



# A Multi-scale Pyramid of 3D Fully Convolutional Networks for Abdominal Multi-organ Segmentation

Holger R. Roth<sup>1</sup>(✉), Chen Shen<sup>1</sup>, Hirohisa Oda<sup>1</sup>, Takaaki Sugino<sup>1</sup>,  
Masahiro Oda<sup>1</sup>, Yuichiro Hayashi<sup>1</sup>, Kazunari Misawa<sup>2</sup>,  
and Kensaku Mori<sup>1,3,4</sup>(✉)

<sup>1</sup> Graduate School of Informatics, Nagoya University, Nagoya, Japan  
rothhr@mori.m.is.nagoya-u.ac.jp, kensaku@is.nagoya-u.ac.jp

<sup>2</sup> Aichi Cancer Center, Nagoya, Japan

<sup>3</sup> Information Technology Center, Nagoya University, Nagoya, Japan

<sup>4</sup> Research Center for Medical Bigdata, National Institute of Informatics,  
Tokyo, Japan

**Abstract.** Recent advances in deep learning, like 3D fully convolutional networks (FCNs), have improved the state-of-the-art in dense semantic segmentation of medical images. However, most network architectures require severely downsampling or cropping the images to meet the memory limitations of today's GPU cards while still considering enough context in the images for accurate segmentation. In this work, we propose a novel approach that utilizes auto-context to perform semantic segmentation at higher resolutions in a multi-scale pyramid of stacked 3D FCNs. We train and validate our models on a dataset of manually annotated abdominal organs and vessels from 377 clinical CT images used in gastric surgery, and achieve promising results with close to 90% Dice score on average. For additional evaluation, we perform separate testing on datasets from different sources and achieve competitive results, illustrating the robustness of the model and approach.

## 1 Introduction

Multi-organ segmentation has attracted considerable interest over the years. The recent success of deep learning-based classification and segmentation methods has triggered widespread applications of deep learning-based semantic segmentation in medical imaging [1, 2]. Many methods focused on the segmentation of single organs like the prostate [1], liver [3], or pancreas [4, 5]. Deep learning-based multi-organ segmentation in abdominal CT has also been approached recently in works like [6, 7]. Most of these methods are based on variants of fully convolutional networks (FCNs) [8] that either employ 2D convolutions on orthogonal cross-sections in a slice-by-slice fashion [3–5, 9] or 3D convolutions [1, 2, 7]. A common feature of these segmentation methods is that they are able to extract features useful for image segmentation directly from the training imaging data,

which is crucial for the success of deep learning. This avoids the need for hand-crafting features that are suitable for detection of individual organs.

However, most network architectures require severely downsampling or cropping the images for 3D processing to meet the memory limitations of today’s GPU cards [1, 7] while still considering enough context in the images for accurate segmentation of organs.

In this work, we propose a multi-scale 3D FCN approach that utilizes a scale-space pyramid with auto-context to perform semantic image segmentation at a higher resolution while also considering large contextual information from lower resolution levels. We train our models on a large dataset of manually annotated abdominal organs and vessels from pre-operative clinical computed tomography (CT) images used in gastric surgery and evaluate them on a completely unseen dataset from a different hospital, achieving a promising performance compared to the state-of-the-art.

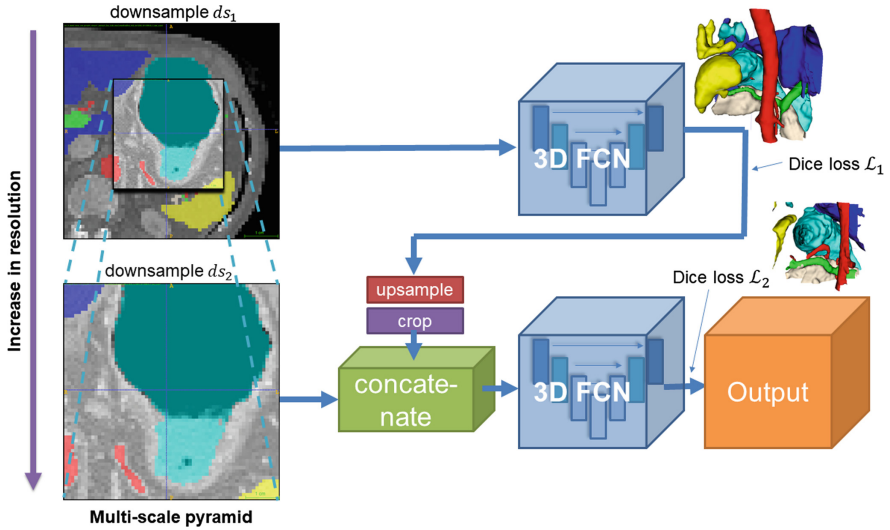
Our approach is shown schematically in Fig. 1. We are influenced by classical scale-space pyramid [10] and auto-context ideas [11] for integrating multi-scale and varying context information into our deep learning-based image segmentation method. Instead of having separate FCN pathways for each scale as explored in other work [12, 13], we utilize the auto-context principle to fuse and integrate the information from different image scales and different amounts of context. This helps the 3D FCN to integrate the information of different image scales and image contexts at the same time. Our model can be trained end-to-end using modern deep learning frameworks. This is in contrast to previous work which utilized auto-context using a separately trained models for brain segmentation [13].

In summary, our contributions are (1) introduction of a multi-scale pyramid of 3D FCNs; (2) improved segmentation of fine structures at higher resolution; (3) end-to-end training of multi-scale pyramid FCNs showing improved performance and good learning properties. We perform a comprehensive evaluation on a large training and validation dataset, plus unseen testing on data from different hospitals and public sources, showing promising generalizability.

## 2 Methods

### 2.1 3D Fully Convolutional Networks

Convolutional neural networks (CNN) have the ability to solve challenging classification tasks in a data-driven manner. Fully convolutional networks (FCNs) are an extension to CNNs that have made it feasible to train models for pixel-wise semantic segmentation in an end-to-end fashion [8]. In FCNs, feature learning is purely driven by the data and segmentation task at hand and the network architecture. Given a training set of images and labels  $\mathbf{S} = \{(X_n, L_n), n = 1, \dots, N\}$ , where  $X_n$  denotes a CT image and  $L_n$  a ground truth label image, the model can train to minimize a loss function  $\mathcal{L}$  in order to optimize the FCN model  $f(I, \Theta)$ , where  $\Theta$  denotes the network parameters, including the convolutional kernel weights for hierarchical feature extraction.



**Fig. 1.** Multi-scale pyramid of 3D fully convolutional networks (FCNs) for multi-organ segmentation. The lower-resolution-level 3D FCN predictions are upsampled, cropped and concatenated with the inputs of a higher resolution 3D FCN. The Dice loss is used for optimization at each level and training is performed end-to-end.

While efficient implementations of 3D convolutions and growing GPU memory have made it possible to deploy FCN on 3D biomedical imaging data [1, 2], image volumes are in practice often cropped and downsampled in order for the network to access enough context to learn an effective semantic segmentation model while still fitting into memory. Our employed network model is inspired by the fully convolutional type 3D U-Net architecture proposed in Çiçek et al. [2].

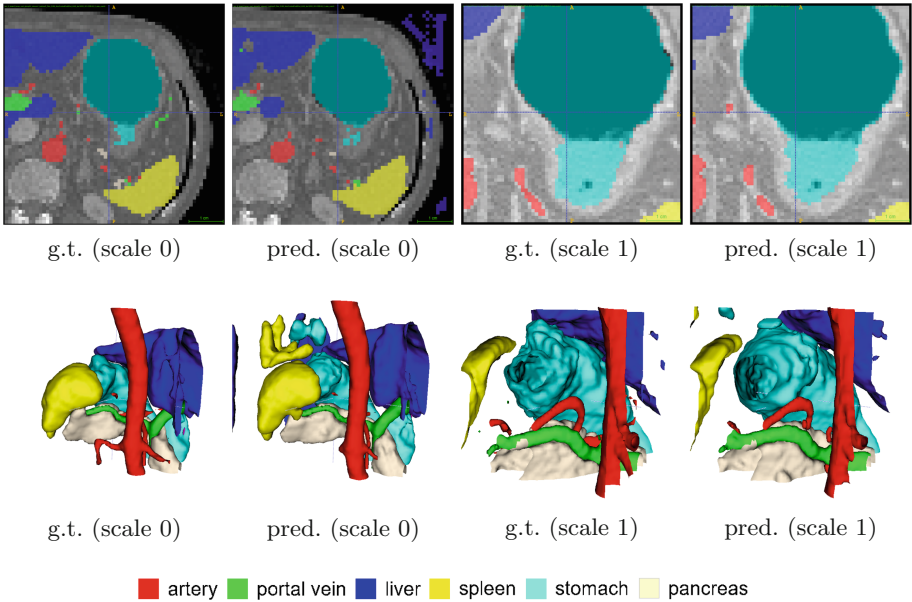
The 3D U-Net architecture is based on U-Net proposed in [14] and consists of analysis and synthesis paths with four resolution levels each. It utilizes deconvolution [8] (also called transposed convolutions) to remap the lower resolution and more abstract feature maps within the network to the denser space of the input images. This operation allows for efficient dense voxel-to-voxel predictions. Each resolution level in the analysis path contains two  $3 \times 3 \times 3$  convolutional layers, each followed by rectified linear units (ReLU) and a  $2 \times 2 \times 2$  max pooling with strides of two in each dimension. In the synthesis path, the convolutional layers are replaced by deconvolutions of  $2 \times 2 \times 2$  with strides of two in each dimension. These are followed by two  $3 \times 3 \times 3$  convolutions, each followed by ReLU activations. Furthermore, 3D U-Net employs shortcut (or skip) connections from layers of equal resolution in the analysis path to provide higher-resolution features to the synthesis path [2]. The last layer contains a  $1 \times 1 \times 1$  convolution that reduces the number of output channels to the number of class labels  $K$ . This architecture has over 19 million learnable parameters and can be trained to minimize the average Dice loss derived from the binary case in [1]:

$$\mathcal{L}(X, \Theta, L) = -\frac{1}{K} \sum_{k=1}^K \left( \frac{2 \sum_i^N p_{i,k} l_{i,k}}{\sum_i^N p_{i,k} + \sum_i^N l_{i,k}} \right). \quad (1)$$

Here,  $p_{i,k} \in [0, \dots, 1]$  represents the continuous values of the *softmax* 3D prediction maps for each class label  $k$  of  $K$  and  $l_{i,k}$  the corresponding ground truth value in  $L$  at each voxel  $i$ .

## 2.2 Multi-scale Auto-Context Pyramid Approach

To effectively process an image at higher resolutions, we propose a method that is inspired by the auto-context algorithm [11]. Our method both captures the context information at lower resolution downsampled images and learns more accurate segmentations from higher resolution images in two levels of a scale-space pyramid  $\mathbf{F} = \{(f_s(X_s, \Theta_s)), s = 1, \dots, S\}$ , with  $S$  being the number of levels  $s$  in our multi-scale pyramid, and  $X_s$  being one of the multi-scale input subvolumes at each level  $s$ .



**Fig. 2.** Axial CT images and 3D surface rendering with ground truth (g.t.) and predictions overlaid. We show the two scales used in our experiments. Each scale’s input is of size  $64 \times 64 \times 64$  in this setting.

In the first level, the 3D FCN is trained on images of the lowest resolution in order to capture the largest amount of context, downsampled with a factor of  $ds_1 = 2S$  and optimized using the Dice loss  $\mathcal{L}_1$ . This can be thought of as a form

of deep supervision [15]. In the next level, we use the predicted segmentation maps as a second input channel to the 3D FCN while learning from the images at a higher resolution, downsampled by a factor of  $ds_2 = ds_1/2$ , and optimized using Dice loss  $\mathcal{L}_2$ . For input to this second level of the pyramid, the previous level prediction maps are upsampled by a factor of 2 and cropped in order to spatially align with the higher resolution levels. These predictions can then be fed together with the appropriately cropped image data as a second channel. This approach can be learned end-to-end using modern multi-GPU devices and deep learning frameworks with the total loss being  $\mathcal{L}_{\text{total}} = \sum_{s=1}^L \mathcal{L}_s(X_s, \Theta_s, L_s)$ . This idea is shown schematically in Fig. 1. The resulting segmentation masks for the two-level case are shown in Fig. 2. It can be observed that the second-level auto-context network markedly outperforms the first-level predictions and is able to segment structures with improved detail, especially at the vessels.

### 2.3 Implementation and Training

We implement our approach in Keras<sup>1</sup> using the TensorFlow<sup>2</sup> backend. The Dice loss [3] is used for optimization with Adam and automatic differentiation for gradient computations. Batch normalization layers are inserted throughout the network, using a mini-batch size of three, sampled from different CT volumes of the training set. We use randomly extracted subvolumes of fixed size during training, such that at least one foreground voxel is at the center of each subvolume. On-the-fly data augmentation is used via random translations, rotations and elastic deformations similar to [2].

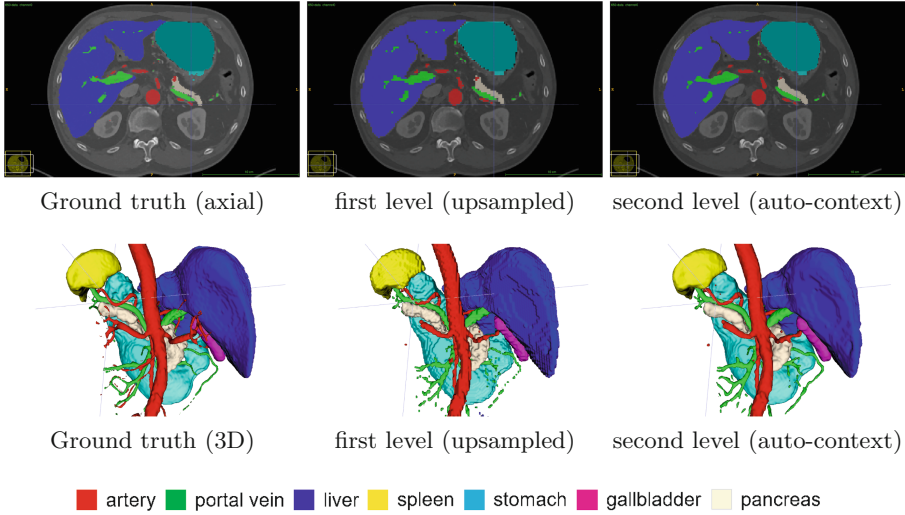
## 3 Experiments and Results

In our implementation, a constant input and output size of  $64 \times 64 \times 64$  randomly cropped subvolumes is used for training in each level. For inference, we employ network reshaping [8] to more efficiently process the testing image with a larger input size while building up the full image in a tiling approach [2]. The resulting segmentation masks for both levels are shown in Fig. 3. It can be observed that the second-level auto-context network markedly outperforms the first-level predictions and is able to segment structures with improved detail. All experiments were performed using a DeepLearning BOX (GDEP Advance) with four NVIDIA Quadro P6000s with 24 GB memory each. Training of 20,000 iterations using this unoptimized implementation took several days, while inference on a full CT scan takes just a few minutes on one GPU card.

*Data:* Our data set includes 377 contrast-enhanced clinical CT images of the abdomen in the portal-venous phase used for pre-operative planning in gastric surgery. Each CT volume consists of 460–1,177 slices of  $512 \times 512$  pixels. Voxel dimensions are [0.59–0.98, 0.59–0.98, 0.5–1.0] mm. With  $S = 2$ , we downsample

<sup>1</sup> <https://keras.io/>.

<sup>2</sup> <https://www.tensorflow.org/>.



**Fig. 3.** Axial CT images and 3D surface rendering of predictions from two multi-scale levels in comparison with ground truth annotations. In particular, the vessels are segmented more completely and in greater detail in the second level, which utilizes auto-context information in its prediction.

each volume by a factor of  $ds_1 = 4$  in the first level and a factor of  $ds_2 = 2$  in the second level. A random 90/10% split of 340/37 patients is used for training and testing the network. We achieve Dice similarity scores for each organ labeled in the testing cases as summarized in Table 1. We list the performance for the first level and second level models when utilizing auto-context trained separately or end-to-end, and compare to using no auto-context in the second level. This shows the impact of using or not using the lower resolution auto-context channel at the higher resolution input while training from the same input resolution from scratch. In our case, each  $L_n$  contains  $K = 8$  labels consisting of the manual annotations of seven anatomical structures (artery, portal vein, liver, spleen, stomach, gallbladder, pancreas), plus background.

Table 2 compares our results to recent literature and also displays the result using an unseen testing dataset from a different hospital consisting of 129 cases from a distinct research study. Furthermore, we test our model on a public data set of 20 contrast-enhanced CT scans.<sup>3</sup>

<sup>3</sup> We utilize the 20 training cases of the VISCERAL data set (<http://www.visceral.eu/benchmarks/anatomy3-open>) as our test set.

**Table 1.** Comparison of different levels of our model. End-to-end training gives a statistically significant improvement ( $p < 0.001$ ).

| Dice (%)                                   | Artery      | Vein        | Liver       | Spleen      | Stomach     | Gall.       | Pancreas    | Avg.        |
|--|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| <i>Level 1: initial (low res)</i>          |             |             |             |             |             |             |             |             |
| Avg  | 75.4        | 64.0        | 95.4        | 94.0        | 93.7        | 80.2        | 79.8        | 83.2        |
| Std  | 3.9         | 5.4         | 1.0         | 0.8         | 7.6         | 15.5        | 8.5         | 06.1        |
| Min  | 67.4        | 41.3        | 91.5        | 92.6        | 48.4        | 27.3        | 49.7        | 59.7        |
| Max  | 82.3        | 70.9        | 96.4        | 95.8        | 96.5        | 93.5        | 90.6        | 89.4        |
| <i>Level 2: auto-context</i>               |             |             |             |             |             |             |             |             |
| Avg  | 82.5        | 76.8        | 96.7        | 96.6        | 95.9        | <b>84.4</b> | 83.4        | 88.1        |
| Std  | 4.1         | 6.4         | 1.0         | 0.7         | 8.0         | 14.0        | 8.4         | 6.1         |
| Min  | 73.3        | 46.3        | 92.9        | 94.4        | 48.1        | 28.0        | 53.9        | 62.4        |
| Max  | 90.0        | 83.5        | 97.9        | 98.0        | 98.7        | 96.0        | 93.4        | 93.9        |
| <i>End-to-end: auto-context (high-res)</i> |             |             |             |             |             |             |             |             |
| Avg  | <b>83.0</b> | <b>79.4</b> | <b>96.9</b> | <b>97.2</b> | <b>96.2</b> | 83.6        | <b>86.7</b> | <b>89.0</b> |
| Std  | 4.4         | 6.7         | 1.0         | 1.0         | 5.9         | 17.1        | 7.4         | 6.2         |
| Min  | 73.2        | 50.2        | 93.5        | 94.9        | 61.4        | 29.7        | 60.0        | 66.1        |
| Max  | 91.0        | 87.7        | 98.3        | 98.7        | 98.7        | 96.4        | 95.2        | 95.1        |
| <i>Level 2: no auto-context (high-res)</i> |             |             |             |             |             |             |             |             |
| Avg  | 69.9        | 72.8        | 86.7        | 90.9        | 3.8         | 73.4        | 77.0        | 67.8        |
| Std  | 6.2         | 7.0         | 6.4         | 5.3         | 1.3         | 22.5        | 10.8        | 8.5         |
| Min  | 59.5        | 47.1        | 69.9        | 75.7        | 0.7         | 7.8         | 36.1        | 42.4        |
| Max  | 82.1        | 82.9        | 95.7        | 97.0        | 7.4         | 95.9        | 90.9        | 78.8        |

\*Best average performance is shown in **bold**.

**Table 2.** We compare our model trained in an end-to-end fashion to recent work on multi-organ segmentation. [9] is using a 2D FCN approach with a majority voting scheme, while [7] employs 3D FCN architectures. Furthermore, we list our performance on an unseen testing dataset from a different hospital and on the public Visceral dataset without any re-training and compare it to the current challenge leaderboard (LB) best performance for each organ. Note that this table is incomprehensive and direct comparison to the literature is always difficult due to the different datasets and evaluation schemes involved.

| Dice (%)                     | Train/Test | Artery | Vein | Liver | Spleen | Stomach | Gall. | Pancreas | Avg. |
|------------------------------|------------|--------|------|-------|--------|---------|-------|----------|------|
| Ours (end-to-end)            | 340/37     | 83.0   | 79.4 | 96.9  | 97.2   | 96.2    | 83.6  | 86.7     | 89.0 |
| Unseen test                  | none/129   | -      | -    | 95.3  | 93.6   | -       | 80.8  | 75.7     | 86.3 |
| Gibson et al. [7]            | 72 (8-CV)  | -      | -    | 92    | -      | 83      | -     | 66       | 80.3 |
| Zhou et al. [9] <sup>a</sup> | 228/12     | 73.8   | 22.4 | 93.7  | 86.8   | 62.4    | 59.6  | 56.1     | 65.0 |
| Hu et al. [6]                | 140 (CV)   | -      | -    | 96.0  | 94.2   | -       | -     | -        | 95.1 |
| Visceral (LB)                | 20/10      | -      | -    | 95.0  | 91.1   | -       | 70.6  | 58.5     | 78.8 |
| Visceral (ours) <sup>b</sup> | none/20    | -      | -    | 94.0  | 87.2   | -       | 68.2  | 61.9     | 77.8 |

<sup>a</sup>Dice score estimated from Intersection over Union (Jaccard index).

<sup>b</sup>At the time of writing, the testing evaluation servers of the challenge were not available anymore for submitting results.

## 4 Discussion and Conclusion

The multi-scale auto-context approach presented in this paper provides a simple yet effective method for employing 3D FCNs in medical-imaging settings. No post-processing was applied to any of the network outputs. The improved performance in our approach is effective for all organs tested (apart from the gallbladder, where the differences are not significant). Note that we used different datasets (from different hospitals and scanners) for separate testing. This experiment illustrates our method's generalizability and robustness to differences in image quality and populations. Running the algorithms at a quarter to half of the original resolution improved performance and efficiency in this application. While this method could be extended to using a multi-scale pyramid with the original resolution as the final level, we found that the added computational burden did not add significantly to the segmentation performance. The main improvement comes from utilizing a very coarse image (downsampled by a factor of four) in an effective manner. In this work, we utilized a 3D U-Net-like model for each level of the image pyramid. However, the proposed auto-context approach should in principle also work well for other 3D CNN/FCN architectures and 2D and 3D image modalities.

In conclusion, we showed that an auto-context approach can result in improved semantic segmentation results for 3D FCNs based on the 3D U-Net architecture. While the low-resolution part of the model is able to benefit from a larger context in the input image, the higher resolution auto-context part of the model can segment the image with greater detail, resulting in better overall dense predictions. Training both levels end-to-end resulted in improved performance.

**Acknowledgments.** This work was supported by MEXT KAKENHI (26108006, 17H00867, 17K20099) and the JPSP International Bilateral Collaboration Grant.

## References

1. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: fully convolutional neural networks for volumetric medical image segmentation. In: 3D Vision (3DV), pp. 565–571. IEEE (2016)
2. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016 Part II. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46723-8\\_49](https://doi.org/10.1007/978-3-319-46723-8_49)
3. Christ, P.F., et al.: Automatic Liver and lesion segmentation in CT using cascaded fully convolutional neural networks and 3D conditional random fields. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016 Part II. LNCS, vol. 9901, pp. 415–423. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46723-8\\_48](https://doi.org/10.1007/978-3-319-46723-8_48)



4. Roth, H.R., Lu, L., Farag, A., Sohn, A., Summers, R.M.: Spatial aggregation of holistically-nested networks for automated pancreas segmentation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016 Part II. LNCS, vol. 9901, pp. 451–459. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46723-8\\_52](https://doi.org/10.1007/978-3-319-46723-8_52)
5. Zhou, Y., Xie, L., Shen, W., Wang, Y., Fishman, E.K., Yuille, A.L.: A fixed-point model for pancreas segmentation in abdominal CT scans. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017 Part I. LNCS, vol. 10433, pp. 693–701. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-66182-7\\_79](https://doi.org/10.1007/978-3-319-66182-7_79)
6. Hu, P., Wu, F., Peng, J., Bao, Y., Chen, F., Kong, D.: Automatic abdominal multi-organ segmentation using deep convolutional neural network and time-implicit level sets. *Int. J. Comput. Assist. Radiol. Surg.* **12**(3), 399–411 (2017)
7. Gibson, E., et al.: Towards image-guided pancreas and biliary endoscopy: automatic multi-organ segmentation on abdominal CT with dense dilated networks. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017 Part I. LNCS, vol. 10433, pp. 728–736. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-66182-7\\_83](https://doi.org/10.1007/978-3-319-66182-7_83)
8. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *IEEE CVPR*, pp. 3431–3440 (2015)
9. Zhou, X., Takayama, R., Wang, S., Hara, T., Fujita, H.: Deep learning of the sectional appearances of 3D CT images for anatomical structure segmentation based on an FCN voting method. *Med. Phys.* **44**, 5221–5233 (2017)
10. Adelson, E.H., Anderson, C.H., Bergen, J.R., Burt, P.J., Ogden, J.M.: Pyramid methods in image processing. *RCA Eng.* **29**(6), 33–41 (1984)
11. Tu, Z., Bai, X.: Auto-context and its application to high-level vision tasks and 3D brain image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(10), 1744–1757 (2010)
12. Chen, L.C., Yang, Y., Wang, J., Xu, W., Yuille, A.L.: Attention to scale: scale-aware semantic image segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3640–3649 (2016)
13. Salehi, S.S.M., Erdogmus, D., Gholipour, A.: Auto-context convolutional neural network (auto-net) for brain extraction in magnetic resonance imaging. *IEEE Trans. Med. Imaging* **36**(11), 2319–2330 (2017)
14. 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 Part III. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)
15. Lee, C.Y., Xie, S., Gallagher, P., Zhang, Z., Tu, Z.: Deeply-supervised nets. In: *Artificial Intelligence and Statistics*, pp. 562–570 (2015)