

Received 22 February 2023, accepted 29 March 2023, date of publication 5 April 2023, date of current version 12 April 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3264582

RESEARCH ARTICLE

Improved Abdominal Multi-Organ Segmentation via 3D Boundary-Constrained Deep Neural Networks

SAMRA IRSHAD^{1,2}, DOUGLAS P. S. GOMES³, AND SEONG TAE KIM¹, (Member, IEEE)

¹Department of Computer Science and Engineering, Kyung Hee University, Yongin-si, Gyeonggi-do 17104, South Korea

²Faculty of Engineering, Science and Technology, Swinburne University of Technology, Hawthorn, VIC 3122, Australia

³College of Engineering and Science, Victoria University, Melbourne, VIC 3008, Australia

Corresponding author: Seong Tae Kim (st.kim@khu.ac.kr)

This work was supported in part by the National Research Foundation of Korea (NRF) Grant funded by the Korea Government (MSIT) under Grant 2021R1G1A1094990, in part by the Institute of Information and Communications Technology Planning and Evaluation (IITP) Grant funded by the Korea Government (MSIT) under Grant 2022-0-00078 (Explainable Logical Reasoning for Medical Knowledge Generation), Grant 2021-0-02068 (Artificial Intelligence Innovation Hub), and Grant RS-2022-00155911 (Artificial Intelligence Convergence Innovation Human Resources Development (Kyung Hee University)).

ABSTRACT Quantitative assessment of the abdominal region from CT scans requires the accurate delineation of abdominal organs. Therefore, automatic abdominal image segmentation has been the subject of intensive research for the past two decades. Recently, deep learning-based methods have resulted in state-of-the-art performance for the 3D abdominal CT segmentation. However, the complex characterization of abdominal organs with weak boundaries prevents the deep learning methods from accurate segmentation. Specifically, the voxels on the boundary of organs are more vulnerable to misprediction due to the highly-varying intensities. This paper proposes a method for improved abdominal image segmentation by leveraging organ-boundary prediction as a complementary task. We train 3D encoder-decoder networks to simultaneously segment the abdominal organs and their boundaries via multi-task learning. We explore two network topologies based on the extent of weights shared between the two tasks within a unified multi-task framework. In the first topology, the whole-organ prediction task and the boundary detection task share all the layers in the network except for the last task-specific layers. The second topology employs a single shared encoder but two separate task-specific decoders. The effectiveness of utilizing the organs' boundary information for abdominal multi-organ segmentation is evaluated on two publically available abdominal CT datasets: Pancreas-CT and the BTCV dataset. The improvements shown in segmentation results reveal the advantage of the multi-task training that forces the network to pay attention to ambiguous boundaries of organs. A maximum relative improvement of 3.5% and 3.6% is observed in Mean Dice Score for Pancreas-CT and BTCV datasets, respectively.

INDEX TERMS Abdominal multi-organ segmentation, fully convolutional neural networks, boundary-constrained segmentation, multi-task learning.

I. INTRODUCTION

Multi-organ segmentation on abdominal Computed Tomography (CT) scans is an essential prerequisite for computer-assisted surgery and organ transplantation [1], [2].

The associate editor coordinating the review of this manuscript and approving it for publication was Byung-Gyu Kim.

Particularly, quantitative assessment of abdominal regions enables accurate organ dose calculation, required in numerous radiotherapy treatment options. Erroneous delineation of abdominal organs prevents harnessing the benefits of radiotherapeutic advancements. In clinical practice, physicians delineate abdominal organs using manual segmentation tools, which are time-consuming, observer-dependent, and




FIGURE 1. Exemplary 2D abdominal CT image showing the visual characteristics of organs. (a) 2D abdominal image, (b) Abdominal organs annotated on CT image: pancreas (orange), spleen (purple), liver (yellow), stomach (blue), gallbladder (green), (c) 3D multi-organ voxel map.

error-prone. With the increased use of imaging facilities and production of a large number of abdominal CT scans, the utilization of automated, robust, and efficient organ-delineation tools has become compulsory [2], [3], [4].

Abdominal CT scans often present weak inter-organ boundaries characterized by regions of similar voxel intensities. Such appearances are usually caused by the representation of abdominal soft tissues in a narrow band of Hounsfield (HU) values. Another factor that enhances the already complex representation of abdominal organs is the existence of artifacts occurring due to blood flow, respiratory, and cardiac motion. Accurate delineation of abdominal organs with unclear boundaries and complex geometrical shapes is one of the ongoing challenges that hurdles the abdominal-related clinical diagnosis.

Earlier methods proposed for the abdominal multi-organ segmentation mainly were based on multi-atlas [5], [6] or statistical models [7], [8]. Some methods also made use of handcrafted or learned features to segment abdominal organs [9], [10]. However, the recent Fully Convolutional Network (FCN) based approaches have presented better results due to the improved organ representation learning [2], [11]. Being able to preserve the image structure and provision of efficient learning as well as inference, FCN-based methods are currently considered state-of-the-art for abdominal multi-organ segmentation [2], [12], [13].

Existing FCN-based methods for abdominal multi-organ segmentation employ either 2D or 3D convolutional architectures [12], [13]. 2D methods process the CT scans in a slice-by-slice fashion and predict the organ labels on individual slices [13]. Despite being memory- and parameter-efficient, 2D methods are unable to make full use of 3D contextual information [2]. 3D methods make use of rich volumetric context by processing the whole CT volume and generating voxel-maps in a single forward propagation pass, leading to better abdominal CT segmentation performance than 2D approaches [15], [16].

The existing 3D methods have primarily focused on designing better architectures for improved abdominal multi-organ representation learning [2], [12]. However, they treat all the anatomical parts within a single organ equally

since they solely rely on voxel-level information and do not specifically focus on improving the segmentation of voxels in vulnerable regions/parts of organs. As an example, we highlight some of the important characteristics of abdominal organs in Fig. 1. From Figures 1a and 1b, it can be noticed that the adjacent organs have weak contours which sometimes touch each other. As an example, observe the low-contrasted and touching boundaries between stomach (blue) and pancreas (orange). Moreover, 3D multi-organ visualization in Fig. 1c shows that the adjacent positioning of organs in the abdominal cavity aggravates the complex spatial relationship among the organs.

The boundaries of anatomical regions in medical scans serve as an important cue for facilitating manual and automated delineation. Numerous existing deep learning-based studies leveraged learning of features corresponding to boundary of regions for improved medical image segmentation via multitask learning paradigm [17], [18], [19], [20], [21]. In recent years, deep multitask learning paradigm has been widely used due to its potential to solve multiple tasks in one forward propagation and ability to learn better representations because of the multiple supervisory signals [22], [23]. In this paper, we propose to improve the segmentation of abdominal organs on CT scans by enhancing the segmentation of boundary of organs. Particularly, we train the 3D deep learning networks to simultaneously predict the boundary and the entire region of organs. The inclusion of boundary information is motivated by the fact that the voxels on the boundary of organs are more vulnerable to misprediction because of their ambiguous appearance and complex relationship with adjacent organs. Specifically, our work makes the following contributions:

- (i) We develop a trainable 3D multi-task learning framework that simultaneously predicts the voxel-labels of abdominal organs and their corresponding boundaries. By integrating the boundary features, our proposed boundary-constrained 3D deep learning framework focuses on the accurate prediction of the edges of organs in addition to whole organs.
- (ii) Instead of relying on a single network topology, we explore and compare two network topologies for

conducting multi-task learning. In the first topology, the whole encoder-decoder network is shared with separate task-specific prediction layers at the end for predicting boundaries and entire organs' maps. In the second topology, an encoder is shared with separate task-specific decoders for decoding the features, jointly learned by the shared encoder to predict the boundary and organ probability maps. With an extensive comparison, we reveal that integration of boundary features invariably improves the multi-organ segmentation performance, independent of the multi-task network design.

- (iii) We utilize three state-of-the-art 3D encoder-decoder architectures, i.e., UNet [24], UNet₊₊ [25], and Attention-UNet [26] as baseline networks for evaluating the effect of incorporating boundary information. We modify each baseline architecture according to our proposed multi-task topologies. We demonstrate significant performance improvements with a negligible increase in trainable parameters.
- (iv) We validate the performance of baseline and counterpart boundary-constrained models on two publically available datasets (Pancreas-CT [27] and BTCV [28]) using Dice Score, Average Hausdorff Distance, Recall, and Precision. Furthermore, we conduct additional experiments to evaluate the improvement in the segmentation of regions around the boundaries. The results show that the boundary-constrained networks learn feature representations that focus on the accurate organs segmentation and the challenging parts around the border of the organs. All source codes are publically available on <https://github.com/samra-irshad/3d-boundary-constrained-networks>.

The rest of the article is organized as follows. In Section II, we review the existing methods for abdominal multi-organ segmentation. Section III describes our framework for incorporating the boundary information into the 3D fully convolutional networks, including the multi-task loss function and the details of boundary-constrained network topologies. Next, we describe the dataset specifications and implementation details in Section IV. We then present the experimental results, comparisons with existing single-task approaches, and in-depth performance analysis of boundary-constrained models in section V. Finally, we discuss the important highlights and some directions for future work in Section VI and present the conclusion in Section VII.

II. RELATED WORK

Segmentation of anatomical structures from abdominal scans is a prerequisite for various high-level CT-based clinical applications. Existing automated state-of-the-art tools for abdominal image segmentation are mostly based on fully convolutional network-based methods. In this section, we first briefly review the fully convolutional network-based methods for abdominal multi-organ segmentation (section II-A). We conclude this section with a discussion on multi-task deep

neural networks being employed for complementary boundary learning task to improve medical image segmentation (section II-B).

A. FULLY CONVOLUTIONAL NETWORKS FOR ABDOMINAL MULTI-ORGAN SEGMENTATION

In recent years, Fully Convolution Network (FCN) and its variants (e.g., UNet [14]) have become a common choice for medical image segmentation. This dominance can be attributed to their ability to learn effective task representations and efficient inference. UNet has an encoder-decoder style architecture and consists of skip connections, joining the encoding and decoding layers on the same level. Despite being trained from scratch, UNet demonstrated state-of-the-art performance for various medical image segmentation tasks [34], [35]. Built on top of UNet, several other modified architectures were subsequently proposed, e.g., UNet₊₊ [25], Attention-UNet [26], etc.

Roth et al. [15] proposed a cascaded architecture based on two 3D UNets where the first UNet is trained to separate the abdominal area from the background, and the latter utilized the output from the first UNet to simultaneously segment the abdominal organs. Peng et al. [33] delineated abdominal organs using 3D UNet with residual-learning based units (ResNets) to calculate patient-specific CT organ dose. In another study [2], abdominal organs are segmented using a 3D FCN with dilated convolutions based densely connected units. Heinrich et al. [11] leveraged 3D deformable convolutions to spatially adapt the receptive field for abdominal multi-organ segmentation. In [37], abdominal scans were segmented using a 3D deeply supervised patch-based UNet with grid-based attention gates to encourage the network to focus on useful salient features propagated through the skip connections. Some existing methods have employed post-processing steps, including level-sets [3] and graph-cut [4] to refine initial segmentation maps obtained from 3D deep convolutional networks.

Through the efforts mentioned above, the existing 3D methods have mostly emphasized developing better deep learning architectures and did not attempt to improve the segmentation of challenging parts of abdominal organs, e.g., voxels that belong to the contour of organs and regions within the vicinity of organ-contour. The fuzzy appearance of the boundary of organs and low contrast between the adjacent abdominal structures makes the voxels belonging to these regions more susceptible to wrong label prediction.

B. BOUNDARY-CONSTRAINED MEDICAL IMAGE SEGMENTATION

Several existing deep learning-based medical image segmentation methods have utilized the boundary information of regions of interest to overcome the misprediction of boundary pixels [17], [18], [19], [38]. In these methods, the networks are trained in a multi-task learning fashion to simultaneously predict the probability maps of entire organs and their

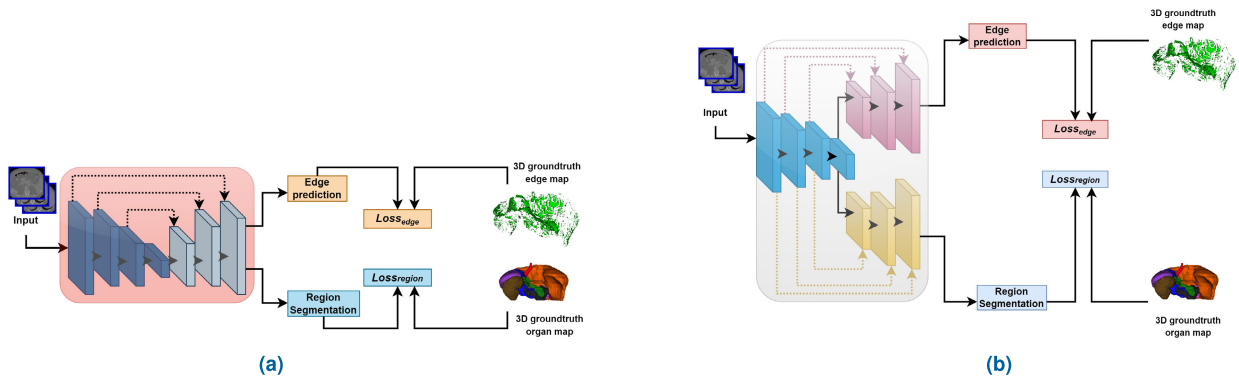


FIGURE 2. Multi-task topologies of 3D boundary-constrained network. (a) Multi-task topology with shared encoder-decoder network and task-specific prediction layers, and (b) Multi-task topology with shared encoder and task-specific decoders.

corresponding boundaries. Most of these methods have resorted to the hard-parameter sharing technique, where a single network contains shared and task-specific parameters and is jointly trained to solve multiple tasks.

Chen et al. [17] segmented the glands and their corresponding boundaries via multi-task training. By training the model to learn the co-representations, the model achieved better gland segmentation performance than the single-task models. In [39], boundary and distance maps were used for improved polyp and optic disk segmentation, respectively. Tan et al. [18] proposed a multi-task medical image segmentation network consisting of a single encoder and separate dedicated arms for decoding regions and boundaries. The study was evaluated on numerous applications, including MR femur and CT kidney segmentation. Lee et al. [38] proposed a framework that predicts boundary keypoint maps and makes use of adversarial loss for improved boundary preserving in medical image segmentation.

Given the challenge presented by voxels on the organs' boundaries and the evidence in the literature that focusing on boundaries is beneficial for performance, we integrate the organs boundary prediction as an auxiliary task into the training of state-of-the-art 3D medical image segmentation networks. Since the design choice of network topology impacts the learning process, we explore two multi-task network designs and analyze their performance. The boundary co-training resulted in improved performance on abdominal CT segmentation tasks compared to the several state-of-the-art 3D fully convolutional baseline architectures.

III. PROPOSED METHOD

In this section, we first describe the boundary-constrained loss for training the 3D encoder-decoder network to simultaneously predict the boundaries and entire abdominal organ regions via multi-task learning (Section III-A), followed by an exhibition of our proposed multi-task network topologies (Section III-B). After that, we discuss the architecture of the 3D networks that we have as baselines in our work (Section III-C). Finally, we present the architectural design of the counterpart 3D boundary-constrained models (Section III-D).

A. BOUNDARY-CONSTRAINED LOSS

Consider a 3D encoder-decoder network trained to predict the voxel labels of the abdominal CT scan with $W \times H \times Z$ dimensions, where W , H , and Z denote the length, width, and depth of the scan, respectively. Such a network takes an abdominal multi-organ CT scan as an input and outputs a labelled voxel map of the same size as the input. To utilize the boundary information of abdominal organs for improved representation learning, we train the network to predict the 3D organ-semantic masks and 3D organ-boundaries in one forward propagation pass. We formulate this problem using a multi-task learning paradigm where multiple tasks are learned jointly using shared and task-specific representations. The loss \mathcal{L} for this multi-task learning problem is a weighted combination of per-task losses, organ segmentation loss \mathcal{L}_{RS} and organ boundary detection loss \mathcal{L}_{BD} . We use multi-class dice loss [40] for evaluating the performance of the multi-organ segmentation task, given as

$$\mathcal{L}_{RS} = \sum_{c=0}^{C-1} \frac{2(\hat{y}_{i,c} \times y_{i,c})}{\hat{y}_{i,c}^2 + y_{i,c}^2} \quad (1)$$

where, $\hat{y}_{i,c}$ and $y_{i,c}$ denote the 3D multi-organ probability map and ground-truth mask, respectively, of the i^{th} abdominal CT scan. C denotes the number of organ classes.

$$\hat{y}_i = \hat{p}(x_i; \theta_s) \quad (2)$$

$$\begin{aligned} \hat{p}(x_i; \theta_s) &= \sum_{n=0}^{N-1} \hat{p}(x_{i,n}; \theta_s) \\ &= \sum_{w=0}^{W-1} \sum_{h=0}^{H-1} \sum_{z=0}^{Z-1} \hat{p}(x_{i,w,h,z}; \theta_s) \end{aligned} \quad (3)$$

where $\hat{p}(x_{i,n})$ represents the label probability of n^{th} voxel in i^{th} scan and N refers to the total number of voxels in a scan.

To evaluate the model's performance in predicting the boundaries, we use binary cross-entropy loss (shown in (4)). Binary cross-entropy loss for predicting 3D boundaries is

given as

$$\begin{aligned} \mathcal{L}_{BD} &= - \sum_{n=0}^{N-1} e_{i,n} \log(\hat{e}_{i,n}) + (1 - e_{i,n}) \log(1 - \hat{e}_{i,n}) \\ &= - \sum_{w=0}^{W-1} \sum_{h=0}^{H-1} \sum_{z=0}^{Z-1} p(x_{i,w,h,z}; \theta_s) \log(\hat{p}(x_{i,w,h,z}; \theta_s)) \\ &\quad + (1 - p(x_{i,w,h,z}; \theta_s)) \log(1 - \hat{p}(x_{i,w,h,z}; \theta_s)) \end{aligned} \quad (4)$$

\hat{e}_i and e_i represent the edge probability map and the corresponding ground-truth. $\hat{p}(x_{i,w,h,z})$ represents the edge probability of the n^{th} voxel in i^{th} scan. θ_s represents the weights of the entire deep multi-task encoder-decoder network.

The combined total loss \mathcal{L} is minimized with respect to the parameters θ_s , as shown in (5). Thus our goal is to evaluate if a network can learn more robust features and subsequently produce improved organ segmentations by being trained to explicitly recognize the boundaries.

$$\mathcal{L}(\theta_s) = \sum_{i=1}^M \mathcal{L}_{RS} + \lambda \sum_{i=1}^M \mathcal{L}_{BD} \quad (5)$$

M and λ represents the total number of CT scans in the training set and the weight assigned to the edge detection loss in (5), respectively.

We hypothesize that the additional boundary loss (\mathcal{L}_{BD}) would impose a larger penalty on erroneous contour voxels, and it subsequently pushes the optimization of the segmentation network towards the solutions with more accurate boundaries. Thus, one would potentialize the ability of a boundary-constrained network to extract features that account for the semantic abdominal organ regions and boundaries.

B. BOUNDARY-CONSTRAINED NETWORK TOPOLOGIES

Multi-task learning is generally formulated via hard-parameter sharing and soft-parameter sharing. In the hard-parameter sharing paradigm, multiple tasks share a subset of jointly optimized parameters, whereas task-specific parameters are optimized separately. In soft-parameter sharing, each task is parameterized using its own set of parameters which are jointly regularized using constraints [41]. In practice, hard-parameter sharing approaches incur much less parameter and computational cost. In our work, we formulate the multi-task learning problem via hard-parameter sharing to train the encoder-decoder network to do multiple tasks, i.e., organ segmentation and boundary detection. For deep neural networks, the hard-parameter sharing approach is realized by sharing some network layers between the tasks while keeping some layers task-specific.

We explore two different network topologies to conduct multi-task training, as shown in Figures 2a and 2b. The motivation to explore multiple topologies is to investigate the impact of sharing the larger and smaller number of parameters in the network between the two tasks. We explain these multi-task topologies below.

1) TASK-SPECIFIC OUTPUT LAYERS (TSOL)

The first multi-task topology that we explore is formulated by appending two separate prediction layers for predicting the boundaries and semantic organ masks. This topology employs an encoder-decoder network whose weights are shared between the tasks, except for the last output layers, as shown in Fig. 2a. Technically, it encourages the use of compact and tightly shared feature representations. As evident, this configuration has negligibly fewer more parameters than the single-task network. We denote this configuration as TSOL.

2) TASK-SPECIFIC DECODERS (TSD)

In second multi-task topology, we modify the 3D encoder-decoder model to have a single shared encoder but two separate decoding arms for predicting the semantic regions and boundaries. The sibling-decoding arms upsample the region and boundary maps separately. This type of formulation ensures sparse representation sharing amongst the two tasks since decoders have been parameterized separately, as shown in Fig. 2b. The presence of two synthesis paths results in having significantly more parameters than its counterpart single-task network. We refer to this configuration as TSD.

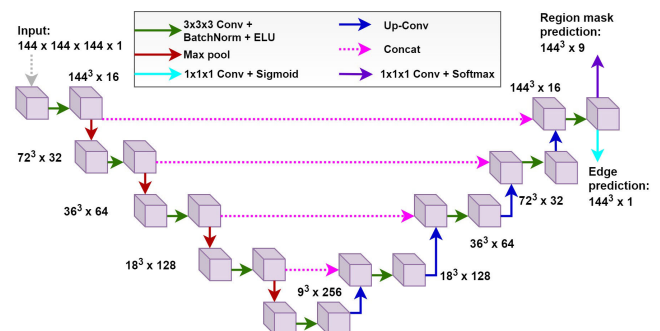


FIGURE 3. 3D UNet-MTL-TSOL: Multi-task learning based 3D UNet with task-specific output layers for simultaneously predicting the organs and their boundaries.

C. 3D BASELINE MODELS

We use UNet [24], UNet++ [25] and Attention-UNet (Att-UNet) [26] as our baseline models. We illustrate these models in Supplementary material (see Figures 15, 16 and 17). These architectures are based on encoder-decoder design and extended to segment 3D volumes by replacing the 2D convolutions, pooling, and upsampling with 3D counterparts. Each baseline model processes a 3D input scan with dimensions $W \times H \times Z$ and outputs a predicted organ-label map of the same size as input. The encoder of the model contains five convolutional blocks with pooling layers, and the decoder comprises four upsampling layers. Each convolutional block in the encoder consists of two convolutional layers with $3 \times 3 \times 3$ filters, followed by batch normalization and Exponential Linear Unit (ELU) activation [42]. We use padded convolutions to keep the output dimensions of convolutional layers the same as the

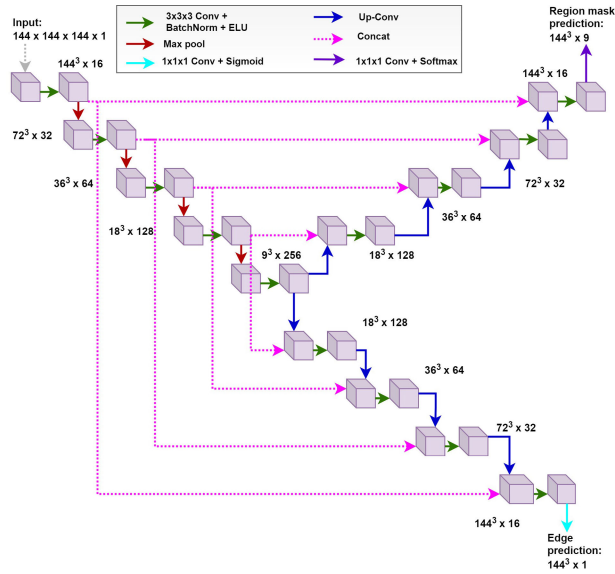


FIGURE 4. 3D UNet-MTL-TSD: Multi-task learning based 3D UNet with task-specific decoders for simultaneously predicting the organs and their boundaries.




FIGURE 6. 3D UNet₊₊-MTL-TSD: Multi-task learning based 3D UNet₊₊ with task-specific decoders for simultaneously predicting the organs and their boundaries.

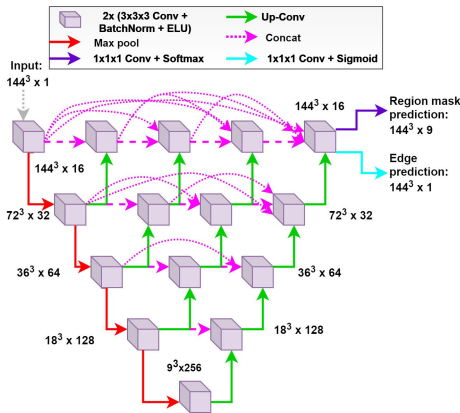


FIGURE 5. 3D UNet₊₊-MTL-TSOL: Multi-task learning based 3D UNet₊₊ with task-specific layers for simultaneously predicting the organs and their boundaries.

input dimensions. A $2 \times 2 \times 2$ max pooling layer with a stride of two in each dimension is sandwiched between every two convolutional blocks for feature maps' downsampling. The bilinear interpolation layers are used in the decoder to upsample the extracted feature maps in each dimension. The feature maps in the decoder are concatenated with the equal-sized representations learned in the encoder via skip connections. The concatenated feature maps are then transformed using convolutional blocks, similar to those used in the encoder. The last $1 \times 1 \times 1$ convolutional layer maps the feature channels to the class labels, followed by a *softmax* activation.

The resolution of the smallest feature map is $9 \times 9 \times 9$ and the minimum and maximum feature count at the first and last encoding stage is 16 and 256, respectively. Note that the original UNet₊₊ model is trained with deep supervision driven by output layers of UNet with varying depths;

however, we train UNet₊₊ without deep supervision to constrain the computational expense.

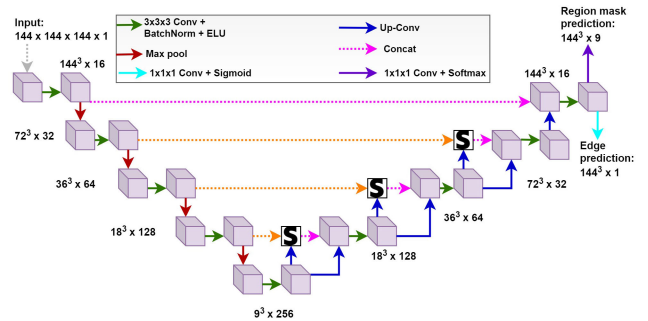


FIGURE 7. 3D Att-UNet-MTL-TSOL: Multi-task learning based 3D Att-UNet with task-specific prediction layers for simultaneously predicting the organs and their boundaries.

D. 3D BOUNDARY-CONSTRAINED MODELS

To utilize the boundary information of the organs, we train the baseline models (given in Section III-C) to predict the organ boundaries along with organs. We propose two multi-task network topologies (shown in Fig. 2) for integrating the boundary information and for that, we modify each baseline model, i.e., 3D UNet, 3D UNet₊₊, and 3D Att-UNet, according to two multi-task learning-based topologies.

To modify the baseline models according to the first multi-task topology (TSOL) (shown in Fig. 2a), we append a separate head at the end to predict boundaries along with the organs. We refer to these models as UNet-MTL-TSOL, UNet₊₊-MTL-TSOL, and Att-UNet-MTL-TSOL, as shown in Figures 3 5, and 7, respectively. To modify the baseline

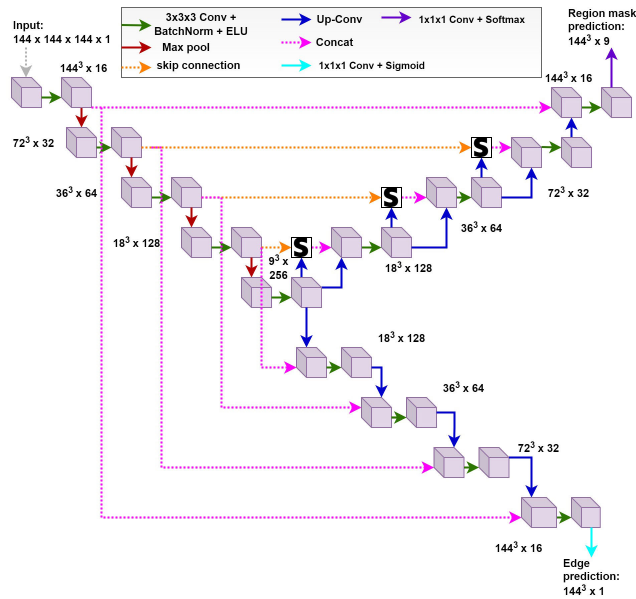


FIGURE 8. 3D Att-UNet-MTL-TSD: Multi-task learning based 3D Att-UNet with task-specific decoders for simultaneously predicting the organs and their boundaries.

models according to the second boundary-constrained multi-task topology (TSD) (shown in Fig. 2b), we modify the baseline models (3D UNet, 3D UNet₊₊, and 3D Att-UNet) to have separate decoding paths followed by prediction layers for predicting boundaries and organs. We refer to the models modified according to this topology as UNet-MTL-TSD, UNet₊₊-MTL-TSD, and Att-UNet-MTL-TSD and show them in Figures 4, 6, and 8, respectively. Note from Fig. 6, we only use the skip connections between the encoder with the greatest depth and boundary-decoder instead of utilizing feature maps extracted by nested-UNets of all depths. This design choice is made to constrain the number of parameters in UNet₊₊-MTL-TSD. Furthermore, observe from Fig. 8, we do not employ the attention mechanism while decoding the boundary-features in Att-UNet-MTL-TSD.

IV. EXPERIMENTAL DETAILS

This section first describes the datasets used to validate our study and the pre-processing we perform on the datasets (Section IV-A), followed by implementation details (Section IV-B). Finally, we conclude this section by discussing the metrics used to evaluate baseline and boundary-constrained models (Section IV-C).

A. DESCRIPTION OF DATASETS AND DATA PREPROCESSING

We utilize two publically available abdominal CT datasets (Pancreas-CT and BTCV) to evaluate baseline and boundary-constrained models. Abdominal scans in Pancreas-CT were acquired at the National Institutes of Health Clinical Center from pre-nephrectomy healthy kidney donors and subjects with neither major abdominal pathologies nor pancreatic cancer lesions [27]. The BTCV dataset consists of abdominal

scans acquired at the Vanderbilt University Medical Center from metastatic liver cancer patients or post-operative ventral hernia patients [31].

1) PANCREAS-CT DATASET (TCIA-43)

The pancreas-CT dataset [27], [43] comprised 82 abdominal contrast-enhanced 3D CT scans and was initially provided with manually drawn contours of the pancreas [43], [44]. Recently, 43 scans from this dataset have been re-annotated to include the segmentation of the liver, duodenum, stomach, esophagus, spleen, gallbladder, and left kidney [45]. Therefore, we use only 43 scans that have been re-annotated to incorporate labels for multiple organs. We randomly divide the available 43 studies into 28, 5, and 10 for training, validation, and test, respectively.

2) BTCV DATASET

BTCV was released [28], [31] as a part of a challenge held in conjunction with MICCAI 2015. The challenge compared the abdominal organs' segmentation algorithms on 3D CT scans. Our work focuses on the segmentation of eight organs from the BTCV dataset, i.e., liver, duodenum, stomach, esophagus, spleen, gallbladder, left kidney, and pancreas.

For the BTCV dataset, we utilize the bounding box coordinates given with the dataset for cropping the region-of-interest for both the CT scans and ground-truth labels [45]. Like the Pancreas-CT dataset, the cropped region-of-interest is then resampled to a common dimension of $144 \times 144 \times 144$ voxels. Finally, we randomly divide the available 47 studies into 32, 5, and 10 for training, validation, and test, respectively.

We pre-process both the datasets in the same way as proposed by [45]. This pre-processing involves cropping a tight bounding box using the bounding box coordinates provided with the dataset [45] and then resampling to a common dimension of $144 \times 144 \times 144$ voxels. This pre-processing step results in irregular voxel spacing [11]. We applied affine random transformations to augment the data but did not observe a significant difference in segmentation performance on the validation set. Therefore, we did not use any data augmentation. We used the same dataset splits for all the experiments. To analyze the occurrence of each organ in the dataset, we present the organs' occupancy ratio in Figures 9 and 10.

B. IMPLEMENTATION DETAILS

All experiments are conducted using Pytorch [46] on two Nvidia Tesla P100, accessed through the HPC platform.¹ We train all the baseline networks with a mini-batch of size 4, except for 3D UNet₊₊, which is trained with one batch size. These choices have been made according to available GPU memory. All baseline (single-task) networks are trained using multi-class dice loss as expressed in Equations (1)–(3). For acquiring multi-class dice loss, the

¹<https://supercomputing.swin.edu.au/>




FIGURE 9. Organ occupancy ratio for Pancreas-ct dataset.

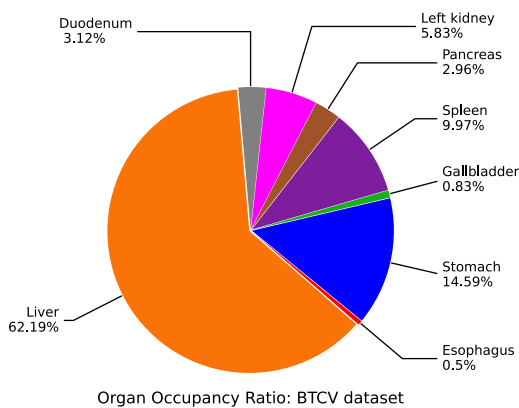


FIGURE 10. Organ occupancy ratio for BTCV dataset.

3D predicted organ-segmentation maps are compared with 3D organ ground-truth maps. All the networks (baseline and boundary-constrained) are optimized via Adam optimizer [47]. The learning rate is initially set to 0.001, decaying by a factor of 0.9 after every epoch. To assess the effect of changing values of training hyperparameters on validation segmentation performance, we conduct experiments to guide us in selecting the optimal settings for baseline models. The results for these experiments are given in Tables 5 and 6 in Supplementary section. We monitor the mean dice score on the validation set during training and utilize the model for testing that results in the highest dice coefficient on the validation set.

We use a combination of multi-class dice loss and binary cross-entropy loss, as illustrated in Equation 5 for training boundary-constrained models. 3D organ-boundary predictions are compared with the 3D boundary ground-truths to obtain the binary cross-entropy loss. Since the datasets do not contain the boundary annotations of organs, we acquire the ground-truth boundaries by first eroding the multi-organ ground-truth labels and then taking the difference from the original ground-truth map. This process gives us boundary annotations of organs. For UNet-MTL-TSOL, UNet-MTL-TSD, Att-UNet-MTL-TSOL and Att-UNet-MTL-TSD, we use a batch of size two, whereas, for

UNet₊₊-MTL-TSOL and UNet₊₊-MTL-TSD, we use a single batch size. These choices are made according to the available GPU memory. Note that the boundary predictions are only used in the training stage. During the validation stage, we only consider the organ predictions.

We conduct a grid search (on the range from 0 to 2 with a step of 0.5) to find the optimal value of λ (responsible for balancing the boundary detection loss). The exact value of λ selected to balance the boundary loss for each boundary-constrained model are given in Tables 7 and 8 in Supplementary material. These tables also present the standard deviation between validation dice scores when different values of λ are used.

C. EVALUATION METRICS

We compare the predicted segmentation masks with the ground-truths to evaluate the segmentation performance of baseline and boundary-constrained models on test set. To do that, we utilize Dice Score, Recall, Precision, and Average Hausdorff Distance as metrics for assessing the quality of predicted segmentation masks. These metrics are calculated for each organ individually and then an average is taken across all subjects. All metrics are calculated by taking a mean of the 5 runs.

V. RESULTS AND ANALYSIS

This section demonstrates the experimental results obtained from boundary-constrained abdominal segmentation models and compares them against the performance given by baseline models. For the sake of brevity, we denote the baseline and boundary-constrained models with the abbreviations below (shown in **bold**).

- 3D UNet**: 3D UNet shown in Fig. 15 in Supplementary material.
- 3D UNet-MTL-TSOL**: Boundary constrained 3D UNet with task specific output layers shown in Fig. 3.
- 3D UNet-MTL-TSD**: Boundary constrained 3D UNet with task specific decoders shown in Fig. 4.
- 3D UNet₊₊**: 3D UNet₊₊ shown in Fig. 16 in Supplementary material.
- 3D UNet₊₊-MTL-TSOL**: Boundary constrained 3D UNet₊₊ with task specific output layers shown in Fig. 5.
- 3D UNet₊₊-MTL-TSD**: Boundary constrained 3D UNet₊₊ with task specific decoders shown in Fig. 6.
- 3D Att-UNet**: 3D Att-UNet shown in Fig. 17 in Supplementary material.
- 3D Att-UNet-MTL-TSOL**: Boundary constrained 3D Att-UNet with task specific output layers shown in Fig. 7.
- 3D Att-UNet-MTL-TSD**: Boundary constrained 3D Att-UNet with task specific decoders shown in Fig. 8.

A. QUANTITATIVE RESULTS

Table 1 summarizes the segmentation results for the Pancreas-CT and BTCV datasets. We report the mean Dice

TABLE 1. Quantitative comparison between baseline and boundary-constrained models for abdominal multi-organ segmentation: Mean values \pm Std. Dev. of Dice Score, Avg. HD, Recall, and Precision are shown. Results in **BOLD indicate the best results corresponding to each baseline.**

Method	Pancreas-CT dataset				BTCV dataset			
	Dice score	Avg. Hd	Recall	Precision	Dice score	Avg. Hd	Recall	Precision
3D UNet	0.799 \pm 0.013	0.795 \pm 0.088	0.786 \pm 0.049	0.850 \pm 0.082	0.753 \pm 0.008	1.085 \pm 0.098	0.760 \pm 0.085	0.769 \pm 0.116
3D UNet-MTL-TSOL	0.813 \pm 0.009	0.748 \pm 0.130	0.807\pm0.055	0.849 \pm 0.079	0.775 \pm 0.004	0.989 \pm 0.095	0.770\pm0.086	0.799 \pm 0.106
3D UNet-MTL-TSD	0.814\pm0.005	0.703\pm0.059	0.798 \pm 0.053	0.859\pm0.079	0.776\pm0.002	0.918\pm0.047	0.764 \pm 0.081	0.802\pm0.107
3D UNet ₊₊	0.772\pm0.009	1.369 \pm 0.114	0.760\pm0.089	0.825 \pm 0.106	0.715 \pm 0.011	1.419 \pm 0.088	0.707 \pm 0.078	0.781 \pm 0.136
3D UNet ₊₊ -MTL-TSOL	0.772\pm0.014	1.170\pm0.161	0.754 \pm 0.073	0.829 \pm 0.110	0.741\pm0.011	1.189\pm0.099	0.733\pm0.080	0.789\pm0.121
3D UNet ₊₊ -MTL-TSD	0.771 \pm 0.008	1.308 \pm 0.269	0.751 \pm 0.077	0.832\pm0.111	0.723 \pm 0.008	1.349 \pm 0.027	0.708 \pm 0.084	0.783 \pm 0.127
3D Att-UNet	0.792 \pm 0.015	0.825 \pm 0.082	0.788 \pm 0.054	0.831 \pm 0.093	0.752 \pm 0.008	1.313 \pm 0.238	0.754 \pm 0.091	0.778 \pm 0.112
3D Att-UNet-MTL-TSOL	0.820\pm0.004	0.673\pm0.035	0.822\pm0.050	0.839 \pm 0.076	0.769 \pm 0.008	1.065 \pm 0.175	0.769\pm0.086	0.790 \pm 0.095
3D Att-UNet-MTL-TSD	0.807 \pm 0.005	0.732 \pm 0.049	0.797 \pm 0.050	0.847\pm0.080	0.778\pm0.004	0.919\pm0.052	0.759 \pm 0.077	0.810\pm0.104

TABLE 2. Comparison of parameter-cost and computational time.

Method	No. of parameters	Inference time (ms)
3D UNet	5.89M (-)	44
3D UNet-MTL-TSOL	5.89M (17)	53
3D UNet-MTL-TSD	8.24M (2.35M)	68
3D UNet ₊₊	6.87M	63
3D UNet ₊₊ - MTL-TSOL	6.87M (17)	79
3D UNet ₊₊ - MTL-TSD	9.22M (2.35M)	81
3D Att-UNet	6.47M	77
3D Att-UNet-MTL-TSOL	6.47M (17)	78
3D Att-MTL-TSD	8.82M (2.35M)	89

Score, mean Average Hausdorff Distance (Avg. HD) in mm, mean Recall, and mean Precision computed by comparing the predicted segmentation against the ground-truth. These measures are calculated on test sets of each dataset. Table 1 shows that the boundary-constrained models achieve improved multi-organ segmentation on the abdominal CT scans. Firstly, the mean Dice scores are improved by 1.8% (UNet vs. UNet-MTL-TSOL) and 3.5% (Att-UNet vs. Att-UNet-MTL-TSOL), for Pancreas-CT dataset. The corresponding values of mean dice scores for BTCV dataset are improved by 3.1% (UNet vs. UNet-MTL-TSD), 3.6% (UNet₊₊ vs. UNet₊₊-MTL-TSOL), and 3.5% (Att-UNet vs. Att-UNet-MTL-TSD). The improved overlap between predicted segmentations and manually annotated masks can be attributed to the enhanced semantic representations learned by boundary-constrained models.

Secondly, we observe that boundary-constrained models achieve a lower Avg. HDs for all the datasets than those obtained from baseline models as shown in Table 1. For example, after adding boundary information, the Avg. HD values of UNet, UNet₊₊, and Att-UNet are decreased by 11.5%, 14.5%, and 18.4%, respectively, for the Pancreas-CT dataset. Likewise, a decrease of 15.4%, 16.2%, and 30%, respectively, is seen for the BTCV dataset. Furthermore, we notice that the Avg. HD is still lower for the cases where boundary-constrained models obtained lower or equivalent mean Dice score, e.g., UNet₊₊-MTL-TSOL vs. UNet₊₊ and UNet₊₊-MTL-TSD vs. UNet₊₊, for Pancreas-CT results. This indicates that even for the equivalent dice overlap, the

performance of the boundary-constrained models in accurately predicting the boundaries is improved.

Thirdly, our boundary-constrained models achieve higher values of mean Recall and mean Precision for all the models and datasets except for UNet₊₊ corresponding to the Pancreas-CT dataset, as shown in Table 1. The utilization of boundary information has caused a decrease in false-negative rates and false-positive rates. Specifically, the mean recall is increased by 1.3% (UNet vs. UNet-MTL-TSOL) and 4.3% (Att-UNet vs. Att-UNet-MTL-TSOL) for Pancreas-CT dataset. For BTCV dataset, an increase of 1.3% (UNet vs. UNet-MTL-TSOL), 3.7% (UNet₊₊-MTL-TSOL vs. UNet₊₊), and 1.9% (Att-UNet vs. Att-UNet-MTL-TSOL) is observed in Mean Recalls. Finally, we see a maximum improvement of 1.9% and 4.3% in values of mean Precision for Pancreas-CT and BTCV datasets, respectively. The improvement in mean Recalls and mean Precisions show the capability of boundary-constrained models in addressing the issues of under-segmentation and over-segmentation.

Finally, we also observe that there does not exist a clear relation between the parameter-complexity of the model (number of parameters it has) and its performance. For example, the baseline UNet₊₊ has greatest number of parameters, yet it perform worse than UNet and Att-UNet.

B. COMPUTATIONAL COMPLEXITY AND ARCHITECTURAL ANALYSIS

Table 2 reports parameter count (in million M) and each method's time to segment a single CT abdominal volume in the test phase. We also highlight the increase in parameter count (given in brackets) compared with the baseline model. Among the single-task baseline models, UNet₊₊ is most parameter-extensive with 6.87×10^6 parameters. The parameter count of boundary-constrained models with TSOL topology have only 17 parameters more than the baseline models; however, these models showed significantly considerable improvements in the segmentation of abdominal organs. The shared encoder and decoder in the TSOL design enable the small parameter count while capacitating the segmentation algorithm to learn the masks and boundaries using separate

TABLE 3. Organ-wise Mean Dice scores \pm Std. Dev. obtained from baseline and best performing boundary-constrained models: **Bold** shows the highest value corresponding to each baseline.

Method	■	■	■	■	■	■	■	■
Pancreas-CT dataset								
3D UNet	0.935\pm0.008	0.700 \pm 0.035	0.927 \pm 0.014	0.793 \pm 0.030	0.714 \pm 0.024	0.934 \pm 0.009	0.829 \pm 0.024	0.566 \pm 0.031
3D UNet-MTL-TSD	0.931 \pm 0.008	0.723\pm0.028	0.937\pm0.004	0.820\pm0.028	0.718\pm0.014	0.947\pm0.003	0.854\pm0.015	0.583\pm0.021
3D UNet ₊₊	0.917\pm0.006	0.613 \pm 0.053	0.897 \pm 0.019	0.802\pm0.026	0.712\pm0.013	0.899 \pm 0.012	0.818\pm0.014	0.516 \pm 0.042
3D UNet ₊₊ - MTL-TSOL	0.915 \pm 0.005	0.621\pm0.042	0.915\pm0.013	0.774 \pm 0.051	0.697 \pm 0.013	0.913\pm0.005	0.818\pm0.032	0.519\pm0.035
3D Att-UNet	0.927 \pm 0.007	0.688 \pm 0.021	0.927\pm0.013	0.772 \pm 0.037	0.700 \pm 0.021	0.929 \pm 0.016	0.827 \pm 0.029	0.565 \pm 0.036
3D Att-UNet-MTL-TSOL	0.937\pm0.008	0.724\pm0.010	0.927\pm0.009	0.858\pm0.017	0.734\pm0.005	0.949\pm0.004	0.847\pm0.018	0.589\pm0.023
BTCV dataset								
3D UNet	0.878 \pm 0.016	0.659 \pm 0.018	0.897 \pm 0.014	0.567 \pm 0.020	0.698 \pm 0.009	0.932 \pm 0.008	0.813 \pm 0.012	0.576 \pm 0.023
3D UNet-MTL-TSD	0.897\pm0.011	0.682\pm0.009	0.912\pm0.004	0.608\pm0.019	0.732\pm0.016	0.937\pm0.007	0.837\pm0.008	0.599\pm0.011
3D UNet ₊₊	0.846 \pm 0.016	0.559 \pm 0.032	0.851 \pm 0.030	0.552 \pm 0.023	0.691 \pm 0.014	0.882 \pm 0.008	0.799 \pm 0.015	0.538 \pm 0.028
3D UNet ₊₊ - MTL-TSOL	0.853\pm0.023	0.625\pm0.018	0.891\pm0.019	0.575\pm0.019	0.694\pm0.014	0.918\pm0.004	0.829\pm0.018	0.549\pm0.032
3D Att-UNet	0.881 \pm 0.020	0.668 \pm 0.016	0.882 \pm 0.036	0.561 \pm 0.009	0.704 \pm 0.023	0.939 \pm 0.006	0.805 \pm 0.007	0.575 \pm 0.007
3D Att-UNet-MTL-TSD	0.913\pm0.009	0.674\pm0.006	0.920\pm0.015	0.603\pm0.034	0.720\pm0.009	0.947\pm0.007	0.832\pm0.022	0.608\pm0.025

TABLE 4. Organ-wise Mean Avg. HDs \pm Std. Dev. obtained from baseline and best performing boundary-constrained models: **Bold** shows the highest value corresponding to each baseline.

Method	■	■	■	■	■	■	■	■
Pancreas-CT dataset								
3D UNet	0.261\pm0.085	0.936 \pm 0.279	0.184 \pm 0.143	0.803 \pm 0.356	0.788 \pm 0.275	0.234\pm0.063	0.834\pm0.234	2.321 \pm 0.316
3D UNet-MTL-TSD	0.317 \pm 0.083	0.911\pm0.244	0.150\pm0.036	0.389\pm0.074	0.683\pm0.072	0.259 \pm 0.122	0.879 \pm 0.304	2.038\pm0.204
3D UNet ₊₊	0.419 \pm 0.069	1.852 \pm 0.554	0.265 \pm 0.153	1.569 \pm 0.649	1.445 \pm 0.494	0.711 \pm 0.218	1.371 \pm 0.457	3.315 \pm 0.637
3D UNet ₊₊ - MTL-TSOL	0.389\pm0.055	1.718\pm0.638	0.167\pm0.048	0.873\pm0.229	1.187\pm0.328	0.675\pm0.104	1.296\pm0.483	3.058\pm0.948
3D Att-UNet	0.363 \pm 0.063	0.928\pm0.085	0.174\pm0.106	0.527 \pm 0.099	0.809 \pm 0.277	0.442 \pm 0.175	0.969 \pm 0.336	2.384 \pm 0.466
3D Att-UNet-MTL-TSOL	0.255\pm0.041	0.993 \pm 0.315	0.192 \pm 0.033	0.305\pm0.075	0.559\pm0.022	0.154\pm0.054	0.766\pm0.274	2.162\pm0.434
BTCV dataset								
3D UNet	0.474 \pm 0.161	1.741\pm0.179	0.337 \pm 0.077	1.750 \pm 0.334	0.770 \pm 0.140	0.636 \pm 0.206	0.941 \pm 0.155	2.032 \pm 0.149
3D UNet-MTL-TSD	0.329\pm0.105	1.787 \pm 0.326	0.253\pm0.059	1.076\pm0.136	0.648\pm0.168	0.619\pm0.226	0.778\pm0.123	1.849\pm0.199
3D UNet ₊₊	0.780\pm0.220	2.423 \pm 0.510	0.574 \pm 0.282	2.103 \pm 0.284	1.034 \pm 0.187	0.984\pm0.136	1.153 \pm 0.148	2.306 \pm 0.206
3D UNet ₊₊ - MTL-TSOL	0.794 \pm 0.089	1.842\pm0.518	0.328\pm0.087	1.630\pm0.479	0.908\pm0.221	1.080 \pm 0.063	0.860\pm0.248	2.073\pm0.265
3D Att-UNet	0.577 \pm 0.245	2.017 \pm 0.418	0.404 \pm 0.154	2.889 \pm 1.387	0.981 \pm 0.659	0.435 \pm 0.162	1.011 \pm 0.133	2.193 \pm 0.249
3D Att-UNet-MTL-TSD	0.273\pm0.055	1.377\pm0.123	0.191\pm0.062	1.631\pm0.456	0.887\pm0.248	0.318\pm0.202	0.789\pm0.149	1.888\pm0.318

task-specific output layers. Observing the second multi-task topology TSD, our model requires approximately 2.4 M more parameters than the baseline.

Furthermore, we can see (from Table 2) that the boundary-constrained models take more time to segment a single volume at the inference stage. This behavior can be associated with the extended parameter size of boundary-constrained models since they are trained to predict a 3D boundary of organs in addition to region masks.

Analyzing the relationship between the multi-task network design and the segmentation performance from Table 1, we note there is not a single/fixed topology that led to the maximum improvement. For example, the TSD showed maximum improvement in mean DC and Avg HD corresponding to the Pancreas-CT dataset. Hd over the baseline UNet. However, when the performance of Att-UNet is compared with boundary-constrained models, we notice that the TSOL demonstrated the best results. This reveals that the multi-task network configuration that offers the best results varies depending on each baseline architecture. All in all, we found that integration of boundary information improved

the multi-organ segmentation, independent of the network topology.

C. ASSESMENT OF ORGAN-WISE SEGMENTATION PERFORMANCE

To assess which specific organs benefitted greatly from incorporating boundary information for the segmentation task, we examine the mean Dice scores and mean Avg. HDs achieved by baseline, and best performing boundary constrained models (from Table 1) for each abdominal organ. We compute the Dice scores and Average Hausdorff distances for each organ individually and then average across all subjects. From Table 3, we can see both baseline and boundary-constrained models have yielded the highest mean Dice scores for liver (■), spleen (■), and kidney (■) and lowest for duodenum (■). However, boundary-based models have led to the maximum relative improvement for the gallbladder (■), pancreas (■), and duodenum (■). From Table 4, we observe that the boundary-constrained models have significantly improved the mean Avg. HD distances for the spleen (■), kidney (■), and gallbladder (■). Finally, relating

the increase in dice overlap to the organ occurrence (shown in Figures 9 and 10), we observe that the most significant improvement has occurred for the underweighted classes. In contrast, the boundary distances are maximally decreased for both small (e.g., gallbladder (■)) and large structures like spleen (■). Furthermore, the boundary-constrained models led to lower standard deviations of the mean Dice scores and Avg. HDs across different subjects which show the robustness of proposed models.




FIGURE 11. Trimap regions along boundary of organs shown on 2D abdominal images. (a) Original abdominal image, (b) Organ label-maps corresponding to 2D image, (c) Trimap region of 5 voxel width around boundaries of organs shown in gray color, and (d) Trimap region of 7 voxel width shown in gray color.

D. SEGMENTATION PERFORMANCE ALONG BOUNDARY OF ORGANS

Unclear boundaries of organs and low inter-organ contrast prevent accurate segmentation of challenging regions around the organ boundaries on abdominal CT scans. To assess if the incorporation of boundary information has improved the segmentation of voxels within the close vicinity of organ boundaries, we evaluate the quality of predicted voxel-labels within these regions and compare them against the ones acquired via baseline methods. To do this, we generate trimaps [48], [49] with different voxel-bands surrounding the boundaries of organs. Trimap is a narrow region along the boundary of an object which is utilized to evaluate the quality of segmentation within a specific distance from the object's contour. First, we generate trimap regions around the boundary of organs for predicted and ground-truth segmentations, and then, we compare the 3D trimaps by computing mean Dice scores between them. We show the exemplary trimap regions on 2D abdominal axial slices in Figure 11. For the sake of compendious presentation, we have computed the trimap Dice scores only for TSOL network topology. Figure 12 shows the mean Dice score plotted against the number of voxels the trimap contains. The top row shows trimap plots for the Pancreas-CT dataset, whereas the bottom row shows the trimap plots for the BTCV dataset. Our proposed boundary-constrained models consistently perform better than the baseline models in predicting the semantic labels within the vicinity of organ-boundaries, except for one case, i.e., (3D UNet₊₊ vs 3D UNet₊₊-MTL-TSOL).

E. QUALITATIVE RESULTS

Figure 13 shows semantic segmentation predictions on a single 2D axial slice of the 3D scans. The first and second row correspond to segmentation results from the

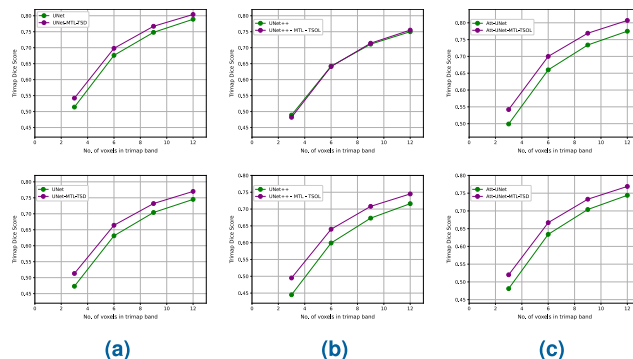


FIGURE 12. Evaluation of boundary segmentation via trimaps. In each subfigure, we have plotted the dice scores (y-axis) obtained by comparing the predicted and groundtruth segmentations in trimap regions with varying number of voxels (x-axis). The dice scores for both the baseline (green) and boundary-constrained counterparts (purple) are plotted. Top row of (a-c) corresponds to trimap segmentation comparison for Pancreas-CT dataset and bottom row corresponds to BTCV dataset.

Pancreas-CT dataset whereas, the third and fourth row corresponds to results from the BTCV dataset. Each column (from left to right) illustrates the original abdominal 2D images (Figure 13a), ground-truth masks (Figure 13b), baseline model (UNet and Att-UNet) segmentations (Figure 13c), and segmentations acquired from the boundary-constrained counterparts (UNet-MTL-TSD and Att-UNet-MTL-TSOL) in (Figure 13d). We observe that the baseline models led to under-segmentations and over-segmentations, indicated in white boxes in Figure 13c. Furthermore, the segmentations generated by single-task baseline models show isolated and biologically implausible organs' parts. Moreover, comparing with the corresponding boundary-constrained segmentations (Figure 13d), the incorporation of boundary information has prevented the issue of mispredictions near boundary of organs and led to generation of biologically plausible segmentations. Figure 14 presents the 3D segmentations generated by baseline (UNet) and boundary-constrained model (UNet-MTL-TSD) along with the ground-truths from the Pancreas-CT (first row) and BTCV dataset (second row). Notice that the boundary-constrained segmentations (third column) are more similar to the ground-truth masks (first column) as compared to the baseline segmentations (second column). These qualitative results show the improvements induced by the use of organs borders in training the 3D fully convolutional models for abdominal organs segmentation.

VI. DISCUSSION

Accurate segmentation of abdominal organs from CT scans is required for numerous advanced clinical procedures such as computer-assisted surgery and organ transplantation. The low-contrasted appearance and weak edges of abdominal organs in CT scans adversely affect the accurate segmentation. In this paper, we propose to leverage boundary information of organs as an additional cue for improved 3D abdominal multi-organ segmentation. The boundary-constrained encoder-decoder network simultaneously learns to delineate the semantic abdominal regions and detect the boundaries of




FIGURE 13. Qualitative results for multi-organ segmentation between baseline and boundary-constrained models are shown. Rows 1-2 correspond to results for the Pancreas-CT dataset, whereas Rows 3-4 show results for the BTCV dataset. Columns 1-2 show the original image and the corresponding ground-truth mask overlaid on the 2D image, i.e., spleen (■), left kidney (■), gallbladder (■), liver (■), stomach (■), and duodenum (■). Columns 3-4 illustrate the segmentation results related to baseline UNet and UNet-MTL-TSOL counterparts. White boxes indicate the segmented regions improved by incorporation of boundary information.

organs. This multi-task learning model exploits the statistics from more than one ground-truth source and subsequently retains features shared between the tasks. The boundary annotations of abdominal organs can be easily generated from the ground-truth masks and provide cost-free additional knowledge about the organs.

As reported by quantitative results in Table 1, the proposed boundary-constrained 3D encoder-decoder models achieve improved multi-organ segmentation across the majority of the baselines (3D UNet, 3D UNet₊₊, and 3D Att-UNet) and datasets (Pancreas-CT and BTCV). We have shown that the significant improvement in segmentation performance evaluated via Dice Score, Avg. HD, Recall, and Precision is caused by the improved segmentation of organs' boundaries and regions surrounding boundaries. Furthermore, significant improvements with a negligible increase in parameter count (0.0002% -TSOL topology) reveal the benefit of regularizing the existing encoder-decoder segmentation models using boundary information.

The reduction in Avg. HD for both datasets across all baselines depicts the advantage of informing the model about the vulnerable regions of the organ. The dramatic decrease in Avg. HD, ranging from 18% to 30%, shows that the model learned feature combinations that were expressive about the entire appearance of organs. We believe that training the

models with auxiliary knowledge encourages learning more generalizable and discriminative features. Notably, the additional experiments that we have conducted to precisely assess the improvements in segmentation of regions in the vicinity of organ-boundaries further verify the superiority of training the segmentation model with complementary boundary information (shown in Figure 14c).

Since there may exist several different configurations through which a fully convolutional architecture can be designed under a multi-task learning paradigm, we explore two network topologies based on the extent of parameter-sharing between the tasks. Through our extensive comparison, we notice that an overly-shared multi-task network (TSOL) performs on par with the network designed to have an increased number of task-specific layers (TSD) (Table 1). This indicates that models with many parameters do not necessarily correspond to higher performance. Most importantly, we also found that incorporation of boundary information improved the multi-organ segmentation performance, regardless of the network topology. One of the critical challenges in designing 3D multi-task deep learning models is determining which layers should be shared while keeping the computational expense reasonable. In the future, we aim to investigate other network topologies for training the encoder-decoder model in a multi-task learning fashion.




FIGURE 14. Visualization of 3D segmentation predictions using Pancreas-CT and BTCV dataset. First column corresponds to abdominal groundtruth annotations, i.e., spleen (■), pancreas (■), left kidney (■), gallbladder (■), esophagus (■), liver (■), stomach (■), and duodenum (■), whereas second and third column corresponds to volume labels predicted by UNet and UNet-MTL-TSD.

Another valuable extension of our work is to develop a mechanism/policy that can automatically dictate the sharing pattern of network layers between the two tasks.

As reported in Tables 3 and 4, the improvement in the segmentation performance of organs rarely occurring in the dataset reveals the efficacy of boundary-regularized models in compensating for their rare presence. The qualitative results shown in Figures 13 and 14 verify the positive impact of making the model aware of organ-boundaries during training. The ability of our model to simultaneously learn the improved representations of multiple organs is indicated by the qualitative examples in Figures 13 and 14, where the boundary-constrain has reduced the occurrence of over- and under-segmented organs.

VII. CONCLUSION

In this paper, we leverage organ boundary information for an improved 3D abdominal multi-organ segmentation by addressing the challenge of unclear boundaries in low-contrasted CT scans. We demonstrate that boundary information can be seamlessly introduced in the training of 3D encoder-decoder models through different multi-task configurations. The experimental results show the boundary-constrained multi-organ segmentation outperforms the ones obtained from several FCN-based baseline models, including 3D UNet, 3D UNet₊₊, 3D Attention-UNet. Furthermore, we found that the multi-task topology that shows maximum improvement is not fixed and varies depending on the baseline architecture. This insight shows that the optimal utilization of auxiliary information cannot always be harvested through a single deep multi-task design but instead requires the exploration of different multi-task topologies. Our findings also reveal that leveraging organs boundary features improves the segmentation of underweighted organs like the gallbladder, pancreas, and duodenum with a negligible parameter-overhead. Additionally, the experimental results disclose that the boundary-constrained models improve the labelling of weak sub-parts of organs in the vicinity of boundaries. We believe the proposed 3D boundary-constrained

models would be valuable for enhancing abdominal organ segmentation and utilizing those segmentations in relevant clinical applications.


REFERENCES

- [1] T. Okada, M. G. Linguraru, M. Hori, Y. Suzuki, R. M. Summers, N. Tomiyama, and Y. Sato, "Multi-organ segmentation in abdominal CT images," in *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, Aug. 2012, pp. 3986–3989.
- [2] E. Gibson, F. Giganti, Y. Hu, E. Bonmati, S. Bandula, K. Gurusamy, B. Davidson, S. P. Pereira, M. J. Clarkson, and D. C. Barratt, "Automatic multi-organ segmentation on abdominal CT with dense V-Networks," *IEEE Trans. Med. Imag.*, vol. 37, no. 8, pp. 1822–1834, Aug. 2018.
- [3] P. Hu, F. Wu, J. Peng, Y. Bao, F. Chen, and D. Kong, "Automatic abdominal multi-organ segmentation using deep convolutional neural network and time-implicit level sets," *Int. J. Comput. Assist. Radiol. Surgery*, vol. 12, no. 3, pp. 399–411, Mar. 2017.
- [4] H. Kim, J. Jung, J. Kim, B. Cho, J. Kwak, J. Y. Jang, S.-W. Lee, J.-G. Lee, and S. M. Yoon, "Abdominal multi-organ auto-segmentation using 3D-patch-based deep convolutional neural network," *Sci. Rep.*, vol. 10, no. 1, p. 6204, Apr. 2020.
- [5] Z. Xu, R. P. Burke, C. P. Lee, R. B. Baucom, B. K. Poulouse, R. G. Abramson, and B. A. Landman, "Efficient multi-atlas abdominal segmentation on clinically acquired CT with SIMPLE context learning," *Med Image Anal.*, vol. 24, no. 1, pp. 18–27, 2015.
- [6] M. Suzuki, M. G. Linguraru, and K. Okada, "Multi-organ segmentation with missing organs in abdominal CT images," in *Proc. MICCAI*, 2012, pp. 418–425.
- [7] J. J. Cerrolaza, M. Reyes, R. M. Summers, M. A. G.-Ballester, and M. G. Linguraru, "Automatic multi-resolution shape modeling of multi-organ structures," *Med Image Anal.*, vol. 25, no. 1, pp. 11–21, 2015.
- [8] T. Okada, M. G. Linguraru, M. Hori, R. M. Summers, N. Tomiyama, and Y. Sato, "Abdominal multi-organ segmentation from CT images using conditional shape-location and unsupervised intensity priors," *Med. Image Anal.*, vol. 26, no. 1, pp. 1–18, Dec. 2015.
- [9] E. Casiraghi, P. Campadelli, S. Pratisoli, and G. Lombardi, "Automatic abdominal organ segmentation from CT images," *ELCVIA Electron. Lett. Comput. Vis. Image Anal.*, vol. 8, no. 1, p. 1, Feb. 2009.
- [10] M. A. Selver, "Segmentation of abdominal organs from CT using a multi-level, hierarchical neural network strategy," *Comput Methods Programs Biomed.*, vol. 113, no. 3, pp. 830–852, 2014.
- [11] M. P. Heinrich, O. Oktay, and N. Bouteldja, "OBELISK-Net: Fewer layers to solve 3D multi-organ segmentation with sparse deformable convolutions," *Med. Image Anal.*, vol. 54, pp. 1–9, May 2019.
- [12] M. F. Bobo, S. Bao, Y. Huo, Y. Yao, J. Virostko, A. J. Plassard, I. Lyu, A. Assad, R. G. Abramson, and M. A. Hilmes, "Fully convolutional neural networks improve abdominal organ segmentation," in *Proc. SPIE*, vol. 10574, 2018, Art. no. 105742V.
- [13] A. Sinha and J. Dolz, "Multi-scale self-guided attention for medical image segmentation," *IEEE J. Biomed. Health Informat.*, vol. 25, no. 1, pp. 121–130, Jan. 2021.
- [14] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. MICCAI*, 2015, pp. 234–241.
- [15] H. R. Roth, H. Oda, X. Zhou, N. Shimizu, Y. Yang, Y. Hayashi, M. Oda, M. Fujiwara, K. Misawa, and K. Mori, "An application of cascaded 3D fully convolutional networks for medical image segmentation," *Computerized Med. Imag. Graph.*, vol. 66, pp. 90–99, Jun. 2018.
- [16] X. Zhou, K. Yamada, R. Takayama, X. Zhou, T. Hara, H. Fujita, S. Wang, and T. Kojima, "Performance evaluation of 2D and 3D deep learning approaches for automatic segmentation of multiple organs on CT images," in *Proc. SPIE*, Feb. 2018, pp. 520–525.
- [17] H. Chen, X. Qi, L. Yu, and P.-A. Heng, "DCAN: Deep contour-aware networks for accurate gland segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2487–2496.
- [18] C. Tan, L. Zhao, Z. Yan, K. Li, D. Metaxas, and Y. Zhan, "Deep multi-task and task-specific feature learning network for robust shape preserved organ segmentation," in *Proc. IEEE 15th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2018, pp. 1221–1224.
- [19] R. Zhao, W. Chen, and G. Cao, "Edge-boosted U-Net for 2D medical image segmentation," *IEEE Access*, vol. 7, pp. 171214–171222, 2019.

- [20] B. Shi, H. Zhang, R. Yan, W. Jing, J. Zang, and F. Zhang, "LR-Net: A multi-task model using relationship-based contour information to enhance the semantic segmentation of cancer regions," in *Proc. IEEE Int. Conf. Bioinf. Biomed. (BIBM)*, Dec. 2020, pp. 2371–2378.
- [21] L. Mei, X. Guo, X. Huang, Y. Weng, S. Liu, and C. Lei, "Dense contour-imbalance aware framework for colon gland instance segmentation," *Biomed. Signal Process. Control*, vol. 60, Jul. 2020, Art. no. 101988.
- [22] A. Amyar, R. Modzelewski, H. Li, and S. Ruan, "Multi-task deep learning based CT imaging analysis for COVID-19 pneumonia: Classification and segmentation," *Comput. Biol. Med.*, vol. 126, Nov. 2020, Art. no. 104037.
- [23] K. Du, X. Lin, Y. Sun, and X. Ma, "CrossInfoNet: Multi-task information sharing based hand pose estimation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 9896–9905.
- [24] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3D U-Net: Learning dense volumetric segmentation from sparse annotation," in *Proc. MICCAI*, 2016, pp. 424–432.
- [25] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang, "UNet++: Redesigning skip connections to exploit multiscale features in image segmentation," *IEEE Trans. Med. Imag.*, vol. 39, no. 6, pp. 1856–1867, Jun. 2020.
- [26] J. Schlemper, O. Oktay, M. Schaap, M. Heinrich, B. Kainz, B. Glocker, and D. Rueckert, "Attention gated networks: Learning to leverage salient regions in medical images," *Med. Image Anal.*, vol. 53, pp. 197–207, Apr. 2019.
- [27] H. R. Roth, A. Farag, E. B. Turkbey, L. Lu, J. Liu, and R. M. Summers, "Data from pancreas-CT," *Cancer Imag. Arch.*, 2016.
- [28] B. Landman, Z. Xu, J. E. Igelsias, M. Styner, T. R. Langerak, and A. Klein, "MICCAI multi-atlas labeling beyond the cranial vault—Workshop and challenge," in *Proc. MICCAI*, 2015, p. 12.
- [29] T. Heimann and H.-P. Meinzer, "Statistical shape models for 3D medical image segmentation: A review," *Med. Image Anal.*, vol. 13, no. 4, pp. 543–563, Aug. 2009.
- [30] T. F. Cootes, G. J. Edwards, and C. J. Taylor, "Active appearance models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 6, pp. 681–685, Jun. 2001.
- [31] Z. Xu, C. P. Lee, M. P. Heinrich, M. Modat, D. Rueckert, S. Ourselin, R. G. Abramson, and B. A. Landman, "Evaluation of six registration methods for the human abdomen on clinically acquired CT," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 8, pp. 1563–1572, Aug. 2016.
- [32] H. Lombaert, D. Zikic, A. Criminisi, and N. Ayache, "Laplacian forests: Semantic image segmentation by guided bagging," in *Proc. MICCAI*, pp. 496–504, 2014.
- [33] Z. Peng, X. Fang, P. Yan, H. Shan, T. Liu, X. Pei, G. Wang, B. Liu, M. K. Kalra, and X. G. Xu, "A method of rapid quantification of patient-specific organ doses for CT using deep-learning-based multi-organ segmentation and GPU-accelerated Monte Carlo dose computing," *Med. Phys.*, vol. 47, no. 6, pp. 2526–2536, Jun. 2020.
- [34] Z. Chen and Z. Chen, "RBNet: A deep neural network for unified road and road boundary detection," in *Proc. ICONIP*, 2017, pp. 677–687.
- [35] H. Dong, G. Yang, F. Liu, Y. Mo, and Y. Guo, "Automatic brain tumor detection and segmentation using U-Net based fully convolutional networks," 2017, *arXiv:1705.03820*.
- [36] L. Zhang, J. Zhang, P. Shen, G. Zhu, P. Li, X. Lu, H. Zhang, S. A. Shah, and M. Bennamoun, "Block level skip connections across cascaded V-Net for multi-organ segmentation," *IEEE Trans. Med. Imag.*, vol. 39, no. 9, pp. 2782–2793, Sep. 2020.
- [37] Y. Liu, Y. Lei, Y. Fu, T. Wang, X. Tang, X. Jiang, W. J. Curran, T. Liu, P. Patel, and X. Yang, "CT-based multi-organ segmentation using a 3D self-attention U-Net network for pancreatic radiotherapy," *Med. Phys.*, vol. 47, no. 9, pp. 4316–4324, Sep. 2020.
- [38] H. J. Lee, J. U. Kim, S. Lee, H. G. Kim, and Y. M. Ro, "Structure boundary preserving segmentation for medical image with ambiguous boundary," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 4817–4826.
- [39] B. Murugesan, K. Sarveswaran, S. M. Shankaranarayana, K. Ram, J. Joseph, and M. Sivaprakasam, "PSI-Net: Shape and boundary aware joint multi-task deep network for medical image segmentation," in *Proc. 41st Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Jul. 2019, pp. 7223–7226.
- [40] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-Net: Fully convolutional neural networks for volumetric medical image segmentation," in *Proc. 4th Int. Conf. 3D Vis. (3DV)*, Oct. 2016, pp. 565–571.
- [41] S. Ruder, "An overview of multi-task learning in deep neural networks," 2017, *arXiv:1706.05098*.
- [42] D. Clevert, T. Unterthiner, and S. Hochreiter, "Fast and accurate deep network learning by exponential linear units (ELUs)," in *Proc. ICLR*, 2016, pp. 1–14.
- [43] H. R. Roth, L. Lu, A. Farag, H. Shin, J. Liu, E. B. Turkbey, and R. M. Summers, "DeepOrgan: Multi-level deep convolutional networks for automated pancreas segmentation," in *Proc. MICCAI*, vol. 9349, 2015, pp. 556–564.
- [44] K. Clark, B. Vendt, K. Smith, J. Freymann, J. Kirby, P. Koppel, S. Moore, S. Phillips, D. Maffitt, M. Pringle, L. Tarbox, and F. Prior, "The cancer imaging archive (TCIA): Maintaining and operating a public information repository," *J. Digit. Imag.*, vol. 26, no. 6, pp. 1045–1057, Dec. 2013, doi: 10.1007/s10278-013-9622-7.
- [45] E. Gibson, F. Giganti, Y. Hu, E. Bonmati, S. Bandula, K. Gurusamy, B. Davidson, S. P. Pereira, M. J. Clarkson, and D. C. Barratt, "Multi-organ abdominal CT reference standard segmentations," Dataset Version 1.0, Feb. 22, 2018. [Online]. Available: <https://zenodo.org/record/1169361>
- [46] A. S. Paszke Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in PyTorch," in *Proc. NIPS-W*, 2017, pp. 1–4.
- [47] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [48] P. Kohli, L. Ladicky, and P. H. S. Torr, "Robust higher order potentials for enforcing label consistency," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–8, doi: 10.1109/CVPR.2008.4587417.
- [49] B. Cheng, "R. B. Girshick, P. Dollár, A. C. Berg, and A. Kirillov, "Boundary IoU: Improving object-centric image segmentation evaluation," in *Proc. CVPR*, 2021, pp. 15329–15337.



SAMRA IRSHAD is currently pursuing the Ph.D. degree with Kyung Hee University, South Korea. Her research interests include medical image analysis, adversarial learning, and evolutionary computation.



DOUGLAS P. S. GOMES received the B.E. degree in electrical engineering from the Federal University of Mato Grosso, Cuiab, Brazil, in 2013, the M.Sc. degree in power systems from the University of São Paulo, São Paulo, Brazil, in 2016, and the Ph.D. degree from Victoria University, Footscray, VIC, Australia. He has published on machine learning for power systems and computer vision applications. His research interests include artificial intelligence systems and signal processing.



SEONG TAE KIM (Member, IEEE) received the B.S. degree from Korea University, Seoul, South Korea, and the M.S. and Ph.D. degrees from the Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea, in 2014 and 2019, respectively. In 2015, he was a Visiting Researcher with the University of Toronto, Toronto, ON, Canada. From 2019 to 2021, he was a Senior Research Scientist with the Chair for Computer Aided Medical Procedures, Technical University of Munich, Munich, Germany. He is currently an Assistant Professor with the Department of Computer Science and Engineering, Kyung Hee University, Yongin-si, South Korea. He authored or coauthored more than 50 peer-reviewed journals and conference papers. His current research interests include deep learning, spatiotemporal learning, explainable deep learning, and data-efficient deep learning. He received the Best Student Paper Award of SPIE Medical Imaging, in 2018.

• • •