




Advancing Pancreas Segmentation in Multi-protocol MRI Volumes Using Hausdorff-Sine Loss Function

Hykoush Asaturyan¹ , E. Louise Thomas², Julie Fitzpatrick²,
Jimmy D. Bell², and Barbara Villarini¹

¹ School of Computer Science and Engineering, University of Westminster,
London, UK

h.asaturyan@my.westminster.ac.uk, b.villarini@westminster.ac.uk

² School of Life Sciences, University of Westminster, London, UK
{[L.Thomas3](mailto:L.Thomas3@westminster.ac.uk),[J.Fitzpatrick](mailto:J.Fitzpatrick@westminster.ac.uk),[J.Bell](mailto:J.Bell@westminster.ac.uk)}@westminster.ac.uk

Abstract. Computing pancreatic morphology in 3D radiological scans could provide significant insight about a medical condition. However, segmenting the pancreas in magnetic resonance imaging (MRI) remains challenging due to high inter-patient variability. Also, the resolution and speed of MRI scanning present artefacts that blur the pancreas boundaries between overlapping anatomical structures. This paper proposes a dual-stage automatic segmentation method: (1) a deep neural network is trained to address the problem of vague organ boundaries in high class-imbalanced data. This network integrates a novel loss function to rigorously optimise boundary delineation using the modified Hausdorff metric and a sinusoidal component; (2) Given a test MRI volume, the output of the trained network predicts a sequence of targeted 2D pancreas classes that are reconstructed as a volumetric binary mask. An energy-minimisation approach fuses a learned digital contrast model to suppress the intensities of non-pancreas classes, which, combined with the binary volume performs a refined segmentation in 3D while revealing dense boundary detail. Experiments are performed on two diverse MRI datasets containing 180 and 120 scans, in which the proposed approach achieves a mean Dice score of $84.1 \pm 4.6\%$ and $85.7 \pm 2.3\%$, respectively. This approach is statistically stable and outperforms state-of-the-art methods on MRI.

Keywords: Automatic pancreas segmentation ·
Energy-minimisation · MRI · Hausdorff loss function

1 Introduction

Segmenting the pancreas in 3D radiological scans (e.g. an MRI volume) could provide significant insight into the severity or progression of type 2 diabetes [1] and ductal adenocarcinoma [2]. However, pancreas segmentation presents

several challenges due to high structural and inter-patient variability in size and location. The greyscale intensity of the pancreas can be very similar to neighbouring tissue, and the boundary contrast can vary depending on the level of surrounding visceral fat. Differing from computer tomography (CT), the low resolution and slower imaging speed of MRI presents edge-based artefacts that blur the imaging boundaries between the pancreas and surrounding organs [3]. In existing research literature, pancreas segmentation tasks have been driven by two major methodologies: multi-atlas based [4,5] coupled with statistical shape modeling [6], and in more recent years, convolutional neural networks (CNNs) or deep learning [3,7,8]. While CNNs have achieved higher quantitative accuracy scores in 2D medical image segmentation, such methods can exhibit discontinuity in predicting pancreatic regions between consecutive slices for an input volume.

This paper presents a novel approach for automatic pancreas segmentation in MRI. As illustrated in Fig. 1, the proposed method consists of two successive stages. First, a CNN specialising in blurred boundary detection is trained to predict targeted pixel-wise pancreas tissue. This deep learning stage firstly identifies the main pancreas region (ROI) in a dataset of MRI volumes [8] by training a random forest on extracted texture and probability-wise features on image patches of 25×25 pixels. Next, inspired by the encoder-decoder architecture of SegNet [9] a new model termed Hausdorff Sine SegNet (HSSN) is developed using the ROI data. A novel loss function incorporates the modified Hausdorff distance metric and a sinusoidal component to capture local boundary information, enforce edge detection and thus raise the true pancreas prediction rate on a 2D (slice-by-slice) basis. The testing stage consists of two phases. First, the output of the trained HSSN for a given test MRI volume encodes spatial information to classify every pixel in each slice, thus forming a volumetric binary mask (VBM). The second phase generates dense contouring by further tackling the low dissimilarity between organ boundaries: a digital contrast enhancement model is utilised to improve the greyscale variation between surrounding background classes within close proximity to the pancreas. A 3D energy-minimising algorithm performs refined segmentation on the enhanced pancreas that is fused with the VBM, producing greater consistency in spatial smoothness and prediction among successive slices.

The proposed method, which is evaluated on two MRI datasets with varying noise, outperforms the state-of-the-art approaches [8,10–12], and moreover, surpasses the performance of readily employed deep learning-based loss functions. Although this approach has been tested on pancreas segmentation, the methodology is reproducible, scalable and generalisable to other organ segmentation tasks. The implementation is available at <https://github.com/med-seg/p>.

2 Methodology

2.1 Training the HSSN

The proposed HSSN model has an encoder-decoder topology, as illustrated in Fig. 2. The decoder network uses max-pooling indices to upsample low-resolution

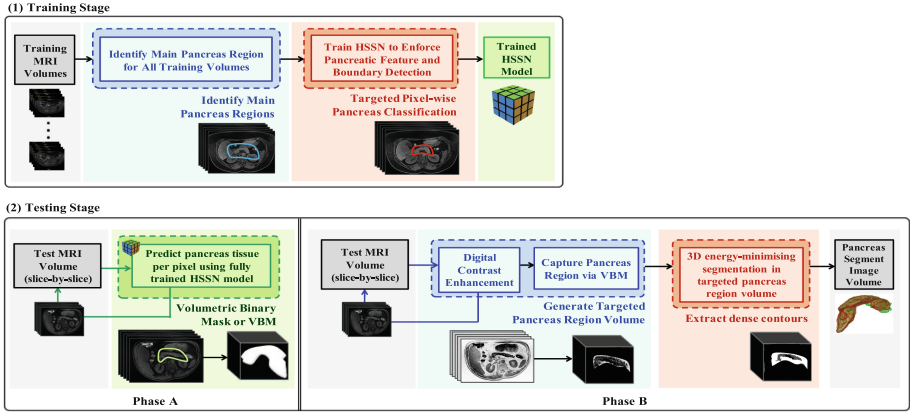


Fig. 1. Overview of proposed approach. (1) develop the HSSN deep learning model using training MRI; and (2) apply the test MRI to generate segmented pancreas volume.

feature maps, consequently retaining high-frequency details to improve pancreatic boundary delineation, and reducing the total number of trainable parameters in the decoders. Unlike other models that have been fine-tuned from pre-trained CNNs using a large number of natural images [3, 13], this network is trained from scratch using exclusively pancreas datasets. Since this organ accounts for $\sim 1\%$ in a scan, there is a need to weight the loss differently based on the true class: Median frequency balancing [14] is utilised, in which the weight assigned to a class in the loss function is the ratio of the median of class frequencies computed on the entire training set divided by the class frequency. The HSSN also employs data augmentation of random reflections and translations to reduce overfitting [15], and further address problems caused by high shape variability.

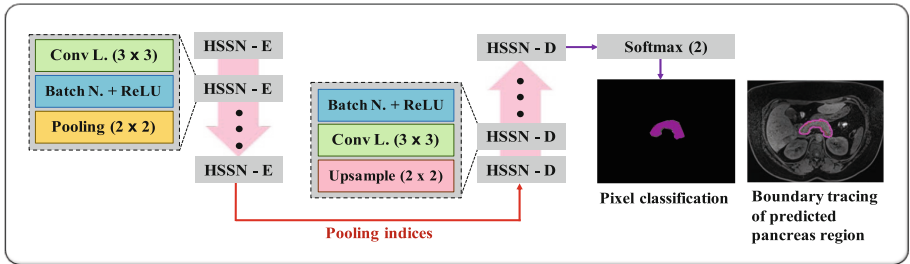


Fig. 2. Overview of HSSN model. An encoder stage (5 blocks of HSSN-E) downsamples the MRI input through convolution, batch normalisation and ReLU. A decoder stage (5 blocks of HSSN-D) upsamples its input using the transferred pooling indices from its corresponding encoder to generate sparse feature maps. From here, convolution is performed with a trainable filter of weights to density the feature map. Resulting decoder output feature maps are fed to soft-max classifier for 2-channel pixel-wise classification of the input image as “pancreas” or “non-pancreas”.

Integrated Hausdorff-Sine Loss Function: A novel loss function is proposed for training the segmentation neural network. The optimisation of the modified Hausdorff distance and a sinusoidal functionality serves to reduce the boundary matching error and “enhance” a resulting pixel-wise pancreas prediction. Let T_H and Y_H represent the ground-truth and network boundary predictions respectively, where $T_H, Y_H \subset \mathbb{R}^n$ such that $|T_H|, |Y_H| < \infty$. Furthermore, t_j and $y_j \in \{0, 1\}$ are indexed pixel values in T_H and Y_H respectively, and can be viewed as boundary points. The Euclidean distance between a point t_j and set of points, Y_H is $s(t_j, Y_H) = \min_{y_j \in Y_H} \|t_j - y_j\|$, and the distance between a point y_j and set of points, T_H is $s(y_j, T_H) = \min_{t_j \in T_H} \|y_j - t_j\|$. If $\varepsilon_Y = \frac{1}{|Y_H|} \sum_{y_j \in Y_H} s(t_j, Y_H)$ and $\varepsilon_T = \frac{1}{|T_H|} \sum_{t_j \in T_H} s(y_j, T_H)$, the modified Hausdorff distance loss, L_{mhd} is:

$$L_{mhd} = \max \{ \varepsilon_Y, \varepsilon_T \} \quad (1)$$

Thus, computing the gradient yields:

$$\frac{\partial L_{mhd}}{\partial Y_H} = \begin{cases} \frac{\partial}{\partial Y_H} (\varepsilon_Y) & \text{if } \varepsilon_Y > \varepsilon_T \\ \frac{\partial}{\partial Y_H} (\varepsilon_T) & \text{if } \varepsilon_T < \varepsilon_Y \\ \text{undefined} & \text{if } \varepsilon_Y = \varepsilon_T \end{cases} \quad (2)$$

An additional sinusoidal component increases non-linearity during network training and, empirically evaluated, raises the true positive predictions. If T and Y represent the ground-truth and network predictions, the loss L_{sine} is defined:

$$L_{sine} = -\frac{1}{|Y|} \sum_{i=1}^{nC} \sin(T_i) \log_2(Y_i) \quad (3)$$

where $nC = 2$ is the number of classes (e.g., Y_1 refers to “pancreas” and Y_2 refers to “non-pancreas”). From here, computing the gradient yields:

$$\frac{\partial L_{sine}}{\partial Y_i} = -\frac{1}{|Y|} \frac{\sin(T_i)}{Y_i \log_{10}(2)} \quad (4)$$

The model is updated via the combined gradients of L_{sine} and L_{mhd} .

2.2 Testing Stage

(A) Targeted Pancreas Binary Mask: The trained HSSN model performs pixel-wise prediction on each slice in a test MRI volume to generate a resulting volumetric binary mask (VBM). Columns (a) and (b) in Fig. 3 displays three sample input slices in three different image volumes, and the corresponding positive pancreas region (white mask) as predicted by the HSSN model. The red contouring in each image in column (b) is the ground-truth.




Fig. 3. Visualising proposed approach.

(B) Achieve Dense Contouring: The test MRI volume undergoes non-local means for denoising, after which a learned intensity model incorporates a sigmoid function to exhaustively differentiate pancreatic tissue against background classes. Every s_i -th slice transforms to $C(s_i) = 1/(1 + \exp[g(c-s_i)])$, where g controls the actual contrast, and c is the cut-off value representing the (normalised) greyscale value about which g is changed [12, 16]. The VBM is applied to the enhanced image volume and processed through a 3D unsupervised energy-minimisation method via continuous max-flow [17], revealing detailed contouring as highlighted in Fig. 3, column (c). The accurate HSSN predictions reduce the level of non-pancreatic tissue carried into the max-flow segmentation stage, as shown in Fig. 3, column (d), eliminating the need for post-processing.

3 Experimental Results and Analysis

3.1 Datasets and Evaluation

Two expert-led annotated pancreas datasets are utilised. MRI-A and MRI-B contain 180 and 120 abdominal MRI scans (T2-weighted, fat suppressed), which have been obtained using a Philips Intera 1.5T and a Siemens Trio 3T scanner, respectively. Every MRI-A scan has 50 slices, each of size 384×384 with spacing 2 mm, and 0.9766 mm pixel interval in the axial and sagittal direction. Every MRI-B scan has 80 slices, each of size 320×260 with 1.6 mm spacing, and 1.1875 mm pixel interval in the axial and sagittal direction. The proposed approach is evaluated using the Dice Similarity Coefficient (DSC), precision (PC), recall (RC) and the Hausdorff distance (HSD) representing the maximum boundary deviation between the segmentation and ground-truth.

3.2 Network Implementation

The training and testing data are randomly split into 160 and 20 (MRI-A) and 100 and 20 (MRI-B). The HSSN model employs stochastic gradient descent with parameters momentum (0.9), initial learning rate (0.001), maximum epochs (300) and mini-batch size (10). The mean time for model training is ~ 11 h and the testing phase is ~ 7.5 min per MRI volume using an i7-59-30k-CPU at 3.50 Ghz. Future work can potentially reduce these run-times by a factor of 10 via multiple GeForce Titan X GPUs.

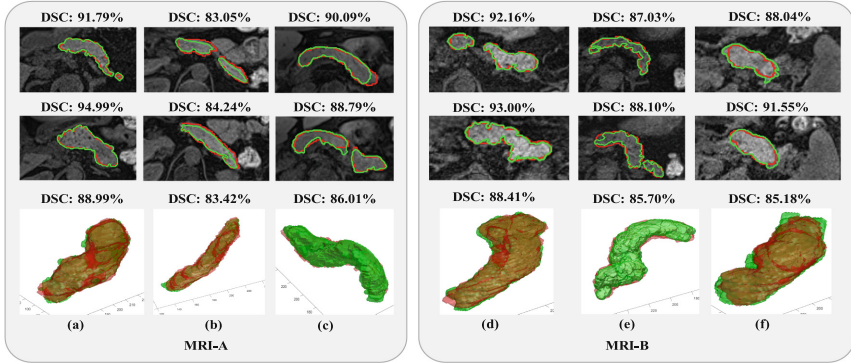


Fig. 4. Segmentation results in six different MRI scans (volumes). Every column corresponds to a single MRI volume. From left, first row displays sample MRI axial slices with segmentation outcome (green) against ground-truth (red), and computed DSC. Second row displays 3D reconstruction of entire pancreas with computed DSC. (Color figure online)




Fig. 5. Box plots of DSC and JI.

3.3 Analysis of Proposed Approach

Figure 4 displays the final segmentation results in six MRI scans, equally split between MRI-A and MRI-B. Columns (a, b, c) are part of MRI-A, yet there is high variation between intensity and contrast in the original axial MRI slices. Columns (d, e, f) corresponds to exemplars from MRI-B. As reflected in Fig. 5, 85% of MRI-A compared to 95% in MRI-B segmentations score above 80% in DSC, demonstrating the robust performance of the approach with respect to poor image quality, intensity distribution and spatial dimensions.

Hausdorff-Sine Loss: Figure 6 compares the segmentation results (in DSC) using Hausdorff-Sine and the loss functions, Hausdorff, Cross-entropy, Dice [18] and Jaccard [19] in the probability range $[0.05, 0.95]$. The cross-entropy penalises true positive predictions, forcing the “optimum” probability to approximately 0.5. Although the Dice loss minimises the class distribution distance, squaring the weights in the backpropagation stage causes instability and a higher rate of false negative predictions. Similarly, the Jaccard loss suffers from low true positive predictions. Empirically tested, the Hausdorff loss minimises the maximum deviation between a prediction and desired outcome; however, the addition of a sinusoidal component increases non-linearity during training, and thus Hausdorff-Sine achieves improved true positive predictions across differing




Fig. 6. DSC across threshold ranges [0.05, 0.95] via multiple loss functions.




Fig. 7. Averaged ROC curves via multiple loss functions.

thresholds while delivering strong discrimination of true negatives. The ROC curves in Fig. 7 highlight the inferior performance of other loss functions in the extremely unbalanced segmentation, whereas Hausdorff-Sine generally improves the true positive accuracy results.

Phase B of Testing Stage: Integrating the second phase (B) produces contextual boundary information that is essential for accurate segmentation in biomedical imaging. Figure 3, column (b) and column (e) visibly highlights the differences in segmentation boundary delineation against the ground-truth before and after this phase. Thus, the mean HSD metric confirms less deviation from the ground-truth (see Tables 1 and 2) by approximately 1 mm, and furthermore, the mean DSC raises by approximately 4% in both MRI-A and MRI-B.

Table 1. Deep learning model performance using state-of-the-art loss functions versus the integrated novel Hausdorff and Hausdorff-Sine loss. Datasets MRI-A and MRI-B are evaluated in 9-fold and 6-fold cross-validation (FCV), respectively. DSC, PC, RC and HSD are presented as mean \pm standard deviation.

MRI-A: Train/Test (160/20) 9-FCV					MRI-B: Train/Test (100/20) 6-FCV				
Loss	DSC(%)	PC(%)	RC(%)	HSD(mm)	Loss	DSC(%)	PC(%)	RC(%)	HSD(mm)
CE	77.9 \pm 3.6	88.4 \pm 6.18	95.6 \pm 2.26	12.4 \pm 5.5	CE	79.9 \pm 4.33	92.6 \pm 6.89	96.3 \pm 2.76	10.5 \pm 3.34
Dice	63.5 \pm 9.1	63.8 \pm 20.0	86.8 \pm 10.5	16.8 \pm 5.3	Dice	67.1 \pm 12.8	77.2 \pm 15.1	85.2 \pm 16.8	21.4 \pm 12.3
Jac	63.2 \pm 9.6	62.5 \pm 19.8	87.1 \pm 10.0	17.0 \pm 5.4	Jac	68.6 \pm 6.96	68.3 \pm 16.9	88.5 \pm 8.34	17.9 \pm 7.58
Haus	78.4 \pm 6.1	89.5 \pm 9.11	96.2 \pm 4.06	12.7 \pm 4.9	Haus	81.0 \pm 4.25	94.8 \pm 3.84	98.3 \pm 2.28	10.2 \pm 4.17
Haus-Sin	79.7\pm4.0	93.2\pm7.46	97.2\pm2.67	11.2\pm3.6	Haus-Sin	82.1\pm2.99	97.7\pm2.50	99.1\pm0.78	10.0\pm6.63

Comparison with the State-of-the-Art: Table 2 highlights the proposed approach outperforming state-of-the-art methods [8, 10–12] in terms of accuracy and statistical stability despite employing non-organ optimised protocol data.

Table 2. DSC, PC, RC and HSD as mean \pm standard deviation [lowest, highest] for automatic segmentation methods. Datasets MRI-A and MRI-B are evaluated in 9-fold and 6-fold cross-validation (FCV), respectively.

MRI-A: Train/Test (160/20) 9-FCV										
Method	DSC(%)			PC(%)			RC(%)			HSD (mm)
U-Net [10]	66.8 \pm 8.8	[42.3, 77.3]		71.3 \pm 4.4	[62.9, 80.5]		85.1 \pm 4.8	[76.9, 88.16]		16.9 \pm 5.8 [8.22, 24.1]
Cascaded-CNN [8]	52.7 \pm 6.9	[34.4, 60.7]		64.0 \pm 4.1	[50.4, 68.0]		75.2 \pm 4.6	[68.1, 78.25]		21.5 \pm 9.3 [15.7, 38.6]
Dense V-Net [11]	73.6 \pm 6.1	[49.6, 78.8]		86.1 \pm 3.3	[78.5, 88.5]		94.6 \pm 3.4	[82.8, 96.37]		14.4 \pm 7.2 [6.63, 20.5]
Geo-Desc [12]	78.2 \pm 5.8	[67.1, 86.3]		85.3 \pm 9.7	[70.8, 98.9]		93.9 \pm 9.5	[52.5, 99.13]		13.8 \pm 4.4 [6.11, 18.4]
Proposed	84.1\pm4.6	[72.1, 89.6]		95.5\pm6.3	[71.7, 99.7]		97.6\pm3.0	[89.9, 100.0]		10.6\pm3.7 [6.184, 18.4]
MRI-B: Train/Test (100/20) 6-FCV										
Method	DSC(%)			PC(%)			RC(%)			HSD (mm)
U-Net [10]	72.8 \pm 6.0	[58.9, 80.8]		83.8 \pm 3.1	[74.2, 87.46]		94.6 \pm 3.5	[82.8, 95.72]		14.0 \pm 8.1 [6.82, 21.7]
Cascaded-CNN [8]	54.8 \pm 5.1	[44.4, 65.7]		64.3 \pm 3.5	[59.5, 67.91]		76.2 \pm 3.7	[69.9, 79.64]		22.3 \pm 8.6 [16.0, 37.5]
Dense V-Net [11]	74.0 \pm 5.3	[65.1, 80.3]		85.4 \pm 3.1	[78.5, 89.74]		93.0 \pm 3.8	[84.9, 96.35]		16.7 \pm 7.0 [8.46, 19.8]
Geo-Desc [12]	81.2 \pm 5.0	[72.6, 85.8]		84.7 \pm 5.8	[73.1, 93.64]		84.6 \pm 8.2	[69.2, 97.28]		14.7 \pm 4.1 [8.13, 17.6]
Proposed	85.7\pm2.3	[79.9, 90.3]		96.1\pm3.6	[86.7, 100.0]		99.3\pm0.7	[99.9, 100.0]		9.08\pm2.0 [4.87, 14.8]

4 Conclusion

This paper presents a novel approach for automatic pancreas segmentation in MRI volumes generated from different scanner protocols. Combined with the proposed Hausdorff-Sine loss, an encoder-decoder network reinforces pancreatic boundary detection in MRI slices, outperforming the rate of true positive predictions compared to multiple loss functions. In the later stage, a 3D hybrid energy-minimisation algorithm addresses the intensity consistency problem that is often the case when segmenting image volumes on a 2D basis. The proposed approach generates quantitative accuracy results that surpass reported state-of-the-art methods, and moreover, preserve detailed contouring.

References

1. Macauley, M., Percival, K., Thelwall, P.E., Hollingsworth, K.G., Taylor, R.: Altered volume, morphology and composition of the pancreas in type 2 diabetes. *PLoS ONE* **10**(5), 1–14 (2015)
2. Omeri, A., et al.: Contour variations of the body and tail of the pancreas: evaluation with MDCT. *J. Radiol.* **35**(6), 310–318 (2017)
3. Cai, J., Lu, L., Zhang, Z., Xing, F., Yang, L., Yin, Q.: Pancreas segmentation in MRI using graph-based decision fusion on convolutional neural networks. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) *MICCAI 2016*. LNCS, vol. 9901, pp. 442–450. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_51
4. Okada, T., et al.: Abdominal multi-organ segmentation of CT images based on hierarchical spatial modeling of organ interrelations. In: Yoshida, H., Sakas, G., Linguraru, M.G. (eds.) *ABD-MICCAI 2011*. LNCS, vol. 7029, pp. 173–180. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28557-8_22
5. Shimizu, A., Kimoto, T., Kobatake, H., Nawano, S., Shinozaki, K.: Automated pancreas segmentation from three-dimensional contrast-enhanced computed tomography. *Int. J. Comput. Assist. Radiol. Surg.* **5**(1), 85–98 (2010)
6. Okada, T., Linguraru, M.G., Hori, M., Summers, R., Tomiyama, N., Sato, Y.: Abdominal multi-organ segmentation from CT images using conditional shape-location and unsupervised intensity priors. *Med. Image Anal.* **26**, 1–18 (2015)

7. Cai, J., Lu, L., Xie, Y., Xing, F., Yang, L.: Improving deep pancreas segmentation in CT and MRI images via recurrent neural contextual learning and direct loss function. CoRR, abs/1707.04912 (2017)
8. Farag, A., Lu, L., Roth, H.R., Liu, J., Turkbey, E., Summers, R.M.: A bottom-up approach for pancreas segmentation using cascaded superpixels and (deep) image patch labeling. *IEEE Trans. Image* **26**(1), 386–399 (2017)
9. Badrinarayanan, V., Kendall, A., Cipolla, R.: SegNet: a deep convolutional encoder-decoder architecture for image segmentation. [arXiv:1511.00561](https://arxiv.org/abs/1511.00561) (2015)
10. 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
11. Gibson, E., et al.: Automatic multi-organ segmentation on abdominal CT with dense V-Networks. *IEEE Trans. Med. Imaging* **37**, 1822–1834 (2018)
12. Asaturyan, H., Gligorievski, A., Villarini, B.: Morphological and multi-level geometrical descriptor analysis in CT and MRI volumes for automatic pancreas segmentation. *Comput. Med. Imaging Graph.* **75**, 1–13 (2019)
13. Xie, S., Tu, Z.: Holistically-nested edge detection. In: Proceedings of the IEEE ICCV, pp. 1395–1403 (2015)
14. Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2650–2658 (2015)
15. Perez, L., Wang, J.: The effectiveness of data augmentation in image classification using deep learning. [arXiv:1712.04621](https://arxiv.org/abs/1712.04621) (2017)
16. Asaturyan, H., Villarini, B.: Hierarchical framework for automatic pancreas segmentation in MRI using continuous max-flow and min-cuts approach. In: Campilho, A., Karray, F., ter Haar Romeny, B. (eds.) ICIAR 2018. LNCS, vol. 10882, pp. 562–570. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-93000-8_64
17. Yuan, J., Bae, E., Tai, X.-C.: A study on continuous max-flow and min-cut approaches. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 2217–2224 (2010)
18. Milletari, F., Navab, N., Ahmadi, S.A.: V-Net: fully convolutional neural networks for volumetric medical image segmentation. In: Fourth Fourth International Conference on 3D Vision, pp. 565–571 (2016)
19. Rahman, M.A., Wang, Y.: Optimizing intersection-over-union in deep neural networks for image segmentation. In: Bebis, G., et al. (eds.) ISVC 2016. LNCS, vol. 10072, pp. 234–244. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-50835-1_22