

Fast and Accurate Lung Tumor Spotting and Segmentation for Boundary Delineation on CT Slices In A Coarse-To-Fine Framework

Shuchao Pang¹, Anan Du², Xiaoli He³, Jorge Diez⁴ and Mehmet Orgun¹

¹ Department of Computing, Macquarie University, Sydney, Australia

² School of Electrical and Data Engineering, University of Technology, Sydney, Australia

³ Department of Internal Medicine, Qingdao Huikang Hospital, Qingdao, China

⁴ Artificial Intelligence Center, University of Oviedo at Gijon, Spain

pangshuchao1212@sina.com, duanan2008@163.com,
dochexiaoli@126.com, jdiez@uniovi.es, mehmet.orgun@mq.edu.au

Abstract. *Label noise* and *class imbalance* are two of the critical challenges when training image-based deep neural networks, especially in the biomedical image processing domain. Our work focuses on how to effectively and accurately address the two challenges in the task of lesion segmentation from biomedical/medical images. To address *the pixel-level label noise problem*, an advanced transfer training and learning approach with a detailed DICOM pre-processing work is proposed in this paper. To address *the tumor/non-tumor class imbalance problem*, we exploit a self-adaptive fully convolutional neural network with an automated weight distribution mechanism to accurately spot the Radiomics lung tumor regions. Furthermore, an improved conditional random field method is utilized to obtain a sophisticated lung tumor contour delineation and segmentation. Finally, our approach has been evaluated using several well-known evaluation metrics on the Radiomics Lung Tumor segmentation dataset used in the 2018 IEEE VIP-CUP Challenge, which totally contains 422 patient records with 51513 Radiomics Lung slices. Experimental results show that our weakly supervised learning algorithm outperforms other deep models and state-of-the-art approaches.

Keywords: Boundary Delineation, Lung Tumor Segmentation, Fully Convolutional Neural Networks.

1 Introduction

With the improvement of clinical diagnostic equipment in terms of their capability, quality and availability in hospitals, biomedical/medical image data analysis has attracted much attention. The Volume, Variety, and Velocity (3V) of these images make it impractical and infeasible for clinicians to analyze them without making subjective errors [1]. Among different imaging devices, computed tomography is the most popular imaging modality because of its high resolution, imaging sensitivity, and isotropic ac-

2

quisition, e.g., locating the lung and its lesions [2]. Moreover, with the successful application of imaging technology in clinical medicine, automated image segmentation has been playing an increasingly important role. The accuracy of lesion region segmentation can be improved further by the consideration of Radiomics feature extraction and its detailed qualification with the ultimate goal of developing predictive models for precise prognosis in clinical medicine [3]. However, due to the variability and diversity during medical imaging processing and the existence of noisy-labelled datasets as well as tumor/non-tumor class imbalance, it is really hard to train a discriminative model for a specific lesion spotting and segmentation task with accurate contour delineation.

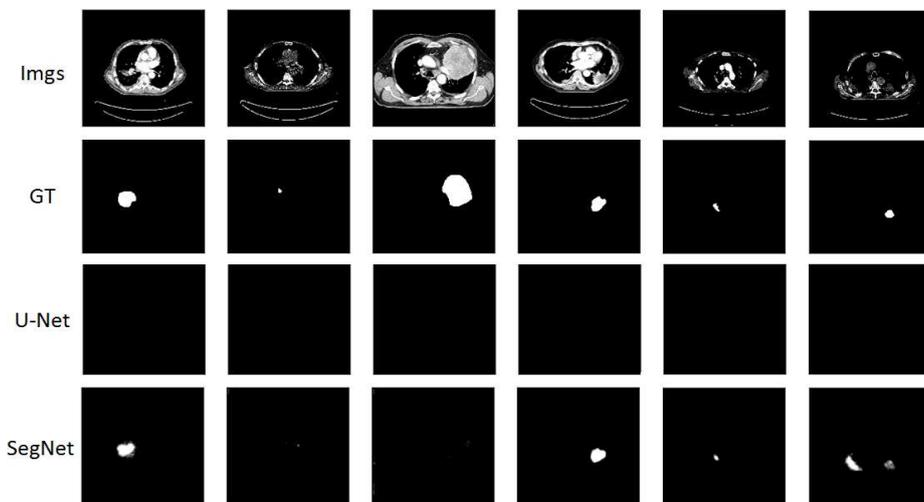


Fig. 1. Comparison of segmentation performance of current popular deep approaches (*U-Net* [17] and *SegNet* [25]) trained on the public NSCLC-Radiomics dataset with severely noisy labels and class imbalance. The row *imgs* shows several randomly selected lung CT slices and *GT* is the ground truth in spite of existing manual labelling mistakes.

1.1 Motivation

In the current vision tasks or competitions, the datasets used are carefully crafted and even balanced with respect to a golden standard, for example, ImageNet [27]. In this way, it will be easier to deal with some detection, classification, segmentation tasks without considering the complications arising from the dataset itself. By contrast, in real applications with raw datasets, many popular methods often fail to perform well or even completely fail [26]. Fig. 1 shows that classical biomedical image segmentation algorithms and other popular image semantic segmentation methods do not perform so well on this task, such as *U-Net* [17] and *SegNet* [25]. Furthermore, it is really difficult and time-consuming to obtain pure data sets and labels in many image processing applications, and we have to directly use the available raw dataset. Therefore it is necessary to develop a unified approach for accurately spotting and delineating the tumor

boundaries from NSCLC-Radiomics datasets with error pixel-level labels and severe class imbalance problems.

In this paper, we mainly focus on how to overcome the two critical challenges of i) noisy pixel-level labels and ii) tumor/non-tumor class imbalance when training a robust pixel-level deep model for biomedical/medical lesion spotting and segmentation with sophisticated gross tumor contour delineation. Specifically, Fig. 2 highlights the diversity of, and challenges that arise from, the lung CT dataset used in this work.

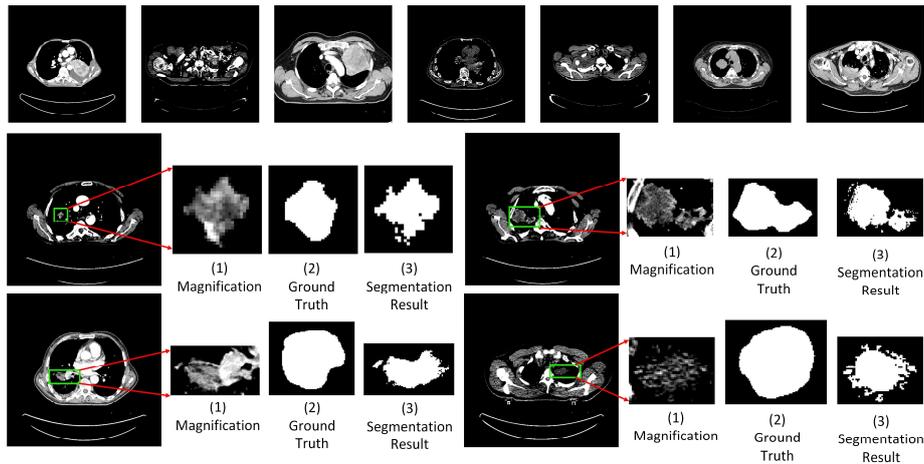


Fig. 2. Some Radiomics Lung CT slices from different patients in the public NSCLC-Radiomics dataset. Note that the first row indicates the diversity of the dataset and the rest of the figure shows the obvious noisy pixel-level label problems and small tumor regions, where the lesion region in each lung CT slice is magnified in contrast with its rough manual ground truth and our segmentation result.

1.2 Contributions

The main contributions of our work can be summarized as follows:

- We first propose a novel framework to simultaneously address noisy labels and class imbalance of a raw dataset for accurate biomedical/medical image lesion segmentation and its boundary delineation in real applications.

- We design an algorithm for reading from raw DICOM meta files as a preprocessing step for training deep neural networks. Then, a fully convolutional neural network is proposed to address pixel-level noisy labels using transfer learning. Furthermore, we develop a self-adapting weight allocation mechanism for addressing severe tumor/non-tumor imbalance problems to establish a differentiable cost function for spotting tumors. Finally, an improved conditional random field is exploited for accurate CT lung tumor segmentation and its contour delineation.

- Several well-known evaluation criteria are used to illustrate the performance of our framework in lesion spotting and boundary delineation compared to other state-of-the-art approaches for biomedical and medical image processing. Moreover, with fast

4

processing time, our average dice coefficient result is higher than the winners of 2018 IEEE VIP-CUP Challenge by a large margin.

2 Related Work

The current biomedical/medical image segmentation methods can be grouped under two categories: the co-segmentation method and the deep learning method.

2.1 Co-segmentation Method

Here, the co-segmentation method means combining two segmentation methods or using two or more types of biomedical/medical image modalities. Because any single approach used in biomedical/medical image segmentation has its own limitations and characteristics in a specific area, how to make full use of the advantages of different algorithms has become the focus recently. Many studies [4-6] have indicated that the co-segmentation method by combining different segmentation methods is treated as an energy minimization problem to delineate the gross tumor contours. Besides, due to the superior contrast of PET images and high spatial resolution of CT images, more recent methods and techniques [7-10] in the field of clinic and lesion segmentation prefer to integrate PET and CT images.

2.2 Deep Learning Method

In the past few years, deep learning has already swept through most research fields of computer vision and has achieved better results than traditional methods, and biomedical/medical image segmentation is also no exception [11, 12]. The existing methods [13-16] strive to obtain more precise and comprehensive tumor features by taking the advantage of deep learning in its superior ability of hierarchical feature representation.

Therefore, in this work, we attempt to fuse the integration capability of the co-segmentation method and the powerful descriptive capability of deep learning method to precisely localize the gross tumor boundaries for computer-aided automatic processing.

3 Proposed Method

3.1 Problem Setup and Preprocessing

In this section, we define the problem of segmenting lung tumors in CT volumes in a coarse-to-fine way. Suppose that there is a testing CT slice image (I) with a lung tumor from any 3D CT scans of a patient and the size of the image is 512×512 pixels. We denote each pixel in image I as $v_i, i = 1, 2, \dots, N$, where N is the total number of pixels in I . And the set of possible labels for each pixel can be represented as $L = \{0, 1, \dots, t\}$. Besides, for each pixel v_i , we define a variable $l_{v_i} \in L$ that indicates the assigned label. The probability of a pixel v_i belonging to label k in the given CT slice image I is represented by $P(l_{v_i} = k | I)$ and it will be calculated by our proposed fully convolutional

neural networks (FCNNs). In this gross tumor segmentation task, we take the set $L = \{0,1\}$ to denote the labels: 0 means non-tumor and 1 means tumor. In the training stage, our work is to train a deep neural network $\Phi(X, \theta)$, where $\Phi(\cdot)$ is the learned network on each training image X and θ indicates all network parameters, with the NSCLC-Radiomics dataset from 2018 IEEE VIP-CUP Challenge to compute the predicted probability maps $P(l_{v_i} | I)$ for any testing image I . In order to predict the result for each testing image, we use all the training data pairs (X_q, Y_q) , $q = 1, 2, \dots, Q$ for supervised training the neural network model $\Phi(X, \theta)$ looking for the best network parameters θ^* . Note that (X_q, Y_q) is the q^{th} training image and its label, and Q is the total number of training dataset. In this way, $P(l_{v_i} | I)$ is equivalent to $\Phi(I, \theta^*)$. To attain it, let $C(\hat{Y}, Y)$ be the loss function to minimize during the training phase, where $\hat{Y} = \Phi(X, \theta)$ and Y is the label of each training image X . Now, the optimal parameter θ^* for gross tumor segmentation can be calculated with the following formula:

$$\theta^* = \arg \min_{\theta} \sum_{s=1}^m C(\hat{Y}_s, Y_s) = \arg \min_{\theta} \sum_{s=1}^m C(\Phi(X_s, \theta), Y_s), \quad (1)$$

where the $m \ll Q$ stands for the number of images in the mini-batch. Then, the predicted pixel probability result $P(l_{v_i} | I, \theta^*)$ for each gross tumor slice can be further refined by using our improved dense conditional random fields as maximum a posteriori inference.

In addition, this subsection also shortly introduces the preprocessing steps from reading RTSTRUCT annotations through an extra DICOM metafile in clinics to transforming the original DICOM images in CT volumes as the inputs for deep neural networks (the workflow can be found in Fig. 3①). After this procedure, the whole NSCLC-Radiomics dataset [22, 23] from 422 patients are further grouped into missing file cases with 104 patients, irrelevant labelling cases with 30 patients checked by our physician and roughly usable cases (in spite of some noisy labels in each slice) with 288 patients. Therefore, all the CT slices from these 288 patients are used to train and test a lung tumor segmentation algorithm in the rest of this paper.

3.2 The Whole Architecture

Adaptive Fully Convolutional Neural Networks. By leveraging the benefits of the U-Net model [17], the strong limitation of 3D CNN [18] and the characteristics of CT lung slices, we improve the process to adapt the lung tumor segmentation task on each CT slice, as is shown in the step ② of Fig. 3. Here, a whole fully convolutional neural network model is illustrated with different stacked layers, which mainly comprises two key modules one of which is the encoder part which aims to capture spatial and context information of tumors and non-tumors, while the other is the decoder model that is used to recover the details and localize the position of tumors and non-tumors.

To address the noisy pixel-level label problem from the provided NSCLC-Radiomics dataset, we adopt a transfer learning strategy among different but interrelated biomedical image segmentation datasets for alleviating noisy pixel-level label interruption. This strategy is inspired by dealing with image-level classification task for

6

2017 CVPR WebVision Challenge in the work [19, 26]. For it, we resort to the Neuronal Structure Segmentation Dataset in Electron Microscopic Stacks at 2015 ISBI Challenge [20], which has an accurate neuronal structure segmentation gold standard in spite of only 30 training images.

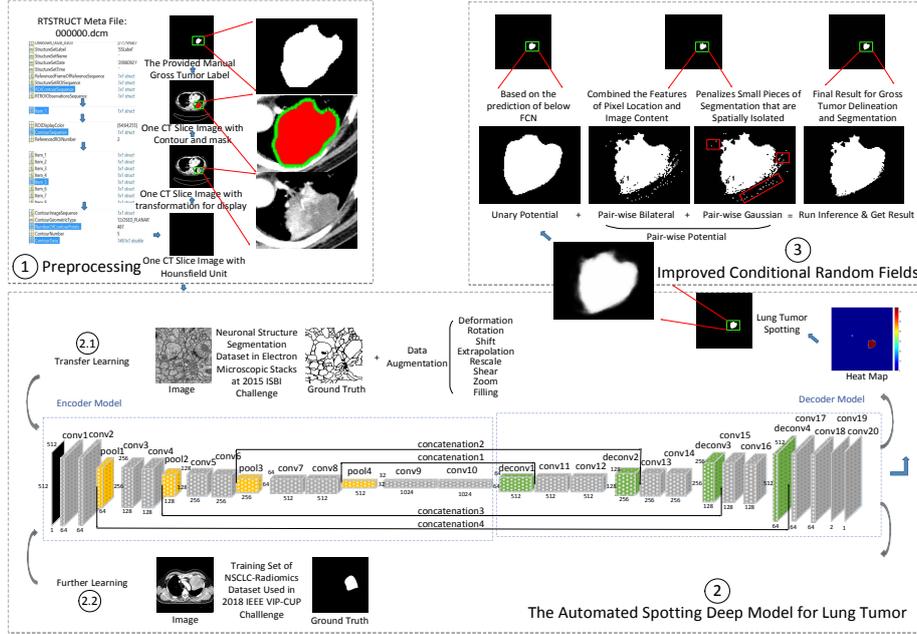


Fig. 3. Overview of the proposed gross tumor spotting and segmentation pipeline for Radiomics Lung CT images from different patients on NSCLC-Radiomics dataset.

When we train any classification network, the class imbalance could make the network recognize the classes with the vast majority of examples and ignore the rarely seen classes. To solve the severe tumor/non-tumor class imbalance problem from the dataset when training the deep model, we improve the binary loss function with a self-adapting weights allocation mechanism for these two categories in each mini-batch:

$$Loss_{mini-batch} = \sum_{s=1}^m C(\hat{Y}_s, Y_s) = \sum_{s=1}^m C(\theta(X_s, \theta), Y_s) = -\frac{1}{m \times N} \sum_{s=1}^m \sum_{i=1}^N [\omega_{tumour}^{class} Y_{s,v_i} \log \hat{Y}_{s,v_i} + \omega_{non-tumour}^{class} (1 - Y_{s,v_i}) \log (1 - \hat{Y}_{s,v_i})], \quad (2)$$

where the mini-batch size for pixel-level segmentation in our task is set to $m = 1$ which can reduce the requirement for GPU capability and improve the speed of training our deep model, and the weighting factors for tumor ω_{tumour}^{class} and non-tumor $\omega_{non-tumour}^{class}$ are defined in the following formulas respectively:

$$\omega_{tumour}^{class} = \frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=1} / \frac{1}{2} \left(\frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=1} + \frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=0} \right), \quad (3)$$

7

$$\omega_{non-tumour}^{class} = \frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=0} \bigg/ \frac{1}{2} \left(\frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=1} + \frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=0} \right). \quad (4)$$

Please note that the two weighting factors could be different for each training image pair (X_s, Y_s) in the mini-batch, which are automatically computed according to the size of the tumor in each CT image, so called self-adapting weights allocation mechanism for coping with serious tumor/non-tumor imbalance problem.

$$\theta^* = \arg \min_{\theta} - \frac{1}{m*N} \sum_{s=1}^m \sum_{i=1}^N \left[\frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=1} \bigg/ \frac{1}{2} \left(\frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=1} + \frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=0} \right) Y_{s,v_i} \log \hat{Y}_{s,v_i} + \frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=0} \bigg/ \frac{1}{2} \left(\frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=1} + \frac{1.0}{\sum_{i=1}^N Y_{s,v_i}=0} \right) (1 - Y_{s,v_i}) \log (1 - \hat{Y}_{s,v_i}) \right]. \quad (5)$$

Then, like [17], we also use Adam, a variant of the stochastic gradient descent algorithm, to optimize the above loss function and update the parameters θ . By minimizing Eq. (5), the parameters θ^* can be finally computed after training the fully convolutional neural network with 100 epochs.

Improved Dense Conditional Random Fields. The main highlight of this step is that we ameliorate the traditional conditional random fields for non-RGB images by adding pair-wise Gaussian potential into old pair-wise potential. The new potential can penalize small pieces of segmentation results that are spatially isolated in the original output results produced by the traditional conditional random fields [21], which can be clearly observed in part ③ of Fig. 3. Here, we adopt the graph model $G = (V, E)$ to represent a CT lung slice image, where $V = \{v_i\}, i = 1, 2, \dots, N$ and $E = \{e_{i,j}\}, i, j = 1, 2, \dots, N, i < j$. For the whole pipeline of step three, the improved dense conditional random fields includes two critical components, which are unary potential and pair-wise potential. For each testing image I , the objective function of our improved dense conditional random fields is defined as follows:

$$Energy(l_v) = \sum_V \varphi_u(l_{v_i}) + \sum_E \varphi_p(l_{v_i}, l_{v_j}), \quad (6)$$

where the first term $\varphi_u(l_{v_i})$ denotes unary potential and it can be equal to the probabilistic output of our trained deep gross tumor spotting model with $P(l_{v_i} | I, \theta^*)$. And the second term $\varphi_p(l_{v_i}, l_{v_j})$ is the pair-wise potential, where we define the bilateral potential and Gaussian potential inside with the following equation (7). Here, $\mu(l_{v_i}, l_{v_j})$ is given with the Potts function that evaluates the label compatibility, and p_v and I_v denote pixel position and intensity content information respectively. Besides, δ_{bil} and δ_{gau} separately stand for the proportion of each kind of pair-wise potential with different effective range α, β, γ .

8

$$\varphi_p(l_{v_i}, l_{v_j}) = \mu(l_{v_i}, l_{v_j}) \left[\delta_{bil} \exp\left(-\frac{|p_{v_i} - p_{v_j}|^2}{2\alpha^2} - \frac{|l_{v_i} - l_{v_j}|^2}{2\beta^2}\right) + \delta_{gau} \exp\left(-\frac{|p_{v_i} - p_{v_j}|^2}{2\gamma^2}\right) \right], \quad (7)$$

At last, to gain more precise gross tumor boundaries with $P(l_{v_i} | I)$, the best label result l_v^* can be computed by the following formula with the efficient approximation inference approach proposed in [21]:

$$l_v^* = \arg \min_{l_v \in L} \text{Energy}(l_v). \quad (8)$$

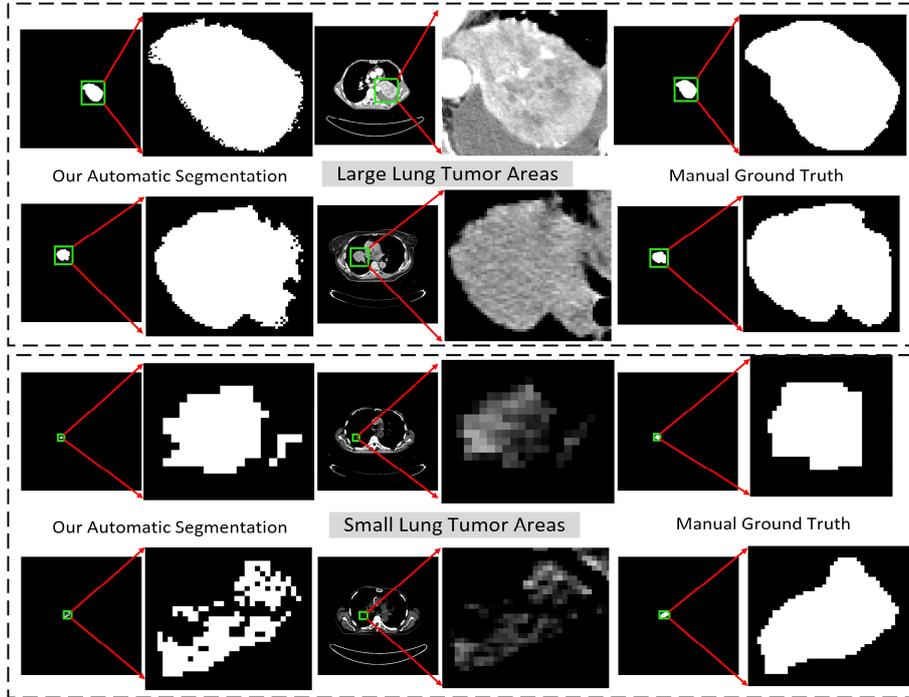


Fig. 4. Results of our proposed gross tumor spotting and segmentation approach for large and small lung tumors on different CT lung slices.

4 Experiment

4.1 Datasets and Evaluation Metrics

Like [20], we also use different data augmentation techniques on the Neuronal Structure Segmentation Dataset for transfer learning, and NSCLC-Radiomics dataset [22, 23] is randomly split into training, validation and test sets with the proportion of 7:1:2 for

lung tumor spotting and segmentation. In order to efficiently evaluate the performance of our proposed approach, several public and widely used image semantic segmentation evaluation metrics are utilized in our experiments, including Dice Coefficient, Hausdorff Distance, Jaccard Index, Precision, Sensitivity (Recall), Specificity and F1.

4.2 Experimental Results and Analysis

Our coarse-to-fine algorithm with three steps can achieve significant lung tumor spotting and segmentation results. In terms of qualitative analysis shown in Fig. 4, the first two rows are to check the performance on large gross tumor areas and we can observe that the boundaries of large lesions are described in more details and clearer than the given roughly manual ground truth. Furthermore, our algorithm does not accurately spot the position of tumor, but also carefully discriminate tumor or non-tumor pixel by pixel. This can be also seen in Table 1.

Table 1. Segmentation results on several test lung tumor slices (above) listed in Fig. 4 and all the test data (below) with different evaluation criteria. Note that the numbers in bold face indicate the best result under different criteria for all the test set.

Test Slices \ Metrics	DICE	HD	JAC	Precision	Sensitivity	Specificity	F1
Large 1	0.9646	4.2426	0.9317	0.9751	0.9544	0.9996	0.9646
Large 2	0.9192	5.3852	0.8504	0.9605	0.8813	0.9997	0.9192
Small 1	0.5035	7.2111	0.3365	0.9817	0.3386	1.0000	0.5035
Small 2	0.5029	9.2195	0.3359	1.0000	0.3359	1.0000	0.5029
All test set \ Methods							
SegNet [25]	0.7518	50.935	0.6260	0.6666	0.9256	0.9958	0.7750
U-Net [17]- w.-Our-Tech.	0.6209	346.71	0.4662	0.5367	0.7869	0.9947	0.6382
Ours-w.-Old- CRF [21]	0.5313	48.436	0.3953	0.8155	0.4425	0.9991	0.5737
Ours	0.7767	15.492	0.6493	0.9247	0.6951	0.9995	0.7936

As annotated in last two rows in Fig. 4 for small lung tumors, the manual labeled gross tumor region is really far away from the real tumor in the original slices and this is why the evaluation results in Table 1 are so bad. In other words, our proposed automatic lung tumor segmentation method is rather accurate to identify precise locations of lesion boundaries. So, this method is useful for assisting the physicians to complete their time-consuming labelling work and reducing inaccurate supervision information. These results also reflect the manual annotation errors in the provided dataset by a radiation oncologist [23]. Furthermore, in order to evaluate boundary location performance of our segmentation method with high quality, a clear tumor contour delineation is shown in Fig. 5 with several local enlargement patches.

As a quantitative analysis of our method, we evaluated our model over the test set by removing the cases with obvious and unrealistic raw data errors by our physician.

Table 1 shows the results and we can see that our method achieves a significant performance for lung tumor spotting and segmentation in CT slices, and it is especially competitive compared to other typical segmentation models and methods. Except original U-Net without any prediction for tumor, we compare our method with SegNet and the latter can obtain a better recall. However, its precision is really worse which means SegNet classifies lots of non-tumor pixels as tumor regions. In particular, it might take more misclassified results under the situation of many noisy ground truths. Next, we improve the U-Net method with our techniques proposed in this paper and we can find it can roughly predict tumor regions better than before. Besides, we also compare our improved conditional random fields with its naive model, which can be observed that our framework can obtain a 92.47% precision than traditional conditional random fields' 81.55%. Furthermore, by comparison, our approach can even achieve much better average dice coefficient on test set with 0.7767 than the reported results of 2018 VIP Cup (Winner: Team Markovian of 0.594 and Runner up: Team NTU_MiRA of 0.521) [24]. Besides, the average processing time for a CT slice is 468 ms, including spotting time with 79 ms and segmentation time with 389 ms, which is rather fast for practical applications.

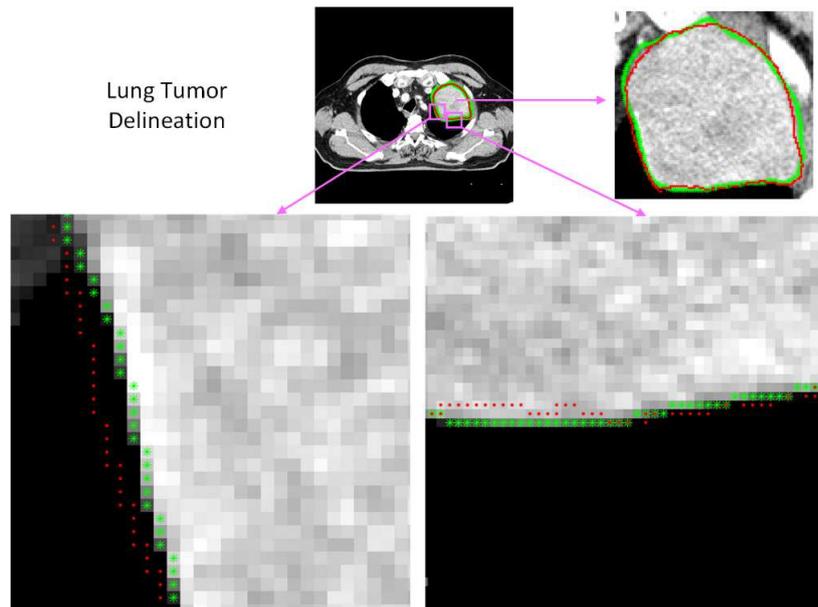


Fig. 5. Comparison of our proposed automatic delineation result (green) and manual ground truth (red) in details.

5 Conclusion

In this work, we have proposed a novel framework to fuse the integration benefits of co-segmentation model and powerful discriminative capability of deep learning method

to precisely and automatically localize the gross tumor boundaries in medical images. Most importantly, we propose a unified framework to successfully overcome these two critical bottlenecks in training a deep model for pixel-level medical image segmentation task: i) noisy pixel-level labels and ii) tumor/non-tumor class imbalance problems. Finally, by designing a coarse-to-fine model via weakly supervised learning, the proposed Radiomics gross tumor segmentation algorithm can achieve more precise contour delineation and segmentation performance than those state of the art methods. Moreover, our proposed approach has fast response times for an accurate clinical diagnosis and a good radiation therapy treatment planning.

References

1. Pang, S., Yu, Z., Orgun, M.A.: A novel end-to-end classifier using domain transferred deep convolutional neural networks for biomedical images. *Computer methods and programs in biomedicine* 140, 283–293 (2017)
2. Skourt, B.A., El Hassani, A., Majda, A.: Lung ct image segmentation using deep neural networks. *Procedia Computer Science* 127, 109–113 (2018)
3. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., Van Der Laak, J.A., Van Ginneken, B., Sánchez, C.I.: A survey on deep learning in medical image analysis. *Medical image analysis* 42, 60–88 (2017)
4. Ju, W., Xiang, D., Zhang, B., Wang, L., Kopriva, I., Chen, X.: Random walk and graph cut for co-segmentation of lung tumor on pet-ct images. *IEEE Transactions on Image Processing* 24(12), 5854–5867 (2015)
5. Hsieh, T.M., Liu, Y.M., Liao, C.C., Xiao, F., Chiang, I.J., Wong, J.M.: Automatic segmentation of meningioma from non-contrasted brain mri integrating fuzzy clustering and region growing. *BMC medical informatics and decision making* 11(1), 54 (2011)
6. Soltaninejad, M., Yang, G., Lambrou, T., Allinson, N., Jones, T.L., Barrick, T.R., Howe, F.A., Ye, X.: Automated brain tumour detection and segmentation using superpixel-based extremely randomized trees in flair mri. *International journal of computer assisted radiology and surgery* 12(2), 183–203 (2017)
7. Song, Q., Bai, J., Han, D., Bhatia, S., Sun, W., Rockey, W., Bayouth, J.E., Buatti, J.M., Wu, X.: Optimal co-segmentation of tumor in pet-ct images with context information. *IEEE transactions on medical imaging* 32(9), 1685–1697 (2013)
8. Wang, X., Ballangan, C., Cui, H., Fulham, M., Eberl, S., Yin, Y., Feng, D.: Lung tumor delineation based on novel tumor-background likelihood models in pet-ct images. *IEEE Transactions on Nuclear Science* 61(1), 218–224 (2014)
9. 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. *Medical image analysis* 17(8), 929–945 (2013)
10. Zhong, Z., Kim, Y., Zhou, L., Plichta, K., Allen, B., Buatti, J., Wu, X.: 3d fully convolutional networks for co-segmentation of tumors on pet-ct images. In: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. pp. 228–231. IEEE (2018)
11. Pang, S., del Coz, J.J., Yu, Z., Luaces, O., Díez, J.: Deep learning to frame objects for visual target tracking. *Engineering Applications of Artificial Intelligence* 65, 406–420 (2017)
12. Salehinejad, H., Valae, S., Dowdell, T., Colak, E., Barfett, J.: Generalization of deep neural networks for chest pathology classification in x-rays using generative adversarial networks.

12

In: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 990–994. IEEE (2018)

13. Dong, H., Yang, G., Liu, F., Mo, Y., Guo, Y.: Automatic brain tumor detection and segmentation using u-net based fully convolutional networks. In: annual conference on medical image understanding and analysis. pp. 506–517. Springer (2017)
14. Christ, P.F., Elshaer, M.E.A., Ettlinger, F., Tatavarty, S., Bickel, M., Bilic, P., Rempfler, M., Armbruster, M., Hofmann, F., DAnastasi, M., et al.: Automatic liver and lesion segmentation in ct using cascaded fully convolutional neural networks and 3d conditional random fields. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 415–423. Springer (2016)
15. Leung, K., Marashdeh, W., Wray, R., Ashrafinia, S., Rahmim, A., Pomper, M., Jha, A.: A deep-learning-based fully automated segmentation approach to delineate tumors in fdg-pet images of patients with lung cancer. *Journal of Nuclear Medicine* 59(supplement 1), 323–323 (2018)
16. Saltz, J., Gupta, R., Hou, L., Kurc, T., Singh, P., Nguyen, V., Samaras, D., Shroyer, K.R., Zh., T., Batiste, R., et al.: Spatial organization and molecular correlation of tumor-infiltrating lymphocytes using deep learning on pathology images. *Cell reports* 23(1), 181–193 (2018)
17. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
18. Prasoorn, A., Petersen, K., Igel, C., Lauze, F., Dam, E., Nielsen, M.: Deep feature learning for knee cartilage segmentation using a triplanar convolutional neural network. In: International conference on medical image computing and computerassisted intervention. pp. 246–253. Springer (2013)
19. Guo, S., Huang, W., Zhang, H., Zhuang, C., Dong, D., Scott, M.R., Huang, D.: Curriculumnet: Weakly supervised learning from large-scale web images. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 135–150 (2018)
20. Arganda-Carreras, I., Turaga, S.C., Berger, D.R., Cire, san, D., Giusti, A., Gambardella, L.M., Schmid, J., Laptev, D., Dwivedi, S., Buh., J.M., et al.: Crowdsourcing the creation of image segmentation algorithms for connectomics. *Frontiers in neuroanatomy* 9, 142 (2015)
21. Kr̄ahenb̄uhl, P., Koltun, V.: Efficient inference in fully connected crfs with gaussian edge potentials. In: Advances in neural information processing systems. pp. 109– 117 (2011)
22. Aerts, H., Rios Velazquez, E., Leijenaar, R.T., Parmar, C., Grossmann, P., Carvalho, S., Lambin, P.: Data from nslc-radiomics. *Cancer Imaging Archive* (2015)
23. Aerts, H.J., Velazquez, E.R., Leijenaar, R.T., Parmar, C., Grossmann, P., Carvalho, S., Bussink, J., Monshouwer, R., Haibe-Kains, B., Rietveld, D., et al.: Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nature communications* 5, 4006 (2014)
24. Mohammadi, A., Afshar, P., Asif, A., Farahani, K., Kirby, J., Oiko., A., Plata., K.N.: Lung cancer radiomics: Highlights from the ieev video and image processing cup 2018 student competition [sp competitions]. *IEEE Signal Processing Magazine* 36(1), 164–173 (2018)
25. Badrinarayanan, V., Kendall, A., Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence* 39(12), 2481–2495 (2017)
26. Li, W., Wang, L., Li, W., Agustsson, E., Van Gool, L.: Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862* (2017)
27. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* 115(3), 211–252 (2015)