

MSS U-Net: 3D segmentation of kidneys and tumors from CT images with a multi-scale supervised U-Net

Wenshuai Zhao^{a,*}, Dihong Jiang^a, Jorge Peña Queralta^b, Tomi Westerlund^b

^a School of Information Science and Technology, Fudan University, China

^b Turku Intelligent Embedded and Robotic Systems Lab, University of Turku, Finland

ARTICLE INFO

Keywords:

Kidney segmentation
Kidney tumor segmentation
Multi-scale supervision
Exponential logarithmic loss
3D U-Net

ABSTRACT

Accurate segmentation of kidneys and kidney tumors is an essential step for radiomic analysis as well as developing advanced surgical planning techniques. In clinical analysis, the segmentation is currently performed by clinicians from the visual inspection of images gathered through a computed tomography (CT) scan. This process is laborious and its success significantly depends on previous experience. We present a multi-scale supervised 3D U-Net, MSS U-Net to segment kidneys and kidney tumors from CT images. Our architecture combines deep supervision with exponential logarithmic loss to increase the 3D U-Net training efficiency. Furthermore, we introduce a connected-component based post processing method to enhance the performance of the overall process. This architecture shows superior performance compared to state-of-the-art works, with the Dice coefficient of kidney and tumor up to 0.969 and 0.805 respectively. We tested MSS U-Net in the KITS19 challenge with its corresponding dataset.

1. Introduction

Renal cell carcinoma is one of the most common genitourinary cancers with the highest mortality rate [3]. An accurate segmentation of kidney and tumor based on medical images, such as images from computed tomography (CT) scans, is the cornerstone to appropriate treatment. In computer-aided therapy, the success of this step is an essential prerequisite to any other processes. The segmentation process therefore is the key for exploring the relationship between the tumor and its corresponding surgical outcome, and aids doctors in making more accurate treatment plans [8]. Nonetheless, the manual segmentation of organs or lesions has the potential to be highly time-consuming, since a radiologist may need to label out target regions in hundreds of slices for one patient. The need for more accurate automatic segmentation tools is thus evident.

In recent years, considerable research efforts have been put towards the automatic segmentation of kidney and kidney tumors from CT images. In particular, novel deep learning techniques have played a key role. Previous works have mostly been focusing on the utilization of unsupervised training methods, including threshold-based methods, region-based methods (e.g., region growing), clustering-based methods (e.g., fuzzy C-means or Markov Random Fields), edge-detection

methods, or deformable model methods [7]. More recently, the application of deep artificial neural networks for medical image segmentation has gained increased momentum, in particular 3D convolutional neural networks [4]. Neural networks providing end-to-end analysis (from raw images to segmented images) are more generic and therefore do not suffer of some of the challenges in previous methods. For instance, threshold-based techniques yield the best results when the regions of interest have a significant difference in intensity with respect to the background, but present problems in more homogeneous images, significantly reducing their performance and limiting their applicability.

Fig. 1 shows abdominal CT images from three different patients. Because the location, size and shape of kidney and tumor vary considerable across patients, the segmentation of kidneys and kidney tumors is challenging. The major challenges can be attributed to the following considerations. First, the location of tumors may vary significantly from patient to patient. The tumor can appear anywhere inside the organs or attached to the kidneys. Trying to predict the location from experience and previous knowledge is unfeasible both from a human's and a computer's perspective. Second, the shape and size of tumors present huge diversity. Tumors in some patients can be very small on the kidneys while others can almost erode the whole kidney. Besides, their shapes might be regular, distorted, or scattered. Third, the tissue of tumors is

* Corresponding author.

E-mail address: wezha@utu.fi (W. Zhao).

<https://doi.org/10.1016/j.imu.2020.100357>

Received 9 April 2020; Received in revised form 20 May 2020; Accepted 20 May 2020

Available online 2 June 2020

2352-9148/© 2020 The Authors.

Published by Elsevier Ltd.

This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).




Fig. 1. Illustration of sample segmented images from three patients in the KiTS19 dataset [13]. The first row is the transverse plane, and the second row is the 3D reconstruction. Red voxels denote kidney while green voxels denote kidney tumor. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

also heterogeneous: the large amount of different subtypes of renal cell carcinoma, coupled with their heterogeneity, can bring diverse intensity attribution in CT images. Finally, the simultaneous segmentation of kidney and kidney tumor from raw full-scale CT images may lead to additional difficulties due to the co-existence of multiple labels and large background sizes.

1.1. Background

In recent years, deep learning methods have emerged as a promising solution for segmentation in medical imaging. Among these methods, convolutional neural network (CNN) architectures have already outperformed traditional algorithms in computer vision tasks in general [16], and in the segmentation of CT images in particular [4]. Fully convolutional network (FCN) architecture is a notably powerful end-to-end training segmentation solution [19]. This type of architecture is able to produce raw-scale, pixel-level labels at the images, and it is the current state-of-the-art across multiple domains. Other recent works have taken FCN as a starting point towards deeper and more complex segmentation architectures, such as SegNet [2] introducing the decoder part to enhance the segmentation performance, the feature pyramid networks (FPN), mainly used in object detection [17], pyramid scene parsing networks (PSPNet), utilized in scene understanding [30], Mark R-CNN, for object instance segmentation [11], or DeepLab series for semantic image segmentation [16].

While there is a wide range of deep learning methods for image segmentation, medical images present significant differences from natural images. Some of the most distinct features are the following: First, medical images are relatively simple, when compared to the wide range of semantic patterns, colors and intensities in natural images. This increased homogeneity across individual images hinders the identification of patterns and regions of interest; Second, the boundary between organs, lesions or other regions of interest is fuzzy, and the images are not obtained through passive observation of the subject but instead through an active stimulus. In consequence, methods and neural network architectures employed in other types of image segmentation cannot be directly extrapolated to the medical field. When taking into account the way in which these images are obtained, the difference becomes even more significant. Medical images are often obtained through a volumetric sampling of the subject's body. This characteristic and key differential aspect can be taken as an advantage towards integrating three-dimensional neural network architectures. Among those, one of the most popular architectures to date is 3D U-Net presented in 2016 [5].

3D U-Net is one of the most well-known methods and widely used

three-dimensional architecture for medical image segmentation [21], inspired by FCNs and encoder-decoder models. The 3D U-Net architecture has been further developed and new solutions are built on top of it, for example: Nabila et al. proposed the use of Tversky loss to enhance the performance of the attention mechanism in U-Net [1], whereas Zhe et al. proposed an improved U-Net architecture in which the authors perform liver segmentation in CT images by minimizing the graph cut energy function [18].

A smaller number of research have been focused on the segmentation of kidneys or kidney tumors. An ideal architecture should further extend the existing end-to-end networks for pixel-wise segmentation. In addition, the nature of three-dimension data could provide higher levels of correlation to employ. In this direction, Yang et al. combined in 2018 a basic 3D FCN and a pyramid pooling module to enhance the feature extraction capability, with a network able to segment kidneys and kidney tumor at the same time [28]. However, the experiment is conducted on the region of interest (ROI) only, rather than on raw CT images. This significantly reduces the complexity of the segmentation task as well as the clinical practicality. Yu et al. proposed in 2019 a novel network architecture combining vertical patch and horizontal patch based sub-models to conduct center pixel prediction of kidney tumors [29]. However, owing to the nature of the data being segmented, its training and inference processes become increasingly challenging, and this type of solution minimizes the problem with a simpler output.

In general, we have found that most algorithms in the field of medical image segmentation take the U-Net architecture as a starting point for further developments. Fabian et al. implemented a well-tuned 3D U-Net (nnU-Net) and demonstrated its applicability and potential with top rankings in multiple medical image segmentation challenges [15]. In this paper, we get inspiration from their work towards an end-to-end framework to perform segmentation of kidneys and kidney tumors simultaneously from CT images. The proposed network architecture is also developed from the original 3D U-Net architecture [5]. Because of a clear similarity of CT images across multiple patients, we make the assumption that the original 3D U-Net is capable of extracting sufficient features for recognition. Therefore, we do not consider the utilization of additional modules or branches behind the main 3D U-Net backbone, such as residual modules [20], FPN [17], or attention gates [24], among others. Instead, we focus on optimizing the training and enhancing the performance of the original 3D U-Net architecture.

1.2. Contribution and structure

In this work, we explore the potential of the 3D U-Net architecture through the combination of deep supervision and the exponential

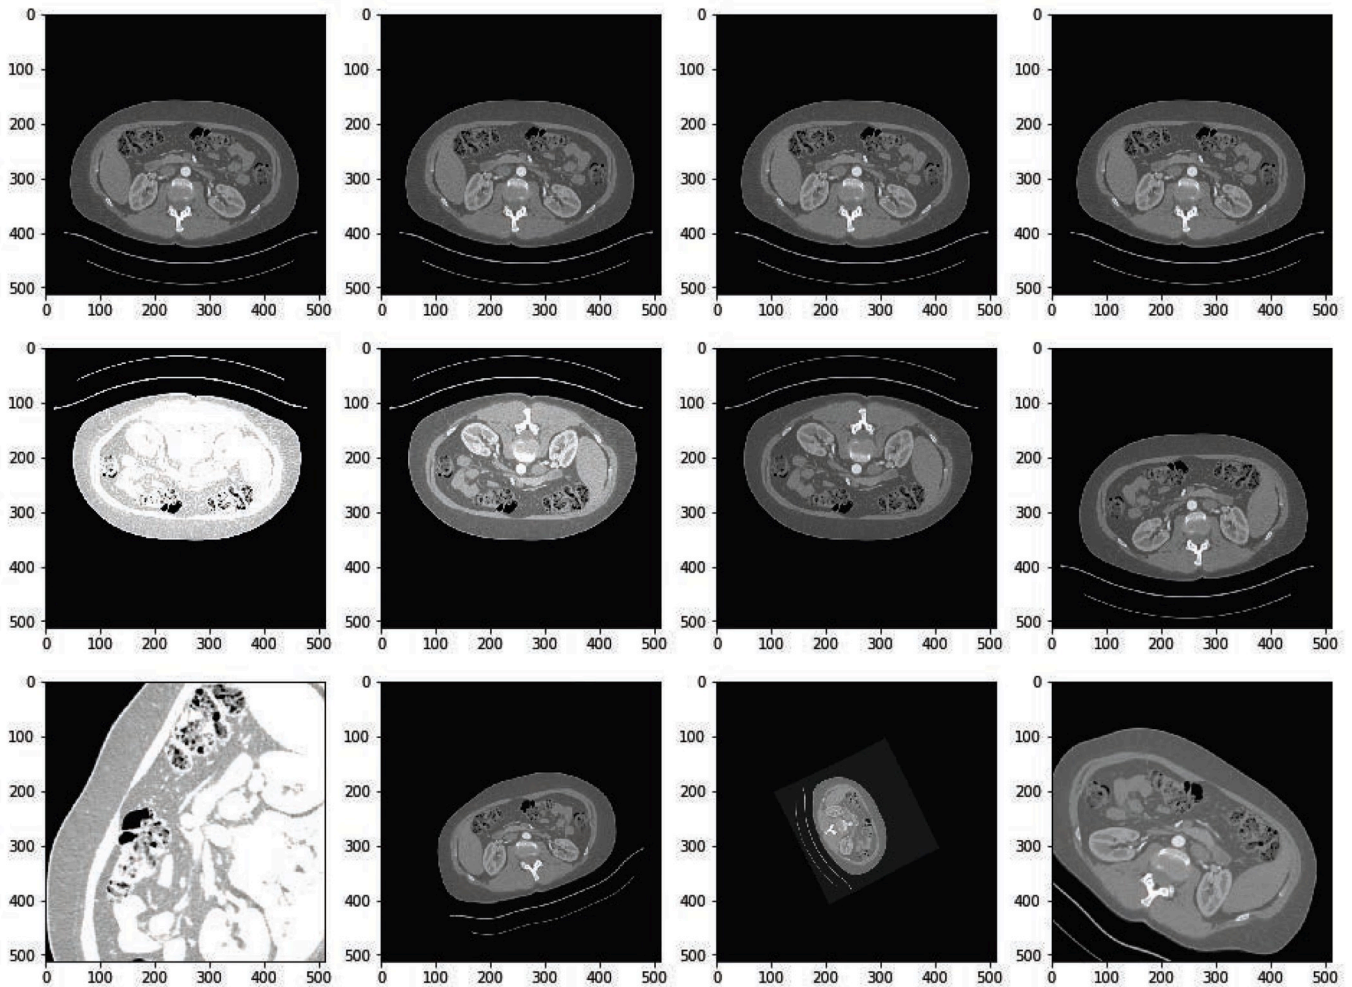


Fig. 2. The effect of data augmentation techniques used in our method. First row is the original images; second row shows the effect of contrast augmentation and mirroring; third row shows elastic deforming, scaling, gamma correction and rotation. Of note, we show 2D images instead of the actual 3D ones for simplicity.

logarithmic loss. With the increased number of hyperparameters, we can train better the network and utilize our results as a baseline to analyze the performance of 3D U-Net. Thus, the main contributions of this work are the following:

1. The introduction of a multi-scale supervision scheme for the 3D U-Net that tunes the network to conduct an accurate prediction from deep layers; and
2. The utilization of exponential logarithmic loss [27] to alleviate the class imbalance problem between foregrounds (kidney and tumor voxels) and background.

By integrating these two approaches, we have improved the performance of the original 3D U-Net architecture. Furthermore, we have designed a connected-component based post processing method to remove disconnected voxels that are detected as evident false positives.

The remainder of this paper is organized as follows. Section 2 introduces the neural network architecture, and the different methods employed in the experiments, after which Section 3 introduces experimental results and analysis of the network's performance. In Section 4, we discuss on the different strategies being taken towards medical image segmentation. Finally, Section 5 concludes the work and outlines future research directions.

2. Methodology

The current approaches to medical image segmentation with deep learning can be roughly classified in two trends. First, those in which the input data to the neural networks is not the raw data but instead the regions of interest (ROIs). This naturally allows higher accuracy and performance across multiple metrics. Nonetheless, the overall performance can be significantly impacted by the uncertainty in the detection and extraction to ROIs. Second, in a more recent trend, end-to-end segmentation architectures have been introduced. In these, raw images are fed to the network and the output of the network is pixel-wise segmentation with images of the same size as the inputs. We follow this end-to-end architecture to segment both kidneys and kidney tumors simultaneously from raw volumetric CT images. This allows for direct application of our methods in clinical settings.

The CT data that is fed to a CNN consists of abdominal CT images with hundreds of slices. Each three-dimensional region is called a voxel. A typical input is $512 \times 512 \times 200$ voxels, where 200 is the number of slices and 512×512 the resolution of each image in pixels. Owing to the large input size, it is unfeasible to feed the data into the network at once especially when graphic processing units (GPUs) are used to accelerate CNNs. In GPUs, the challenges stem from both the limited amount of memory and the required computing power in general. Therefore, we follow the recent trend in both patch-based training and inference for our network architecture.

The rest of this section describes the different steps taken for training

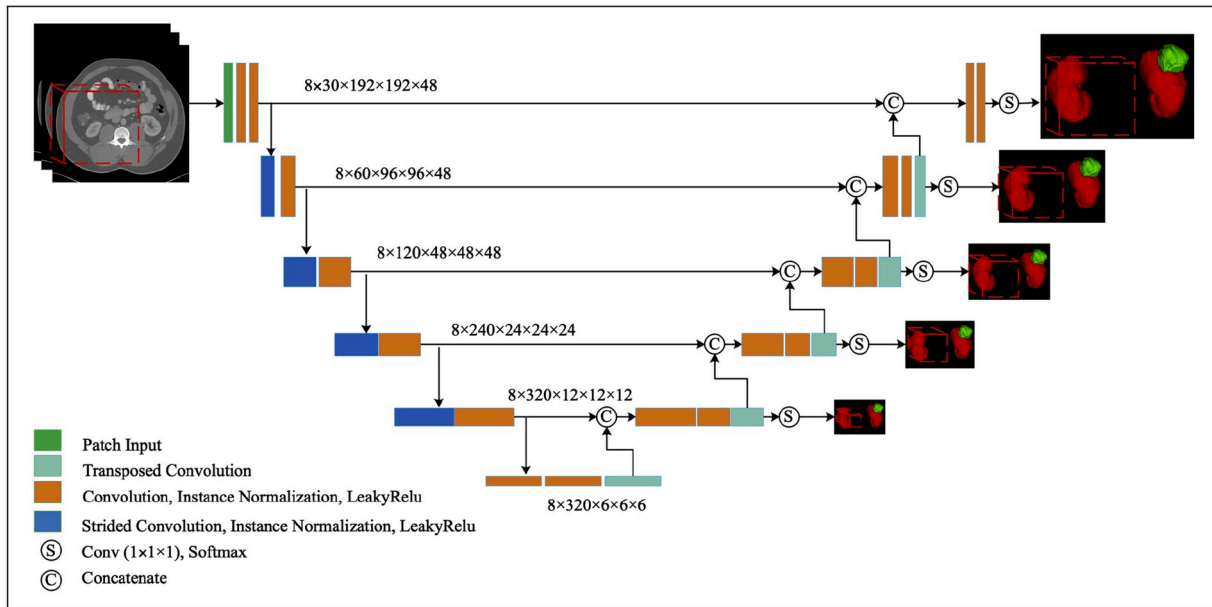


Fig. 3. The architecture of our proposed multi-scale supervised 3D U-Net. For simplicity, we use 2D icons instead of actual 3D ones, and it is best viewed in color. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

the network and processing the data.

2.1. Data preprocessing

Preprocessing raw CT images before feeding them to the network is an essential step to enable an efficient training. The first aspect to consider is the existence of unexpected materials that might appear inside patients' bodies. In particular, it is a well-known fact that metallic artifacts have a significant negative effect on the quality of CT images. The main problem with artifacts is when these create regions in the images with abnormal intensity values, either much higher or much lower than in pixels corresponding to organic tissue. Due to data-driven models that deep learning algorithms are built upon, the learning process can be significantly affected by outlier voxels corresponding to non-organic artifacts. To reduce the impact of non-organic artifacts, we perform a unified preprocessing of the complete dataset, both for training and testing data. In all images, we only consider valid the intensity range between the 0.5th and 99.5th percentiles, and clip outlier values accordingly. After preprocessing, the data is normalized with a normal foreground mean and standard deviation to improve the training of the three-dimensional network.

Another preprocessing technique necessary to appropriately train a three-dimensional network from the KiTS19 dataset is the unification of the voxel space across different image sets. This is necessary because even though the transverse plane is always formed by a constant number of pixels, the corresponding voxel size may change. Therefore, failing to unify the voxel space will lead to different data inputs representing different volumes. Such anisotropy of 3D data might diminish the advantage of using 3D convolution, and end up leading to a worse performance than 2D networks can achieve. Therefore, we chose to resample the original CT images into the same voxel space if they are not, even though the resampled images often end up being of different sizes.

2.2. Data augmentation

Annotating a medical image dataset can often be a lengthy and challenging task. This has so far limited the availability and size of labelled datasets. At the same time, the more complex deep learning methods become, the more data is needed to train the networks. This

becomes even more critical with three dimensional networks due to the inherent increase of parameters that comes with the extra dimension. Insufficient training data would lead to overfitting and devalue the advantage of deep learning. To solve this, a typical step is to utilize different augmentation techniques to increase the amount of available data while avoiding overfitting as much as possible. We conduct a variety of data augmentation techniques on our limited training data in order to obtain an enhanced performance of the trained network. These techniques are implemented based on the batchgenerator framework¹ including random rotations, random scaling, random elastic deformations, gamma correction augmentation and mirroring. Visualization of their effect is shown in Fig. 2.

2.3. Network architecture

The network architecture defined in this paper has been designed taking the nnU-Net neural network framework as a starting point [15]. The nnU-Net framework, unlike other recently published methods, does not add complex submodules and instead is mostly based on the original U-Net architecture. In addition to the default framework, we utilize multi-scale supervision to enhance the network's segmentation performance. The architecture of our proposed network is illustrated in Fig. 3, where two-dimensional images have been utilized for illustrative purposes even though the network's layers are three-dimensional. Following the main structure of 3D U-Net, the network implements decoder (left side) and encoder (right side) elements. The encoder layers are utilized to learn feature representations from the input data. The decoder is then employed for retrieving voxel locations and for determining their categories based on the semantic information extracted from encoder path.

We adopt strided convolution instead of common pooling operation to implement downsampling, which could avoid significant negative effect on the fusion of position information. Furthermore, we replace trilinear interpolation with transposed convolution to enable adaptive upsampling. In multiple previous works, normalization is usually deployed between any two layers to obtain fixed input distribution. However, since the batch size we used is limited by the GPU memory

¹ <https://github.com/MIC-DKFZ/batchgenerators/>.




Fig. 4. The principle of exponential logarithmic loss. Compared with linear loss, exponential logarithmic loss shows specific nonlinearity providing better indication for the learning process.

capacity, in this work we employ instance normalization instead of batch normalization. Batch normalization naturally suits better bigger batch sizes.

The depth of the network is often decided as a trade-off between the amount of semantic features and the spatial accuracy of the segmentation. The deeper the layer, the more semantic information can be extracted. At the same time, deeper layers tend to lose location information due to the continuous decrease of resolution of the feature maps. Therefore, we have set the architecture of our network to only 6 layers, including the bottleneck, so that the deepest feature map is not smaller than $8 \times 8 \times 8$. At the same time, to reduce the model volume, we set the basic kernel number to 30. In addition, a short connection between the encoder and decoder layers is built in the U-Net to enable the utilization at the decoder of more exact location information embedded in the encoder part.

In contrast with the original 3D U-Net architecture, we construct multi-scale supervision to encourage every layer in the decoder path to achieve exact location as well as semantic information. The motivation behind of our proposed multi-scale supervision is comprehensively depicted in the next subsection.

2.4. Multi-scale supervision

In traditional deep learning based segmentation, the model outputs the probability map from the top layer, which does not fully utilize the deeper feature maps even though they might contain more semantic information. As the top layer is upsampled from the deep layers, it is reasonable to in advance guarantee the deep layers with correct predictions. This would be expected to provide a better basis to up layers in turn. Therefore, we add labels with corresponding resolution to each layer in the decoder path, and compare them with the side outputs from deep layers. With the loss calculated from different layers, more effective gradient backpropagation can be obtained and thus increase the learning efficiency. The loss function for each iteration from multi-scale supervision is given by (1):

$$Loss_t = \sum_{l=1}^N Loss_l w_l \quad (1)$$

where $Loss_t$ denotes the total loss achieved by the multi-scale supervision, while $Loss_l$ refers to the loss calculated at the l -th layer. The total number of layers with N , excluding the bottleneck layer in the network. Finally, the weight of each individual layer's loss is given by w_l .

2.5. Loss function

Cross entropy (CE) is a widely used pixel-wise loss as it computes the entropy of prediction probability and ground truth based on each pixel. However, such property might lead to severe sample imbalance since the background occupies most of the CT images. Its definition is given by (2):

$$CE = - \sum_{c \in \text{classes}} w_c y_{true} \log(y_{pred}) \quad (2)$$

where y_{true} represents the ground truth and y_{pred} the predicted probability. The weight w_c for each of the classes is utilized to adjust the global cross entropy value.

Another key loss function is the Dice coefficient (Dice). The Dice differs from the cross entropy in that it is especially useful for tiny target segmentation as it calculates the similarity of predicted result and ground truth regardless of the target's relative size. Its definition for a class is given by (3):

$$Dice = \frac{2|U_{pred} \cap U_{true}|}{|U_{pred}| + |U_{true}|} \quad (3)$$

where U_{pred} and U_{true} represent the segmentation results set (predictions) and the ground truth set, respectively, and $|U_i|$ denotes the cardinality of set U_i . The Dice is proportional to intersection of both sets, and therefore is affected by both false positives and false negatives. The higher the value of Dice, the better the segmentation effect is. Additionally, we

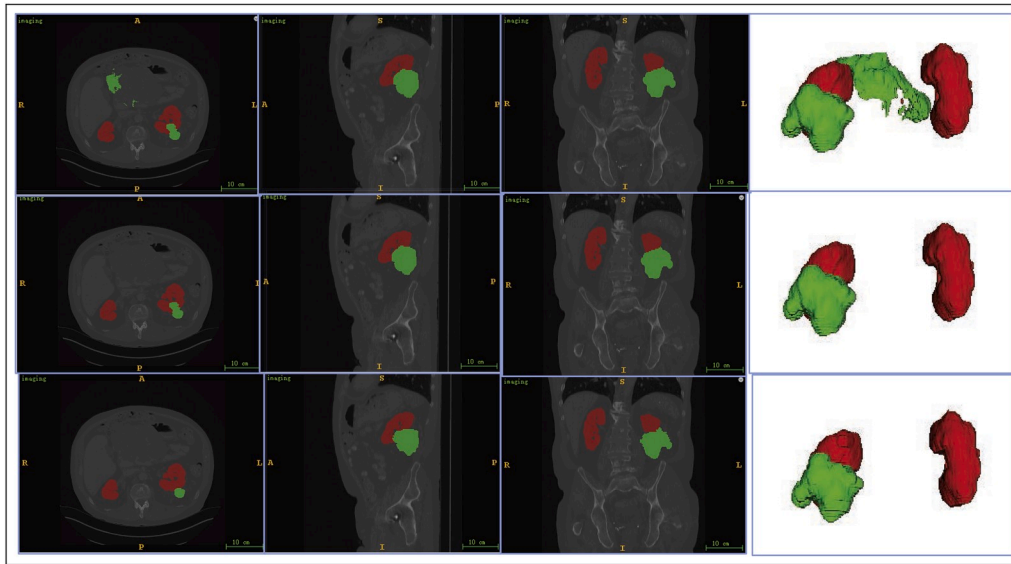


Fig. 5. The effect of our postprocessing method. The top row is original output from the network; the middle is after post processing and the bottom is the ground truth.

utilize the Soft Dice (SD) coefficient, as it is beneficial to directly utilize the prediction probabilities and not the binary mask that the softmax layer generates. This allows us to have a cost function that is more sensitive to the learning process, and adjust the network weights more efficiently during the learning process. In order to have a function to minimize, $1 - Dice$ is often utilized as a cost function. Thus, in practice, (3) turns into (4) for each class mask:

$$SD_{cost} = 1 - \frac{2 \sum_{pixels} y_{pred} y_{true}}{\sum_{pixels} y_{pred}^2 + \sum_{pixels} y_{true}^2} \quad (4)$$

The nnU-Net implementation utilizes both cross-entropy and $1 - Dice$ for training. However, in kidney and kidney tumor segmentation additional challenges arise leading us to choose a different cost function. First, the number tumor samples in the CT images is significantly smaller than the number of background and kidney samples. Second, the morphological heterogeneity of tumor voxels is significantly larger than that of kidney voxels. This imbalance of both samples and difficulty is a generic problem in medical image segmentation and would cause tendency of the network to incorrectly categorize tumor voxels.

Therefore, in this paper, we modify the nnU-Net to perform exponential logarithmic loss on our Soft Dice loss to alleviate this imbalance. Exponential logarithmic loss shows specific nonlinear property as the objective value increases, as shown in Fig. 4. In the initial training phase, the loss has a dramatic variety which could better indicate the correct learning direction. In the middle phase, it shows stable decrease as the linear function. As the learning goes, the exponential logarithmic keeps higher loss than the linear loss. Until the objective value is rather close to one it shows a fast drop. The details of the modification appear on the public GitHub repository where we have uploaded the code utilized in this paper. The cost function utilized for minimization during the training is given by (5):

$$SD_{ell} = (-\log SD_{kidney})^{0.3} \times 0.4 + (-\log SD_{tumor})^{0.3} \times 0.6 \quad (5)$$

where SD_{ell} represents the Soft Dice modified by exponential logarithmic loss. SD_{kidney} and SD_{tumor} denote the original Soft Dice calculated on kidney and tumor respectively:

$$SD_{kidney} = \frac{2 \sum_{kidney\ pixels} y_{pred} y_{true}}{\sum_{kidney\ pixels} y_{pred}^2 + \sum_{kidney\ pixels} y_{true}^2}, \quad SD_{tumor} = \frac{2 \sum_{tumor\ pixels} y_{pred} y_{true}}{\sum_{tumor\ pixels} y_{pred}^2 + \sum_{tumor\ pixels} y_{true}^2} \quad (6)$$

This operation of nonlinearity could enable the network to achieve higher loss when the samples are very difficult to recognize, i.e. the prediction result is very bad, and only when the prediction is good above a certain threshold will the loss decline dramatically. In this way, the network with exponential logarithmic loss will potentially obtain more efficient gradient updates than when using linear loss functions. In addition, to further induce the network into increasing the significance of tumor samples during training, we attribute different multipliers, 0.4 and 0.6, as weights to kidney and tumor, respectively.

Finally, we combine the Soft Dice with exponential logarithmic loss and CE as the loss function for each layer, shown in (7):

$$Loss_{layer} = SD_{ell} + CE \quad (7)$$

where the $Loss_{layer}$ denotes the loss we obtain from each layer. With multi-scale supervision, we further assign different weights on different layers. As a result, from the top down, the total 5 layers excluding the bottleneck get 0.4, 0.2, 0.2, 0.1, 0.1 respectively and we attribute 0.28, 0.28 and 0.44 for the CE weights of background, kidney and tumor, respectively, to further emphasize tumor samples. Our method brings such hyperparameters combining the multi-scale supervision with exponential logarithmic loss, which can be further optimized to extract more potential out of the 3D U-Net. While there is plenty of room for optimization of these parameters, our experiments already show the benefits and an enhanced segmentation performance.

2.6. Inference and postprocessing

Because network training and inference are conducted patch-wise, the accuracy of the border of patches is decreased compared with the patch center. Therefore, we adopt overlap prediction and weigh more on the center values when aggregating predictions across patches. Patches are chosen to overlap by one half of the patch size. Furthermore, we employ test-time data augmentation by mirroring testing patches along all valid axes to aggregate more predictions as well as adding Gaussian noise. Thus, for every voxel there are multiple predictions aggregated to

Table 1
Properties of the KiTS19 dataset.

Property	Value
Number of Patients	210
Modality	CT (late-arterial phase)
Training dataset size	134
Validation dataset size	34
Test dataset size	42
Min Patient Size in Voxels	[434, 434, 69]
Max Patient Size in Voxels	[639, 639, 182]
Median Patient Size in Voxels	[523, 523, 116]
Affine	$\begin{pmatrix} 0 & 0 & -0.7816 & 0 \\ 0 & -0.7816 & 0 & 0 \\ -3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$

post-processing effectively matches the ground truth sample. It should be accounted that if there is only one kidney, the patient might have previously received a nephrectomy.

3. Experimental results

In order to evaluate the proposed three-dimensional network architecture, and to validate its usability with real data, we conduct experiments on the public KiTS19 dataset. We separate the KiTS19 dataset into the training, validation and test datasets, which we utilize to train and evaluate our model's performance with both visualization metrics and quantification metrics.

3.1. KiTS19 dataset




Fig. 6. Loss and Dice evolution during the network training. The red and blue lines represent the validation and training loss, respectively. The green line represents the average Dice of kidneys and kidney tumor. The total training time was approximately five days. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

determine its optimal value. In the center of patient data, up to 64 predictions from several overlaps and mirroring are aggregated up.

After the inference by the network, we utilize the basic knowledge to further improve the performance: most humans have two kidneys, and the kidney tumor should be attached on or embedded in kidneys. Even though this is very elementary information, it can be employed in postprocessing to remove disconnected voxels that are detected as false positive. False positive kidney voxels are removed from the output after either one or two kidney components have been found. Thus, all tumor components attached to a kidney are maintained as valid segmentation result while others are removed from the output as well. The 3D connected components are detected following a similar procedure to Ref. [22], and utilizing the *connected-components-3d*² Python library [9].

Fig. 5 shows how our postprocessing method increases the quality of the segmentation. In the top most figure, we have two kidneys and a disconnected tumor. The segmentation keeps the top two biggest kidney components as shown in the middle of the figure. The result after the

The KiTS19 dataset contains volumetric CT scans from 210 patients. These scans are all preoperative abdominal CT imaging in the late-arterial phase, with unambiguous definition of kidney tumor voxels in the ground truth images. The images are provided in Neuroimaging Informatics Technology Initiative (Nifti) format. Scans of different patients have different properties. Therefore, there is a heterogeneity in the raw data, including the voxel size along the three plans and their affine. As diverse voxel spacings may have a significantly negative impact on the learning process of deep neural networks, we use instead the interpolated dataset in KiTS19, which interpolates the original dataset to achieve the same affine for every patient.

The statistic properties of the dataset are shown in Table 1. We randomly select 20% of the patients as the independent test dataset. Among the rest, we partition the records into validation dataset (20% of scans) and training dataset (80% of scans).

3.2. Implementation details

We utilize Adam as the network's optimizer function and set the initial learning rate to be 3×10^{-4} . In addition, we conduct an adaptive

² <https://github.com/seung-lab/connected-components-3d>.

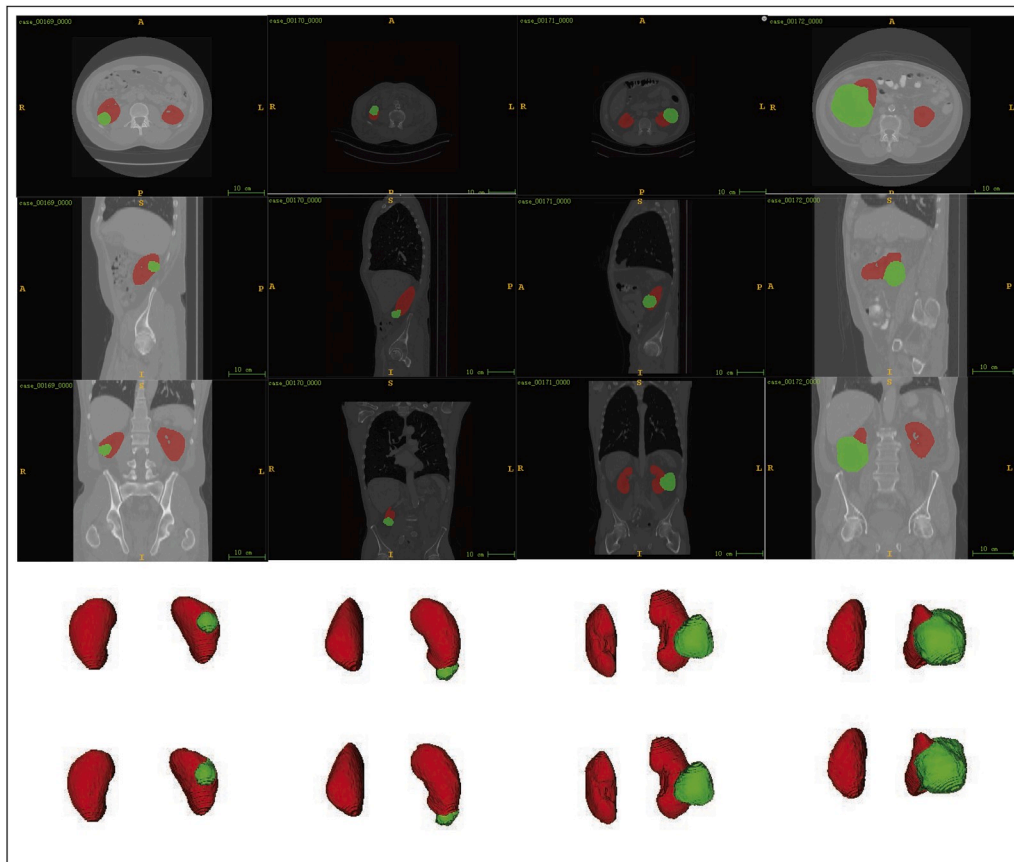


Fig. 7. Examples of segmentation results. Each column denotes one patient, and from top down, they are scan images in transverse plane, sagittal plane, coronal plane, 3D reconstruction of our predictions and the ground truth. The red mask represents kidney area, while the green mask represents tumor area. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

adjustment strategy for the learning rate during the training process. This results in the learning rate dropping by the factor of 0.2 whenever the training loss is not improved over 30 epochs. Similarly, we consider that the training is over whenever no improvements in the loss are identified over 50 epochs. The complete process is implemented in Python utilizing the PyTorch framework.

In the experiments, the training is carried out with two Nvidia Tesla 32 GB GPUs. Owing to the limited amount of the GPU memory, we adopt

a patch size of $192 \times 192 \times 48$ and set the batch size to eight. Since the training is patch-based, the patch is randomly sampled from the data loader and we each epoch is set to 250 iterations. This translates into each epoch effectively selecting 250×8 patches from the training data.

3.3. Experiment results

The training time until convergence was achieved with our network

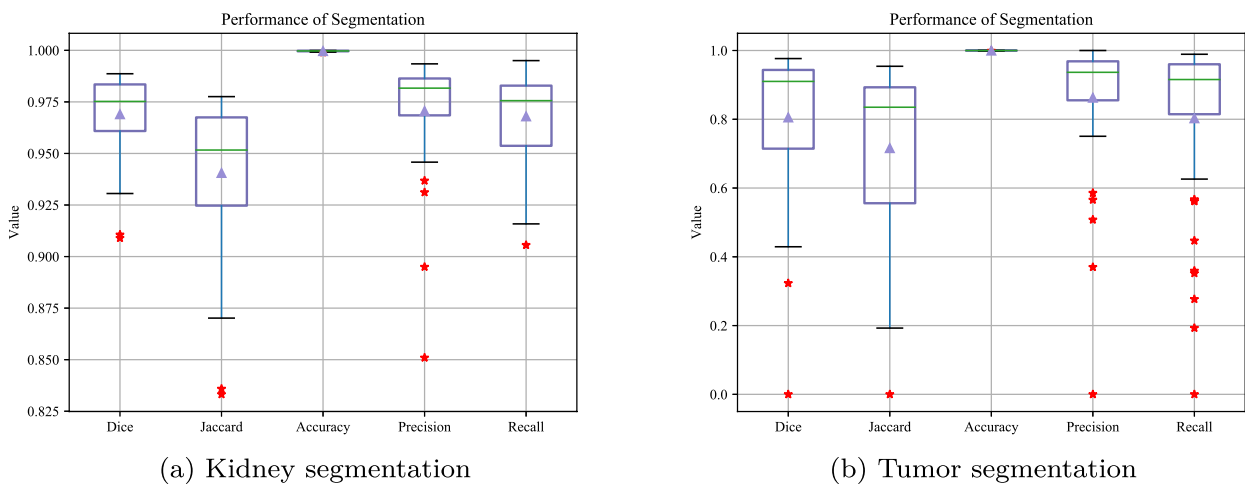


Fig. 8. Quantitative metrics.

Table 2

Comparison between the proposed MSS U-Net and our implementation of the classic 3D U-Net (kidney), both on the KiTS19 dataset.

Metric	MSS U-Net	Classic 3D U-Net
Dice	0.969	0.962
Jaccard	0.941	0.930
Accuracy	0.999	0.999
Precision	0.971	0.961
Recall	0.968	0.965
Hausdorff (mm)	19.188	38.945

Table 3

Comparison between the proposed MSS U-Net and classic 3D U-Net (tumor).

Metric	MSS U-Net	Classic 3D U-Net
Dice	0.805	0.781
Jaccard	0.716	0.699
Accuracy	0.999	0.999
Precision	0.863	0.841
Recall	0.802	0.810
Hausdorff (mm)	33.469	50.808

Table 4

Comparison of Dice coefficient of the proposed MSS U-Net and state-of-the-art methods from Ref. [28] (PSPNET, 3D U-NET and FCN-PPM). The datasets utilized in Ref. [28] and our paper (MSS U-NET) are different.

Method	Kidney	Tumor
2D PSPNET (RoI)	0.902	0.638
3D U-NET (RoI)	0.927	0.751
3D FCN PPM (RoI)	0.931	0.802
MSS U-NET (Raw Images)	0.969	0.805

architecture was approximately five days. The loss variation is shown in Fig. 6, from which we can observe that the loss decreases stably as the number of training epoch increases.

After the model is trained and convergence achieved, we test it using the independent test dataset containing 42 patient scans. Several selected samples of output segmentation results are shown in Fig. 7. Even though the location, intensity and texture of kidneys and kidney tumor can differ significantly, the predicted regions are in good agreement with the ground truth.

We evaluate all of the 42 test patient scans using the six quantitative metrics to evaluate our method objectively and comprehensively; Dice, Jaccard, Accuracy, Precision, Recall and Hausdorff are computed for both kidney and kidney tumor segmentation. To clearly observe the distribution of all the test patients, the indicators of Dice, Jaccard, Accuracy, Precision and Recall are gathered into two box plots, shown in Fig. 8a and Fig. 8b for kidney and kidney tumor segmentation, respectively.

Our method is further elaborated from the basic 3D U-Net. We do not consider nor compare our approach to other complicated architecture modifications, as they are often only effective for specific cases or metrics. Therefore, in order to investigate the effectiveness of our strategies, we compare the performance of the basic 3D U-Net with our multi-scale supervised 3D U-Net. The implementation of the 3D U-Net is identical to our multi-scale supervised U-Net except for the three strategies defined in this paper. The comparison results are listed in Table 2 and Table 3. These two tables present the average values of the six indicators and prove the improvement after using our three-fold enhancing strategies. The most notable differences are in terms of tumor segmentation.

After comparing with the basic U-Net architecture, we compare our approach with three recent methods listed in Table 4. In this comparison, it should be noted that our proposed method processes raw-size CT images, which adds significant complexity to the segmentation model

Table 5

Dice of the proposed method and other algorithms in KiTS19 challenge. These results are obtained from the competition's test dataset. This test dataset is public, but the ground truth labels were not made public.

Method	Kidney	Tumor
Fabian et al. [14], 1st place	0.974	0.851
Xiaoshuai et al. [12], 2nd place	0.967	0.845
Guangrui et al. [22], 3rd place	0.973	0.832
Andriy et al. [23], 9th place	0.974	0.810
MSS U-Net, 7th place	0.974	0.818

when compared to the other three based on smaller ROIs. Nevertheless, our method still outperforms them in multiple metrics regarding both kidney and kidney tumor segmentation.

In addition, to compare our methods with other current state-of-the-art architectures and approaches, we participated in the KiTS19 challenge. During the challenge, we were able to perform a broader evaluation with the other state-of-the-art methods. The proposed network architecture, MSS U-Net, ranked in 7th position out of 106 teams [13] achieving a kidney and kidney tumor Dice of 0.974 and 0.818, respectively, measured using the data of the 90 test patients. This demonstrates both the applicability of the approach described in this paper for real data and its improved performance compared to previous methods relying only on the 3D-UNet architecture for medical image segmentation.

We list several top winners of KiTS19 challenge in Table 5. Fabian et al. [14], the first-ranked solution, proposed a residual 3D U-Net to enhance the segmentation performance. In this case, the authors modified part of the training data in order to gain a unique advantage. Xiaoshuai et al. [12] and Guangrui et al. [22] constructed their methods using 3-stage and 2-stage segmentation respectively, which followed a different strategy compared to our motivation for having an end-to-end method. Finally, Andriy et al. [23] followed a more similar strategy to ours because they also focused on how to better train a basic 3D U-Net and employed the fashionable boundary-aware loss [10]. However, they used a much bigger input ($176 \times 176 \times 176$). In general terms, we believe that our method can be further optimized by conducting more adaptive experiments due to the larger number of hyperparameters directed to training the 3D U-Net more efficiently.

It is worth noting that, due to the limited size of the dataset, the results could vary significantly if only one tumor was not detected. This is particularly evident with particularly small tumors that can pass undetected simply because of sampling or data handling issues. Nonetheless, this is an important aspect that must be considered because detecting small tumors or lesions can be key in early disease detection. Thus, we believe that formal solutions need to be proposed to tackle this specific problem of small tumor segmentation.

4. Discussion

In recent years, deep learning based methods have accounted for the largest fraction of research papers in the field of medical image segmentation. A wide variety of networks with many new architectures have been proposed, with innovative properties and significantly superior results in multiple aspects when compared to more traditional methods. Nonetheless, among all these, and to the best of our knowledge, the original U-Net architecture is still able to achieve results comparable to the state-of-the-art and even outperform more recent architectures in certain aspects related to medical image segmentation [15]. This has been further exemplified in this paper, as we have demonstrated the capabilities of the original architecture if more efficient training techniques are introduced as proved by our experimental results.

Based on the 3D U-Net architecture, we have designed and trained a six-layer network and proposed three effective strategies including both




Fig. 9. Illustration of the worst-case segmentation from the proposed network. The top row shows the predicted segmentation while the bottom row shows the ground truth. The tumor, in green, is not identified by our network, which is only able to segment the kidney tissue instead, in red. The small size of the tumor in this case may have played a significant role in the segmentation failure. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

training techniques and data augmentation. First, we have employed multi-scale supervision to increase the probability of the network predicting low-resolution labels correctly from deeper layers. This concept can be better understood with an analogy of human behaviour: the action of first zooming out an image to label or identify coarse contours and then zoom into precisely label the image at a finer level. Second, in order to alleviate the inherent imbalance in samples and segmentation difficulty across organs, we have introduced the use of exponential logarithmic loss to induce the network into paying more attention to tumor samples and more difficult samples. Third, we have designed a connected-component based post processing method to remove the clearly mistaken voxels that have been identified by the network.

The experiments that we have carried out and the comparisons with existing methods have shown the advantages of the proposed architecture and the effectiveness of the enhancement strategies over the original 3D U-Net. Nevertheless, several aspects remain challenging and require further investigation. In particular, from the segmentation statistics of the test patients, we have found that two patients had a low tumor Dice with one of them being as low as zero. The data corresponding to the latter patient are shown in Fig. 9. In the figure, we can see that the prediction mask does not map the tumor, a consistent result with the Dice coefficient. The cause of this low Dice might be the particularly small size of the tumor in the CT images. We attribute this phenomenon to the fixed receptive field of classic convolution. Therefore, we will consider in our future work to employ deformable convolution as a potential solution to this problem [6].

Even though we have achieved rather good performance on the kidney segmentation, it should be noted that all the experiments are performed on the same dataset, in which the training data are randomly split and would have identical distribution with test data. However, when applied to real clinical settings, the trained model has to adapt to CT images from different vendors and settings. This has the potential to bring different distribution and hence degrades its performance. This is a general problem in machine learning, and we hope to employ data from multiple vendors and centers to improve its generalization in the future [26].

In summary, the proposed architecture and chosen methodology in this paper are based on the assumption that the basic 3D U-Net architecture is capable of extracting sufficient features for segmentation. Therefore, we have directed our efforts towards the training procedures while discarding the complex architecture modifications with marginal effectiveness. At the time of finalizing the experiments reported in this paper, a modified U-Net (mU-Net) has been proposed by Seo et al. [25]. In their paper, the authors propose the utilization of residual path to the skip connections of the U-Net for the segmentation of livers and liver tumors. This paper shares a similar motivation with ours, aiming and innovative ways of mining more effective information from low-resolution features for accurate segmentation of medical images. The main difference from the architectural point of view is that Seo et al.

introduce additional blocks, while we focus on effective but simpler strategies to achieve the same goal.

5. Conclusion

In this paper, we have proposed an end-to-end multi-scale supervised 3D U-Net to simultaneously segment kidneys and kidney tumors from raw-scale computed tomography images. Extending the original 3D U-Net architecture, we have combined a multi-scale supervision approach with exponential logarithmic loss. This has enabled further optimization of the U-Net architecture, extending its possibilities, and hence obtaining better performance. Compared with a current trend in deep neural networks with complex architectures and multiple different sub-modules, we have taken a more generalized approach yet obtained results comparable to the state of the art. A simpler architecture has the advantage of higher reproducibility and wider generalization of results, in contrast with the potentially highly inflated models and poor reproducibility of more complex architectures.

In general, we have directed our efforts towards a more efficient training of the original 3D U-Net architecture by incorporating the multi-scale supervision and the exponential logarithmic loss. We have demonstrated the advantages of this approach with our experiments and comparisons with the state-of-the-art. While our architecture can be outperformed by others in specific metrics on the KiTS19 dataset, we have argued that having a simpler architecture still leaves room for further optimization and discussed the advantages from the point of view of applicability and extendability.

Finally, the code leading to this work has been made public through a GitHub repository,³ with the code that has been used on the KiTS19 dataset. In the future work, we expect to extend the application of our architecture towards segmentation of other organs and modalities, such as magnetic resonance imaging (MRI) or positron-emission tomography (PET).

Ethical statement

Hereby we certify that this manuscript is original and has not been published and will not be submitted elsewhere for publication while being considered by Informatics in Medicine Unlocked. All analyses were based on public datasets and previous published studies, thus no additional ethical approval and patient consent are required.

Declaration of competing interest

The author(s) declare no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

³ <https://github.com/LINGYUNFDU/MSSU-Net>.

Acknowledgement

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- [1] Abraham N, Khan NM. A novel focal tvsky loss function with improved attention u-net for lesion segmentation. In: 2019 IEEE 16th international symposium on biomedical imaging (ISBI 2019). IEEE; 2019. p. 683–7.
- [2] Badrinarayanan V, Kendall A, Cipolla R. Segnet: a deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans Pattern Anal Mach Intell* 2017;39(12):2481–95.
- [3] Cairns P. Renal cell carcinoma. *Canc Biomarkers* 2011;9(1–6):461–73.
- [4] Christ PF, Ettliger F, Grün F, Elshaera MEA, Lipkova J, Schlecht S, Ahmaddy F, Tatavarty S, Bickel M, Bilic P, et al. Automatic liver and tumor segmentation of ct and mri volumes using cascaded fully convolutional neural networks. 2017. arXiv preprint arXiv:1702.05970.
- [5] Çiçek Ö, Abdulkadir A, Lienkamp SS, Brox T, Ronneberger O. 3d u-net: learning dense volumetric segmentation from sparse annotation. In: International conference on medical image computing and computer-assisted intervention. Springer; 2016. p. 424–32.
- [6] Dai J, Qi H, Xiong Y, Li Y, Zhang G, Hu H, Wei Y. Deformable convolutional networks. In: Proceedings of the IEEE international conference on computer vision; 2017. p. 764–73.
- [7] Fasihi MS, Mikhael WB. Overview of current biomedical image segmentation methods. In: 2016 international conference on computational science and computational intelligence (CSCI). IEEE; 2016. p. 803–8.
- [8] Gillies RJ, Kinahan PE, Hricak H. Radiomics: images are more than pictures, they are data. *Radiology* 2016;278(2):563–77.
- [9] Grana C, Borghesani D, Cucchiara R. Optimized block-based connected components labeling with decision trees. *IEEE Trans Image Process* 2010;19(6): 1596–609.
- [10] Hatamizadeh A, Terzopoulos D, Myronenko A. Boundary aware networks for medical image segmentation. 2019. arXiv preprint arXiv:1908.08071.
- [11] He K, Gkioxari G, Dollár P, Girshick R. Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision; 2017. p. 2961–9.
- [12] Heller N, Isensee F, Maier-Hein KH, Hou X, Xie C, Li F, Nan Y, Mu G, Lin Z, Han M, et al. The state of the art in kidney and kidney tumor segmentation in contrast-enhanced ct imaging: results of the kits19 challenge. 2019. arXiv preprint arXiv: 1912.01054.
- [13] Heller N, Sathianathan N, Kalapara A, Walczak E, Moore K, Kaluzniak H, Rosenberg J, Blake P, Rengel Z, Oestreich M, et al. The kits19 challenge data: 300 kidney tumor cases with clinical context, ct semantic segmentations, and surgical outcomes. 2019. arXiv preprint arXiv:1904.00445.
- [14] Isensee F, Maier-Hein KH. An attempt at beating the 3d u-net. 2019. arXiv preprint arXiv:1908.02182.
- [15] Isensee F, Petersen J, Klein A, Zimmerer D, Jaeger PF, Kohl S, Wasserthal J, Koehler G, Norajitra T, Wirkert S, et al. nnu-net: self-adapting framework for u-net-based medical image segmentation. 2018. arXiv preprint arXiv:1809.10486.
- [16] Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems; 2012. p. 1097–105.
- [17] Lin T-Y, Dollár P, Girshick R, He K, Hariharan B, Belongie S. Feature pyramid networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition; 2017. p. 2117–25.
- [18] Liu Z, Song Y-Q, Sheng VS, Wang L, Jiang R, Zhang X, Yuan D. Liver ct sequence segmentation based with improved u-net and graph cut. *Expert Syst Appl* 2019; 126:54–63.
- [19] Long J, Shelhamer E, Darrell T. Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition; 2015. p. 3431–40.
- [20] Milletari F, Navab N, Ahmadi S-A. V-net: fully convolutional neural networks for volumetric medical image segmentation. In: 2016 fourth international conference on 3D vision (3DV). IEEE; 2016. p. 565–71.
- [21] Minaee S, Boykov Y, Porikli F, Plaza A, Kehtarnavaz N, Terzopoulos D. Image segmentation using deep learning: a survey. 2020. arXiv preprint arXiv: 2001.05566.
- [22] Mu G, Lin Z, Han M, Yao G, Gao Y. Segmentation of kidney tumor by multi-resolution vb-nets. 2019.
- [23] Myronenko A, Hatamizadeh A. 3d kidneys and kidney tumor semantic segmentation using boundary-aware networks. 2019. arXiv preprint arXiv: 1909.06684.
- [24] Oktay O, Schlemper J, Folgoc LL, Lee M, Heinrich M, Misawa K, Mori K, McDonagh S, Hammerla NY, Kainz B, et al. Attention u-net: learning where to look for the pancreas. 2018. arXiv preprint arXiv:1804.03999.
- [25] Seo H, Huang C, Bassenne M, Xiao R, Xing L. Modified u-net (mu-net) with incorporation of object-dependent high level features for improved liver and liver-tumor segmentation in ct images. In: IEEE transactions on medical imaging; 2019.
- [26] Tao Q, Yan W, Wang Y, Paiman EH, Shamonin DP, Garg P, Plein S, Huang L, Xia L, Sramko M, et al. Deep learning-based method for fully automatic quantification of left ventricle function from cine mr images: a multivendor, multicenter study. *Radiology* 2019;290(1):81–8.
- [27] Wong KC, Moradi M, Tang H, Syeda-Mahmood T. 3d segmentation with exponential logarithmic loss for highly unbalanced object sizes. In: International conference on medical image computing and computer-assisted intervention. Springer; 2018. p. 612–9.
- [28] Yang G, Li G, Pan T, Kong Y, Wu J, Shu H, Luo L, Dillenseger J-L, Coatrieux J-L, Tang L, et al. Automatic segmentation of kidney and renal tumor in ct images based on 3d fully convolutional neural network with pyramid pooling module. In: 2018 24th international conference on pattern recognition (ICPR). IEEE; 2018. p. 3790–5.
- [29] Yu Q, Shi Y, Sun J, Gao Y, Zhu J, Dai Y. Crossbar-net: a novel convolutional neural network for kidney tumor segmentation in ct images. *IEEE Trans Image Process* 2019;28(8):4060–74.
- [30] Zhao H, Shi J, Qi X, Wang X, Jia J. Pyramid scene parsing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition; 2017. p. 2881–90.