Three-Dimensional CT Image Segmentation by Combining 2D Fully Convolutional Network with 3D Majority Voting

Xiangrong Zhou[1(✉)], Takaaki Ito[1], Ryosuke Takayama[1],
Song Wang[2], Takeshi Hara[1], and Hiroshi Fujita[1]

[1] Department of Intelligent Image Information, Graduate School of Medicine,
Gifu University, Gifu 501-1194, Japan
`zxr@fjt.info.gifu-u.ac.jp`
[2] Department of Computer Science and Engineering,
University of South Carolina, Columbia SC 29208, USA

Abstract. We propose a novel approach for automatic segmentation of anatomical structures on 3D CT images by voting from a fully convolutional network (FCN), which accomplishes an end-to-end, voxel-wise multiple-class classification to map each voxel in a CT image directly to an anatomical label. The proposed method simplifies the segmentation of the anatomical structures (including multiple organs) in a CT image (generally in 3D) to majority voting for the semantic segmentation of multiple 2D slices drawn from different viewpoints with redundancy. An FCN consisting of "convolution" and "de-convolution" parts is trained and re-used for the 2D semantic image segmentation of different slices of CT scans. All of the procedures are integrated into a simple and compact all-in-one network, which can segment complicated structures on differently sized CT images that cover arbitrary CT scan regions without any adjustment. We applied the proposed method to segment a wide range of anatomical structures that consisted of 19 types of targets in the human torso, including all the major organs. A database consisting of 240 3D CT scans and a humanly annotated ground truth was used for training and testing. The results showed that the target regions for the entire set of CT test scans were segmented with acceptable accuracies (89 % of total voxels were labeled correctly) against the human annotations. The experimental results showed better efficiency, generality, and flexibility of this end-to-end learning approach on CT image segmentations comparing to conventional methods guided by human expertise.

Keywords: CT images · Anatomical structure segmentation · Fully convolutional network (FCN) · 3D majority voting · End-to-end learning

1 Introduction

Three-dimensional (3D) computerized tomography (CT) images are important resources that provide useful internal information about the human body to support diagnosis, surgery, and therapy [1]. Fully automatic image segmentation is a fundamental part of

the applications based on 3D CT images by mapping the physical image signal to a useful abstraction. Conventional approaches to CT image segmentation usually try to transfer human knowledge directly to a processing pipeline, including numerous hand-crafted signal processing algorithms and image features [2–5]. In order to further improve the accuracy and robustness of image segmentation, we need to be able to handle a larger variety of ambiguous image appearances, shapes, and relationships of anatomical structures. It is difficult to achieve this goal by defining and considering human knowledge and rules explicitly. Instead, a data-drive approach using big image data—such as a deep convolutional neural network (deep CNN)—is expected to be better for solving this segmentation problem.

Recently, several studies were reported that applied deep CNNs to medical image analysis. Many of these used deep CNNs for lesion detection or classification [6, 7]. Studies of this type usually divide CT images into numerous small 2D/3D patches at different locations, and then classify these patches into multiple pre-defined categories. Deep CNNs are used to learn a set of optimized image features (sometimes combined with a classifier) to achieve the best classification rate for these image patches. Similarly, deep CNNs have also been embedded into conventional organ-segmentation processes to reduce the FPs in the segmentation results or to predict the likelihoods of the image patches [8–10]. However, the anatomical segmentation of CT images over a wide region of the human body is still challenging because of the image appearance similarities between different structures, as well as the difficulty of ensuring global spatial consistency in the labeling of patches in different CT cases.

This paper proposes a novel approach based on deep CNNs that naturally imitate the thought processes of radiologists during CT image interpretation for image segmentation. Our approach models CT image segmentation in a way that can best be described as "multiple 2D proposals with a 3D integration." This is very similar to the way that a radiologist interprets a CT scan as many 2D sections, and then reconstructs the 3D anatomical structure as a mental image. Unlike previous work on medical image segmentation that labels each voxel/pixel by a classification based on its neighborhood information (i.e., either an image patch or a "super-pixel") [8–10], our work uses rich information from the entire 2D section to directly predict complex structures (multiple labels on images). Furthermore, the proposed approach is based on an end-to-end learning without using any conventional image-processing algorithms such as smoothing, filtering, and level-set methods.

2 Methods

2.1 Overview

As shown in Fig. 1, the input is a 3D CT case (the method can also handle a 2D case, which can be treated as a degenerate 3D case), and the output is a label map of the same size and dimension, in which the labels are a pre-defined set of anatomical structures. Our segmentation process is repeated to sample 2D sections from the CT case, pass them to a fully conventional network (FCN) [11] for 2D image segmentation, and stack the 2D labeled results back into 3D. Finally, the anatomical structure label at each

Fig. 1. Pipeline of proposed anatomical structure segmentation for 3D CT scan. See Fig. 2 for the details of FCN structure.

voxel is decided based on majority voting at the voxel. The core part of our segmentation is an FCN that is used for the anatomical segmentation of the 2D sections. This FCN is trained based on a set of CT cases, with the human annotations as the ground truth. All of the processing steps of our CT image segmentation are integrated into an all-in-one network under a simple architecture with a global optimization.

2.2   3D-to-2D Image Sampling and 2D-to-3D Label Voting

In the proposed approach, we decompose a CT case (a 3D matrix, in general) into numerous sections (2D matrices) with different orientations, segment each 2D section, and finally, assemble the outputs of the segmentation (labeled 2D maps) back into 3D. Specifically, each voxel in a CT case (a 3D matrix) can lie on different 2D sections that pass through the voxel with different orientations. Our idea is to use the rich image information of the entire 2D section to predict the anatomical label of this voxel, and to increase the robustness and accuracy by redundantly labeling this voxel on multiple 2D sections with different orientations. In this work, we select all the 2D sections in three orthogonal directions (axial, sagittal, and coronal-body); this ensures that each voxel in a 3D case is located on three 2D CT sections.

After the 2D image segmentation, each voxel is redundantly annotated three times from these three 2D CT sections. The annotated results for each voxel should ideally be identical, but may be different in practice because of mislabeling during the 2D image segmentation. A label fusion by majority voting for the three labels is then introduced to improve the stability and accuracy of the final decision. Furthermore, a *prior* for each organ type (label) is estimated by calculating voxel appearance frequency of the organ region within total image based on training samples. In the case of no consensus between three labels during the majority voting process, our method simply selects the label with the biggest *prior* as the output.

2.3   FCN-Based 2D Image Segmentation via Convolution and de-Convolution Networks

We use an FCN for semantic segmentation in each 2D CT slice by labeling each pixel. Convolutional networks are constructed using a series of connected basic components (convolution, pooling, and activation functions) with translation invariance that depends only on the relative spatial coordinates. Each component acts as a nonlinear filter that operates (e.g., by matrix multiplication for convolution or maximum pooling) on the local input image, and the whole network computes a general nonlinear transformation from the input image. These features of the convolutional network provide the capability to adapt naturally to an input image of any size and any scan range of the human body, producing an output with the corresponding spatial dimensions.

Our convolutional network is based on the VGG16 net structure (16 layers of 3×3 convolution interleaved with maximum pooling plus 3 fully connected layers) [12], but with a change in the VGG16 architecture by replacing its fully connected layers (FC6 and 7 in Fig. 2) with convolutional layers (Conv 6 and 7 in Fig. 2). Its final fully connected classifier layer (FC 8 in Fig. 2) is then changed to a 1×1 convolution layer (Conv 8 in Fig. 2) whose channel dimension is fixed at the number of labels (the total number of segmentation targets was 20 in this work, including the background). This network is further expanded by docking a de-convolution network (the right-hand side in Fig. 2). Here, we use idea of the de-convolution in [11], and reinforce the network structure by adding five de-convolution layers, each of which consists of up-sampling, convolution, and crop (summation) layers as shown in Fig. 2.

FCN training: The proposed network (both convolution and de-convolution layers) is trained with numerous CT cases of humanly annotated anatomical structures. All of the 2D CT sections (corresponding to the label maps) along the three body orientations are shuffled, and used to train the FCN. The training process repeats feed-forward computation and back-propagation to minimize the loss function, which is defined as the sum of the pixel-wise losses between the network prediction and the label map annotated by the human experts. The gradients of the loss are propagated from the end to the start of the network, and the method of stochastic gradient descent with momentum is used to refine the parameters of each layer.

The FCN is trained sequentially by adding de-convolution layers [11]. To begin with, a coarse prediction (by a 32-pixel stride) is trained for the modified VGG16 network with one de-convolution layer (called FCN32s). A finer training is then added after adding one further de-convolution layer at the end of the network. This is done by using skips that combine the final prediction layer with a lower layer with a finer stride in the modified VGG16 network. This fine-training is repeated with the growth of the network layers to build FCN16s, 8s, 4s, and 2s which are trained from the predictions of 16, 8, 4, 2 strides on the CT images, respectively. The output of FCN 2s acts as the 2D segmentation result.

2D CT segmentation using trained FCN: The density resolution of the CT images is reduced from 12 to 8 bits using linear interpolation. The trained FCN is then applied to each 2D section independently, and each pixel is labeled automatically. The labels from

Fig. 2. Semantic image segmentation of 2D CT slice using fully convolutional network (FCN) [11]. Conv: convolution, Deconv: deconvolution, and FC: fully connected.

each 2D section are then projected back to their original 3D locations for the final vote-based labeling, as described above.

3   Experiment and Results

Our experiment used a CT image database that was produced and shared by a research project entitled "Computational Anatomy [13]". This database included 640 3D volumetric CT scans from 200 patients at Tokushima University Hospital. The anatomical ground truth (a maximum of 19 labels that included Heart, right/left Lung, Aorta, Esophagus, Liver, Gallbladder, Stomach and Duodenum (lumen and contents), Spleen,

left/right Kidney, Inferior Vein Cava, region of Portal Vein, Splenic Vein, and Superior Mesenteric Vein, Pancreas, Uterus, Prostate, and Bladder) in 240 CT scans was also distributed with the database. Our experimental study used all of the 240 ground-truth CT scans, comprising 89 torso, 17 chest, 114 abdomen, and 20 abdomen-with-pelvis scans. Furthermore, our research work was conducted with the approval of the Institutional Review Boards at Gifu and Tokushima Universities.

We picked 10 CT scans at random as the test samples, using the remaining 230 CT scans for training. As previously mentioned, we took 2D sections along the axial, sagittal, and coronal body directions. For the training samples, we obtained a dataset of 84,823 2D images with different sizes (width: 512 pixels; height: 80–1141 pixels). We trained a single FCN based on the ground-truth labels of the 19 target regions. Stochastic gradient descent (SGD) with momentum was used for the optimization. A mini-batch size of 20 images, learning rate of 10^{-4}, momentum of 0.9, and weight decay of 2^{-4} were used as the training parameters. All the 2D images were used directly as the inputs for FCN training, without any patch sampling.

We tested the proposed FCN network (Fig. 1) using 10 CT cases that were not used in the FCN training. An example of the segmentation result for a 3D CT case covering the human torso is shown in Fig. 3. The accuracy of the segmentation was evaluated per organ type and per image. We measured the intersection over union (IU) (also known as the Jaccard similarity coefficient) between the segmentation result and the ground truth. Because each CT case may contain different anatomical structures—with the information about these unknown before the segmentation—we performed a comprehensive evaluation of multiple segmentation results for all the images in the test dataset by considering the variance of the target numbers and volume. Two measures (voxel accuracy: true positive for multiple label prediction on all voxels in a CT case; frequency-weighted IU: mean value of IUs that normalized by target volumes and numbers in a CT case [11]) were employed for the evaluations. The evaluation results


Fig. 3. Left: example of segmentation in 3D CT case, with segmented regions labeled with different colors for one 2D CT slice and 3D visualization based on surface-rendering method. Right: corresponding ground truth segmentation.

for the voxel accuracy, frequency-weighted IU were 0.89 and 0.84, respectively, when averaged over all the segmentation results of the test dataset. These results show that 89 % of the voxels within the anatomical structures (constructed using multiple target regions) were labeled correctly, with a mean overlap ratio of 84 % for 19 target regions in the test dataset. The mean IU values in each organ type are listed in Table 1 for both training and test data.

Table 1. Accuracy evaluations in terms of mean value of IUs per target type between segmentation results of FCN-8s and ground truth in 230 training and 10 test CT scans after voting in 3D [14].

Target name	Mean value of IUs	
	Training samples (230)	Test samples (10)
Right Lung	0.92	0.87
Left Lung	0.91	0.88
Heart	0.87	0.87
Aorta	0.72	0.63
Esophagus	0.18	0.27
Liver	0.91	0.91
Gallbladder	0.58	0.48
Stomach and Duodenum (2nd pos.)	0.48	0.43
Stomach and Duodenum Lumen	0.59	0.61
Contents inside of Stomach and Duodenum	0.21	0.10
Spleen	0.85	0.86
Right Kidney	0.85	0.86
Left Kidney	0.85	0.84
Inferior Vena Cava	0.56	0.51
Portal Vein, Splenic Vein, and Superior Mesenteric Vein	0.32	0.03
Pancreas	0.48	0.45
Uterus	0.23	0.09
Prostate	0.48	0.35
Bladder	0.67	0.72

4   Discussion

We found that the target organs were recognized and extracted correctly in all the test CT images, except for oversights of the portal vein, splenic vein, and superior mesenteric vein in two CT cases. Because our segmentation targets covered a wide range of shapes, volumes, and sizes, either with or without contrast enhancement, and at different locations in the human body, these experimental results demonstrated the potential capability of our approach to recognize whole anatomical structures appearing in CT images. The IUs of the organs with larger volumes (e.g., liver: 0.91, heart: 0.87) were comparable to the accuracies reported from the previous state-of-the-art methods

[2–5]. For some smaller organs (e.g., gallbladder) or line structures (e.g., portal vein, splenic vein, and superior mesenteric vein) that have not been reported in previous work, our segmentation did not show particularly high IUs, but this performance was deemed reasonable because the IU tends to be lower for those organs with smaller volumes. The physical CT image resolution is the major cause of this limited performance, rather than the segmentation method. Our evaluation showed that the average segmentation accuracy of all the targets over all the test CT images was approximately 84 % in terms of the frequency weighted IUs. The segmentation result of each deconvolution layer (FCN 32 s to FCN 2 s) was also investigated. We confirmed the frequency weighted IUs were monotonically increasing (about 0.16, 0.03 and 0.01) from FCN 32s, 16s, 8s and 4s, and no further improvement was observed by FCN 2s. This result showed diminishing returns of gradient descent from the training stage of FCN 8s, which was also mentioned in [11]. From experimental results, we see that our approach can recognize and extract all types of major organs simultaneously, achieving a reasonable accuracy according to the organ volume in the CT images. Furthermore, our approach can deal automatically with segmentation in 2D or 3D CT images with a free scan range (chest, abdominal, whole body region, etc.), which was impossible in previous work [2–5].

Our segmentation process has a high computational efficiency because of its simple structure and GPU-based implementation. The segmentation of one 2D CT slice takes approximately 30 ms (roughly 1 min for a 3D CT scan with 512 slices) when using the Caffe software package [15] and CUDA Library on a GPU (NVIDIA GeForce TITAN-X with 12 GB of memory). The efficiency in terms of system development and improvement is much better than that of previous work that attempted to incorporate human specialist experience into complex algorithms for segmenting different organs. Furthermore, neither the target organ type, number of organs within the image, nor image size limits the CT images that are used for the training process.

For the future work, network performance by using different training parameters as well as cost functions needs to be investigated, especially for de-convolution network. We plan to expand the range of 3D voting process from more than three directions of 2D image sections to improve the segmentation accuracy. Furthermore, bounding box of each organ [16] will be introduced into the network to overcome the insufficient image resolution for segmenting small-size of organ types. A comparison against 3D CNNs will also be investigated.

5 Conclusions

We proposed a novel approach for the automatic segmentation of anatomical structures (multiple organs and interesting regions) in CT images, by majority voting the results from a fully convolutional network. This approach was applied to segment 19 types of targets in 3D CT cases, demonstrating highly promising results. Our work is the first to tackle anatomical segmentation (with a maximum of 19 targets) on scale-free CT scans (both 2D and 3D images) through a deep CNN. Compared with previous work [2–5, 8–10], the novelty and advantages of our study are as follows. (1) Our approach uses an end-to-end, voxel-to-voxel labeling, with a global optimization of parameters, which

has the advantage of better performance and flexibility in accommodating the large variety of anatomical structures in different CT cases. (2) It can automatically learn a set of image features to represent all organ types collectively, using an "all-in-one" architecture (a simple structure for both model training and implementation) for image segmentation. This approach leads to more robust image segmentation that is easier to implement and extend. Image segmentation using our approach has more advantages in terms of usability (it can be used to segment any type of organ), adaptability (it can handle 2D or 3D CT images over any scan range), and efficiency (it is much easier to implement and extend) than those of previous work.

References

1. Doi, K.: Computer-aided diagnosis in medical imaging: historical review, current status and future potential. Comput. Med. Imaging Graph. **31**, 198–211 (2007)
2. Lay, N., Birkbeck, N., Zhang, J., Zhou, S.K.: Rapid multi-organ segmentation using context integration and discriminative models. In: Gee, J.C., Joshi, S., Pohl, K.M., Wells, W.M., Zöllei, L. (eds.) IPMI 2013. LNCS, vol. 7917, pp. 450–462. Springer, Heidelberg (2013)
3. Wolz, R., Chu, C., Misawa, K., Fujiwara, M., Mori, K., Rueckert, D.: Automated abdominal multi-organ segmentation with subject-specific atlas generation. IEEE Trans. Med. Imaging **32**(9), 1723–1730 (2013)
4. Okada, T., Linguraru, M.G., Hori, M., Summers, R.M., Tomiyama, N., Sato, Y.: Abdominal multi-organ segmentation from CT images using conditional shape-location and unsupervised intensity priors. Med. Image Anal. **26**(1), 1–18 (2015)
5. Bagci, U., Udupa, J.K., Mendhiratta, N., Foster, B., Xu, Z., Yao, J., Chen, X., Mollura, D.J.: Joint segmentation of anatomical and functional images: applications in quantification of lesions from PET, PET-CT, MRI-PET, and MRI-PET-CT images. Med. Image Anal. **17**(8), 929–945 (2013)
6. Shin, H.C., Roth, H.R., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D., Summers, R.M.: Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. IEEE Tran. Med. Imaging **35**(5), 1285–1298 (2016)
7. Ciompi, F., de Hoop, B., van Riel, S.J., Chung, K., Scholten, E., Oudkerk, M., de Jong, P., Prokop, M., van Ginneken, B.: Automatic classification of pulmonary peri-fissural nodules in computed tomography using an ensemble of 2D views and a convolutional neural network out-of-the-box. Med. Image Anal. **26**(1), 195–202 (2015)
8. de Brebisson, A., Montana, G.: Deep neural networks for anatomical brain segmentation. In: Proceedings of CVPR, Workshops, pp. 20–28 (2015)
9. Roth, H.R., Farag, A., Lu, L., Turkbey, E.B., Summers, R.M.: Deep convolutional networks for pancreas segmentation in CT imaging. In: Proceedings of SPIE, Medical Imaging 2016: Image Processing, vol. 9413, pp. 94131G-1–94131G-8 (2015)

10. Cha, K.H., Hadjiiski, L., Samala, R.K., Chan, H.P., Caoili, E.M., Cohan, R.H.: Urinary bladder segmentation in CT urography using deep-learning convolutional neural network and level sets. Med. Phys. **43**(4), 1882–1896 (2016)
11. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of CVPR, pp. 3431–3440 (2015)
12. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: Proceedings of ICLR. arXiv:1409.1556 (2015)
13. http://www.comp-anatomy.org/wiki/
14. Zhou, X., Ito, T., Takayama, R., Wang, S., Hara, T., Fujita, H.: First trial and evaluation of anatomical structure segmentations in 3D CT images based only on deep learning. In: Medical Image and Information Sciences (2016, in press)
15. http://caffe.berkeleyvision.org
16. Zhou, X., Morita, S., Zhou, X., Chen, H., Hara, T., Yokoyama, R., Kanematsu, M., Hoshi, H., Fujita, H.: Automatic anatomy partitioning of the torso region on CT images by using multiple organ localizations with a group-wise calibration technique. In: Proceedings of SPIE Medical Imaging 2015: Computer-Aided Diagnosis, vol. 9414, pp. 94143K-1–94143K-6 (2015)