

RUNTIME FREEZING: DYNAMIC CLASS LOSS FOR MULTI-ORGAN 3D SEGMENTATION

James Willoughby¹ Irina Voiculescu¹

¹Oxford University Department of Computer Science

ABSTRACT

Segmentation has become a crucial pre-processing step to many refined downstream tasks, and particularly so in the medical domain. Even with recent improvements in segmentation models, many segmentation tasks remain difficult. When multiple organs are segmented simultaneously, difficulties are due not only to the limited availability of labelled data, but also to class imbalance. In this work we propose dynamic class-based loss strategies to mitigate the effects of highly imbalanced training data. We show how our approach improves segmentation performance on a challenging Multi-Class 3D Abdominal Organ dataset.

1. INTRODUCTION

With the advent of isometric volumes of 3D scan data, 3D methods for automated medical image analysis are becoming increasingly common. Analysing 3D scan data is also key to crucial clinical tasks such as assessing tumor volume and borders. In much the same way that clinicians view slices with the understanding that they are part of a known complex 3D structure, a 3D segmentation network is able to make predictions more effectively by relying on the information encoded across several slices. This is most obvious in the case of organs that appear discontinuous in some 2D slices but continuous in 3D (the Pancreas in abdominal organ segmentation).

The processing of large volumes of 3D data exacerbates a problem common to automated medical image analysis: class (or organ) imbalance. Within the same volume, some organs (liver) occupy the majority of voxels, with other organs (gallbladder) being significantly smaller. Most multi-class medical image segmentation tasks present a significant imbalance in the class distributions. Whether this is the case of a multi-disease scenario or identifying multiple anatomical structures it is generally true that a few classes will be very common or represent most of the labelled pixels and the rest will be comparatively rare.

In this work we propose an improvement to 3D segmentation techniques using a strategy for adjusting per-class loss dynamically. We assess its performance compared to a basic multi-class Dice loss in three distinct 3D segmentation algorithms. As part of this work we also generalise to 3D the ‘lambda layer’ method [1], and assess its potential.




Fig. 1. Abdominal 3D Segmentation Output.

2. METHODS

2.1. Data

The data used for assessing 3D segmentation performance has been a publicly available set of abdominal CT scans from the MICCAI 2015 challenge [2]. This contains 120 different scans, of about 80 slices each, within which thirteen different organs are labelled. An additional dataset has been the TotalSegmentator dataset [3], a large dataset of 1204 CT scans containing 104 anatomical structures, from which we took any CT scans which contained segmentations for six or more of the thirteen organs segmented in the challenge dataset, totalling 923 scans of 80 slices each.

2.2. Network Structures

The 3DUNet [4] is a simple extension to 3D of the basic UNet concept: the accepted input dimensions are expanded into the convolutional layers whilst keeping the main structure the same. The UNETR [5] concept takes the basic structure of the UNet and replace the convolutional blocks with blocks more similar to the Vision Transformer (ViT) [6]. In ViT the input image is divided into smaller ‘patches’ which are projected into an embedding space, mimicking the nodes in the original Transformer concept used in natural language processing [7].

Self-Attention is then used to construct relations between all the nodes, which is advantageous in image processing as it allows for long range interactions between sections of the image to be captured better than in other methods. The UNETR network uses a three dimensional patch, dividing the input up into large voxels instead to allow for the processing of 3D data types such as MRI and CT scans.

2.2.1. *Lambda3D*

The Lambda UNet network is similar to the UNETR except it replaces the transformer blocks with ‘lambda layers’ [1]. These layers are also inspired by the Transformer concept except they replace the self-attention step with a more inexpensive linear attention function, which allows for the network to be trained on smaller patch sizes. This results in similar overall runtimes while still maintaining the Transformer’s ability to capture long-range interactions well. The Lambda network has been tested on 2D data and in a 2.5D capacity [8].

To our knowledge, there has been no literature on a 3D extension. To add to our collection of 3D networks for the proposed dynamic scheduling experiments, we have made the additional novel contribution of extending the existing 2D Lambda network concept to full 3D. Although fine tuning it has been of no immediate interest for this particular work, its performance as a 3D segmentation network is loosely compared to the UNETR and 3DUNet.

The original Lambda network mirrors the functionality of a Transformer, but with improved computational efficiency via a linear attention function. Self-attention for 3D is even more computationally expensive than for 2D, due to the exponential increase in patch number. However self-attention is able to capture the complex relationships of the 3D patches well. Including this novel Lambda3D network in the experiments is an opportunity to analyse how well its basic version performs in 3D segmentation in comparison to other 3D segmentation networks. Since it responds well to the proposed strategies, its potential can then be exploited in further work.

2.3. Loss Scheduling Strategies

Multiclass Medical imaging datasets tend to be considerably imbalanced in terms of the class distributions. The class imbalance in the Multi-organ Abdominal CT Dataset is clearly shown in the leftmost column of Table 1. On average, the liver represents 55.2% of the labelled pixels whereas the adrenal glands each only represent 0.1% of labelled pixels. We want our network to perform as well as possible on all classes but some classes are dramatically harder for the network to learn than others simply due to the quantity of available training data.

The loss function in a multiclass problem usually calculates a mean or sum of individual class losses. Most methods to address the impact of class imbalance work by weighting

the class components of the loss function. In our case the ubiquitous Dice loss function is used throughout:

$$DiceLoss = 1 - \frac{2 TP}{2 TP + FP + FN} \quad (1)$$

Weighting the class losses equally promotes a training pathway which works well for the more prevalent classes and not necessarily for the less prevalent ones. This is because the loss contributions for the more prevalent classes will decrease far more quickly. Weighting the loss contributions based on class prevalence, we will remove this tendency, as the loss will remain high after the prevalent classes have improved. This approach improves the network’s learning for less prevalent classes. However, there are a few flaws. Firstly it can create an incredibly unstable loss function in very unbalanced cases as the loss contribution for a very small class will be weighted very highly. If the network labels even one instance of this class correctly the loss function will be hugely impacted, which makes optimisation difficult as the loss is unstable.

Another issue with this method is that we make the assumption that the difficulty of learning a class is proportional to the quantity of class data in the dataset. This assumption ignores a number of factors which can make classes easier or harder to learn independent of class prevalence. Some datasets will have a majority class which is very difficult to distinguish from background pixels but distinctive minority classes, for example disseminated peritoneal metastasis vs lung nodules. In such a case the class prevalence weighting method is clearly ill-suited to the task because there is a mismatch between class “difficulty” and class prevalence. What we actually want to weight the classes based on is “relative difficulty” which is challenging to quantify a priori. Class-prevalence weighting also makes it artificially harder for the network to learn the majority classes. This is disadvantageous as classifying non-majority classes is easier with a well learned majority class due to the reduced possibility space.

2.4. Dynamic Weighting of Class Loss

We propose a loss scheduling method to address the challenge of imbalanced multi-class learning while minimising the above issues. We propose dynamically adjusting the class contributions to the loss function during the learning process in order to optimise overall performance. We define a set of class weights

$$\{w_i\}_t, i \in 1, \dots, N, \quad w_i \in [0, 1], \quad \{w_i\}_0 = 1, \forall i \in 1, \dots, N$$

where N is the number of classes and t is the epoch, and class performance scores:

$$\{S_i\}_t, i \in 1, \dots, N, \quad S_i \in [0, 1], \quad \{S_i\}_0 = 1, \forall i \in 1, \dots, N$$

at the end of each epoch we re-weight the $\{w_i\}$. This method aims to address the above concerns in the following ways.

Table 1. The effect of applying loss scheduling on 3D network structures with plain Dice loss. Overall mean Dice improves in each case. Bold indicates improvement over plain loss for individual classes. Star (*) indicates poorly performing classes.

CLASS	PIXELS(%)	UNETR	+TCF	+PCF	+CBS	3DUNET	+CBS	LAMBDA3D	+CBS
MEAN		0.777	0.787	0.765	0.792	0.484	0.588	0.707	0.725
LIVER	55.2%	0.959	0.959	0.939	0.952	0.932	0.945	0.914	0.934
STOMACH	14.1%	0.779*	0.814	0.748	0.784	0.739*	0.777	0.663*	0.678
SPLEEN	9.1%	0.916	0.925	0.916	0.917	0.831	0.788	0.797*	0.932
R KIDNEY	5.1%	0.932	0.935	0.935	0.932	0.814	0.920	0.811	0.862
L KIDNEY	5.1%	0.932	0.928	0.931	0.923	0.623*	0.912	0.831	0.818
AORTA	3.1%	0.873	0.865	0.870	0.863	0.689*	0.692	0.840	0.869
IVC	2.8%	0.797*	0.792	0.775	0.804	0.441*	0.711	0.751*	0.771
PANCREAS	2.8%	0.738*	0.735	0.698	0.745	0.474*	0.354	0.526*	0.505
VEINS	1.1%	0.669*	0.680	0.688	0.680	0.455*	0.483	0.663*	0.667
GALLBLADDER	0.9%	0.590*	0.617	0.547	0.718	0.269*	0.423	0.514*	0.484
OESOPHAGUS	0.5%	0.675*	0.704	0.713	0.707	0.157*	0.303	0.678*	0.756
R ADRENAL	0.1%	0.634*	0.661	0.631	0.644	0.348*	0	0.653*	0.630
L ADRENAL	0.1%	0.615*	0.622	0.558	0.629	0*	0	0.554*	0.517

Firstly it is closer to weighting by ‘difficulty’ because the weighting is based on current performance which is the best metric we have for ‘difficulty’ during training. Secondly, by changing the loss dynamically throughout the process we still allow the network to learn ‘easy’ classes quickly before re-weighting to promote learning for more ‘difficult’ ones. We want to be careful, however, not to re-weight in such a way that the loss function becomes unstable and difficult to optimise. We first propose the update rule *TCF*:

$$w_i = \begin{cases} 0.1, & \text{if } S_i > 0.9 \\ 1, & \text{otherwise} \end{cases} \quad (2)$$

TCF, or *Threshold Class Freezing* is designed to ‘freeze’ learning of classes that pass an empirically determined threshold of good performance, here set to 0.9. This update rule allows for fast learning of “easy” classes and does not promote instability as we reduce the contribution of classes which already contribute little to the overall loss. As a trade-off we hypothesise that any positive impact on the overall performance must be minimal.

In the current form the re-weighting is conditional on class performance reaching 0.9, which is never guaranteed. A method with a more universally applicable trigger criterion can be defined for current iteration $t = j$ as *PCF*:

$$\begin{aligned} C &= |\max(\{S_{i_t}\}, t \in [\max(j - \delta, u), j]) \\ &\quad - \max(\{S_{i_t}\}, t \in [\max(0, u), j])| \\ i &= \operatorname{argmax}(\{S_i\} / \{S_i, w_i = 0.1\}) \\ w_{i_{t=j}} &= 0.1 \quad \text{if } C > \epsilon \end{aligned} \quad (3)$$

Where u is the iteration where the update rule was last triggered. This update rule requires the local maximum of performance within an epoch window, δ , to be within some small ϵ

of the global performance maximum for a given class else the weighting of the class is decreased to 0.1. *PCF*, or *Plateau Class Freezing* is designed to freeze the learning of well performing classes based on the plateau trigger condition outlined above. Tying the weight update to performance stagnation rather than a given value allows for classes which have reached their maximum at any point to be deprioritised in favour of classes which are still improving. We restrict the i to only the best performing class with a weighting over 0.1 in order to allow the adjusted loss function to have an effect on the remaining classes.

To further develop the method we experiment with a reweighting that is more focused on addressing the poorer performing classes. We propose *CBS*, which uses the same plateau trigger condition from (2). We then define:

$$\begin{aligned} k &= \{S_k\} / \{S_k, w_k = 0.1\}) \\ w_{k_{t=j}} &= \min\left(\frac{1}{S_{i_{t=j}} + \zeta}, 2\right) \quad \text{if } C > \epsilon \end{aligned} \quad (4)$$

With some small ζ to ensure our values are defined. *CBS*, or *Class Boost Strategy* is designed to boost the class weighting based on relative performance of each class, with the trigger condition being the same as in *PCF* for more general application. This variation on the update rule relates the weighting much more closely to the relative performance for all the classes which have not plateaued. This promotes a learning path which improves performance on these classes but also keeps the range of weights within $[0, 2]$ to ensure we do not end up with an unstable loss function.

3. RESULTS AND DISCUSSION

Our experiments investigate when the strategies help improve the segmentation of challenging or poorly-performing

classes, and the impact that these strategies have on the segmentation of better-performing classes. A class is considered to be performing poorly if its Dice metric scores under 0.8. We report two sets of experiments: one where each strategy is applied to UNETR; another where *CBS* is applied to different network structures (3DUNet, Lambda3D, UNETR). Table 1 evaluates performance in all cases using the Dice metric. This is not necessarily the most appropriate metric in segmentation tasks but is universally used, and consistent with the loss function.

3.1. Scheduling strategy comparison

This experiment compares UNETR in its original form with training in the presence of one of the loss scheduling methods *TCF*, *PCF*, *CBS*. The network’s attention is focused dynamically on poorer performing classes. Performance on the veins was improved by each of the strategies. Every other poorly performing class gets segmented more effectively by at least one strategy.

Of the three, *CBS* fulfils the original intent of improving on *all* the poorly performing classes. Some classes which originally performed well have also been improved (spleen), whereas some have decreased in performance, albeit only slightly (liver, kidneys).

PCF was the worst-performing, generally decreasing the performance of the classes, this is potentially due to the plateau freezing destabilizing training. Freezing of poorer performing classes disrupts the loss far more than freezing the better performing classes, which may be making learning difficult. *TCF* performed well for a simple method, serving as a good indicator that the hypothesis is correct and that the dynamic attention is able to improve performance of more difficult classes even with a very simple weight update rule.

3.2. Scheduling strategy generalisability

This experiment evaluates the effect of *CBS*, the most successful of the three scheduling strategies, on three different 3D segmentation networks: UNETR, 3DUNet and Lambda3D (§2.2). Every network’s overall mean Dice score is improved. This effect is shown most dramatically in 3DUNet, which improved in almost every class when the loss scheduling was included and improved dramatically overall. This illustrates that dynamic multi-class scheduling can be applied to any model architecture. It is also not dependent on the choice of loss function (omitted here).

4. CONCLUSIONS AND FUTURE WORK

Prevalent classes (stomach) are shown to perform poorly. Class prevalence must never be assumed as an indicator of *class difficulty*, which is the quality that we are in reality concerned with estimating. The proposed strategies focus

the attention of a network towards poorer performing classes, balancing the segmentation performance per-class. They are generic, and can be implemented irrespective of the choice of model or loss function, allowing for this method to be applied to a wide range of problems very easily. The most sophisticated, *CBS*, brings an overall improvement to all poorly performing classes.

Our Lambda3D extension of the Lambda network is a viable 3D segmentation network. It outperforms the best method (UNETR) on the spleen and oesophagus. It is outperformed on other classes but we plan to further develop the model with other strategies such as varying patch size and domain-knowledge-inspired priors on the attention function.

The loss strategies improve segmentation where the baseline performs averagely; however, models which fail to identify any pixels (adrenal glands) do not recover from failing to generate a good enough embedding. This finding requires further investigation but, intuitively, it is hard to draw attention to an absent feature. The *CBS* weight update rule has significant scope for further iteration; the effect of cumulative weighting should be explored further; as should the use of evaluation metrics other than Dice.

5. COMPLIANCE WITH ETHICAL STANDARDS

This study was conducted retrospectively using publicly available human subject data. The authors acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility <http://dx.doi.org/10.5281/zenodo.22558>.

6. REFERENCES

- [1] Irwan Bello, “LambdaNetworks: Modeling long-range interactions without attention,” *CoRR*, 2021.
- [2] MICCAI 2015 Challenge, “Multi-atlas labeling beyond the cranial vault - workshop and challenge,” 2013.
- [3] Jakob Wasserthal et al., “TotalSegmentator: Robust segmentation of 104 anatomic structures in CT images,” *Radiology: Artificial Intelligence*, 2023.
- [4] Ö. Çiçek et al., “3D U-Net: Dense volumetric segmentation,” *MICCAI*, 2016.
- [5] A. Hatamizadeh, Y. Tang, and V. Nath et al., “UNETR: Transformers for 3D medical image segmentation,” *WACV*, 2022.
- [6] A. Dosovitskiy et al., “An image is worth 16×16 words: Transformers for image recognition,” *CoRR*, 2020.
- [7] A. Vaswani et al., “Attention is all you need,” *CoRR*, 2017.
- [8] Yanglan Ou et al., “LambdaUNet: 2.5D stroke lesion segmentation,” *arXiv preprint arXiv:2104.13917*, 2021.