

Comparison of 2D vs. 3D U-Net Organ Segmentation in abdominal 3D CT images

Nico Zettler
Aalen University
Beethovenstr. 1
Burren Campus
Germany, 73430, Aalen,
Baden-Württemberg
nico.zettler@hs-aalen.de

Andre Mastmeyer
Aalen University
Beethovenstr. 1
Burren Campus
Germany, 73430, Aalen,
Baden-Württemberg
andre.mastmeyer@hs-aalen.de

ABSTRACT

A two-step concept for 3D segmentation on 5 abdominal organs inside volumetric CT images is presented. First each relevant organ's volume of interest is extracted as bounding box. The extracted volume acts as input for a second stage, wherein two compared U-Nets with different architectural dimensions re-construct an organ segmentation as label mask. In this work, we focus on comparing 2D U-Nets vs. 3D U-Net counterparts. Our initial results indicate Dice improvements of about 6% at maximum. In this study to our surprise, liver and kidneys for instance were tackled significantly better using the faster and GPU-memory saving 2D U-Nets. For other abdominal key organs, there were no significant differences, but we observe highly significant advantages for the 2D U-Net in terms of GPU computational efforts for all organs under study.

Keywords

Organ bounds, U-Net, Architecture, Abdomen, Segmentation, 3D CT images

1 INTRODUCTION

Manual segmentation of liver, kidneys, spleen and the low-contrast pancreas in axial CT slices is very time consuming for radiologists. The automated just-in-time (JIT) segmentation of 3D patient organ models continues to be an unsolved challenge. Automatic and JIT reconstruction procedures of 3D models are highly relevant for surgery planning and navigation. Another important aspect is the saving of computational power and memory, which differ greatly among various Convolutional Neural Network (CNN) architectures for image segmentation. The 2D U-Net architecture[1] showed promising results in 2015, outperforming conventional models on 2D biomedical segmentation problems and being effective on smaller images. Unlike previously known CNN models, the U-Net concept uses down- and up-sampling steps to resample the condensed feature map to the original size of the input image. By incorporating higher resolution feature information in each upsampling step, semantic segmentation of the input images can be achieved very efficiently by skip con-

nections. The resulting U-shaped architecture was extended to a 3D U-Net edition in 2016 by Cicek et al.[2] by replacing 2D operations with their 3D equivalents. The research focus of this paper targets the question, is the 3D U-Net really better suited for 3D data?

In Big Data studies we expect, parallel segmentation of thousands of 3D volumetric images requires high computational time due to the limited amount of processing nodes and sub-processes for parallel processing. We suppose, 3D U-Nets demand even higher GPU computational and memory overhead for 3D image processing. In this work, the focus lies on proposing solutions with costs as low as possible while keeping or even improving segmentation quality. In our concept, we (1) detect and let U-Nets work in local sub-images for each task (liver, kidneys, spleen, pancreas). We (2) compare two U-Net architectures for quality and GPU-performance carrying out semantic segmentations. (3) Based on a detailed statistical analysis, we recommend the more suitable architectural model to the interested reader. Again, subject to our image data base of 80, could the simpler 2D U-Net perform better than complex and theoretically more powerful 3D U-Nets?

2 RECENT SOLUTIONS

3D segmentation is especially relevant for the emerging field of intervention training and planning using Virtual Reality (VR) techniques. Time-variant 4D VR simulations with breathing simulation are readily available for

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

training and planning of liver interventions[3, 4, 5]. The JIT and high quality reconstruction of all the necessary 3D models remains a main hindrance, i.e. especially abdominal organs are rarely easy to segment. This is due to varying imaging conditions such as contrast agent administration, structure variations and noise.

To this aim in the literature, 3D U-Nets are very often proposed as mightier[6, 7, 8, 9, 10]. However, there are some studies indicating benefits of 2D U-nets in 3D segmentation tasks such as[11, 12, 13]. Nemoto et al.[11] state that low computational burden 2D U-Nets are effective and on par with 3D U-Nets for semantic lung segmentation excluding the trachea and bronchi.

Chang et al.[14] suggested a hybrid approach of 3D and 2D inputs for the evaluation of hemorrhage on head CT scans, which was later adapted by Ushinsky et al.[13] for application to prostate segmentation in MRI images, showing that 2D U-nets are very effective for 3D data. Christ et al.[15] demonstrated a slice-wise application of 2D U-Nets to liver segmentation in combination with 3D random fields. Similarly Meine et al.[12] proposed liver segmentation methods with the assistance of 3D, 2D and three fused 2D U-Net sectional (axial, coronal, sagittal) results in a 2.5D ensemble approach. They find the 2.5D U-Net ensemble results statistically superior for liver segmentation, especially for images with pathologies.

Other recent approaches aim to combine multiple stacked 2D U-Nets[16] and further improve information flow through semantic connections between different components[17]. This concept has also been used in application areas such as colon polyp segmentation[18] and face recognition[19].

In our 2D U-Net setup for the abdominal organs liver, spleen, kidneys and pancreas, we favor axial section training inside of pre-detected organ-specific VOIs. This is for significant radiation dosage savings by high scan pitch. Yet image resolution in axial CT slices is very high.

New aspects in this work are the organ-specific VOI approach for the organs under study such as liver, spleen, kidneys and pancreas. Concluding, an organ-specific architecture dimensionality recommendation for each of the organs is given.

3 PROPOSED SOLUTIONS

Eighty CT-scans and related label data found in various public sources¹ were considered for training and tests of our U-Nets. Their image information differs not only in quality, noise and field of view. But also, the patient-individual volume appearance varies in the amount of

¹ <http://visceral.eu>,
<http://sliver07.org>,
<http://competitions.codalab.org/competitions/17094>

slices (64 to 861), pathological lesions, slice width (1 to 5 mm) and applied contrast agent.

The issue relative to data annotation was mainly the lowly contrasted pancreas organ in the CTs, unavailable with some image sets consisting of congruent intensity and label images from the public sources. We coped with that by four eye reviewed manual segmentation of this occasionally missing structure in the label maps. However, most organ reference segmentation were readily available in the public data sources.

Preprocessing

The orientation of images was changed to Right-Anterior-Inferior (RAI) and zero origin (0.0, 0.0, 0.0). As CNNs are not able to interpret voxel spacings natively, isotropic image resampling with 2.0^3 mm as trade-off for varying image xyz-spacings ($xy \leq 2$ mm, $z \leq 5$ mm) was performed.

Our concept is composed of two different machine learning approaches, bounding box detection using random regression forests (RRF) and U-Nets for semantic segmentation. The RRFs can be used to detect organ VOIs in CT data. The U-Nets use the detected VOIs and segment the contained organs.

Volume of interest bounding box detection

The ground truth corresponding organ VOI bounding box (BB) - needed in this work's evaluation - of each of the organs can be created by scanning the reference segmentation maps for labeled voxel coordinate extremes. To create a three-dimensional BB vector for each organ, for each coordinate direction (x, y, z), we iterate in orthogonal slices through the organ's label map and save the extreme limits in a 6D BB vector to serve as ground truth BBs.

Alternatively, VOI-based organ extraction[20, 21, 22] is readily available for organ-specific VOI detection of the organ ensemble (liver, still left kidney, right kidney, spleen, pancreas). This step also solves the FoV problem in the scanner z-direction, as CT scans cover variable body portions. Our currently investigated scheme to learn bounding boxes of VOIs using RRFs is summarized in the following.

In our concept, we use Random Regression Forests (RRF) to determine the location and extent of abdominal organs[22]. As seen in Fig. 1, the RRF training step expects scans and ground truth VOIs as input.

A three-dimensional VOI b_c of an organ c can be described by using a 6D vector $b_c = (b_c^{Left}, b_c^{Right}, b_c^{Anterior}, b_c^{Posterior}, b_c^{Head}, b_c^{Foot})$ with coordinates in mm[22]. We run over all voxels $p = (x_p, y_p, z_p)$, which are within a specified radial distance ($r = 5$ cm) from the scan medial axis. The distance d between such a voxel and each of the VOI walls,

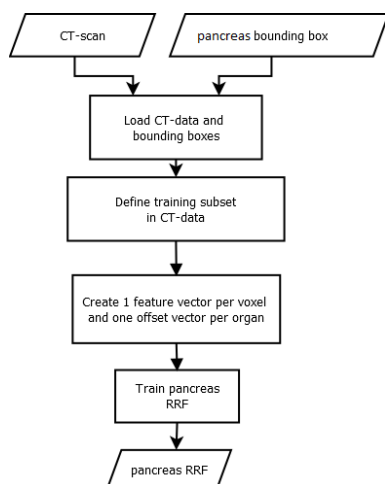


Figure 1: The inputs for the training process are CT scans and ground truth VOIs of a targeted pancreas. We create one feature vector and one mm offset vector for each voxel that is part of a predefined medial cylinder subset in the scan[22]. The trained RRF is able to predict the offset between a voxel and an organ’s VOI walls.

can be calculated by using $d(p) = (x_p - x^{Left}, x_p - x^{Right}, y_p - y^{Anterior}, y_p - y^{Posterior}, z_p - z^{Head}, z_p - z^{Foot})$ and is saved as the offset-vector to be learned. In contrast to Criminisi et al.[22], we use only 50 feature boxes, that are evenly distributed on three spheres ($r = 5\text{ cm}, 2.5\text{ cm}, 1.25\text{ cm}$) to generate the input feature vector. The feature boxes F_j are intended to capture the spatial and intensity context of the current voxel. For this purpose, the mean intensities of the feature boxes are calculated and saved in the feature vector.

Within the first step of our application concept (Fig. 2, left) for every single organ, the RRF locates the VOI for a particular organ as BB.

A detected VOI’s information (intensities, labels) is finally resampled to $96 \times 96 \times 96$ as CNN input.

Training and application of 2D and 3D U-Net architectures

Within the 2nd stage, as shown in the bird’s eye view in Fig. 2, right, the obtained VOI is either passed to a 2D[1] or a 3D U-Net[2] for semantic labeling.

The training data for our U-Net consists of the expert segmentations and ground truth bounding box contents, as schematized in Figs. 2, 3. The VOIs are then used to locally extract the intensity and label data from the CT scans. As input, a U-Net receives a VOI from the intensity data, while the corresponding label data is connected to the output. We use lightly modified 2D and 3D U-Net architectures vs. Ronneberger et al.[1]. Our architecture consists of four down- and up-scaling steps connected by skip connections and max-pooling while down-sampling. For the 3D U-Net, an additional layer

depth dimension is added while all other design elements are kept constant.

The U-Net uses the data contained inside a given VOI bounding box to segment the corresponding organ. The output is a segmentation map of the full target organ in a probability range from 0 to 100%. The threshold for all organs except pancreas was chosen as 50%. For pancreas 30% was empirically found best.

The U-Nets were trained using batches of size 8 over 100 epochs. In addition, Adam optimization and a cross entropy loss function were used. We train one U-Net for each organ, using a ReLU activation function.

Evaluation, metrics and statistics

For fair comparison, in each cross-validation iteration, we use the same 64 training images for both U-Nets, and test with the same 16. This way, we prevent data contamination by separate training and test sets. The difference between the U-Net architectures is input, output and filter kernel layer dimension only. For 2D U-Nets, input and output size is 96×96 and 2D filter kernels are used in the layers. Accordingly for 3D U-Nets, we have $96 \times 96 \times 96$ in- and outputs and use 3D convolutions in the architecture layers.

A 5-fold randomized cross-validation using 4:1 splits was used. Five iterations yield 80 quality measures for each organ being used in our statistics (Fig. 1). This means five new models for each organ are trained with randomly selected training data and used in the evaluation. To purely concentrate on the influence associated with U-Net dimensionality, the reference VOIs were utilized in the evaluation of this study. The metric used for analysis is the Dice similarity coefficient:

$$DSC = \frac{2 \cdot |U \cap G|}{|U| + |G|}, \quad (1)$$

where U represents the voxel set from U-Net object segmentation and G the ground truth voxels. A DSC value of 1 indicates perfect segmentation, a value of near 0 poor segmentation. From the DSC results, we calculate means and standard deviations, medians and Inter-Quartile-Ranges (IQR) as measures of accuracy and precision. Statistical assessments using paired T-tests and Wilcoxon-Signed-Rank(WSR)-test were performed with GNU-R 4.0.3. We finally recommend the dimensionality of the U-Net architecture based on accuracy, i.e. greater mean or median, and precision, i.e. smaller standard deviation or IQR and smaller number of outliers. We also keep an eye on the ratio between quality and GPU efforts in time and memory.

4 EXPERIMENTAL RESULTS

Qualitative liver results (Fig. 4) show the 2D U-Net (top) superior with a more evenly distributed segmentation surface coverage. The 3D U-Net (bottom) obviously suffers from under-segmentation as more brown

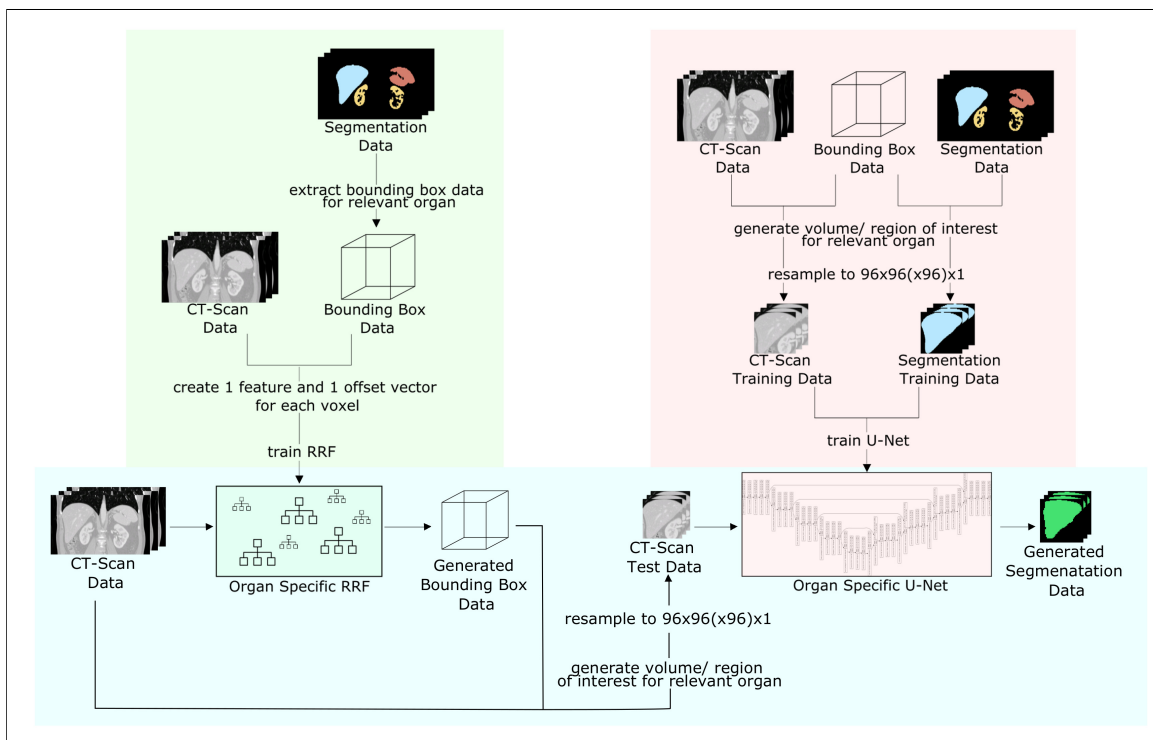


Figure 2: Bird's eye perspective on our current segmentation concept with RRF for VOI detection and subsequent U-Net. Both training (green, red) and application (blue) concepts are shown. In this work, the right part with regards to the U-Nets is focused in the evaluation

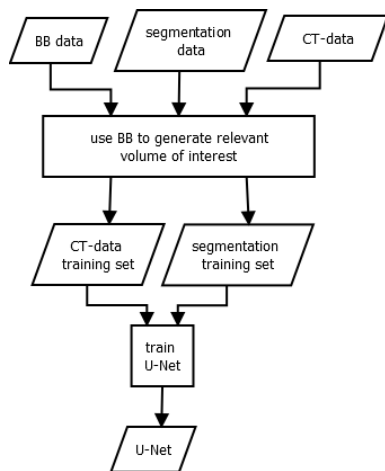


Figure 3: The inputs for the training process are ground truth bounding boxes VOIs, CT scans and their corresponding segmentation maps. The box cuts out the CT- and segmentation data to extract the relevant image region. Inside the organ VOIs, the segmentation is learned. The process results in organ-wise U-Nets

ground truth surface is visible in Fig. 4 (bottom). Referring to Tab. 1, precision is also better for the 2D U-Net due to lower std. deviations and IQRs.

For kidneys, the 2D U-Net is highly significantly better in both T- and WSR-tests. The boxplots with a small x for the mean in Figs. 5, 6 confirm visually: the 2D

U-Net is significantly better for liver and highly significantly better for kidneys (Figs. 5(a), 6(a), 5(b), 6(b)). Spleen segmentation shows an advantage for the 2D U-Net by trend as shown in Figs. 5(c), 6(c)) and Tab. 1.

The 3D U-Net possesses higher accuracy and precision for the low-contrast pancreas. Higher precision is indicated by smaller standard deviations and less outliers in Tab. 1 and Figs. 5(d) and 6(d). With regards to accuracy measured by medians, the 3D U-Net has an edge by trend of 0.02 over the 2D U-Net, while precision measured by IQR is the same.

The GPU-Performance evaluation in Tab. 2 shows the superiority of the 2D U-Net in this study's scale with an 80 images data base. The ratio of quality to GPU resources is always better for the 2D U-Net. In Tab. 2 GPU memory saving is >6-fold in the training stage and >5-fold in the model application, trivially a highly significant result as there are no variations, i.e. no standard deviation and IQR. With respect to GPU calculation times measured in training, the 2D U-Net is precisely 40 seconds or 7% faster on average, again a highly significant mean result. In the U-Net model application, we can observe a weakly significant ($p < 0.1$) advantage for the 2D U-Net of 37% to 75%.

5 CONCLUSION

Surprisingly in this study, 2D U-Nets are favorable regarding the ratio of quality vs. computing costs.

DSCs: Organ	Mean±Std.		Median~IQR	
	2D U-Net	3D U-Net	2D U-Net	3D U-Net
Liver	0.94±0.03*	0.93±0.04	0.95~0.02**	0.94~0.03
R. kidney	0.91±0.05***	0.89±0.05	0.92~0.03***	0.90~0.05
L. kidney	0.92±0.05***	0.86±0.14	0.93~0.03***	0.89~0.08
Spleen	0.93±0.04	0.92±0.04	0.94~0.03	0.93~0.03
Pancreas	0.57±0.19	0.59±0.15	0.60~0.21	0.61~0.21

Table 1: Mean DSCs with standard deviations (Mean±Std.) and Median DSCs with Inter-Quartile-Range (IQR) (Median~IQR) of 2D and 3D U-Nets from 5-fold randomized cross-validation experiments using 4:1 splits of the 80 images into training and test data.

Legend: We use statistical standard notation for found significances: *, **, ***: $p < 0.05$; $p < 0.01$; $p < 0.001$ from T-tests (paired) and Wilcoxon-Signed-Rank-Tests on the right of the **favorable U-Net result**.

GPU-Performance: Organ	Memory				Time			
	Training [MiB]		Application [MiB]		Training [min:sec]		Application [sec]	
	2D U-Net	3D U-Net	2D U-Net	3D U-Net	2D U-Net	3D U-Net	2D U-Net	3D U-Net
Liver	1693	10957	1693	9117	9:26	10:07	1.47	3.18
R. kidney	1693	10957	1693	9117	9:28	10:07	0.40	0.55
L. kidney	1693	10957	1693	9117	9:26	10:07	0.40	0.55
Spleen	1693	10957	1693	9117	9:27	10:08	0.40	0.55
Pancreas	1693	10957	1693	9117	9:28	10:08	0.40	0.55
Mean±Std.	1693±0***	10957±0	1693±0***	9117±0	9:27±0:01***	10:07±0:01	0.61±0.43	1.07±1.05
Median~IQR	1693~0***	10957~0	1693~0***	9117~0	9:27~0:01+	10:07~0:01	0.40~0.003+	0.55~0.001
Mean Improvement	647.19%	N/A	538.51%	N/A	107.13%	N/A	175.22%	N/A
Median Improvement	647.19%	N/A	538.51%	N/A	107.07%	N/A	137.10%	N/A

Table 2: GPU-Performance table with memory consumption for training and application to the left and training and application times to the right.

Legend: 1 MiB=1.048581024 MB=1024² bytes. We use statistical standard notation for found significances: +; *, **, ***: $p < 0.10$; $p < 0.05$; $p < 0.01$; $p < 0.001$ from T-tests (paired) and Wilcoxon-Signed-Rank-Tests on the right of the **favorable U-Net result**.

The liver and spleen hold the greatest volume in our abdominal organ group. The liver and especially kidney competition is significantly won by the 2D U-Net ($p < 0.05$). The liver is a difficult organ often with a variably filled stomach as a neighbor.

Kidneys can be regarded as easy organs lighted by contrast agent and inside fatty tissue with low CT intensity.

A better posed training for the 2D U-Net could be the reason for the 2D U-Nets' better results for liver and kidney tasks. A higher relative number training elements is used, i.e. axial slice pixels, vs. the number of net weights.

The spleen results are in favor of the 2D U-Net regarding the mean and medians by trend. The 2D U-Net is also favorable for a less number of outliers (Fig. 5(c)).

The difficult pancreas does not provide many axial training slices useful for the 2D U-Net, as its elongation is not prominent on the z-axis. We suppose this is the reason, why the 3D U-Net wins here by mean trend in terms of accuracy and precision, i.e. higher mean and lower standard deviation and lower number of outliers (Fig. 5(d)). This win is supported by higher accuracy for the 3D U-Net in terms of medians. However, the race is not decided by significant differences making the 2D U-Net still very attractive for some users with GPU performance and memory concerns.

Abdominal volumetric CT images and key organ segmentation were analyzed. This new study shows interesting results from competing U-Net architectures, especially focusing different dimensionalities of net filter bank kernels and quality vs. GPU performance. The interested reader can now select a particular U-Net architecture, primarily whether to use a computational inexpensive design. Finally in this study's scope, a humble recommendation for the 3D U-Net could be given for the pancreatic organ in terms of better accuracy by trend and smaller standard deviation only. The IQRs as an alternative measure of precision are on par, and regarding the median the 3D U-Net wins just by a small trend. We co-conclude, because of the deeper layering structure and thus more trainable weights, a 3D U-Net needs significantly more training data vs. possible overfitting to outpace 2D U-Nets.

As training volumes are normalized to a square or cube of 96 voxels, the GPU memory consumption in Tab. 2 is always constant. Therefore differences are trivially highly significant, as no varying results occur. We observe consistently lower memory consumption for 2D U-Nets. The memory effort in training is higher for the 3D U-Net including more space for administrative overhead data. Thus, 2D U-Nets can run on affordable 2-4 GB GPUs for 3D CT volume segmentation. 3D U-




Figure 4: Two liver segmentations (anterior view) with 2D U-Net (top) and 3D U-Net (bottom). The coverage of the 2D U-Net result (top) appears more equally distributed, i.e. more sensitive. Brown: Reference and purple: U-Net CNN. N.B., for the 2D model result on top, the oscillation pattern between reference (brown) and 2D CNN segmentation (purple) is much denser, thus showing out better mean local quality.
Legend: reference (brown) and U-Net segmentation (purple)

Nets definitely need currently totally overpriced 12 GB GPUs.

The timing measures in Tab. 2 clearly speak out for the 2D U-Net. Training is highly significantly 40 seconds or 7% faster on average using the 2D U-Net. However regarding application, in the trained model prediction, the differences are not so striking with a weak significance by median. 37% to 75% improvement can be

achieved, however in the range of one second, which is practically unimportant.

As final and bold conclusion regarding our study design and results, we can recommend using the 3D U-Net subject to the amount of data we used here - for pancreas only. The conclusions are justified by statistically significant or by trend quality and GPU-computation performance results for all organs under study. We suppose, 3D U-Nets may overcome in quality when

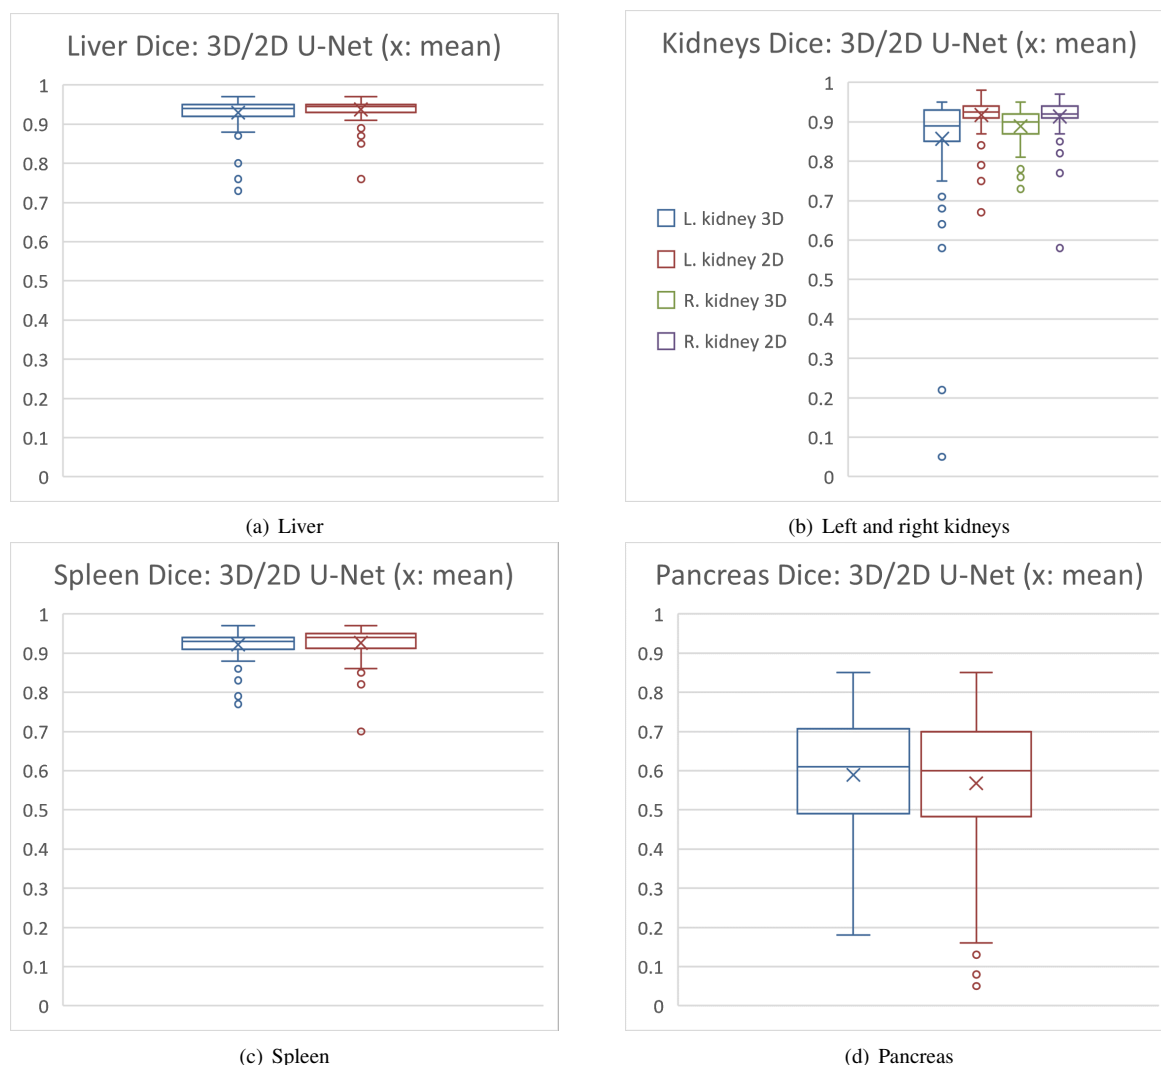


Figure 5: Overview DSC boxplots: 3D (blue, green) and 2D U-Net (brown, purple) on the x-axis vs. DSCs on the y-axis: 2D U-Nets in favor for (a) liver and (b) kidneys (left k. and 3D U-Net: blue, left k. and 2D U-Net: brown; right k. and 3D U-Net: green, right k. and 2D U-Net: purple). (c) Spleen comes out with an edge for the 2D U-Net by trend. (d) Mixed results: 2D U-Net wins the accuracy contest, but loses the precision contest in terms of lower standard deviation (cf. Tab. 1) and regarding less outliers for pancreas.

Legend for (a), (c) and (d): blue boxplots: 3D U-Net; red boxplots: 2D U-Net.

using several hundreds of training images. However, this comes approximately with an order of magnitude higher additional computational burden.

For the first time, an original and significant comparison of U-Net architecture dimension is provided to the reader focusing the key abdominal organs of liver, spleen, kidneys and pancreas. The reader can decide, which approach is appropriate for his concrete target organ, amount of training data and used GPU or cluster nodes, parallel process design, e.g. for atlas-based usage of U-Nets as encountered in multi-classifier fusion[12].

Regarding the difficult pancreas with mixed DSC results in this study, we plan to train 2D U-Nets using a elongation optimized algorithm[23, 24, 25] to provide

slices oriented axially along its main central curve, to better reflect its orientation to generate more training slices for 2D U-Nets. On the other hand, more training images shall be used to explore the 3D U-Nets' theoretical advantage under better conditions, as we have discussed, for its training and application.

The scientific explanation of the U-Net methods is given detailed in the original works of Ronneberger and Cicek et al.[2, 1], and we lift these methods here to compare them. The current limitations of this study will manifest as improvements in the future. Development will cover improved bounding box detection. At this state of our research, our current RRF bounding box detection would confound the core message of this

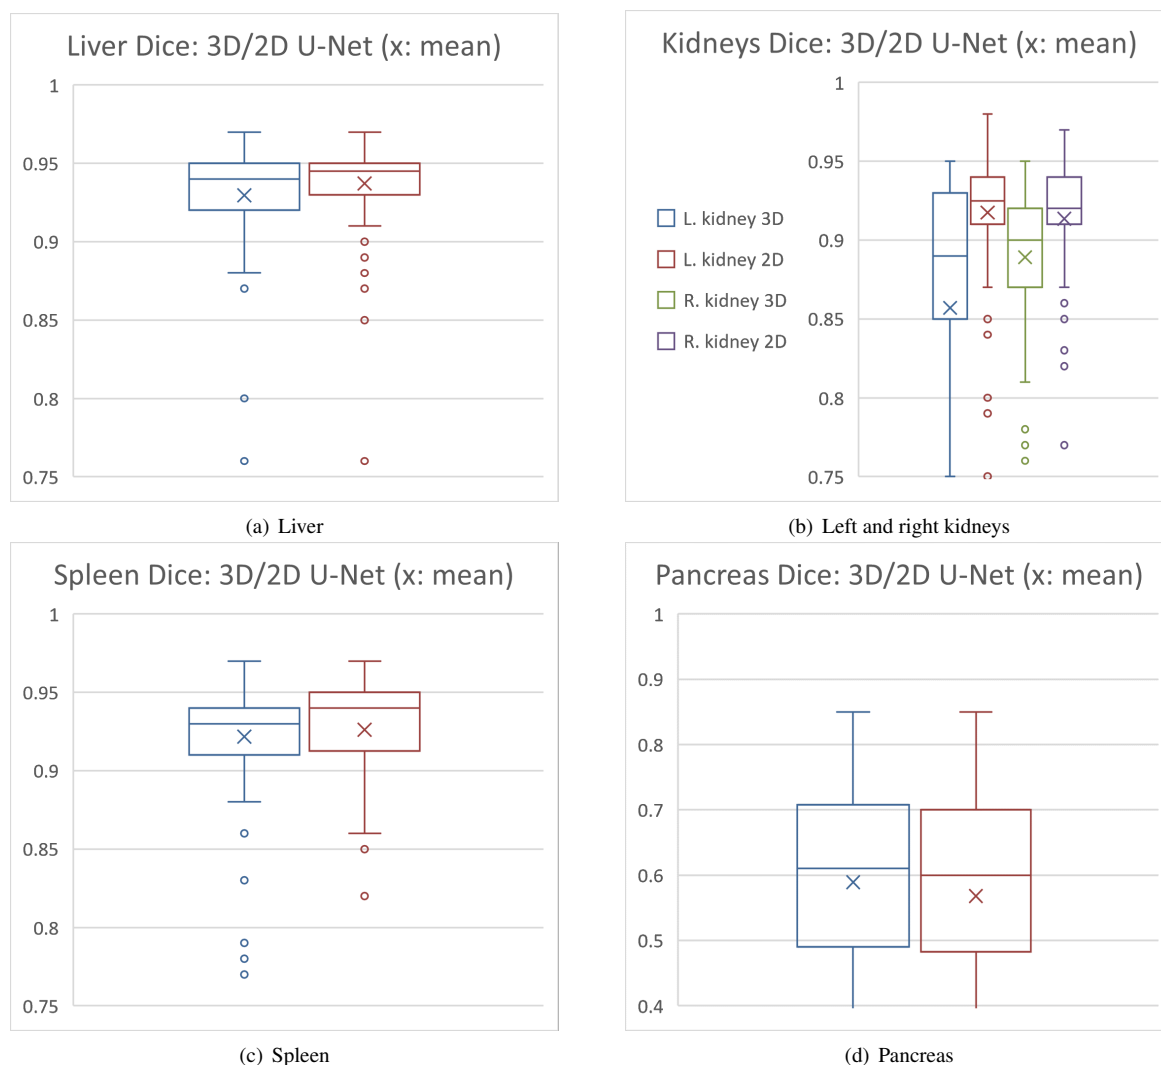


Figure 6: Zoomed DSC boxplots: 3D (blue, green) and 2D U-Net (brown, purple) on the x-axis vs. DSCs on the y-axis: 2D U-Nets in favor for (a) liver and (b) kidneys (left k. and 3D U-Net: blue, left k. and 2D U-Net: brown; right k. and 3D U-Net: green, right k. and 2D U-Net: purple). (c) Spleen comes out with an edge for the 2D U-Net by trend. (d) Mixed results: 2D U-Net wins the accuracy contest, but loses the precision contest in terms of lower standard deviation (cf. Tab. 1) and regarding less outliers for pancreas.

Legend for (a), (c) and (d): blue boxplots: 3D U-Net; red boxplots: 2D U-Net.

paper, the aim of which was to focus purely on U-Net performances.

6 ACKNOWLEDGMENTS

German Research Foundation: DFG MA 6791/1-1;
Nvidia GPU grant;
Foundation Kessler+Co. for Education and Research:
EXPLOR19-AM.

REFERENCES

- [1] Olaf Ronneberger et al. “U-Net: Convolutional Networks for Biomedical Image Segmentation”. In: *Medical Image Computing and Computer Assisted Intervention - MICCAI 2015*, 2015.
- [2] Ozgun Cicek et al. “3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation”. In: *Medical Image Computing and Computer Assisted Intervention - MICCAI 2016*, 2016.
- [3] Andre Mastmeyer et al. “Population-based respiratory 4D motion atlas construction and its application for VR simulations of liver punctures”. In: *Proc. SPIE Medical Imaging: Image Processing*. Vol. 10574. International Society for Optics and Photonics. 2018, p. 1057417.
- [4] Andre Mastmeyer et al. “Interpatient Respiratory Motion Model Transfer for Virtual Reality Simulations of Liver Punctures”. In: *Journal of*

- World Society of Computer Graphics - WSCG 25.1* (2017), pp. 1–10.
- [5] Andre Mastmeyer et al. “Real-Time Ultrasound Simulation for Training of US-Guided Needle Insertion in Breathing Virtual Patients”. In: *Studies in Health Technology and Informatics*. Vol. 220. IOS Press, 2016, p. 219.
- [6] Getao Du et al. “Medical image segmentation based on u-net: A review”. In: *Journal of Imaging Science and Technology* 64.2 (2020), pp. 20508–1.
- [7] Nahian Siddique et al. “U-Net and its variants for medical image segmentation: theory and applications”. In: *arXiv preprint arXiv:2011.01118* (2020).
- [8] Pavlo Radiuk. “Applying 3D U-Net architecture to the task of multi-organ segmentation in computed tomography”. In: *Applied Computer Systems* 25.1 (2020), pp. 43–50.
- [9] Nima Tajbakhsh et al. “Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation”. In: *Medical Image Analysis* 63 (2020), p. 101693.
- [10] Geert Litjens et al. “A survey on deep learning in medical image analysis”. In: *Medical image analysis* 42 (2017), pp. 60–88.
- [11] Takafumi Nemoto et al. “Efficacy evaluation of 2D, 3D U-Net semantic segmentation and atlas-based segmentation of normal lungs excluding the trachea and main bronchi”. In: *Journal of radiation research* 61.2 (2020), pp. 257–264.
- [12] Hans Meine et al. “Comparison of u-net-based convolutional neural networks for liver segmentation in ct”. In: *arXiv:1810.04017* (2018).
- [13] Alexander Ushinsky et al. “A 3D-2D Hybrid U-Net Convolutional Neural Network Approach to Prostate Organ Segmentation of Multiparametric MRI”. In: *American Journal of Roentgenology* 216.1 (2021), pp. 111–116.
- [14] P.D. Chang et al. “Hybrid 3D/2D Convolutional Neural Network for Hemorrhage Evaluation on Head CT”. In: *American Journal of Neuroradiology* 39 (July 2018). DOI: 10.3174/ajnr.A5742.
- [15] Patrick Ferdinand Christ et al. “Automatic Liver and Lesion Segmentation in CT Using Cascaded Fully Convolutional Neural Networks and 3D Conditional Random Fields”. In: *Lecture Notes in Computer Science* (2016), 415–423. ISSN: 1611-3349. DOI: 10.1007/978-3-319-46723-8_48. URL: http://dx.doi.org/10.1007/978-3-319-46723-8_48.
- [16] Debesh Jha et al. “DoubleU-Net: A Deep Convolutional Neural Network for Medical Image Segmentation”. In: July 2020, pp. 558–564. DOI: 10.1109/CBMS49503.2020.00111.
- [17] Zhiqiang Tang et al. *CU-Net: Coupled U-Nets*. Aug. 2018.
- [18] Dinh Viet Sang et al. *AG-CUResNeSt: A Novel Method for Colon Polyp Segmentation*. 2021. arXiv: 2105.00402 [eess.IV].
- [19] In Seop Na et al. “Facial UV map completion for pose-invariant face recognition: a novel adversarial approach based on coupled attention residual UNets”. In: *Human-centric Computing and Information Sciences* 10.1 (Nov. 2020). ISSN: 2192-1962. DOI: 10.1186/s13673-020-00250-w. URL: <http://dx.doi.org/10.1186/s13673-020-00250-w>.
- [20] Daria Kern et al. “3D Bounding Box Detection in Volumetric Medical Image Data: A Systematic Literature Review”. In: 2020. arXiv: 2012.05745 [eess.IV].
- [21] Andre Mastmeyer et al. “Random forest classification of large volume structures for visuo-haptic rendering in CT images”. In: *Proc. SPIE Medical Imaging: Image Processing*. 2016, 97842H.
- [22] Antonio Criminisi et al. *Decision Forests for Computer Vision and Medical Image Analysis*. London: Springer, 2013.
- [23] Christian Reinbacher et al. “Variational segmentation of elongated volumetric structures”. In: *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE. 2010, pp. 3177–3184.
- [24] Maria Kallergi et al. “Automatic Segmentation of Pancreatic Tumors in Computed Tomography”. In: *Handbook of Biomedical Image Analysis*. Springer, 2005, pp. 183–228.
- [25] Andre Mastmeyer et al. “Model-based Catheter Segmentation in MRI-Images”. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention - MICCAI*. 2015.

