



City Research Online

City St George's, University of London

Citation: Marimont, S. N. & Tarroni, G. (2022). Implicit U-Net for Volumetric Medical Image Segmentation. Paper presented at the Medical Image Understanding and Analysis 26th Annual Conference, MIUA 2022, 27-29 Jul 2022, Cambridge, UK. doi: 10.1007/978-3-031-12053-4_29

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/29079/>

Link to published version: https://doi.org/10.1007/978-3-031-12053-4_29

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Implicit U-Net for volumetric medical image segmentation

Sergio Naval Marimont¹[0000-0002-7075-5586] and Giacomo Tarroni^{1,2}[0000-0002-0341-6138]

¹ CitAI Research Centre, City, University of London, London, UK

² BioMedIA, Imperial College, London, UK

{sergio.naval-marimont,giacomo.tarroni}@city.ac.uk

Abstract. U-Net has been the go-to architecture for medical image segmentation tasks, however computational challenges arise when extending the U-Net architecture to 3D images. We propose the Implicit U-Net architecture that adapts the efficient Implicit Representation paradigm to supervised image segmentation tasks. By combining a convolutional feature extractor with an implicit localization network, our implicit U-Net has 40% less parameters than the equivalent U-Net. Moreover, we propose training and inference procedures to capitalize sparse predictions. When comparing to an equivalent fully convolutional U-Net, Implicit U-Net reduces by approximately 30% inference and training time as well as training memory footprint while achieving comparable results in our experiments with two different abdominal CT scan datasets.

Keywords: efficient segmentation · supervised learning · volumetric segmentation · CT.

1 Introduction

U-Net [18] is the go-to architecture for medical image segmentation tasks [10]. A U-Net consists of two convolutional networks an encoder or feature extraction network and a decoder or localization network. U-Net incorporates *skip connections* that share feature maps directly from encoder to decoder layers with the same spatial resolution.

Different approaches have been proposed to extend U-Nets to volumetric images. 3D convolutions, as used in V-Net [13], are very computationally challenging at training and inference time [6]. Despite this limitation, U-Nets with fully 3D convolutional architectures have been the building block of general purpose segmentation techniques such as [8,7] that have been state-of-the art until recently. Lately substantial changes to the network design have been made: e.g. Cotr [22] and UNETR [5] improve U-Net performance by replacing some of the convolutions with Transformers in the architecture.

In a related research direction, there is a growing interest in reducing computational requirements towards improved practical application of deep learning in medical image analysis. Lighter and faster models lead to faster experiments

to test achievable segmentation performance, faster hyper-parameter tuning; faster inference time is highly beneficial when processing large medical datasets [20]; smaller memory footprint translates to better model portability and lower hardware requirements for hospitals and companies running actual segmentation tasks. [23] won the MICCAI FLARE 2021 ³ challenge with a U-Net architecture in which 3D decoder convolutions were replaced by anisotropic convolutions and a two step, coarse-to-fine segmentation framework. A different strategy to improve efficiency is to integrate features extracted with 2D convolutions from three orthogonal views of a volumetric image. This strategy is generally referred to as 2.5D convolution and it has been widely explored in segmentation methods [17,14,19]. However, 2.5D approaches are lagging behind in performance and they do not leverage the availability of volumetric data [6]. [4] proposes to improve efficiency by training a 3D encoder network to predict the transformations required to compare the volume against an atlas and thus associate a segmentation mask. This approach is largely efficient given that it removes most of the computation associated with the 3D decoder branch, however it requires additional validation in image modalities with high inter-subject variance such as abdominal CT scans. [1] proposed to use depth-wise separable convolutions in the encoder network to improve the overall efficiency.

Implicit representation: Implicit Field learning (or occupancy networks) has been recently proposed for 3D shape representation. Implicit networks can reconstruct 3D shapes using their implicit surface representation [11,3]. Instead of using convolutional architectures to generate dense voxel outputs, a linear neural network learns to classify as background or object the spatial coordinates [11,3]. In a related approach, [16] learns the signed distance function with respect to the object surface instead of performing object/background classification.

Implicit approaches can be more efficient than voxel representations in 3D images because they can generate sparse outputs (i.e. for only a subset of points), while in voxel representations, number of parameters and computation grow with a cubic function of the image size. Some recent research has proposed to leverage Implicit Fields in medical image analysis, including unsupervised anomaly detection [15] and super resolution [21].

Contributions: The Implicit Field formalism has allowed to avoid the use of convolutional layers and still produce smooth(er) outputs for shape reconstruction. Following the previously described research direction of making changes to the UNet architecture, we explored what could be achieved by removing the Decoding branch of the UNet, aiming for a more lightweight and faster to train design, without sacrificing accuracy.

We propose a new architecture, named Implicit U-Net, that adapts the principles of the efficient implicit representation paradigm to supervised volumetric

³ [Fast and Low GPU memory Abdominal oRgan sEgmentation](https://flare.grand-challenge.org/FLARE21/)
<https://flare.grand-challenge.org/FLARE21/>

medical image segmentation. Given that the Implicit decoder enables prediction of only a subset of points, we evaluate mechanisms to leverage this feature to improve the efficiency of both training and inference. Our main contributions are the following:

- We propose an adaptation of the implicit architecture for segmentation tasks whereby the Implicit decoder receives features from multiple spatial resolutions, replicating the function of skip-connections in a standard U-Net;
- We introduce training and inference procedures that leverage sparse predictions as a strategy to improve training and inference efficiency, obtaining important reductions in time and memory footprint;
- We tested this approach in two datasets from the Medical Segmentation Decathlon [2] and achieved accuracy comparable to the standard U-Net while reducing the inference and training time as well as memory footprint by approximately 30%.

2 Methods

Implicit field representation: 3D images are commonly represented as a dense set of voxels. Implicit field networks represent images by learning a continuous mapping f between spatial coordinates $\mathbf{p} = (x, y, z) \in \mathbb{R}^3$ and a target variable T . Therefore, implicit networks receive as inputs and make predictions for a sparse set of points. We leverage this feature to reduce computation during training and inference time.

In addition to spatial coordinates, features $\mathbf{z} \in \mathbb{R}^D$ describing the image are also received as inputs by the Implicit decoder network:

$$f : \mathbb{R}^3 \times \mathbb{R}^D \rightarrow T \quad (1)$$

In a segmentation task f learns the posterior probability over C objective classes for continuous spatial coordinates \mathbf{p} and latent features \mathbf{z} . In binary segmentation, we model the posterior probability using the logistic sigmoid activation function. Training an implicit network generally implies minimizing the training loss \mathcal{L} over a set of k points which are randomly sampled from N images. For an implicit network parametrized by θ , 2 describes the training objective:

$$\arg \min_{\theta} \sum_{i=1}^N \left(\sum_{j=1}^K \mathcal{L}(f_{\theta}(\mathbf{z}_i, \mathbf{p}_{i,j}), t_{i,j}) \right) \quad (2)$$

In our implicit decoder we used the architecture proposed in [16]. The decoder is a feed-forward network composed of 8 fully-connected layers with all hidden layers with 512 units, ReLU as activation, weight normalization and dropout 0.2 in all layers. Prior to feeding coordinates (x, y, z) to the decoder, the coordinates are normalized to the range $[-1, 1]$ and then encoded using function described in [12] (we used $L = 10$ in our experiments).

Implicit U-Net architecture: We propose to use a standard 3D convolutional neural network (CNN) to generate features \mathbf{z} that are typically passed to the decoder. In a standard CNN encoder, spatial dimensions of feature maps are contracted progressively. In our implementation we use 2-strided convolutions for pooling. Consequently, given an image with dimensions (W, H, D) , the feature map dimensions in the downward block (identified with index $b = \{0, 1, \dots, n\}$) are $(W//2^b, H//2^b, D//2^b)$, $//$ being the floor division operation.

In the original implicit network implementations, it is proposed to obtain features only from the deepest CNN encoder layer, which is expected to contain global features. However, for segmentation tasks we hypothesized that both global and local features are required to make voxel-wise predictions. Consequently, we propose an architecture that extracts features from multiple spatial resolutions at the same time. In order to achieve this, we concatenate features at each resolution in the encoder network. Specifically, for a point \mathbf{p} with spatial coordinates (x, y, z) , we concatenate the vectors with coordinates $(x//2^b, y//2^b, z//2^b)$ from the b feature maps in the CNN encoder. Intuitively, we gather the feature vector in the same relative position as the point in the original image (see Fig. 1). With this approach we intend to not only give the decoder access to local features but to provide signal directly at multiple depths in the CNN, similarly to how *deep supervision* [9] in standard U-Nets operate.

We adapted the encoder architecture from [7], using in all our experiments 2 blocks of 2 convolutions followed by 2 blocks of 4 convolutions. We also eval-

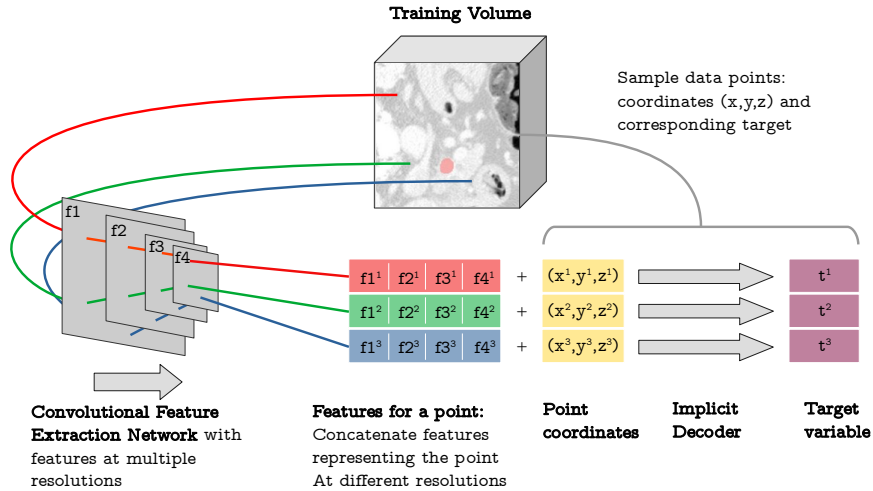


Fig. 1: Diagram of the gather feature vector operation. For example, for \mathbf{p} with coordinates $(8, 16)$, in the 2nd feature map get $(4, 8)$, in 3rd feature map get $(2, 4)$, etc. Features extracted at different resolutions are concatenated with the spatial coordinates and passed to the Implicit decoder.

uated the original architecture of 5 blocks of 2 convolutions but found that it produced *patchy* artefacts. We hypothesised that the sharp edges in the artefacts are related to the gather feature vectors operation and that the sharp straight edges are produced in the image areas where there is a transition from one deep latent vector to the next. We therefore use a fewer number of blocks but with more convolutions per block to effectively increase the receptive field of the 4th block. This subtle change in the encoder architecture reduced the artefacts because each of the deepest feature vectors is used in the predictions of smaller areas (16x16x16 voxels) compared to the architecture with 5 blocks (32x32x32 voxels) (Fig. 2).

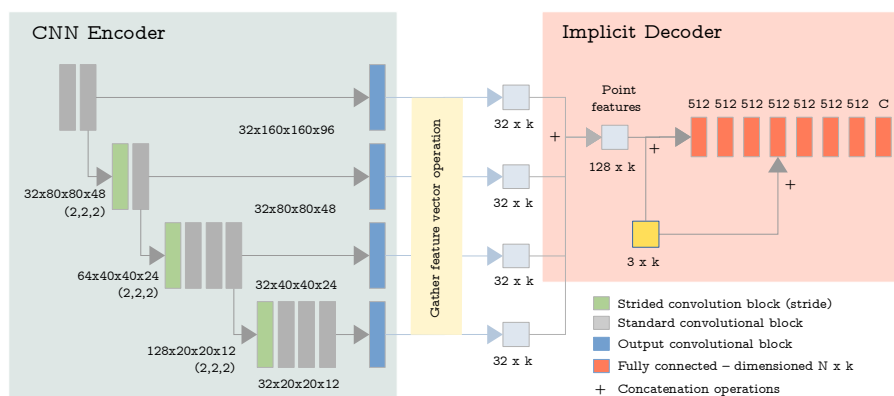


Fig.2: Implicit U-Net architecture with CNN encoder and Implicit decoder blocks. Gather feature vector operation selects the feature vectors at the relative position of the point with respect the original image.

Training procedure: The CNN encoder and Implicit decoder networks are trained simultaneously end-to-end. The training consist of 1) CNN encoder forward pass, 2) random sampling of k points from the image, 3) gather feature vector operation for the sampled points, 4) Implicit decoder forward pass with feature vectors and point coordinates and 5) backward pass.

In reference to step 2, we hypothesized that model training and final accuracy could be improved by oversampling points near label boundaries. We implemented this strategy with two hyper-parameters. $\alpha = [0, 1]$ specifies the proportion of points that are sampled from the label boundary, with $(1 - \alpha)$ being the proportion of points being sampled uniformly from the full image. Secondly, σ is used to control the distance from the boundary of the $k \times \alpha$ sampled points. The final points are obtained by adding displacement sampled from $N(0, \sigma)$ to the $k \times \alpha$ points sampled from the boundary. We performed hyper-parameter tuning for k , α and σ (see results on Fig. 3 for the Lungs dataset). The final experiments were run with $k = 30,000$, $\alpha = 0.5$ and $\sigma = 5$.

In our experiments we used a patch-based training pipeline, with the spatial coordinate system defined for each patch as opposed to the full image. We sampled patches using 1 to 2 positive to negative ratio. Patch-sizes are described in the Experiments section.

Inference procedure: With the objective of improving inference efficiency, we capitalize the feature of Implicit Fields that allows to make predictions for sparse points and propose a three-stage inference process:

- **Broad prediction:** this stage makes predictions for a subset of points in the image forming a broad mesh. In our final implementation the broad mesh was created selecting one in every 4 voxels across all dimensions, thus extracting predictions for $1/4^3 = 1/64$ of the input voxels. Predictions for the remaining points are obtained through nearest neighbour interpolation;
- **Fine boundary prediction:** it consists of 2 steps: 1) we localise a predicted segmentation boundary using the initial broad prediction; and 2) If a boundary is identified, the points near the broad boundary are predicted by the Implicit decoder to obtain the finer details around the segmentation boundary.

The proposed inference procedure adds as hyper-parameter the spacing of voxels taken in the broad prediction. We evaluated 2, 4 and 8 spacings. With a spacing of 4 we did not observe any difference in DICE score compared to predicting every single voxel and was the fastest of the three values evaluated.

Also inference is patch-based and uses a sliding-window approach with an overlap of 0.3. Patch predictions are consolidated using Gaussian weighting. Predictions are finally smoothed using a 3D average pooling filter with a kernel of size 3.




Fig. 3: Mean DICE score in Lungs dataset for different sets of hyper-parameters, starting from the baseline $k = 30,000$, $\alpha = 0.5$ and $\sigma = 5$. Results are from preliminary experiments with a batch-size 3 and 1,200 epochs.

3 Experiments and Results

3.1 Experimental set-up:

We evaluated our Implicit-U-Net on two datasets from the Medical Segmentation Decathlon [2], specifically Pancreas and Lung CT scans:

- **Task06 Lung dataset:** 64 abdominal CT scans. The task is to identify the lung tumor mass. Images were resampled to a common 1.25mm resolution along the the z-axis.
- **Task07 Pancreas dataset:** 282 contrast-enhanced, abdominal CT scans. The task is to identify pancreas (label 1) and pancreatic tumor mass (label 2). Images were resampled to a common 2.5mm resolution along the z-axis.

Image intensities were capped at the 95th percentile dataset-wise. Subsequently we applied z-score intensity normalization. Training and evaluation were both performed patch-wise with 160x160x96 patches in both abdominal CT datasets. Training augmentations included elastic transforms, scaling, rotations, axis flip, gaussian noise and random contrast.

We compared our model with an equivalent standard 3D U-Net architecture as implemented in [7]. Both standard U-Net and ours were trained using a combined DICE and Cross-Entropy loss, and AdamW optimizer with a learning rate of 10^{-4} and a batch-size of 6 for a fixed number of epochs (2,000 in Lungs and 800 in Pancreas).

5-fold cross validation was used in the two datasets. The reported results show the mean validation set performance for the best performing models fold-wise.

Network implementation, training and testing procedures are made publicly available in ⁴. Experiments were run across multiple hardware, including Nvidia 2 x RTX2080Ti, TITAN RTX and A100.

3.2 Results and Discussion:

Table 1 contains quantitative results for the experiments in the two datasets. Mean DICE score for each label and standard deviation across subjects are reported as performance metrics. Additionally, mean inference time, training time and peak GPU memory usage for 100 training steps are reported.

Fig. 4 shows qualitative comparisons in both dataset between segmentation outputs obtained with the the proposed Implicit U-Net (im-UNet) and the standard U-Net.

Implicit U-Net achieves comparable performance to U-Net in both datasets. Our technique seems to outperform the U-Net in the Lungs dataset while underperforming in tumor segmentation in the Pancreas dataset, however in all datasets performance differences are well below one standard deviation. Of note, the relatively low DICE scores and high standard deviations (e.g. in the Lungs dataset) reflect the difficulty of performing the required segmentation task. These

⁴ https://github.com/snavalm/imunet_miua22

Table 1: Quantitative results on Segmentation Decathlon datasets.

Lung					
Method	DICE	Inf t ¹	Tr t ²	Tr Mem ²	
UNet ³	64.1 ± 27.3	36.7	58.2	16,060	
im-UNet	65.7 ± 23.3	25.2	41.7	11,728	

Pancreas					
Method	DICE 1	DICE 2	Inf t ¹	Tr t ²	Tr Mem ²
UNet ³	75.8 ± 9.1	35.5 ± 29.6	9.2	59.7	16,151
im-UNet	75.8 ± 9.1	33.3 ± 29.2	6.7	44.4	11,459

1 - Mean inference time for N=10 images in seconds using sliding window.

2 - GPU time in seconds and peak memory usage for 100 training steps on 2xRTX2080Ti using mixed-precision.

3 - 3D UNet architecture with deep-supervision as in [7].




Fig. 4: Qualitative comparison of Implicit U-Net and standard U-Net in 2 different lungs images and pancreas image. For each image, columns show image and ground truth, Implicit U-Net prediction vs ground truth and U-Net prediction and ground truth. In Pancreas, first row correspond to the label 0 (Pancreas) and second row to label 1 (Tumor).

datasets where in fact selected to provide a good testbed for our comparisons. Differently from DICE scores, inference, training time and training memory show instead a clear advantage of the Implicit U-Net over the standard U-Net, with reductions in the range of 30 - 35%. It is important to note that inference time

in implicit U-Net depends on the characteristics of the lesion because of the number of fine-boundary prediction steps required at inference.

With regards to the hyper-parameters added by the model, k , α and σ for training and the broad-mesh scale for inference, we found that the final parameters proposed appeared optimal in the two Abdominal CT scan datasets evaluated. Further evaluation would be required in very different image modalities or targets.

4 Conclusion

We introduced a new strategy to improve efficiency of deep learning architectures in 3D segmentation tasks, which consists in leveraging sparse predictions in an Implicit network that replaces the standard convolutional decoder network. Our experiments show that our method achieves competitive results when compared to the reference architecture for this task (i.e. the 3D U-Net) while improving training and inference times by 30%. Training time and memory advantages lead to faster research iterations and hyper-parameter tuning. Faster inference makes Implicit U-Net relevant in the current practice dealing with larger datasets, larger image sizes and for hospitals and companies running actual segmentation tasks.

We propose Implicit U-Net in the context of growing research interest to improve U-Net architecture replacing some of the computationally expensive 3D convolutions and we believe that is complementary to other methods such as UNETR that focuses on the encoder architecture.

Future research will focus on further uncovering the power of the implicit field representations and testing the proposed approach on other datasets/settings. We will also exploring combinations of implicit decoders with transformer encoders, potentially speeding up and improving the model footprint of high-parameter architectures like UNETR.

References

1. Alalwan, N., Abozeid, A., ElHabshy, A.A.A., Alzahrani, A.: Efficient 3d deep learning model for medical image semantic segmentation. *Alexandria Engineering Journal* **60** (2021)
2. Antonelli, M., Reinke, A., Bakas, S., Farahani, K., Landman, B.A., Litjens, G., Menze, B., Ronneberger, O., Summers, R.M., van Ginneken, B., et al.: The medical segmentation decathlon. *arXiv preprint arXiv:2106.05735* (2021)
3. Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5939–5948 (2019)
4. Dong, S., Luo, G., Tam, C., Wang, W., Wang, K., Cao, S., Chen, B., Zhang, H., Li, S.: Deep atlas network for efficient 3d left ventricle segmentation on echocardiography. *Medical Image Analysis* **61** (2020)

5. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: Unetr: Transformers for 3d medical image segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 574–584 (2022)
6. Hesamian, M.H., Jia, W., He, X., Kennedy, P.: Deep learning techniques for medical image segmentation: Achievements and challenges. *Journal of Digital Imaging* **32** (2019)
7. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods* **18**, 203–211 (2021)
8. Jin, Q., Meng, Z., Sun, C., Cui, H., Su, R.: Ra-unet: A hybrid deep attention-aware network to extract liver and tumor in ct scans. *Frontiers in Bioengineering and Biotechnology* **8** (2020)
9. Lee, C.Y., Xie, S., Gallagher, P., Zhang, Z., Tu, Z.: Deeply-supervised nets. In: Artificial intelligence and statistics. pp. 562–570. PMLR (2015)
10. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., Van Der Laak, J.A., Van Ginneken, B., Sánchez, C.I.: A survey on deep learning in medical image analysis. *Medical image analysis* **42**, 60–88 (2017)
11. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4460–4470 (2019)
12. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European Conference on Computer Vision. pp. 405–421. Springer (2020)
13. 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). pp. 565–571. IEEE (2016)
14. Moeskops, P., Wolterink, J.M., van der Velden, B.H., Gilhuijs, K.G., Leiner, T., Viergever, M.A., Išgum, I.: Deep learning for multi-task medical image segmentation in multiple modalities. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). vol. 9901 LNCS (2016)
15. Naval Marimont, S., Tarroni, G.: Implicit field learning for unsupervised anomaly detection in medical images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 189–198. Springer (2021)
16. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 165–174 (2019)
17. Prason, A., Petersen, K., Igel, C., Lauze, F., Dam, E., Nielsen, M.: Deep feature learning for knee cartilage segmentation using a triplanar convolutional neural network. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). vol. 8150 LNCS (2013)
18. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). vol. 9351 (2015)

19. Roth, H.R., Lu, L., Lay, N., Harrison, A.P., Farag, A., Sohn, A., Summers, R.M.: Spatial aggregation of holistically-nested convolutional neural networks for automated pancreas localization and segmentation. *Medical Image Analysis* **45** (2018)
20. Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., et al.: Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine* **12**(3), e1001779 (2015)
21. Wu, Q., Li, Y., Xu, L., Feng, R., Wei, H., Yang, Q., Yu, B., Liu, X., Yu, J., Zhang, Y.: Irem: High-resolution magnetic resonance image reconstruction via implicit neural representation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 65–74. Springer (2021)
22. Xie, Y., Zhang, J., Shen, C., Xia, Y.: Cotr: Efficiently bridging cnn and transformer for 3d medical image segmentation. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. vol. 12903 LNCS (2021)
23. Zhang, F., Wang, Y., Yang, H.: Efficient context-aware network for abdominal multi-organ segmentation. arXiv preprint arXiv:2109.10601 (2021)