



Learning to Segment Anatomical Structures Accurately from One Exemplar

Yuhang Lu^{1,2(✉)}, Weijian Li^{1,3}, Kang Zheng¹, Yirui Wang¹,
Adam P. Harrison¹, Chihung Lin⁴, Song Wang², Jing Xiao⁵, Le Lu¹,
Chang-Fu Kuo⁴, and Shun Miao¹

¹ PAII Inc., Bethesda, MD, USA
yuhang@email.sc.edu

² University of South Carolina, Columbia, SC, USA

³ University of Rochester, Rochester, NY, USA

⁴ Chang Gung Memorial Hospital, Linkou, Taiwan, R.O.C.

⁵ Ping An Technology, Shenzhen, China

Abstract. Accurate segmentation of critical anatomical structures is at the core of medical image analysis. The main bottleneck lies in gathering the requisite expert-labeled image annotations in a scalable manner. Methods that permit to produce accurate anatomical structure segmentation without using a large amount of fully annotated training images are highly desirable. In this work, we propose a novel contribution of *Contour Transformer Network* (CTN), a one-shot anatomy segmentor including a naturally built-in human-in-the-loop mechanism. Segmentation is formulated by learning a contour evolution behavior process based on graph convolutional networks (GCN). Training of our CTN model requires only one labeled image exemplar and leverages additional unlabeled data through newly introduced loss functions that measure the global shape and appearance consistency of contours. We demonstrate that our one-shot learning method significantly outperforms non-learning-based methods and performs competitively to the state-of-the-art fully supervised deep learning approaches. With minimal human-in-the-loop editing feedback, the segmentation performance can be further improved and tailored towards the observer desired outcomes. This can facilitate the clinician designed imaging-based biomarker assessments (to support personalized quantitative clinical diagnosis) and outperforms fully supervised baselines.

Keywords: Contour Transformer Network · One-shot segmentation · Graph convolutional network

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-59710-8_66) contains supplementary material, which is available to authorized users.

1 Introduction

Obtaining manual image segmentation labels or masks has often been an obstacle in scaling up medical image segmentation applications. Without abundant pixel-level fully annotated image data, the state-of-the-art CNN-based segmentation methods cannot achieve their best performances [5, 7, 10, 21, 22, 27, 29]. However, annotating segmentation masks for medical images is very time-consuming and requires specialized expertise on human anatomy and its variations [28]. How to train an accurate segmentation model with less labeled data demands prompt solutions. In this paper, we tackle the problem on high-resolution anatomical structure X-ray images and propose *Contour Transformer Network*, allowing to learn from only one labeled exemplar image.

Several one-shot or few-shot segmentation methods have been proposed for natural images [8, 12, 19, 24, 33] by extracting information from a few support images to guide the segmentation of query images in testing. Nevertheless, the training process still relies on large-scale annotated datasets such as PASCAL VOC [9] and MS-COCO [16]. This condition renders them inapplicable directly to the medical imaging domain, because such equivalent datasets do not exist yet. Our problem is defined under a very different setting that only one pixel-level annotated training image instance is available. Other problem settings could be found in [20, 34] where they attempt to alleviate the label shortage problem via data augmentation. In contrast, our method is to train the segmentation model with only one labeled exemplar and a set of unlabeled images.

The main challenge of one-shot segmentation is the lack of ground truth image mask or contour labels. Regular training strategies of comparing predictions with ground truth labels are no longer applicable. We adopt a new training scheme in CTN. Because of the inherent regularized nature of anatomical structures, the same anatomy in different (X-ray) images may share some common features or properties, such as the anatomical structure’s *shape*, *appearance* and *gradients* along the structural object boundary. Although different images are not directly comparable, we can compare their common features only and use the exemplar segmentation to guide other unlabeled images partially, thus making CTN trainable in a one-shot setting.

To leverage these shared anatomical properties, we represent the image segmentation problem as learning a contour evolution behavior. Thus three differentiable contour-based loss functions are proposed to describe the common features. For each unlabeled image, CTN takes the exemplar contour as an initialization, then gradually evolves it to minimize the weighted loss. Furthermore, we offer a naturally built-in human-in-the-loop mechanism to allow CTN to learn from extra partial labels. If any part in the predicted contour is inaccurate, users can correct them by drawing line segments, then CTN will format these corrections as partial contours and incorporate them back into the training via an additional Chamfer loss. In this way, we can improve and refine the segmentation performance with minimum annotation costs.

In summary, our contributions are three folds. (1) We propose a CNN-based image segmentation framework that could be trained with only one labeled

image. (2) We describe the contour perceptual loss and the contour bending loss as two new optimization loss functions, to measure the similarity of two contours in terms of the appearance or shape cues, respectively. (3) We demonstrate that CTN achieves the state-of-the-art one-shot segmentation results; performs competitively when compared to fully supervised alternatives; and can outperform them with minimal human-in-the-loop feedback, on three datasets.

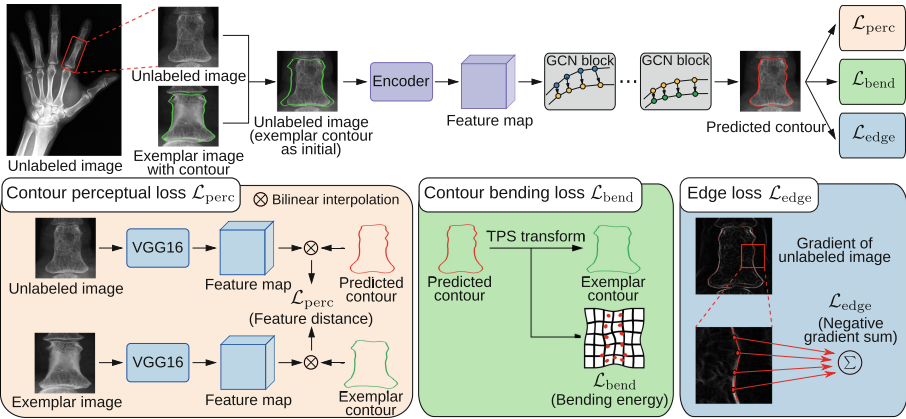


Fig. 1. Contour Transformer Network. CTN is trained to fit a contour to the object boundary by learning from one labeled exemplar. It takes the exemplar and an unlabeled image as input, and predict a contour that has similar contour features with the exemplar. Three losses are proposed to make this network “one-shot” trainable.

2 Methods

The problem of anatomical structure segmentation on images can be decomposed into two steps: ROI (Region of Interest) cropping; and ROI segmentation. ROI detection has been well-studied in past literature [4, 6, 15, 29, 31, 32], so we focus on achieving very high segmentation accuracy by taking the detected/cropped ROI (with noise and errors) as input.

Assuming that a set of images \mathbf{I} contains the same type of anatomical structure and only one of them is labeled, called the *exemplar*. Our goal is to learn a segmentation model for this structure from \mathbf{I} . As mentioned above, we frame image segmentation as a process of contour evolution. Each contour is represented by N uniformly spaced vertices. Denote the exemplar image and its contour by I_E and C_E , respectively. For any unlabeled image $I \in \mathbf{I}$, its contour is $C = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N\}$. The exemplar contour is placed at the center of I as the initial location of C ; next CTN is employed to estimate the point-wise offsets from the initial to the correct location, which is formulated by

$F_{\theta}(I, I_E, C_E) = \{\Delta\mathbf{p}_1, \Delta\mathbf{p}_2, \dots, \Delta\mathbf{p}_N\}$ where F_{θ} denotes the CTN model with weights θ .

Inspired by [17], we use a *CNN-GCN network architecture* to model the contour evolution in CTN. From Fig. 1, CTN consists of two parts of an image encoding block and cascaded contour evolution blocks. It takes an unlabeled image I , an exemplar image I_E and its ground truth contour C_E as input, and predicts the contour C of I . (1) We first place C_E at the center of I as the initial location of C , then the encoder outputs a feature map encoding the local image appearance of I . ResNet-50 [11] is used as the backbone of CNN encoder. (2) The cascaded GCN blocks are then employed to evolve the contour C step by step. The GCN takes the contour graph with vertex features as input. Each vertex in the contour is connected to four neighboring vertices, two on each side. These vertex features are extracted from the feature map of I at the vertex locations via interpolation. Each GCN block takes the output contour of the previous block and updates it by predicting the point-wise coordinate offsets. We use five GCN blocks with the same multi-layer GCN architecture, although weights are not shared. The output of the 5th block is the predicted contour of CTN. (3) Three one-shot trainable losses are utilized to optimize CTN, as the contour perceptual loss L_{perc} , the contour bending loss L_{bend} and the edge loss L_{edge} . The total loss of CTN in the one-shot setting is written as: $L = \lambda_1 L_{perc} + \lambda_2 L_{bend} + \lambda_3 L_{edge}$ where $\lambda_1, \lambda_2, \lambda_3$ are the weighting factors of the three losses. We describe the three employed losses in detail as follows.

Contour Perceptual Loss. We propose a new contour perceptual loss to measure the *appearance dissimilarity* between the visual patterns of the exemplar contour C_E and the predicted contour C , on the exemplar image I_E or the target image I , respectively. Partially motivated by the original perceptual loss [14] developed for image super-resolution, modeling the image perceptual similarities in the feature space of VGG-Net [26], we measure the contour perceptual similarities in the graph feature space. In particular, graph features are extracted from the ImageNet pre-trained VGG-16 feature maps of the two images along the two contours, and their L1 distance is calculated as the contour perceptual loss: $L_{perc} = \sum_{i=1, \dots, N} \|P(\mathbf{p}_i) - P_E(\mathbf{p}'_i)\|_1$ where $\mathbf{p}_i \in C$, $\mathbf{p}'_i \in C_E$, and P and P_E denote the VGG-16 features of I and I_E , respectively. The VGG-16 baseline network weights are trained on ImageNet dataset [23].

The contour perceptual loss is used to guide the evolution of the contour in CTN, by having several advantages. (1) Since VGG-16 network features can capture the image pattern of a neighboring area with spatial contexts (*i.e.*, network receptive field), the contour perceptual loss enjoys a relatively large capturing range (*i.e.*, the convex region around the minimum), making the CTN training optimization easier. (2) The backbone VGG-16 model is trained on ImageNet [23] for classification tasks, so that its learned features are less sensitive to noises and illumination variations, which also benefits the training of CTN.

Contour Bending Loss. If we operate under the assumption that an exemplar contour is broadly informative to other data samples, then it should be beneficial to use the exemplar shape to ground any predictions on other samples. To this

end, we propose a novel contour bending loss to measure the *shape dissimilarity* between contours. The loss is calculated as the bending energy of thin-plate spline (TPS) warping [1] that maps C_E to C . It is worth noting that TPS warping achieves the minimum bending energy among all warpings that map C_E to C . Since the bending energy measures the magnitude of the second order derivatives of the warping function, the contour bending loss penalizes more on the local and acute shape changes, often associated with mis-segmentations. Given C and C_E , the TPS bending energy can be calculated as follows. Define $\mathbf{p}_i = (x_i, y_i)$, $\mathbf{p}'_i = (x'_i, y'_i)$, and $\mathbf{K} = \left(\|\mathbf{p}'_i - \mathbf{p}'_j\|_2^2 \cdot \log \|\mathbf{p}'_i - \mathbf{p}'_j\|_2 \right)$, $\mathbf{P} = (\mathbf{1}, \mathbf{x}', \mathbf{y}')$, $\mathbf{L} = \begin{bmatrix} \mathbf{K} & \mathbf{P} \\ \mathbf{P}^T & \mathbf{0} \end{bmatrix}$ where $\mathbf{x}' = \{x'_1, x'_2, \dots, x'_N\}^T$, $\mathbf{y}' = \{y'_1, y'_2, \dots, y'_N\}^T$. The TPS bending energy is written as $L_{bend} = \max \left[\frac{1}{8\pi} (\mathbf{x}^T \mathbf{H} \mathbf{x} + \mathbf{y}^T \mathbf{H} \mathbf{y}), 0 \right]$ where $\mathbf{x} = \{x_1, x_2, \dots, x_N\}^T$, $\mathbf{y} = \{y_1, y_2, \dots, y_N\}^T$, and \mathbf{H} is the $N \times N$ upper left submatrix of \mathbf{L}^{-1} [30].

Edge Loss. Although the contour perceptual and bending losses can achieve robust segmentation, they are inherently insensitive to (very) small segmentation fluctuations, such as minimal deviations from the correct boundary by a few pixels. Therefore, in order to obtain desirably high segmentation accuracy and adequately facilitate the downstream workflows like rheumatoid arthritis quantification [13], we employ an edge loss that measures the image gradient magnitude along the computed contour and attracts the contour toward edges in the image naturally. The edge loss is written as: $L_{edge} = -\frac{1}{N} \sum_{\mathbf{p} \in C} \|\nabla I(\mathbf{p})\|_2$ where ∇ is the gradient operator.

2.1 Human-in-the-Loop

More labels are always helpful to enhance the model’s generalization ability and robustness, if available. Benefiting from the contour-based setting, CTN offers a natural way to incorporate additional user labels with a human-in-the-loop mechanism. Assuming that we have a CTN model trained with one exemplar image, we intend to finetune it with more segmentation annotations. We run this model on a set of unlabeled images first, and select any number of images with inaccurate predictions as new instances. Instead of drawing the whole contour from scratch on these new images, the annotator only needs to redraw some partial contours, to correct the previously undesirable predictions. The point-wise training of CTN makes it feasible to learn from these partial corrections.

A **partial contour matching loss** is proposed to utilize the partial ground truth contours during the CTN training. Denote $\hat{\mathbf{C}}$ as a set of partial contours in image I , each element of which is an individual contour segment. For each contour segment $\hat{C}_i \in \hat{\mathbf{C}}$, we build the point correspondence between \hat{C}_i and C . For each \hat{C}_i , we find two points in the predicted contour C , closest to the start and end points of \hat{C}_i , then each predicted point between the two points are assigned to the closest corrected point. Denote the corresponding predicted contour segment by C_i ($C_i \in C$). We define the distance between C and \hat{C}_i as the Chamfer distance from C_i to \hat{C}_i : $D(\hat{C}_i, C) = \sum_{\mathbf{p} \in C_i} \min_{\hat{\mathbf{p}} \in \hat{C}_i} \|\mathbf{p} - \hat{\mathbf{p}}\|_2$ and the partial matching loss of C is defined as $L_{pcm} = \frac{1}{N} \sum_{\hat{C}_i \in \hat{\mathbf{C}}} D(\hat{C}_i, C)$. In the

human-in-the-loop scenario, we combine all losses to train the CTN, and rewrite the loss function as $\hat{L} = \lambda_1 L_{perc} + \lambda_2 L_{bend} + \lambda_3 L_{edge} + \lambda_4 L_{pcm}$, which allows CTN to be trained with fully labeled, partially labeled and unlabeled images simultaneously and seamlessly. Whenever new labels are available, we can use \hat{L} to finetune the one-shot CTN model.

3 Experimental Results

Datasets. We evaluate our method on three X-ray image datasets of knee, lung and phalanx, respectively. The *knee* dataset contains 212 knee X-ray images from the Osteoarthritis Initiative (OAI) database¹, 100 for training and 112 for testing. The *lung* dataset is the public JSRT [25] of 247 posterior-anterior chest radiographs, 124 for training and 123 for testing. The *phalanx* dataset comes from hand X-ray images of patients with rheumatoid arthritis. 202 ROIs of proximal phalanx are extracted automatically from hand joint detection [13], randomly split into 100 training and 102 testing images.

Table 1. Performances of CTN and seven existing methods on three datasets.

Method		Knee		Lung		Phalanx	
		IoU(%)	HD(px)	IoU(%)	HD(px)	IoU(%)	HD(px)
Non- learning	MorphACWE [2, 18]	65.89	54.07	76.09	55.35	74.33	69.13
	MorphGAC [3, 18]	87.42	15.78	70.79	45.67	82.15	24.73
One-shot	CANet [33]	29.22	175.86	56.90	73.46	60.90	67.13
	Brainstorm [34]	90.17	29.07	77.13	43.28	80.05	30.30
	CTN (Ours)	97.32	6.01	94.75	12.16	96.96	8.19
Fully supervised	UNet [21]	96.60	7.14	95.38	12.48	95.76	10.10
	DeepLab [5]	97.18	5.41	96.18	10.81	97.63	6.52
	HRNet [29]	96.99	5.18	95.99	10.44	97.47	7.03

We evaluate the accuracy of segmentation masks by Intersection-over-Union (IoU) metric and the corresponding object contour distance by the Hausdorff distance (HD). For comparative methods not explicitly outputting the anatomy contours, we extract the external contour of the largest segmented image region. The hyperparameters are $N = 1000$, $\lambda_1 = 1$, $\lambda_2 = 0.25$, $\lambda_3 = 0.1$, $\lambda_4 = 1$. All networks are trained using Adam optimizer with a learning rate of 1×10^{-4} , a weight decay of 1×10^{-4} and a batch size of 12 for 500 epochs. The one-shot training and human-in-the-loop finetuning settings are the same.

The proposed CTN is compared with seven previous methods. The quantitative results are reported in Table 1; qualitative results are given in Fig. 2.

¹ <https://nda.nih.gov/oai/>.

Comparison with Non-Learning-Based Methods. We first compare with two non-learning based methods of MorphACWE [2,18] and MorphGAC [3,18]. Both approaches are based on active contour models (ACMs), which evolves an initial contour to the object by minimizing an energy function. The initial contour is the same as ours. Quantitative results in Table 1 show that our method significantly outperforms them. Specifically, in average CTN achieves 16.22% higher IoU and 19.94 pixels less in HD than MorphGAC, the better of the two. Visualization results in Fig. 2 show that these two approaches cannot localize anatomical structures accurately, especially when the boundary of such structures are not clearly contrasted, such as in the lung image. Both methods are based on ACMs and predict contours by minimizing some hand-crafted energy functions for a *single* image. In contrast, CTN learns from an exemplar contour to guide the contour transformation for all images in the entire training set.




Fig. 2. Segmentation results of three example images using eight methods. From top to bottom, these images are from the knee, lung and phalanx testing sets, respectively. The ground truth boundaries are drawn in **green** line for the ease of comparison. (Color figure online)

Comparison With Other One-Shot Methods. Next, we compare with two state-of-the-art one-shot segmentation methods: CANet [33] and Brainstorm [34]. All one-shot approaches (including ours) use the same exemplar image. For each training set, we compare the distance of each image to all other images in the VGG network feature space and the exemplar is selected to be the image with the smallest distance. CANet is proposed to perform one-shot segmentation on unseen objects in testing, but it requires a fully annotated dataset to train. Such that, we use the specific model [33] trained on PASCAL VOC 2012 dataset for comparison. From Table 1, CANet achieves only 49.01% IoU on average. We speculate that the poor performance is caused by the domain gap between natural images and medical images. Brainstorm [34] learns an image augmentation model and is one-shot trainable. We follow its default procedure

to train the segmentation models on three datasets. It yields reasonable results with averaged 82.45% IoU and 34.22 HD, but still dramatically lower than ours.

Comparison with Fully Supervised Methods. Last, we evaluate and compare the performance of three fully supervised methods: UNet [21], DeepLab v3+ [5] and HRNet W18 [29]. All of them are trained for 500 epochs using all training image annotations in three datasets. On average, CTN performs comparably or better than UNet, and falls behind DeepLab (the best of the three), by 0.66% in IoU and 1.21 pixels in HD, respectively. It demonstrates that CTN while using only one training image can usually compete head-to-head with the state-of-the-art fully supervised baselines [5, 21, 29]. Note that the heatmap-based segmentation methods predict the per-pixel labels, which could cause the loss of integrity of object boundaries, *e.g.*, some small “islands” in the lung masks of Fig. 2. On the other hand, CTN naturally retains the integrity of the object segmentation, as an important aspect in assessing visual segmentation quality.

Incorporating Simulated Human Corrections. To evaluate the effectiveness of the human-in-the-loop mechanism, we empirically simulate different degrees of human-computer interactions. For each dataset, we first train a CTN model with the default exemplar and run inference on the training set. We sort all training images by their HD segmentation errors from high to low. Three subsets are formed by selecting the top 10%, 25% or 100% training images; and fine-tune the initial one-shot CTN model using these training subsets augmented by the ground truth contours, respectively. This protocol results in four CTN models. From Fig. 3, we observe that CTN consistently improves with more human corrections. Specifically, when using 25% such corrected samples, CTN starts to outperform DeepLab using all training images (IoUs of 97.17% vs 97.0%, and HDs of 7.01 vs 7.58). With all samples, CTN reaches 97.33% on IoU and 6.5 on HD. These results indicate that the human-in-the-loop mechanism can potentially help CTN achieves better performance than fully supervised methods with considerably less annotation efforts.




Fig. 3. Using 0, 10%, 25% and 100% human corrections to finetune the one-shot CTN model, respectively (“0” means no finetuning).

Ablation Study. We conduct an ablation experiment to validate the effectiveness of the three proposed losses. The results are shown in Table 2 where the performance of our method is indeed impaired if any loss is removed, with mean

IoU reductions of 4.32%, 4.12%, and 1.72% for L_{perc} , L_{bend} , and L_{edge} , respectively. This validates the necessity of all three losses. An exception is the knee dataset when L_{bend} is removed. Knee X-ray images share similar appearance features along the contour so that they can be segmented robustly with just the contour perceptual loss and edge loss. Thus, adding contour bending loss leads to statistically insignificant decreases (*i.e.*, IoUs of 97.32% vs 97.50%, HDs of 6.01 vs 5.87) in this particular scenario. However, such a regularization effect by the contour bending loss is generally desired to alleviate the worst-case scenarios and proves useful in the other two datasets.

Table 2. Ablation study. Remove one loss each time and re-train the model.

L_{perc}	L_{bend}	L_{edge}	Knee		Lung		Phalanx	
			IoU(%)	HD(px)	IoU(%)	HD(px)	IoU(%)	HD(px)
	✓	✓	94.62	8.28	87.45	26.51	94.01	15.81
✓		✓	97.50	5.87	84.93	36.74	94.24	26.13
✓	✓		94.43	11.90	92.99	16.22	96.45	9.84
✓	✓	✓	97.32	6.01	94.75	12.17	96.96	8.19

4 Conclusion

In this paper, we propose a novel one-shot segmentation method, Contour Transformer Network, which takes one labeled exemplar and a set of unlabeled images to train a segmentation model for anatomical structures in medical images. The key idea that enables one-shot training is to guide the segmentation of unlabeled images by utilizing their shared features with the exemplar image but not ground truth masks. Experiments on three datasets demonstrate that CTN performs competitively to state-of-the-art fully supervised approaches and outperforms them with minimal human corrections. Although CTN is for anatomical structures, the idea of one-shot training is also applicable to other images with shared features. In the future, we will explore its application in more medical image analysis problems.

References

1. Bookstein, F.L.: Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.* **11**(6), 567–585 (1989)
2. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. *Int. J. Comput. Vis.* **22**(1), 61–79 (1997)
3. Chan, T.F., Vese, L.A.: Active contours without edges. *IEEE Trans. Image Process.* **10**(2), 266–277 (2001)
4. Chen, H., et al.: Anatomy-aware Siamese network: exploiting semantic asymmetry for accurate pelvic fracture detection in x-ray images. *arXiv preprint arXiv:2007.01464* (2020)

5. Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11211, pp. 833–851. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01234-2_49
6. Chen, R., Ma, Y., Chen, N., Lee, D., Wang, W.: Cephalometric landmark detection by attentive feature pyramid fusion and regression-voting. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11766, pp. 873–881. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32248-9_97
7. Dolz, J., Gopinath, K., Yuan, J., Lombaert, H., Desrosiers, C., Ayed, I.B.: Hyperdense-net: a hyper-densely connected CNN for multi-modal image segmentation. *IEEE Trans. Med. Imaging* **38**(5), 1116–1126 (2018)
8. Dong, N., Xing, E.: Few-shot semantic segmentation with prototype learning. In: Proceedings of the British Machine Vision Conference, vol. 1, p. 6 (2018)
9. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.* **88**(2), 303–338 (2010)
10. Harrison, A.P., Xu, Z., George, K., Lu, L., Summers, R.M., Mollura, D.J.: Progressive and multi-path holistically nested neural networks for pathological lung segmentation from CT images. In: Descoteaux, M., Maier-Hein, L., Franz, A., Janin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10435, pp. 621–629. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_71
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
12. Hu, R., Dollár, P., He, K., Darrell, T., Girshick, R.: Learning to segment every thing. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4233–4241 (2018)
13. Huo, Y., Vincken, K.L., van der Heijde, D., De Hair, M.J., Lafeber, F.P., Viergever, M.A.: Automatic quantification of radiographic finger joint space width of patients with early rheumatoid arthritis. *IEEE Trans. Biomed. Eng.* **63**(10), 2177–2186 (2015)
14. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9906, pp. 694–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_43
15. Li, W., et al.: Structured landmark detection via topology-adapting deep graph learning. arXiv preprint [arXiv:2004.08190](https://arxiv.org/abs/2004.08190) (2020)
16. Lin, T., Maire, M., et al.: Microsoft COCO: common objects in context. arXiv preprint [arXiv:1405.0312](https://arxiv.org/abs/1405.0312) (2014)
17. Ling, H., Gao, J., Kar, A., Chen, W., Fidler, S.: Fast interactive object annotation with curve-GCN. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5257–5266 (2019)
18. Marquez-Neila, P., Baumela, L., Alvarez, L.: A morphological approach to curvature-based evolution of curves and surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(1), 2–17 (2013)
19. Michaelis, C., Ustyuzhaninov, I., Bethge, M., Ecker, A.S.: One-shot instance segmentation. arXiv preprint [arXiv:1811.11507](https://arxiv.org/abs/1811.11507) (2018)
20. Oliveira, A., Pereira, S., Silva, C.A.: Augmenting data when training a CNN for retinal vessel segmentation: How to warp? In: IEEE 5th Portuguese Meeting on Bioengineering, pp. 1–4 (2017)

21. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
22. Roth, H.R., et al.: Spatial aggregation of holistically-nested convolutional neural networks for automated pancreas localization and segmentation. *Med. Image Anal.* **45**, 94–107 (2018)
23. Russakovsky, O., et al.: ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**(3), 211–252 (2015)
24. Shaban, A., Bansal, S., Liu, Z., Essa, I., Boots, B.: One-shot learning for semantic segmentation. arXiv preprint [arXiv:1709.03410](https://arxiv.org/abs/1709.03410) (2017)
25. Shiraishi, J., et al.: Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists' detection of pulmonary nodules. *Am. J. Roentgenol.* **174**(1), 71–74 (2000)
26. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
27. Sinha, A., Dolz, J.: Multi-scale guided attention for medical image segmentation. arXiv preprint [arXiv:1906.02849](https://arxiv.org/abs/1906.02849) (2019)
28. Tajbakhsh, N., Jeyaseelan, L., Li, Q., Chiang, J., Wu, Z., Ding, X.: Embracing Imperfect Datasets: A Review of Deep Learning Solutions for Medical Image Segmentation (2019)
29. Wang, J., et al.: Deep high-resolution representation learning for visual recognition. arXiv preprint [arXiv:1908.07919](https://arxiv.org/abs/1908.07919) (2019)
30. Wang, S., Munsell, B., Richardson, T.: Correspondence establishment in statistical shape modeling: Optimization and evaluation. In: *Statistical Shape and Deformation Analysis*, pp. 67–87. Elsevier (2017)
31. Wang, Y., et al.: Weakly supervised universal fracture detection in pelvic x-rays. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11769, pp. 459–467. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32226-7_51
32. Wu, W., Qian, C., Yang, S., Wang, Q., Cai, Y., Zhou, Q.: Look at boundary: a boundary-aware face alignment algorithm. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2129–2138 (2018)
33. Zhang, C., Lin, G., Liu, F., Yao, R., Shen, C.: CANet: class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5217–5226 (2019)
34. Zhao, A., Balakrishnan, G., Durand, F., Gutttag, J.V., Dalca, A.V.: Data augmentation using learned transformations for one-shot medical image segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8543–8553 (2019)