CNN yields a large number of proposals, we first use hard threshold (-3 before sigmoid function) to remove proposals of small confident probability, then employ non-maximum suppression (NMS) with intersection over union (IoU) as 0.1. We then employ two schemes to approximately infer the latent variable \mathbf{H}_m : maximum a posteriori (MAP) or sampling.

DeepEM with MAP We only use the proposal of maximal posterior probability to calculate the expectation.

$$\hat{\mathbf{H}}_m = \arg \max_{\mathbf{H}_m} P(\mathbf{H}_m|\mathbf{I}; \boldsymbol{\theta}')P(z_m, loc_m|\mathbf{H}_m; \boldsymbol{\theta}') \quad (5.6)$$

DeepEM with Sampling We approximate the distribution by sampling \hat{M} proposals $\hat{\mathbf{H}}_m$ according to normalized equation 5.5. The expected log-likelihood function in equation 5.4 becomes

$$Q(\boldsymbol{\theta}; \boldsymbol{\theta}') = \frac{1}{M\hat{M}} \sum_{m=1}^M \sum_{\hat{\mathbf{H}}_m}^{\hat{M}} (\log P(\hat{\mathbf{H}}_m|\mathbf{I}; \boldsymbol{\theta}) + \log P(z_m, loc_m|\hat{\mathbf{H}}_m; \boldsymbol{\theta})) \quad (5.7)$$

$$+ \mathbb{E}_{Q(\bar{\mathbf{H}}_n|z)} [\log P(\bar{\mathbf{H}}_n|\mathbf{I}; \boldsymbol{\theta})].$$

After obtaining the expectation of complete-data log-likelihood function in equation 5.4, we can update the parameters $\boldsymbol{\theta}$ by

$$\hat{\boldsymbol{\theta}} = \arg \max Q(\boldsymbol{\theta}; \boldsymbol{\theta}'). \quad (5.8)$$

The M-step in equation 5.8 can be conducted by stochastic gradient descent commonly used in deep network optimization for equation 5.1. Our entire algorithm is outlined in algorithm

Algorithm 1 DeepEM for Weakly Supervised Detection

- Input:** Fully supervised dataset $D_F = \{(\mathbf{I}, \mathbf{H})_i\}_{i=1}^{N_F}$, weakly supervised dataset $D_W = \{(\mathbf{I}, \mathbf{X})_i\}_{i=1}^{N_W}$, 3D Faster R-CNN and logistic regression parameters θ .
- 1: *Initialization:* Update weights θ by maximizing equation 5.1 using data from D_F .
 - 2: *for epoch = 1 to #TotalEpochs:*
 - ▷▷▷ **Weakly supervised training**
 - 3: Use Faster R-CNN model θ' to obtain proposal probability $P(\mathbf{H}_m|\mathbf{I}; \theta')$ for weakly supervised data sampled from D_W .
 - 4: Remove proposals with small probabilities and NMS.
 - 5: *for m = 1 to M:* ▷▷▷ Each weak label
 - 6: Calculate $P(z_m, loc_m|\mathbf{H}_m; \theta)$ for each proposal by equation 5.3.
 - 7: Estimate posterior distribution $P(\mathbf{H}_m|\mathbf{I}, z_m, loc_m; \theta')$ by equation 5.5 with normalization.
 - 8: Employ MAP by equation 5.6 or Sampling to obtain the inference of \mathbf{H}_m .
 - 9: Obtain the expect log-likelihood function by equation 5.4 using the estimated proposal (MAP) or by equation 5.7 (Sampling).
 - 10: Update parameter by equation 5.8.
 - ▷▷▷ **Fully supervised training**
 - 11: Update weights θ by maximizing equation 5.1 using fully supervised data D_F .
-

1.

5.3 Experiments

We used 3 datasets, LUNA16 dataset [82] as fully supervised nodule detection, NCI NLST¹ dataset as weakly supervised detection, Tianchi Lung Nodule Detection² dataset as holdout dataset for test only. LUNA16 dataset is the largest publicly available dataset for pulmonary nodules detection [82]. LUNA16 dataset removes CTs with slice thickness greater than 3mm, slice spacing inconsistent or missing slices, and consist of 888 low-dose lung CTs which have explicit patient-level 10-fold cross validation split. NLST dataset consists of hundreds of thousands of lung CT images associated with electronic medical records (EMR). In this work, we focus on nodule detection based on image modality and only use the central slice

¹<https://biometry.nci.nih.gov/cdas/datasets/nlst/>

²<https://tianchi.aliyun.com/>

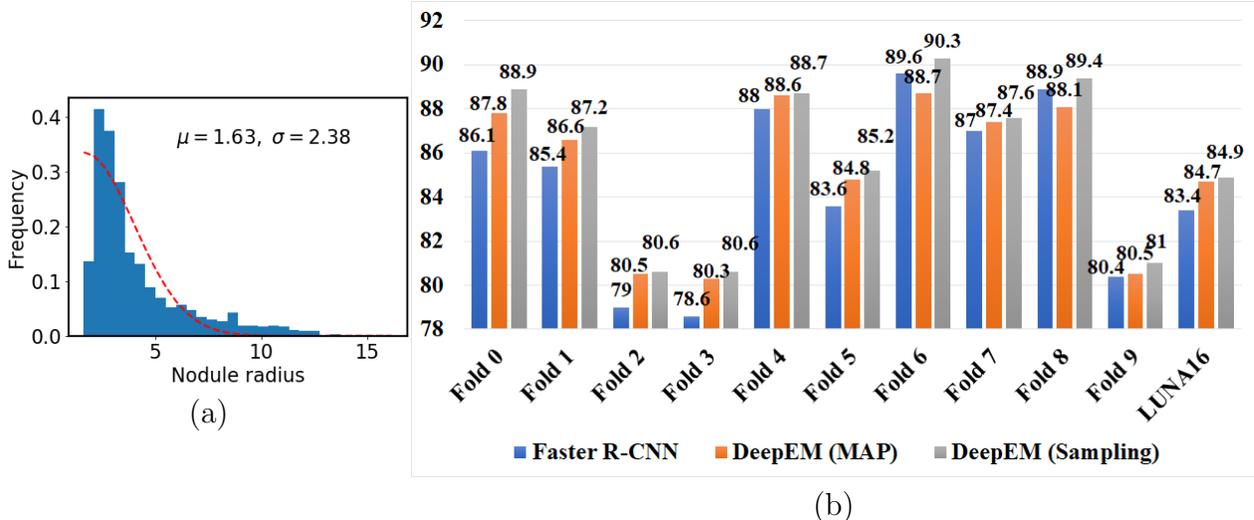


Figure 5.2: (a) Empirical estimation of half-Gaussian model for $P(z_m | \mathbf{H}_m; \theta)$ on LUNA16. (b) FROC (%) comparison among Faster R-CNN, DeepEM with MAP, DeepEM with Sampling on LUNA16.

and nodule location as weak supervision from the EMR. As part of data cleansing, we remove negative CTs, CTs with slice thickness greater than 3mm and nodule diameter less than 3mm. After data cleaning, we have 17,602 CTs left with 30,951 weak annotations. In each epoch, we randomly sample $\frac{1}{16}$ CT images for weakly supervised training because of the large numbers of weakly supervised CTs. Tianchi dataset contains 600 training low-dose lung CTs and 200 validation low-dose lung CTs for nodule detection. The annotations are location centroids and diameters of the pulmonary nodules, and do not have less than 3mm diameter nodule, which are the same with those on LUNA16 dataset.

Parameter estimation in $P(z_m | \mathbf{H}_m; \theta)$ If the current z_m is within the nodule, it is a true positive proposal. We can model $|z_m - z_{\mathbf{H}_m}|$ using a half-Gaussian distribution shown as the red dash line in Fig. 5.2(a). The parameters of the half-Gaussian is estimated from the LUNA16 data empirically. Because LUNA16 removes nodules of diameter less than 3mm, we use the truncated half-Gaussian to model the central slice z_m as $\max(|z_m - z_{\mathbf{H}_m}| - \mu, 0)$, where μ is the mean of related Gaussian as the minimal nodule radius with 1.63.

Performance comparisons on LUNA16 We conduct 10-fold cross validation on LUNA16 to validate the effectiveness of DeepEM. The baseline method is Faster R-CNN with 3D Res18 network denoted as **Faster R-CNN** [76, 118]. Then we employ it to model $P(\mathbf{H}_m|\mathbf{I}; \theta')$ for weakly supervised detection scenario. Two inference scheme for \mathbf{H}_m are used in DeepEM denoted as **DeepEM (MAP)** and **DeepEM (Sampling)**. In the proposal inference of DeepEM with Sampling, we sample two proposals for each weak label because the average number of nodules each CT is 1.78 on LUNA16. The evaluation metric, Free receiver operating characteristic (FROC), is the average recall rate at the average number of false positives at 0.125, 0.25, 0.5, 1, 2, 4, 8 per scan, which is the official evaluation metric for LUNA16 and Tianchi [82].

From Fig. 5.2(b), DeepEM with MAP improves about 1.3% FROC over Faster R-CNN and DeepEM with Sampling improves about 1.5% FROC over Faster R-CNN on average on LUNA16 when incorporating weakly labeled data from NLST. We hypothesize the greater improvement of DeepEM with Sampling over DeepEM with MAP is that MAP inference is greedy and can get stuck at a local minimum while the nature of sampling may allow DeepEM with Sampling to escape these local minimums during optimization.

Performance comparisons on holdout test set from Tianchi We employed a holdout test set from Tianchi to validate each model from 10-fold cross validation on LUNA16. The results are summarized in Table 5.1. We can see DeepEM utilizing weakly supervised data improves 3.9% FROC on average over Faster R-CNN. The improvement on holdout test data validates DeepEM as an effective model to exploit potentially large amount of weak data from electronic medical records (EMR) which would not require further costly annotation by expert doctors and can be easily obtained from hospital associations.

Visualizations We compare Faster R-CNN with the proposed DeepEM visually in Fig. 5.2(b). We randomly choose nodules from Tianchi. From Fig. 5.2(b), DeepEM yields better detection for nodule center and tighter nodule diameter which demonstrates DeepEM

Table 5.1: FROC (%) comparisons among Faster R-CNN with 3D ResNet18, DeepEM with MAP, DeepEM with Sampling on Tianchi.

Fold	0	1	2	3	4	5	6	7	8	9	Average
Faster R-CNN	72.8	70.8	69.8	71.9	76.4	73.0	71.3	74.7	72.9	71.3	72.5
DeepEM (MAP)	77.2	75.8	75.8	74.9	77.0	75.5	77.2	75.8	76.0	74.7	76.0
DeepEM (Sampling)	77.4	75.8	75.9	75.0	77.3	75.0	77.3	76.8	77.7	75.8	76.4

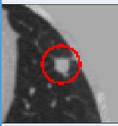
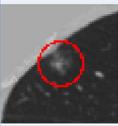
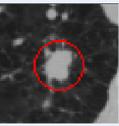
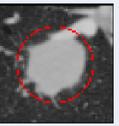
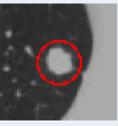
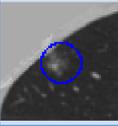
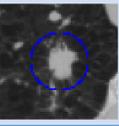
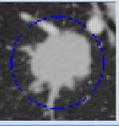
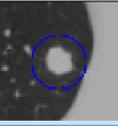
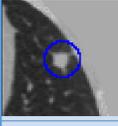
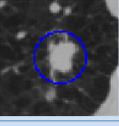
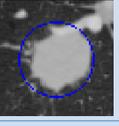
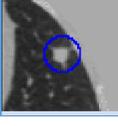
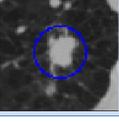
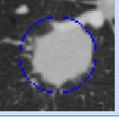
Ground truth center slice		slice 83 center (48,240) diam. 9		slice 114 center (40,88) diam. 7		slice 129 center (129,99) diam. 14		slice 114 center (137,75) diam. 29		slice 47 center (122,202) diam. 12
Faster R-CNN center slice		slice 84 center (48,242) diam. 10		slice 113 center (41,89) diam. 9		slice 126 center (127,98) diam. 19		slice 111 center (138,74) diam. 36		slice 47 center (122,202) diam. 17
DeepEM (MAP) center slice		slice 84 center (48,241) diam. 8		slice 113 center (41,88) diam. 5		slice 128 center (129,98) diam. 16		slice 113 center (137,73) diam. 28		slice 48 center (122,202) diam. 15
DeepEM (sampling) center slice		slice 84 center (48,241) diam. 8		slice 113 center (41,88) diam. 5		slice 128 center (129,98) diam. 16		slice 113 center (137,73) diam. 29		slice 46 center (121,201) diam. 13

Figure 5.3: Detection visual comparison among Faster R-CNN, DeepEM with MAP and DeepEM with Sampling on nodules randomly sampled from Tianchi. DeepEM provides more accurate detection (central slice, center and diameter) than Faster R-CNN.

improves the existing detector by exploiting weakly supervised data.

5.4 Conclusion

In this paper, we have focused on the problem of detecting pulmonary nodules from lung CT images, which previously has been formulated as a supervised learning problem and requires a large amount of training data with the locations and sizes of nodules precisely labeled. Here we propose a new framework, called DeepEM, for pulmonary nodule detection by taking advantage of abundantly available weakly labeled data extracted from EMRs. We treat each nodule proposal as a latent variable, and infer the posterior probabilities of proposal nodules

being true ones conditioned on images and weak labels. The posterior probabilities are further fed to the nodule detection module for training. We use an EM algorithm to train the entire model end-to-end. Two schemes, maximum a posteriori (MAP) and sampling, are used for the inference of proposals. Extensive experimental results demonstrate the effectiveness of DeepEM for improving current state of the art nodule detection systems by utilizing readily available weakly supervised detection data. Although our method is built upon the specific application of pulmonary nodule detection, the framework itself is fairly general and can be readily applied to other medical image deep learning applications to take advantage of weakly labeled data.

Chapter 6

AnatomyNet: Deep Learning for Fast and Fully Automated Whole-volume Segmentation of Head and Neck Anatomy

6.1 Introduction

Head and neck cancer is one of the most common cancers around the world [96]. Radiation therapy is the primary method for treating patients with head and neck cancers. The planning of the radiation therapy relies on accurate organs-at-risks (OARs) segmentation [41], which is usually undertaken by radiation therapists with laborious manual delineation. Computational tools that automatically segment the anatomical regions can greatly alleviate doctors' manual efforts if these tools can delineate anatomical regions accurately with a reasonable amount of time [83].

There is a vast body of literature on automatically segmenting anatomical structures from CT or MRI images. Here we focus on reviewing literature related to head and neck (HaN) CT anatomy segmentation. Traditional anatomical segmentation methods use primarily atlas-based methods, producing segmentations by aligning new images to a fixed set of manually labelled exemplars [75]. Atlas-based segmentation methods typically undergo a few steps, including preprocessing, atlas creation, image registration, and label fusion. As a consequence, their performances can be affected by various factors involved in each of these steps, such as methods for creating atlas [14, 33, 36, 41, 52, 88, 99, 100, 101], methods for label fusion [26, 26, 32], and methods for registration [11, 26, 33, 36, 41, 61, 74, 101, 114]. Although atlas-based methods are still very popular and by far the most widely used methods in anatomy segmentation, their main limitation is the difficulty to handle anatomy variations among patients because they use a fixed set of atlas. In addition, it is computationally intensive and can take many minutes to complete one registration task even with most efficient implementations [110].

Instead of aligning images to a fixed set of exemplars, learning-based methods trained to directly segment OARs without resorting to reference exemplars have also been tried [73, 92, 95, 106, 108]. However, most of the learning-based methods require laborious preprocessing steps, and/or hand-crafted image features. As a result, their performances tend to be less robust than registration-based methods.

Recently, deep convolutional models have shown great success for biomedical image segmentation [78], and have been introduced to the field of HaN anatomy segmentation [35, 42, 51, 77]. However, the existing HaN-related deep-learning-based methods either use sliding windows working on patches that cannot capture global features, or rely on atlas registration to obtain highly accurate small regions of interest in the preprocessing. What is more appealing are models that receive the whole-volume image as input without heavy-duty preprocessing, and then directly output the segmentations of all interested anatomies.

In this work, we study the feasibility and performance of constructing and training a deep neural net model that jointly segment all OARs in a fully end-to-end fashion, receiving raw whole-volume HaN CT images as input and generating the masks of all OARs in one shot. The success of such a system can improve the current performance of automated anatomy segmentation by simplifying the entire computational pipeline, cutting computational cost and improving segmentation accuracy.

There are, however, a number of obstacles that need to overcome in order to make such a deep convolutional neural net based system successful. First, in designing network architectures, we ought to keep the maximum capacity of GPU memories in mind. Since whole-volume images are used as input, each image feature map will be 3D, limiting the size and number of feature maps at each layer of the neural net due to memory constraints. Second, OARs contain organs/regions of variable sizes, including some OARs with very small sizes. Accurately segmenting these small-volumed structures is always a challenge. Third, existing datasets of HaN CT images contain data collected from various sources with non-standardized annotations. In particular, many images in the training data contain annotations of only a subset of OARs. How to effectively handle missing annotations needs to be addressed in the design of the training algorithms.

Here we propose a deep learning based framework, called AnatomyNet, to segment OARs using a single network, trained end-to-end. The network receives whole-volume CT images as input, and outputs the segmented masks of all OARs. Our method requires minimal pre- and post-processing, and utilizes features from all slices to segment anatomical regions. We overcome the three major obstacles outlined above through designing a novel network architecture and utilizing novel loss functions for training the network.

More specifically, our major contributions include the following. First, we extend the standard U-Net model for 3D HaN image segmentation by incorporating a new feature extraction component, based on squeeze-and-excitation (SE) residual blocks [46]. Second, we propose a

new loss function for better segmenting small-volumed structures. Small volume segmentation suffers from the imbalanced data problem, where the number of voxels inside the small region is much smaller than those outside, leading to the difficulty of training. New classes of loss functions have been proposed to address this issue, including Tversky loss [79], generalized Dice coefficients [15, 89], focal loss [64], sparsity label assignment deep multi-instance learning [119], and exponential logarithm loss. However, we found none of these solutions alone was adequate to solve the extremely data imbalanced problem (1/100,000) we face in segmenting small OARs, such as optic nerves and chiasm, from HaN images. We propose a new loss based on the combination of Dice scores and focal losses, and empirically show that it leads to better results than other losses. Finally, to tackle the missing annotation problem, we train the AnatomyNet with masked and weighted loss function to account for missing data and to balance the contributions of the losses originating from different OARs.

To train and evaluate the performance of AnatomyNet, we curated a dataset of 261 head and neck CT images from a number of publicly available sources. We carried out systematic experimental analyses on various components of the network, and demonstrated their effectiveness by comparing with other published methods. When benchmarked on the test dataset from the MICCAI 2015 competition on HaN segmentation, the AnatomyNet outperformed the state-of-the-art method by **3.3%** in terms of Dice coefficient (DSC), averaged over nine anatomical structures.

The rest of the paper is organized as follows. Section 6.2.2 describes the network structure and SE residual block of AnatomyNet. The designing of the loss function for AnatomyNet is present in Section 6.2.3. How to handle missing annotations is addressed in Section 6.2.4. Section 6.3 validates the effectiveness of the proposed networks and components. Discussions and limitations are in Section 6.4. We conclude the work in Section 6.5.

6.2 Materials and Methods

Next we describe our deep learning model to delineate OARs from head and neck CT images. Our model receives whole-volume HaN CT images of a patient as input and outputs the 3D binary masks of all OARs at once. The dimension of a typical HaN CT is around $178 \times 512 \times 512$, but the sizes can vary across different patients because of image cropping and different settings. In this work, we focus on segmenting nine OARs most relevant to head and neck cancer radiation therapy - brain stem, chiasm, mandible, optic nerve left, optic nerve right, parotid gland left, parotid gland right, submandibular gland left, and submandibular gland right. Therefore, our model will produce nine 3D binary masks for each whole volume CT.

6.2.1 Data

Before we introduce our model, we first describe the curation of training and testing data. Our data consists of whole-volume CT images together with manually generated binary masks of the nine anatomies described above. There were collected from four publicly available sources: 1) DATASET 1 (38 samples) consists of the training set from the MICCAI Head and Neck Auto Segmentation Challenge 2015 [75]. 2) DATASET 2 (46 samples) consists of CT images from the Head-Neck Cetuximab collection, downloaded from The Cancer Imaging Archive (TCIA)¹ [13]. 3) DATASET 3 (177 samples) consists of CT images from four different institutions in Québec, Canada [97], also downloaded from TCIA [13]. 4) DATASET 4 (10 samples) consists of the test set from the MICCAI HaN Segmentation Challenge 2015. We combined the first three datasets and used the aggregated data as our training data, altogether yielding 261 training samples. DATASET 4 was used as our final evaluation/test dataset so that we can benchmark our performance against published results

¹<https://wiki.cancerimagingarchive.net/>

evaluated on the same dataset. Each of the training and test samples contains both head and neck images and the corresponding manually delineated OARs.

In generating these datasets, We carried out several data cleaning steps, including 1) mapping annotation names named by different doctors in different hospitals into unified annotation names, 2) finding correspondences between the annotations and the CT images, 3) converting annotations in the radiation therapy format into usable ground truth label mask, and 4) removing chest from CT images to focus on head and neck anatomies. We have taken care to make sure that the four datasets described above are non-overlapping to avoid any potential pitfall of inflating testing or validation performance.

6.2.2 Network architecture

We take advantage of the robust feature learning mechanisms obtained from squeeze-and-excitation (SE) residual blocks [46], and incorporate them into a modified U-Net architecture for medical image segmentation. We propose a novel three dimensional U-Net with squeeze-and-excitation (SE) residual blocks and hybrid focal and dice loss for anatomical segmentation as illustrated in Fig. 6.1.

The AnatomyNet is a variant of 3D U-Net [78, 118, 120], one of the most commonly used neural net architectures in biomedical image segmentation. The standard U-Net contains multiple down-sampling layers via max-pooling or convolutions with strides over two. Although they are beneficial to learn high-level features for segmenting complex, large anatomies, these down-sampling layers can hurt the segmentation of small anatomies such as optic chiasm, which occupy only a few slices in HaN CT images. We design the AnatomyNet with only one down-sampling layer to account for the trade-off between GPU memory usage and network learning capacity. The down-sampling layer is used in the first encoding block so that the feature maps and gradients in the following layers occupy less GPU memory than other net-

work structures. Inspired by the effectiveness of squeeze-and-excitation residual features on image object classification, we design 3D squeeze-and-excitation (SE) residual blocks in the AnatomyNet for OARs segmentation. The SE residual block adaptively calibrates residual feature maps within each feature channel. The 3D SE Residual learning extracts 3D features from CT image directly by extending two-dimensional squeeze, excitation, scale and convolutional functions to three-dimensional functions. It can be formulated as

$$\begin{aligned}
\mathbf{X}^r &= \mathbf{F}(\mathbf{X}), \\
z_k &= \mathbf{F}_{sq}(\mathbf{X}^r_k) = \frac{1}{S \times H \times W} \sum_{s=1}^S \sum_{h=1}^H \sum_{w=1}^W x_k^r(s, h, w), \\
\mathbf{z} &= [z_1, z_2, \dots, z_k, \dots, z_K], \\
\mathbf{s} &= \mathbf{F}_{ex}(\mathbf{z}, \mathbf{W}) = \sigma(\mathbf{W}_2 \mathbf{G}(\mathbf{W}_1 \mathbf{z})), \\
\tilde{\mathbf{X}}_k &= \mathbf{F}_{scale}(\mathbf{X}^r_k, s_k) = s_k \mathbf{X}^r_k, \\
\tilde{\mathbf{X}} &= [\tilde{\mathbf{X}}_1, \tilde{\mathbf{X}}_2, \dots, \tilde{\mathbf{X}}_k, \dots, \tilde{\mathbf{X}}_K], \\
\mathbf{Y} &= \mathbf{G}(\tilde{\mathbf{X}} + \mathbf{X}),
\end{aligned} \tag{6.1}$$

where $\mathbf{X}^r_k \in \mathbb{R}^3$ denotes the feature map of one channel from the residual feature \mathbf{X}^r . \mathbf{F}_{sq} is the squeeze function, which is global average pooling here. S, H, W are the number of slices, height, and width of \mathbf{X}^r respectively. \mathbf{F}_{ex} is the excitation function, which is parameterized by two layer fully connected neural networks here with activation functions \mathbf{G} and σ , and weights \mathbf{W}_1 and \mathbf{W}_2 . The σ is the sigmoid function. The \mathbf{G} is typically a ReLU function, but we use LeakyReLU in the AnatomyNet [67]. We use the learned scale value s_k to calibrate the residual feature channel \mathbf{X}^r_k , and obtain the calibrated residual feature $\tilde{\mathbf{X}}$. The SE block is illustrated in the upper right corner in Fig. 6.1.

The AnatomyNet replaces the standard convolutional layers in the U-Net with SE residual blocks to learn effective features. The input of AnatomyNet is a cropped whole-volume head and neck CT image. We remove the down-sampling layers in the second, third, and fourth

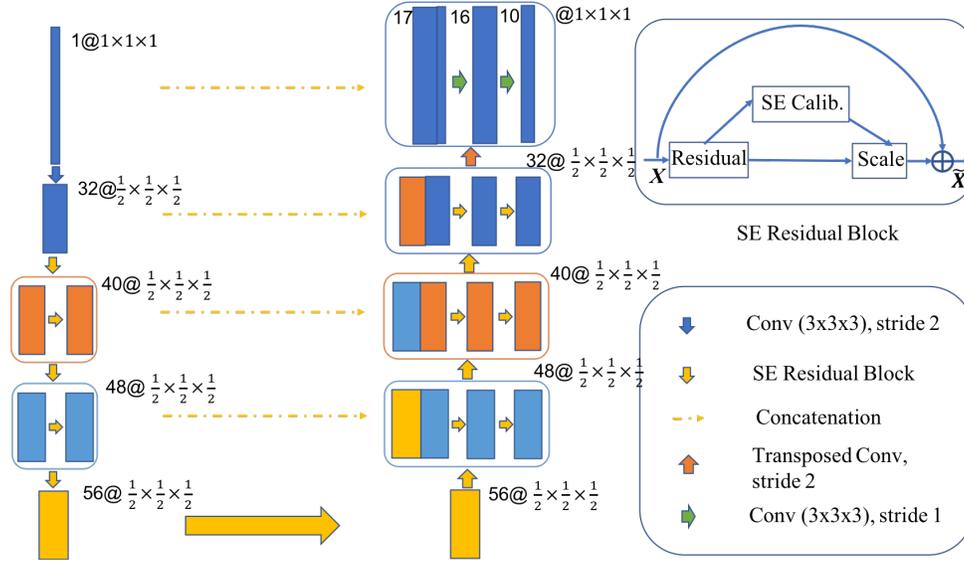


Figure 6.1: The AnatomyNet is a variant of U-Net with only one down-sampling and squeeze-and-excitation (SE) residual building blocks. The number before symbol @ denotes the number of output channels, while the number after the symbol denotes the size of feature map relative to the input. In the decoder, we use concatenated features. Hybrid loss with dice loss and focal loss is employed to force the model to learn not-well-classified voxels. Masked and weighted loss function is used for ground truth with missing annotations and balanced gradient descent respectively. The decoder layers are symmetric with the encoder layers. The SE residual block is illustrated in the upper right corner.

encoder blocks to improve the performance of segmenting small anatomies. In the output block, we concatenate the input with the transposed convolution feature maps obtained from the second last block. After that, a convolutional layer with 16 $3 \times 3 \times 3$ kernels and LeakyReLU activation function is employed. In the last layer, we use a convolutional layer with 10 $3 \times 3 \times 3$ kernels and soft-max activation function to generate the segmentation probability maps for nine OARs plus background.

6.2.3 Loss function

Small object segmentation is always a challenge in semantic segmentation. From the learning perspective, the challenge is caused by imbalanced data distribution, because image semantic

segmentation requires pixel-wise labeling and small-volumed organs contribute less to the loss. In our case, the small-volumed organs, such as optic chiasm, only take about 1/100,000 of the whole-volume CT images from Fig. 6.2. The dice loss, the minus of dice coefficient (DSC), can be employed to partly address the problem by turning pixel-wise labeling problem into minimizing class-level distribution distance [79].

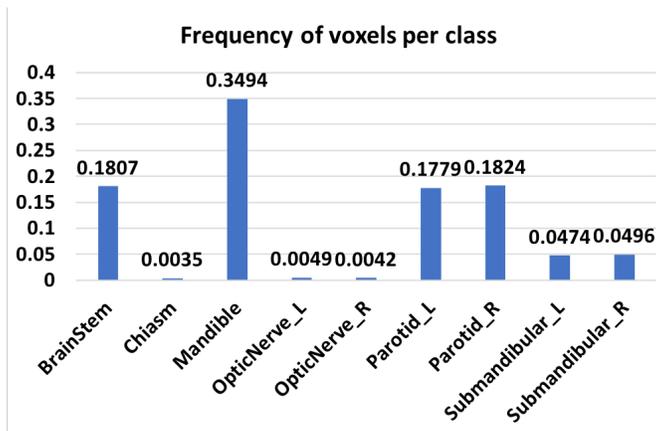


Figure 6.2: The frequency of voxels for each class on MICCAI 2015 challenge dataset. Background takes up 98.18% of all the voxels. Chiasm takes only 0.35% of the foreground which means it only takes about 1/100,000 of the whole-volume CT image. The huge imbalance of voxels in small-volumed organs causes difficulty for small-volumed organ segmentation.

Several methods have been proposed to alleviate the small-volumed organ segmentation problem. The generalized dice loss uses squared volume weights. However, it makes the optimization unstable in the extremely unbalanced segmentation [89]. The exponential logarithmic loss is inspired by the focal loss for class-level loss as $\mathbb{E}[(-\ln(D))^\gamma]$, where D is the dice coefficient (DSC) for the interested class, γ can be set as 0.3, and \mathbb{E} is the expectation over classes and whole-volume CT images. The gradient of exponential logarithmic loss w.r.t. DSC D is $-\frac{0.3}{D \ln(D)^{0.7}}$. The absolute value of gradient is getting bigger for well-segmented class (D close to 1). Therefore, the exponential logarithmic loss still places more weights on well-segmented class, and is not effective in learning to improve on not-well-segmented class.

In the AnatomyNet, we employ a hybrid loss consisting of contributions from both dice loss

and focal loss [64]. The dice loss learns the class distribution alleviating the imbalanced voxel problem, where as the focal loss forces the model to learn poorly classified voxels better. The total loss can be formulated as

$$\begin{aligned}
TP_p(c) &= \sum_{n=1}^N p_n(c)g_n(c) \\
FN_p(c) &= \sum_{n=1}^N (1 - p_n(c))g_n(c) \\
FP_p(c) &= \sum_{n=1}^N p_n(c)(1 - g_n(c))
\end{aligned} \tag{6.2}$$

$$\begin{aligned}
\mathcal{L} &= \mathcal{L}_{Dice} + \lambda \mathcal{L}_{Focal} \\
&= C - \sum_{c=0}^{C-1} \frac{TP_p(c)}{TP_p(c) + \alpha FN_p(c) + \beta FP_p(c)} \\
&\quad - \lambda \frac{1}{N} \sum_{c=0}^{C-1} \sum_{n=1}^N g_n(c)(1 - p_n(c))^2 \log(p_n(c)),
\end{aligned}$$

where $TP_p(c)$, $FN_p(c)$ and $FP_p(c)$ are the true positives, false negatives and false positives for class c calculated by prediction probabilities respectively, $p_n(c)$ is the predicted probability for voxel n being class c , $g_n(c)$ is the ground truth for voxel n being class c , C is the total number of anatomies plus one (background), λ is the trade-off between dice loss \mathcal{L}_{Dice} and focal loss \mathcal{L}_{Focal} , α and β are the trade-offs of penalties for false negatives and false positives which are set as 0.5 here, N is the total number of voxels in the CT images. λ is set to be 0.1, 0.5 or 1 based on the performance on the validation set. Because of size differences for different HaN whole-volume CT images, we set the batch size to be 1.

6.2.4 Handling missing annotations

Another challenge in anatomical segmentation is due to missing annotations common in the training datasets, because annotators often include different anatomies in their annotations.

For example, we collect 261 head and neck CT images with anatomical segmentation ground truths from 5 hospitals, and the numbers of nine annotated anatomies are very different as shown in Table 6.1. To handle this challenge, we mask out the background (denoted as class 0) and the missed anatomy. Let $c \in \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$ denote the index of anatomies. We employ a mask vector \mathbf{m}_i for the i th CT image, and denote background as label 0. That is $\mathbf{m}_i(c) = 1$ if anatomy c is annotated, and 0 otherwise. For the background, the mask is $\mathbf{m}_i(0) = 1$ if all anatomies are annotated, and 0 otherwise. The missing annotations for some

Table 6.1: The numbers of the nine annotated anatomies from 261 head and neck training CT images.

Anatomy names	#Annotations
Brain Stem	196
Chiasm	129
Mandible	227
Opt Ner L	133
Opt Ner R	133
Parotid L	257
Parotid R	256
Submand. L	135
Submand. R	130

anatomical structures cause imbalanced class-level annotations. To address this problem, we employ weighted loss function for balanced weights updating of different anatomies. The weights \mathbf{w} are set as the inverse of the number of annotations for class c , $\mathbf{w}(c) = 1/\sum_i \mathbf{m}_i(c)$, so that the weights in deep networks are updated equally with different anatomies. The dice loss for i th CT image in equation 6.2 can be written as

$$\tilde{\mathcal{L}}_{Dice} = C - \sum_{c=0}^{C-1} \mathbf{m}_i(c) \mathbf{w}(c) \frac{TP_p(c)}{TP_p(c) + \alpha FN_p(c) + \beta FP_p(c)}. \quad (6.3)$$

The focal loss for missing annotations in the i th CT image can be written as

$$\tilde{\mathcal{L}}_{Focal} = -\frac{1}{N} \sum_{c=0}^{C-1} \mathbf{m}_i(c) \mathbf{w}(c) \sum_{n=1}^N g_n(c) (1 - p_n(c))^2 \log(p_n(c)). \quad (6.4)$$

We use loss $\tilde{\mathcal{L}}_{Dice} + \lambda \tilde{\mathcal{L}}_{Focal}$ in the AnatomyNet.

6.2.5 Implementation details and performance evaluation

We implemented AnatomyNet in PyTorch, and trained it on NVIDIA Tesla P40. Batch size was set to be 1 because of different sizes of whole-volume CT images. We first used RMSprop optimizer [94] with learning rate being 0.002 and the number of epochs being 150. Then we used stochastic gradient descend with momentum 0.9, learning rate 0.001 and the number of epochs 50. During training, we used affine transformation and elastic deformation for data augmentation, implemented on the fly.

We use Dice coefficient (DSC) as the final evaluation metric, defined to be $2TP/(2TP + FN + FP)$, where TP , FN , and FP are true positives, false negatives, false positives, respectively.

6.3 Results

We trained our deep learning model, AnatomyNet, on 261 training samples, and evaluated its performance on the MICCAI head and neck segmentation challenge 2015 test data (10 samples, DATASET 4) and compared it to the performances of previous methods benchmarked on the same test dataset. Before we present the final results, we first describe the rationale behind several designing choices under AnatomyNet, including architectural designs and model training.

6.3.1 Determining down-sampling scheme

The standard U-Net model has multiple down-sampling layers, which help the model learn high-level image features. However, down-sampling also reduces image resolution and makes it harder to segment small OARs such as optic nerves and chiasm. To evaluate the effect of the number of down-sampling layers on the segmentation performance, we experimented with four different down-sampling schemes shown in Table 6.2. Pool 1 uses only one down-sampling step, while Pool 2, 3, and 4 use 2, 3 and 4 down-sampling steps, respectively, distributed over consecutive blocks. With each down-sampling, the feature map size is reduced by half. We incorporated each of the four down-sampling schemes into the standard U-Net model, which was then trained on the training set and evaluated on the test set. For fair comparisons, we used the same number of filters in each layer. The decoder layers of each model are set to be symmetric with the encoder layers.

Table 6.2: Sizes of encoder blocks in U-Nets with different numbers of down-samplings. The number before symbol @ denotes the number of output channels, while the number after the symbol denotes the size of feature map relative to the input.

Nets	1st block	2nd block	3rd block	4th block
Pool 1	32@(1/2) ³	40@(1/2) ³	48@(1/2) ³	56@(1/2) ³
Pool 2	32@(1/2) ³	40@(1/4) ³	48@(1/4) ³	56@(1/4) ³
Pool 3	32@(1/2) ³	40@(1/4) ³	48@(1/8) ³	56@(1/8) ³
Pool 4	32@(1/2) ³	40@(1/4) ³	48@(1/8) ³	56@(1/16) ³

The DSC scores of the four down-sampling schemes are shown in Table 6.3. On average, one down-sampling block (Pool 1) yields the best average performance, beating other down-sampling schemes in 6 out of 9 anatomies. The performance gaps are most prominent on three small-volumed OARs - optic nerve left, optic nerve right and optic chiasm, which demonstrates that the U-Net with one down-sampling layer works better on small organ segmentation than the standard U-Net. The probable reason is that small organs reside in

only a few slices and more down-sampling layers are more likely to miss features for the small organs in the deeper layers. Based on these results, we decide to use only one down-sampling layer in AnatomyNet (Fig. 6.1).

Table 6.3: Performances of U-Net models with different numbers of down-sampling layers, measured with Dice coefficients.

Anatomy	Pool 1	Pool 2	Pool 3	Pool 4
Brain Stem	85.1	85.3	84.3	84.9
Chiasm	50.1	48.7	47.0	45.3
Mand.	91.5	89.9	90.6	90.1
Optic Ner L	69.1	65.7	67.2	67.9
Optic Ner R	66.9	65.0	66.2	63.7
Paro. L	86.6	86.6	85.9	86.6
Paro. R	85.6	85.3	84.8	85.9
Subm. L	78.5	77.9	77.9	77.3
Subm. R	77.7	78.4	76.6	77.8
Average	76.8	75.9	75.6	75.5

6.3.2 Choosing network structures

In addition to down-sample schemes, we also tested several other architecture designing choices. The first one is on how to combine features from horizontal layers within U-Net. Traditional U-Net uses concatenation to combine features from horizontal layers in the decoder, as illustrated with dash lines in Fig. 6.1. However, recent feature pyramid network (FPN) recommends summation to combine horizontal features [63]. Another designing choice is on choosing local feature learning blocks with each layer. The traditional U-Net uses simple 2D convolution, extended to 3D convolution in our case. To learn more effective features, we tried two other feature learning blocks: a) residual learning, and b) squeeze-and-excitation residual learning. Altogether, we investigated the performances of the following six architec-

tural designing choices:

1. **3D SE Res UNet**, the architecture implemented in AnatomyNet (Fig. 6.1) with both squeeze-excitation residual learning and concatenated horizontal features.
2. **3D Res UNet**, replacing the SE Residual blocks in 3D SE Res UNet with residual blocks.
3. **Vanilla U-Net**, replacing the SE Residual blocks in 3D SE Res UNet with 3D convolutional layers.
4. **3D SE Res UNet (sum)**, replacing concatenations in 3D SE Res UNet with summations. When the numbers of channels are different, one additional $1 \times 1 \times 1$ 3D convolutional layer is used to map the encoder to the same size as the decoder.
5. **3D Res UNet (sum)**, replacing the SE Residual blocks in 3D SE Res UNet (sum) with residual blocks.
6. **Vanilla U-Net (sum)**, replacing the SE Residual blocks in 3D SE Res UNet (sum) with 3D convolutional layers.

The six models were trained on the same training dataset with identical training procedures. The performances measured by DSC on the test dataset are summarized in Table 6.4. We notice a few observations from this study. First, feature concatenation shows consistently better performance than feature summation. It seems feature concatenation provides more flexibility in feature learning than the fixed operation through feature summation. Second, 3D SE residual U-Net with concatenation yields the best performance. It demonstrates the power of SE features on 3D semantic segmentation, because the SE scheme learns the channel-wise calibration and helps alleviate the dependencies among channel-wise features as discussed in Section 6.2.2.

The SE residual block learning incorporated in AnatomyNet results in 2-3% improvements in DSC over the traditional U-Net model, outperforming U-Net in 6 out of 9 anatomies.

Table 6.4: Performance comparison on different network structures

Anatomy	Vanilla UNet	3D Res UNet	3D SE Res UNet	Vanilla UNet (sum)	3D Res UNet (sum)	3D SE Res UNet (sum)
Brain Stem	85.1	85.9	86.4	85.0	85.8	86.0
Chiasm	50.1	53.3	53.2	50.8	49.8	53.5
Mand.	91.5	90.6	92.3	90.2	90.9	91.3
Optic Ner L	69.1	69.8	72.0	66.5	70.4	68.9
Optic Ner R	66.9	67.5	69.1	66.1	67.4	45.6
Paro. L	86.6	87.8	87.8	86.8	87.5	87.4
Paro. R	85.6	86.2	86.8	86.0	86.7	86.9
Subm. L	78.5	79.9	81.1	79.3	79.1	82.4
Subm. R	77.7	80.2	80.8	77.6	78.4	80.5
Average	76.8	77.9	78.8	76.5	77.3	75.8

6.3.3 Choosing loss functions

We also validated the effects of different loss functions on training and model performance. To differentiate the effects of loss functions from network design choices, we used only the vanilla U-Net and trained it with different loss functions. This way, we can focus on studying the impact of loss functions on model performances. We tried four loss functions, including Dice loss, exponential logarithmic loss, hybrid loss between Dice loss and focal loss, and hybrid loss between Dice loss and cross entropy. The trade-off parameter in hybrid losses (λ in Eq. 6.2) was chosen from either 0.1, 0.5 or 1, based on the performance on a validation set. For hybrid loss between Dice loss and focal loss, the best λ was found to be 0.5. For

hybrid loss between Dice loss and cross entropy, the best λ was 0.1.

Table 6.5: Comparisons of test performances of models trained with different loss functions, evaluated with Dice coefficients.

Anatomy	Dice loss	Exp. Log. Dice	Dice + focal	Dice + cross entropy
Brain Stem	85.1	85.0	86.1	85.2
Chiasm	50.1	50.0	52.2	48.8
Mand.	91.5	89.9	90.0	91.0
Optic Ner L	69.1	67.9	68.4	69.6
Optic Ner R	66.9	65.9	69.1	67.4
Paro. L	86.6	86.4	87.4	88.0
Paro. R	85.6	84.8	86.3	86.9
Subm. L	78.5	76.3	79.6	77.8
Subm. R	77.7	78.2	79.8	78.4
Average	76.8	76.0	77.7	77.0

The performances of the model trained with the four loss functions described above are shown in Table 6.5. The performances are measured in terms of the average DSC on the test dataset. We notice a few observations from this experiment. First, the two hybrid loss functions consistently outperform simple Dice or exponential logarithmic loss, beating the other two losses in 8 out of 9 anatomies. This suggests that taking the voxel-level loss into account can improve performance. Second, between the two hybrid losses, Dice combined with focal loss has better performances. In particular, it leads to significant improvements (2-3%) on segmenting two small anatomies - optic nerve R and optic chiasm, consistent with our motivation discussed in the Section 6.2.3.

Based on the above observations, the hybrid loss with Dice combined with focal loss was used to train AnatomyNet, and benchmark its performance against previous methods.

6.3.4 Comparing to state-of-the-art methods

After having determined the structure of AnatomyNet and the loss function for training it, we set out to compare its performance with previous state-of-the-art methods. For consistency purpose, all models were evaluated on the MICCAI head and neck challenge 2015 test set. The average DSC of different methods are summarized in Table 6.6. The best result for each anatomy from the MICCAI 2015 challenge is denoted as MICCAI 2015 [75], which may come from different teams with different methods.

Table 6.6: Performance comparisons with state-of-the-art methods, showing the average DSC on the test set.

Anatomy	MICCAI 2015 [75]	Fritscher et al. 2016 [35]	Ren et al. 2018 [77]	Wang et al. 2018 [106]	AnatomyNet
Brain Stem	88	N/A	N/A	90±4	86.65±2
Chiasm	55	49±9	58±17	N/A	53.22±15
Mand.	93	N/A	N/A	94±1	92.51±2
Optic Ner L	62	N/A	72±8	N/A	72.10±6
Optic Ner R	62	N/A	70±9	N/A	70.64±10
Paro. L	84	81±4	N/A	83±6	88.07±2
Paro. R	84	81±4	N/A	83±6	87.35±4
Subm. L	78	65±8	N/A	N/A	81.37±4
Subm. R	78	65±8	N/A	N/A	81.30±4
Average	76	N/A	N/A	N/A	79.25

MICCAI 2015 competition merged left and right paired organs into one target, while we treat them as two separate anatomies. As a result, MICCAI 2015 competition is a seven (6 organs + background) class segmentation and ours is a ten-class segmentation, which makes the segmentation task more challenging. Nonetheless, the AnatomyNet achieves an

average Dice coefficient of 79.25, which is 3.3% better than the best result from MICCAI 2015 Challenge (Table 6.6). In particular, the improvements on optic nerves are about 9-10%, suggesting that deep learning models are better equipped to handle small anatomies with large variations among patients. The AnatomyNet also outperforms the atlas based ConvNets in [35] on all classes, which is likely contributed by the fact that the end-to-end structure in AnatomyNet for whole-volume HaN CT image captures global information for relative spatial locations among anatomies. Compared to the interleaved ConvNets in [77] on small-volumed organs, such as chiasm, optic nerve left and optic nerve right, AnatomyNet is better on 2 out of 3 cases. The interleaved ConvNets achieved higher performance on chiasm, which is likely contributed by the fact that its prediction was operated on small region of interest (ROI), obtained first through atlas registration, while AnatomyNet operates directly on whole-volume slices.

Aside from the improvement on segmentation accuracy, another advantage of AnatomyNet is that it is orders of magnitude faster than traditional atlas-based methods using in the MICCAI 2015 challenge. AnatomyNet takes about 0.12 seconds to fully segment a head and neck CT image of dimension $178 \times 302 \times 225$. By contrast, the atlas-based methods can take a dozen minutes to complete one segmentation depending on implementation details and the choices on the number of atlases.

6.3.5 Visualizations on MICCAI 2015 test

In Fig. 6.3 and Fig. 6.4, we visualize the segmentation results by AnatomyNet on four cases from the test dataset. Each row represents one (left and right) anatomy or 3D reconstructed anatomy. Each column denotes one sample. The last two columns show cases where AnatomyNet did not perform well. The discussions of these cases are presented in Section 6.4.2. Green denotes the ground truth. Red represents predicted segmentation results. Yellow

denotes the overlap between ground truth and prediction. We visualize the slices containing the largest area of each related organ. For small OARs such as optic nerves and chiasm (shown in Fig. 6.4), only cross-sectional slices are shown.

6.3.6 Visualizations on independent samples

To check the generalization ability of the trained model, we also visualize the segmentation results of the trained model on a small internal dataset in Fig. 6.5 and Fig. 6.6. Visual inspection suggests that the trained model performed well on this independent test set. In general, the performances on larger anatomies are better than small ones (such as optic chiasm), which can be attributed by both manual annotation inconsistencies and algorithmic challenges in segmenting these small regions.

6.4 Discussion

6.4.1 Impacts of training datasets

The training datasets we collected come from various sources with annotations done by different groups of physicians with different guiding criteria. It is unclear how the different datasets might contribute the model performance. For this purpose, we carried out an experiment to test the model performance trained with two different datasets: a) using only the training data provided in the MICCAI head and neck segmentation challenge 2015 (DATASET 1, 38 samples), and b) the combined training data with 216 samples (DATASET 1-3 combined). In terms of annotations, the first dataset is more consistent with the test dataset, therefore less likely to suffer from annotational inconsistencies. However, on the other hand, the size of the dataset is much smaller, posing challenges to training deep

learning models.

Table 6.7 shows the test performances of a 3D Res U-Net model trained with the above-mentioned two datasets after applying the same training procedure of minimizing Dice loss. We notice a few observations. First, overall the model trained with the larger dataset (DATASET 1-3) achieves better performance with a 2.5% improvement over the smaller dataset, suggesting that the larger sample size does lead to a better performance. Second, although the larger dataset improves performances on average, there are some OARs on which the smaller dataset actually does better, most noticeably, on mandible and optic nerves. This suggests that there are indeed significant data annotation inconsistencies between different datasets, whose impact on model performance cannot be neglected. Third, to further check the generalization ability of the model trained with DATASET 1 only, we checked its performance on DATASETS 2-3 and found its performance was generally poor. Altogether, this suggests both annotation quality and data size are important for training deep learning models. How to address inconsistencies in existing datasets is an interesting open question to be addressed in the future.

Table 6.7: Performance comparisons of models trained with different datasets.

Datasets	DATASET 1	DATASET 1,2,3
Brain Stem	58.60	85.91
Chiasm	39.93	53.26
Mandible	94.16	90.59
Opt Ner L	74.62	69.80
Opt Ner R	73.77	67.50
Parotid L	88.83	87.84
Parotid R	87.24	86.15
Submand. L	78.56	79.91
Submand. R	81.83	80.24
Average	75.28	77.91

6.4.2 Limitations

There are a couple of limitations in the current implementation of AnatomyNet. First, AnatomyNet treats voxels equally in the loss function and network structure. As a result, it cannot model the shape prior and connectivity patterns effectively. The translation and rotation invariance of convolution are great for learning appearance features, but suffer from the loss of spatial information. For example, the AnatomyNet sometimes misclassifies a small background region into OARs (Fig. 6.3,6.4). The mis-classification results in a partial anatomical structures, which can be easily excluded if the overall shape information can also be learned. A network with multi-resolution outputs from different levels of decoders, or deeper layers with bigger local receptive fields should help alleviate this issue.

Table 6.8: Average 95th percentile Hausdorff distance (unit: mm) comparisons with state-of-the-art methods

Anatomy	MICCAI 2015 Range [75]	Ren et al. 2018 [77]	AnatomyNet
Brain Stem	4-6	N/A	6.42±2.38
Chiasm	3-4	2.81±1.56	5.76±2.49
Mand.	2-13	N/A	6.28±2.21
Optic Ner L	3-8	2.33±0.84	4.85±2.32
Optic Ner R	3-8	2.13±0.96	4.77±4.27
Paro. L	5-8	N/A	9.31±3.32
Paro. R	5-8	N/A	10.08±5.09
Subm. L	4-9	N/A	7.01±4.44
Subm. R	4-9	N/A	6.02±1.78

Second, our evaluation of the segmentation performance is primarily based on the Dice coefficient. Although it is a common metric used in image segmentation, it may not be the most relevant one for clinical applications. Identifying a new metric in consultation

with the physicians practicing in the field would be an important next step in order for real clinical applications of the method. Along this direction, we quantitatively evaluated the geometric surface distance by calculating the average 95th percentile Hausdorff distance (unit: mm, detailed formulation in [75]) (Table 6.8). We should note that this metric imposes more challenges to AnatomyNet than other methods operating on local patches (such as the method by Ren et al. [77]), because AnatomyNet operates on whole-volume slices and a small outlier prediction outside the normal range of OARs can lead to drastically bigger Hausdorff distance. Nonetheless, AnatomyNet is roughly within the range of the best MICCAI 2015 challenge results on six out of nine anatomies [75]. Its performance on this metric can be improved by considering surface and shape priors into the model as discussed above [70, 121].

6.5 Conclusion

In summary, we have proposed an end-to-end atlas-free and fully automated deep learning model for anatomy segmentation from head and neck CT images. We propose a number of techniques to improve model performance and facilitate model training. To alleviate highly imbalanced challenge for small-volumed organ segmentation, a hybrid loss with class-level loss (dice loss) and focal loss (forcing model to learn not-well-predicted voxels better) is employed to train the network, and one single down-sampling layer is used in the encoder. To handle missing annotations, masked and weighted loss is implemented for accurate and balanced weights updating. The 3D SE block is designed in the U-Net to learn effective features. Our experiments demonstrate that our model provides new state-of-the-art results on head and neck OARs segmentation, outperforming previous models by 3.3%. It is significantly faster, requiring only a fraction of a second to segment nine anatomies from a head and neck CT. In addition, the model is able to process a whole-volume CT and delineate all OARs in one pass. All together, our work suggests that deep learning offers

a flexible and efficient framework for delineating OARs from CT images. With additional training data and improved annotations, it would be possible to further improve the quality of auto-segmentation, bringing it closer to real clinical practice.

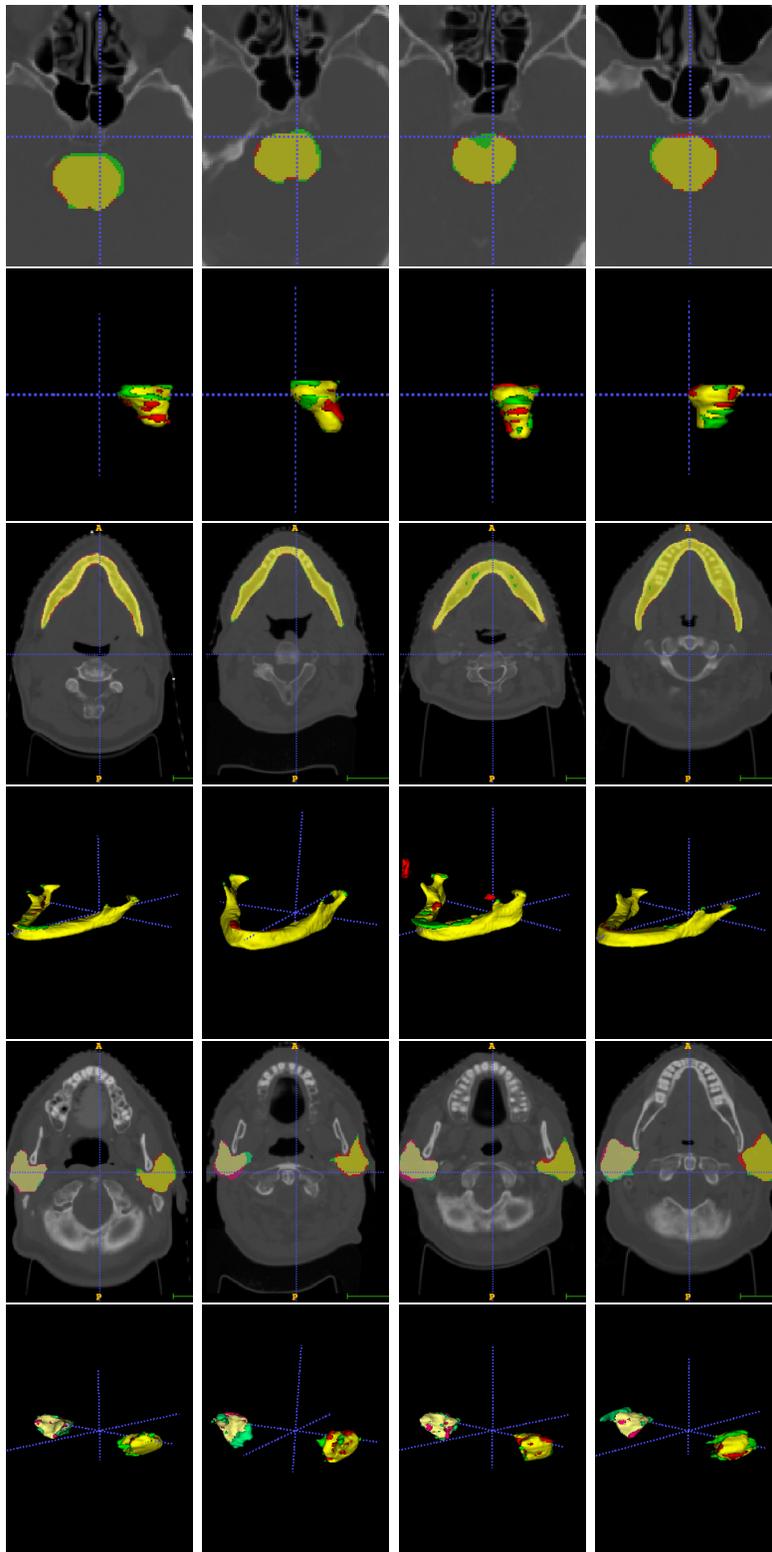


Figure 6.3: Visualizations on four test CT images. Rows from top to bottom represent brain stem, brain stem 3D, mandibular, mandibular 3D, parotid left and right, and parotid left and right 3D. Each column represents one CT image. Green represents ground truths, and red denotes predictions. Yellow is the overlap.

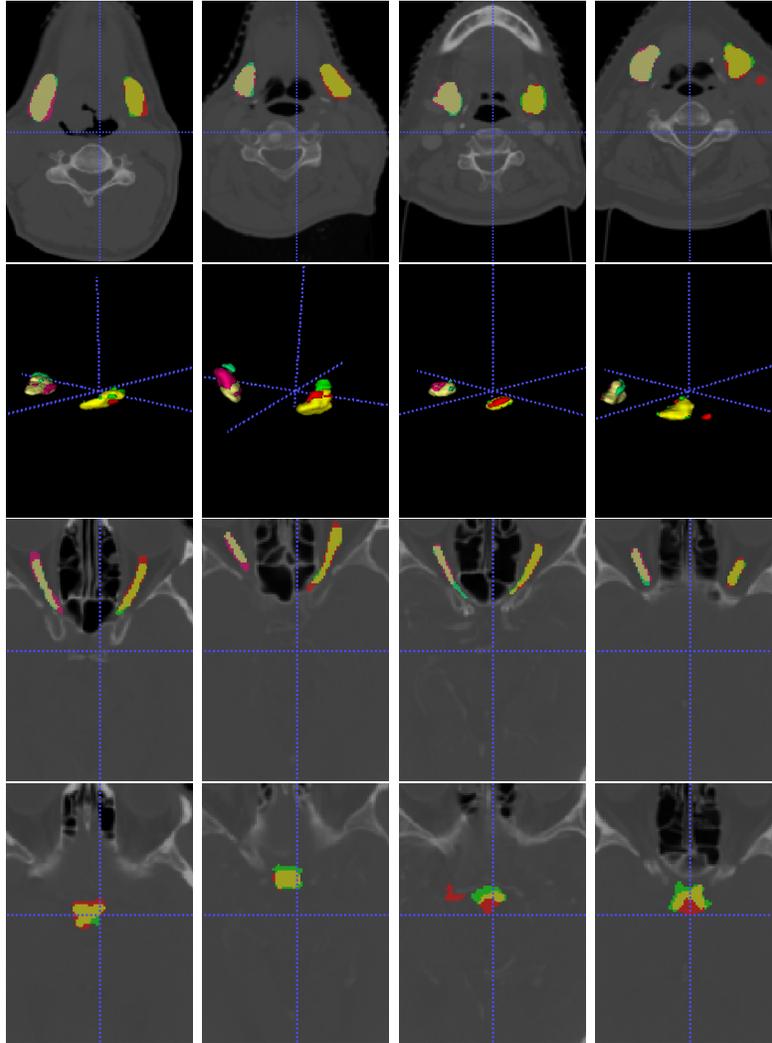


Figure 6.4: Visualizations on four test CT images. Rows from top to bottom represent submandibular left and right, submandibular left and right 3D, optic nerve left and right, and chiasm. Each column represents one CT image. Green represents ground truths, and red denotes predictions. Yellow is the overlap. The AnatomyNet performs well on small-volumed anatomies.

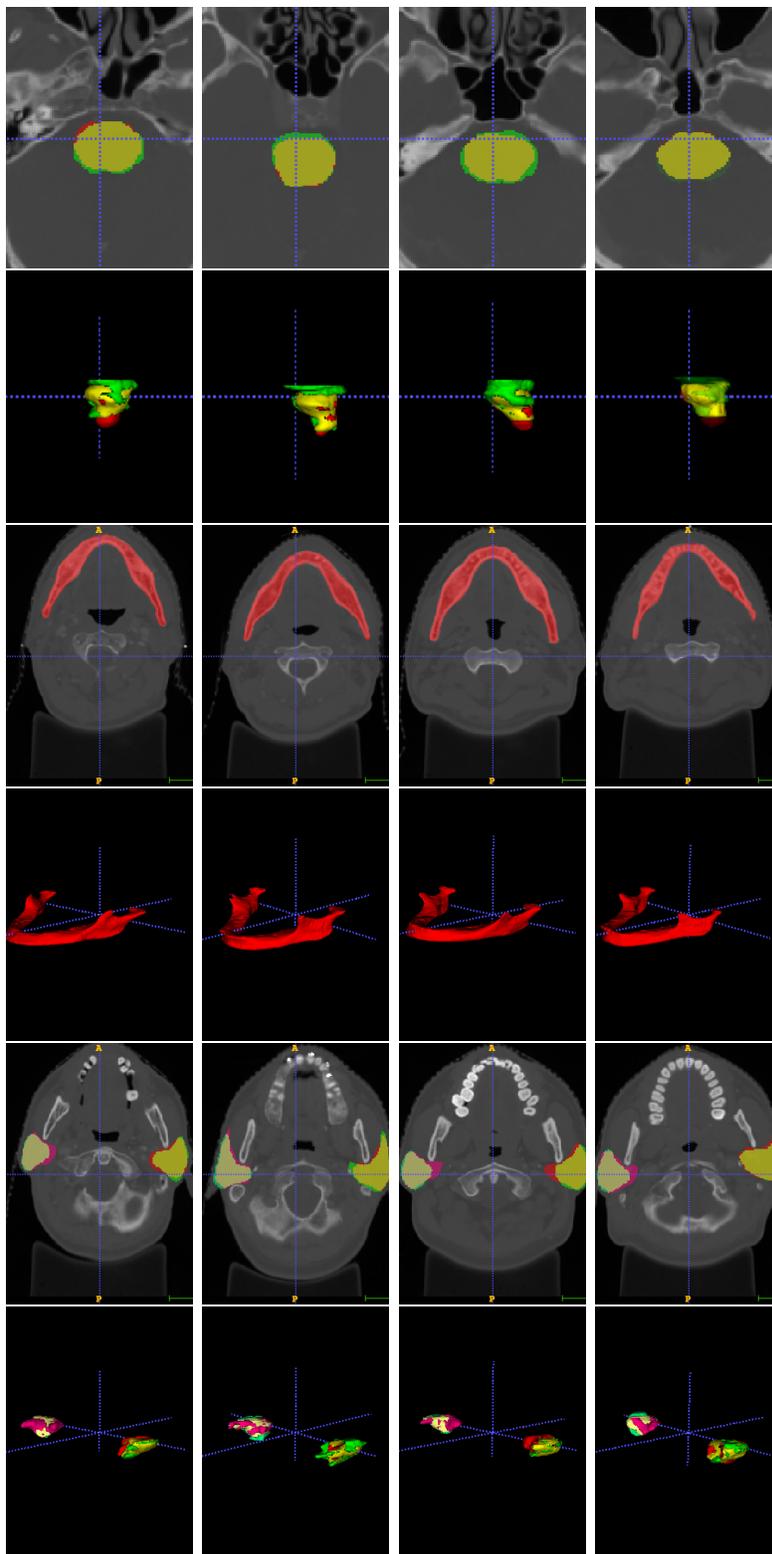


Figure 6.5: Visualizations for the first four anatomy on the first four holdout CT images. There is no ground truth for mandible and submandibular glands. Because this is a different source from MICCAI 2015, the annotations of brain stem and chiasm are inconsistent with those from MICCAI 2015. The AnatomyNet generalizes well for hold out test set.

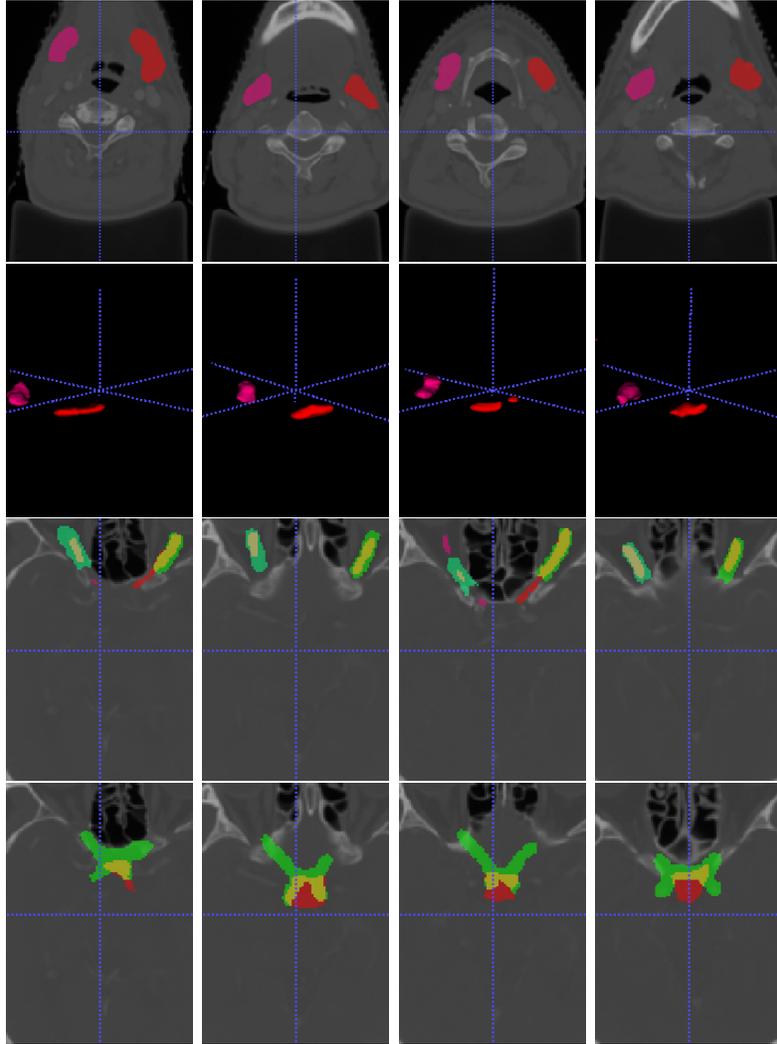


Figure 6.6: Visualizations for the rest five anatomies on the first four holdout CT images.

Chapter 7

Conclusion

The rapid progress of deep learning on medical image analysis has undoubtedly revolutionized the common clinical practice including breast cancer mammogram screening, pulmonary cancer CT screening, and radiotherapy in head and neck cancer treatment planning. In this thesis, we introduce how to develop a mass segmentation system based on mammogram in chapter 2, how to develop a breast cancer screening system in chapter 3, how to use 3D convolutional neural networks for lung computed tomography (CT) nodule detection and classification system in chapter 4, how to use electronic medical report (EMR) to further improve the nodule detection system in chapter 5, and how to design 3D convolutional neural networks to significantly reduce the delineation time of organ at risks (OARs) for radiotherapy planning in chapter 6. These methods reduce the false positives, improve sensitivities or reduce the delay time in the current clinical practice, and can be potentially great tools for radiologists.

The main challenge for deep learning based medical image analysis is that, deep learning models typically require a lot of training data but the datasets of medical image computing are typically small because 1) patient privacy prevents building large scale datasets, 2) the

ground-truth requires experienced radiologists laborious annotating or even biopsy results, 3) some anatomies are small naturally. To alleviate the difficulty, we 1) use adversarial concept to generate the hard example to improve the generalization ability of mass segmentation [121], 2) formulate whole mammogram classification as a deep multi-instance learning and propose three inference schemes to employ image-level label for mammogram diagnosis system, instead of requiring detection and segmentation ground-truth [119], 3) propose a deep learning based probabilistic graphic model employing weak labels from existing in-house EMR to further improve the lung nodule detection system [120], 4) proposed a novel encoding framework with weighted hybrid loss of class-distribution loss (Dice loss) and adaptively weighted pixel-wise loss (focal loss) for small anatomy segmentation in radiotherapy planning.

However, there are still a few challenges in the deep learning based medical image analysis. First, medical images, such as ultrasound images, are typically of low signal-to-noise ratio compared with natural images, and echo-gram is a cheap way for screening. How to design a framework that can bring breakthrough for echo image analysis is a significant task for medical image computing research. Second, medical images, such as CT or MRI, are typically very large. The spatial 3D data requires 3D models naturally which occupy a lot of GPU memory. How to design light models both for training and inference is an interesting direction for the community. Third, the deep learning models are exposed to be of great capacity and the big complexity leads to low generalization ability on outlier samples. However, the shape of abnormal tissues are typically of big variations, and some landmarks can not be taken off during screening or therapy. Using the deep generative models to improve the generalization ability of deep models is another interesting direction for medical image computing.

Bibliography

- [1] American cancer society. what are the key statistics about breast cancer? <https://www.cancer.org/cancer/breast-cancer/about/how-common-is-breast-cancer.html>.
- [2] H. J. Aerts et al. Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nature communications*, 5, 2014.
- [3] S. G. Armato et al. The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics*, 38(2):915–931, 2011.
- [4] J. Ba and D. Kingma. Adam: A method for stochastic optimization. ICLR, 2015.
- [5] J. E. Ball and L. M. Bruce. Digital mammographic computer aided diagnosis (cad) using adaptive level set segmentation. In *EMBS*, 2007.
- [6] M. Beller, R. Stotzka, T. O. Müller, and H. Gemmeke. An example-based system to support the segmentation of stellate lesions. In *Bildverarbeitung für die Medizin 2005*, pages 475–479. Springer, 2005.
- [7] H. Bilen et al. Weakly supervised deep detection networks. In *CVPR*, 2016.
- [8] K. Bowyer, D. Kopans, W. Kegelmeyer, R. Moore, M. Sallam, K. Chang, and K. Woods. The digital database for screening mammography. In *IWDM*, 1996.
- [9] J. S. Cardoso, I. Domingues, and H. P. Oliveira. Closed shortest path in the original coordinates with an application to breast cancer. *International Journal of Pattern Recognition and Artificial Intelligence*, 29(01):1555002, 2015.
- [10] G. Carneiro, J. Nascimento, and A. P. Bradley. Unregistered multiview mammogram analysis with pre-trained deep learning models. In *MICCAI*, pages 652–660. Springer, 2015.
- [11] A. Chen, M. A. Deeley, K. J. Niermann, L. Moretti, and B. M. Dawant. Combining registration and active shape models for the automatic segmentation of the lymph node regions in head and neck ct images. *Medical physics*, 37(12):6338–6346, 2010.
- [12] Y. Chen, J. Li, H. Xiao, X. Jin, S. Yan, and J. Feng. Dual path networks. In *NIPS*, 2017.

- [13] K. Clark, B. Vendt, K. Smith, J. Freymann, J. Kirby, P. Koppel, S. Moore, S. Phillips, D. Maffitt, M. Pringle, et al. The cancer imaging archive (tcia): maintaining and operating a public information repository. *Journal of digital imaging*, 26(6):1045–1057, 2013.
- [14] O. Commowick, V. Grégoire, and G. Malandain. Atlas-based delineation of lymph node levels in head and neck computed tomography images. *Radiotherapy and Oncology*, 87(2):281–289, 2008.
- [15] W. R. Crum, O. Camara, and D. L. Hill. Generalized overlap measures for evaluation and validation in medical image analysis. *IEEE TMI*, 2006.
- [16] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*. IEEE, 2009.
- [17] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [18] N. Dhungel, G. Carneiro, and A. P. Bradley. Deep learning and structured prediction for the segmentation of mass in mammograms. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 605–612. Springer, 2015.
- [19] N. Dhungel, G. Carneiro, and A. P. Bradley. Deep structured learning for mass segmentation from mammograms. In *ICIP*, pages 2950–2954. IEEE, 2015.
- [20] N. Dhungel, G. Carneiro, and A. P. Bradley. Tree re-weighted belief propagation using deep learning potentials for mass segmentation from mammograms. In *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pages 760–763. IEEE, 2015.
- [21] N. Dhungel, G. Carneiro, and A. P. Bradley. The automated learning of deep features for breast mass classification from mammograms. In *MICCAI*, pages 106–114. Springer, 2016.
- [22] T. G. Dietterich, R. H. Lathrop, and T. Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1):31–71, 1997.
- [23] J. Ding, A. Li, Z. Hu, and L. Wang. Accurate pulmonary nodule detection in computed tomography images using deep convolutional neural networks. In *MICCAI*, 2017.
- [24] I. Domingues, E. Sales, J. Cardoso, and W. Pereira. Inbreast-database masses characterization. *XXIII CBEB*, 2012.
- [25] Q. Dou, H. Chen, Y. Jin, H. Lin, J. Qin, and P.-A. Heng. Automated pulmonary nodule detection via 3d convnets with online sample filtering and hybrid-loss residual learning. In *MICCAI*, 2017.

- [26] H. Duc, K. Albert, G. Eminowicz, R. Mendes, S.-L. Wong, J. McClelland, M. Modat, M. J. Cardoso, A. F. Mendelson, C. Veiga, et al. Validation of clinical acceptability of an atlas-based segmentation algorithm for the delineation of organs at risk in head and neck cancer. *Medical physics*, 42(9):5027–5034, 2015.
- [27] A. El-Baz et al. 3d shape analysis for early diagnosis of malignant lung nodules. In *IPMI*, 2011.
- [28] D. Y. et al. Automatic liver segmentation using an adversarial image-to-image network. In *MICCAI*. Springer, 2017.
- [29] M. J. et al. Mammographic mass segmentation with online learned shape and appearance priors. In *MICCAI*, 2016.
- [30] W. Z. et al. Deep multi-instance networks with sparse label assignment for whole mammogram classification. *MICCAI*, 2017.
- [31] X. Feng et al. Discriminative localization in cnns for weakly-supervised segmentation of pulmonary nodules. In *MICCAI*, 2017.
- [32] V. Fortunati, R. F. Verhaart, W. J. Niessen, J. F. Veenland, M. M. Paulides, and T. van Walsum. Automatic tissue segmentation of head and neck mr images for hyperthermia treatment planning. *Physics in Medicine & Biology*, 60(16):6547, 2015.
- [33] V. Fortunati, R. F. Verhaart, F. van der Lijn, W. J. Niessen, J. F. Veenland, M. M. Paulides, and T. van Walsum. Tissue segmentation of head and neck ct images for treatment planning: a multiatlas approach combined with intensity modeling. *Medical physics*, 2013.
- [34] J. H. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, 2001.
- [35] K. Fritscher, P. Raudaschl, P. Zaffino, M. F. Spadea, G. C. Sharp, and R. Schubert. Deep neural networks for fast segmentation of 3d medical images. In *MICCAI*, 2016.
- [36] K. D. Fritscher, M. Peroni, P. Zaffino, M. F. Spadea, R. Schubert, and G. Sharp. Automatic segmentation of head and neck ct images for radiotherapy treatment planning using multiple atlases, statistical appearance models, and geodesic active contours. *Medical physics*, 41(5), 2014.
- [37] Y. Gao and J. A. Noble. Detection and characterization of the fetal heartbeat in free-hand ultrasound sweeps with weakly-supervised two-streams convolutional networks. In *MICCAI*, pages 305–313. Springer, 2017.
- [38] R. Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015.
- [39] H. Greenspan, B. van Ginneken, and R. M. Summers. Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique. *IEEE TMI*, 2016.

- [40] F. Han, G. Zhang, H. Wang, B. Song, H. Lu, D. Zhao, H. Zhao, and Z. Liang. A texture feature analysis for diagnosis of pulmonary nodules using lidc-idri database. In *Medical Imaging Physics and Engineering*. IEEE, 2013.
- [41] X. Han, M. S. Hoogeman, P. C. Levendag, L. S. Hibbard, D. N. Teguh, P. Voet, A. C. Cowen, and T. K. Wolf. Atlas-based auto-segmentation of head and neck ct images. In *MICCAI*, 2008.
- [42] A. Hänsch, M. Schwier, T. Gass, T. Morgas, B. Haas, J. Klein, and H. K. Hahn. Comparison of different deep learning approaches for parotid gland segmentation from ct images. In *SPIE*, 2018.
- [43] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [44] M. Heath, K. Bowyer, D. Kopans, P. Kegelmeyer Jr, R. Moore, K. Chang, and S. Munishkumaran. Current status of the digital database for screening mammography. In *Digital mammography*, pages 457–460. Springer, 1998.
- [45] L. Hou, D. Samaras, T. M. Kurc, Y. Gao, J. E. Davis, and J. H. Saltz. Patch-based convolutional neural network for whole slide tissue image classification. *arXiv:1504.07947*, 2015.
- [46] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. In *IEEE CVPR*, 2018.
- [47] G. Huang, Z. Liu, K. Q. Weinberger, and L. van der Maaten. Densely connected convolutional networks. In *CVPR*, 2017.
- [48] J. Huang et al. Speed/accuracy trade-offs for modern convolutional object detectors. In *CVPR*, 2017.
- [49] S. Hussein, K. Cao, Q. Song, and U. Bagci. Risk stratification of lung nodules using 3d cnn-based multi-task learning. In *IPMI*, 2017.
- [50] S. Hwang and H.-E. Kim. Self-transfer learning for weakly supervised lesion localization. In *MICCAI*, pages 239–246. Springer, 2016.
- [51] B. Ibragimov and L. Xing. Segmentation of organs-at-risks in head and neck ct images using convolutional neural networks. *Medical physics*, 2017.
- [52] A. Isambert, F. Dhermain, F. Bidault, O. Commowick, P.-Y. Bondiaou, G. Malandain, and D. Lefkopoulos. Evaluation of an atlas-based automatic segmentation software for the delineation of brain organs at risk in a radiation therapy clinical context. *Radiotherapy and oncology*, 87(1):93–99, 2008.
- [53] C. Jacobs et al. Automatic detection of subsolid pulmonary nodules in thoracic computed tomography images. *Medical image analysis*, 2014.

- [54] A. Jesson et al. Cased: Curriculum adaptive sampling for extreme data imbalance. In *MICCAI*, 2017.
- [55] Z. Jiao, X. Gao, Y. Wang, and J. Li. A deep feature based framework for breast masses classification. *Neurocomputing*, 2016.
- [56] P. Kohli, P. H. Torr, et al. Robust higher order potentials for enforcing label consistency. *IJCV*, 82(3):302–324, 2009.
- [57] T. Kooi, G. Litjens, B. van Ginneken, A. Gubern-Mérida, C. I. Sánchez, R. Mann, A. den Heeten, and N. Karssemeijer. Large scale deep learning for computer aided detection of mammographic lesions. *Medical image analysis*, 35:303–312, 2017.
- [58] P. Krähenbühl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, pages 109–117, 2011.
- [59] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [60] J. R. Landis and G. G. Koch. The measurement of observer agreement for categorical data. *biometrics*, 1977.
- [61] C. Leavens, T. Vik, H. Schulz, S. Allaire, J. Kim, L. Dawson, B. O’Sullivan, S. Breen, D. Jaffray, and V. Pekar. Validation of automatic landmark identification for atlas-based segmentation for radiation treatment planning of the head-and-neck region. In *SPIE*, 2008.
- [62] F. Liao, M. Liang, Z. Li, X. Hu, and S. Song. Evaluate the malignancy of pulmonary nodules using the 3d deep leaky noisy-or network. *arXiv preprint arXiv:1711.08324*, 2017.
- [63] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *CVPR*, volume 1, page 4, 2017.
- [64] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar. Focal loss for dense object detection. In *CVPR*, 2017.
- [65] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015.
- [66] E. Lopez Torres et al. Large scale validation of the m5l lung cad on heterogeneous ct datasets. *Medical physics*, 2015.
- [67] A. L. Maas, A. Y. Hannun, and A. Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3, 2013.
- [68] I. C. Moreira, I. Amaral, I. Domingues, A. Cardoso, M. J. Cardoso, and J. S. Cardoso. Inbreast: toward a full-field digital mammographic database. *Academic radiology*, 2012.

- [69] K. Murphy et al. A large-scale evaluation of automatic pulmonary nodule detection in chest ct using local image features and k-nearest-neighbour classification. *Medical image analysis*, 2009.
- [70] S. Nikolov, S. Blackwell, R. Mendes, J. De Fauw, C. Meyer, C. Hughes, H. Askham, B. Romera-Paredes, A. Karthikesalingam, C. Chu, et al. Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy. *arXiv preprint arXiv:1809.04430*, 2018.
- [71] K. C. Oeffinger, E. T. Fontham, R. Etzioni, A. Herzig, J. S. Michaelson, Y.-C. T. Shih, L. C. Walter, T. R. Church, C. R. Flowers, S. J. LaMonte, et al. Breast cancer screening for women at average risk: 2015 guideline update from the american cancer society. *Jama*, 2015.
- [72] G. Papandreou et al. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *ICCV*, 2015.
- [73] G. V. Pednekar, J. K. Udupa, D. J. McLaughlin, X. Wu, Y. Tong, C. B. Simone, J. Camaratta, and D. A. Torigian. Image quality and segmentation. In *SPIE*, 2018.
- [74] A. A. Qazi, V. Pekar, J. Kim, J. Xie, S. L. Breen, and D. A. Jaffray. Auto-segmentation of normal and target structures in head and neck ct images: A feature-driven model-based approach. *Medical physics*, 2011.
- [75] P. F. Raudaschl, P. Zaffino, G. C. Sharp, M. F. Spadea, A. Chen, B. M. Dawant, T. Albrecht, T. Gass, C. Langguth, M. Lüthi, et al. Evaluation of segmentation methods on head and neck ct: Auto-segmentation challenge 2015. *Medical physics*, 2017.
- [76] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [77] X. Ren, L. Xiang, D. Nie, Y. Shao, H. Zhang, D. Shen, and Q. Wang. Interleaved 3d-cnn s for joint segmentation of small-volume structures in head and neck ct images. *Medical physics*, 2018.
- [78] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [79] S. S. M. Salehi, D. Erdogmus, and A. Gholipour. Tversky loss function for image segmentation using 3d fully convolutional deep networks. In *International Workshop on MLMI*, 2017.
- [80] A. A. A. Setio et al. Pulmonary nodule detection in ct images: false positive reduction using multi-view convolutional networks. *IEEE TMI*, 2016.
- [81] A. A. A. Setio et al. Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the luna16 challenge. *Medical Image Analysis*, 2017.

- [82] A. A. A. Setio et al. Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the luna16 challenge. *Medical image analysis*, 42:1–13, 2017.
- [83] G. Sharp, K. D. Fritscher, V. Pekar, M. Peroni, N. Shusharina, H. Veeraraghavan, and J. Yang. Vision 20/20: perspectives on automated image segmentation for radiotherapy. *Medical physics*, 41(5), 2014.
- [84] W. Shen, M. Zhou, F. Yang, D. Dong, C. Yang, Y. Zang, and J. Tian. Learning from experts: Developing transferable deep features for patient-level lung cancer prediction. In *MICCAI*, 2016.
- [85] W. Shen, M. Zhou, F. Yang, D. Dong, C. Yang, Y. Zang, and J. Tian. Learning from experts: Developing transferable deep features for patient-level lung cancer prediction. In *MICCAI*, 2016.
- [86] W. Shen, M. Zhou, F. Yang, C. Yang, and J. Tian. Multi-scale convolutional neural networks for lung nodule classification. In *IPMI*, 2015.
- [87] W. Shen, M. Zhou, F. Yang, D. Yu, D. Dong, C. Yang, Y. Zang, and J. Tian. Multi-crop convolutional neural networks for lung nodule malignancy suspiciousness classification. *Pattern Recognition*, 2017.
- [88] R. Sims, A. Isambert, V. Grégoire, F. Bidault, L. Fresco, J. Sage, J. Mills, J. Bourhis, D. Lefkopoulos, O. Commowick, et al. A pre-clinical assessment of an atlas-based automatic segmentation tool for the head and neck. *Radiotherapy and Oncology*, 2009.
- [89] C. H. Sudre, W. Li, T. Vercauteren, S. Ourselin, and M. J. Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. 2017.
- [90] K. Suzuki, F. Li, S. Sone, and K. Doi. Computer-aided diagnostic scheme for distinction between benign and malignant nodules in thoracic low-dose ct by use of massive training artificial neural network. *IEEE TMI*, 2005.
- [91] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *ICLR*, 2014.
- [92] C. Tam, X. Yang, S. Tian, X. Jiang, J. Beitler, and S. Li. Automated delineation of organs-at-risk in head and neck ct images using multi-output support vector regression. In *SPIE*, 2018.
- [93] H. Tang, D. Kim, and X. Xie. Automated pulmonary nodule detection using 3d deep convolutional neural networks. *ISBI*, 2018.
- [94] T. Tieleman and G. Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31, 2012.

- [95] Y. Tong, J. K. Udupa, X. Wu, D. Odhner, G. Pednekar, C. B. Simone, D. McLaughlin, C. Apinorasethkul, G. Shammo, P. James, et al. Hierarchical model-based object localization for auto-contouring in head and neck radiation therapy planning. In *SPIE*, 2018.
- [96] L. A. Torre, F. Bray, R. L. Siegel, J. Ferlay, J. Lortet-Tieulent, and A. Jemal. Global cancer statistics, 2012. *CA: a cancer journal for clinicians*, 2015.
- [97] M. Vallières, E. Kay-Rivest, L. J. Perrin, X. Liem, C. Furstoss, H. J. Aerts, N. Khaouam, P. F. Nguyen-Tan, C.-S. Wang, K. Sultanem, et al. Radiomics strategies for risk assessment of tumour failure in head-and-neck cancer. *Scientific reports*, 7(1):10117, 2017.
- [98] C. Varela, S. Timp, and N. Karssemeijer. Use of border information in the classification of mammographic masses. *Physics in Medicine and Biology*, 51(2):425, 2006.
- [99] R. F. Verhaart, V. Fortunati, G. M. Verduijn, A. Lugt, T. Walsum, J. F. Veenland, and M. M. Paulides. The relevance of mri for patient modeling in head and neck hyperthermia treatment planning: A comparison of ct and ct-mri based tissue segmentation on simulated temperature. *Medical physics*, 41(12), 2014.
- [100] P. W. Voet, M. L. Dirkx, D. N. Teguh, M. S. Hoogeman, P. C. Levendag, and B. J. Heijmen. Does atlas-based autosegmentation of neck levels require subsequent manual contour editing to avoid risk of severe target underdosage? a dosimetric analysis. *Radiotherapy and Oncology*, 98(3):373–377, 2011.
- [101] C. Wachinger, K. Fritscher, G. Sharp, and P. Golland. Contour-driven atlas-based segmentation. *IEEE transactions on medical imaging*, 34(12):2492–2505, 2015.
- [102] X. Wang et al. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *CVPR*, 2017.
- [103] Z. Wang et al. Exploring fisher vector and deep networks for action spotting. In *CVPRW*, 2015.
- [104] Z. Wang et al. Weakly supervised patchnets: Describing and aggregating local patches for scene recognition. *IEEE TIP*, 2017.
- [105] Z. Wang et al. Structured triplets learning with pos-tag guided attention for visual question answering. In *WACV*, 2018.
- [106] Z. Wang, L. Wei, L. Wang, Y. Gao, W. Chen, and D. Shen. Hierarchical vertex regression-based segmentation of head and neck ct images for radiotherapy planning. *IEEE TIP*, 2018.
- [107] T. W. Way, L. M. Hadjiiski, B. Sahiner, H.-P. Chan, P. N. Cascade, E. A. Kazerooni, N. Bogot, and C. Zhou. Computer-aided diagnosis of pulmonary nodules on ct scans: Segmentation and classification using 3d active contours. *Medical Physics*, 2006.

- [108] X. Wu, J. K. Udupa, Y. Tong, D. Odhner, G. V. Pednekar, C. B. Simone, D. McLaughlin, C. Apinorasethkul, J. Lukens, D. Mihailidis, et al. Auto-contouring via automatic anatomy recognition of organs at risk in head and neck cancer on ct images. In *SPIE*, 2018.
- [109] W. Xiong, L. Wu, F. Alleva, J. Droppo, X. Huang, and A. Stolcke. The microsoft 2017 conversational speech recognition system. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5934–5938. IEEE, 2018.
- [110] H. Xu, A. Arsene Henry, M. Robillard, M. Amessis, and Y. M. Kirova. The use of new delineation tool mirada at the level of regional lymph nodes, step-by-step development and first results for early-stage breast cancer patients. *The British journal of radiology*, 2018.
- [111] X. Yan, J. Pang, H. Qi, Y. Zhu, C. Bai, X. Geng, M. Liu, D. Terzopoulos, and X. Ding. Classification of lung nodule malignancy risk on computed tomography images using convolutional neural network: A comparison between 2d and 3d strategies. In *ACCV*, 2016.
- [112] Z. Yan, Y. Zhan, Z. Peng, S. Liao, Y. Shinagawa, S. Zhang, D. N. Metaxas, and X. S. Zhou. Multi-instance deep learning: Discover discriminative local anatomies for bodypart recognition. *IEEE transactions on medical imaging*, 35(5):1332–1343, 2016.
- [113] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus. Deconvolutional networks. In *CVPR*, pages 2528–2535. IEEE, 2010.
- [114] T. Zhang, Y. Chi, E. Meldolesi, and D. Yan. Automatic delineation of on-line head-and-neck computed tomography images: toward on-line adaptive radiotherapy. *International Journal of Radiation Oncology* Biology* Physics*, 68(2):522–530, 2007.
- [115] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. Torr. Conditional random fields as recurrent neural networks. In *ICCV*, pages 1529–1537, 2015.
- [116] W. Zhu, Y. Huang, L. Zeng, X. Chen, Y. Liu, Z. Qian, N. Du, W. Fan, and X. Xie. Anatomynet: Deep learning for fast and fully automated whole-volume segmentation of head and neck anatomy. *Medical physics*, 2018.
- [117] W. Zhu, C. Lan, J. Xing, W. Zeng, Y. Li, L. Shen, and X. Xie. Co-occurrence feature learning for skeleton based action recognition using regularized deep lstm networks. In *AAAI*, 2016.
- [118] W. Zhu, C. Liu, W. Fan, and X. Xie. Deeplung: Deep 3d dual path nets for automated pulmonary nodule detection and classification. In *IEEE WACV*, 2018.
- [119] W. Zhu, Q. Lou, Y. S. Vang, and X. Xie. Deep multi-instance networks with sparse label assignment for whole mammogram classification. In *MICCAI*, 2017.

- [120] W. Zhu, Y. S. Vang, Y. Huang, and X. Xie. Deepem: Deep 3d convnets with em for weakly supervised pulmonary nodule detection. *MICCAI*, 2018.
- [121] W. Zhu, X. Xiang, T. D. T. Tran, G. D. H. Hager, and X. Xie. Adversarial deep structured nets for mass segmentation from mammograms. *IEEE International Symposium on Biomedical Imaging*, 2018.
- [122] W. Zhu and X. Xie. Adversarial deep structural networks for mammographic mass segmentation. *arXiv:1612.05970*, 2016.