

Masked Supervised Learning for Semantic Segmentation

Hasib Zunair
hasibzunair@gmail.com

Concordia University
Montreal, QC, Canada

A. Ben Hamza
hamza@ciise.concordia.ca

Abstract

Self-attention is of vital importance in semantic segmentation as it enables modeling of long-range context, which translates into improved performance. We argue that it is equally important to model short-range context, especially to tackle cases where not only the regions of interest are small and ambiguous, but also when there exists an imbalance between the semantic classes. To this end, we propose Masked Supervised Learning (MaskSup), an effective single-stage learning paradigm that models both short- and long-range context, capturing the contextual relationships between pixels via random masking. Experimental results demonstrate the competitive performance of MaskSup against strong baselines in both binary and multi-class segmentation tasks on three standard benchmark datasets, particularly at handling ambiguous regions and retaining better segmentation of minority classes with no added inference cost. In addition to segmenting target regions even when large portions of the input are masked, MaskSup is also generic and can be easily integrated into a variety of semantic segmentation methods. We also show that the proposed method is computationally efficient, yielding an improved performance by 10% on the mean intersection-over-union (mIoU) while requiring $3\times$ less learnable parameters.

1 Introduction

The basic goal of semantic segmentation, or simply segmentation, is to classify each pixel in an image into one of the pre-defined semantic categories or classes. Its real-world applications are abound, ranging from medical image analysis [19] to robotics [12]. In medical imaging, for instance, semantic segmentation can enable physicians to analyze regions of interests (ROIs) more effectively and efficiently for morphological analysis in cancer treatment, especially in high-resolution images [19], and information retrieval in diagnosis and surgery [10]. It also extends to visual scene understanding, which disentangles a scene into objects (e.g. chair), surfaces (e.g. wall) and their relations for robotic object recognition, navigation, manipulation and interaction [12].

Previous works on semantic segmentation include FCN [15] and U-Net [18], which are convolutional-based encoder-decoder networks, and have been further extended for improved performance [2, 9, 22, 30, 8]. Some recent works have also demonstrated that modeling long-range context, typically via self-attention mechanism [24], translates into better

segmentation performance [17, 21, 25, 29]. Despite the effectiveness of self-attentive models, in this paper we argue that semantic image segmentation is still a challenging problem due to a number of reasons. First, there is diversity in the size and texture of the ROIs [19]. Second, the same type of ROIs may have different sizes and colors due to the label acquisition protocol. Moreover, there are cases of ambiguity, where the boundary between the ROI and the background cannot easily be distinguished [2, 11]. Third, in the case of natural scenes, there are multiple class instances (i.e. objects and surfaces are cluttered) and there exists imbalance in the semantic classes (e.g. pixels in an image associated with the class `wall` are more than the class `chair`), as well as different lighting conditions, making the task much more difficult [11]. Examples of these challenging images are shown in Figure 1.




Figure 1: Examples demonstrating challenges in semantic segmentation. **Left to right:** First two from GLaS [19] show the variation in appearance; middle two from CVC-ClinicDB [2] show difference in scale and ambiguous (ROI); last two from NYU Depth V2 [16] show many classes with heavy imbalance under different lighting conditions.

While powerful, most of these self-attentive methods for semantic segmentation tend to over-segment ROIs, output noisy and discontinuous predictions, fail to accurately predict the boundary regions, and poorly segment minority classes. Moreover, they tend to yield misclassification in multi-class image segmentation due in part to the imbalance that exists between the semantic classes and the large number of semantic classes. We argue that the short-range context is equally important to predict small ROIs in medical images, as well as to accurately segment ROIs and reduce misclassification of minority classes in images of natural scenes, where the class instances are dense and cluttered.

To address the above limitations, we propose Masked Supervised Learning (MaskSup), a novel single-stage learning paradigm for semantic segmentation to effectively learn rich and discriminative representations. MaskSup follows a Siamese style network [8], where the two branches are identical and share weights. Given an image and its randomly masked version, MaskSup models short-range context among neighboring pixels as the context branch is tasked with predicting the semantic class of masked pixels; thereby leveraging information from non-masked pixels. MaskSup also models global or long-range context by a task similarity constraint, where the similarity of the outputs of the two branches is maximized in order to better learn the shape of class instances, and we find is especially useful in multi-class settings. The main contributions of this work can be summarized as follows:

- We propose a learning paradigm that aims to model both short- and long-range context via random masking for image segmentation without incurring any additional inference cost.
- We show through experimental results and ablation studies for binary and multi-class semantic segmentation tasks on three public datasets that MaskSup yields competitive performance in comparison with single and multi-task learning baselines.

- We demonstrate that MaskSup is robust to large masked corruptions and is computationally efficient, especially in multi-class segmentation as it improves by over 10% mIoU while at the same time requires $3\times$ less learnable parameters.

2 Related work

Single-task semantic segmentation. Most state-of-the-art segmentation methods usually follow an encoder-decoder network structure, where the image is first downsampled by the encoder subnetwork to a latent representation and then the decoder subnetwork is used to semantically project the latent representation into a pixel space for precise localization. Convolutional-based methods include FCN [15] and U-Net [18], and their variants such as U-Net++ [30], ResU-Net [27], ResU-Net++ [9]. Due to the inability of convolutional-based methods to model long-range context [25], self-attention [24] has become a core building block in various attention-based methods such as Attention U-Net [17] and Axial Attention U-Net [25]. To better segment ROIs at boundaries, Selective Feature Aggregation (SFA) [9] employs area-boundary constraints for polyp segmentation. KiU-Net [22] leverage overcomplete convolutional architectures to better segment very small ROIs and distinguish between ROI and background accurately. More recently, the advent of Vision Transformers [8] has accelerated research in the direction of transformer-based segmentation methods, which also build upon self-attention [24]. These transformer-based methods include MedT [21] and LeViT-UNet [29], which aim to learn long-range context. Our proposed framework differs from previous work in that it captures both short- and long-range context, while learning fewer parameters without compromising performance. In fact, masking enables the base segmentation model to learn short-range context among nearby pixels, as the model is tasked to make a pixel-level prediction for a masked input. This forces the network to leverage information from the nearby pixels in order to make a prediction.

Multi-task semantic segmentation. Semantic segmentation can be jointly optimized with other visual scene understanding tasks such as depth estimation and edge detection. Hybrid-Net A2 [13] is a multi-task learning method, which employs a hybrid convolutional neural network to jointly tackle the task of image segmentation and depth estimation using a single network. PAD-Net [28] is a multi-task learning and distillation based network, which jointly predicts a segmentation, depth, surface normal and edge map by multi-modal data fusion. This is extended in MTI-Net [23], where interactions between segmentation and depth estimation are captured at multiple scales when distilling information based on multi-modal distillation, in which the tasks mutually benefit from each other. Unlike multi-task learning methods, our method does not require additional training data, and hence reduces the need for intense manual labeling of additional data.

3 Proposed Method

We consider the problem of learning an encoder-decoder network \mathbf{f}_θ that classifies each pixel of an image \mathbf{I} into its semantic class category. The output is an image $\mathbf{M}_p = \mathbf{f}_\theta(\mathbf{I})$ of the same size as the input image. A gland segmentation task, for example, can be thought of as a binary segmentation problem with two semantic classes: `gland` and `background`.

3.1 Masked Supervised Learning

We present a masked supervised learning framework for effectively learning rich and discriminative representations for semantic segmentation. The proposed MaskSup method follows a Siamese network [9] style architecture, in which the segmentation branch (SB) and context branch (CB) are identical and share weights. Given an image \mathbf{I} and its randomly masked version $\mathbf{I}_{\text{masked}}$, we first employ a base segmentation network \mathbf{f}_θ to output the predictions \mathbf{M}_p and \mathbf{M}_{pm} , respectively, followed by computing an overall loss function. We use a loss function $\mathcal{L}_{\text{context}}$ to learn short-range context among neighboring pixels, as this branch **predicts the semantic class of masked pixels** by leveraging pixel information from the non-masked parts in $\mathbf{I}_{\text{masked}}$. We also use a loss term $\mathcal{L}_{\text{tasksim}}$ to learn long-range context, as the similarity of the two outputs \mathbf{M}_p and \mathbf{M}_{pm} is maximized, and hence enables us to better learn the shape of the semantic classes. Overall, MaskSup enables better representation learning for semantic segmentation by capturing the contextual relationships between pixels by predicting a segmentation map for the masked version of the input. MaskSup can tackle cases where the ROIs are ambiguous at boundary, various scales, shape, appearance and also for images that have multiple class instances with imbalance, resulting in fewer pixel-level misclassification of minority classes.

After training, we employ the network \mathbf{f}_θ to infer a new image \mathbf{I} that outputs a prediction \mathbf{M}_p . It is important to mention that at test time the images are not masked, and random masking is only used during training. The overall framework is illustrated in Figure 2.

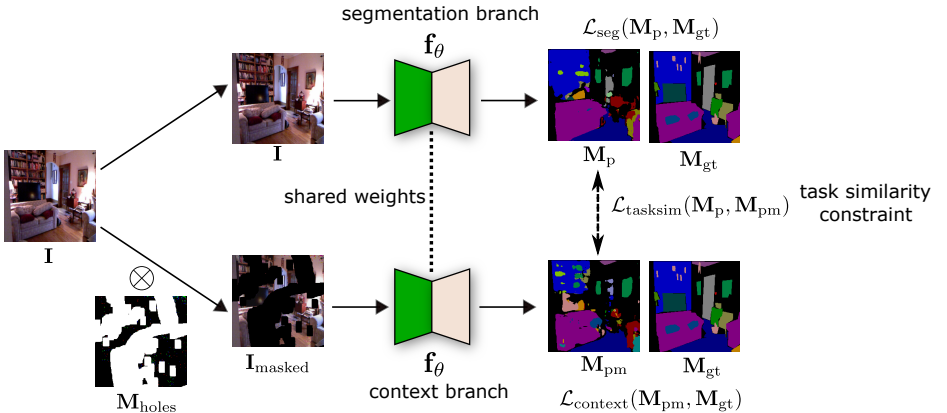


Figure 2: **Overview of MaskSup training:** joint prediction architecture with context branch and task similarity constraint for semantic segmentation, where \mathbf{f}_θ is a base segmentation network. The segmentation and context branches are identical and share weights.

Context Branch. The goal of the context branch (CB) is to learn short-range context among nearby pixels as this branch outputs predictions for masked pixels by leveraging information from the non-masked pixels in $\mathbf{I}_{\text{masked}}$. We take inspiration from the idea of image inpainting, which refers to the task of filling holes in an image, and is commonly used in image editing in order to remove image content such as text in videos or objects [14, 20]. We follow the masking procedure in [14] to construct a masked image $\mathbf{I}_{\text{masked}}$ from \mathbf{I} using masks of random streaks and holes of arbitrary shapes

$$\mathbf{I}_{\text{masked}} = \mathbf{I} \odot \mathbf{M}_{\text{holes}} \quad (1)$$

where \odot denotes element-wise multiplication and $\mathbf{M}_{\text{holes}}$ is a binary mask of random streaks and holes of arbitrary shapes. Intuitively, $\mathbf{I}_{\text{masked}}$ has a similar layout as \mathbf{I} , but randomly removes (i.e. pixel values set to 0) roughly over 50% of the pixels in the image.

Given an input image \mathbf{I} and the masked image $\mathbf{I}_{\text{masked}}$, we train an encoder-decoder network \mathbf{f}_{θ} to predict the output of the segmentation branch \mathbf{M}_p and the context branch (CB) \mathbf{M}_{pm} . Note that \mathbf{f}_{θ} is a Siamese network [9] like architecture where the branches are identical and share weights. In most of our experiments, \mathbf{f}_{θ} is either a LeViT-UNet-384 [29] for gland and polyp segmentation or U-Net++ [30] for indoor scene segmentation. We train \mathbf{f}_{θ} by minimizing the cross-entropy loss of both the segmentation branch and the context branch for all samples in the training set. The context branch loss is given by

$$\mathcal{L}_{\text{CB}} = \mathcal{L}_{\text{seg}}(\mathbf{M}_p, \mathbf{M}_{\text{gt}}) + \mathcal{L}_{\text{context}}(\mathbf{M}_{\text{pm}}, \mathbf{M}_{\text{gt}}) \quad (2)$$

where \mathcal{L}_{seg} and $\mathcal{L}_{\text{context}}$ are cross-entropy losses between the output and ground truth. It is worth mentioning that other application-specific loss functions such as the focal Tversky loss [31] and distance-based losses [32, 33] can also be used.

Task Similarity Constraint. The goal of the task similarity constraint is to model long-range context by maximizing the similarity between the output from the segmentation branch and the output from context branch in order to predict the semantic classes of masked pixels; thereby enabling us to better learn the shape of the class instances (e.g. gland, polyp, wall etc.). More specifically, we aim to maximize the similarity between the predictions made by the segmentation branch output \mathbf{M}_p and the context branch output \mathbf{M}_{pm} by minimizing the L_2 error $\mathcal{L}_{\text{tasksim}} = \|\mathbf{M}_p - \mathbf{M}_{\text{pm}}\|_2$. Therefore, the overall loss function of MaskSup is a weighted sum of the segmentation, context and task similarity loss terms

$$\mathcal{L}_{\text{total}} = \alpha_1 \mathcal{L}_{\text{seg}}(\mathbf{M}_p, \mathbf{M}_{\text{gt}}) + \alpha_2 \mathcal{L}_{\text{context}}(\mathbf{M}_{\text{pm}}, \mathbf{M}_{\text{gt}}) + \alpha_3 \mathcal{L}_{\text{tasksim}}(\mathbf{M}_p, \mathbf{M}_{\text{pm}}) \quad (3)$$

where α_1 , α_2 and α_3 are nonnegative regularization hyper-parameters, which control the contribution of each loss term. In our experiments, we empirically set them to 1.

During training, the total loss $\mathcal{L}_{\text{total}}$ is minimized for several epochs to learn the parameters of \mathbf{f}_{θ} using a labeled training set $\mathcal{D} = \{(\mathbf{I}_1, \mathbf{M}_1), \dots, (\mathbf{I}_n, \mathbf{M}_n)\}$, where \mathbf{M}_i is the ground truth segmentation mask of the input image \mathbf{I}_i . During testing, the network \mathbf{f}_{θ} is used for semantic segmentation, outputting a segmentation mask prediction \mathbf{M}_p given an input image \mathbf{I} . The architecture and the different loss terms of MaskSup are illustrated in Figure 2.

4 Experiments

In this section, we present our experimental setup and results in comparison with competing single and multi-task learning baselines for semantic segmentation. Details on datasets, implementation, architecture, training, and additional results are included in the supplementary material. Code is available at: <https://github.com/hasibzunair/masksup-segmentation>

4.1 Experimental Setup

Datasets. We demonstrate and analyze the performance of our method on Gland Segmentation (GLaS) [29], Kvasir [34] & CVC-ClinicDB [35] and NYUDv2 [36] datasets. While GLaS and Kvasir & CVC-ClinicDB are for medical image segmentation tasks, NYUDv2 is for indoor scene segmentation tasks. These datasets cover a wide range of challenges in

semantic segmentation, and they represent both binary and multi-class segmentation. They also cover both natural and medical image modalities. In medical images, ROIs are usually very small compared to background. In addition, these datasets have their own challenges such as variation in appearance, scale, ambiguous ROIs, and many class instances with imbalance that are densely cluttered.

Baselines. We evaluate the performance of our method against several state-of-the-art convolutional-based methods including FCN [15], U-Net [18], U-Net++[30], ResU-Net [27], ResU-Net++ [9], SFA [0], KiU-Net [22] and attention-based methods such as Attention U-Net [17], Axial Attention U-Net [25]. We also compare with more recent transformer-based methods MedT [21] and LeViT-UNet [29], and multi-task learning methods Hybrid-Net A2 [13] and PAD-Net [28] and MTI-Net [23] with HRNet-18 [26] as backbone.

Evaluation Metric. We report results using the Mean Intersection-Over-Union (mIoU), which is a commonly used metric in semantic segmentation [18, 21, 22, 25, 29]. The values of mIoU range from 0 to 1, with 1 indicating perfect match between the true and predicted labels, while 0 indicates a complete mismatch between them.

4.2 Results

Comparison with State-Of-The-Art. We compare the performance of MaskSup against several state-of-the-art methods and report the results in Table 1. All mIoU scores are averaged over 3 runs. As can be seen, MaskSup consistently outperforms all baselines, achieving relative improvements of 1.26%, 3.45% and 4.85% over the strongest baseline in terms of mIoU on GLaS, Kvasir & CVC-ClinicDB and NYUDv2 datasets, respectively.

MaskSup yields significant relative improvements of 18.7% over Axial Attention U-Net [25] and 2.8% over KiU-Net [22] on GLaS. MaskSup also outperforms transformer-based methods such as MedT [21] and LeViT-UNets [29] with relative improvements of 1.26% and 3.45% over LeViT-UNet-384 [29] on GLaS and Kvasir & CVC-ClinicDB, respectively. MaskSup performs better than multi-task learning methods PAD-Net [28] and HybridNet A2 [13] with relative improvements of 18.76% and 14.6%. In addition, MaskSup outperforms MTI-Net [23] with a relative improvement of 4.85% on NYUDv2. This improvement is significant because MTI-Net [23] is a multi-task learning method that jointly learns four different tasks (i.e. semantic segmentation, depth estimation, edge detection and surface normal estimation), and hence requires additional annotated data for training. In contrast, MaskSup only requires images and the pixel level annotations (i.e. segmentation masks); thereby reducing the need for intense manual labeling of additional data. The results demonstrate the effectiveness and capability of MaskSup in modeling short- and long-range context, yielding improved segmentation.

In Table 2, we report the performance comparison results of MaskSup and masked autoencoders (MAE) [8]. For MAE, we pre-train for 800 epochs on the images and fine-tune for 50 epochs on images and labels, while MaskSup only requires a single stage of training for 200 epochs. MAE uses patch-based masking to predict visual tokens similar to image inpainting, whereas MaskSup outputs a prediction label and not the full inpainted image.

Qualitative Results. In Figures 3 and 4, we visually compare MaskSup predictions against the baselines U-Net [18], U-Net++[30] and LeViT-UNets [29] on GLaS, Kvasir & CVC-ClinicDB and NYUDv2. In the first row of Figure 3, when there is a change in the overall shape and appearance of the glands, the baseline methods tend to over-segment the regions

Table 1: Performance comparison of MaskSup and baselines on GLaS, Kvasir & CVC-ClinicDB and NYUDv2 test sets using mIoU. Boldface numbers indicate the best performance, whereas the best baselines are underlined. \triangle indicates a multi-task learning method.

Method	GLaS, mIoU (\uparrow)	CVC-Clinic-DB, mIoU (\uparrow)	NYUDv2 (\uparrow)
U-Net [18]	67.41	69.74	33.60
FCN [15]	50.84	-	29.20
U-Net++ [30]	69.10	72.90	34.74
HRNet-18 [26]	-	-	33.18
ResU-Net [27]	65.95	-	-
ResU-Net++ [9]	-	79.60	-
SFA [9]	-	60.70	-
Attention U-Net [17]	-	82.70	-
Axial Attention U-Net [25]	63.03	-	-
MedT [21]	69.61	-	-
KiU-Net [22]	72.78	-	-
LeViT-UNet-128 [29]	70.45	-	-
LeViT-UNet-192 [29]	71.83	79.16	-
LeViT-UNet-384 [29]	<u>73.88</u>	<u>81.38</u>	-
PAD-Net [28] \triangle	-	-	33.10
HybridNet A2 [13] \triangle	-	-	34.30
MTI-Net [23] \triangle	-	-	<u>37.49</u>
MaskSup (Ours)	76.06	84.02	39.31

Table 2: Performance comparison of MaskSup and MAE. MaskSup is efficient and achieves better performance.

Method	GLaS, mIoU (\uparrow)	CVC-Clinic-DB, mIoU (\uparrow)	NYUDv2 (\uparrow)
MAE [8]	75.04	82.50	37.42
MaskSup (Ours)	76.06	84.02	39.31

and also produce noisy outputs as they fail to capture the global structure and semantics of the glands. The second row of Figure 3 shows the case of ambiguous ROIs of polyps, where the baselines fail to accurately segment the ROI. This is largely attributed to the limited capability of the learned representations used in the baselines. Interestingly the LeViT-UNets baseline fail to segment ROIs that are ambiguous at boundaries and vary in size and color, albeit transformers are quite strong in modeling long-range context [6, 29]. Self-attention can be regarded as a form of the non-local means [9], and it captures long-range dependencies, resulting in over-segmentation as shown in Figure 3.

Figure 4 shows that the baselines fail to accurately segment multiple class instances, output discontinuous predictions and misclassify instances. In the last row of Figure 4, we can see that the baselines fail to segment the minority class (i.e. pillow). Overall, the baselines fail to capture context of target regions, resulting in over-segmentation, noisy and discontinuous predictions as well as misclassification of instances, leading to unsatisfactory predictions.

By comparison, MaskSup is able to better capture the shape and appearance of instances

due, in large part, to the context branch, which models short-range context among pixels, resulting in better representation learning. Moreover, the task similarity constraint leads to long-range context invariance, enabling MaskSup to better learn the shape of the ROI, which in turn translates into better output predictions (see Figure 6 for more comparative results). Overall, learning with the context branch and task similarity constraint helps in cases of segmenting ambiguous ROIs at varying size and color, and also better segment minority class instances in cases of multi-class segmentation.




Figure 3: Visual comparison of MaskSup and baselines on the GLaS and Kvasir & CVC-ClinicDB test sets. MaskSup outputs better predictions in cases of variation in overall appearance and also very small and ambiguous ROIs.




Figure 4: Visual comparison of MaskSup and baselines on the NYU Depth V2 test set. MaskSup is able to output better predictions for minority classes (e.g. pillows).

4.3 Ablation study

Effectiveness of Context Branch. Figure 5 illustrates the benefit of using the the context branch on both convolutional and transformer-based methods. Using the the context branch leads to better modeling of local semantics, as it is tasked to output pixel-wise predictions for **masked regions** in the input; thereby leveraging information from neighboring pixels. We can see that the context branch improves performance of different segmentation methods across the three datasets. This shows that MaskSup is generic and can be easily integrated

into existing image segmentation methods. However, it is important to mention that a higher performance improvement is observed when the architecture is a transformer-based method due to its key characteristic of modeling long-range context [8, 29].

Effectiveness of Task Similarity Constraint. Figure 5 shows the benefit of using the task similarity constraint. We observe that the task similarity constraint further improves performance of both convolutional and transformer-based methods. The use of the task similarity constraint results in learning long-range context invariant representations, which help capture the shape of the ROI. This, in turn, leads to accurate predictions even in cases of different shapes and appearance of instances, ambiguous ROIs at different sizes, and imbalance among multiple class instances in multi-class segmentation.




Figure 5: Ablation study of different modules of MaskSup on GLaS, Kvasir & CVC-ClinicDB and NYU Depth V2 test sets. MaskSup (CB & TS) consistently improves performance of various baselines in both binary and multi-class image segmentation tasks.

Amount of Masking. We performed an ablation study of high- and low-masked pixels regions during MaskSup training, and the results are reported in Table 3, which shows that masking the images heavily during training yields better performance of MaskSup across all three datasets.

Table 3: Ablation study of high- and low-masked pixels regions during MaskSup training.

Masking	GLaS, mIoU (↑)	CVC-Clinic-DB, mIoU (↑)	NYUDv2 (↑)
Low	75.65	81.80	35.33
High	76.06	84.02	39.31

4.4 Analysis

Robustness to Masked Corruptions. Figure 6 shows the robustness of MaskSup to masked corruptions. As can be seen, MaskSup is able to better predict the ROI even when a large portion of the image is masked, demonstrating its capability in modeling short- and long-range context. Using both the context branch and task similarity constraint, MaskSup is able to learn context invariant representations to better segment and preserve the ROI shape.

Computational Efficiency. In Table 4, we report the number of parameters in millions (M), as well as mIoU for MaskSup and baseline methods. MaskSup with LeViT-192 network outperforms LeViT-384 [29] on both GLaS and Kvasir & CVC-ClinicDB, while having almost $2.6\times$ fewer learnable parameters. MaskSup with U-Net also outperforms U-Net++ [50], which has almost $3\times$ more learnable parameters on NYUDv2, with a relative improvement of 10.91% in terms of mIoU. Hence, there is no trade-off between segmentation accuracy and computational efficiency when using MaskSup in comparison with scaled versions of the networks.




Figure 6: Visual comparison of predictions made by MaskSup and baseline for images with **masked** regions.

Table 4: Comparison of MaskSup and baselines on GLaS, Kvasir & CVC-ClinicDB and NYUDv2 test sets. MaskSup is computationally efficient and achieves superior performance with fewer parameters. Boldface numbers indicate better performance.

Method	Params (M) (↓)	GLaS, mIoU (↑)	CVC-Clinic-DB, mIoU (↑)	NYUDv2 (↑)
LeViT-384 [49]	51	73.88	81.38	-
MaskSup (LeViT-192)	19(2.6x)	74.44(+0.75)	82.17(+0.97)	-
U-Net++ [60]	9	-	-	34.74
MaskSup (U-Net)	3(3x)	-	-	38.54(+10.91)

5 Conclusion

We introduced a new learning paradigm, called Masked Supervised Learning, for semantic segmentation. By constructing a randomly masked version of the input image, we first make predictions using a base segmentation network on the two inputs. Then, we maximize the predictions between the two outputs to model both short- and long-range context. MaskSup can be easily integrated into any existing semantic segmentation method. We show that MaskSup achieves better performance than state-of-the-art single and multi-task learning baselines in both binary and multi-class semantic segmentation tasks, especially in tackling small, ambiguous regions and minority class instances. In addition, MaskSup is robust to masked corruptions and is computationally efficient without compromising performance.

For future work, we aim to investigate what type of masking strategies works best in MaskSup. Since MaskSup is a generic paradigm, we plan to adapt it to other computer vision tasks such as multi-label recognition, object detection and human pose estimation. We also plan to explore high-resolution segmentation (e.g. Cityscapes, ADE20K datasets) using MaskSup.

References

- [1] Nabila Abraham and Naimul Mefraz Khan. A novel focal Tversky loss function with improved attention U-Net for lesion segmentation. In *Proc. IEEE International Symposium on Biomedical Imaging*, pages 683–687, 2019.
- [2] Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Debora Gil, Cristina Rodríguez, and Fernando Vilariño. WM-DOVA maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Computerized Medical Imaging and Graphics*, 43:99–111, 2015.
- [3] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a "siamese" time delay neural network. In *Advances in Neural Information Processing Systems*, 1993.
- [4] Antoni Buades, Bartomeu Coll, and J-M Morel. A non-local algorithm for image denoising. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2005.
- [5] Francesco Caliva, Claudia Iriondo, Alejandro Morales Martinez, Sharmila Majumdar, and Valentina Pedoia. Distance map loss penalty term for semantic segmentation. *arXiv preprint arXiv:1908.03679*, 2019.
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [7] Yuqi Fang, Cheng Chen, Yixuan Yuan, and Kai-yu Tong. Selective feature aggregation network with area-boundary constraints for polyp segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 302–310. Springer, 2019.
- [8] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022.
- [9] Debesh Jha, Pia H Smedsrud, Michael A Riegler, Dag Johansen, Thomas De Lange, Pål Halvorsen, and Håvard D Johansen. ResUNet++: An advanced architecture for medical image segmentation. In *Proc. IEEE International Symposium on Multimedia*, pages 225–2255. IEEE, 2019.
- [10] Debesh Jha, Pia H Smedsrud, Michael A Riegler, Pål Halvorsen, Thomas de Lange, Dag Johansen, and Håvard D Johansen. Kvasir-SEG: A segmented polyp dataset. In *International Conference on Multimedia Modeling*, pages 451–462. Springer, 2020.
- [11] Davood Karimi and Septimiu E Salcudean. Reducing the hausdorff distance in medical image segmentation with convolutional neural networks. *IEEE Transactions on Medical Imaging*, 39(2):499–513, 2019.
- [12] Kevin Lai, Liefeng Bo, Xiaofeng Ren, and Dieter Fox. A large-scale hierarchical multi-view RGB-D object dataset. In *Proc. IEEE Conference on Robotics and Automation*, pages 1817–1824, 2011.

- [13] Xiao Lin, Dalila Sánchez-Escobedo, Josep R Casas, and Montse Pardàs. Depth estimation and semantic segmentation from a single RGB image using a hybrid convolutional neural network. *Sensors*, 19(8):1795, 2019.
- [14] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proc. European Conference on Computer Vision*, pages 85–100, 2018.
- [15] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
- [16] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from RGBD images. In *Proc. European Conference on Computer Vision*, 2012.
- [17] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention U-Net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*, 2018.
- [18] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.
- [19] Korsuk Sirinukunwattana, Josien PW Pluim, Hao Chen, Xiaojuan Qi, Pheng-Ann Heng, Yun Bo Guo, Li Yang Wang, Bogdan J Matuszewski, Elia Bruni, Urko Sanchez, et al. Gland segmentation in colon histology images: The GLAS challenge contest. *Medical Image Analysis*, 35:489–502, 2017.
- [20] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *Proc. IEEE Winter Conference on Applications of Computer Vision*, pages 2149–2159, 2022.
- [21] Jeya Maria Jose Valanarasu, Poojan Oza, Ilker Hacihaliloglu, and Vishal M Patel. Medical transformer: Gated axial-attention for medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 36–46. Springer, 2021.
- [22] Jeya Maria Jose Valanarasu, Vishwanath A Sindagi, Ilker Hacihaliloglu, and Vishal M Patel. KiU-Net: Overcomplete convolutional architectures for biomedical image and volumetric segmentation. *IEEE Transactions on Medical Imaging*, 41(4):965–976, 2021.
- [23] Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. MTI-Net: Multi-scale task interaction networks for multi-task learning. In *Proc. European Conference on Computer Vision*, pages 527–543. Springer, 2020.
- [24] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017.

- [25] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-DeepLab: Stand-alone axial-attention for panoptic segmentation. In *Proc. European Conference on Computer Vision*, pages 108–126. Springer, 2020.
- [26] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Minghui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(10):3349–3364, 2020.
- [27] Xiao Xiao, Shen Lian, Zhiming Luo, and Shaozi Li. Weighted Res-UNet for high-quality retina vessel segmentation. In *Proc. Information Technology in Medicine and Education*, pages 327–331. IEEE, 2018.
- [28] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. PAD-Net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 675–684, 2018.
- [29] Guoping Xu, Xingrong Wu, Xuan Zhang, and Xinwei He. LeViT-UNet: Make faster encoders with transformer for medical image segmentation. *arXiv preprint arXiv:2107.08623*, 2021.
- [30] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. UNet++: Redesigning skip connections to exploit multiscale features in image segmentation. *IEEE Transactions on Medical Imaging*, 39(6):1856–1867, 2019.
- [31] Hasib Zunair and A Ben Hamza. Sharp U-Net: Depthwise convolutional network for biomedical image segmentation. *Computers in Biology and Medicine*, 136:104699, 2021.