PEFAT: Boosting Semi-supervised Medical Image Classification via Pseudo-loss Estimation and Feature Adversarial Training [CVPR 2023]

<u>Problems:</u>

1. Finding samples with high-confidence pseudo-labels may lead to the inclusion of incorrectly pseudo-labeled data.

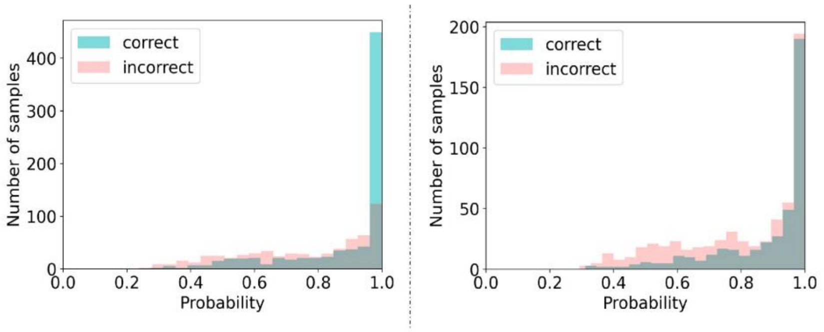

(b) Probability distribution of labeled data (left) and validation data (right), when using the warm-upped model on ISIC2018 dataset.

2. Low-confidence probability samples are frequently disregarded and not employed to their full potential.

To solve problem 1:

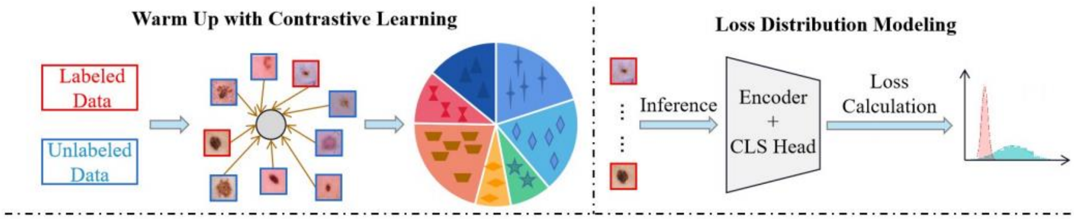

<u>Loss Distribution Modeling on D_l:</u>

It is hard to regard the predicted probability as threshold to collect a clean pseudo-labeled set. Alternatively, wrongly pseudo-labeled samples tend to have a higher loss during the early training, which makes it possible to **distinguish correct and incorrect samples by loss distribution.**

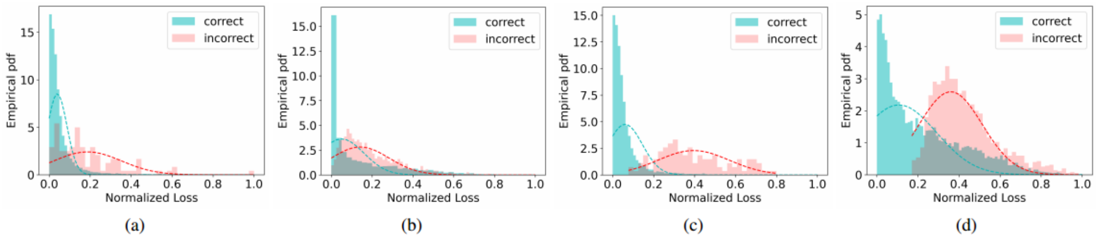

Figure 3. Empirical probability density function (PDF) of the fitted GMM for loss distribution. (a) Training with FixMatch and loss distribution on labeled data; (b) Training with FixMatch and loss distribution on validation data; (c) Training with PEFAT and loss distribution on labeled data; (d) Training with PEFAT and loss distribution on validation data; (a) and (b) show zero-biased loss distribution, which is mainly attributed to over-confident prediction, while (c) and (d) present dividable distribution for pseudo-labeled data with correct and incorrect pseudo-labels, validating the effectiveness of cross pseudo-loss estimation.

Based on the above observation, we assume that the overall loss distribution is composed of two normal distributions and further utilize the Gaussian Mixture Model (GMM) to fit the loss distribution on D_l. Formally, the instance-wise loss and probability density function (pdf) of GMM on loss ℓ_i can be formulated as:

$$\mathcal{L}(\mathcal{D}_l|h_\theta) = \{-y_i \log(h_\theta(\hat{y}_i|x_i)), x_i \in \mathcal{D}_l\}$$

$$\mathcal{I}(\ell_i) = \sum_{k=0}^{K-1} \pi_k I_k(\ell_i|\mu_k, \Sigma_k), \ell_i \in \mathcal{L}(\mathcal{D}_l|h_\theta)$$

We use the Expectation Maximization (EM) algorithm to fit the GMM with the loss observation on D_l, and the optimization procedure is maximizing the log-likelihood:

$$\hat{\theta}_{GMM} = \arg\max_{\theta_{GMM}}[\log \prod_{i=1}^{N_l} \mathcal{I}(\ell_i|\theta_{GMM})]$$

where $\theta_{GMM} = \{\pi_k, \mu_k, \Sigma_k\}, 0 \leq k \leq K-1.$

GMM perceives the prior loss distribution on D_l, and is able to distinguish trustworthy pseudo-labeled samples by pseudo-loss distribution.

<u>Trustworthy Pseudo-labeled Data Selection:</u>

Cross Pseudo-loss Estimation on D_u.

$$\hat{y}_i^{1\to2} = \arg\max(h_\theta(A_{s1}(u_i))) \quad \ell_i^{1\to2} = -\hat{y}_i^{1\to2}\log(h_\theta(A_{s2}(u_i)))$$

$$\hat{y}_i^{2\to1} = \arg\max(h_\theta(A_{s2}(u_i))) \quad \ell_i^{2\to1} = -\hat{y}_i^{2\to1}\log(h_\theta(A_{s1}(u_i)))$$

Pseudo-labeled Sample Selection.

Based on $\ell_i^{1\to2}$, $\ell_i^{2\to1}$ and the fitted GMM, we can select trustworthy pseudo-labeled sample u_i by the posterior probability.

$$p_{gmm} = \mathcal{I}(I_k|(\eta\ell_i^{1\to2} + (1-\eta)\ell_i^{2\to1}))$$

where k = 0(1) stands for correct (incorrect) pseudo-loss component.

To solve problem 2:

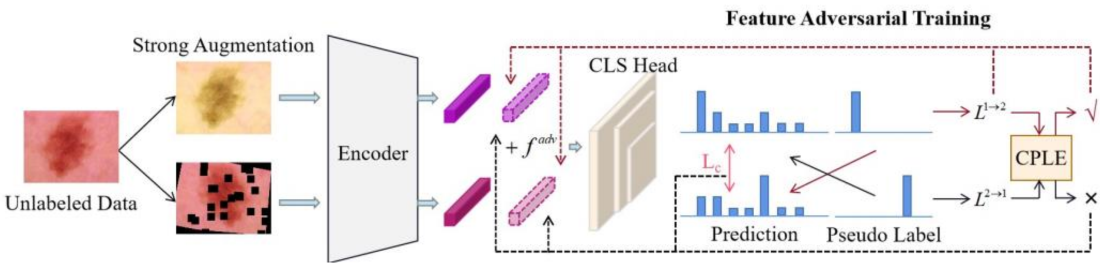

$$L_{FAT} = J(h_\theta(p_i|z_i + r_{i1}^{adv}), h_\theta(p_i^+|z_i^+ + r_{i2}^{adv}))$$