
High Resolution Medical Image Analysis with Spatial Partitioning

Le Hou^{1,2} Youlong Cheng^{1*} Noam Shazeer¹ Niki Parmar¹ Yeqing Li¹
Panagiotis Korfiatis³ Travis M. Drucker⁴ Daniel J. Blezek³ Xiaodan Song¹

¹Google Brain, Mountain View, CA

²Department of Computer Science, Stony Brook University, Stony Brook, NY

³Department of Radiology, Mayo Clinic, Rochester MN

⁴Enterprise Architecture, Mayo Clinic, Rochester MN

Abstract

Medical images such as 3D computerized tomography (CT) scans and pathology images, have 10^8 to 10^9 voxels/pixels. It is infeasible to train CNN models directly on such high resolution images, due to the memory limitation of a single GPU/TPU, and naïve data and model parallelism approaches do not work. Existing image analysis approaches alleviate this problem by cropping or down-sampling input images, which leads to complicated implementation and sub-optimal performance due to information loss. In this paper, we implement spatial partitioning, which internally distributes the input and output of convolutional layers across GPUs/TPUs. Our implementation is based on the Mesh-TensorFlow framework and the computation distribution is transparent to end users. With this technique, we train a 3D Unet on up to $512 \times 512 \times 512$ resolution data. To the best of our knowledge, this is the first work for handling such high resolution images end-to-end.

1 Introduction

Applying neural networks models [13, 3, 12, 11] on high resolution image data, such as computerized tomography (CT) scans, satellite imagery, or histopathology slides is very computational extensive. For example, CT scans are typically acquired with sub-millimeter resolution resulting in image data sizes of $512 \times 512 \times 512$ or more voxels. Furthermore, higher and higher resolution medical images are being collected [14]. Assuming that neural activations are stored in half-precision floating point numbers (2 bytes) and the batch size is 8, a 1-layer Convolutional Neural Network (CNN) with 64 filters requires more than 137GB of GPU/TPU memory. To handle such high resolution input, existing models use a combination of down-sampling, dividing, and/or coarse-to-fine schemes [8, 2, 10, 19, 4]. An obvious drawback of these methods is that potentially useful information such as contextual features, small pathological volumes or high resolution details are lost.

To scale up neural network models, the data parallelism approach divides an input batch of instances into sub-batches and distribute them across multiple GPU/TPUs. Model parallelism [6] approaches split and distribute model parameters, typically network layers, across multiple GPU/TPUs. These approaches do not solve the problem, since each GPU/TPU still need to process at least one high resolution image which results in more than 16GB of neural activations (in some cases > 32 GB) in a single GPU/TPU. More recent approaches [9, 15] are able to split the input batch along multiple dimensions, in addition to the batch dimension, and distribute parts across GPU/TPUs. However, it is not straightforward to split input images and distribute non-overlapping patches, since convolutions may take input across multiple patches. Splitting images into overlapping patches is not computationally efficient, because for most network architectures the overlap is so large such that almost the

*ylc@google.com

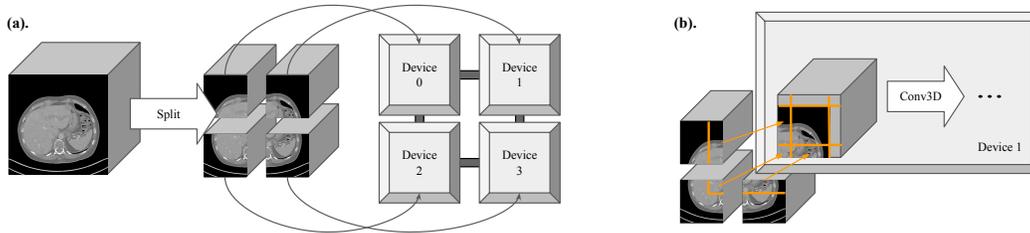


Figure 1: An illustration of spatial partitioning and halo exchange. **(a)**, Spatial partitioning: we split very high resolution medical images into non-overlapping patches. Each computational device (GPU/TPU) processes one or more patches. **(b)**, Halo exchange: before *every* convolution operation, devices exchange patch margins (half the size of the convolution kernel) with each other.

entire image is broadcast to every GPU/TPU. To overcome this, we implement spatial partitioning with halo exchange in tensorflow-TPU, to Mesh-TensorFlow. Halo exchange is a process during which GPU/TPU devices exchange data (patch margins) before convolution operations [7].

We evaluate our approach on the Liver Tumor Segmentation (LiTS) benchmark [1] data (131 CT abdominal scans with liver segmentations). Directly training neural network models on high resolution images is prone to overfitting. To address this problem, we propose a synthesis-based data augmentation method for this application.

Our contributions are: **(1)**. An open source framework² for training neural network models on high resolution images. It supports training and evaluation on both GPUs and TPUs and is highly efficient: spatial partitioning adds around 5% to the total training time. To the best of our knowledge, we are the first to train neural networks on $512 \times 512 \times 512$ resolution CT scans, in an end-to-end fashion. **(2)**. A data augmentation method for the liver tumor segmentation task using CT scans.

2 Mesh-TensorFlow with Spatial Partitioning and Halo Exchange

Mesh-TensorFlow [15] is a framework for large scale data and model parallelism. Given a collection of computational devices (*e.g.* 8 GPUs), an end user defines how to map data dimensions to these devices. For example, splitting a batch of data along its batch dimension to 8 sub-batches and send each sub-batch to 1 GPU, is an 8-way data parallelism; splitting the parameters of a fully connected layer to 2 parts and send each part to 4 GPUs, is a 2-way model parallelism. Multiple such mappings can be defined for a model, and thus data and model parallelism can happen simultaneously. Mesh-TensorFlow internally transfers data across computational devices when necessary.

Mesh-TensorFlow is successful for training giant language models such as the transformer [18]. For image analysis tasks, high resolution images consume hundreds of gigabytes of GPU/TPU memory. To address this problem, one should split images along spatial dimensions to enable model parallelism in Mesh-TensorFlow. However, the current Mesh-TensorFlow framework does not support convolutional layers on spatially split images, due to the sliding-window nature of convolutions. Modeling a convolutional layer as many fully-connected layers in Mesh-TensorFlow is unacceptable.

To enable convolutional layers on spatially split images, splitting the input into overlapping patches is not computationally efficient, since the overlap might be very large, when the network is deep. We propose to simply exchange margins of patches across computational devices, then pad the patches with received margins, and finally apply convolution. We illustrate this process in Fig. 1. With this method, we are able to train a 3D Unet model [13, 5] on $512 \times 512 \times 512$ resolution CT scans.

Data augmentation via Synthesizing As opposed to training on 2D slices or 2D+ “slabs” of 3D scans, training on full resolution CT data directly is prone to overfitting, since there are fewer instances and more features per instance. To alleviate this problem, we propose a simple yet effective data augmentation method. We assume that on the training set, segmentation ground truth of liver and tumor are given. For each training image, we compute the intensity difference between tumor and non-tumor voxels in the liver region. We then “remove” tumor by subtracting the computed intensity difference. Finally, we synthesize tumor volumes by adding the computed intensity difference on

²Code: github.com/tensorflow/mesh. Example under mesh/mesh_tensorflow/experimental.

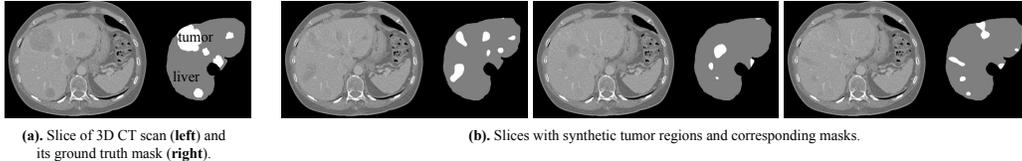


Figure 2: Given a real 3D CT scan with its ground truth mask (example of a 2d slice shown in (a)), we generate new training examples with corresponding masks (examples of 2d slices shown in (b)).

Table 1: Validation results on the LiTS benchmark [1]. We compute Dice scores of 4 models and show their mean +/- standard deviation below. Note that the results are not directly comparable to the results on the LiTS challenge leader board, due to different evaluation metrics.

	Data resolution	Dice per case	Dice global
2D	512×512×5	0.4072 +/- 0.0081	0.6432 +/- 0.0579
3D	64×64×64	0.2513 +/- 0.0024	0.5364 +/- 0.0108
	128×128×128	0.3589 +/- 0.0147	0.6494 +/- 0.0445
	256×256×256	0.4359 +/- 0.0126	0.5783 +/- 0.1654
	512×512×512	0.4547 +/- 0.0475	0.7180 +/- 0.0446

random 3D volumes inside the liver. Boundaries around synthetic tumor are blurred (Fig. 2). Without this data augmentation method, Dice scores (detailed in Sec. 3) drop at least 10%.

3 Experimental Results

We use the Liver Tumor Segmentation (LiTS) benchmark [1] for the evaluation of our implementation. This dataset has 131 3D CT images. We randomly use 99 images for training and the remaining 32 for validation. We train 3D Unet [5] models on four **3D resolutions: 64³, 128³, 256³, and 512³**. For the 64³ resolution, we train a 3D Unet with 3 blocks (each block consists of 4 convolutional layers and 1 max-pooling) in the down-sampling (encoding) part. The up-sampling (decoding) part is symmetric to the down-sampling part. Each convolutional layer in the first block has 256 filters. The number of filters doubles after each max-pooling layer. For the 128³ resolution, we attach another block with 128 filters, right after the input layer. For the 256³ resolution, we further attach a block with 64 filters. Finally, for the 512³ resolution, we attached a block with 32 filters. Thus networks at different resolutions have similar receptive field sizes. Since state-of-the-art methods on the LiTS benchmark work on 2D+ “slabs” instead of 3D data, we also test a 2D Unet on **512×512×5 resolution**. The 2D Unet has the same architecture as the 3D Unet except 3D convolution/max-pooling layers are changed to 2D. We apply data augmentation stated in Sec. 2.

We train models on a cluster of TPUs. Each TPU has 2 cores. For 64³ and 128³ data, we use 2-way data parallelism and 16-way spatial partitioning on a TPU pod of 4×4 TPUs. For 256³ data, we use 2-way data parallelism and 128-way spatial partitioning on a TPU pod of 8×16 TPUs. For 512³ data, we use 2-way data parallelism and 256-way spatial partitioning on a TPU pod of 16×16 TPUs. We use a batch size of 8 on 512³ data, and 16 on other resolutions. We use the Adafactor [16] optimizer with a learning rate of 0.003. We use $0.9 \times \text{Dice} + 0.1 \times \text{CrossEntropy}$ as the loss function [17].

We compute two Dice scores as our evaluation criterion: Dice per case, and Dice global. For Dice per case, we compute the Dice score per validation image, then average the scores across all 32 validation scans. For Dice global, we combine all scans as if there was only one scan and compute a Dice score on the whole volume [1]. We average the Dice scores of 4 models.

From the results in Tab. 1, we conclude that higher resolution data yields better Dice scores. Note that the results are not comparable to the results on the LiTS challenge leader board, due to different evaluation metrics. The Dice global scores have large standard deviations. This is because that a few CT scans in the validation set contain very large volumes of tumor, and hence prediction results of a few CT scans dominate the Dice global score. On the LiTS challenge leader board, our model trained on 99 images (out of 131 images in the training set) achieves a Dice per case score of 0.6420 and a Dice global score of 0.7660 (user: le.hou). In our experiments, operations introduced by spatial partitioning **only add around 5% to the total training time**.

References

- [1] P. Bilic, P. F. Christ, E. Vorontsov, G. Chlebus, H. Chen, Q. Dou, C.-W. Fu, X. Han, P.-A. Heng, J. Hesser, et al. The liver tumor segmentation benchmark (lits). *arXiv preprint arXiv:1901.04056*, 2019.
- [2] A. Biswas, P. Bhattacharya, and S. P. Maity. A smart system of 3d liver tumour segmentation. *International Journal of Product Development*, 23(2-3), 2019.
- [3] H. Chen, X. Qi, L. Yu, Q. Dou, J. Qin, and P.-A. Heng. Dcan: Deep contour-aware networks for object instance segmentation from histology images. *Medical Image Analysis*, 2017.
- [4] G. Chlebus, A. Schenk, J. H. Moltz, B. van Ginneken, H. K. Hahn, and H. Meine. Deep learning based automatic liver tumor segmentation in ct with shape-based post-processing. *Unpublished*, 2018.
- [5] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *MICCAI*, 2016.
- [6] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, A. Senior, P. Tucker, K. Yang, Q. V. Le, et al. Large scale distributed deep networks. In *NeurIPs*, 2012.
- [7] N. Dryden, N. Maruyama, T. Benson, T. Moon, M. Snir, and B. Van Essen. Improving strong-scaling of cnn training by exploiting finer-grained parallelism. *arXiv preprint arXiv:1903.06681*, 2019.
- [8] L. Hou, D. Samaras, T. M. Kurc, Y. Gao, J. E. Davis, and J. H. Saltz. Patch-based convolutional neural network for whole slide tissue image classification. In *CVPR*, 2016.
- [9] Z. Jia, M. Zaharia, and A. Aiken. Beyond data and model parallelism for deep neural networks. *arXiv preprint arXiv:1807.05358*, 2018.
- [10] X. Li, H. Chen, X. Qi, Q. Dou, C.-W. Fu, and P.-A. Heng. H-denseunet: hybrid densely connected unet for liver and tumor segmentation from ct volumes. *IEEE transactions on medical imaging*, 37(12), 2018.
- [11] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. In *ICCV*, 2017.
- [12] P. Mobadersany, S. Yousefi, M. Amgad, D. A. Gutman, J. S. Barnholtz-Sloan, J. E. V. Vega, D. J. Brat, and L. A. Cooper. Predicting cancer outcomes from histology and genomics using convolutional networks. *Proceedings of the National Academy of Sciences*, 115(13), 2018.
- [13] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [14] R. A. Ruddle, R. G. Thomas, R. Randell, P. Quirke, and D. Treanor. The design and evaluation of interfaces for navigating gigapixel images in digital pathology. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 23(1), 2016.
- [15] N. Shazeer, Y. Cheng, N. Parmar, D. Tran, A. Vaswani, P. Koanantakool, P. Hawkins, H. Lee, M. Hong, C. Young, et al. Mesh-tensorflow: Deep learning for supercomputers. In *NeurIPs*, 2018.
- [16] N. Shazeer and M. Stern. Adafactor: Adaptive learning rates with sublinear memory cost. *arXiv preprint arXiv:1804.04235*, 2018.
- [17] C. H. Sudre, W. Li, T. Vercauteren, S. Ourselin, and M. J. Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*. Springer, 2017.
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPs*, 2017.
- [19] E. Vorontsov, A. Tang, C. Pal, and S. Kadoury. Liver lesion segmentation informed by joint liver segmentation. In *International Symposium on Biomedical Imaging (ISBI)*, 2018.