

Shape-Aware Complementary-Task Learning for Multi-Organ Segmentation

Fernando Navarro^{*1}, Suprosanna Shit^{*1}, Ivan Ezhov¹, Johannes C. Paetzold¹, Andrei Gafita³, Jan C. Peeken², Stephanie E. Combs², and Bjoern H. Menze¹

¹ Department of Informatics, Technische Universität München, Germany

² Department of Radiotherapy, Klinikum rechts der Isar, Germany

³ Department of Nuclear Medicine, Klinikum rechts der Isar, Germany

`fernando.navarro@tum.de`

Abstract. Multi-organ segmentation in whole-body computed tomography (CT) is a constant pre-processing step which finds its application in organ-specific image retrieval, radiotherapy planning, and interventional image analysis. We address this problem from an organ-specific shape-prior learning perspective. We introduce the idea of complementary-task learning to enforce shape-prior leveraging the existing target labels. We propose two complementary-tasks namely i) distance map regression and ii) contour map detection to explicitly encode the geometric properties of each organ. We evaluate the proposed solution on the public VISCERAL dataset containing CT scans of multiple organs. We report a significant improvement of overall dice score from 0.8849 to 0.9018 due to the incorporation of complementary-task learning.

Keywords: Multi-task learning · Complementary-task · Multi-organ segmentation.

1 Introduction

In representation learning, auxiliary-tasks are often designed to leverage *free-of-cost* supervision which is derived from existing *target labels*. The purpose of including auxiliary tasks is not only to learn a shared representation but also to learn efficiently by solving the common *meta-objective*. A group of auxiliary-tasks driven by a common *meta-objective* often have *complementary objectives*. For example, to detect an orange, one can define two sub-tasks: i) learn only the shape of the orange, and, ii) learn only the color of the orange. Here, learning the shape and the color complement each other to learn the common *meta-objective*: how an orange looks. In this context we leverage complementary-tasks learning by jointly optimizing the common *meta-objective* of multiple-tasks.

Solving multiple tasks simultaneously [5] is known to improve each task’s performance when compared to learning them independently. Mutual information exchange between multiple tasks such as detection and segmentation drive a

* The authors contributed equally to the work.

neural network towards learning a generalized shared representation [6]. A more recent success in multi-task learning is mask-RCNN[4] which benefited from the combined object detection and instance segmentation tasks. Uslu et al. [12] shows that learning vessel interior, centerline and edges as a set of complementary-tasks improves junction detection in the retinal vasculature. There is a new body of research which aims to efficiently balance losses of different competing tasks [2,5]. In contrast, our main objective is to leverage multiple complementary-tasks which shares a common *meta-objective* and do not need additional annotated targets. Given that medical imaging problems such as multi-organ segmentation have to be solved with limited annotated data, complementary-task learning is a promising alternative solution.

Multiple approaches have been proposed for multi-organ segmentation, which can be classified into registration-based and machine-learning based approaches. In registration-based methods, an atlas is registered to a test volume to obtain the segmentation map. [13,10]. However this is a time-consuming method and its performance also suffers from inter-subject variability. Alternatively, the state-of-the-art methods for multi-organ segmentation are based on deep learning architectures such as fully convolutional neural networks FCN [9] and U-Net [8]. Deep learning-based methods can successfully segment large anatomical structures but are prone to miss small organs. Recently, Zhao *et al.*[14] have explored the idea of combining both registration-based approaches and deep learning to segment small organs.

However a contemporary study by Geirhos et al. [3] shows that convolutional neural networks are inherently biased towards texture information over the shape of an object. In human anatomy, all organs have a discriminative shape feature which has not been addressed at its full potential by previous approaches. Our hypothesis is that learning *shape-prior* as a complementary-task improves the performance of the segmentation task. For learning the meta-objective of multi-organ representation, we define two complementary-tasks for shapes-prior learning: i) inferring geometric shape properties of an organ, and, ii) detecting the exterior contour of an organ. For the former, we propose the distance transform [1] of shape as geometric shape properties. For the latter, we leverage the binary edge-map of each organ.

In summary, our key contributions in this work are as follows:

- To the best of our knowledge, this is the first work for multi-organ segmentation leveraging complementary-task learning.
- We introduce two complementary-tasks in context of organ-specific shape-prior learning. We show that the inclusion of these complementary-tasks alongside the segmentation task improves its overall performance.

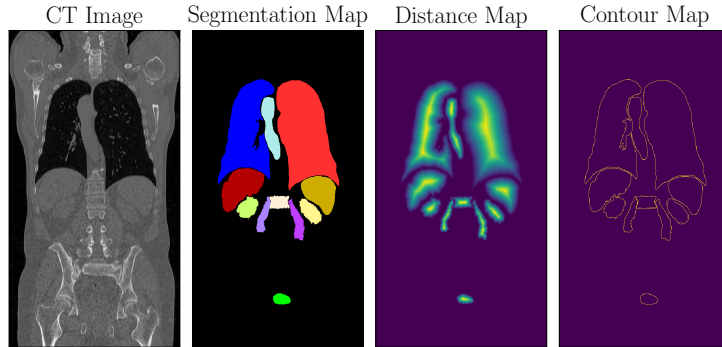


Fig. 1: Target generation for complementary-tasks: *Distance Map* is cumulative sum of the normalized distance transform of each organs’ segmentation map. *Contour Map* is the binary edge of each organ. Shape representation learning is jointly enforced in the distance and contour maps.

2 Methodology

In this section, we present the proposed complementary-task learning for multi-organ segmentation and, subsequently describe the network architecture and the associated loss functions.

In general, learning shape prior is itself a difficult task for a convolutional network due to the variety of shapes in different organs. We identify that the task of learning shape-priors can be broken down into multiple easier and quantifiable sub-task. Motivated by this fact, we propose two complementary-tasks to explicitly enforce shape and anatomical positional prior to the network.

Learning Distance-transform: Euclidean distance transform of a shape is commonly used to find the inscribed circle having largest radius within an arbitrary shape. It maps a boundary regularity measure of a shape with respect to an interior point. Hence we argue that the distance transform regression learns geometric properties of shapes. In addition to that, we find that Gaussian heat-map regression is a common task in localizing anatomical landmarks [7]. Parallels can be drawn between Gaussian heat-map regression and Euclidean distance map regression for soft-localization of organ-specific landmarks.

Learning Organ Contour: Detecting organ boundaries is the most challenging task for a segmentation network. Hence we propose to explicitly learn the organ contour as the second complementary-task alongside the distance map regression. The hypothesis is that the distance transform aids the network to accurately localize the organ from the learned anatomical prior whereas contour learning penalizes for boundary miss-classification, and, thus fine-tuning the or-




Fig. 2: Network architecture for complementary-task learning. The input is a CT slice. The network consists of an encoding and decoding part with skip connections resembling the U-Net architecture with three branches diverging at the end of the last up-convolution. The outputs of the network are the segmentation map, the distance map, and the contour map.

gan shape.

Loss Function: The loss function consist of three optimization objectives of the corresponding task: segmentation, distance regression and contour detection. Given the probability $p_l(x)$ of a pixel in location x to belong to class l and the ground truth by $g_l(x)$ the multi-organ segmentation loss \mathcal{L}_{seg} and the contour map loss $\mathcal{L}_{contour}$ are defined as:

$$\mathcal{L}_{seg}, \mathcal{L}_{contour} = \underbrace{-\sum_x g_l(x) \log p_l(x)}_{\text{Cross-Entropy Loss}} - \underbrace{\frac{2 \sum_x p_l(x) g_l(x)}{\sum_x p_l^2(x) + \sum_x g_l^2(x)}}_{\text{Dice Loss}} \quad (1)$$

where $l = \#$ of organs + 1 in the multi-organ segmentation loss, and $l = 2$ for the contour map. Dice loss is incorporated to handle the class-imbalance. For the distance map regression, the mean square error loss is optimized. Given the estimated distance map $p(x)$ of pixel x and the ground truth value $g(x)$ the distance map loss function \mathcal{L}_{dist} is:

$$\mathcal{L}_{dist} = -\frac{1}{n} \sum_x (g(x) - p(x))^2 \quad (2)$$

where $n = \#$ of pixels and the final loss $\mathcal{L}_{total} = \mathcal{L}_{seg} + \mathcal{L}_{contour} + \mathcal{L}_{dist}$ is given by the summation of all losses.




Fig. 3: Qualitative comparison of an axial slice of a test sample between ground truth (GT) and the proposed model prediction (Pred). Our proposed model efficiently learn to i) segment the organs, ii) regress on distance map and iii) detect the organ contour map at the same time.

Network Architecture: The network architecture is inspired by encoding-decoding architectures with skip connections [8,9]. A generalized shared feature representation is learned throughout the encoding blocks, the bottleneck, and a part of the decoding blocks of the network. Three different branches are added to the common feature representation to produce three output maps(c.f. Fig. 2).

3 Experiments and Discussion

To validate the performance of the proposed approach, we perform the segmentation of 16 organs. The following four different experiments are performed:

- U-Net: baseline using only \mathcal{L}_{seg} .
- U-Net + distance: $\mathcal{L}_{seg} + \mathcal{L}_{dist}$.
- U-Net + contour: $\mathcal{L}_{seg} + \mathcal{L}_{contour}$.
- U-Net + distance,contour: $\mathcal{L}_{seg} + \mathcal{L}_{dist} + \mathcal{L}_{contour}$.

Data-set: The data-set used in the experiments consists of CT scans from the gold corpus and silver corpus in VISCERAL dataset [11]. 74 CT scans from the silver corpus data were used for training, the annotations in this set are automatically labeled by fusing the results of multiple algorithms, yielding noisy labels. 23 CT scans from the gold corpus are used for testing, which contains manually annotated labels. Multi-organ segmentation imposes several challenges including the different fields of view (whole-body, thorax), modalities (contrast and non-contrast), severe class-imbalance across the target classes due to the diversity of organ’s size and shapes and the capability to generalize when trained with noisy labels.

All models were trained with Adam optimizer with a decaying learning rate initialized at 0.001, mini batch size of 4. Training was continued till validation loss converged. All the experiments were conducted on an NVIDIA Quadro P6000 GPU with 24GB vRAM. The code was developed in TensorFlow.

Table 1: Quantitative results: mean and standard deviation of the global dice scores shows that complementary-task learning achieves the best result.

Model	Dice
U-Net	0.8849 ± 0.120
U-Net + distance	0.8868 ± 0.116
U-Net + contour	0.8791 ± 0.118
U-Net + distance,contour	0.9018 ± 0.116




Fig. 4: Box plots of different organs show consistent improvement of the dice score using our proposed model (U-Net+distance+contour) over the baseline model(U-net). We also report the p -value obtained using Wilcoxon signed-rank statistical test between our proposed model and the baseline. The statistically significant p -values < 0.001 are shown in red.

To evaluate the overall accuracy of the proposed segmentation approach, we report the average dice score between the ground truth and the predicted segmentation map. Organ-specific dice scores are reported using box plots. We also report the Wilcoxon signed-rank test to find the statistical significance of the results.

Quantitative Results: We observe from Table 1 that adding distance task improves the overall dice while adding contour task alone does not show any improvement. We attribute this behavior to the fact that the training data-set contains noisy labels and therefore the external part of the organs are not exactly represented. On the other hand, adding both tasks drives the segmentation network to better generalization resulting in an improvement of 2% compared to the baseline. From the box plots in Fig 4 we observe that dice score for big organs does not improve significantly compared to the baseline. In contrast, complementary-task learning provides statistically significant improvement for small organs such as the spleen, thyroid gland, and trachea.




Fig. 5: Qualitative results: Each image shows the mid-slice in the coronal view. a) the input CT slice, b) ground truth segmentation, c) prediction from U-Net baseline, d) complementary-task learning. ROIs are indicated by green box highlighting regions where complementary-tasks learning improves the segmentation.

Qualitative Results: Fig 5 shows a qualitative comparison between the U-Net baseline and the proposed approach. We highlight ROIs with a green box, to show regions where complementary-task learning improves the segmentation. We can visually assess that particularly the small organs result in better segmentation. Additionally, Fig 3 compares the predicted distance map and the contour map to the ground truth. This confirms that the network is also able to solve the complementary-tasks.

4 Conclusions

In this work, we have proposed complementary-task learning to provide a novel and alternative solution to the challenging task of multi-organ segmentation. We have validated our method in a public benchmark data set which shows consistent improvement in dice score, especially for the small organs. In medical image segmentation where large data sets are scarce and corresponding dense annotation is expensive, designing complementary-task by leveraging existing target label could be beneficial to learn a generalized representation.

Acknowledgements

Fernando Navarro gratefully acknowledge the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - GRK for the financial support. Suprosanna Shit and Ivan Ezhov are supported by the Translational Brain Imaging Training Network (TRABIT) under the European Unions Horizon 2020 research & innovation programme (Grant agreement ID: 765148). The authors gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

References

1. Bischke, B., et al.: Multi-task learning for segmentation of building footprints with deep neural networks. arXiv preprint arXiv:1709.05932 (2017)
2. Chen, Z., et al.: Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In: Proceedings of the ICML (2018)
3. Geirhos, R., et al.: Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In: Proceedings of the ICLR (2019)
4. He, K., et al.: Mask r-cnn. In: Proceedings of the IEEE CVPR. pp. 2961–2969 (2017)
5. Kendall, A., et al.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: Proceedings of the IEEE CVPR. pp. 7482–7491 (2018)
6. Misra, I., et al.: Cross-stitch networks for multi-task learning. In: Proceedings of the IEEE CVPR. pp. 3994–4003 (2016)
7. Payer, C., et al.: Integrating spatial configuration into heatmap regression based cnns for landmark localization. *Medical Image Analysis* **54**, 207–219 (2019)
8. Ronneberger, O., et al.: U-net: Convolutional networks for biomedical image segmentation. In: Proceedings of the MICCAI (2015)
9. Shelhamer, E., et al.: Fully convolutional networks for semantic segmentation. *IEEE TPAMI*. (2017)
10. Song, Y., et al.: Progressive multi-atlas label fusion by dictionary evolution. *Medical Image Analysis* (2017)
11. Jimenez-del Toro, O., et al.: Cloud-based evaluation of anatomical structure segmentation and landmark detection algorithms: VISCERAL anatomy benchmarks. *IEEE TMI* **35**(11), 2459–2475 (2016)
12. Uslu, F., Bharath, A.A.: A multi-task network to detect junctions in retinal vasculature. In: Proceedings of the MICCAI. pp. 92–100. Springer (2018)
13. Wang, H., et al.: Multi-atlas segmentation with joint label fusion. *IEEE TMI* (2012)
14. Zhao, Y., et al.: Knowledge-aided convolutional neural network for small organ segmentation. *IEEE JBHI* (2019)