Class and Region-Adaptive Constraints for Network Calibration

Balamurali Murugesan*, Julio Silva-Rodríguez, Ismail Ben Ayed, and Jose Dolz

ETS Montréal, Canada

*balamurali.murugesan.1@ens.etsmtl.ca

Le génie pour l'industrie

Take-Home Message

Proposed method:

- We extend NACL, a spatial aware segmentation calibration to introduce a region and class wise penalty term, contrasting with NACL which uses a uniform scalar.
- We develop Augmented Lagrangian method (ALM) based solution to obtain the optimal penalty weights instead of empirically finding each one.

Results:

- Our solution provides better segmentation and calibration metrics across benchmarks.
- Our method follows the logit distribution specific to the difficulty of class and region.
- Our method is stable compared to NACL and agnostic to segmentation backbones.

Introduction

- **Motivation:** Deep learning models are poorly calibrated [1].
  - They assign high confidence to incorrect predictions.
  - Inaccurate uncertainty estimates can have severe consequences.
  - High capacity models overfit with the cross entropy loss.
- **Related works:** Model uncertainty can be improved by introducing specialized objective functions [2][3].
  - SVLS [4] proposed a soft labeling technique that used a Gaussian filter smoothed label to capture the spatial uncertainty.
  - Neighbor Aware CaLibration (NACL) [5] provided an alternative to SVLS with a simple linear penalty.
- **Limitations:** NACL employs a single uniform penalty disregarding differences across individual categories, and different regions.
- **Contribution:**
  - We propose a solution that considers the specificities of each category and different regions by introducing independent class and region-wise penalty weights.
  - We find the optimal values of penalties using an Augmented Lagrangian method (ALM).
  - Comprehensive experiments on popular segmentation benchmarks (ACDC, FLARE) and backbones (UNet, nnUNet) demonstrate the superiority of our approach.

Method

- **Preliminaries:**

For a K class segmentation task, let s be the softmax prediction, y be the ground truth, and Ω be the spatial image domain.

NACL minimizes the cross entropy while constraining the logits ($\mathbf{l}$) based on the spatial prior ($\boldsymbol{\tau}$) by the constrained optimization problem below:

$$\min_{\boldsymbol{\theta}} \quad \mathcal{L}_{CE} \quad \text{s.t.} \quad \boldsymbol{\tau} = \mathbf{l},$$

$$\mathcal{L}_{NACL} = \sum_{i \in \Omega} \sum_{k \in K} (-y_k^{(i)} \log(s_k^{(i)}) + \lambda |\tau_k^{(i)} - l_k^{(i)}|).$$

Let us consider the penalty weight for each class in the *Inner* ($\mathcal{I}$) and *Outer* ($\mathcal{O}$) regions, where $\mathcal{I}$ is the surrounding ground truth patch of a given pixel that contains only one category, and $\mathcal{O}$ the patch that has more than one category (typically along the boundaries). The new loss can be then formally defined as:

$$\min_{\boldsymbol{\theta}} \quad \sum_{i \in \Omega} \mathcal{H}(\mathbf{y}^{(i)}, \mathbf{s}^{(i)}) + \sum_{i \in \mathcal{I}} \sum_{k \in K} \lambda_{k,0} |\tau_k^{(i)} - l_k^{(i)}| + \sum_{i \in \mathcal{O}} \sum_{k \in K} \lambda_{k,1} |\tau_k^{(i)} - l_k^{(i)}|,$$

Finding $k \times 2$ optimal λ values manually poses optimization challenges.

- **Solution:**

ALM approaches are optimization techniques that integrate penalties and primal-dual updates to efficiently tackle constrained optimization problems. We use ALM to find the balancing weights (Lagrange multipliers) by reformulating the loss function presented above as:

$$\min_{\boldsymbol{\theta}, \boldsymbol{\lambda}_0, \boldsymbol{\lambda}_1} \quad \sum_{i \in \Omega} \mathcal{H}(\mathbf{y}^{(i)}, \mathbf{s}^{(i)}) + \sum_{i \in \mathcal{I}} \sum_{k \in K} P(\tau_k^{(i)} - l_k^{(i)}, \rho_{k,0}, \lambda_{k,0}) + \sum_{i \in \mathcal{O}} \sum_{k \in K} P(\tau_k^{(i)} - l_k^{(i)}, \rho_{k,1}, \lambda_{k,1}).$$

Algorithm

- Penalty function P is parameterized by (ρ, λ) which are obtained on the validation set after a single training epoch approximately minimizes the target loss function.
- The overall algorithm to simultaneously optimize the cross entropy loss with class and region aware penalty constraints is given below:

Algorithm 1 Augmented Lagrangian Multiplier

for $j = 0, \dots, T$ **do**
  $\boxed{\text{Optimize for } \theta \text{ in the inner loop}}$
  for each mini-batch $\{(\boldsymbol{x}^{(i)}, y^{(i)})\}_{i=1}^{B}$ in $\mathcal{D}_{\text{train}}$ **do**
    $\mathcal{L}_{\text{c}} = \mathcal{H}(\mathbf{y}^{(i)}, \mathbf{s}^{(i)})$ {Cross-entropy}
    $\mathcal{L}_{\text{p}}^{I} = \sum_{i \in \mathcal{I}} \sum_{k \in K} P(\tau_k^{(i)} - l_k^{(i)}, \rho_{k,0}^{(j)}, \lambda_{k,0}^{(j)})$ {Penalty for Inner region}
    $\mathcal{L}_{\text{p}}^{O} = \sum_{i \in \mathcal{O}} \sum_{k \in K} P(\tau_k^{(i)} - l_k^{(i)}, \rho_{k,1}^{(j)}, \lambda_{k,1}^{(j)})$ {Penalty for Outer region}
    $\mathcal{L} = \frac{1}{B}(\mathcal{L}_{\text{c}} + \mathcal{L}_{\text{p}}^{I} + \mathcal{L}_{\text{p}}^{O})$ {Total loss}
    $\theta^{(t+1)} \leftarrow \theta^{(t)} - \alpha \cdot \nabla_{\theta} \mathcal{L}$ {Adjust θ by gradient descent}
  end for
  $\boxed{\text{Optimize for } \lambda \text{ and } \rho \text{ in the outer loop}}$
  $\lambda_{k,r}^{(j+1)} = \frac{1}{|\mathcal{D}_{val}|} \sum_{(\mathbf{x},\mathbf{y}) \in \mathcal{D}_{val}} P'\left(\tau_k - l_k, \rho_{k,r}^{(j)}, \lambda_{k,r}^{(j)}\right)$ {Adjust $\lambda_{k,r}$}
  if $j \geq 1$ **and** $|\tau_k^{(i)} - l_k^{(i)}| > \mu * |\tau_k^{(i)} - l_k^{(i)}|$ **then**
    $\rho_{k,r}^{(j+1)} \leftarrow \gamma \rho_{k,r}^{(j)}$ {Adjust $\rho_{k,r}$}
  else
    $\rho_{k,r}^{(j+1)} \leftarrow \rho_{k,r}^{(j)}$
  end if
end for

Repository: https://github.com/Bala93/MarginLoss

Results

(1) Quantitative performance: Our method consistently yields the best compromise between segmentation and calibration metrics.

	ACDC				FLARE			
	DSC	HD	ECE	TACE	DSC	HD	ECE	TACE
FL	0.620	7.30	0.153	0.224	0.834	6.65	0.053	0.145
ECP	0.782	4.44	0.130	0.151	0.860	5.30	0.037	0.134
LS	0.809	3.30	0.083	0.093	0.860	5.33	0.055	0.050
SVLS	0.824	2.81	0.091	0.138	0.857	5.72	0.039	0.144
MbLS	0.827	2.99	0.103	0.081	0.836	5.75	0.046	0.041
NACL	0.854	2.93	0.068	0.073	0.868	**5.12**	0.033	**0.031**
BWCR	0.841	2.69	**0.051**	0.075	0.848	5.39	**0.029**	0.059
Ours	**0.877**	**1.72**	0.057	**0.058**	**0.876**	5.52	**0.029**	0.033

(2) Benefits compared to NACL: Performance of NACL significantly varies with the value of its penalty weight.

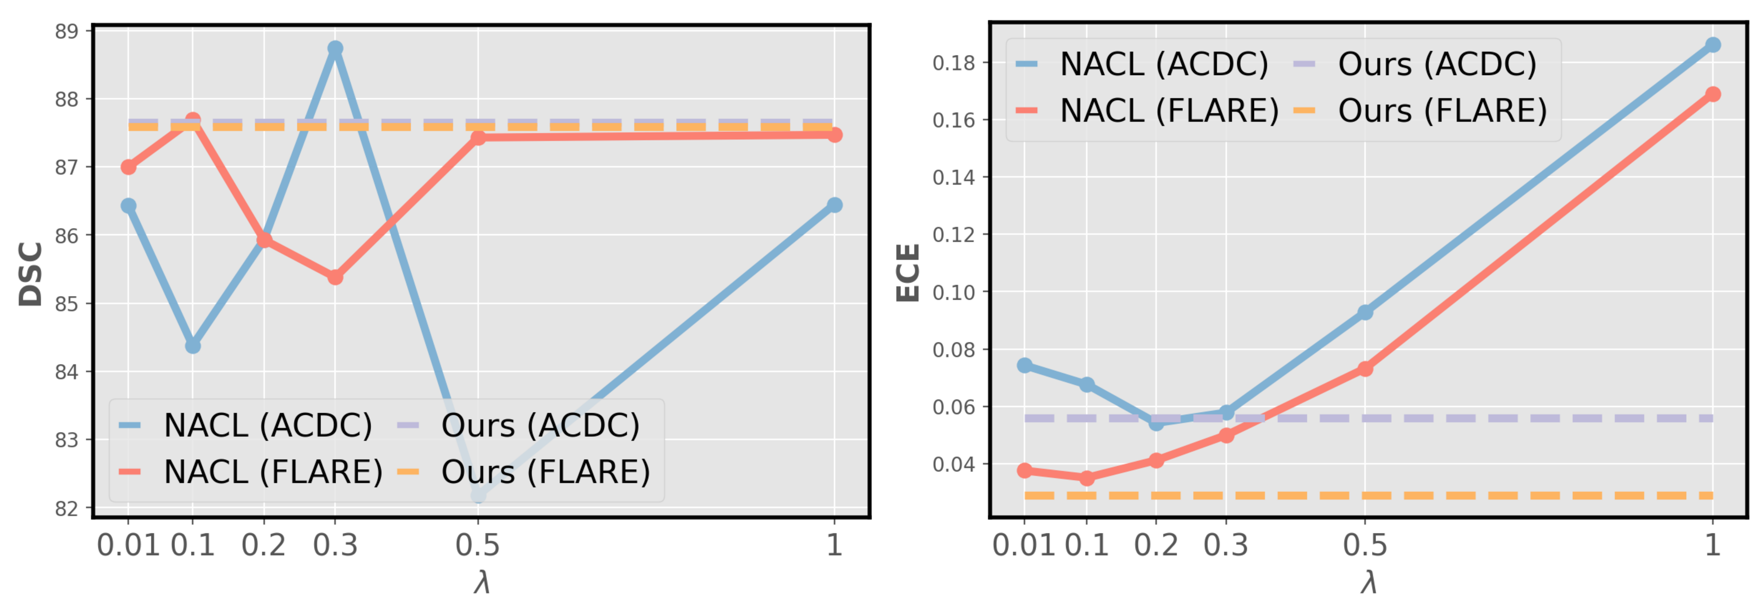

(3) Logit Magnitudes: Our approach provides desirable logit distribution without hyper-parameter optimization as that of NACL.

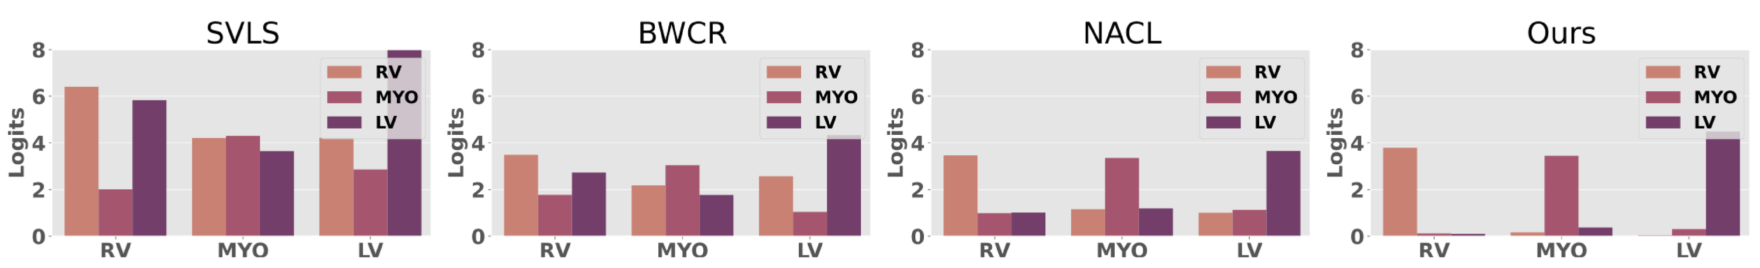

References

[1] Guo et al. On calibration of modern neural networks. In *ICML*, 2017.
[2] Murugesan et al. Calibrating segmentation networks with margin-based label smoothing. *MediA*, 2023.
[3] Karani et al. Boundary-weighted logit consistency for calibration in segmentation. In *MICCAI*, 2023.
[4] Islam et al. Spatially varying label smoothing. In *IPMI*, 2021.
[5] Murugesan et al. Penalty-based constraints for model calibration. In *MICCAI*, 2023.