

Spatial Aggregation of Holistically-Nested Networks for Automated Pancreas Segmentation

Holger R. Roth, Le Lu, Amal Farag,
Andrew Sohn, and Ronald M. Summers[✉]

Imaging Biomarkers and Computer-Aided Diagnosis Laboratory,
Radiology and Imaging Sciences, National Institutes of Health,
Bethesda, MD 20892-1182, USA
rms@nih.gov

Abstract. Accurate automatic organ segmentation is an important yet challenging problem for medical image analysis. The pancreas is an abdominal organ with very high anatomical variability. This inhibits traditional segmentation methods from achieving high accuracies, especially compared to other organs such as the liver, heart or kidneys. In this paper, we present a holistic learning approach that integrates semantic mid-level cues of deeply-learned organ interior and boundary maps via robust spatial aggregation using random forest. Our method generates boundary preserving pixel-wise class labels for pancreas segmentation. Quantitative evaluation is performed on CT scans of 82 patients in 4-fold cross-validation. We achieve a (mean \pm std. dev.) Dice Similarity Coefficient of $78.01\% \pm 8.2\%$ in testing which significantly outperforms the previous state-of-the-art approach of $71.8\% \pm 10.7\%$ under the same evaluation criterion.

1 Introduction

Pancreas segmentation in computed tomography (CT) is challenging for current computer-aided diagnosis (CAD) systems. While automatic segmentation of numerous other organs in CT scans such as liver, heart or kidneys achieves good performance with Dice Similarity Coefficients (DSC) of $>90\%$ [1–3], the pancreas’ variable shape, size and location in the abdomen limits segmentation accuracy to $<73\%$ DSC being reported in the literature [3–6]. Deep convolutional Neural Networks (CNNs) have successfully been applied to many high-level tasks in medical imaging, such as recognition and object detection. The main advantage of CNNs comes from the fact that end-to-end learning of salient feature representations for the task at hand is more effective than hand-crafted features with heuristically tuned parameters [7]. Similarly, CNNs demonstrate promising performance for pixel-level labeling problems, e.g., semantic segmentation [6, 8–10]. Recent work in computer vision and biomedical imaging including fully convolutional neural networks (FCN) [9], DeepLab [10] and U-Net [8],

The rights of this work are transferred to the extent transferable according to title 17 § 105 U.S.C.

have gained significant improvements in performance over previous methods by applying state-of-the-art CNN based image classifiers and representation to the semantic segmentation problem in both domains. Semantic organ segmentation involves assigning a label to each pixel in the image. On one hand, features for classification of single pixels (or patches) play a major role, but on the other hand, factors such as edges (i.e., organ boundaries), appearance consistency and spatial consistency, could greatly impact the overall system performance [7]. Furthermore, there are indications of semantic vision tasks requiring hierarchical levels of visual perception and abstraction [11]. As such, generating rich feature hierarchies for both the interior and the boundary of the organ could provide important “mid-level visual cues” for semantic segmentation. Subsequent spatial aggregation of these mid-level cues then has the prospect of improving semantic segmentation methods by enhancing the accuracy and consistency of pixel-level labeling.

2 Methods

In this work, we present a holistic semantic segmentation method for organ segmentation in CT which incorporates deeply learned organ interior and boundary mid-level cues with subsequent spatial aggregation. This approach to organ segmentation is completely based on machine-learning principles. No multi-atlas registration and label fusion methods are employed. Our methods are evaluated on CT scans of 82 patients in 4-fold cross-validation (instead of “leave-one-patient-out” evaluation often used in other work [1–3]).

Candidate Region Generation: As a form of initialization, we employ a previously proposed method based on random forest (RF) classification [6] to compute a candidate bounding box regions. We only operate the RF labeling at a low probability threshold of >0.5 which is sufficient to reject the vast amount of non-pancreas from the CT images. This initial candidate generation is sufficient to extract bounding box regions that completely surround the pancreases in all used cases with nearly 100 % recall. All candidate regions are computed during the testing phase of cross-validation (CV) as in [6]. Note that candidate region proposal is not the focus of this work and assumed to be fixed for the rest of this study. This part could be replaced by other means of detecting an initial bounding box for pancreas detection, e.g., by RF regression [12] or sliding-window CNNs [6].

Semantic Mid-Level Segmentation Cues: We show that organ segmentation can benefit from multiple mid-level cues, like organ interior and boundary predictions. We investigate deep-learning based approaches to independently learn the pancreas’ interior and boundary mid-level cues. Combining both cues via learned spatial aggregation can elevate the overall performance of this semantic segmentation system. Organ boundaries are a major mid-level cue for defining

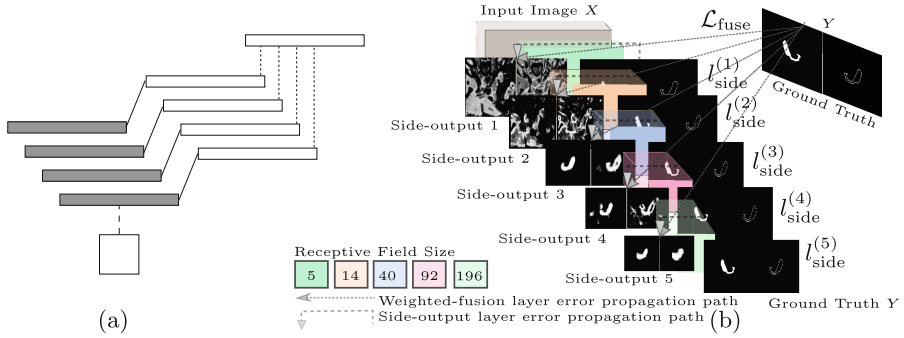


Fig. 1. Schematics of (a) the holistically-nested nets, in which multiple side outputs are added, and (b) the **HNN-I/B** network architecture for both interior (left images) and boundary (right images) detection pathways. We highlight the error back-propagation paths to illustrate the deep supervision performed at each side-output layer after the corresponding convolutional layer. As the side-outputs become smaller, the receptive field sizes get larger. This allows HNN to combine multi-scale and multi-level outputs in a learned weighted fusion layer (Figures adapted from [11] with permission).

and delineating the anatomy of interest. It could prove to be essential for accurate semantic segmentation of an organ.

Holistically-Nested Nets: In this work, we explicitly learn the pancreas’ interior and boundary image-labeling models via Holistically-Nested Networks (HNN). Note that this type of CNN architecture was first proposed by [11] under the name “holistically-nested edge detection” as a deep learning based general image edge detection method. We however find that it can be a suitable method for segmenting the interior of organs as well (see Sect. 3). HNN tries to address two important issues: (1) training and prediction on the whole image end-to-end (holistically) using a per-pixel labeling cost; and (2) incorporating multi-scale and multi-level learning of deep image features [11] via auxiliary cost functions at each convolutional layer. HNN computes the image-to-image or pixel-to-pixel prediction maps (from any input raw image to its annotated labeling map). The per-pixel labeling cost function [9, 11] offers the good feasibility that HNN/FCN can be effectively trained using only several hundred annotated image pairs. This enables the automatic learning of rich hierarchical feature representations (contexts) that are critical to resolve spatial ambiguity in the segmentation of organs. The network structure is initialized based on an ImageNet pre-trained VGGNet model [13]. It has been shown that fine-tuning CNNs pre-trained on the general image classification task (ImageNet) is helpful to low-level tasks, e.g., edge detection [11].

Network Formulation: Our training data $S^{I/B} = \{(X_n, Y_n^{I/B}), n = 1, \dots, N\}$ is composed of cropped axial CT images X_n (rescaled to within $[0, \dots, 255]$ with a soft-tissue window of $[-160, 240]$ HU); and $Y_n^I \in \{0, 1\}$ and $Y_n^B \in \{0, 1\}$ denote the (binary) ground truths of the interior and boundary map of the pancreas, respectively, for any corresponding X_n . Each image is considered holistically and independently as in [11]. The network is able to learn features from these images alone from which interior (**HNN-I**) boundary (**HNN-B**) predication maps can be produced. HNN can efficiently generate multi-level image features due to its deep architecture. Furthermore, multiple stages with different convolutional strides can capture the inherent scales of (organ edge/interior) labeling maps. However, due to the difficulty of learning such deep neural networks with multiple stages from scratch, we use the pre-trained network provided by [11] and fine-tuned to our specific training data sets $S^{I/B}$ with a relatively smaller learning rate of 10^{-6} . We use the HNN network architecture with 5 stages, including strides of 1, 2, 4, 8 and 16, respectively, and with different receptive field sizes as suggested by the authors¹. In addition to standard CNN layers, a HNN network has M side-output layers as shown in Fig. 1. These side-output layers are also realized as classifiers in which the corresponding weights are $\mathbf{w} = (\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(M)})$. For simplicity, all standard network layer parameters are denoted as \mathbf{W} . Hence, the following objective function can be defined²: $\mathcal{L}_{\text{side}}(\mathbf{W}, \mathbf{w}) = \sum_{m=1}^M \alpha_m l_{\text{side}}^{(m)}(\mathbf{W}, \mathbf{w}^{(m)})$. Here, l_{side} denotes an image-level loss function for side-outputs, computed over all pixels in a training image pair X and Y . Because of the heavy bias towards non-labeled pixels in the ground truth data, [11] introduces a strategy to automatically balance the loss between positive and negative classes via a per-pixel class-balancing weight β . This allows to offset the imbalances between edge/interior ($y = 1$) and non-edge/exterior ($y = 0$) samples. Specifically, a class-balanced cross-entropy loss function can be used with j iterating over the spatial dimensions of the image:

$$l_{\text{side}}^{(m)}(\mathbf{W}, \mathbf{w}^{(m)}) = -\beta \sum_{j \in Y_+} \log \Pr(y_j = 1|X; \mathbf{W}, \mathbf{w}^{(m)}) - (1 - \beta) \sum_{j \in Y_-} \log \Pr(y_j = 0|X; \mathbf{W}, \mathbf{w}^{(m)}). \quad (1)$$

Here, β is simply $|Y_-|/|Y|$ and $1 - \beta = |Y_+|/|Y|$, where $|Y_-|$ and $|Y_+|$ denote the ground truth set of *negatives* and *positives*, respectively. The class probability $\Pr(y_j = 1|X; \mathbf{W}, \mathbf{w}^{(m)}) = \sigma(a_j^{(m)}) \in [0, 1]$ is computed on the activation value at each pixel j using the sigmoid function $\sigma(\cdot)$. Now, organ edge/interior map predictions $\hat{Y}_{\text{side}}^{(m)} = \sigma(\hat{A}_{\text{side}}^{(m)})$ can be obtained at each side-output layer, where $\hat{A}_{\text{side}}^{(m)} \equiv \{a_j^{(m)}, j = 1, \dots, |Y|\}$ are activations of the side-output of layer m . Finally, a “weighted-fusion” layer is added to the network that can be simultaneously learned during training. The loss function at the fusion layer L_{fuse} is defined as $\mathcal{L}_{\text{fuse}}(\mathbf{W}, \mathbf{w}, \mathbf{h}) = \text{Dist}(Y, \hat{Y}_{\text{fuse}})$, where $\hat{Y}_{\text{fuse}} \equiv \sigma\left(\sum_{m=1}^M h_m \hat{A}_{\text{side}}^{(m)}\right)$

¹ <https://github.com/s9xie/hed>.

² We follow the notation of [11].

with $h = (h_1, \dots, h_M)$ being the fusion weight. $Dist(.,.)$ is a distance measure between the fused predictions and the ground truth label map. We use cross-entropy loss for this purpose. Hence, the following objective function can be minimized via standard stochastic gradient descent and back propagation:

$$(\mathbf{W}, \mathbf{w}, \mathbf{h})^* = \operatorname{argmin} (\mathcal{L}_{\text{side}}(\mathbf{W}, \mathbf{w}) + \mathcal{L}_{\text{fuse}}(\mathbf{W}, \mathbf{w}, \mathbf{h})) \quad (2)$$

Testing Phase: Given image X , we obtain both interior (HNN-I) and boundary (HNN-B) predictions from the models’ side output layers and the weighted-fusion layer as in [11]:

$$\left(\hat{Y}_{\text{fuse}}^{I/B}, \hat{Y}_{\text{side}}^{I_1/B_1}, \dots, \hat{Y}_{\text{side}}^{I_M/B_M} \right) = \mathbf{HNN-I/B} (X, (\mathbf{W}, \mathbf{w}, \mathbf{h})). \quad (3)$$

Learning Organ-Specific Segmentation Object Proposals: “Multiscale Combinatorial Grouping” (MCG³) [14] is one of the state-of-the-art methods for generating segmentation object proposals in computer vision. We utilize this approach to generate organ-specific superpixels based on the learned boundary prediction maps **HNN-B**. Superpixels are extracted via continuous oriented watershed transform at three different scales ($\hat{Y}_{\text{side}}^{B_2}, \hat{Y}_{\text{side}}^{B_3}, \hat{Y}_{\text{fuse}}^B$) supervisedly learned by **HNN-B**. This allows the computation of a hierarchy of superpixel partitions at each scale, and merges superpixels across scales thereby, efficiently exploring their combinatorial space [14]. This, then, allows MCG to group the merged superpixels toward object proposals. We find that the first two levels of object MCG proposals are sufficient to achieve $\sim 88\%$ DSC (see Table 1 and Fig. 2), with the optimally computed superpixel labels using their spatial overlapping ratios against the segmentation ground truth map. All *merged* superpixels \mathcal{S} from the first two levels are used for the subsequently proposed spatial aggregation of **HNN-I** and **HNN-B**.

Spatial Aggregation with Random Forest: We use the superpixel set \mathcal{S} generated previously to extract features for spatial aggregation via random forest classification⁴. Within any superpixel $s \in \mathcal{S}$ we compute simple statistics including the 1st-4th order moments and 8 percentiles [20%, 30%, ..., 90%] on CT, **HNN-I**, and **HNN-B**. Additionally, we compute the mean x , y , and z coordinates normalized by the range of the 3D candidate region (Sect. 2). This results in 39 features describing each superpixel and are used to train a random forest classifier on the training *positive* or *negative* superpixels at each round of 4-fold CV. Empirically, we find 50 trees to be sufficient to model our feature set. A final 3D pancreas segmentation is simply obtained by stacking each slice prediction back into the space of the original CT images. No further post-processing is employed and spatial aggregation of **HNN-I** and **HNN-B** maps for superpixel classification is already of high quality. This complete pancreas segmentation model is denoted as **HNN-I/B-RF** or **HNN-RF**.

³ <https://github.com/jponttuset/mcg>.

⁴ Using MATLAB’s `TreeBagger()` class.




Fig. 2. Combinatorial Grouping” (MCG) [14] on three different scales of learned boundary predication maps from **HNN-B**: $\hat{Y}_{side}^{B_2}$, $\hat{Y}_{side}^{B_3}$, and \hat{Y}_{fuse}^B using the original CT image as input (shown with ground truth delineation of pancreas). MCG computes superpixels at each scale and produces a set of merged superpixel-based object proposals. We only visualize the boundary probabilities $p > 10\%$.

3 Results and Discussion

Data: Manual tracings of the pancreas for 82 contrast-enhanced abdominal CT volumes were provided by a publicly available dataset⁵ [6], for the ease of comparison. Our experiments are conducted on random splits of ~ 60 patients for training and ~ 20 for unseen testing in 4-fold cross-validation. Most previous work [1–3] use the leave-one-patient-out (LOO) protocol which is computationally expensive (e.g., ~ 15 h to process one case using a powerful workstation [1]) and may not scale up efficiently towards larger patient populations.

Evaluation: Table 1 shows the improvement from **HNN-I** to using spatial aggregation via **HNN-RF** based on thresholded probability maps (calibrated based on the training data), using DSC and average minimum distance. The average DSC is increased from 76.99 % to 78.01 % statistically significantly ($p < 0.001$, Wilcoxon signed-rank test). In contrast, using dense CRF (DCRF) optimization [10] (with **HNN-I** as the unary term and the pairwise term depending on the CT values) as a means of introducing spatial consistency does not improve upon **HNN-I** noticeably (avg. DSC of 77.14 %, see Table 1). To the best of our knowledge, our result comprises the highest reported average DSC (in testing folds) under the same 4-fold CV evaluation metric [6]. Strict comparison to previous methods (except for [6]) is not directly possible due to different datasets utilized. Our holistic segmentation approach with spatial aggregation advances the current state-of-the-art quantitative performance to an average DSC of 78.01 % in testing. To the best of our knowledge, this is the highest DSC ever reported in the literature. Previous state-of-the-art results range from $\sim 68\%$ to $\sim 73\%$ [3–5]. In particular, DSC drops from 68 % (150 patients) to 58 % (50 patients) under the leave-one-out protocol as reported in [3]. Our methods also perform with

⁵ <http://dx.doi.org/10.7937/K9/TCIA.2016.tNB1kqBU>.

Table 1. 4-fold cross-validation: The DSC [%] and average minimum distance (Dist) [mm] performance of our implementation of [6], optimally achievable super-pixels, **HNN-I**, and **HNN-RF** spatial aggregation, and DCRF (best performance in bold).

DSC	[6]	Opt.	HNN-I	HNN-RF	DCRF	Dist	[6]	Opt.	HNN-I	HNN-RF	DCRF
Mean	71.42	88.08	76.99	78.01	77.14	Mean	1.53	0.15	0.70	0.60	0.69
Std	10.11	2.10	9.45	8.20	10.58	Std	1.60	0.08	0.73	0.55	0.76
Min	23.99	81.24	24.11	34.11	16.10	Min	0.20	0.08	0.17	0.15	0.15
Max	86.29	92.00	87.78	88.65	88.30	Max	10.32	0.81	5.91	4.37	5.71




Fig. 3. Examples of the **RF** pancreas segmentation (green) using the proposed approach in testing with the manual ground truth annotation (red). Case with DSC close to the data set mean and the maximum are shown. The percentage of total cases that lie above a certain DSC with **RF** are shown on the right. 80 % of the cases achieve a minimum DSC of 74.13 %, and 90 % of the cases achieve a DSC of 69.0 % and higher.

the better statistical stability (i.e., comparing 8.2 % versus 18.6 % [1], 15.3 % [2] in the standard deviation of DSCs). The minimal DSC value is 34.11 % for **HNN-RF**, whereas [1–3, 6] all report patient cases with DSC <10 %. A typical patient result achieving a DSC close to the data set mean is shown in Fig. 3. Furthermore, we apply our trained **HNN-I** model on a different CT data set⁶ with 30 patients, and achieve a mean DSC of 62.26 % without any re-training on the new data cases, but if we average the outputs of our 4 **HNN-I** models from cross-validation, we achieve 65.66 % DSC. This demonstrates that **HNN-I** may be highly generalizable in cross-dataset evaluation. Performance on that separated data will likely improve with further fine-tuning.

4 Conclusion

In this paper, we present a holistic deep CNN approach for pancreas segmentation in abdominal CT scans, combining interior and boundary mid-level cues via spatial aggregation. Holistically-Nested Networks (**HNN-I**) alone already achieve good performance on the pixel-labeling task for segmentation. However, we show a significant improvement ($p < 0.001$) by incorporating the organ boundary responses from the **HNN-B** model. **HNN-B** can improve supervised object

⁶ 30 training data sets at <https://www.synapse.org/#!/Synapse:syn3193805/wiki/217789>.

proposals via superpixels and is beneficial to train **HNN-RF** that spatially aggregates information on organ interior, boundary and location. The highest reported DSCs of $78.01\% \pm 8.2\%$ in testing is achieved, at the computational cost of 2~3 min, not hours as in [1–3]. Our deep learning based organ segmentation approach could be generalizable to other segmentation problems with large variations and pathologies, e.g., tumors.

Acknowledgments. This work was supported by the Intramural Research Program of the NIH Clinical Center.

References

1. Wang, Z., Bhatia, K.K., Glocker, B., Marvao, A., Dawes, T., Misawa, K., Mori, K., Rueckert, D.: Geodesic patch-based segmentation. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) MICCAI 2014. LNCS, vol. 8673, pp. 666–673. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10404-1_83](https://doi.org/10.1007/978-3-319-10404-1_83)
2. Chu, C., et al.: Multi-organ segmentation based on spatially-divided probabilistic atlas from 3D abdominal CT images. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013. LNCS, vol. 8150, pp. 165–172. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-40763-5_21](https://doi.org/10.1007/978-3-642-40763-5_21)
3. Wolz, R., Chu, C., Misawa, K., Fujiwara, M., Mori, K., Rueckert, D.: Automated abdominal multi-organ segmentation with subject-specific atlas generation. *TMI* **32**(9), 1723–1730 (2013)
4. Tong, T., Wolz, R., Wang, Z., Gao, Q., Misawa, K., Fujiwara, M., Mori, K., Hajnal, J.V., Rueckert, D.: Discriminative dictionary learning for abdominal multi-organ segmentation. *Med. Image Anal.* **23**(1), 92–104 (2015)
5. 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)
6. Roth, H.R., Lu, L., Farag, A., Shin, H.-C., Liu, J., Turkbey, E.B., Summers, R.M.: DeepOrgan: multi-level deep convolutional networks for automated pancreas segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9349, pp. 556–564. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-24553-9_68](https://doi.org/10.1007/978-3-319-24553-9_68)
7. Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., Huang, C., Torr, P.H.: Conditional random fields as recurrent neural networks. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1529–1537 (2015)
8. 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, Heidelberg (2015). doi:[10.1007/978-3-319-24574-4_28](https://doi.org/10.1007/978-3-319-24574-4_28)
9. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3431–3440 (2015)
10. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint* (2014). [arXiv:1412.7062](https://arxiv.org/abs/1412.7062)

11. Xie, S., Tu, Z.: Holistically-nested edge detection. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1395–1403 (2015)
12. Criminisi, A., Robertson, D., Konukoglu, E., Shotton, J., Pathak, S., White, S., Siddiqui, K.: Regression forests for efficient anatomy detection and localization in computed tomography scans. *Med. Image Anal.* **17**(8), 1293–1303 (2013)
13. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. ICLR 2015, arXiv preprint (2014). [arXiv:1409.1556](#)
14. Pont-Tuset, J., Arbeláez, P., Barron, J., Marques, F., Malik, J.: Multiscale combinatorial grouping for image segmentation and object proposal generation (2015). [arXiv:1503.00848](#)