

## Tilted cross-entropy (TCE)

### Promoting fairness in semantic segmentation

Szabo, Attila; Jamali-Rad, Hadi; Mannava, Siva Datta

#### DOI

[10.1109/CVPRW53098.2021.00261](https://doi.org/10.1109/CVPRW53098.2021.00261)

#### Publication date

2021

#### Document Version

Final published version

#### Published in

Proceedings - 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPRW 2021

#### Citation (APA)

Szabo, A., Jamali-Rad, H., & Mannava, S. D. (2021). Tilted cross-entropy (TCE): Promoting fairness in semantic segmentation. In *Proceedings - 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPRW 2021* (pp. 2305-2310). (IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops). IEEE. <https://doi.org/10.1109/CVPRW53098.2021.00261>

#### Important note

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

#### Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

#### Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.

# Tilted Cross-Entropy (TCE): Promoting Fairness in Semantic Segmentation

Attila Szabó<sup>1\*</sup> Hadi Jamali-Rad<sup>1,2\*</sup> Siva-Datta Mannava<sup>1</sup>

<sup>1</sup>Shell Global Solutions International B.V., Amsterdam, The Netherlands

<sup>2</sup>Delft University of Technology (TU Delft), Delft, The Netherlands

{attila.szabo,hadi.jamali-rad,siva-datta.mannava}@shell.com,h.jamalirad@tudelft.nl

## Abstract

*Traditional empirical risk minimization (ERM) for semantic segmentation can disproportionately advantage or disadvantage certain target classes in favor of an (unfair but) improved overall performance. Inspired by the recently introduced tilted ERM (TERM), we propose tilted cross-entropy (TCE) loss and adapt it to the semantic segmentation setting to minimize performance disparity among target classes and promote fairness. Through quantitative and qualitative performance analyses, we demonstrate that the proposed Stochastic TCE for semantic segmentation can offer improved overall fairness by efficiently minimizing the performance disparity among the target classes of Cityscapes.*

## 1. Introduction and Related Work

Semantic segmentation is a fundamental problem in computer vision and a pivotal step towards content-based image analysis and scene understanding. It has received an upsurge of attention recently owing to its wide variety of applications in medical imaging [23, 21], autonomous driving [15, 3], satellite image processing [26, 7], and robotics [5, 24], to name a few. Early segmentation methodologies are mostly developed with clustering algorithms at their core [10, 18, 20, 16]. Recent advances in deep learning have revolutionized this field resulting in state-of-the-art (SoTA) image segmentation algorithms such as DeepLabv3+ [1], PS and Panoptic DeepLab [11, 2], HRNet [27] and many other elegant architectures that considerably outperformed the traditional signal processing-based methods.

The choice of loss function is essential for applicability of semantic segmentation in different contexts. Therefore, several studies have investigated the impact of modified loss functions on the performance of semantic segmentation models in generic and custom-built datasets. A comprehensive list of some of these recent loss functions is provided in four different categories in [8]: i) distribution-based losses such as binary class cross-entropy (BCE), ii) region-based losses such as Dice loss [25] and its variants, iii) boundary based

losses such as Hausdorff distance [22, 9], and iv) compound losses such as Exponential Logarithmic loss [28]. Generic empirical risk minimization (ERM) loss functions such as BCE or Dice can disproportionately advantage or disadvantage some classes in favor of an improved average performance. This results in models that treat certain classes (e.g. with less presence in training dataset) in an *unfair* fashion. Somewhat related to this concern, significant research has been conducted on proposing alternatives to the cross-entropy (CE) loss for semantic segmentation to handle imbalanced data using weighted cross entropy [19], balanced cross entropy [30], and the Focal loss [14]. However, none of these approaches directly address the *fairness* problem. Being the most relevant approach among the mentioned solutions, Focal loss [14] down-weights the contribution of easy examples and enables the model to focus more on learning hard ones. In a broader scope of optimization for machine learning, [6] proposes a solution to ensure different subgroups within a population are treated fairly and [4] develops a solution with favorable out-of-sample performance. Most recently, tilted empirical risk minimization (TERM) has been proposed in [12] to flexibly address the deficiencies of traditional ERM with respect to handling outliers and treating subgroups fairly. It is demonstrated in [12] that TERM can not only efficiently promote fairness in a multitude of applications, but also outperforms the likes of Focal Loss [14] and RobustRegRisk [4] in different settings. Besides, [12] provides efficient batch and stochastic first-order optimization methods for solving TERM. Inspired by the flexibility of TERM in promoting fairness, we have studied the impact of adopting a similar approach in a new context, i.e., semantic segmentation. To reiterate, *fairness* in this context means a less varied performance across target segmentation classes.

Our contributions are as follows: i) We propose to employ tilted cross-entropy (TCE) as a novel loss to promote fairness in semantic segmentation. ii) We adapt the derivations of [12] to fit into semantic segmentation setting and reformulate the commonly used CE loss as TCE. We then propose a stochastic non-hierarchical optimization algorithm for solving TCE. iii) We demonstrate the effectiveness of Stochastic

\*Equal contributions.

TCE on promoting fairness in semantic segmentation for Cityscapes (and ADE20k in supplementary material).<sup>1</sup>

## 2. Tilted CE (TCE) for Semantic Segmentation

In this section, we introduce TCE: an adapted and reformulated version of TERM [12] for semantic segmentation. We then propose a stochastic non-hierarchical batch solution of TCE. In the following, we denote the cardinality of  $\mathcal{X}$  as  $|\mathcal{X}|$ , and the set  $\{1, \dots, n\}$  as  $[n]$ . Let  $\mathcal{D}_t = \{(\mathbf{X}, \mathbf{Y})_1, \dots, (\mathbf{X}, \mathbf{Y})_M\}$  be the training dataset containing  $M$  samples with  $\mathbf{X}_m$  and  $\mathbf{Y}_m$  respectively denoting the  $m$ th image and its corresponding label map (also called mask). Here,  $\mathbf{X}$  is of size  $H \times W \times 3$  for RGB images with a total of  $H \times W = N$  pixels. The corresponding label map  $\mathbf{Y}$  is of size  $H \times W$  with elements in  $[K]$ .  $\mathbf{X}_m$  contains a maximum of  $K$  classes each occupying  $n_{m,c}$  pixels, where  $\sum_{c \in [K]} n_{m,c} = H \times W = N, \forall m$ . The most commonly used ERM loss for semantic segmentation is the pixel-wise loss  $\mathcal{L} = \sum_{m=1}^M \text{CE}(\mathbf{Y}_m, \hat{\mathbf{Y}}_m)$  [1, 29, 27], which is computed using a multi-class CE between the 1-hot encoded versions of the original label map  $\mathbf{Y}$  and the inferred one  $\hat{\mathbf{Y}}$ :

$$\mathcal{L} = -\frac{1}{MN} \sum_{m=1}^M \sum_{c=1}^K \sum_{i=1}^{n_c} y_{m,c,i} \log(\hat{y}_{m,c,i}), \quad (1)$$

where  $y_{m,c,i}$  denotes the  $i$ th pixel in the  $c$ th class of  $\mathbf{Y}_m$ . [12] proposes to *tilt* ERM  $\mathcal{R}(\theta) = 1/M \sum_{i \in [M]} f(x_i, \theta)$  as  $\tilde{\mathcal{R}}(\theta, t) = 1/t \log(1/M \sum_{i \in [M]} e^{t f(x_i, \theta)})$ , with loss function  $f(x_i, \theta)$  and model parameters  $\theta$ . The tilt parameter  $t$  can be tuned to flexibly promote robustness or fairness. In theory, setting  $t = 0$  recovers ERM (i.e.,  $\tilde{\mathcal{R}}(\theta, 0) = \mathcal{R}(\theta)$ ), and as  $t \rightarrow \infty$ , TERM minimizes the worst loss, thus ensuring the model is a reasonable fit for all samples [12]. With this in mind, there are (at least) two levels around which a sensible tilting of multi-class CE (MCCE) (1) can be implemented. More specifically, we can tilt (1) at i) image (or sample) level and ii) class level. Let us start with the image level. To tilt at *image level*, we need to reformulate (1) as

$$\tilde{\mathcal{L}}^{img} = \frac{1}{t} \log \left( \frac{1}{M} \sum_{m=1}^M e^{t \mathcal{L}_m} \right), \quad (2)$$

$$\text{where: } \mathcal{L}_m = -\frac{1}{N} \sum_{c=1}^K \sum_{i=1}^{n_c} y_{m,c,i} \log(\hat{y}_{m,c,i}).$$

Following the same strategy, to tilt at *class level* per image we need to reformulate (1) as

$$\tilde{\mathcal{L}}^{cls} = \frac{1}{M} \sum_{m=1}^M \frac{1}{t} \log \left( \frac{1}{K} \sum_{c=1}^K e^{t \mathcal{L}_{m,c}} \right), \quad (3)$$

$$\text{where: } \mathcal{L}_{m,c} = -\frac{1}{n_c} \sum_{i=1}^{n_c} y_{m,c,i} \log(\hat{y}_{m,c,i}).$$

<sup>1</sup>The code will be publicly available soon.

---

### Algorithm 1: Stochastic TCE for Segmentation

---

**Initialize:**  $w_c = \tilde{\mathcal{L}}_c = 0, \forall c \in [C], \theta$   
**Require:**  $\gamma, t$   
1 Divide  $\mathcal{D}^t$  into  $\mathcal{D}_c^t, \forall c \in [C]$   
2 **while** stopping criteria not met **do**  
3     sample class  $c \in [C]$  from a categorical  
    distribution with probabilities  $w_c \in W$   
4     sample minibatch  $B$  within  $\mathcal{D}_c^t$   
5      $\mathcal{L}_B \leftarrow$  compute the loss (1) on  $B$   
6     tilt the batch loss:  $\tilde{\mathcal{L}}_B \leftarrow e^{t \mathcal{L}_B}$   
7      $\tilde{\mathcal{L}}_c \leftarrow (1 - \gamma) \tilde{\mathcal{L}}_c + \gamma \tilde{\mathcal{L}}_B$   
8      $w_c \leftarrow \tilde{\mathcal{L}}_c / (\sum_{l=1}^K \tilde{\mathcal{L}}_l), \forall c \in [C]$   
9     Update model parameters:  $\theta \leftarrow \theta - \nabla \mathcal{L}_B$   
10 **end**

---

### 2.1. Solving TCE for Semantic Segmentation

Depending on the tilt level, one has to replace the pixel-wise MCCE part of the semantic segmentation loss with one of the proposed losses in (2) and (3). In our experience, directly plugging these loss functions into semantic segmentation optimization problem could lead to convergence issues and caution has to be put in place. An alternative approach to solve TCE (applicable to both sample and hierarchical levels) is to follow along the *stochastic* approach proposed in [12]. It is proven in [12] that the gradient of the tilted loss  $\tilde{\mathcal{L}}$  is a weighted average of the gradients of the original individual losses, where each data point is weighted exponentially proportional to the value of its loss. This is the key idea behind the proposed dynamic weight updating and sampling strategy of Stochastic TCE laid out in Algorithm 1. Let us dive deeper and walk through the algorithm. For the sake of simplicity, here we drop the superscript denoting the tilting level of  $\tilde{\mathcal{L}}$  in (2) and (3), and use a subscript to refer to the class  $\tilde{\mathcal{L}}_c$  and batch of data within the class  $\tilde{\mathcal{L}}_B$ . The class weights  $w_c$  are stored/updated in  $W$ .

The algorithm starts by dividing the training dataset  $\mathcal{D}^t$  into  $C$  subsets  $\mathcal{D}_c^t$  each containing the images corresponding to individual classes. Note that an image can contain multiple classes, and thus, these sets can overlap. One can also consider forming non-overlapping sets based on  $\mathcal{D}^t$ . Per propagation round, one class (let us say  $c$ ) will be selected from the categorical distribution  $[C]$  with probabilities (weights)  $W$  (line 3). These weights are dynamically updated (line 8). Next, a minibatch  $B$  is sampled from the training data of the selected class  $\mathcal{D}_c^t$  and the tilted batch loss  $\tilde{\mathcal{L}}_B$  is calculated on  $B$  (lines 5 and 6). Line 7 proposes a linear dynamic with rate  $\gamma$  to update the tilted loss of the selected class  $\tilde{\mathcal{L}}_c$  based on its previous value and the current batch estimate  $\tilde{\mathcal{L}}_B$ . The weight of class  $c$ ,  $w_c$ , will then be updated using a normalization applied to all the tilted losses

Table 1: Performance comparison on Cityscapes *validation* set, sorted based on DLv3+ with MCCE.

| Method          | wall         | train        | rider        | fence        | terrain      | truck        | m.cycle      | pole         | bus          | t. light     |
|-----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| MCCE [1]        | 48.46        | 53.66        | 61.12        | <b>62.23</b> | 62.98        | 66.21        | 67.72        | 69.19        | 71.53        | 74.33        |
| Focal loss [14] | 49.11        | <b>79.55</b> | <b>67.56</b> | 60.75        | 62.31        | <b>73.16</b> | 67.71        | 64.53        | <b>85.02</b> | 68.92        |
| TCE $t = .1$    | 49.36        | 76.75        | 66.64        | 60.38        | <b>65.69</b> | 72.03        | 69.09        | 69.34        | 75.61        | 73.64        |
| TCE $t = 1$     | <b>53.47</b> | <b>79.32</b> | 65.67        | 59.25        | 63.74        | 64.32        | <b>69.45</b> | <b>69.54</b> | 66.05        | <b>74.51</b> |

| continued       | bicycle      | t. sign      | person       | sidewalk     | vegetation   | building     | sky          | car          | road         | mIoU         |
|-----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| MCCE [1]        | 79.23        | <b>81.22</b> | 83.19        | 84.95        | <b>92.52</b> | 93.04        | 95.00        | <b>95.45</b> | 98.06        | 75.79        |
| Focal loss [14] | 77.71        | 77.62        | 81.23        | 81.56        | 91.54        | 91.89        | 93.71        | 94.67        | 97.14        | 77.14        |
| TCE $t = .1$    | <b>79.89</b> | 81.77        | <b>84.00</b> | <b>86.24</b> | 92.43        | <b>93.16</b> | <b>95.34</b> | <b>95.47</b> | <b>98.27</b> | <b>78.16</b> |
| TCE $t = 1$     | <b>79.92</b> | 81.04        | <b>83.96</b> | <b>86.22</b> | 92.44        | 92.92        | 95.11        | 94.53        | <b>98.28</b> | 77.35        |

Table 2: Performance comparison on Cityscapes *validation* set, sorted based on SoTA DLv3+.

| Method        | wall         | fence        | rider        | terrain      | m.cycle      | pole         | t. light     | bicycle      | t. sign      | train        |
|---------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| SoTA MCCE [1] | <b>57.26</b> | <b>62.18</b> | 62.76        | 63.38        | 64.50        | 65.11        | 68.41        | 77.26        | 78.78        | <b>80.90</b> |
| TCE $t = .1$  | 49.36        | 60.38        | <b>66.64</b> | <b>65.69</b> | <b>69.09</b> | <b>69.34</b> | <b>73.64</b> | <b>79.89</b> | <b>81.77</b> | 76.75        |

| continued     | person       | sidewalk     | truck        | bus          | vegetation   | building     | sky          | car          | road         | mIoU         |
|---------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| SoTA MCCE [1] | 82.14        | 84.7         | <b>85.31</b> | <b>89.07</b> | <b>92.65</b> | 92.69        | 95.29        | 95.31        | 98.13        | <b>78.73</b> |
| TCE $t = .1$  | <b>84.00</b> | <b>86.24</b> | 72.03        | 75.61        | 92.43        | <b>93.16</b> | <b>95.34</b> | <b>95.47</b> | <b>98.27</b> | 78.16        |

(line 8). These dynamically updated weights,  $w_c \in W$ , will be used in the next iteration to decide from which class to sample. Finally, model parameters in  $\theta$  are updated.

### 3. Experimental Setup

Here, we assess the impact of TCE on one of the most commonly adopted datasets for semantic segmentation, Cityscapes [3]. The evaluation results for ADE20k [31] can be found in the appendix. Cityscapes contains 2,975 train and 500 validation images from 19 main target classes.

**Training strategy and baselines.** Our trainings are run separately on standard Microsoft Azure 4-GPU P100 Tesla nodes, each with 16GB of memory. For experiments on Cityscapes, we used DeepLabv3+ (also referred to as DLv3+) [1] with ResNet-101 backbone as our reference implementation of multi-class CE (MCCE), and on top of that we have implemented TCE. DLv3+ is among the top performing model architectures for Cityscapes. Following [1], we used minibatch SGD with learning rate  $l_r = 0.01$  and momentum 0.9 for all models, and adjusted for a total minibatch size of 8 (2 per GPU). The reported results of [1] are based on our own trainings, for the sake of a fair comparison. Image crop size and other pre/post-processing parameters are set per default as suggested in [1]. We also compare our performance against the Focal loss for semantic segmentation [14] with the best parameters  $\gamma = 2$  and (class weights)  $\alpha$  set to the inverse (normalized) class pixel counts computed across the whole dataset.

**Fairness and its evaluation criteria.** The notion of fairness in this setting is promoting a more consistent (and less varied) performance across different classes. This is to ensure that there are less (or ideally no) classes that have been significantly disadvantaged, due to for instance less presence

or difficult characteristic features, for the sake of a higher average performance. Promoting fairness by minimizing performance disparity is also a core idea of TERM [12], and resonates with other recent approaches to fairness [6, 13, 17]. More concretely, i) best worst-case performance [6, 17], and ii) least variance across clients/classes [13] are recently proposed to promote/evaluate fairness across a set of tasks or networked clients. We investigate both measures as our key criteria. More specifically, besides the overall mean-intersection-over-union (mIoU), we compare the models on: i) sorted (w.r.t MCCE) bottom and top 25% mIoUs; ii) bottom and top 25th percentiles; and iii) standard deviation and worst case performance (in IoU) across classes. Note that the overall mIoU does not have to be improved when applying TCE; the goal is to minimize performance disparity by boosting the low-performing and difficult classes. This can sometimes come at the cost of lower overall mIoU and/or degradation of the high-performing target classes.

### 4. Evaluation Results

Table 1 compares the sorted mIoU breakdown of DLv3+ with ResNet-101 backbone [1] trained with the standard multi-class cross-entropy (MCCE) against the same model retrained with the Focal loss and the proposed TCE for  $t = 0.1$  and 1. Even though we are not *necessarily* expecting an improvement in the overall mIoU, here for all  $t$ 's, the overall performance with TCE has improved beyond same model trained with MCCE (+2% for  $t = 0.1$ ) and Focal loss (about 1% for  $t = 0.1$ ). TCE with  $t = 1$  is expected to push more towards minimizing performance disparity and thus promoting fairness (as is also shown in Table 3), and it shows the best improvement in the least performing classes such as “wall” and “train”. The same model architecture, DeepLabv3+ (ab-




Figure 1: Impact of TCE on improving low-performing classes of MCCE. Best view in color with 300% zoom.

Table 3: Fairness criteria for Cityscapes [3]

| Method          | sorted 25%   |       | (25 <sup>th</sup> perc., mIoU) |                | overall      |              |
|-----------------|--------------|-------|--------------------------------|----------------|--------------|--------------|
|                 | bottom       | top   | bottom                         | top            | worst        | std.         |
| MCCE [1]        | 57.69        | 94.81 | (64.60, 57.69)                 | (88.74, 94.81) | 48.46        | 14.96        |
| Focal loss [14] | 63.86        | 93.79 | (67.64, 60.85)                 | (88.28, 93.79) | 49.11        | <b>13.35</b> |
| TCE $t = .1$    | 63.76        | 94.93 | (69.22, <b>62.23</b> )         | (89.34, 94.93) | 49.36        | <b>13.34</b> |
| TCE $t = 1$     | <b>64.29</b> | 94.66 | (65.86, 61.29)                 | (89.33, 94.66) | <b>53.47</b> | 13.57        |

breviated to DLv3+ here), reports better performance results with a more complex backbone Xception-65 [1]. We could not reproduce those results because we did not have access to large enough nodes on Microsoft Azure to accommodate this backbone with batch sizes larger than 8. Nonetheless, we were curious to know how TCE implemented on top of a model with a weaker backbone would compare with the the state-of-the-art (SoTA) results reported in [1]. This comparison is summarized in Tables 2. Interestingly, even compared to the SoTA model with improved backbones, TCE is still improving on several low-performing classes (such as “rider”, “terrain”, “motorcycle”, etc.) in favor of promoting fairness across target classes.

These two tables illustrate promising improvement in low-performing classes the extent of which is further investigated and is summarized in Table 3. Here, the following three measures are presented. First, the sorted bottom and top 25% mIoU. To compute this, the target classes are sorted based on the IoU performance breakdown of the the model trained with MCCE. Then for each model the mIoU of the bottom and top 25% (5 classes out of 19) are taken into account. This is to demonstrate the impact of TCE compared to the reference MCCE. Here, improved mIoU for bottom classes (even at cost the of lower mIoU for the top ones) indicates more fairness. Second, in a tuple, the IoU threshold corresponding to (bottom and top) 25th percentile of each model and the mIoU of the classes falling within the percentile are

presented. Note that in this case each model will be sorted according to its own target class IoU’s. The idea is that improvement in bottom percentile threshold and corresponding mIoU could be indicative of improved fairness. Again, we do not expect improvement but potential drop in the top percentiles; they are reported to provide a more complete picture. Third, the overall fairness measures [6, 13, 17], i.e., the worst performance among target classes and the standard deviation across class IoU’s (denoted respectively as worst and std. in the table) are presented.

The results show that on all three metrics TCE offers improved fairness when compared to MCCE and Focal loss. Let us focus on  $t = 1$ . In sorted bottom 25%, we gain +6% and +0.4% beyond MCCE and Focal loss, respectively. For 25th percentiles, +5% and +0.4% beyond MCCE and Focal loss, respectively. The TCE with  $t = 0.1$  seems to do better (+1% beyond Focal) which could be due to different sorting per model, and thus, different classes falling in those percentiles. In the overall metrics [6, 13, 17], the worst-case IoU (among target classes) is 5% and +4% better (higher) than MCCE and Focal loss, respectively. The standard deviation across classes is improved (decreased) by +1% compared to MCCE and remains in the same regime as Focal loss. Finally, qualitative results in Fig. 1 further corroborate the impact of TCE in comparison with MCCE and Focal loss. The top two rows highlight improvement in “rider”, “bus”, and “truck”, and the next two rows show improvement in “tram” and “sidewalk”, most of which associate to the low-performing classes of the model trained with MCCE.

**Concluding remarks.** There is plenty of room for extending this work. An avenue to explore is different implementations of TCE, especially (in a non-stochastic fashion) by directly plugging in (2) and (3) as the optimization loss function. In doing so, remedies have to be put in place to circumvent convergence issues. Further results on ADE20k dataset are provided in the supplementary material.

## References

- [1] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.
- [2] Bowen Cheng, Maxwell D Collins, Yukun Zhu, Ting Liu, Thomas S Huang, Hartwig Adam, and Liang-Chieh Chen. Panoptic-deeplab. *arXiv preprint arXiv:1910.04751*, 2019.
- [3] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [4] John Duchi and Hongseok Namkoong. Variance-based regularization with convex objectives. *arXiv preprint arXiv:1610.02581*, 2016.
- [5] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [6] Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*, pages 1929–1938. PMLR, 2018.
- [7] Corentin Henry, Seyed Majid Azimi, and Nina Merkle. Road segmentation in sar satellite images with deep fully convolutional neural networks. *IEEE Geoscience and Remote Sensing Letters*, 15(12):1867–1871, 2018.
- [8] 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. IEEE, 2020.
- [9] 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.
- [10] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *International journal of Computer Vision*, 1(4):321–331, 1988.
- [11] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9404–9413, 2019.
- [12] Tian Li, Ahmad Beirami, Maziar Sanjabi, and Virginia Smith. Tilted empirical risk minimization. *arXiv preprint arXiv:2007.01162*, 2020.
- [13] Tian Li, Maziar Sanjabi, Ahmad Beirami, and Virginia Smith. Fair resource allocation in federated learning. *arXiv preprint arXiv:1905.10497*, 2019.
- [14] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [15] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 3061–3070, 2015.
- [16] Shervin Minaee and Yao Wang. An ADMM approach to masked signal decomposition using subspace representation. *IEEE Transactions on Image Processing*, 28(7):3192–3204, 2019.
- [17] Mehryar Mohri, Gary Sivek, and Ananda Theertha Suresh. Agnostic federated learning. *arXiv preprint arXiv:1902.00146*, 2019.
- [18] Richard Nock and Frank Nielsen. Statistical region merging. *IEEE Transactions on pattern analysis and machine intelligence*, 26(11):1452–1458, 2004.
- [19] Vasyli Pihur, Susmita Datta, and Somnath Datta. Weighted rank aggregation of cluster validation measures: a monte carlo cross-entropy approach. *Bioinformatics*, 23(13):1607–1615, 2007.
- [20] Nils Plath, Marc Toussaint, and Shinichi Nakajima. Multi-class image segmentation using conditional random fields and global classification. In *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)*, pages 817–824, 2009.
- [21] Mina Rezaei, Konstantin Harmuth, Willi Gierke, Thomas Kellermeier, Martin Fischer, Haojin Yang, and Christoph Meinel. A conditional adversarial network for semantic segmentation of brain tumor. In *International MICCAI Brainlesion Workshop*, pages 241–252. Springer, 2017.
- [22] Javier Ribera, David Güera, Yuhao Chen, and Edward Delp. Weighted hausdorff distance: A loss function for object localization. *arXiv preprint arXiv:1806.07564*, 2, 2018.
- [23] 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.
- [24] Alexey A Shvets, Alexander Rakhlin, Alexandr A Kalinin, and Vladimir I Iglovikov. Automatic instrument segmentation in robot-assisted surgery using deep learning. In *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 624–628. IEEE, 2018.
- [25] 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 *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pages 240–248. Springer, 2017.
- [26] Michele Volpi and Vittorio Ferrari. Semantic segmentation of urban scenes by learning local class interactions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1–9, 2015.
- [27] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 2020.
- [28] Ken CL Wong, Mehdi Moradi, Hui Tang, and Tanveer Syeda-Mahmood. 3d segmentation with exponential logarithmic loss for highly unbalanced object sizes. In *International*

*Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 612–619. Springer, 2018.

- [29] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 418–434, 2018.
- [30] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *Proceedings of the IEEE international conference on computer vision*, pages 1395–1403, 2015.
- [31] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127(3):302–321, 2019.