DeepCADx : Automated Prostate Cancer Detection and Diagnosis in mp-MRI based on Multimodal Convolutional Neural Networks

Zhiwei Wang, Chaoyue Liu, Xiang Bai, Xin Yang*
School of Electronic Information and Communications, Huazhong University of Science and Technology
Wuhan, Hubei, China 430074
{zhiweiwang,chaoyueliu,xbai,xinyang2014}@hust.edu.cn

ABSTRACT

In this paper, we present *DeepCADx*, a computer-aided prostate detection and diagnosis (CADx) system powered by a novel deep convolutional neural networks (CNNs). Specifically, the developed DeepCADx system processes multi-parametric magnetic resonance imaging (mp-MRI) sequences in three major steps: 1) pre-processing which registers images from different modalities and detect prostates, 2) multimodal CNNs which jointly identifies images containing prostate cancers (PCa) and generate cancer response maps (CRM) with each pixel indicating the probability to be cancerous, and 3) post-processing which localize lesion in CRMs and assess the aggressiveness (i.e. Gleason score) of each localized lesion using multimodal CNN features and a 5-class SVM classifier.

CCS CONCEPTS

• **Computing methodologies** → **Object detection**; **Neural networks**;

KEYWORDS

Prostate Cancer Detection and Diagnosis, Multimodal CNNs, Multimodal Fusion

1 INTRODUCTION

PCa detection and diagnosis in mp-MRI is an active research area due to high incidence rate and lethal rate of PCa. Interpreting mp-MRI sequences manually by radiologists requires substantial expertise and efforts, and usually suffers from large inter-/intra-observer variations, and low sensitivity and specificity [4]. Therefore, automated and accurate PCa detection and diagnosis in mp-MRI is of high demand for alleviating requirement for expertise in radiology reading, reducing risk of over-/under-treatment, and enabling large-scale PCa screening. Most of existing CADx systems [1, 2], if not all, rely on two-stage classification, i.e. voxel-level classification for candidate generation and region-level classification for verification and aggressiveness assessment. Such two-stage design has a consequence that miss detection of lesions in voxel-level classification could propagate to the subsequent step, degrading the overall detection and diagnosis accuracy.

In this work, we introduce *DeepCADx*, an automated PCa detection and diagnosis system based on multimodal deep CNNs. The DeepCADx can concurrently identify the presence of PCa in an image, localize lesions and predicts the Gleason Score (GS). In particular, the key contributions of our DeepCADx includes:

- **Weakly-supervised Learning for CRM Generation:** We developed a weakly-supervised CNN which is trained for distinguishing between images containing cancers and those not containing cancers, yet can realize the PCa localization task by generating a cancer response map (CRM) with each pixel indicating the likelihood to be cancerous. Our weakly-supervised CNN can effectively learn discriminative visual patterns of lesions from prostate clutters without any prior knowledge of lesions' locations, and thus is much more convenient for clinical usage in which manual annotation of lesions' regions are quite expensive.
- **In-depth Multimodal Fusion:** We implicitly fused multimodal information of mp-MRI in our CNN network by designing a new inconsistency loss function. The inconsistency loss enforces consistent CRMs for different modalities, i.e. Apparent Diffusion Coefficients (ADC), and T2 weighted images (T2w), yielding more consistent and representative PCa-relevant features.

Once the prostate is identified as cancerous, points with local maximum values in the CRM are considered as lesions' locations and features extracted by the multimodal CNNs are further used for GS classification (i.e. GS$\leq$ 6, GS = 3 + 4, 4 + 3, 8, and 9) by a 5-class SVM classifier.

2 OUR SYSTEM

As shown in Fig. 1, the DeepCADx consists of three key components: pre-processing for multimodal registration and prostate detection, multimodal CNNs for joint PCa classifcation and CRM generation and post-processing for PCa localization and GS grading.

2.1 Pre-processing

1) Registration of ADC and T2w Images. We first register ADC and T2w sequences so as to remove prostate distortion between the two sequences arising from patients' movement during MRI acquisition. We applied 2D affine registration method based on mutual information to the corresponding ADC and T2w axial slices due to its sufficient accuracy and efficiency.

2) Prostate Detection. To automatically detect prostate region, we train a simple CNN regressor by providing a list of T2w slices

*Corresponding author. This work was supported by the National Natural Science Foundation of China grant 61502188 and Wuhan Science and Technology Bureau award 2017010201010111

MM'17, October 23–27, 2017, Mountain View, CA, USA.
 ISBN 978-1-4503-4906-2/17/10.
DOI: https://doi.org/10.1145/3123266.3127914

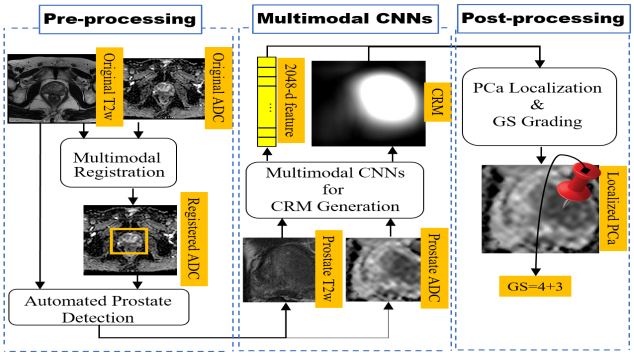

Figure 1: The framework of the DeepCADx.

associated with the corresponding manually labeled squared bounding boxes denoting the prostate regions. The CNN is then trained to regress the height- and width-normalized center coordinates and length of the bounding box encompassing the entire prostate region. In testing, the regressed square by the CNN in the T2w slice is then propagated to the corresponding registered ADC slices.

2.2 Multimodal CNNs

1) Weakly-Supervised Learning for Generating CRM of single MRI Modality. We modified GoogLeNet by replacing its last fully-connected layers with a fully-convolutional layer to explicitly generate a single feature map M. Positions with large response values on the feature map M indicate spatial locations most relevant to each category (i.e. cancer or normal tissues) and hence the feature map M can also considered be used directly as a CRM. We then apply the global average pooling operation to M to obtain a single slice-level score, which is further projected to a probability $p \in [0, 1]$ by the sigmoid function as the final output of the CNN, indicating an estimation of likelihood whether a 2D MRI slice contains PCa or not. Given slice-level labels (i.e., 0 for not containing PCa and 1 for containing PCa) for training the CNN, we treat the classification task as a logistic regression problem, whose loss function is therefore a cross entropy error function.

Since the prediction p is obtained by directly averaging all entries of M, when conducting the back propagation (BP) algorithm, CNN weights will be updated for both decreasing intensities of entries in M for normal slices and increasing intensities of entries in M whose receptive fields are discriminative visual patterns (mainly PCa lesions) for slices containing PCa.

2) Fusion of ADC and T2w Sequences. Multimodal CNNs for multimodal fusion could be highly beneficial because lesions appear at the same locations on T2w and ADC slices, which implies that the cancer response map of an ADC slice (i.e. M_{ADC}) should be consistent with that of a T2w slice (i.e.M_{T2w}). However, training two independent CNN models for ADC and T2w images respectively (i.e. CNN_{ADC} and CNN_{T2w}) cannot guarantee consistency between the two modalities since with weakly-supervised learning, CNN 'sees' not only PCa-relevant patterns but also irrelevant visual patterns when distinguishing slices containing PCa from the normal ones. To address this problem, we enforce the CNN models of ADC and T2w to generate similar CRMs. To encode the enforcement into our CNN learning, we define a new loss called

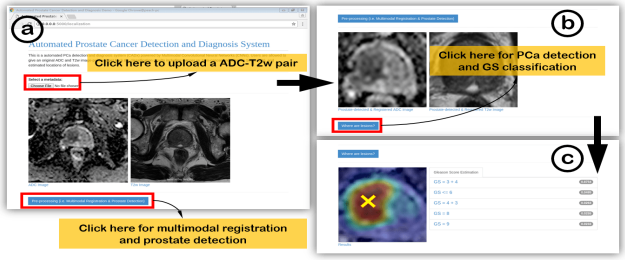

Figure 2: The Web-based UI of the DeepCADx.

inconsistency loss, representing the differences between M_{ADC} and M_{T2w} as expressed in Eq. (1):

$$\ell(M_{ADC}, M_{T2w}) = \frac{1}{N} \|\sigma(M_{ADC}) - \sigma(M_{T2w})\|^2 \quad (1)$$

where $\sigma(\cdot)$ is the sigmoid function and $\|\sigma(M_{ADC}) - \sigma(M_{T2w})\|$ calculates the euclidean distance between the M_{ADC} and M_{T2w} and N is total number of pixels in an image.

2.3 Post-processing

1) PCa Localization. The localization procedure proceeds as follows. We first find a local maximum point with highest intensity in each connected component on M_{ADC} as a candidate. Then, we calculate an adaptive threshold from M_{ADC} using the Otsu's algorithm [3]. Candidates whose intensity values are below the threshold are excluded as false alarms. The remaining points are the locations of lesions for biopsy guidance in clinical practice.

2) GS Grading. For localized lesions, we train a 5-class SVM classifier using RBF kernel to classify them into five categories: GS ≤ 6, GS $= 3 + 4$, GS $= 4 + 3$, GS $= 8$, and GS $= 9$. Specifically, for each category, we randomly choose 5 lesions for training. For each lesion, five 2D axial slices containing the lesion are used, cropped, resized to 224×224 and input to the multimodal CNNs. For each slice, a 2048 high-dimensional feature vector is extracted.

3 WEB-BASED USER INTERFACE

We deploy the DeepCADx on a server and allow users to upload an original ADC-T2w slice pair and get back the prostate-detected registered ADC and T2w images as well as locations of lesions and their GS scores. Fig. 2 illustrates the web-based user interface of the DeepCADx. Images in Fig. 2(a) are the original ADC-T2w slice pair, and images in Fig. 2(b) are the prostate images extracted by the pre-processing, and the image in Fig. 2(c) is the ADC slice overlaid with the corresponding CRM and localized lesions.

REFERENCES

[1] Guillaume Lemaitre. 2016. *Computer-Aided Diagnosis for Prostate Cancer using Multi-Parametric Magnetic Resonance Imaging.* Ph.D. Dissertation. Universite de Bourgogne; Universitat de Girona.

[2] Geert Litjens, Oscar Debats, Jelle Barentsz, Nico Karssemeijer, and Henkjan Huisman. 2014. Computer-aided detection of prostate cancer in MRI. *IEEE transactions on medical imaging* 33, 5 (2014), 1083–1092.

[3] Nobuyuki Otsu. 1975. A threshold selection method from gray-level histograms. *Automatica* 11, 285-296 (1975), 23–27.

[4] Fritz H Schröder, Jonas Hugosson, Monique J Roobol, Teuvo LJ Tammela, Stefano Ciatto, Vera Nelen, Maciej Kwiatkowski, Marcos Lujan, Hans Lilja, Marco Zappa, et al. 2009. Screening and prostate-cancer mortality in a randomized European study. *New England Journal of Medicine* 360, 13 (2009), 1320–1328.