

Self-adversarial Multi-scale Contrastive Learning for Semantic Segmentation of Thermal Facial Images

Jitesh Joshi

jitesh.joshi.20@ucl.ac.uk

Nadia Bianchi-Berthouze

nadia.berthouze@ucl.ac.uk

Youngjun Cho*

youngjun.cho@ucl.ac.uk

Department of Computer Science

University College London

London, UK

Abstract

Segmentation of thermal facial images is a challenging task. This is because facial features often lack salience due to high-dynamic thermal range scenes and occlusion issues. Limited availability of datasets from unconstrained settings further limits the use of the state-of-the-art segmentation networks, loss functions and learning strategies which have been built and validated for RGB images. To address the challenge, we propose Self-Adversarial Multi-scale Contrastive Learning (SAM-CL) framework as a new training strategy for thermal image segmentation. SAM-CL framework consists of a SAM-CL loss function and a thermal image augmentation (TiAug) module as a domain-specific augmentation technique. We use the Thermal-Face-Database to demonstrate effectiveness of our approach. Experiments conducted on the existing segmentation networks (UNET, Attention-UNET, DeepLabV3 and HRNetv2) evidence the consistent performance gains from the SAM-CL framework. Furthermore, we present a qualitative analysis with UBCComfort and DeepBreath datasets to discuss how our proposed methods perform in handling unconstrained situations.

1 Introduction

Thermal infrared imaging of human skin enables remote physiological sensing and affective, psychological states monitoring [1, 2, 3]. Studies have indicated temperature patterns over specific facial regions as important psychophysiological signatures. For instance, temperature changes over the nostril region can be converted into breathing signals [4, 5, 6]. Another example is the nosetip temperature pattern associated with vasomotor activity which is related to mental stress states [7, 8]. Automated computational pipelines for processing thermal images therefore require identification of regions of interest (ROIs) which are either defined by a fixed bounding box or by an anatomical mask, with the latter being more appropriate for reliable extraction of physiological signals [9]. The automated identification of an anatomical mask requires every pixel in a thermal image to be labelled according to its respective anatomical region.

Such semantic segmentation task is particularly complex in thermal imaging. In comparison with typical RGB or gray-scale facial images, thermal facial images possess much less prominent facial features. This is because thermal images represent the temperature distribution over the skin surface, which is affected by dynamically varying physiological state as well as ambient temperature [13, 23]. In addition, variations over the thermal surface are very low and occlusions such as forehead hairs and eye-glasses further make it challenging to reliably segment the ROIs in unconstrained settings.

The state-of-the-art in semantic segmentation has been under continuous progression following the foundational work of Fully Convolutional Networks (FCN) [39]. The methodological contributions in the research of semantic segmentation can be broadly categorized into i) model architectures, ii) loss functions or learning strategies, and iii) data augmentation techniques. Significant development has been made towards deep-learning architectures for semantic segmentation with some notable ones such as: UNET [48] and its variants [50], the family of DeepLab networks [6, 7, 8], HRNet [53], and the more recent transformer based approaches such as “HRNet + OCR + SegFix” [59]. Pretrained network backbones such as ResNet [26], Xception [17], and HRNet [53], have further accelerated the progress owing to the availability of large-scale RGB datasets [19, 27, 37]. On the other side, the widely used loss functions for semantic segmentation include softmax cross-entropy loss [18], DICE loss [67] and region mutual information (RMI) loss [63], among others.

The existing challenges in semantic segmentation of thermal images include the lack of availability of large-scale bench-marking datasets. Furthermore, the studies validating the effectiveness of the segmentation networks, data-augmentation techniques and loss functions for RGB images have not sufficiently addressed the segmentation challenges of thermal images acquired in unconstrained settings. Specifically, data-augmentation techniques developed for RGB images [6, 49] do not consider thermal ambient conditions and therefore not suitable to augment thermal images. In addition, thermal data is single channel, and the basic properties such as transparency in RGB change to opacity in thermal images (e.g. clear glass). As thermal infrared wavelength is not transmissive for most of the objects, occlusions are observed more frequently. In addition, the variations in thermal ambient conditions result in varying appearances and can not be related to the brightness variations in RGB images [14].

This work addresses the challenge of training segmentation network with datasets of limited size using a novel self-adversarial multi-scale contrastive learning (SAM-CL) framework (§3). SAM-CL framework introduces a SAM-CL loss function (§3.1.2) and a thermal image augmentation (TiAug) module (§3.2), while utilizing existing segmentation networks. The TiAug module serves as domain specific augmentation, while the SAM-CL loss function provides enhanced supervision in learning inter-class separation and intra-class proximity in the presence of adversarial-attacks by TiAug. We compare the performance of SAM-CL framework with the existing segmentation loss functions, supervised contrastive learning (CL) [54] and Generative Adversarial Network (GAN) based approaches for segmentation [57, 62] in §4.1.2. Our contributions are:

- Self-Adversarial Multi-scale Contrastive-Learning (SAM-CL) framework that introduces following to train existing segmentation networks:
 - Self-adversarial multi-scale contrastive loss function for semantic segmentation, to efficiently achieve intra-class proximity and inter-class separation of the feature representations by utilizing the adversarial attacks generated by the TiAug module.
 - a Thermal image augmentation (TiAug) module to generate representations of unconstrained thermal settings by applying domain-specific transformations to the thermal

images acquired in controlled settings.

- Performance benchmark with semantic segmentation of facial regions on Thermal Face Database [63, 64] as well as qualitative analysis on UBCComfort [45] and DeepBreath datasets [12] demonstrating the performance gains by the proposed techniques.

2 Related Work

Semantic Segmentation: Segmentation networks largely follow encoder-decoder schemes [39, 40, 43, 48, 65]. To learn cross-pixel dependencies, several models apply attention mechanisms [9] for semantic segmentation [22, 28, 29, 32, 44, 66]. Atrous convolutions along with pyramid pooling in *DeepLab* networks [6, 7, 8] enable learning of the multi-scale features. *HRNet*, as proposed in [63], shows performance gain on the semantic segmentation task, among other tasks, by making use of high-resolution and multilevel representations. Furthermore, the more recent development using transformer networks such as *HRNet + OCR + SegFix* has achieved competitive performances on multiple benchmarking datasets [59]. While increased network complexity and deeper layers prove effective for training models with large-scale RGB datasets, it is challenging to benefit from the same with a limited dataset size as typically observed in the case of thermal infrared imaging.

Existing semantic segmentation approaches for thermal images are not equipped to handle real-world scenarios, such as occlusion and varying thermal ambient conditions. One earlier study on occlusion removal in thermal imaging [63] proposes a modelling-based method using kernel principal component analysis for removing a specific occluding object (eyeglasses). This method requires the use of a registered color image to reconstruct the occluded thermal image, limiting its generalizability unless a large-scale dataset with pair of color-images and thermal-images is available. Large scale datasets allow capturing diverse representations, though acquiring a large scale dataset with thermal imaging and performing pixel wise annotations for semantic segmentation remains impractical. Furthermore, the thermal imaging datasets that are currently available with facial images have been acquired in highly controlled settings [10, 33, 64, 66]. It is therefore required to review the data augmentation techniques that can allow achieving robust performance in real-world scenarios.

Image Augmentation Techniques: The commonly used augmentation techniques include geometric transformations as well as learning or modelling based methods [49]. While geometric transformations are relevant for thermal images, augmentation techniques pertaining to variations in brightness and contrast in RGB images cannot be directly mapped to thermal images. Among the learning based methods, GAN [24] and self-adversarial training (SAT) [25] have shown promising performance. SimGAN as proposed in [60] utilizes simulator generated images and a GAN to synthesize realistic augmented eye images. This method relies on the effectiveness of a simulator in synthesizing images, which may not generalize for different scenes and image modalities. YOLOv4 [9] showed the effectiveness of SAT based augmentation technique called Fast Gradient Sign Method (FGSM) [25] in which an original image gets updated instead of the network weights in one forward pass, and this altered image is then used as an adversarial attack to improve the robustness of the trained model. A more recent work on localization of image forgery [67] also highlights the usefulness of FGSM based self-adversarial attacks in augmenting the data. While existing SAT approaches increase the robustness of the model for subtle changes in an image, they are insufficient in modelling range of real-world scenarios. Unlike gradient based update of images in SAT approaches, our proposed method (§3.2) models plausible variations of unconstrained settings

in thermal images for adversarial attacks to enhance the robustness of the trained model in such settings.

Loss Functions or Learning Strategies: In addition to segmentation network and augmentation techniques, loss function or learning strategy plays a crucial role in achieving higher performance. For semantic segmentation tasks, cross-entropy loss [13] and weighted cross-entropy loss functions have been widely used [2, 61]. In cases of unequal class distribution, due to imbalanced distribution of pixels between semantic classes, focal loss [58] and DICE loss functions [62] have been reported to show better performance. In a more recent development, researchers proposed a loss function based on mutual information between pixels and semantic regions [63], and showed substantial improvements on benchmarking datasets. Learning strategies such as GAN [10, 20, 24, 57, 60, 62] and CL [52, 61, 64] have also shown to be effective for the segmentation task. A recent study using GAN [42] proposes multi-class segmentation approach, though, it is limited to handle only the occluding objects learnt at training time. The corner stone for the success of GAN as well as CL approaches is the availability of large scale datasets, which limits their effective deployment in thermal imaging. Our work takes inspiration from GAN as well as CL, however unlike the critic network in GAN, the auxiliary network in our approach does not compete with the generator (segmentation network). In addition, unlike the use of feature space for sampling anchors in CL, our approach uses predicted segmentation masks or logits as anchors.

3 Proposed Method: Self-Adversarial Multi-scale Contrastive Learning (SAM-CL)

Figure 1 provides an overview of our proposed Self-Adversarial Multi-scale Contrastive Learning (SAM-CL) framework. One of the highlights of our framework is that it is used only during the training, resulting in no computation overhead during the inference. Though

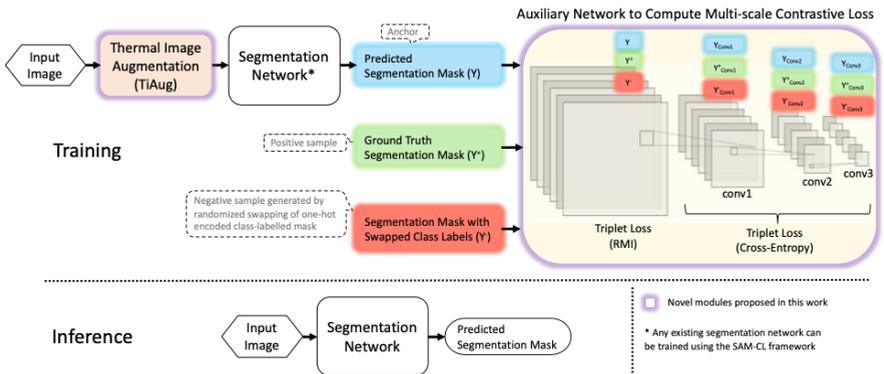


Figure 1: Proposed SAM-CL framework for Semantic Segmentation. In SAM-CL framework, the multi-scale contrastive-loss is computed by the four layer auxiliary network. TiAug generates augmented thermal images representing real-world scenarios.

this framework can be generically applied to segmentation tasks, the key objective in this work is to train a segmentation network on thermal facial images for: i) segmentation of facial regions including eyes, eyebrows, nose, mouth and chin area and, ii) resilience to varying

thermal ambient conditions as well as occlusions in unconstrained settings. The proposed SAM-CL framework utilizes existing segmentation networks and introduces a SAM-CL loss function and a thermal image augmentation (TiAug) module to achieve the objective without requiring thermal images acquired in unconstrained scenarios.

3.1 Loss Function

3.1.1 Preliminaries

For a semantic segmentation task, the segmentation network SEG learns a function $f_{SEG}(I)$ that maps input image I to the ground-truth mask Y , that specifies a semantic class $c \in C$ for every pixel $i \in I$. A limitation of the most commonly used pixel-wise segmentation loss functions such as cross-entropy loss, is their inability to capture the relationships between pixels. To address this limitation, a recent work proposed the mutual information based loss function [63] that combines cross-entropy and structural similarity loss. Furthermore, to learn the relationship between pixels of multiple images and to supervise the representations within pixel-embedding, supervised contrastive loss for semantic segmentation is proposed in [64]:

$$\mathcal{L}_y^{NCE} = \frac{1}{|P_y|} \sum_{y^+ \in P_y} -\log \frac{\exp(y \cdot y^+ / \tau)}{\exp(y \cdot y^+ / \tau) + \sum_{y^- \in N_y} \exp(y \cdot y^- / \tau)} \quad (1)$$

where P_y and N_y are positive and negative samples of pixel-embedding, belonging to classes dissimilar to that of the anchor pixel y . This learning strategy is very effective in maximizing inter-class separation, while minimizing intra-class distance within pixel-embedding, specially when large-scale dataset and corresponding pretrained weights are available. However, feature representations within pixel-embeddings remain transient while training a segmentation network without pretrained weights and with a limited dataset size. This limits the effectiveness of CL in training the segmentation network. To address this, we resort to a CL strategy that uses logits instead of pixel-embedding, while remaining effective in maximizing inter-class separation and minimizing intra-class proximity.

3.1.2 SAM-CL Loss Function

In a one-hot encoded ground-truth segmentation mask (Y_{oh}^+), each channel represents a binary mask for the respective classes. Class swapped mask (Y_{oh}^-) is generated by randomized swapping of channels of the (Y_{oh}^+) with the constraint that no channels of Y_{oh}^+ and Y_{oh}^- match. With logits or one-hot predicted mask Y_{oh} representing an anchor, Y_{oh}^+ and Y_{oh}^- representing positive and negative samples respectively, the first triplet loss is computed as shown in Equation (2):

$$\mathcal{L}_{s0}(Y_{oh}, Y_{oh}^+, Y_{oh}^-) = \max\{d(Y_{oh}, Y_{oh}^+) - d(Y_{oh}, Y_{oh}^-) + \text{margin}, 0\} \quad (2)$$

Equation (2) allows learning inter-class separation as well as intra-class proximity without requiring to compute the contrastive loss with pixel-embedding. As Y_{oh}^- preserves spatial features at mask-level, the optimization results in effective inter-class separation of the spatial features. Y_{oh} , Y_{oh}^+ , and Y_{oh}^- are passed through a 4-layered auxiliary network in three different forward passes to compute the feature maps y_{Comv1} , y_{Comv2} and y_{Comv3} ; ($y = Y_{oh}, Y_{oh}^+, Y_{oh}^-$) for each layer. Down-scaling of 2 is applied at each layer with the number of channels in every layer held constant and equal to the number of classes, consistent with the first layer input

channels.

$$\mathcal{L}_{SAM-CL} = \mathcal{L}_{s0}(Y_{oh}, Y_{oh}^+, Y_{oh}^-) + \mathcal{L}_{s1}(Y_{Conv1}, Y_{Conv1}^+, Y_{Conv1}^-) + \mathcal{L}_{s2}(Y_{Conv2}, Y_{Conv2}^+, Y_{Conv2}^-) + \mathcal{L}_{s3}(Y_{Conv3}, Y_{Conv3}^+, Y_{Conv3}^-) \quad (3)$$

The overall SAM-CL loss function as formulated in Equation (3), therefore offers supervision to maximize inter-class separation at multiple-scales. To compute the distance function $d(x, y)$ as mentioned in Equation (2), we use RMI [63] on logits, and cross-entropy loss for down-convolved feature-maps in the auxiliary network (see Figure 1).

3.2 Thermal Image Augmentation (TiAug) Module

The thermal image augmentation module (TiAug), as illustrated in Figure 2, transforms a thermal image acquired in controlled settings into an image resembling one acquired in unconstrained ambient settings. This is an important step as there often exist high-dynamic thermal range scenes in the real world settings [13]. Inspired by the Optimal Quantisation technique in [13], this module is designed to first add synthesized objects with diverse parameters in an occluding as well as a non-occluding manner. These parameters include size, shape, temperature, position and configuration (i.e. single, dual or dual-connected objects). In addition, a random temperature value is added as thermal noise to every pixel. The maximum magnitude of the noise is set as per the noise equivalent temperature difference (NETD), a sensitivity parameter of thermal infrared imaging camera, that provides the minimum value of temperature difference that can be sensed reliably by a camera. While a high-sensitive thermal camera has lower magnitude of NETD, it is higher for the low-cost mobile thermal imaging camera. TiAug sets the maximum NETD value ($Th_{NETD}^{max} = 0.1^\circ\text{C}$) considering the low-cost thermal imaging camera.

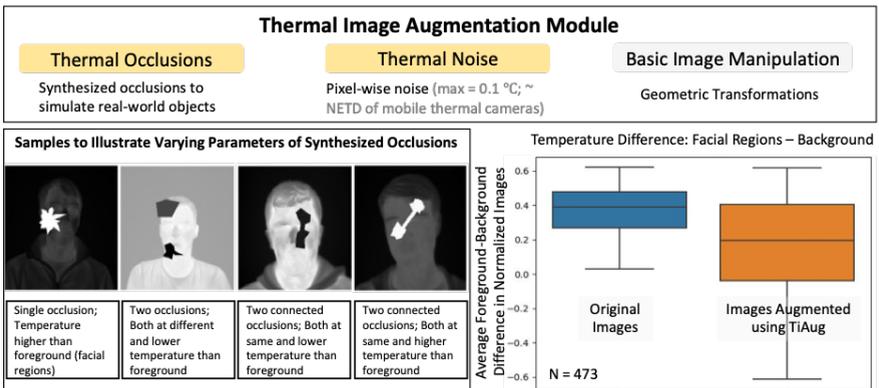


Figure 2: Thermal Image Augmentation (TiAug) Module. TiAug module consists of modules for thermal data specific occlusion generation as well as additive thermal noise modules along with the commonly used geometric transformation modules.

As expressed in Equation (4), an augmented image ($I_{aug}^{H \times W}$) is generated from an original image ($I_{org}^{H \times W}$) of height (H) and width (W) by applying an occlusion transformation (f_{occ}) and adding thermal noise ($\eta^{H \times W}$: $0 < \kappa < Th_{NETD}^{max} \forall \kappa \in \eta^{H \times W}$). The parameters characterizing occluding objects in f_{occ} are size (ϑ_{sz} : 1-40% of facial region), shape (ϑ_{sh}) characterized by

number of vertices (2-30) and the corresponding irregularity (0 to 10%) as well as spikeness (0 to 50%), temperature ($\vartheta_{temp}: I_{org}^{min} - (I_{org}^{max} - I_{org}^{min}) \leq \vartheta_{temp} \leq I_{org}^{max} + (I_{org}^{max} - I_{org}^{min})$), position ($\vartheta_{xy}: (x, y) \in R^{H \times W}$), and configurations (ϑ_{config} : single, dual and dual-connected). The function $g(\dots)$ uniquely combines the values of each parameter from the respective range, allowing the synthesis of endless new representations. f_{occ} replaces the specific pixels of $I_{org}^{H \times W}$ with the synthesized objects which may or may not occlude the facial regions. This leads to real-world histogram distribution of thermal images, unlike bimodal distribution when thermal images are acquired in controlled settings. In a bimodal histogram distribution, a peak at a lower temperature value depicts background and a peak at a higher temperature corresponds to the facial regions. Several segmentation algorithms assume bimodal histogram distribution for an automated segmentation [30], which affects their performance in unconstrained settings.

$$I_{aug}^{H \times W} = f_{occ}(I_{org}^{H \times W}, g(\vartheta_{sz}, \vartheta_{sh}, \vartheta_{temp}, \vartheta_{xy}, \vartheta_{config})) + \eta I^{H \times W} \quad (4)$$

Real-world scenarios may include objects at temperatures higher than the facial regions (e.g., sun, hot beverages), as well as objects at lower temperatures [13]. In such real-world scenarios, histogram distribution is not generally bimodal. To simulate real-world variations in histogram distribution pattern, TiAug adds synthesized objects both at temperatures higher and lower than that of the facial regions. This approach prevents the deep-learning network from over-fitting to bimodal-distributions. Hence, the obtained $I_{aug}^{H \times W}$ represents real-world scenarios both in terms of spatial characteristics of ambient objects as well as the histogram distribution of temperature values. $I_{aug}^{H \times W}$ is further normalized and passed as input to the segmentation network. While the examples in the Figure 2 show changes in spatial characteristics, the box-plot analysis signifies the effectiveness of TiAug in altering average foreground (facial-regions) temperature and average background temperature in the normalized images. Furthermore, TiAug applies the following geometric transformations: horizontal flip, vertical flip, rotation, Gaussian blur, and resizing (0.5X to 2X).

4 Experiments

We perform various experiments to compare the proposed SAM-CL framework with existing loss functions and learning strategies, using the following segmentation networks: U-NET [43], Attention UNET [42], DeepLabV3 [9] and HRNet [53]. Our code uses PyTorch [46] and is available on [GitHub](#). It is built upon a prior work on contrastive learning for semantic segmentation [54]. Given a lack of benchmark segmentation performance reports on thermal facial datasets as well as pretrained models, we implement and train the aforementioned prior-art segmentation networks with Xavier uniform initialization. We use a batch size of 16 along with an SGD optimiser with a weight decay of $1e-8$ and betas set to 0.9 and 0.999.

For our experiments, it is required to use datasets of raw thermal matrices (ie. absolute temperature value assigned to each pixel). The available datasets of thermal facial images include: Thermal Face Database [33, 35], UBComfort dataset [45], and DeepBreath dataset [12, 13]. For training and quantitative evaluation (§4.2), we mainly use the Thermal Face Database [33, 35] as it provides, together with the data, the ground-truth labels. A limitation of this dataset is that it is acquired in controlled setting. For this reason, we have use the other two datasets collected in unconstrained settings. However, due to the lack of ground truth in these two datasets, we use them for qualitative analysis only (§4.2) to demonstrate the effectiveness of SAM-CL framework in such settings.

4.1 Quantitative Evaluation on Thermal Face Database

4.1.1 Dataset Description

Thermal Face Database [33, 35] consists of 2935 images of 90 individuals with 68 manually annotated facial landmark points. We derive segmentation masks from the landmarks points for each anatomical region (chin, mouth, nose, eyes and eye-brows). Data is split into training (85%) and validation (15%) sets based on subject ids.

4.1.2 Results

To investigate the generalizability of the SAM-CL framework, we train UNET [48], Attention UNET [44], DeepLabV3 [9] and HRNet [53] segmentation networks. In all the experiments, augmentation with basic geometric transformations including horizontal flip, vertical flip, rotation, Gaussian blur, and resizing (0.5X to 2X), is uniformly applied. The loss functions used for bench-marking includes weighted binary cross-entropy loss (BCE), DICE loss, and region mutual information (RMI) loss [53]. As SAM-CL framework relates to CL and GAN, we additionally compare performance with SegAN [56, 57], SegGAN [62], along with a recent work on the supervised CL applied to semantic segmentation [64].

Table 1: Performance Evaluation of SAM-CL Framework

Segmentation Network	Learning Strategy (Loss Function)	mIoU (%)	Segmentation Network	Learning Strategy (Loss Function)	mIoU (%)
UNET [48]	Pixel-wise Segmentation (BCE)	67.64	Attention UNET [44]	Pixel-wise Segmentation (BCE)	66.61
	Pixel-wise Segmentation (DICE)	75.00		Pixel-wise Segmentation (DICE)	75.14
	GAN (SegAN) [56]	76.79		GAN (SegAN) [56]	76.75
	GAN (SegGAN) [62]	75.50		GAN (SegGAN) [62]	76.24
	RMI [53]	81.35		RMI [53]	81.39
	ContrastiveSeg [64]	81.24		ContrastiveSeg [64]	81.50
	SAM-CL (Ours)	82.11 (+0.76)		SAM-CL (Ours)	82.85 (+1.35)
DeepLabV3+ResNet101 [9, 44]	RMI [53]	75.85	HRNetV2-W48 [53]	RMI [53]	78.46
	ContrastiveSeg [64]	74.45	ContrastiveSeg [64]	78.36	
	SAM-CL (Ours)	79.29 (+3.44)	SAM-CL (Ours)	78.97 (+0.61)	

Table 1 shows the comparison of performances with percentage mean IoU metric. We observe consistent performance gains with the use of DICE loss when compared against BCE loss. GAN based learning strategy [56, 57] is found effective for UNET and Attention UNET, while the performance drops for DeepLabV3 network when comparing against the respective performance with the DICE loss. Consistent performance improvements from DICE loss function as well as GAN based learning strategy is evident for the models trained with RMI loss function. The performance gains across all the segmentation networks can be noted when deploying the SAM-CL framework.

Table 2: Ablation Study for TiAug and SAM-CL Loss Function

Segmentation Network	mIoU (%) Performance		
	RMI	RMI + TiAug	RMI + TiAug + SAM-CL
UNET [48]	81.36	81.91	82.11
Attention UNET [44]	81.39	82.29	82.85
HRNetV2-W48 [53]	78.13	78.87	78.97
DeepLabV3+ResNet101 [9, 44]	75.85	78.07	78.12
DeepLabV3+Xception [9, 44]	76.55	77.31	77.85

We performed an ablation study to examine the individual contribution of the TiAug module

and the SAM-CL loss function. From Table 2, we observe performance gains from the baseline, trained with RMI loss, for all the segmentation networks when the data is augmented using the TiAug module. Similarly additional performance gains for the respective segmentation networks are observed when the SAM-CL loss function is used along with the TiAug module. The TiAug module presents a segmentation network with the adversaries such as occlusions and varying ambient temperature levels in the input thermal images, while the SAM-CL loss function maximizes the inter-class separation using the class-swapped negative sample Y^- and its down-scaled representations in the auxiliary network.

4.2 Qualitative Analysis

As Thermal Face Database does not include real-world occlusions, we extend the evaluation of our approach with a qualitative analysis on the datasets acquired in unconstrained settings (see Figure 3). The UBComfort dataset [45] was acquired from in-the-wild car users with varying thermal ambient conditions using a high-resolution thermal camera. Instead, the

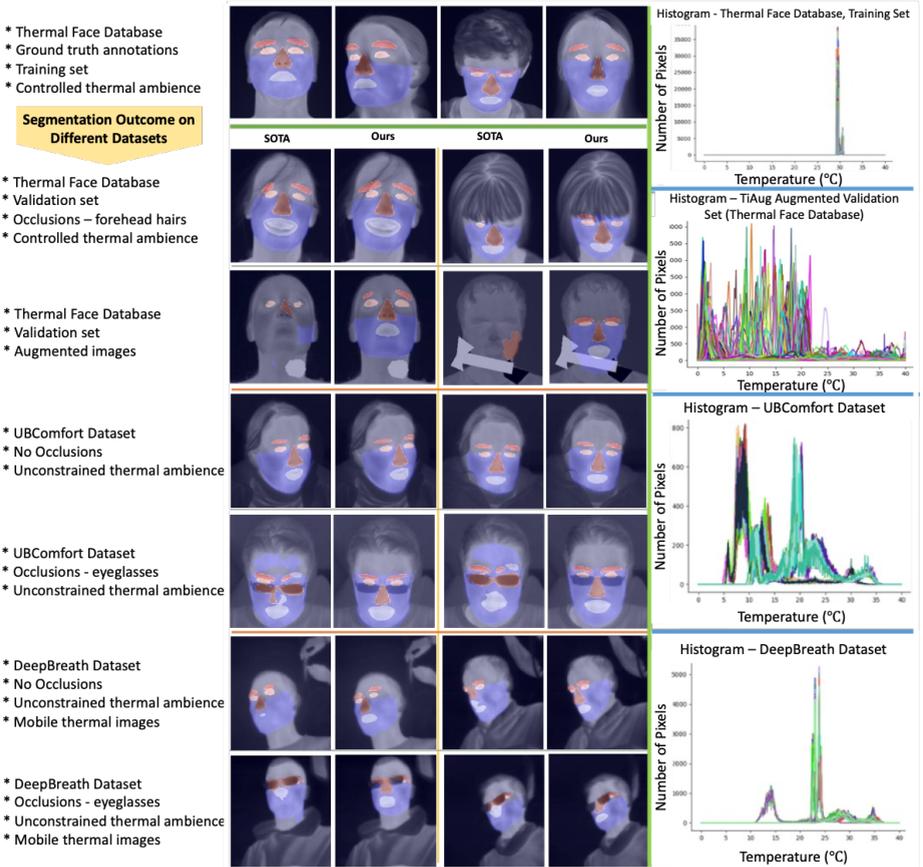


Figure 3: Qualitative Performance Analysis on Different Thermal Imaging Datasets. Attention UNET is trained using RMI loss to compare the performance of the proposed SAM-CL framework. Please refer to §4.2 for the discussion on qualitative performance analysis.

DeepBreath dataset [12] was acquired from participants under induced stress levels using a low-resolution mobile thermal camera.

The thermal images in the first row of the Figure 3 present ground-truth labelmask overlaid with color-coded class-labels. The histogram plot on the top-right shows the temperature distribution across the images in the Thermal Face database, highlighting the highly controlled laboratory settings. We further identify a few samples (row-2) within the Thermal Face Database in which hairs occlude a small part of the thermal image. The superiority of the model trained jointly with the SAM-CL loss function and the TiAug module in reliably handling forehead hairs occlusions is evidenced by the segmentation outcome. In the following row, we present samples generated by the TiAug module, which after min-max normalization, appear significantly different, as would be the case when objects that are either too hot or too cold appear in an image. The state-of-the-art (SOTA) model fails in these scenarios (row-3) as it has not been trained with such variations.

The thermal images of individuals without and with eye-glasses, seated in a car [15] are shown in rows 4 and 5 of Figure 3 respectively. Though the training set does not include such images, the model trained with our proposed method shows resilience towards performing reliable segmentation in the presence of eye-glasses, while the SOTA sub-performs. Similarly, the thermal images of individuals performing cognitive tasks [12], without and with eye-glasses, are shown in rows 6 and 7 respectively. While SOTA sub-performs on both cases, SAM-CL framework shows reliable performance. It can be noted that the thermal images in the DeepBreath dataset [12] are acquired using mobile thermal camera (FLIR One), highlighting the robustness of the model trained using SAM-CL framework across for different thermal camera specifications.

5 Conclusion

The proposed SAM-CL framework that introduces the SAM-CL loss function and the TiAug module, is shown to be effective in training segmentation networks with datasets of limited size. The TiAug module transforms thermal images acquired in controlled laboratory environment into ones representing real-world scenarios. This transformation considers a range of plausible ambient temperature, geometric properties of common occluding objects as well as noise specification of widely used thermal cameras. This makes the TiAug module suitable to be applied on thermal images for various learning-based computer-vision tasks including classification, object detection, instance and panoptic segmentation to train deep-learning networks to handle common real-world scenarios, without explicitly requiring a thermal dataset to be acquired in such scenarios.

Furthermore, the SAM-CL loss function benefits from the TiAug module that presents the segmentation network with adversaries (e.g. occluded images), resulting in a portion of predicted logits to overlap with the synthesized negative sample (from class-swapped label-mask). This overlap of logits with negative sample results in a higher loss for the incorrect label predicted by the segmentation network in the corresponding region. This explains the effectiveness of the SAM-CL loss function in conjunction with the TiAug module in offering consistent performance gains across different segmentation networks. SAM-CL loss function can be extended to other imaging modalities to train segmentation networks with datasets of limited size, by devising appropriate augmentation technique for self-adversarial learning.

Acknowledgements: Jitesh Joshi was fully supported by Physiological Computing Studentship from UCL Computer Science (<https://github.com/PhysiologicAILab>).

References

- [1] Madina Abdrakhmanova, Askat Kuzdeuov, Sheikh Jarju, Yerbolat Khassanov, Michael Lewis, and Huseyin Atakan Varol. SpeakingFaces: A Large-Scale Multimodal Dataset of Voice Commands with Visual and Thermal Video Streams. *Sensors*, 21(10): 3465, January 2021. doi: 10.3390/s21103465. URL <https://www.mdpi.com/1424-8220/21/10/3465>. ZSCC: 0000004 Number: 10 Publisher: Multidisciplinary Digital Publishing Institute.
- [2] Saeid Asgari Taghanaki, Kumar Abhishek, Joseph Paul Cohen, Julien Cohen-Adad, and Ghassan Hamarneh. Deep semantic segmentation of natural and medical images: a review. *Artificial Intelligence Review*, 54(1):137–178, January 2021. ISSN 1573-7462. doi: 10.1007/s10462-020-09854-1. URL <https://doi.org/10.1007/s10462-020-09854-1>. ZSCC: 0000115.
- [3] Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. Multiple Object Recognition with Visual Attention. *arXiv:1412.7755 [cs]*, April 2015. URL <http://arxiv.org/abs/1412.7755>. arXiv: 1412.7755.
- [4] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. YOLOv4: Optimal Speed and Accuracy of Object Detection. *http://arxiv.org/abs/2004.10934*, April 2020. URL <http://arxiv.org/abs/2004.10934>.
- [5] Alexander Buslaev, Vladimir I. Iglovikov, Eugene Khvedchenya, Alex Parinov, Mikhail Druzhinin, and Alexandr A. Kalinin. Albumentations: Fast and Flexible Image Augmentations. *Information*, 11(2):125, February 2020. ISSN 2078-2489. doi: 10.3390/info11020125. URL <https://www.mdpi.com/2078-2489/11/2/125>. Number: 2 Publisher: Multidisciplinary Digital Publishing Institute.
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *arXiv:1606.00915 [cs]*, May 2017. URL <http://arxiv.org/abs/1606.00915>. arXiv: 1606.00915.
- [7] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking Atrous Convolution for Semantic Image Segmentation. *arXiv:1706.05587 [cs]*, December 2017. URL <http://arxiv.org/abs/1706.05587>. ZSCC: 0003216 arXiv: 1706.05587.
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, April 2018. ISSN 1939-3539. doi: 10.1109/TPAMI.2017.2699184. ZSCC: 0000003 Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.

- [9] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. pages 801–818, 2018. URL https://openaccess.thecvf.com/content_ECCV_2018/html/Liang-Chieh_Chen_Encoder-Decoder_with_Atrous_ECCV_2018_paper.html.
- [10] Anoop Cherian and Alan Sullivan. Sem-GAN: Semantically-Consistent Image-to-Image Translation. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1797–1806, January 2019. doi: 10.1109/WACV.2019.00196. ZSCC: 0000029 ISSN: 1550-5790.
- [11] Youngjun Cho and Nadia Bianchi-Berthouze. Physiological and Affective Computing through Thermal Imaging: A Survey. *arXiv:1908.10307 [cs]*, August 2019. URL <http://arxiv.org/abs/1908.10307>. arXiv: 1908.10307.
- [12] Youngjun Cho, Nadia Bianchi-Berthouze, and Simon J. Julier. DeepBreath: Deep learning of breathing patterns for automatic stress recognition using low-cost thermal imaging in unconstrained settings. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 456–463, October 2017. doi: 10.1109/ACII.2017.8273639. ISSN: 2156-8111.
- [13] Youngjun Cho, Simon J. Julier, Nicolai Marquardt, and Nadia Bianchi-Berthouze. Robust tracking of respiratory rate in high-dynamic range scenes using mobile thermal imaging. *Biomedical Optics Express*, 8(10):4480–4503, October 2017. ISSN 2156-7085. doi: 10.1364/BOE.8.004480. URL <https://www.osapublishing.org/boe/abstract.cfm?uri=boe-8-10-4480>. Number: 10 Publisher: Optical Society of America.
- [14] Youngjun Cho, Nadia Bianchi-Berthouze, Nicolai Marquardt, and Simon J. Julier. Deep Thermal Imaging: Proximate Material Type Recognition in the Wild through Deep Learning of Spatial Surface Temperature Patterns. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pages 1–13, New York, NY, USA, April 2018. Association for Computing Machinery. ISBN 978-1-4503-5620-6. doi: 10.1145/3173574.3173576. URL <https://doi.org/10.1145/3173574.3173576>.
- [15] Youngjun Cho, Nadia Bianchi-Berthouze, Manuel Oliveira, Catherine Holloway, and Simon Julier. Nose Heat: Exploring Stress-induced Nasal Thermal Variability through Mobile Thermal Imaging. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 566–572, September 2019. doi: 10.1109/ACII.2019.8925453. ISSN: 2156-8111.
- [16] Youngjun Cho, Simon J. Julier, and Nadia Bianchi-Berthouze. Instant Stress: Detection of Perceived Mental Stress Through Smartphone Photoplethysmography and Thermal Imaging. *JMIR Mental Health*, 6(4):e10140, 2019. doi: 10.2196/10140. URL <https://mental.jmir.org/2019/4/e10140/>. Number: 4 Company: JMIR Mental Health Distributor: JMIR Mental Health Institution: JMIR Mental Health Label: JMIR Mental Health Publisher: JMIR Publications Inc., Toronto, Canada.
- [17] Francois Chollet. Xception: Deep Learning With Depthwise Separable Convolutions. pages 1251–1258, 2017. URL <https://openaccess.thecvf.com/>

content_cvpr_2017/html/Chollet_Xception_Deep_Learning_CVPR_2017_paper.html.

- [18] Pieter-Tjerk de Boer, Dirk P. Kroese, Shie Mannor, and Reuven Y. Rubinfeld. A Tutorial on the Cross-Entropy Method. *Annals of Operations Research*, 134(1):19–67, February 2005. ISSN 1572-9338. doi: 10.1007/s10479-005-5724-z. URL <https://doi.org/10.1007/s10479-005-5724-z>. ZSCC: 0002154.
- [19] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, June 2009. doi: 10.1109/CVPR.2009.5206848. ISSN: 1063-6919.
- [20] Xue Dong, Yang Lei, Tonghe Wang, Matthew Thomas, Leonardo Tang, Walter J. Curran, Tian Liu, and Xiaofeng Yang. Automatic multiorgan segmentation in thorax CT images using U-net-GAN. *Medical Physics*, 46(5):2157–2168, 2019. ISSN 2473-4209. doi: 10.1002/mp.13458. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mp.13458>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mp.13458>.
- [21] A. Duarte, L. Carrão, M. Espanha, T. Viana, D. Freitas, P. Bártolo, P. Faria, and H. A. Almeida. Segmentation Algorithms for Thermal Images. *Procedia Technology*, 16:1560–1569, January 2014. ISSN 2212-0173. doi: 10.1016/j.protcy.2014.10.178. URL <http://www.sciencedirect.com/science/article/pii/S2212017314004058>. ZSCC: 0000042.
- [22] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual Attention Network for Scene Segmentation. pages 3146–3154, 2019. URL https://openaccess.thecvf.com/content_CVPR_2019/html/Fu_Dual_Attention_Network_for_Scene_Segmentation_CVPR_2019_paper.html. ZSCC: 0001493.
- [23] Rikke Gade and Thomas B. Moeslund. Thermal cameras and applications: a survey. *Machine Vision and Applications*, 25(1):245–262, January 2014. ISSN 1432-1769. doi: 10.1007/s00138-013-0570-5. URL <https://doi.org/10.1007/s00138-013-0570-5>. ZSCC: 0000492.
- [24] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL <https://proceedings.neurips.cc/paper/2014/hash/5ca3e9b122f61f8f06494c97b1afccf3-Abstract.html>. ZSCC: 0035398.
- [25] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples, March 2015. URL <http://arxiv.org/abs/1412.6572>. arXiv:1412.6572 [cs, stat].
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. pages 770–778, 2016. URL https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html. ZSCC: 0109350.

- [27] Derek Hoiem, Santosh K. Divvala, and James H. Hays. Pascal VOC 2008 Challenge, 2009. ZSCC: 0000034.
- [28] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-Excitation Networks. pages 7132–7141, 2018. URL https://openaccess.thecvf.com/content_cvpr_2018/html/Hu_Squeeze-and-Excitation_Networks_CVPR_2018_paper. ZSCC: 0008551.
- [29] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. CCNet: Criss-Cross Attention for Semantic Segmentation. pages 603–612, 2019. URL https://openaccess.thecvf.com/content_ICCV_2019/html/Huang_CCNet_Criss-Cross_Attention_for_Semantic_Segmentation_ICCV_2019_paper.html. ZSCC: 0000643.
- [30] Akshay Isalkar and K. Manikandan. Analysis of Image Segmentation Algorithms for Infrared Images. In Amit Dhawan, Vijay Shanker Tripathi, Karm Veer Arya, and Kshirasagar Naik, editors, *Recent Trends in Electronics and Communication*, Lecture Notes in Electrical Engineering, pages 639–646, Singapore, 2022. Springer. ISBN 9789811627613. doi: 10.1007/978-981-16-2761-3_57.
- [31] Shruti Jadon. A survey of loss functions for semantic segmentation. In *2020 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB)*, pages 1–7, October 2020. doi: 10.1109/CIBCB48159.2020.9277638. ZSCC: 0000165.
- [32] Qiangguo Jin, Zhaopeng Meng, Changming Sun, Hui Cui, and Ran Su. RA-UNet: A Hybrid Deep Attention-Aware Network to Extract Liver and Tumor in CT Scans. *Frontiers in Bioengineering and Biotechnology*, 0, 2020. ISSN 2296-4185. doi: 10.3389/fbioe.2020.605132. URL <https://www.frontiersin.org/articles/10.3389/fbioe.2020.605132/full>. ZSCC: 0000073 Publisher: Frontiers.
- [33] M. Kopaczka, R. Kolk, and D. Merhof. A fully annotated thermal face database and its application for thermal facial expression recognition. In *2018 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, pages 1–6, May 2018. doi: 10.1109/I2MTC.2018.8409768. ZSCC: 0000021.
- [34] M. Kopaczka, R. Kolk, J. Schock, F. Burkhard, and D. Merhof. A Thermal Infrared Face Database With Facial Landmarks and Emotion Labels. *IEEE Transactions on Instrumentation and Measurement*, 68(5):1389–1401, May 2019. ISSN 1557-9662. doi: 10.1109/TIM.2018.2884364. ZSCC: NoCitationData[s0] Conference Name: IEEE Transactions on Instrumentation and Measurement.
- [35] Marcin Kopaczka, Lukas Breuer, Justus Schock, and Dorit Merhof. A Modular System for Detection, Tracking and Analysis of Human Faces in Thermal Infrared Recordings. *Sensors*, 19(19):4135, January 2019. doi: 10.3390/s19194135. URL <https://www.mdpi.com/1424-8220/19/19/4135>. ZSCC: 0000003 Number: 19 Publisher: Multidisciplinary Digital Publishing Institute.
- [36] M. Kowalski and A. Grudzień. High-resolution thermal face dataset for face and expression recognition. *Metrology and Measurement Systems*, Vol. 25(nr 2), 2018. ISSN 0860-8229. doi: 10.24425/119566. URL

<http://yadda.icm.edu.pl/baztech/element/bwmeta1.element.baztech-d00dadcc-8eaf-4883-9b6e-ala66fd44ba6>. ZSCC: 0000006.

- [37] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common Objects in Context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, Lecture Notes in Computer Science, pages 740–755, Cham, 2014. Springer International Publishing. ISBN 978-3-319-10602-1. doi: 10.1007/978-3-319-10602-1_48.
- [38] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal Loss for Dense Object Detection. pages 2980–2988, 2017. URL https://openaccess.thecvf.com/content_iccv_2017/html/Lin_Focal_Loss_for_ICCV_2017_paper.html.
- [39] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. pages 3431–3440, 2015. URL https://openaccess.thecvf.com/content_cvpr_2015/html/Long_Fully_Convolutional_Networks_2015_CVPR_paper.html.
- [40] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 565–571, October 2016. doi: 10.1109/3DV.2016.79. ZSCC: 0004038.
- [41] R. Murthy and I. Pavlidis. Noncontact measurement of breathing function. *IEEE Engineering in Medicine and Biology Magazine*, 25(3):57–67, May 2006. ISSN 1937-4186. doi: 10.1109/MEMB.2006.1636352. Conference Name: IEEE Engineering in Medicine and Biology Magazine.
- [42] David Müller, Andreas Ehlen, and Bernd Valeske. Convolutional Neural Networks for Semantic Segmentation as a Tool for Multiclass Face Analysis in Thermal Infrared. *Journal of Nondestructive Evaluation*, 40(1):9, January 2021. ISSN 1573-4862. doi: 10.1007/s10921-020-00740-y. URL <https://doi.org/10.1007/s10921-020-00740-y>. ZSCC: 0000004.
- [43] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning Deconvolution Network for Semantic Segmentation. pages 1520–1528, 2015. URL https://openaccess.thecvf.com/content_iccv_2015/html/Noh_Learning_Deconvolution_Network_ICCV_2015_paper.html.
- [44] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y. Hammerla, Bernhard Kainz, Ben Glocker, and Daniel Rueckert. Attention U-Net: Learning Where to Look for the Pancreas. *arXiv:1804.03999 [cs]*, May 2018. URL <http://arxiv.org/abs/1804.03999>. ZSCC: 0000934 arXiv: 1804.03999.
- [45] Temitayo Olugbade, Youngjun Cho, Zak Morgan, Mohamed Abd El Ghani, and Nadia Bianchi-Berthouze. Toward Intelligent Car Comfort Sensing: New Dataset and Analysis of Annotated Physiological Metrics. In *2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 1–8, September 2021. doi: 10.1109/ACII52823.2021.9597393. ZSCC: 0000000 ISSN: 2156-8111.

- [46] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html>. ZSCC: NoCitationData[s0].
- [47] Carina Barbosa Pereira, Xinchu Yu, Michael Czaplik, Rolf Rossaint, Vladimir Blazek, and Steffen Leonhardt. Remote monitoring of breathing dynamics using infrared thermography. *Biomedical Optics Express*, 6(11):4378–4394, October 2015. ISSN 2156-7085. doi: 10.1364/BOE.6.004378. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4646547/>.
- [48] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Lecture Notes in Computer Science, pages 234–241, Cham, 2015. Springer International Publishing. ISBN 978-3-319-24574-4. doi: 10.1007/978-3-319-24574-4_28.
- [49] Connor Shorten and Taghi M. Khoshgoftaar. A survey on Image Data Augmentation for Deep Learning. *Journal of Big Data*, 6(1):60, July 2019. ISSN 2196-1115. doi: 10.1186/s40537-019-0197-0. URL <https://doi.org/10.1186/s40537-019-0197-0>.
- [50] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning From Simulated and Unsupervised Images Through Adversarial Training. In *IEEE Conference on CVPR*, pages 2107–2116, 2017. URL https://openaccess.thecvf.com/content_cvpr_2017/html/Shrivastava_Learning_From_Simulated_CVPR_2017_paper.html.
- [51] Nahian Siddique, Paheding Sidike, Colin Elkin, and Vijay Devabhaktuni. U-Net and its variants for medical image segmentation: theory and applications. *IEEE Access*, 9: 82031–82057, 2021. ISSN 2169-3536. doi: 10.1109/ACCESS.2021.3086020. URL <http://arxiv.org/abs/2011.01118>. ZSCC: 0000002 arXiv: 2011.01118.
- [52] Carole H. Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M. Jorge Cardoso. Generalised Dice Overlap as a Deep Learning Loss Function for Highly Unbalanced Segmentations. In M. Jorge Cardoso, Tal Arbel, Gustavo Carneiro, Tanveer Syeda-Mahmood, João Manuel R.S. Tavares, Mehdi Moradi, Andrew Bradley, Hayit Greenspan, João Paulo Papa, Anant Madabhushi, Jacinto C. Nascimento, Jaime S. Cardoso, Vasileios Belagiannis, and Zhi Lu, editors, *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, Lecture Notes in Computer Science, pages 240–248, Cham, 2017. Springer International Publishing. ISBN 978-3-319-67558-9. doi: 10.1007/978-3-319-67558-9_28. ZSCC: NoCitationData[s0].

- [53] Ke Sun, Yang Zhao, Borui Jiang, Tianheng Cheng, Bin Xiao, Dong Liu, Yadong Mu, Xinggang Wang, Wenyu Liu, and Jingdong Wang. High-Resolution Representations for Labeling Pixels and Regions. *arXiv:1904.04514 [cs]*, April 2019. URL <http://arxiv.org/abs/1904.04514>. ZSCC: 0000335 arXiv: 1904.04514.
- [54] Wenguan Wang, Tianfei Zhou, Fisher Yu, Jifeng Dai, Ender Konukoglu, and Luc Van Gool. Exploring Cross-Image Pixel Contrast for Semantic Segmentation. pages 7303–7313, 2021. URL https://openaccess.thecvf.com/content/ICCV2021/html/Wang_Exploring_Cross-Image_Pixel_Contrast_for_Semantic_Segmentation_ICCV_2021_paper.html.
- [55] W. K. Wong and Haitao Zhao. Eyeglasses removal of thermal image based on visible information. *Information Fusion*, 14(2):163–176, April 2013. ISSN 1566-2535. doi: 10.1016/j.inffus.2011.09.002. URL <https://doi.org/10.1016/j.inffus.2011.09.002>. ZSCC: 0000026.
- [56] Yuan Xue, Tao Xu, and Xiaolei Huang. Adversarial learning with multi-scale loss for skin lesion segmentation. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 859–863, April 2018. doi: 10.1109/ISBI.2018.8363707. ISSN: 1945-8452.
- [57] Yuan Xue, Tao Xu, Han Zhang, L. Rodney Long, and Xiaolei Huang. SegAN: Adversarial Network with Multi-scale L1 Loss for Medical Image Segmentation. *Neuroinformatics*, 16(3):383–392, October 2018. ISSN 1559-0089. doi: 10.1007/s12021-018-9377-x. URL <https://doi.org/10.1007/s12021-018-9377-x>.
- [58] Cheng Yang, Menghan Hu, Guangtao Zhai, and Xiao-Ping Zhang. Graph-based Denoising for Respiration and Heart Rate Estimation during Sleep in Thermal Video. *IEEE Internet of Things Journal*, pages 1–1, 2022. ISSN 2327-4662. doi: 10.1109/JIOT.2022.3150147. Conference Name: IEEE Internet of Things Journal.
- [59] Yuhui Yuan, Xiaokang Chen, Xilin Chen, and Jingdong Wang. Segmentation Transformer: Object-Contextual Representations for Semantic Segmentation, April 2021. URL <http://arxiv.org/abs/1909.11065>. arXiv:1909.11065 [cs].
- [60] Chaoyi Zhang, Yang Song, Sidong Liu, Scott Lill, Chenyu Wang, Zihao Tang, Yuyi You, Yang Gao, Alexander Klistorner, Michael Barnett, and Weidong Cai. MS-GAN: GAN-Based Semantic Segmentation of Multiple Sclerosis Lesions in Brain Magnetic Resonance Imaging. In *2018 Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–8, December 2018. doi: 10.1109/DICTA.2018.8615771. ZSCC: 0000016.
- [61] Feihu Zhang, Philip Torr, Rene Ranftl, and Stephan R. Richter. Looking Beyond Single Images for Contrastive Semantic Segmentation Learning. May 2021. URL <https://openreview.net/forum?id=MSV1SMBbBt>. ZSCC: 0000000.
- [62] Xinming Zhang, Xiaobin Zhu, 3rd Xiao-Yu Zhang, Naiguang Zhang, Peng Li, and Lei Wang. SegGAN: Semantic Segmentation with Generative Adversarial Network. In *2018 IEEE Fourth International Conference on Multimedia Big Data (BigMM)*, pages 1–5, September 2018. doi: 10.1109/BigMM.2018.8499105. ZSCC: 0000018.

- [63] Shuai Zhao, Yang Wang, Zheng Yang, and Deng Cai. Region Mutual Information Loss for Semantic Segmentation. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/a67c8c9a961b4182688768dd9ba015fe-Abstract.html>.
- [64] Xiangyun Zhao, Raviteja Vemulapalli, Philip Andrew Mansfield, Boqing Gong, Bradley Green, Lior Shapira, and Ying Wu. Contrastive Learning for Label Efficient Semantic Segmentation. pages 10623–10633, 2021. URL https://openaccess.thecvf.com/content/ICCV2021/html/Zhao_Contrastive_Learning_for_Label_Efficient_Semantic_Segmentation_ICCV_2021_paper.html.
- [65] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang. UNet++: Redesigning Skip Connections to Exploit Multiscale Features in Image Segmentation. *IEEE Transactions on Medical Imaging*, 39(6):1856–1867, June 2020. ISSN 1558-254X. doi: 10.1109/TMI.2019.2959609. Conference Name: IEEE Transactions on Medical Imaging.
- [66] Juntang Zhuang. LadderNet: Multi-path networks based on U-Net for medical image segmentation. *arXiv:1810.07810 [cs, eess]*, August 2019. URL <http://arxiv.org/abs/1810.07810>. ZSCC: NoCitationData[s0] arXiv: 1810.07810.
- [67] Long Zhuo, Shunquan Tan, Bin Li, and Jiwu Huang. Self-Adversarial Training Incorporating Forgery Attention for Image Forgery Localization. *IEEE Transactions on Information Forensics and Security*, 17:819–834, 2022. ISSN 1556-6021. doi: 10.1109/TIFS.2022.3152362.