

Solution to the Kidney Tumor Segmentation Challenge 2019

Jun Ma

School of Science, Nanjing University of Science and Technology, China
junma@njjust.edu.cn

Abstract. Kidney and kidney tumor segmentation are essential steps in kidney cancer surgery. In this paper, we focus on addressing hard cases and exploring the kidney tumor shape prior rather than developing new convolution neural network architectures. Specifically, we train additional tumor segmentation networks to bias the ensemble classifier to tumor. Moreover, we propose the compact loss function to constrain the shape of the tumor segmentation results. Experiments on KiTS challenge show that both hard mining and compact can improve the performance of U-Net baseline.

Keywords: Kidney segmentation · U-Net · Compact loss.

1 Introduction

There are more than 400,000 new cases of kidney cancer each year¹, and surgery is its most common treatment [2]. Due to the wide variety in kidney and kidney tumor morphology, there is currently great interest in how tumor morphology relates to surgical outcomes, [3,4] as well as in developing advanced surgical planning techniques [5]. Automatic semantic segmentation is a promising tool for these efforts, but morphological heterogeneity makes it a difficult problem.

2 Method

There are lots of convolutional neural network (CNN) architectures for medical image segmentation tasks. In this work, instead of exploring new or existing CNN architectures, we focus on dealing with the hard class and exploring the shape prior of kidney tumor.

2.1 Baselines: Vanilla 3D U-Net and cascaded 3D U-Net

We use the 3D U-Net [4] as the main network architecture as it has been proved to be a strong baseline for many segmentation tasks [7]. We build two baselines.

¹ Kidney Cancer Statistics. World Cancer Research Fund, 12 Sept. 2018, www.wcrf.org/dietandcancer/cancer-trends/kidney-cancer-statistics.

Specifically, the first baseline trains a 3D U-Net to segment both kidney and tumor. The second baseline employs cascaded pipeline which has achieved state-of-the-art performance in recent atrial segmentation challenge [9] and thoracic organs segmentation challenge [5]. In particular, we firstly segment the kidney by a 3D U-Net and crop the kidney region-of-interest (ROI), and then a new 3D U-Net is used to segment both kidney and tumor in the ROI.

2.2 Biased U-Net and Compact loss for tumor segmentation

Biased U-Net After hard mining, we find the accuracy of tumor is still significantly worse than the accuracy of kidney. The main reason is class imbalance. Many methods deal with label imbalance from data level, such as oversampling the minor class data. Instead of oversampling the tumor patch for training, we “oversample” the model to eliminate class imbalance. Specifically, we train a U-Net to only segment kidney tumor, termed as biased U-Net.

Compact loss Most kidney lesions have a ball shape in 3D or a nearly circle shape in 2D in the training sets. The commonly used loss functions (e.g., cross entropy, Dice loss [8]) do not have shape constrain. Thus, we propose the compact loss function to impose explicitly compact shape regularization. Shape compactness is usually defined as

$$C = \frac{P^2}{A} \quad (1)$$

where C is the shape compactness, P and A are the shape perimeter and area, respectively. We use g, s denote the ground truth and predicted segmentation, the compact loss is defined as

$$L_C = \frac{(\text{Length of } s)^2}{\text{Area of } s} \approx \frac{(\int_{\Omega} \nabla s \, dx)^2}{\int_{\Omega} s \, dx} \quad (2)$$

To avoid the training stuck in trivial local minima, in practice, we combine the proposed compact loss function with region-based active contour loss [3].

$$L_{AC} = \text{Region}_{in} + \text{Region}_{out} = \int_{\Omega} [s(1-g)^2 + (1-s)g^2] dx \quad (3)$$

Like tradition active contour models which need an initialization, the proposed compact loss uses the trained model as initialization. In other words, compact loss is used to finetune and refine the segmentation generated by commonly used loss functions (in this paper, it is the sum between Dice loss and cross entropy).

3 Experiments and Results

Data: We only use official KiTS challenge dataset [6] and do not use any external dataset. All the training data are interpolated to a common spacing (0.7816, 0.7816, 3). We randomly divide the training set into five folds.

Training protocol: All the U-Nets are trained based on nnU-Net [7], which is a powerful implementation of U-Net. We use the sum between Dice loss and cross entropy as the default loss function. For cascaded U-Net, we train a new version with the sum between Dice and TopK loss. For the biased U-Net, we use the default loss function firstly, and then finetune with compact loss. We use adam optimizer with an initial learning rate $3e4$. For finetuning, the learning rate is decreased by 10. We train the model on local GPU server with four TiTan-XP GPUs and two Intel Xeon E5-2650V4 CPUs. Training each model costs about four days.

Results on training set: We summarize the quantitative results in Tab 1. There is no significant difference for kidney Dice score between different segmentation methods. In fact, a vanilla U-Net performs quite well on kidney segmentation. For tumor segmentation, cascaded pipeline performs better than segmenting kidney and tumor directly. Biased U-Net with compact loss can further improve segmentation performance of tumor.

Table 1. quantitative results of kidney and kidney tumor segmentation. (All the results are based on 5-fold cross validation.)

Model	Kidney Dice	Tumor Dice
Vanilla U-Net with Dice + CE loss	0.9663	0.7778
Cascaded U-Net with Dice + CE loss	0.9682	0.8243
Cascaded U-Net with Dice + TopK loss	0.9753	0.8192
Biased U-Net and Compact Loss	-	0.8328

Testing: All the test data are also re-sampled to (0.7816, 0.7816, 3). Then, each test case is passed the trained networks. Test time augmentation (mirroring) is also used to boost performance.

Ensemble and Post-processing Ensemble is an effective way to further obtain performance gain which has been proved by many participants in previous segmentation challenges ([1],[2]). In our work, the final kidney segmentation results are obtained by the majority vote of two cascades U-Nets and a single kidney segmentation U-Net. Simply combining tumor segmentation results by the majority vote of the last three models will miss some detected lesions. Thus, we group all the tumor results in the last three models, and further fuse with the tumor results in the first model (vanilla U-Net) by reconstruction². This is because the first model has better sensitivity while the combination of the last three models has better specificity for tumor segmentation.

² skimage.morphology.reconstruction

Moreover, there are two types of outlier, small isolated points and big over-segmentation, in the final segmentation. Fortunately, these outliers do not touch with the true kidney. Thus, we can remove the isolated points smaller 20,000 voxels by `skimage.morphology.remove_small_objects`, and remove the big over-segmentation by the relative position relationship. Specifically, the central points of two kidneys is close in axis plane, while the central points of outliers will have obvious deviation.

4 Discussion and Conclusion

In this paper, we use several U-Nets to deal with kidney and tumor segmentation. The main scientific contribution is the compact loss, which is a regularization to constrain the tumor segmentation with compact shape. In the near future, we will investigate the compact loss with more combinations, such as coupled with Dice loss or cross entropy loss.

We acknowledge that using various U-Nets ensemble is not a concise way for segmentation tasks. Although ensemble can obtain a few performance improvements, it only has limited usage in clinical practice. It is desirable to build a new challenge branch that only allows to use single model.

Open questions: 1) Is Dice score enough to evaluate the segmentation performance? A major concern is that Dice score is sensitive to small tumor, as changing a few voxels can change Dice score significantly. In addition, if a patient has both a big tumor in left kidney and a very small tumor in right kidney, but an algorithm only segments the big tumor and misses the small tumor. In this case, although a high Dice score may be obtained, we can not think the algorithm has great performance. 2) What accuracy is clinically acceptable? 3) What the effects of different segmentation accuracy have on clinical practice? For example, are there any differences between kidney Dice score 0.95 and 0.97 for clinical usage? 4) If the current segmentation accuracy is not enough for clinical usage, how long it will take if radiologists manually refine the current segmentation results to make it clinically acceptable? Does it really save time compared with manual segmentation?

Acknowledgement Importantly, we would like to express our gratitude to all the KiTS organizers. Everything has been considered, e.g., avoiding test set optimization, test case-wise refinement and unexpected bugs. This is really an excellent challenge, and we enjoy it very much. We would also like to thank the nnU-Net[7] team which make nnU-Net, a very flexible and powerful tool, publicly available. We also very much appreciate Nicholas Heller and Fabian Isensee for their kindly help in kits19 and nnunet github repositories.

References

1. Bakas, S., Reyes, M., Jakab, A., Bauer, S., Rempfler, M., Crimi, A., Shinohara, R.T., Berger, C., Ha, S.M., Rozycki, M., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. arXiv preprint arXiv:1811.02629 (2018)
2. Bilic, P., Christ, P.F., Vorontsov, E., Chlebus, G., Chen, H., Dou, Q., Fu, C.W., Han, X., Heng, P.A., Hesser, J., et al.: The liver tumor segmentation benchmark (lits). arXiv preprint arXiv:1901.04056 (2019)
3. Chen, X., Williams, B.M., Vallabhaneni, S.R., Czanner, G., Williams, R., Zheng, Y.: Learning active contour models for medical image segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 11632–11640 (2019)
4. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: International conference on medical image computing and computer-assisted intervention. pp. 424–432 (2016)
5. Han, M., Yao, G., Zhang, W., Mu, G., Zhan, Y., Zhou, X., Gao, Y.: Segmentation of ct thoracic organs by multi-resolution vb-nets. In: Segmentation of THoracic Organs at Risk in CT images Proceeding (2019)
6. Heller, N., Sathianathan, N., Kalapara, A., Walczak, E., Moore, K., Kaluzniak, H., Rosenberg, J., Blake, P., Rengel, Z., Oestreich, M., et al.: arXiv preprint arXiv:1904.00445 (2019)
7. Isensee, F., Petersen, J., Kohl, S.A., Jäger, P.F., Maier-Hein, K.H.: nnu-net: Breaking the spell on successful medical image segmentation. arXiv preprint arXiv:1904.08128 (2019)
8. Milletari, F., Navab, N., Ahmadi, S.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 Fourth International Conference on 3D Vision (3DV). pp. 565–571 (2016)
9. Xia, Q., Yao, Y., Hu, Z., Hao, A.: Automatic 3d atrial segmentation from ge-mris using volumetric fully convolutional networks. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 211–220 (2018)