



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Average Calibration Error: A Differentiable Loss for Improved Reliability in Image Segmentation

Theodore Barfoot¹, Luis C. Garcia Peraza Herrera¹, Ben Glocker², and Tom Vercauteren¹

¹ King's College London, London, UK

theodore.d.barfoot@kcl.ac.uk, luis_c.garcia_peraza_herrera@kcl.ac.uk,
tom.vercauteren@kcl.ac.uk

² Imperial College London, London, UK

b.glocker@imperial.ac.uk

Abstract. Deep neural networks for medical image segmentation often produce overconfident results misaligned with empirical observations. Such miscalibration challenges their clinical translation. We propose to use marginal L1 average calibration error (mL1-ACE) as a novel auxiliary loss function to improve pixel-wise calibration without compromising segmentation quality. We show that this loss, despite using hard binning, is directly differentiable, bypassing the need for approximate but differentiable surrogate or soft binning approaches. Our work also introduces the concept of *dataset reliability histograms* which generalises standard reliability diagrams for refined visual assessment of calibration in semantic segmentation aggregated at the dataset level. Using mL1-ACE, we reduce average and maximum calibration error by 45% and 55% respectively, maintaining a Dice score of 87% on the BraTS 2021 dataset. We share our code here: <https://github.com/cai4cai/ACE-DLIRIS>.

1 Introduction

Deep neural networks (DNNs) have significantly advanced the field of semantic segmentation. However, their large capacity makes them susceptible to overfitting, leading to overconfident predictions that do not accurately reflect the underlying uncertainties inherent to the task [8]. In medical image segmentation, where the confidence of predictions can be as crucial as the predictions themselves, and overfitting is exacerbated by smaller datasets, such overconfidence poses significant risks, particularly when this miscalibration is not communicated to the end-user. Appropriate management of confidence can also have benefits for patient outcomes, for example in radiotherapy contouring [2].

The Dice Similarity Coefficient (DSC) loss is a popular choice in medical image segmentation due to its robustness against class imbalance [15], but it is known to yield poorly calibrated, overly confident predictions [16].

Reliability diagrams are the reference means of assessing and visualising the calibration performance of a predictor over a dataset and this has extensively

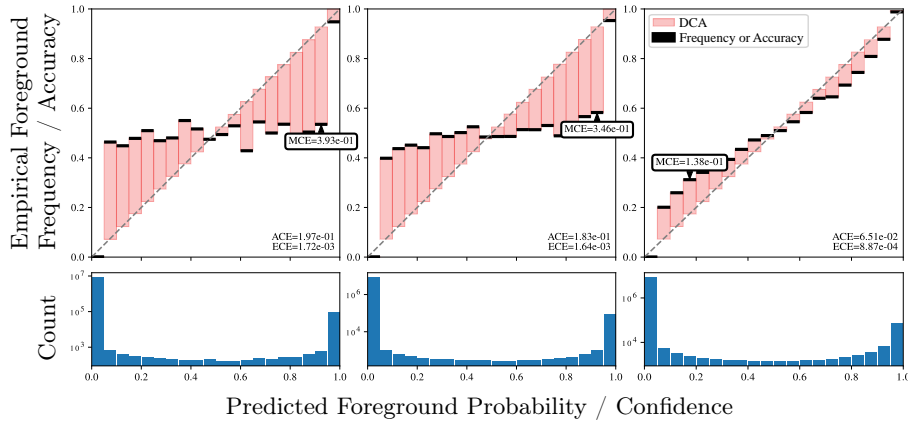


Fig. 1. Image-specific reliability diagrams for whole tumour segmentation (BraTS case 00054) under different training losses. Left: DSC; Middle: DSC + temperature scaling (Ts); Right: DSC + mLI-ACE (proposed).

been used for classification tasks [5]. They are constructed by binning the samples according to a discretisation of the predicted probabilities, also referred to as confidence. Within each bin, the average confidence and average accuracy is computed. The reliability diagram is obtained by plotting the Difference between Confidence and Accuracy (DCA) as a function of confidence. An example is shown in Fig. 1. An aggregation of these DCA values can then be used to estimate measures of calibration error. While reliability diagrams and associated calibration error metrics are typically computed at the dataset level, in semantic segmentation, these diagrams can also be constructed per image, due to multiple pixels/voxels being present. Capturing both image-specific and dataset-wide reliability is of particular relevance in medical imaging, due to the safety critical nature of the task. Box-plot diagrams have been used in combination with image-specific reliability diagrams to represent calibration variability within an image segmentation dataset [12].

Traditional post-hoc calibration strategies such as temperature scaling (Ts) [8,21] offer a partial remedy by adjusting confidence levels uniformly. While more nuanced post-hoc calibration methods offer class-specific adjustments [13,6,11], they still fail to exploit the extensive parameter space of DNNs [14,9].

Train-time calibration methods are more effective in improving calibration than post-hoc methods [8]. They leverage the large model capacity in DNNs to provide fine-grained calibration adjustment with no detriments for the primary task. For image classification tasks, DCA-based auxiliary losses such as expected calibration error (ECE) have been shown to improve calibration [14]. However, the accuracy term in the DCA-based losses is not differentiable due to the argmax operator. Furthermore, the usage of hard binning in standard DCA has led to previous work seeking for alternative differentiable calibration

losses [14]. Multi-class DCA (MDCA) avoids binning, and its associated differentiability challenges, by calculating DCA averaged over each class within a mini-batch rather than averaged within confidence bins [9]. A differentiable surrogate for expected calibration error (DECE), has been proposed [3]. However, the DECE approach relies on an approximation of accuracy, as well as soft binning to avoid the non-differentiability of hard binning. The DECE work also exploits non-trivial meta-learning rather than direct optimisation of the loss thereby complicating its use and limiting its wider applicability.

In a different stream of work, authors have introduced specialised loss functions which do not rely on DCA methods to tackle overconfidence in biomedical image segmentation. For example, DSC++ loss, which uses a modified DSC loss to enhance calibration by adjusting the false positive (FP) and false negative (FN) terms of DSC based on class frequency [23]. Neighbor-Aware Calibration (NACL) is another example focusing on spatial consistency, via penalty constraints on logits, to improve reliability and performance in segmentation tasks [17,18]. While these methods have demonstrated some improvement in calibration, their indirect approach make them more difficult to interpret and potentially lead to suboptimal results.

In this work, we present a novel auxiliary loss function, marginal L1 Average Calibration Error (mL1-ACE), designed to enhance calibration accuracy in medical image segmentation without compromising on segmentation quality. This represents the first clear attempt to use established calibration metrics directly as an auxiliary loss in the context of medical image segmentation. Despite presumptions related in previous DCA-based methods, our approach is inherently differentiable, even with hard-binning of probabilities, eliminating the need for surrogates or soft-binning techniques. mL1-ACE employs a class-wise calibration strategy, opting for ACE over ECE to better address the class imbalances in medical datasets.

While the concept of average calibration error (ACE) is not new, our key contribution lies in adapting it as a directly differentiable auxiliary loss for semantic segmentation, a novel approach not addressed in prior literature. This provides a major distinction from the use of ACE as a metric in classification tasks. The transition from classification to segmentation is a key aspect which we are the first to identify, enabling ACE to be used as a loss rather than a metric. Classification training pipelines provide too few samples within a mini-batch to meaningfully compute its ACE, in contrast to segmentation where each voxel provides a sample.

When combined with standard loss functions, in particular DSC, mL1-ACE demonstrates superior calibration performance on the BraTS2021 dataset [1], significantly surpassing conventional post-hoc calibration methods like Ts. Additionally, we introduce dataset reliability histograms, offering a comprehensive visual tool for evaluating and communicating model calibration across image segmentation datasets, enhancing transparency in reporting model reliability beyond predictive performance.

2 Method

For calculation of image-specific calibration metrics, and our subsequent mL1-ACE loss, we discretise the continuous predicted probabilities of each of the C classes into M bins. For any given image, and class c , the expected foreground probability e_m^c within bin m (also referred to as confidence in DCA) is compared to the observed frequency o_m^c within that same bin (also referred to as accuracy in DCA). This provides us with an analogous representation of DCA but extended to multi-class setting. An ideally calibrated model is one where the gaps, $|o_m^c - e_m^c|$ equal zero over all M bins and all C classes. That is, a perfectly calibrated model is one where the predicted foreground probabilities align precisely with the observed frequency of that class across numerous samples. For example, if a well-calibrated brain tumour segmentation model segments 100 voxels, each with a probability/confidence of 0.7, then we expect 70 of those voxels to be tumour.

Expected Calibration Error (ECE) [19] distils the reliability diagram described above into a single scalar value that can be used to assess the calibration of the model, by considering the weighted sum of the absolute difference of the gaps over all the bins. We extend this approach to a multi-class and image-specific variant, which we refer to as marginal L1 expected calibration error:

$$\text{mL1-ECE} = \frac{1}{C} \sum_c \sum_m^M \frac{n_m^c}{N} |o_m^c - e_m^c| \quad (1)$$

where n_m^c is the number of in bin B_m^c , the m^{th} bin for class c , and N is the total number of voxels for an image. More formally, we define the observed foreground frequency o_m^c and the expected foreground probability e_m^c as:

$$o_m^c = \mathbb{P}[\mathbf{Y}_i^c | i \in B_m^c], \quad e_m^c = \mathbb{E}[f_{\theta}(\mathbf{X})_i^c | i \in B_m^c] \quad (2)$$

where \mathbf{X} is the input image, \mathbf{Y} the associated ground truth label-map, and f_{θ} is our predictor with trainable weights θ .

Average Calibration Error (ACE) [20], performs this distillation in a similar way, but considers an unweighted sum over all bins. We extend this approach to a multi-class and image-specific variant, which we refer to as marginal L1 average calibration error:

$$\text{mL1-ACE} = \frac{1}{CM} \sum_c \sum_m^M |o_m^c - e_m^c| \quad (3)$$

Finally, we report Maximum Calibration Error (MCE) [19] which is only concerned with the maximum gap but extend it to a multi-class and image-specific variant while referring to it as marginal L1 maximum calibration error:

$$\text{mL1-MCE} = \frac{1}{C} \sum_c \max_m |o_m^c - e_m^c| \quad (4)$$




Fig. 2. Dataset reliability histograms for BraTS 2021 dataset, with different loss functions, without (top row) and with (bottom row) mL1-ACE.

We anticipate that all three measures, as shown in equations (1), (3), and (4), should be differentiable almost everywhere. Combined with their image-specific nature, this characteristic makes them trivially usable as auxiliary losses that can be paired with any image segmentation loss. Specifically, useful differentiability arises due to the large number of voxels that have membership to each bin with averaging being performed within each bin rather than across the bins themselves. However, given that mL1-MCE only considers the maximum gap, for a single bin, it is less likely to provide a sufficient supervisory signal compared to mL1-ECE and mL1-ACE.

Since mL1-ECE is weighted by the relative proportion of voxels in each bin, if a model is performing sufficiently well, mL1-ECE is heavily skewed towards the first and last bin as is evident in Fig. 1. Therefore, we argue that mL1-ACE is the most appropriate measure for calibration, and therefore an excellent candidate for an auxiliary loss targeted at model calibration, in the context of medical image segmentation. Indeed, mL1-ACE will be much more sensitive to uncertain predictions such as those found on the periphery of segmented regions.

To capture reliability information at the dataset-level rather than image-level, we suggest to aggregate the image-specific metrics rather than computing a dataset-level DCA. This not only provides an opportunity to gain insights in the variability within the dataset but also offers a pragmatic solution to deal with large datasets. To support visual assessment of dataset-level reliability, we propose an extension of reliability diagrams illustrated in Fig. 2 which we

refer to as dataset reliability histograms. By binning the empirical foreground frequencies within each predicted foreground probability bin across all cases in a dataset, we create a joint histogram representation of reliability. This provides an informative visualisation of calibration across an entire dataset.

3 Experiments

Training, validation and testing was performed using the Multimodal Brain Tumour Segmentation (BraTS) Benchmark training dataset from 2021 [1]. Segmentations consist of three tumour sub-regions. These sub-regions then form clinical tumour sub-regions, summarised as follows: Enhancing tumour (ET) captures the gadolinium enhancement. Tumour core (TC) is comprised of the necrotic tumour core (NCR) and ET. Whole tumour (WT) is comprised of the TC and the peritumoral edematous and infiltrated tissue (ED).

The training dataset of 1251 cases was randomly split into training, validation and testing with a ratio of 1000:51:200 cases. A basic UNet architecture [22] implemented in MONAI [4] was used with 8, 16, 32, and 64 channels and strides of 2. We use the same data augmentation pipeline, for training, as in [7], closely matching the one used for the nnUNet [10].

A training patch size of $224 \times 224 \times 144$ was used, with a batch size of 8. Training was ran for 1000 epochs, with 16 iterations per epoch. The final model for evaluation was the one with the best validation DSC. Different training strategies were compared, using combinations of popular segmentation losses with and without the mL1-ACE loss. We considered cross-entropy (CE), DSC, and CE + DSC loss functions. Equal weighting for each loss function component was used. For all calibration metrics and mL1-ACE loss, 20 bins were used to discretise the predicted probability space. For all runs, an Adam optimiser with a learning rate of 0.001 was used. Post-hoc Ts was performed for each model, using CE loss function, with the temperature parameter optimised using the validation dataset. All experiments were performed using an NVidia DGX cluster with NVidia V100 (32GB) and Nvidia A100 (40GB) GPUs.

4 Results

We first consider the impact of our auxiliary loss in terms of segmentation performance and then focus on calibration. The DSC scores for the different loss functions are summarised in Table 1, showing that despite using a basic network with no hyper-parameter tuning, we obtain competitive DSC values on BraTS 2021. The addition of mL1-ACE as an auxiliary loss with CE or DSC-CE losses, results in a small reduction in DSC scores ($p < 0.01$) However, for DSC loss, the addition of mL1-ACE maintains DSC performance, showing no statistically significant difference ($p = 0.61$). The DSC values of models with post-hoc Ts are not included as Ts has no practically significant effect on DSC.

We show the effect of Ts and mL1-ACE on the calibration of the models in Table 2. Our auxiliary loss has little effect on the calibration of the model trained




Fig. 3. Comparative analysis of DSC against calibration metrics for average tumour component of BraTS 2021 dataset. Statistically significant difference ($p < 0.01$) for all metrics.

Table 1. Dice Score across Loss Functions, higher is better

Loss Function	ET	TC	WT	Avg
CE	0.82 ± 0.19	0.87 ± 0.18	0.90 ± 0.11	0.86 ± 0.14
CE + L1-ACE	0.80 ± 0.20	0.86 ± 0.19	0.87 ± 0.12	0.84 ± 0.15
DSC	0.84 ± 0.17	0.87 ± 0.20	0.89 ± 0.12	0.87 ± 0.14
DSC + L1-ACE	0.83 ± 0.18	0.88 ± 0.18	0.89 ± 0.11	0.87 ± 0.13
DSC-CE	0.85 ± 0.17	0.89 ± 0.17	0.90 ± 0.11	0.88 ± 0.13
DSC-CE + L1-ACE	0.82 ± 0.20	0.86 ± 0.21	0.88 ± 0.13	0.85 ± 0.16

with CE. Our loss shows the biggest improvement in calibration when paired with the DSC loss function, with a 45% reduction in ACE ($p < 0.01$). Overall the use of Ts shows marginal improvement in calibration, with the exception of CE loss where it has a small detrimental effect. Similar conclusion can be drawn for MCE and ECE metrics, shown in Table S1 and Table S2 in the supplementary material. We also note that the number of bins used to calculate our calibration metrics, has little effect, with ACE increasing by less than 3% when increasing the bins from 10 to 100 ($p < 0.01$), as shown in supplementary Fig. S1.

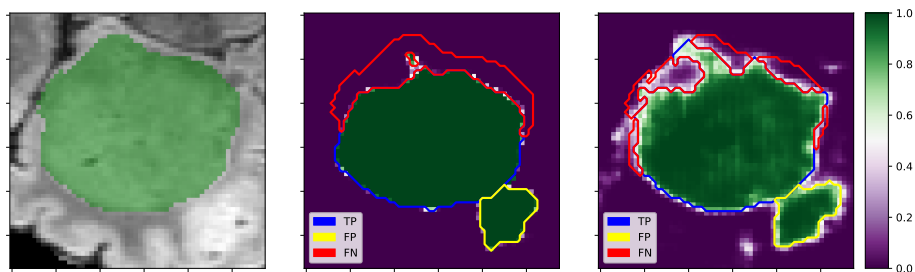
The relationship between DSC and calibration metrics is shown in Fig. 3, for the DSC loss model, showing that the addition of mL1-ACE as an auxiliary loss maintains DSC performance with a substantial improvement in MCE and ACE ($p < 0.01$).

Fig. 1 shows a reliability diagram for the DSC loss model, with and without mL1-ACE and Ts, for the WT segmentation component. As can be seen, Ts has a limited effect on the calibration for an individual case, whereas the addition of mL1-ACE has a significant effect, clearly improving calibration.

Similar, but more pronounced, results are shown in Fig. 2, where the addition of mL1-ACE has a significant qualitative improvement on calibration of the model across the entire dataset, for all loss functions, showing a marked reduction in under- and over-confident cases, by the narrowing of the distribution along the diagonal. We supplement this visual qualitative analysis, by showing an individual case in Fig. 4, where the addition of mL1-ACE has a significant effect

Table 2. Average Calibration Error across Loss Functions, lower is better

Loss Function	ET	TC	WT	Avg
CE	0.13 ± 0.09	0.14 ± 0.10	0.12 ± 0.07	0.13 ± 0.06
CE + Ts	0.15 ± 0.07	0.13 ± 0.08	0.13 ± 0.06	0.14 ± 0.05
CE + L1-ACE	0.12 ± 0.10	0.14 ± 0.10	0.11 ± 0.09	0.12 ± 0.07
CE + L1-ACE + Ts	0.12 ± 0.07	0.13 ± 0.08	0.11 ± 0.07	0.12 ± 0.05
DSC	0.25 ± 0.07	0.26 ± 0.07	0.23 ± 0.06	0.25 ± 0.05
DSC + Ts	0.24 ± 0.07	0.24 ± 0.07	0.22 ± 0.06	0.24 ± 0.05
DSC + L1-ACE	0.14 ± 0.10	0.15 ± 0.09	0.12 ± 0.08	0.14 ± 0.07
DSC + L1-ACE + Ts	0.12 ± 0.09	0.13 ± 0.09	0.11 ± 0.07	0.12 ± 0.06
DSC-CE	0.20 ± 0.08	0.21 ± 0.08	0.18 ± 0.09	0.20 ± 0.06
DSC-CE + Ts	0.18 ± 0.09	0.18 ± 0.09	0.15 ± 0.09	0.17 ± 0.07
DSC-CE + L1-ACE	0.15 ± 0.11	0.15 ± 0.11	0.11 ± 0.08	0.14 ± 0.08
DSC-CE + L1-ACE + Ts	0.13 ± 0.10	0.14 ± 0.10	0.11 ± 0.06	0.13 ± 0.07

**Fig. 4.** Segmentation of BraTS 2021 case 00014, with DSC loss function. Left: Ground truth WT on FLAIR image, Middle: Predicted WT foreground probabilities, Right: WT foreground probabilities with mL1-ACE. True positive, false positive, and false negative contours are shown in blue, yellow, and red, respectively.

on the calibration of the model, with a marked reduction in the confidence of false positive and false negative predictions.

5 Conclusion and discussion

Prior work assumed that hard binning leads to non-differentiable DCA-based metrics. We demonstrate that it is sufficiently differentiable and provide a multi-class and image-specific variant for use in challenging image segmentation tasks. By performing averaging inside of each bin across a large enough number of voxels our approach indeed provides a sufficiently informative gradient updates. We further show that the introduction of the proposed mL1-ACE as an equally weighted auxiliary loss function to DSC loss, can improve the calibration of the model, with a 45% reduction in ACE and 55% reduction in MCE metrics, without any significant change in the Dice coefficient metric, remaining at competitive 87%, when evaluated on the BraTS 2021 dataset.

We also introduced the concept of dataset reliability histograms to visualise the reliability of a semantic segmentation models across an entire dataset, providing valuable insights into modes of failure.

Additionally, we show that models trained with cross-entropy loss tend to have much better calibration than those trained with a DSC-based loss, as has been reported in prior literature [23]. Our method has limited benefit for CE-trained models. However, with DSC-based losses outperforming CE over a large set of tasks [15], the need for well calibrated DSC-based losses remains. We also show that the post-hoc calibration method of temperature scaling has limited effect on the calibration of the model.

While our experiments were focused on BraTS with a baseline U-Net architecture, we believe that our mL1-ACE auxiliary loss has the potential to offer an easy-to-use and transferable solution to improve model calibration in any medical image segmentation approach. Exploring the use with other datasets, base losses, and architectures, including transformers, is left for future work.

Acknowledgments. The authors would like to acknowledge support for TB from EPSRC CDT [EP/S022104/1] and Intel, UK. TV is supported by a Medtronic / RAEng Research Chair [RCSR1819\7\34]. This work was supported by core funding from the Wellcome/EPSRC [WT203148/Z/16/Z; NS/A000049/1]. For the purpose of open access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission. This article only uses publicly available datasets. Their re-use did not require any ethical approval.

Disclosure of Interests. TV is co-founder and shareholder of Hypervision Surgical. BG is a part-time employee of Kheiron Medical Technologies and HeartFlow with stock options as part of the standard compensation package. The authors have no other competing interests to declare that are relevant to the content of this article.

References

1. Baid, U., Ghodasara, S., et al., S.M.: The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification. arXiv (Jul 2021)
2. Bernstein, D., Taylor, A., Nill, S., Oelfke, U.: New target volume delineation and PTV strategies to further personalise radiotherapy. *Physics in Medicine & Biology* **66**(5), 055024 (Feb 2021)
3. Bohdal, O., Yang, Y., Hospedales, T.: Meta-calibration: Learning of model calibration using differentiable expected calibration error. *Transactions on Machine Learning Research* (2023)
4. Cardoso, M.J., Li, W., Brown, R., et al.: MONAI: An open-source framework for deep learning in healthcare. arXiv preprint arXiv:2211.02701 (2022)
5. DeGroot, M.H., Fienberg, S.E.: The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society. Series D (The Statistician)* **32**(1/2), 12–22 (1983)
6. Ding, Z., Han, X., Liu, P., Niethammer, M.: Local temperature scaling for probability calibration. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 6889–6899 (Oct 2021)

7. Fidon, L., Shit, S., Ezhov, I., Paetzold, J.C., Ourselin, S., Vercauteren, T.: Generalized Wasserstein Dice loss, test-time augmentation, and transformers for the BraTS 2021 challenge. In: Crimi, A., Bakas, S. (eds.) *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*. pp. 187–196. Springer International Publishing, Cham (2022)
8. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: Precup, D., Teh, Y.W. (eds.) *Proceedings of the 34th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 70, pp. 1321–1330. PMLR (Aug 2017)
9. Hebbalaguppe, R., Prakash, S., et al.: A stitch in time saves nine: A train-time regularizing loss for improved neural network calibration. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022)
10. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nat Methods* **18**(2), 203–211 (Feb 2021)
11. Islam, M., Seenivasan, L., Ren, H., Glocker, B.: Class-distribution-aware calibration for long-tailed visual recognition. In: *International Conference on Machine Learning* (2021)
12. Kock, F., Thielke, F., Chlebus, G., Meine, H.: Confidence histograms for model reliability analysis and temperature calibration. In: Konukoglu, E., Menze, B., Venkataraman, A., Baumgartner, C., Dou, Q., Albarqouni, S. (eds.) *Proceedings of The 5th International Conference on Medical Imaging with Deep Learning*. *Proceedings of Machine Learning Research*, vol. 172, pp. 741–759. PMLR (Jul 2022)
13. Kull, M., Perello Nieto, M., Kängsepp, M., Silva Filho, T., Song, H., Flach, P.: Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with Dirichlet calibration. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 32. Curran Associates, Inc. (2019)
14. Liang, G., Zhang, Y., Wang, X., Jacobs, N.: Improved trainable calibration method for neural networks on medical imaging classification. In: *British Machine Vision Conference (BMVC)* (2020)
15. Ma, J., Chen, J., Ng, M., Huang, R., Li, Y., Li, C., Yang, X., Martel, A.L.: Loss odyssey in medical image segmentation. *Medical Image Analysis* **71**, 102035 (2021)
16. Mehrtash, A., Wells, W.M., Tempny, C.M., Abolmaesumi, P., Kapur, T.: Confidence calibration and predictive uncertainty estimation for deep medical image segmentation. *IEEE Transactions on Medical Imaging* **39**(12), 3868–3878 (2020)
17. Murugesan, B., Adiga Vasudeva, S., Liu, B., Lombaert, H., Ben Ayed, I., Dolz, J.: Trust your neighbours: Penalty-based constraints for model calibration. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 572–581. Springer Nature Switzerland, Cham (2023)
18. Murugesan, B., Vasudeva, S.A., Liu, B., Lombaert, H., Ben Ayed, I., Dolz, J.: Neighbor-aware calibration of segmentation networks with penalty-based constraints. *arXiv preprint arXiv:2401.14487* (2024)
19. Naeini, M.P., Cooper, G.F., Hauskrecht, M.: Obtaining well calibrated probabilities using bayesian binning. In: Bonet, B., Koenig, S. (eds.) *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, January 25–30, 2015, Austin, Texas, USA. pp. 2901–2907. AAAI Press (2015)

20. Neumann, L., Zisserman, A., Vedaldi, A.: Relaxed softmax: efficient confidence auto-calibration for safe pedestrian detection. In: 2018 NIPS MLITS Workshop: Machine Learning for Intelligent Transportation System. OpenReview (2018)
21. Platt, J.: Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In: Advances in Large Margin Classifiers (1999)
22. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)
23. Yeung, M., Rundo, L., Nan, Y., Sala, E., Schönlieb, C., Yang, G.: Calibrating the Dice loss to handle neural network overconfidence for biomedical image segmentation. *Journal of Digital Imaging* (2021)