FREEMATCH: Self-adaptive Thresholding For SSL [ICLR 2023]

use q_b and Q_b to denote abbreviation of $p_m(y|\omega(u_b))$ and $p_m(y|\Omega(u_b))$, respectively.

1. Self-Adaptive Thresholding (SAT)

Global: (1) related to the model's confidence on unlabeled data (2) stably increase

$$\tau_t = \begin{cases} \frac{1}{C}, & \text{if } t = 0, \\ \lambda \tau_{t-1} + (1-\lambda) \frac{1}{\mu B} \sum_{b=1}^{\mu B} \max(q_b), & \text{otherwise,} \end{cases}$$

Local: modulate the global threshold in a class-specific fashion to account for the intra-class diversity and the possible class adjacency

$$\tilde{p}_t(c) = \begin{cases} \frac{1}{C}, & \text{if } t = 0, \\ \lambda \tilde{p}_{t-1}(c) + (1-\lambda) \frac{1}{\mu B} \sum_{b=1}^{\mu B} q_b(c), & \text{otherwise,} \end{cases}$$

where $\tilde{p}_t = [\tilde{p}_t(1), \tilde{p}_t(2), \ldots, \tilde{p}_t(C)]$ is the list containing all $\tilde{p}_t(c)$.

Final: integrate the global and local thresholds

$$\tau_t(c) = \text{MaxNorm}(\tilde{p}_t(c)) \cdot \tau_t = \frac{\tilde{p}_t(c)}{\max\{\tilde{p}_t(c) : c \in [C]\}} \cdot \tau_t, \quad (7)$$

where MaxNorm is the Maximum Normalization (i.e., $x' = \frac{x}{\max(x)}$). Finally, the unsupervised training objective $\mathcal{L}_u$ at the t-th iteration is:

$$\mathcal{L}_u = \frac{1}{\mu B} \sum_{b=1}^{\mu B} \mathbb{1}(\max(q_b) > \tau_t(\arg\max(q_b))) \cdot \mathcal{H}(\hat{q}_b, Q_b). \quad (8)$$

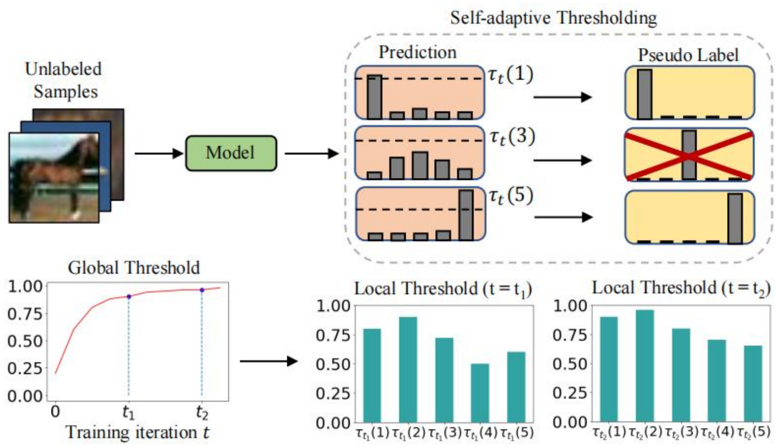

2. Self-Adaptive Fairness

Encourage the model to make diverse predictions for each class (Fairness)

(1) Instead of using a uniform prior, we use the EMA of model predictions from Eq.6 as an estimate of the expectation of prediction distribution over unlabeled data

(2) Considering that the underlying pseudo label distribution may not be uniform, we propose to modulate the fairness objective in a self-adaptive way, i.e., normalizing the expectation of probability by the histogram distribution of pseudo labels to counter the negative effect of imbalance as

$$\overline{p} = \frac{1}{\mu B} \sum_{b=1}^{\mu B} \mathbb{1}\left(\max\left(q_b\right) \geq \tau_t(\arg\max\left(q_b\right)\right) Q_b,$$

$$\overline{h} = \text{Hist}_{\mu B}\left(\mathbb{1}\left(\max\left(q_b\right) \geq \tau_t(\arg\max\left(q_b\right)\right) \hat{Q}_b\right).$$

$$\tilde{h}_t = \lambda \tilde{h}_{t-1} + (1-\lambda)\,\text{Hist}_{\mu B}\left(\hat{q}_b\right).$$

The self-adaptive fairness (SAF) L_f at the t-th iteration is formulated as:

$$\mathcal{L}_f = -\mathcal{H}\left(\text{SumNorm}\left(\frac{\tilde{p}_t}{\tilde{h}_t}\right), \text{SumNorm}\left(\frac{\bar{p}}{\bar{h}}\right)\right),$$

The overall objective for FreeMatch at t-th iteration is:

$$\mathcal{L} = \mathcal{L}_s + w_u \mathcal{L}_u + w_f \mathcal{L}_f,$$