Preface

This book was partially motivated by the recent rapid progress on deep convolutional and recurrent neural network models and the abundance of important applications in computer vision, where quantitative performance has significantly improved in object recognition, detection, and automatic image caption. However publicly available image database with generally well-annotated image or object labels (but with some labeling noise), such as ImageNet (1.2 million images), PASCAL Visual Object Classes (VOC) Dataset and Challenge (VOC) (~12,000 images), and Microsoft Common Objects in Context (COCO; ~300,000 images), have been the essential resource to fuel the learning process of the deep neural networks that work well in practical and challenging scenarios.

These types of large-scale annotated image datasets are, unfortunately, not available in medical image analysis tasks yet, even though improving medical imaging applications through deep neural network models and imaging data is highly valuable. This is partly explained by the fact that labels or annotations for medical image database are much harder or expensive to obtain. The general practice of collecting labels for ImageNet, as for example using Google image search engine for pre-selection, followed by manual label refinement through crowd-sourcing (e.g., Amazon Mechanical Turk), is largely nonfeasible due to the formidable difficulties of medical annotation tasks for those who are not clinically trained.

Indeed, employing deep neural networks, especially convolutional neural network models, requires a large amount of annotated training instances. This concern was reflected in that only four papers (out of 250 total publications) in MICCAI 2014, the 17th International Conference on Medical Image Computing and Computer Assisted Intervention, were based on deep learning models while at least 20% of papers at IEEE Conference on CVPR 2014 were related to deep neural networks. Even after that, this situation has drastically changed. We have had nearly 10% of the publications (23 papers) in MICCAI 2015 that are built upon deep neural network models for a variety of medical imaging problems: fetal ultrasound standard plane detection, vertebrae localization and identification, multi-view mammogram analysis, mass segmentation, glaucoma detection, nucleus localization in microscopy images, lymph node detection and segmentation, organ segmentation in CT/MRI scans, coronary calcium scoring in cardiac CT angiography, etc. We certainly predict this uprising trend will continue for MICCAI 2016.

To answer the question of how to learn powerful and effective deep neural networks with often sub-hundred of patient scans (where significant quantitative performance gains have been reported), many chapters in this book will precisely address this promising trend by describing detailed technical contents on data preparation, network designs, and evaluation strategies. Moreover, what can be learned from this early success and how to move forward rapidly are the other two main topics to be discussed in details for this book.

Overview and Goals

Deep learning, in particular Convolutional Neural Networks (CNN), is a validated image representation and classification technique for medical image analysis and applications. We have observed many encouraging work that report new and newer state-of-the-art performance on quite challenging problems in this domain. The main reason behind this stream of work, we believe, is that effective task-dependent image features can be directly or intrinsically learned through the hierarchy of convolutional kernels inside CNN. Hand-crafted image feature engineering research was a relatively weak subtopic in medical image analysis, compared to the extensive evaluation and studies on image features for computer vision. This sometimes limits the generality and applicability of well-tested natural image feature descriptors, such as SIFT (scale-invariant feature transform), HOG (histogram of oriented gradients) and others into medical imaging tasks.

On the other hand, CNN models have been proved to have much higher modeling capacity, compared to the previous image recognition mainstream pipelines, e.g., HAAR, SIFT, HOG image features followed by spatial feature encoding, then random forest or support vector classifiers. Given millions of parameters to fit during model training (much more than previous pipelines), CNN representation empowers and enables computerized image recognition models, with a good possibility to be able to handle more challenging imaging problems. The primary risk is overfitting since model capacity is generally high in deep learning but often very limited datasets are available (that are with good quality of labels to facilitate supervised training). The core topics of this book are represented by examples on how to address this task-critical overfitting issue with deep learning model selection, dataset resampling and balancing, and the proper quantitative evaluation protocols or setups.

Furthermore, with deep neural networks (especially CNNs) as *De facto* building blocks for medical imaging applications (just as previous waves of Boosting, Random Forest), we would argue that it is of course important to use them to improve existing problems, which has been widely studied before, but more critically, it is the time to consider exploiting new problems and new experimental, clinical protocols that will foster the development of preventative and precision medicine tools in imaging to impact modern clinical practices. Without loss of generality, we give the following three examples: *Holistic CT slice based interstitial lung disease (ILD) prediction* via deep multi-label regression and unordered pooling (significantly improving the current status of the mainstream image patch based ILD classification approaches with several built-in prerequisites that nevertheless prevents clinically desirable diagnosis protocols for ILD pre-screening from ultra-low dose CT scans to be a reality); *Precise deep organ/tumor segmentation based volumetric measurements as new imaging bio-markers* (remarkably enabling to provide more precise imaging measurement information to better assist physicians than the current popular RECIST metric, for high-fidelity patient screening, tumor management and patient-adaptive therapy); and *Unsupervised category discovery and joint text-image deep mining using large-scale radiology image database* (opening the door to compute, extract, and mine meaningful clinical imaging semantics via modern hospitals' PACS/RIS database systems that could include millions of patient imaging and text report cases, overcoming the limitation of lack of strong annotations).

Lastly, from our perspective, this is just the beginning of embracing and employing new deep neural network models and representations for many medical image analysis and medical imaging applications. We hope this book will help to get you more prepared and ready to exploit old and new problems. Happy reading and learning!

Organization and Features

This book covers a range of topics from reviews of the recent state-of-the-art progresses, to deep learning for semantic object detection, segmentation and large-scale radiology database mining. In the following, we give a brief overview of the contents of the book.

Chapter 1 describes a brief review of nearly 20 years of research by Dr. Ronald M. Summers (MD/Ph.D.) in medical imaging based computer-aided diagnosis (CAD) where he won the presidential early career award for scientists and engineers in 2000 and his personal take and insights on the recent development of deep learning techniques for medical image interpretation problems?

Chapter 2 lists a relatively comprehensive review of the recent methodological progress and related literature of deep learning for medical image analysis, in the topics of abdominal, chest and cardiology imaging, histopathology cell imaging, and chest X-ray and mammography.

Chapters 3 to 10 cover all various topics using deep learning for object or landmark detection tasks in 2D and 3D medical imaging. Particularly, we present a random view resampling and integration approach for three CAD problems (Chap. 3); 3D volumetric deep neural networks for efficient and robust landmark detection (Chap. 4); a restricted views based pulmonary embolism detection (Chap.5); a new method on cell detection (Chap.6); tumor cell anaplasia and multi-nucleation detection (Chap. 7) in microscopy images; Predicting interstitial lung diseases and segmentation label propagation (Chap. 8); an in-depth study on CNN architectures, dataset characteristics and transfer learning for CAD problems (Chap. 9); followed by a computationally scalable method of accelerated cell detection with sparse convolution kernels in histopathology images (Chap. 10).

Chapters 11 to 16 discuss several representative works in semantic segmentation using deep learning principles in medical imaging. Specifically, this book describes automatic carotid intima-media thickness analysis (Chap. 11); deep segmentation via distance regularized level set and deep-structured inference (Chap. 12); structured prediction for segmenting masses in mammograms (Chap. 13); pathological kidney segmentation in CT via local versus global image context (Chap. 14); skeletal muscle cell segmentation in Microscopy (Chap. 15); and a bottom-up approach for deep pancreas segmentation in contrasted CT images (Chap. 16).

Chapter 17 discusses a novel work on interleaved text/image deep mining on a large-scale radiology image database for automated image interpretation which was not technically feasible in the pre-deep learning era. Finally, we would like to point out the way we organized this book in roughly three blocks of detection, segmentation, and text-image joint learning aligned with the status how computer vision is progressing with deep learning. These three blocks of research topics are probably essential for imaging understanding and interpretation of all imaging problems.

Target Audience

The intended reader of this book is a professional or a graduate student who is able to apply computer science and math principles into problem solving practices. It may be necessary to have some level of familiarity with a number of more advanced subjects: image formation, processing and understanding, computer vision, machine learning, and statistical learning.

Bethesda, USA — Le Lu
Princeton, USA — Yefeng Zheng
Adelaide, Australia — Gustavo Carneiro
Gainesville, USA — Lin Yang

http://www.springer.com/978-3-319-42998-4

Deep Learning and Convolutional Neural Networks for
Medical Image Computing
Precision Medicine, High Performance and Large-Scale
Datasets
Le, L.; Zheng, Y.; Carneiro, G.; Yang, L. (Eds.)
2017, XIII, 326 p. 117 illus., 100 illus. in color.,
Hardcover
ISBN: 978-3-319-42998-4