# Cost-Sensitive Self-Training for Optimizing Non-Decomposable Metrics 

Anonymous Author(s)<br>Affiliation

## Contents

## A Limitations and Negative Societal Impacts 2

A. 1 Limitations of our Work ..... 2
A. 2 Negative Societal Impact ..... 2
B Connection between Minimization of Weighted Consistency Regularizer Loss (Eq. (6)) and Theoretical Weighted Consistency ( $R_{\mathcal{B}, w}(F)$ in Sec. 3.4) ..... 2
C Additional Examples and Proof of Theorem 5 ..... 3
C. 1 Examples for Theoretical Assumptions ..... 3
C. 2 Proof of Theorem 5 Assuming a Lemma ..... 4
C. 3 Proof of Lemma 10 ..... 5
C. 4 Miscellaneous Proofs for Examples ..... 7
D All-Layer Margin Generalization Bounds ..... 7
E Proof of Proposition 6 ..... 8
F Notation ..... 9
G Code, License, Assets and Computation Requirements ..... 9
G. 1 Code and Licenses of Assets ..... 9
G. 2 Computational Requirements ..... 9
H Objective ..... 9
H. 1 Logit Adjusted Weighted Consistency Regularizer ..... 9
H. $2 \operatorname{CSST}($ FixMatch) ..... 11
H. 3 CSST(UDA) ..... 12
I Threshold mechanism for diagonal Gain Matrix ..... 13
J Dataset ..... 13
L Details of Experiments and Hyper-parameters ..... 14
M Statistical Analysis ..... 14
N Additional Details ..... 15
N. 1 Formal Statement Omitted in Sec. 2.2 ..... 15
N. 2 Comparison with the ( $a, \widetilde{c}$ )-expansion Property in [14] ..... 15
N. 3 Comparison of the Theoretical Assumptions with that of [14] ..... 16

## A Limitations and Negative Societal Impacts

## A. 1 Limitations of our Work

At this point we only consider optimizing objectives through CSST which can be written as a linear combination of entries of confusion matrix. Although there are important metrics like Recall, Coverage etc. which can be expressed as a linear form of confusion matrix. However, there do exist important metrics like Intersection over Union (IoU), Q-Mean etc. which we don't consider in the current work. We leave this as an open direction for further work.
Also in this work we considered datasets where unlabeled data distribution doesn't significantly differ from the labeled data distribution, developing robust methods which can also take into account the distribution shift between unlabeled and labeled is an interesting direction for future work.

## A. 2 Negative Societal Impact

Our work has application in fairness domain [3], where it can be used to improve performance of minority sub-groups present in data. These fairness objectives can be practically enforced on neural networks through the proposed CSST framework. However these same algorithms can be tweaked to artifically induce bias in decision making of trained neural networks, for example by ignoring performance of models on certain subgroups. Hence, we suggest deployment of these models after through testing on all sub groups of data.

## B Connection between Minimization of Weighted Consistency Regularizer Loss (Eq. (6)) and Theoretical Weighted Consistency ( $R_{\mathcal{B}, w}(F)$ in Sec. 3.4)

In this section, we show that minimization of weighted consistency regularizer Eq. (6) and that of theoretical weighted consistency regularizer $R_{\mathcal{B}, w}(F)$ can be related to CSL.
First, we consider $R_{\mathcal{B}, w}(F)$. Using strong augmentation $\mathcal{A}$, (theoretical) weighted consistency regularizer $R_{\mathcal{B}, w}(F)$ is approximated as $R_{\mathcal{B}, w}(F) \approx \sum_{i, j \in[K]} w_{i j} \mathbf{E}_{x \sim P_{i}}[\mathbf{1}(F(\mathcal{A}(x)) \neq F(x))]$. Noting that $\mathbf{1}(F(\mathcal{A}(x)) \neq F(x)) \leq \mathbf{1}(F(\mathcal{A}(x)) \neq j)+\mathbf{1}(F(x) \neq j)$ for any $j, x$, this value is bounded as follows:

$$
\begin{aligned}
R_{\mathcal{B}, w}(F) & \approx \sum_{i, j \in[K]} w_{i j} \mathbf{E}_{x \sim P_{i}}[\mathbf{1}(F(\mathcal{A}(x)) \neq F(x))] \\
& \leq \sum_{i, j \in[K]} w_{i j} \mathbf{E}_{x \sim P_{i}}[\mathbf{1}(F(\mathcal{A}(x) \neq j))]+\sum_{i, j \in[K]} w_{i j} \mathbf{E}_{x \sim P_{i}}[\mathbf{1}(F(x) \neq j)] .
\end{aligned}
$$

If we focus on samples $x$ with high confidence in model predictions, then the latter term $\sum_{i, j \in[K]} w_{i j} \mathbf{E}_{x}[\mathbf{1}(F(x) \neq j)]$ is negligible. Therefore, minimization of (an empirical approximation of) $R_{\mathcal{B}, w}(F)$ on these samples is approximately equivalent to CSL, i.e., the following problem:

$$
\begin{equation*}
\left.\min _{F} \sum_{i, j \in[K]} w_{i j} \mathbf{E}_{x}[\mathbf{1}(F(\mathcal{A}(x)) \neq j))\right] . \tag{8}
\end{equation*}
$$

The above CSL is shown to be calibrated with the loss [9] $\ell^{\mathrm{wt}}\left(y, p_{m}(x)\right)$ (also used in Prop. 7) given below:

$$
\ell^{\mathrm{wt}}\left(y, p_{m}(x)\right)=-\sum_{i \in[K]} G_{y i} \log \left(p_{m}(x)_{i}\right)
$$

Next, we relate Eq. (6) to CSL. If we denote pseudo label $\hat{p}_{m}(x)$ by $y$ in Eq. (6), then we see that Eq. (6) is identical to $\ell^{\mathrm{wt}}\left(y, p_{m}(\mathcal{A}(x))\right)$. By (the proof of) [8, Proposition 4], by minimizing $\ell^{\mathrm{wt}}\left(y, p_{m}(\mathcal{A}(x))\right)$, we obtain a Bayes optimal classifier $F(\mathcal{A}(x))$, where $F$ is the classifier defined by the model $p_{m}$. If $w$ is the corresponding weight to the gain matrix $\mathbf{G}$, then classifier $F$ gives a solution to the CSL (8). Thus, we can relate minimization of weighted consistency regularizer Eq. (6) to that of theoretical weighted consistency regularizer $R_{\mathcal{B}, w}(F)$ through the CSL (8).

## C Additional Examples and Proof of Theorem 5

In this section, we provide some examples for assumptions introduced in Sec. 3 and a proof of Theorem 5. We provide proof of these examples in Sec. C.4.

## C. 1 Examples for Theoretical Assumptions

The following example (Example 3) shows that the $c$-expansion (Definition 2) property is satisfied for mixtures of Gaussians and mixtures of manifolds.

EXAMPLE 3. By [14, Examples 3.4, 3.5], the $c$-expansion property is satisfied for mixtures of isotropic Gaussian distributions and mixtures of manifolds. More precisely, in the case of mixtures of isotropic Gaussian distributions, i.e., if $Q$ is given as mixtures of $\mathcal{N}\left(\tau_{i}, \frac{1}{d} I_{d \times d}\right)$ for $i=1, \ldots, n$ with some $n \in \mathbb{Z}_{\geq 1}$ and $\tau_{i} \in \mathbb{R}^{d}$, and $\mathcal{B}(x)$ is an $\ell_{2}$-ball with radius $r$ then by [1, (13)] and [14, Section B.2], $Q$ satisfies the $c$-expansion property with $c(p)=R_{h}(p) / p$ for $p>0$ and $h=2 r \sqrt{d}$ (c.f., [14, section B.2]). Here $R_{h}(p)=\Phi\left(\Phi^{-1}(p)+h\right)$ and $\Phi$ is the cumulative distribution function of the standard normal distribution on $\mathbb{R}$.

In Sec. 3.4, we required the assumption that $\gamma>3$ (Assumption 4) and remarked that it roughly requires $\operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)$ is "small". The following example provides explicit conditions for $\operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)$ that satisfy the assumption using a toy example.

EXAMPLE 8. Using a toy example provided in Example 3, we provide conditions that satisfy the assumption $\gamma>3$ approximately. To explain the assumption $\gamma>3$, we assume that $\mathcal{P}_{w}$ is given as a mixture of isotropic Gaussians and $\mathcal{B}(x)$ is $\ell_{2}$-ball with radius $r$ as in Example 3. Furthermore, we assume that $|w|_{1}=1$ and $\operatorname{Err}_{w}\left(F^{*}\right)$ is sufficiently small compared to $\operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)$. Then, $p_{w}=\operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)+\operatorname{Err}_{w}\left(F^{*}\right) \approx \operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)$. Using this approximation, since $\mathcal{P}_{w}$ satisfies the $c$-expansion property with $c(p)=R_{2 r \sqrt{d}}(p) / p$, if $r=\frac{1}{2 \sqrt{d}}$ then, the condition $\gamma>3$ is satisfied when $\operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)<0.17$. If $r=\frac{3}{2 \sqrt{d}}$ then, the condition $\gamma>3$ is satisfied when $\operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)<0.33$.

In Assumption 1, we assumed that both of $\operatorname{Err}_{w}\left(F^{*}\right)$ and $R_{\mathcal{B}, w}\left(F^{*}\right)$ are small. The following example suggests the validity of this assumption.

ExAMPLE 9. In this example, we assume $w$ is a diagonal matrix $\operatorname{diag}\left(w_{1}, \ldots, w_{K}\right)$. For simplicity, we normalize $w$ so that $\sum_{i \in[K]} w_{i}=1$. As in [14, Example 3.4], we assume that $P_{i}$ is given as isotropic Gaussian distribution $\mathcal{N}\left(\tau_{i}, \frac{1}{d} I_{d \times d}\right)$ with $\tau_{i} \in \mathbb{R}^{d}$ for $i=1, \ldots, K$ and $\mathcal{B}(x)$ is an $\ell^{2}$-ball with radius $\frac{1}{2 \sqrt{d}}$. Furthermore, we assume $\inf _{1 \leq i<j \leq K}\left\|\tau_{i}-\tau_{j}\right\|_{2} \gtrsim \frac{\sqrt{\log d}}{\sqrt{d}}$ and $\sup _{i, j \in[K]} \frac{w_{i}}{w_{j}}=$ $o(d)$, where the latter assumption is valid for high dimensional datasets (e.g., image datasets). Then it can be proved that there exists a classifier $F$ such that $R_{\mathcal{B}, w}(F)=O\left(\frac{1}{d^{c}}\right)$ and $\operatorname{Err}_{w}(F)=O\left(\frac{1}{d^{c}}\right)$, where $c>0$ is a constant (we can take $F$ as the Bayes-optimal classifier for $\operatorname{Err}_{w}$ ). Thus, this suggests that Assumption 1 is valid for datasets with high dimensional instances.

The statement of Example 8 follows from numerical computation of $R_{2 r \sqrt{d}}(p) / p$. We provide proofs of Examples 3 and 9 in Sec. C.4.

## C. 2 Proof of Theorem 5 Assuming a Lemma

Theorem 5 can be deduced from the following lemma (by taking $H=F^{*}$ and $\mathcal{L}_{Q, H}(\widehat{F}) \leq$ $\mathcal{L}_{Q, H}\left(F^{*}\right)$ ), which provides a similar result to [14, Lemma A.8].
Lemma 10. Let $H$ be a classifier and $Q$ a probability measure on $\mathcal{X}$ satisfying c-expansion property. We put $\gamma_{H}=c\left(Q\left(\left\{x \in \mathcal{X}: F_{\mathrm{pl}}(x) \neq H(x)\right\}\right)\right)$. For a classifier $F$, we define $\mathcal{S}_{\mathcal{B}}(F)$ by $\mathcal{S}_{\mathcal{B}}(F)=\left\{x \in \mathcal{X}: F(x)=F\left(x^{\prime}\right) \quad \forall x^{\prime} \in \mathcal{B}(x)\right\}$. For a classifier $F$, we define $\mathcal{L}_{Q, H}(F)$ by

$$
\begin{aligned}
\frac{\gamma_{H}+1}{\gamma_{H}-1} Q(\{x \in & \left.\left.\mathcal{X}: F(x) \neq F_{\mathrm{pl}}(x)\right\}\right) \\
& +\frac{2 \gamma_{H}}{\gamma_{H}-1} Q\left(\mathcal{S}_{\mathcal{B}}^{c}(F)\right)+\frac{2 \gamma_{H}}{\gamma_{H}-1} Q\left(\mathcal{S}_{\mathcal{B}}^{c}(H)\right)-Q\left(\left\{x \in \mathcal{X}: F_{\mathrm{pl}}(x) \neq H(x)\right\}\right),
\end{aligned}
$$

where $\mathcal{S}_{\mathcal{B}}^{c}(F)$ denotes the complement of $\mathcal{S}_{\mathcal{B}}(F)$. Then, we have $Q(\{x \in \mathcal{X}: F(x) \neq H(x)\}) \leq$ $\mathcal{L}_{Q, H}(F)$ for any classifier $F$.

In this subsection, we provide a proof of Theorem 5 assuming Lemma 10. We provide a proof of the lemma in the next subsection. For a classifier $F$, we define $\mathcal{M}(F)$ as $\left\{x \in \mathcal{X}: F(x) \neq F^{*}(x)\right\}$ and $\mathcal{M}_{\mathrm{pl}}(F)$ as $\left\{x \in \mathcal{X}: F(x) \neq F_{\mathrm{pl}}(x)\right\}$. We define $\widetilde{\mathcal{L}}_{w}(F)$ by

$$
\widetilde{\mathcal{L}}_{w}(F)=\mathcal{L}_{w}(F)+\frac{2 \gamma}{\gamma-1} R_{\mathcal{B}, w}\left(F^{*}\right)-\mathcal{P}_{w}\left(\left\{x \in \mathcal{X}: F_{\mathrm{pl}}(x) \neq F^{*}(x)\right\}\right)
$$

We note that $\widetilde{\mathcal{L}}_{w}(F)-\mathcal{L}_{w}(F)$ does not depend on $F$.

Proof of Theorem 5. We let $Q=\mathcal{P}_{w}$ and $H=F^{*}$ in Lemma 10 and denote $\gamma_{H}$ in the lemma by $\gamma^{\prime}$. Since $w_{i j} \geq 0$ and $\mathbf{E}_{x \sim P_{i}}\left[\mathbb{1}\left(F_{\mathrm{pl}}(x) \neq F^{*}(x)\right)\right] \leq \mathbf{E}_{x \sim P_{i}}\left[\mathbb{1}\left(F_{\mathrm{pl}}(x) \neq j\right)\right]+\mathbf{E}_{x \sim P_{i}}\left[\mathbb{1}\left(F^{*}(x) \neq j\right)\right]$ for any $i, j$, we have the following:

$$
\begin{aligned}
|w|_{1} \mathcal{P}_{w}\left(\mathcal{M}\left(F_{\mathrm{pl}}\right)\right) & =\sum_{i, j \in[K]} w_{i j} \mathbf{E}_{x \sim P_{i}}\left[\mathbb{1}\left(F_{\mathrm{pl}}(x) \neq F^{*}(x)\right)\right] \\
& \leq \sum_{i, j \in[K]} w_{i j}\left\{\mathbf{E}_{x \sim P_{i}}\left[\mathbb{1}\left(F_{\mathrm{pl}}(x) \neq j\right)\right]+\mathbf{E}_{x \sim P_{i}}\left[\mathbb{1}\left(F^{*}(x) \neq j\right)\right]\right\} \\
& =\operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)+\operatorname{Err}_{w}\left(F^{*}\right) .
\end{aligned}
$$

Thus, we obtain $\mathcal{P}_{w}\left(\mathcal{M}\left(F_{\mathrm{pl}}\right)\right) \leq p_{w}$. Because $c$ is non-increasing, we have $\gamma \leq \gamma^{\prime}$. We note that

$$
\begin{aligned}
\operatorname{Err}_{w}(F) & =\sum_{i, j \in[K]} w_{i j} \mathbf{E}_{x \sim P_{i}}[\mathbb{1}(F(x) \neq j)] \\
& \leq \sum_{i, j \in[K]} w_{i j} \mathbf{E}_{x \sim P_{i}}\left[\mathbb{1}\left(F(x) \neq F^{*}(x)\right)\right]+\sum_{i, j \in[K]} w_{i j} \mathbf{E}_{x \sim P_{i}}\left[\mathbb{1}\left(F^{*}(x) \neq j\right)\right] \\
& =|w|_{1} \mathcal{P}_{w}(\mathcal{M}(F))+\operatorname{Err}_{w}\left(F^{*}\right) .
\end{aligned}
$$

By this inequality and Lemma 10, the error is upper bounded as follows:

$$
\begin{aligned}
& \operatorname{Err}_{w}(F) \leq \operatorname{Err}_{w}\left(F^{*}\right)+\frac{\gamma^{\prime}+1}{\gamma^{\prime}-1}|w|_{1} \mathcal{P}_{w}\left(\mathcal{M}_{\mathrm{pl}}(F)\right) \\
& \quad+\frac{2 \gamma^{\prime}}{\gamma^{\prime}-1}|w|_{1} \mathcal{P}_{w}\left(\mathcal{S}_{\mathcal{B}}^{c}(F)\right)+\frac{2 \gamma^{\prime}}{\gamma^{\prime}-1}|w|_{1} \mathcal{P}_{w}\left(\mathcal{S}_{\mathcal{B}}^{c}\left(F^{*}\right)\right)-|w|_{1} \mathcal{P}_{w}\left(\mathcal{M}\left(F_{\mathrm{pl}}\right)\right)
\end{aligned}
$$

Since $\gamma \leq \gamma^{\prime}$, we obtain

$$
\begin{align*}
\operatorname{Err}_{w}(F) \leq & \operatorname{Err}_{w}\left(F^{*}\right)+\frac{\gamma+1}{\gamma-1}|w|_{1} \mathcal{P}_{w}\left(\mathcal{M}_{\mathrm{pl}}(F)\right) \\
& +\frac{2 \gamma}{\gamma-1}|w|_{1} \mathcal{P}_{w}\left(\mathcal{S}_{\mathcal{B}}^{c}(F)\right)+\frac{2 \gamma}{\gamma-1}|w|_{1} \mathcal{P}_{w}\left(\mathcal{S}_{\mathcal{B}}^{c}\left(F^{*}\right)\right)-|w|_{1} \mathcal{P}_{w}\left(\mathcal{M}\left(F_{\mathrm{pl}}\right)\right) . \tag{9}
\end{align*}
$$

Therefore, we have

$$
\begin{aligned}
\mathcal{N}(V) \cap V^{c} & =\mathcal{N}(V) \cap V^{c} \cap\left(S \cup S^{c}\right) \\
& =\left(\mathcal{N}(V) \cap V^{c} \cap S\right) \cup\left(\mathcal{N}(V) \cap V^{c} \cap S^{c}\right) \\
& \subseteq \mathcal{M}_{\mathrm{pl}}(F) \cup S^{c} .
\end{aligned}
$$

By definition of $\mathcal{L}_{w}$ and letting $F=\widehat{F}$, we have the following:

$$
\begin{aligned}
& \operatorname{Err}_{w}(\widehat{F}) \leq \operatorname{Err}_{w}\left(F^{*}\right)+\mathcal{L}_{w}(\widehat{F})+\frac{2 \gamma}{\gamma-1} R_{\mathcal{B}, w}\left(F^{*}\right)-|w|_{1} \mathcal{P}_{w}\left(\left\{x \in \mathcal{X}: F_{\mathrm{pl}}(x) \neq F^{*}(x)\right\}\right) \\
& \leq \operatorname{Err}_{w}\left(F^{*}\right)+\mathcal{L}_{w}\left(F^{*}\right)+\frac{2 \gamma}{\gamma-1} R_{\mathcal{B}, w}\left(F^{*}\right)-|w|_{1} \mathcal{P}_{w}\left(\left\{x \in \mathcal{X}: F_{\mathrm{pl}}(x) \neq F^{*}(x)\right\}\right) \\
& =\operatorname{Err}_{w}\left(F^{*}\right)+\frac{2}{\gamma-1}|w|_{1} \mathcal{P}_{w}\left(\mathcal{M}\left(F_{\mathrm{pl}}\right)\right)+\frac{4 \gamma}{\gamma-1} R_{\mathcal{B}, w}\left(F^{*}\right) \\
& \leq \operatorname{Err}_{w}\left(F^{*}\right)+\frac{2}{\gamma-1}\left(\operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)+\operatorname{Err}_{w}\left(F^{*}\right)\right)+\frac{4 \gamma}{\gamma-1} R_{\mathcal{B}, w}\left(F^{*}\right) \\
& =\frac{2}{\gamma-1} \operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)+\frac{\gamma+1}{\gamma-1} \operatorname{Err}_{w}\left(F^{*}\right)+\frac{4 \gamma}{\gamma-1} R_{\mathcal{B}, w}\left(F^{*}\right)
\end{aligned}
$$

Here, the second inequality holds since $\widehat{F}$ is a minimizer of $\mathcal{L}_{w}$, the third inequality follows from $\mathbb{1}\left(F^{*}(x) \neq F_{\mathrm{pl}}(x)\right) \leq \mathbb{1}\left(F^{*}(x) \neq j\right)+\mathbb{1}\left(F_{\mathrm{pl}}(x) \neq j\right)$ for any $j$. Thus, we have the assertion of the theorem.

## C. 3 Proof of Lemma 10

We decompose $\mathcal{M}(F) \cap \mathcal{S}_{\mathcal{B}}(F) \cap \mathcal{S}_{\mathcal{B}}(H)$ into the following three sets:

$$
\begin{aligned}
& \mathcal{N}_{1}=\left\{x \in \mathcal{S}_{\mathcal{B}}(F) \cap \mathcal{S}_{\mathcal{B}}(H): F(x)=F_{\mathrm{pl}}(x), \text { and } F_{\mathrm{pl}}(x) \neq H(x)\right\} \\
& \mathcal{N}_{2}=\left\{x \in \mathcal{S}_{\mathcal{B}}(F) \cap \mathcal{S}_{\mathcal{B}}(H): F(x) \neq F_{\mathrm{pl}}(x), F_{\mathrm{pl}}(x) \neq H(x), \text { and } F(x) \neq H(x)\right\}, \\
& \mathcal{N}_{3}=\left\{x \in \mathcal{S}_{\mathcal{B}}(F) \cap \mathcal{S}_{\mathcal{B}}(H): F(x) \neq F_{\mathrm{pl}}(x) \text { and } F_{\mathrm{pl}}(x)=H(x)\right\}
\end{aligned}
$$

LEMMA 11. Let $S=\mathcal{S}_{\mathcal{B}}(F) \cap \mathcal{S}_{\mathcal{B}}(H)$ and $V=\mathcal{M}(F) \cap \mathcal{M}\left(F_{\mathrm{pl}}\right) \cap S$. Then, we have $\mathcal{N}(V) \cap$ $\mathcal{M}^{c}(F) \cap S=\emptyset$ and $\mathcal{N}(V) \cap \mathcal{M}^{c}\left(F_{\mathrm{pl}}\right) \cap S \subseteq \mathcal{M}_{\mathrm{pl}}(F)$. Here $\mathcal{M}_{\mathrm{pl}}(F)$ is defined as $\{x \in \mathcal{X}$ : $\left.F(x) \neq F_{\mathrm{pl}}(x)\right\}$.

Proof. We take any element $x$ in $\mathcal{N}(V) \cap S$. Since $x \in \mathcal{N}(V)$ and definition of neighborhoods, there exists $x^{\prime} \in \mathcal{M}\left(F_{\mathrm{pl}}\right) \cap \mathcal{M}(F) \cap S$ such that $\mathcal{B}(x) \cap \mathcal{B}\left(x^{\prime}\right) \neq \emptyset$. Since $x, x^{\prime} \in \mathcal{S}_{\mathcal{B}}(F), F$ takes the same values on $\mathcal{B}(x)$ and $\mathcal{B}\left(x^{\prime}\right)$. By $\mathcal{B}(x) \cap \mathcal{B}\left(x^{\prime}\right) \neq \emptyset, F$ takes the same value on $\mathcal{B}(x) \cup \mathcal{B}\left(x^{\prime}\right)$. It follows that $F(x)=F\left(x^{\prime}\right)$. Since we have $x, x^{\prime} \in \mathcal{S}_{\mathcal{B}}(H)$, similarly, we see that $H(x)=H\left(x^{\prime}\right)$. By $x^{\prime} \in \mathcal{M}(F)$, we have $F(x)=F\left(x^{\prime}\right) \neq H\left(x^{\prime}\right)=H(x)$. Thus, we see that $\mathcal{N}(V) \cap \mathcal{M}^{c}(F) \cap S=\emptyset$. We assume $x \in \mathcal{N}(V) \cap \mathcal{M}^{c}\left(F_{\mathrm{pl}}\right) \cap S$. Then, we have $F(x) \neq H(x)$ and $F_{\mathrm{pl}}(x)=H(x)$. Therefore, we obtain $F(x) \neq F_{\mathrm{pl}}(x)$. This completes the proof.
Lemma 12. Suppose that assumptions of Lemma 10 hold. We define $q$ as follows:

$$
q=\frac{Q\left(\mathcal{M}_{\mathrm{pl}}(F) \cup \mathcal{S}_{\mathcal{B}}^{c}(F) \cup \mathcal{S}_{\mathcal{B}}^{c}(H)\right)}{\gamma_{H}-1}
$$

Then, we have $Q\left(\mathcal{S}_{\mathcal{B}}(F) \cap \mathcal{S}_{\mathcal{B}}(H) \cap \mathcal{M}\left(F_{\mathrm{pl}}\right) \cap \mathcal{M}(F)\right) \leq q$. In particular, noting that $\mathcal{N}_{1} \cup \mathcal{N}_{2} \subseteq$ $\mathcal{S}_{\mathcal{B}}(F) \cap \mathcal{S}_{\mathcal{B}}(H) \cap \mathcal{M}\left(F_{\mathrm{pl}}\right) \cap \mathcal{M}(F)$, we have $Q\left(\mathcal{N}_{1} \cup \overline{\mathcal{N}}_{2}\right) \leq q$.

Proof. We let $S=\mathcal{S}_{\mathcal{B}}(F) \cap \mathcal{S}_{\mathcal{B}}(H)$ and $V=\mathcal{M}(F) \cap \mathcal{M}\left(F_{\mathrm{pl}}\right) \cap S$ as before. Then by Lemma 11, we have

$$
\begin{aligned}
\mathcal{N}(V) \cap V^{c} \cap S & =\left(\mathcal{N}(V) \cap \mathcal{M}^{c}(F) \cap S\right) \cup\left(\mathcal{N}(V) \cap \mathcal{M}^{c}\left(F_{\mathrm{pl}}\right) \cap S\right) \\
& \subseteq \emptyset \cup \mathcal{M}_{\mathrm{pl}}(F)=\mathcal{M}_{\mathrm{pl}}(F)
\end{aligned}
$$

Thus, by the $c$-expansion property, we have

$$
\begin{aligned}
Q\left(\mathcal{M}_{\mathrm{pl}}(F) \cup S^{c}\right) & \geq Q\left(\mathcal{N}(V) \cap V^{c}\right) \\
& \geq Q(\mathcal{N}(V))-Q(V) \\
& \geq(c(Q(V))-1) Q(V) .
\end{aligned}
$$

Since $V \subseteq \mathcal{M}\left(F_{\mathrm{pl}}\right), c$ is non-increasing, and $\gamma_{H}>1$, we have $Q(V) \leq Q\left(\mathcal{M}_{\mathrm{pl}}(F) \cup S^{c}\right) /\left(\gamma_{H}-\right.$ $1) \leq q$. This completes the proof.

The following lemma provides an upper bound of $Q\left(\mathcal{N}_{3}\right)$.
Lemma 13. Suppose that the assumptions of Lemma 10 hold. We have

$$
Q\left(\mathcal{N}_{3}\right) \leq q+Q\left(\mathcal{S}_{\mathcal{B}}^{c}(F) \cup \mathcal{S}_{\mathcal{B}}^{c}(H)\right)+Q\left(\mathcal{M}_{\mathrm{pl}}(F)\right)-Q\left(\mathcal{M}\left(F_{\mathrm{pl}}\right)\right)
$$

where $q$ is defined by (10).

Proof. We let $S=\mathcal{S}_{\mathcal{B}}(F) \cap \mathcal{S}_{\mathcal{B}}(H)$. First, we prove

$$
\begin{equation*}
\mathcal{N}_{3} \sqcup\left(\mathcal{M}_{\mathrm{pl}}^{c}(F) \cap S\right)=\mathcal{N}_{1} \sqcup\left(\mathcal{M}^{c}\left(F_{\mathrm{pl}}\right) \cap S\right) . \tag{11}
\end{equation*}
$$

Here, for sets $A, B$, we denote union $A \cup B$ by $A \sqcup B$ if the union is disjoint. By definition, we have $\mathcal{N}_{1}=S \cap \mathcal{M}_{\mathrm{pl}}^{c}(F) \cap \mathcal{M}\left(F_{\mathrm{pl}}\right)$ and $\mathcal{N}_{3}=S \cap \mathcal{M}_{\mathrm{pl}}(F) \cap \mathcal{M}^{c}\left(F_{\mathrm{pl}}\right)$. Thus, we have

$$
\begin{aligned}
& \mathcal{N}_{3} \cup\left(\mathcal{M}_{\mathrm{pl}}^{c}(F) \cap S\right) \\
& =\left(S \cap \mathcal{M}_{\mathrm{pl}}(F) \cap \mathcal{M}^{c}\left(F_{\mathrm{pl}}\right)\right) \cup\left(\mathcal{M}_{\mathrm{pl}}^{c}(F) \cap S\right) \\
& =S \cap\left\{\left(\mathcal{M}_{\mathrm{pl}}(F) \cap \mathcal{M}^{c}\left(F_{\mathrm{pl}}\right)\right) \cup \mathcal{M}_{\mathrm{pl}}^{c}(F)\right\} \\
& =S \cap\left(\mathcal{M}^{c}\left(F_{\mathrm{pl}}\right) \cup \mathcal{M}_{\mathrm{pl}}^{c}(F)\right) .
\end{aligned}
$$

Similarly, we have the following:

$$
\begin{aligned}
& \mathcal{N}_{1} \cup\left(\mathcal{M}^{c}\left(F_{\mathrm{pl}}\right) \cap S\right) \\
& =\left(S \cap \mathcal{M}_{\mathrm{pl}}^{c}(F) \cap \mathcal{M}\left(F_{\mathrm{pl}}\right)\right) \cup\left(\mathcal{M}^{c}\left(F_{\mathrm{pl}}\right) \cap S\right) \\
& =S \cap\left(\mathcal{M}^{c}\left(F_{\mathrm{pl}}\right) \cup \mathcal{M}_{\mathrm{pl}}^{c}(F)\right) .
\end{aligned}
$$

Since disjointness is obvious by definition, we obtain (11). Next, we note that the following holds:

$$
\begin{align*}
Q\left(\mathcal{M}_{\mathrm{pl}}^{c}(F) \cap S\right) & =Q\left(\mathcal{M}_{\mathrm{pl}}^{c}(F)\right)-Q\left(\mathcal{M}_{\mathrm{pl}}^{c}(F) \cap S^{c}\right) \\
& \geq Q\left(\mathcal{M}_{\mathrm{pl}}^{c}(F)\right)-Q\left(S^{c}\right) . \tag{12}
\end{align*}
$$

By (11), we obtain the following:

$$
\begin{aligned}
Q\left(\mathcal{N}_{3}\right) & =Q\left(\mathcal{N}_{1}\right)+Q\left(\mathcal{M}^{c}\left(F_{\mathrm{pl}}\right) \cap S\right)-Q\left(\mathcal{M}_{\mathrm{pl}}^{c}(F) \cap S\right) \\
& \leq Q\left(\mathcal{N}_{1}\right)+Q\left(\mathcal{M}^{c}\left(F_{\mathrm{pl}}\right)\right)-Q\left(\mathcal{M}^{c}\left(F_{\mathrm{pl}}\right) \cap S\right) \\
& \leq Q\left(\mathcal{N}_{1}\right)+Q\left(\mathcal{M}^{c}\left(F_{\mathrm{pl}}\right)\right)-Q\left(\mathcal{M}_{\mathrm{pl}}^{c}(F)\right)+Q\left(S^{c}\right) \\
& \leq q+Q\left(\mathcal{M}^{c}\left(F_{\mathrm{pl}}\right)\right)-Q\left(\mathcal{M}_{\mathrm{pl}}^{c}(F)\right)+Q\left(S^{c}\right) \\
& =q-Q\left(\mathcal{M}\left(F_{\mathrm{pl}}\right)\right)+Q\left(\mathcal{M}_{\mathrm{pl}}(F)\right)+Q\left(S^{c}\right) .
\end{aligned}
$$

Here the second inequality follows from (12) and the third inequality follows from Lemma 12. This completes the proof.

Now, we can prove Lemma 10 as follows.

Proof of Lemma 10.

$$
\begin{aligned}
Q(\mathcal{M}(F)) & =Q\left(\mathcal{M}(F) \cap \mathcal{S}_{\mathcal{B}}(F) \cap \mathcal{S}_{\mathcal{B}}(H)\right)+Q\left(\mathcal{M}(F) \cap\left(\mathcal{S}_{\mathcal{B}}^{c}(F) \cup \mathcal{S}_{\mathcal{B}}^{c}(H)\right)\right) \\
& \leq Q\left(\mathcal{N}_{1} \cup \mathcal{N}_{2}\right)+Q\left(\mathcal{N}_{3}\right)+Q\left(\mathcal{S}_{\mathcal{B}}^{c}(F) \cup \mathcal{S}_{\mathcal{B}}^{c}(H)\right) \\
& \leq 2 q+2 Q\left(\mathcal{S}_{\mathcal{B}}^{c}(F) \cup \mathcal{S}_{\mathcal{B}}^{c}(H)\right)+Q\left(\mathcal{M}_{\mathrm{pl}}(F)\right)-Q\left(\mathcal{M}\left(F_{\mathrm{pl}}\right)\right)
\end{aligned}
$$

Here, the last inequality follows from Lemmas 12 and 13. Since $q$ satisfies $q \leq$ $\frac{Q\left(\mathcal{M}_{\mathrm{pl}}(F)\right)+Q\left(\mathcal{S}_{\mathcal{B}}^{c}(F)+Q\left(\mathcal{S}_{\mathcal{B}}^{c}(H)\right)\right)}{\gamma_{H}-1}$ by (10), we have our assertion.

## C. 4 Miscellaneous Proofs for Examples

Proof of Example 3. In Example 3, we stated that $p \mapsto R_{h}(p) / p$ is non-increasing. This follows from the concavity of $R_{h}$ and $\lim _{p \rightarrow+0} R_{h}(p)=0$. In fact, we can prove the concavity of $R_{h}$ by $\frac{d^{2} R_{h}}{d p^{2}}(p)=-h \exp \left(\frac{\xi^{2}}{2}-h \xi-\frac{1}{2} h^{2}\right) \leq 0$, where $\xi=\Phi^{-1}(p)$.

Proof of Example 9. For each $i, j \in[K], w_{i} P_{i}(x) \geq w_{j} P_{j}(x)$ is equivalent to $\left(x-\tau_{i}\right) \cdot v_{j i} \leq$ $\frac{\left\|\tau_{i}-\tau_{j}\right\|}{2}+\frac{2\left(\log w_{i}-\log w_{j}\right)}{d\left\|\tau_{i}-\tau_{j}\right\|}$, where $v_{j i}=\frac{\tau_{j}-\tau_{i}}{\left\|\tau_{i}-\tau_{j}\right\|}$. Thus, for each $i \in[K]$, we have $\bigcap_{j \in[K] \backslash\{i\}} X_{i j} \subseteq$ $\mathcal{S}_{\mathcal{B}}\left(F_{\text {opt }}\right)$. Here $X_{i j}$ is defined as $\left\{x \in \mathcal{X}:\left(x-\tau_{i}\right) \cdot v_{j i} \leq \frac{\left\|\tau_{i}-\tau_{j}\right\|}{2}+\frac{2\left(\log w_{i}-\log w_{j}\right)}{d\left\|\tau_{i}-\tau_{j}\right\|}-\frac{r}{2}\right\}$. For any $w \in \mathbb{R}^{d}$ with $\|w\|_{2}=\sqrt{d}$ and $a>0$, we have $P_{i}\left(\left\{x \in \mathcal{X}:\left(x-\tau_{i}\right) \cdot w>a\right\}\right)=$ $1-\Phi(a) \leq \frac{1}{2} \exp \left(-a^{2} / 2\right)$ (c.f., [2]). Thus, $P_{i}\left(X_{i j}^{c}\right) \leq \frac{1}{2} \exp \left(-d a_{i j}^{2} / 2\right)$, where $a_{i j}=\frac{\left\|\tau_{i}-\tau_{j}\right\|}{2}+$ $\frac{2\left(\log w_{i}-\log w_{j}\right)}{d\left\|\tau_{i}-\tau_{j}\right\|}-\frac{r}{2}$. By assumptions, we have $a_{i j} \sqrt{d} \gtrsim \sqrt{\log d}$. Therefore, $P_{i}\left(X_{i j}^{c}\right)=O\left(\frac{1}{p o l y(d)}\right)$. It follows that $P_{i}\left(\mathcal{S}_{\mathcal{B}}^{c}\left(F_{\mathrm{opt}}\right)\right) \leq \sum_{j \in[K] \backslash\{i\}} P_{i}\left(X_{i j}^{c}\right)=O\left(\frac{1}{\text { poly(d) }}\right)$. Thus, we have $R_{\mathcal{B}, w}\left(F_{\mathrm{opt}}\right)=$ $O\left(\frac{1}{p o l y(d)}\right)$. By the same way, we can prove that $\operatorname{Err}_{w}\left(F_{\mathrm{opt}}\right)=O\left(\frac{1}{\operatorname{poly}(d)}\right)$.

## D All-Layer Margin Generalization Bounds

Following [14, 13], we introduce all layer margin of neural networks and provide generalization bounds of CSST. In this section, we assume that classifier $F(x)$ is given as $F(x)=$ $\operatorname{argmax}_{1 \leq i \leq K} \Phi_{i}(x)$, where $\Phi$ is a neural network of the form

$$
\Phi(x)=\left(f_{p} \circ f_{p-1} \circ \cdots \circ f_{1}\right)(x)
$$

Here $f_{i}: \mathbb{R}^{d_{i}} \rightarrow \mathbb{R}^{d_{i+1}}$ with $d_{1}=d$ and $d_{p}=K$. We assume that each $f_{i}$ belongs to a function class $\mathcal{F}_{i} \subset \operatorname{Map}\left(\mathbb{R}^{d_{i}}, \mathbb{R}^{d_{i+1}}\right)$. We define a function class $\mathcal{F}$ to which $\Phi$ belongs by

$$
\mathcal{F}=\left\{\Phi: \mathbb{R}^{d} \rightarrow \mathbb{R}^{K}: \Phi(x)=\left(f_{p} \circ \cdots \circ f_{1}\right)(x), \quad f_{i} \in \mathcal{F}_{i}, \forall i\right\}
$$

For example, for $b>0, \mathcal{F}_{i}$ is given as $\left\{h \mapsto W \phi(h): W \in \mathbb{R}^{d_{i} \times d_{i+1}},\|W\|_{\text {fro }} \leq b\right\}$ if $i>1$ and $\left\{h \mapsto W h: W \in \mathbb{R}^{d_{1} \times d_{2}},\|W\|_{\text {fro }} \leq b\right\}$ if $i=1$, where $\phi$ is a link function (applied on $\mathbb{R}^{d_{i}}$ entry-wise) with bounded operator norm (i.e., $\|\phi\|_{\text {op }}:=\sup _{x \in \mathbb{R}^{d_{i}} \backslash\{0\}}\|\phi(x)\|_{2} /\|x\|_{2}<\infty$ ) and $\|W\|_{\text {fro }}$ denotes the Frobenius norm of the matrix. However, we do not assume the function class $\mathcal{F}_{i}$ does not have this specific form. We assume that each function class $\mathcal{F}_{i}$ is a normed vector space with norm $\|\cdot\|$. In the example above, we consider the operator norm, i.e., if $f(h)=\phi(W h)$, $\|f\|$ is defined as $\|f\|_{\text {op }}$. Let $x_{1}, \ldots, x_{n}$ be a finite i.i.d. sequence of samples drawn from $\mathcal{P}_{w}$. We denote the corresponding empirical distribution by $\widehat{P}_{w}$, i.e., for a measurable function $f$ on $\mathcal{X}$, $\mathbf{E}_{x \sim \widehat{P}_{w}}[f]=\sum_{i=1}^{n} f\left(x_{i}\right)$.
For $\xi=\left(\xi_{1}, \ldots, \xi_{p}\right) \in \prod_{i=1}^{p} \mathbb{R}^{d_{i}}$, we define the perturbed output $\Phi(x, \xi)$ as $\Phi(x, \xi)=h_{p}(x, \xi)$, where

$$
\begin{aligned}
h_{1}(x, \xi) & =f_{1}(x)+\xi_{1}\|x\|_{2} \\
h_{i}(x, \xi) & =f_{i}\left(h_{i-1}(x, \xi)\right)+\xi_{i}\left\|h_{i-1}(x, \xi)\right\|_{2}, \quad \text { for } 2 \leq i \leq p
\end{aligned}
$$

Let $x \in \mathcal{X}$ and $y \in[K]$. We define $\Xi(\Phi, x, y)$ by $\left\{\xi \in \prod_{i=1}^{p} \mathbb{R}^{d_{i}}: \operatorname{argmax}_{i} \Phi_{i}(x, \xi) \neq y\right\}$. Then, the all-layer margin $m(\Phi, x, y)$ is defined as

$$
m(\Phi, x, y)=\min _{\xi \in \Xi(\Phi, x, y)}\|\xi\|_{2}
$$

where $\|\xi\|_{2}$ is given by $\sqrt{\sum_{i=1}^{p}\left\|\xi_{i}\right\|_{2}^{2}}$. Following [14], we define a variant of the all-layer margin that measures robustness of $\Phi$ with respect to input transformations defined by $\mathcal{B}(x)$ as follows:

$$
m_{\mathcal{B}}(\Phi, x):=\min _{x^{\prime} \in \mathcal{B}(x)} m\left(F, x^{\prime}, \operatorname{argmax}_{i} \Phi_{i}(x)\right)
$$

AsSumption 14 (c.f. [13], Condition A.1). Let $\mathcal{G}$ be a normed space with norm $\|\cdot\|$ and $\epsilon>0$. We say $\mathcal{G}$ satisfies the $\epsilon^{-2}$ covering condition with complexity $\mathcal{C}_{\|\cdot\|}(\mathcal{G})$ if for all $\epsilon>0$, we have

$$
\log \mathcal{N}_{\|\cdot\|}(\epsilon, \mathcal{G}) \leq \frac{\mathcal{C}_{\|\cdot\|}(\mathcal{G})}{\epsilon^{2}}
$$

Here $\mathcal{N}_{\|\cdot\|}(\epsilon, \mathcal{G})$ the $\epsilon$-covering number of $\mathcal{G}$. We assume function class $\mathcal{F}_{i}$ satisfies the $\epsilon^{-2}$ covering condition with complexity $\mathcal{C}_{\|\cdot\|}\left(\mathcal{F}_{i}\right)$ for each $1 \leq i \leq p$.

Throughout this section, we suppose that Assumption 14 holds. Essentially, the following two propositions follows were proved by Wei et al. [14]:
Proposition 15 (c.f., [14], Lemma D.6). With probability at least $1-\delta$ over the draw of the training data, for all $t \in(0, \infty)$, any $\Phi \in \mathcal{F}$ satisfies the following:

$$
R_{\mathcal{B}, w}(F)=\mathbf{E}_{\widehat{P}_{w}}\left[\mathbf{1}\left(m_{\mathcal{B}}(\Phi, x) \leq t\right)\right]+\widetilde{O}\left(\frac{\sum_{i=1}^{p} \mathcal{C}_{\|\cdot\|_{\text {op }}}\left(\mathcal{F}_{i}\right)}{t \sqrt{n}}\right)+\zeta
$$

where $\zeta=O\left(\sqrt{\frac{\log (1 / \delta)+\log n}{n}}\right)$ is a lower order term and $F(x)=\operatorname{argmax}_{i \in[K]} \Phi_{i}(x)$.
Proposition 16 (c.f., [14], Theorem D.3). With probability at least $1-\delta$ over the draw of the training data, for all $t \in(0, \infty)$, any $\Phi \in \mathcal{F}$ satisfies the following:

$$
L_{w}\left(F, F_{\mathrm{pl}}\right)=\mathbf{E}_{\widehat{P}_{w}}\left[\mathbf{1}\left(m\left(\Phi, x, F_{\mathrm{pl}}(x)\right) \leq t\right)\right]+\widetilde{O}\left(\frac{\sum_{i=1}^{p} \mathcal{C}_{\|\cdot\|_{\mathrm{op}}}\left(\mathcal{F}_{i}\right)}{t \sqrt{n}}\right)+\zeta
$$

where $\zeta=O\left(\sqrt{\frac{\log (1 / \delta)+\log n}{n}}\right)$ is a lower order term and $F(x)=\operatorname{argmax}_{i \in[K]} \Phi_{i}(x)$.
REMARK. Although we have proved Theorem 5 following [14], we had to provide our own proof due to some differences in theoretical assumptions (e.g., non-existence of the ground-truth classifier, a difference mentioned in the remark just after Assumption 4). On the other hand, the proofs of [14, Lemma D.6] and [14, Theorem D.3] work for any distribution $P$ on $\mathcal{X}$ and its empirical distribution $\widehat{P}$. Since $\|w\|_{1} \mathcal{P}_{w}\left(\mathcal{S}_{\mathcal{B}}^{c}(F)\right)=R_{\mathcal{B}, w}(F)$ and $\|w\|_{1} \mathcal{P}_{w}\left(\left\{x: F(x) \neq F_{\mathrm{pl}}(x)\right\}\right)=L_{w}\left(F, F_{\mathrm{pl}}\right)$, Proposition 15 and Proposition 16 follow from the corresponding results in [14].

Theorem 17. Suppose Assumption 4 and Assumption 14 hold. Then, with probability at least $1-\delta$ over the draw of the training data, for all $t_{1}, t_{2} \in(0, \infty)$, and any neural network $\Phi$ in $\mathcal{F}$, we have the following:

$$
\begin{aligned}
& \operatorname{Err}_{w}(F)=\frac{\gamma+1}{\gamma-1} \mathbf{E}_{\widehat{P}_{w}}\left[\mathbf{1}\left(m\left(\Phi, x, F_{\mathrm{pl}}(x)\right) \leq t_{1}\right)\right]+\frac{2 \gamma}{\gamma-1} \mathbf{E}_{\widehat{P}_{w}}\left[\mathbf{1}\left(m_{\mathcal{B}}(\Phi, x) \leq t_{2}\right)\right] \\
&-\operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)+2 \operatorname{Err}_{w}\left(F^{*}\right)+\frac{2 \gamma}{\gamma-1} R_{\mathcal{B}, w}\left(F^{*}\right) \\
&+\widetilde{O}\left(\frac{\sum_{i=1}^{p} \mathcal{C}_{\|\cdot\|_{\mathrm{op}}}\left(\mathcal{F}_{i}\right)}{t_{1} \sqrt{n}}\right)+\widetilde{O}\left(\frac{\sum_{i=1}^{p} \mathcal{C}_{\|\cdot\|_{\mathrm{op}}}\left(\mathcal{F}_{i}\right)}{t_{2} \sqrt{n}}\right)+\zeta
\end{aligned}
$$

where $\zeta=O\left(\sqrt{\frac{\log (1 / \delta)+\log n}{n}}\right)$ is a lower order term and $F(x)=\operatorname{argmax}_{i \in[K]} \Phi_{i}(x)$.

Proof. By (9) with $H=F^{*}$ and $-\mathbf{1}\left(F_{\mathrm{pl}}(x) \neq F^{*}(x)\right) \leq-\mathbf{1}\left(F_{\mathrm{pl}}(x) \neq j\right)+\mathbf{1}\left(F^{*}(x) \neq j\right)$ for any $x \in \mathcal{X}$ and $j \in[K]$, we obtain the following inequality:
$\operatorname{Err}_{w}(F) \leq \frac{\gamma+1}{\gamma-1} L_{w}\left(F, F_{\mathrm{pl}}\right)+\frac{2 \gamma}{\gamma-1} R_{\mathcal{B}, w}(F)+\frac{2 \gamma}{\gamma-1} R_{\mathcal{B}, w}\left(F^{*}\right)-\operatorname{Err}_{w}\left(F_{\mathrm{pl}}\right)+2 \operatorname{Err}_{w}\left(F^{*}\right)$.
Then, the statement of the theorem follows from Proposition 15 and Proposition 16.

## E Proof of Proposition 6

Proof. Let the average weighted consistency loss be $\mathcal{L}_{u}^{w t}=\frac{1}{|B|} \sum_{x \in B} \ell_{u}^{w t}\left(\hat{p}_{m}(x), p_{m}(\mathcal{A}(x)), \mathbf{G}\right)$ this will be minimized if for each of $x \in B$ the $\ell_{u}^{w t}\left(\hat{p}_{m}(x), p_{m}(\mathcal{A}(x)), \mathbf{G}\right)$ is minimized. This
expression can be expanded as:

$$
\begin{aligned}
\ell_{u}^{w t}\left(\hat{p}_{m}(x), p_{m}(\mathcal{A}(x)), \mathbf{G}\right) & =-\sum_{i=1}^{K}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right)_{i} \log \left(p_{m}(\mathcal{A}(x))_{i}\right) \\
& =-C \sum_{i=1}^{K} \frac{\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right)_{i}}{\left.\sum_{j=1}^{m} \mathbf{G}^{T} \hat{p}_{m}(x)\right)_{j}} \log \left(p_{m}(\mathcal{A}(x))_{i}\right) \\
& =C \times \mathrm{H}\left(\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right) \| p_{m}(\mathcal{A}(x))\right) .
\end{aligned}
$$

Here we use H to denote the cross entropy between two distributions. As we don't backpropogate gradients from the $\hat{p}_{m}(x)$ (pseudo-label) branch of prediction network we can consider $\left.C=\sum_{j=1}^{m} \mathbf{G}^{T} \hat{p}_{m}(x)\right)_{j}$ as a constant in our analysis. Also adding a constant term of entropy $\mathrm{H}\left(\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right)\right)$ to cross entropy term and dropping constant $C$ doesn't change the outcome of minimization. Hence we have the following:

$$
\begin{aligned}
& \min _{p_{m}} \mathrm{H}\left(\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right) \| p_{m}(\mathcal{A}(x))\right)= \min _{p_{m}} \mathrm{H}\left(\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right) \| p_{m}(\mathcal{A}(x))\right) \\
&+\mathrm{H}\left(\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right)\right) \\
&=\min _{p_{m}} \mathcal{D}_{K L}\left(\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right) \| p_{m}(\mathcal{A}(x))\right) .
\end{aligned}
$$

This final term is the $\mathcal{D}_{K L}\left(\mathbf{G}^{T} \hat{p}_{m}(x) \| p_{m}(\mathcal{A}(x))\right.$ which is obtained by using the identity $\mathcal{D}_{K L}(p, q)=\mathrm{H}(p, q)+\mathrm{H}(p)$ where $p, q$ are the two distributions.

## F Notation

We provide the list of notations commonly used in the paper in Table 1.

## G Code, License, Assets and Computation Requirements

## G. 1 Code and Licenses of Assets

In this work, we use the open source implementation of FixMatch [12] ${ }^{1}$ in PyTorch, which is licensed under MIT License for educational purpose. Also for NLP experiments we make use of DistillBERT [11] pretrained model available in the HuggingFace [15] library. We promise to release the code and checkpoints at the time of acceptance of our submission.

## G. 2 Computational Requirements

All experiments were done on a variety of GPUs, with primarily Nvidia A5000 (24GB) with occasional use of Nvidia A100 (80GB) and Nvidia RTX3090 (24GB). For finetuning DistilBERT and all experiments with ImageNet-100 dataset we used PyTorch data parallel over 4 A5000s. Training was done till no significant change in metrics was observed. The detailed list of computation used per experiment type and dataset have been tabulated in Table 2 and Table 3.

## H Objective

## H. 1 Logit Adjusted Weighted Consistency Regularizer

As we have introduced weighted consistency regularizer in Eq. 6 for utilizing unlabeled data, we now provide logit adjusted variant of it for training deep networks in this section. We provide logit

[^0]Table 1: Table of Notations used in Paper

| $\mathcal{Y}$ | Label space |
| :---: | :---: |
| $\mathcal{X}$ | Instance space |
| K | Number of classes |
| $\pi_{i}$ | prior for class i |
| $F$ | a classifier model |
| $s$ | a scoring function, $\mathcal{X} \rightarrow \mathbb{R}^{K}$ |
| D | data distribution |
| $\boldsymbol{\lambda}$ | Lagrange multiplier |
| $\lambda_{u}$ | coefficient of unlabeled loss |
| $\operatorname{rec}_{i}[F]$ | recall of $i^{\text {th }}$ class for a classifier $F$ |
| $\operatorname{acc}[F]$ | accuracy for a classifier $F$ |
| $\operatorname{prec}_{i}[F]$ | precision of $i^{\text {th }}$ class for a classifier $F$ |
| $\operatorname{cov}_{i}[F]$ | coverage for $i^{\text {th }}$ class for a classifier $F$ |
| G | a $K \times K$ matrix |
| D | a $K \times K$ diagonal matrix |
| M | a $K \times K$ matrix |
| $\mu$ | ratio of labelled to unlabelled samples |
| $B$ | batch size for FixMatch |
| $\ell_{u}^{\text {wt }}$ | loss for unlabelled data using pseudo label |
| $\ell_{s}^{\text {hyb }}$ | loss for labelled data |
| $\mathcal{L}_{u}^{\text {hyb }}$ | average loss for unlabelled data using pseudo label on a batch of samples |
| $\mathcal{L}_{s}^{\text {hyb }}$ | average loss for labelled data on a batch of samples |
| H | cross entropy function |
| $\mathcal{A}$ | a $\mathcal{X} \rightarrow \mathcal{X}$ function that is stochastic in nature and applies a strong augmentation to it |
| $\alpha$ | a $\mathcal{X} \rightarrow \mathcal{X}$ function that is stochastic in nature and applies a weak augmentation to it |
| $\rho$ | imbalance factor |
| $B$ | batch size of samples |
| $B_{s}$ | batch of labelled samples |
| $B_{u}$ | batch of unlabelled samples |
| $x$ | an input sample, $x \in \mathcal{X}$ |
| $\hat{p}_{m}$ | a pseudo label generating function |
| $p_{m}$ | distribution of confidence for a model's prediction on a given sample |
| w | a $K \times K$ weight matrix that corresponds to a gain matrix $\mathbf{G}$ |
| $\operatorname{Err}_{w}(F)$ | weighted error of $F$ that corresponds to the objective of CSL |
| $\mathcal{P}_{w}$ | weighted distribution on $\mathcal{X}$ |
| $P_{i}$ | class conditional distribution of samples for class $i$ |
| $R_{\mathcal{B}, w}(F)$ | theoretical weighted (cost sensitive) consistency regularizer |
| $F_{\mathrm{pl}}$ | a pseudo labeler |
| $L_{w}\left(F, F_{\mathrm{pl}}\right)$ | weighted error between $F$ and $F_{\text {pl }}$ |
| $\mathcal{L}_{w}(F)$ | theoretical CSST loss |
| c | a non-increasing function used in the definition of the $c$-expansion property (Definition 2) |
| $\gamma$ | a value of $c$ defined in Assumption 4 |
| $\beta$ | an upper bound of $R_{\mathcal{B}, w}(F)$ in the optimization problem (4) |
| $S^{c}$ | the complement of a set $S$ |


| Method | CIFAR-10 | CIFAR-100 | ImageNet-100 |
| :---: | :---: | :---: | :---: |
| ERM | A5000 | A5000 | RTX3090 |
|  | 49 m | $6 \mathrm{~h} \mathrm{47m}$ | 15 h 8 m |
| LA | RTX3090 | A5000 | A5000 |
|  | 39 m | 6 h 9 m | 15 h 7 m |
| CSL | A5000 | A5000 | A5000 |
| CSST(FixMatch) | 4 X A5000 | 6h 4 A10 | 12h |
| w/o KL-Threshold | 21h 0m | 2d 19h 16 m | 2d A5000 |
| CSST(FixMatch) | 4 X A5000 | 4 X A5000 | 4 X A5000 |
|  | 21h 41m | 2d 11h 52m | 2d 4m |

Table 2: Computational requirements and training time (d:days, h:hours, m:minutes) for experiments relevant to vision datasets. As we can see some of the experiments on the larger datasets such as ImageNet requires long compute times.

| Method | $\operatorname{IMDb}(\rho=10)$ | $\operatorname{IMDb}(\rho=100)$ | DBpedia-14 |
| :---: | :---: | :---: | :---: |
| ERM | 4 X A5000 | 4 X A5000 | 4 X A5000 |
|  | 25 m | 29m | 2h 44m |
| UDA | 4 X A5000 | 4 X A5000 | 4 X A5000 |
|  | 44 m | 32 m | 10h 18m |
| CSST(UDA) | 4 X A5000 | 4 X A5000 | 4 X A5000 |
|  | 49m | 35 m | 13h 12m |

Table 3: Computational requirements and training time(d:days, h:hours, m:minutes) for experiments done on NLP datasets. The DistilBERT model which we are using is pretrained on a language modeling task, hence it requires much less time for training in comparison to vision models which are trained from scratch.
adjusted term for $\ell_{u}^{\mathrm{wt}}\left(\hat{p}_{m}(x), p_{m}(\mathcal{A}(x), \mathbf{G})\right.$ below:

$$
\begin{aligned}
\ell_{u}^{\mathrm{wt}}\left(\hat{p}_{m}(x), p_{m}(\mathcal{A}(x), \mathbf{G})\right. & =-\sum_{i=1}^{K}\left(\mathbf{G}^{\mathbf{T}} \hat{p}_{m}(x)\right)_{i} \log \left(p_{m}(\mathcal{A}(x))_{i}\right) \\
& =-\sum_{i=1}^{K}\left(\mathbf{G}^{\mathbf{T}} \hat{p}_{m}(x)\right)_{i} \log \left(\frac{\exp \left(\mathbf{s}(\mathcal{A}(x))_{i}\right)}{\sum_{j=1}^{K} \exp \left(\mathbf{s}(\mathcal{A}(x))_{j}\right)}\right) \\
& =-\sum_{i=1}^{K}\left(\mathrm{D}^{\mathbf{T}} \mathrm{M}^{\mathbf{T}} \hat{p}_{m}(x)\right)_{i} \log \left(\frac{\exp \left(\mathbf{s}(\mathcal{A}(x))_{i}\right)}{\sum_{j=1}^{K} \exp \left(\mathbf{s}(\mathcal{A}(x))_{j}\right)}\right)
\end{aligned}
$$

The above expression comes from the decomposition $\mathbf{G}=$ MD. The above loss function can be converted into it's logit adjusted equivalent variant by following transformation as suggested by Narasimhan and Menon [8] which is equivalent in terms of optimisation of deep neural networks:

$$
\begin{equation*}
\ell_{u}^{\mathrm{wt}}\left(\hat{p}_{m}(x), p_{m}(\mathcal{A}(x), \mathbf{G}) \equiv-\sum_{i=1}^{K}\left(\mathrm{M}^{\mathbf{T}} \hat{p}_{m}(x)\right)_{i} \log \left(\frac{\exp \left(\mathbf{s}(\mathcal{A}(x))_{i}-\log \left(\boldsymbol{D}_{i i}\right)\right)}{\sum_{j=1}^{K} \exp \left(\mathbf{s}(\mathcal{A}(x))_{j}-\log \left(\boldsymbol{D}_{j j}\right)\right)}\right)\right. \tag{13}
\end{equation*}
$$

The above loss is the consistency loss $\ell_{u}^{\mathrm{wt}}$ that we practically implement for CSST. Further in case $\hat{p}_{m}(x)$ is a hard pseudo label $y$ as in FixMatch, the above weighted consistency loss reduces to $\ell^{\text {hyb }}(y, \mathbf{s}(\mathcal{A}(x)))$. Further in case the gain matrix $G$ is diagonal the above loss will converge to $\ell^{\mathrm{LA}}(y, \mathbf{s}(\mathcal{A}(x)))$. Thus the weighted consistency regularizer can be converted to logit adjusted variants $\ell^{\mathrm{LA}}$ and $\ell^{\text {hyb }}$ based on $\mathbf{G}$ matrix.

## H. 2 CSST(FixMatch)

In FixMatch, we use the prediction made by the model on a sample $x$ after applying a weak augmentation $\alpha$ and is used to get a hard pseudo label for the models prediction on a strongly
augmented sample i.e. $\mathcal{A}(x)$. The set of weak augmentations include horizontal flip, We shall refer to this pseudo label as $\hat{p}_{m}(x)$. The list of strong augmentations are given in Table 12 of Sohn et al. [12]. Weak augmentations include padding, random horizontal flip and cropping to the desired dimensions (32X32 for CIFAR and 224X224 for ImageNet). Given a batch of labeled and unlabeled samples $B_{s}$ and $B_{u}$, CSST modifies the supervised and un-supervised component of the loss function depending upon the non-decomposable objective and its corresponding gain matrix $\mathbf{G}$ at a given time during training. We assume that in the dataset, a sample $x$, be it labeled or unlabeled is already weakly augmented. vanilla FixMatch's supervised componenet of the loss function is a simple cross entropy loss whereas in our CSST(FixMatch) it is replaced by $\ell_{s}^{\text {hyb }}$.

$$
\begin{equation*}
\mathcal{L}_{s}^{\mathrm{hyb}}=\frac{1}{\left|B_{s}\right|} \sum_{x, y \in B_{s}} \ell^{\mathrm{hyb}}(y, s(x)) \tag{14}
\end{equation*}
$$

$$
\begin{equation*}
\left.\mathcal{L}_{u}^{\mathrm{wt}}=\frac{1}{\left|B_{u}\right|} \sum_{x \in B_{u}} \mathbb{1}_{\left(\mathcal{D}_{K L}\left(\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right) \| p_{m}(x)\right) \leq \tau\right)} \ell_{u}^{\mathrm{wt}}\left(\hat{p}_{m}(x), p_{m}(\mathcal{A}(x)), \mathbf{G}\right)\right) \tag{15}
\end{equation*}
$$

In the above expression $p_{m}(x)=\operatorname{softmax} \mathbf{s}(x)$. The component of the loss function for unlabeled data (i.e. consistency regularization) is where one of our contributions w.r.t the novel thresholding mechanism comes into light. vanilla FixMatch selects unlabeled samples for which consistency loss is non-zero, such that the model's confidence on the most likely predicted class is above a certain threshold. We rather go for a threshold mechanism that select based on the basis of degree of distribution match to a target distribution based on $\mathbf{G}$. The final loss function $\mathcal{L}=\mathcal{L}_{s}^{h y b}+\lambda_{u} \mathcal{L}_{u}^{\mathrm{wt}}$, i.e. a linear combination of $\mathcal{L}_{s}^{\mathrm{hyb}}$ and $\mathcal{L}_{u}^{w t}$. Since for FixMatch we are dealing with Wide-ResNets and ResNets which are deep networks, as mentioned in Section H.1, we shall use the alternate logit adjusted formulation as mentioned in Eq. 13 as substitute for $\ell_{u}^{\mathrm{wt}}$ in Eq. 15.

## H. $3 \operatorname{CSST}($ UDA)

The loss function of UDA is a linear combination of supervised loss and consistency loss on unlabeled samples. The former is the cross entropy (CE) loss, while the latter for the unlabeled samples minimizes the KL-divergence between the model's predicted label distribution on an input sample and its augmented sample. Often the predicted label distribution on the unaugmented sample is sharpened. The augmentation we used was a English-French-English backtranslation based on the MarianMT [4] fast neural machine translation model. In UDA supervised component of the loss is annealed using a method described as Training Signal Annealing (TSA), where the CE loss is considered only for those labeled samples whose $\max _{i} p_{m}(x)_{i}<\tau_{t}$, where $t$ is a training time step. We observed that using TSA in a long tailed setting leads to overfitting on the head classes and hence chose to not include the same in our final implementation.

CSST modifies the supervised and unsupervised component of the loss function in UDA depending upon a given objective and its corresponding gain matrix $\mathbf{G}$ at a given time during training. The supervised component of the loss function for a given constrained optimisation problem and a gain matrix $\mathbf{G}$, is the hybrid loss $\ell_{s}^{\text {hyb }}$. For the consistency regularizer part of the loss function, we minimize the KL-divergence between a target distribution and the model's prediction label distribution on its augmented version. The target distribution is norm $\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right)$, where $\hat{p}_{m}(x)$ is the sharpened prediction of the label distribution by the model. Given a batch of labeled and unlabeled samples $B_{s}$ and $B_{u}$, the final loss function in $\operatorname{CSST}(\mathrm{UDA})$ is a linear combination of $\mathcal{L}_{s}^{h y b}$ and $\mathcal{L}_{u}^{w t}$, i.e $\mathcal{L}=\mathcal{L}_{s}^{h y b}+\lambda_{u} \mathcal{L}_{u}^{w t}$.

$$
\begin{gather*}
\mathcal{L}_{s}^{\mathrm{hyb}}=\frac{1}{\left|B_{s}\right|} \sum_{x, y \in B_{s}} \ell^{\mathrm{hyb}}\left(p_{m}(x), y\right)  \tag{16}\\
\left.\left.\mathcal{L}_{u}^{w t}=\frac{1}{\left|B_{u}\right|} \sum_{x \in B_{u}} \mathbb{1}_{\left(\mathcal { D } _ { K L } \left(\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right)\right.\right.} \| p_{m}(x)\right) \leq \tau\right) \ell_{u}^{\mathrm{wt}}\left(\hat{p}_{m}(x), p_{m}(\mathcal{A}(x), \mathbf{G})\right) \tag{17}
\end{gather*}
$$

Since for UDA, we are dealing with DistilBERT, as mentioned in Section H.1, we shall use the alternate formulation as mentioned in Eq. 13 as substitute for $\ell_{u}^{\mathrm{wt}}$ in Eq. above.

## I Threshold mechanism for diagonal Gain Matrix

Consider the case when the gain matrix is a diagonal matrix. The loss function $\mathcal{L}_{u}^{w t}\left(B_{u}\right)$ as defined in (7) makes uses of a threshold function that selects samples based on the KL divergence based threshold between the target distribution as defined by the gain matrix $G$ and the models predicted distribution of confidence over the classes.

$$
\begin{equation*}
\text { Threshold function }:=\mathbb{1}_{\left(\mathcal{D}_{K L}\left(\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right) \| p_{m}(x)\right) \leq \tau\right)} \tag{18}
\end{equation*}
$$

Since $\mathbf{G}$ is a diagonal matrix and the pseudo-label $\hat{p}_{m}(x)$ is one hot, the $\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right)$ is a one-hot vector. The threshold function's KL divergence based criterion can be expanded as follows where $\hat{y}$ is the pseudo-label's maximum class's index:

$$
\begin{equation*}
\mathcal{D}_{K L}\left(\operatorname{norm}\left(\mathbf{G}^{T} \hat{p}_{m}(x)\right) \| p_{m}(x)\right)=-\log p_{m}(x)_{\hat{y}}<\tau \tag{19}
\end{equation*}
$$

The above equations represents a threshold on the negative log-confidence of the model's prediction for a given unlabeled sample, for the pseudo-label class $(\hat{y})$. This can be further simplified to $p_{m}(x)_{\hat{y}} \geq \exp (-\tau)$ which is simply a threshold based on the model's confidence. Since pseudolabel is generated from the model's prediction, this threshold is nothing but a selection criterion to select only those samples whose maximum confidence for a predicted hard pseudo-label is above a fixed threshold. This is identical to the threshold function which is used in Fixmatch [12] i.e. $\max \left(p_{m}(x)\right) \geq \exp (-\tau)$. In FixMatch this $\exp (-\tau)$ is set to 0.95 .

## J Dataset

CIFAR-10 and CIFAR-100 [5]. are image classification datasets of images of size 32 X 32 . Both the datasets have a size of 50 k samples and by default, they have a uniform sample distribution among its classes. CIFAR-10 has 10 classes while CIFAR-100 has 100 classes. The test set is a balanced set of 10k images.
ImageNet-100 [10]. is an image classification dataset carved out of ImageNet-1k by selecting the first 100 classes. The distribution of samples is uniform with 1.3 k samples per class. The test set contains 50 images per class. All have a resolution of 224 X 224 , the same as the original ImageNet-1k dataset.
IMDb[7]. dataset is a binary text sentiment classification dataset. The data distribution is uniform by default and has a total 25 k samples in both trainset and testset. In this work, we converted the dataset into a longtailed version of $\rho=10,100$ and selected 1 k labeled samples while truncating the labels of the rest and using them as unlabeled samples.
DBpedia-14[6]. is a topic classification dataset with a uniform distribution of labeled samples. The dataset has 14 classes and has a total of 560 k samples in the trainset and 70 k samples in the test set. Each sample, apart from the content, also has title of the article that could be used for the task of topic classification. In our experiments, we only make use of the content.

## K Algorithms

We provide a detailed description of algorithms used for optimizing non decomposable objectives through CSST(FixMatch) ans CSST(UDA). Algorithm 1 is used for experiments in Section 5 for maximizing worst-case recall (i.e. min recall using CSST(FixMatch) and CSST(UDA)). Algorithm 2 is used for experiments in Section 5 for maximizing recall under coverage constraints (i.e. min coverage experiments on CIFAR10-LT, CIFAR100-LT and ImageNet100-LT).

```
Algorithm 1 CSST-based Algorithm for Maximizing Worst-case Recall
    Inputs: Training set \(S_{s}\) (labeled) and \(S_{u}\) (unlabeled), Validation set \(S^{\text {val }}\), Step-size \(\omega \in \mathbb{R}_{+}\), Class
    priors \(\pi\)
    Initialize: Classifier \(h^{0}\), Multipliers \(\boldsymbol{\lambda}^{0} \in \Delta_{K-1}\)
    for \(t=0\) to \(T-1\) do
        Update \(\lambda\) :
            \(\lambda_{i}^{t+1}=\lambda_{i}^{t} \exp \left(-\omega \cdot \operatorname{recall}_{i}\left[F^{t}\right]\right), \forall i\),
            \(\boldsymbol{\lambda}=\operatorname{norm}(\boldsymbol{\lambda})\)
            \(\mathbf{G}=\operatorname{diag}\left(\lambda_{1}^{t+1} / \pi_{1}, \ldots, \lambda_{K}^{t+1} / \pi_{K}\right)\)
        Compute \(\ell_{u}^{\mathrm{wt}}, \ell_{s}^{\text {hyb }}\) using \(\mathbf{G}\)
        Cost-sensitive Learning (CSL) for FixMatch:
            \(B_{u} \sim S_{u}, B_{s} \sim S_{s} / /\) Sample batches of data
            \(F^{t+1} \in \arg \min _{F} \sum_{B_{u}, B_{s}} \lambda_{u} \mathcal{L}_{u}^{\mathrm{wt}}+\mathcal{L}_{s}^{\text {hyb }} / /\) Replaced by few steps of SGD
    end for
    return \(F^{T}\)
```

```
Algorithm 2 CSST-based Algorithm for Maximizing Mean Recall s.t. per class coverage \(>0.95 / \mathrm{K}\)
    Inputs: Training set \(S_{s}\) (labeled) and \(S_{u}\) (unlabeled), Validation set \(S^{\text {val }}\), Step-size \(\omega \in \mathbb{R}_{+}\), Class
    priors \(\pi\)
    Initialize: Classifier \(h^{0}\), Multipliers \(\boldsymbol{\lambda}^{0} \in \mathbb{R}_{+}^{K}\)
    for \(t=0\) to \(T-1\) do
        Update \(\boldsymbol{\lambda}\) :
            \(\lambda_{i}^{t+1}=\lambda_{i}^{t}-\omega\left(\operatorname{cov}_{i}\left[F^{t}\right]-\frac{0.95}{K}\right), \forall i\)
            \(\lambda_{i}^{t+1}=\max \left\{0, \lambda_{i}^{t+1}\right\}, \forall i \in[K] \quad / /\) Projection to \(\mathbb{R}_{+}\)
            \(\mathbf{G}=\operatorname{diag}\left(\lambda_{1}^{t+1} / \pi_{1}, \ldots, \lambda_{K}^{t+1} / \pi_{K}\right)+\mathbf{1}_{\mathbf{K}} \boldsymbol{\lambda}^{\top}\)
        Compute \(\ell_{u}^{\mathrm{wt}}, \ell_{s}^{\text {hyb }}\) using \(\mathbf{G}\)
        Cost-sensitive Learning (CSL) for FixMatch:
            \(B_{u} \sim S_{u}, B_{s} \sim S_{s} / /\) Sample batches of data
            \(F^{t+1} \in \arg \min _{F} \sum_{B_{u}, B_{s}} \lambda_{u} \mathcal{L}_{u}^{\mathrm{wt}}+\mathcal{L}_{s}^{\mathrm{hyb}} / /\) Replaced by few steps of SGD
    end for
    return \(F^{T}\)
```


## L Details of Experiments and Hyper-parameters

The experiment of $\max _{F} \min _{i} \operatorname{recall}_{i}[f]$ and $\max _{F} \operatorname{recall}[F]$ s.t. $\operatorname{cov}_{i}[F]>\frac{0.95}{K}, \forall i \in[K]$ was performed on the long tailed version of CIFAR-10, $\operatorname{IMDb}(\rho=10,100)$ and DBpedia-14 datasets. This was because the optimisation of the aforementioned 2 objectives is stable for cases with low number of classes. Hence the objective of $\max _{F} \min \left(\operatorname{recall}_{\mathcal{H}}[F]\right.$, recall $\left.\mathcal{T}[F]\right)$ and $\max _{F} \operatorname{recall}[F]$ s.t. $\min _{\mathcal{H}, \mathcal{T}} \operatorname{cov}_{\mathcal{H}, \mathcal{T}}[F]>\frac{0.95}{K}$ is a relatively easier objective for datasets with large number of classes, hence were the optimisation objectives for CIFAR-100 and ImageNet-100 long tailed datasets. For all experiments for a given dataset, we used the same values for a given common hyperparameter. We ablated the threshold for our novel unlabeled sample selection criterion $(\tau)$ and the ratio of labeled and unlabeled samples, given fixed number of unlabeled samples $(\mu)$ and are available in Fig. 4b.

## M Statistical Analysis

We establish the statistical soundness and validity of our results we ran our experiments on 3 different seeds. Due to the computational requirements for some of the experiments ( $\approx 2$ days) we chose to run the experiments on multiple seeds for a subset of tasks i.e. for maximising the minimum recall among all classes for CIFAR-10 LT. We observe that the std. deviation is significantly smaller than the average values for mean recall and min. recall and our performance metrics fall within our std. deviation hence validating the stability and soundness of training.

| Parameter | CIFAR-10 | CIFAR-100 | ImageNet-100 | IMDb <br> $(\rho=10)$ | IMDb <br> $(\rho=100)$ | DBpedia-14 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\tau$ | 0.05 | 0.05 | 0.05 | 0.1 | 0.1 | 0.1 |
| $\lambda_{u}$ | 1.0 | 1.0 | 1.0 | 0.1 | 0.1 | 0.1 |
| $\mu$ | 4.0 | 4.0 | 4.0 | 13.8 | 12.6 | 133 |
| $\left\|B_{s}\right\|$ | 64 | 64 | 64 | 32 | 32 | 32 |
| $\left\|B_{u}\right\|$ | 256 | 256 | 256 | 128 | 128 | 128 |
| lr | $3 \mathrm{e}-3$ | $3 \mathrm{e}-3$ | 0.1 | $1 \mathrm{e}-5$ | $1 \mathrm{e}-5$ | $1 \mathrm{e}-5$ |
| $\omega$ | 0.25 | 0.25 | 0.1 | 0.5 | 0.5 | 0.5 |
| SGD steps | 32 | 100 | 500 | 50 | 50 | 100 |
| before eval | SGD | SGD | SGD | AdamW | AdamW | AdamW |
| optimizer | 0.95 | 0.95 | 0.95 | 0.9 | 0.9 | 0.9 |
| KL-Thresh | $05-4$ | $1 \mathrm{e}-3$ | $1 \mathrm{e}-4$ | $1 \mathrm{e}-2$ | $1 \mathrm{e}-2$ | $1 \mathrm{e}-2$ |
| Weight Decay | $1 \mathrm{e}-4$ | 10 | 10 | 100 | 100 |  |
| $\rho$ | 100 | 10 | 1.0 | 0.1 | 0.1 | 0.1 |
| $\lambda_{u}$ | 1.0 | 1.0 | ResNet50 | DistilBERT | DistilBERT | DistilBERT |
| Arch. | WRN-28-2 | WRN-28-8 | RAD |  |  |  |

Table 4: This table shows us the detailed hyper parameters used for $\operatorname{CSST}$ (FixMatch) for the long tailed datasets CIFAR-10, CIFAR-100, ImageNet-100 and CSST(UDA) on IMDb, DBpedia-14. All the datasets were converted to their respective long tailed versions based on the imbalance factor $\rho$, and a fraction of the samples were used along with their labels for supervision.

Table 5: Avg. and std. deviation of Mean Recall and Min. Recall for CIFAR-10 LT

| Method | Mean Recall | Min Recall |
| :--- | :---: | :---: |
| ERM | $0.52 \pm 0.01$ | $0.27 \pm 0.02$ |
| LA | $0.54 \pm 0.02$ | $0.37 \pm 0.01$ |
| CSL | $0.63 \pm 0.01$ | $0.43 \pm 0.04$ |
| Vanilla (FixMatch) | $0.78 \pm 0.01$ | $0.47 \pm 0.02$ |
| CSST(FixMatch) | $0.75 \pm 0.01$ | $0.72 \pm 0.01$ |

## N Additional Details

## N. 1 Formal Statement Omitted in Sec. 2.2

In Sec. 2.2, we stated that learning with the hybrid loss $\ell^{\text {hyb }}$ gives the Bayes-optimal classifier for the CSL (2). However, due to space constraint, we did not provide a formal statement. In this section, we provide a formal statement of it for clarity.
Proposition ([8] Proposition 4). For any diagonal matrix $\mathbf{D} \in \mathbb{R}^{K \times K}$ with $D_{i i}>0, \forall i$, $\mathbf{M} \in \mathbb{R}^{K \times K}$, and $\mathbf{G}=\mathbf{M D}$, the hybrid loss $\ell^{\text {hyb }}$ is calibrated for $\mathbf{G}$. That is, for any score function $\widehat{\mathbf{s}}: \mathcal{X} \rightarrow \mathbb{R}^{K}$ that minimizes $\mathbf{E}_{(x, y) \sim D}\left[\ell^{\mathrm{hyb}}(y, \mathbf{s}(x))\right]$, the associated classifier $F(x)=\operatorname{argmax}_{y \in[K]} \widehat{s}_{i}(x)$ is the Bayes-optimal classifier for CSL (2).

## N. 2 Comparison with the ( $a, \widetilde{c}$ )-expansion Property in [14]

We compare the $c$-expansion property with $(a, \widetilde{c})$-expansion property proposed by [14], where $a \in(0,1)$ and $\widetilde{c}>1$. Here we say a distribution $Q$ on $\mathcal{X}$ satisfies the $(a, \widetilde{c})$-expansion property if $Q(\mathcal{N}(S)) \geq \widetilde{c}$ for any $S \subset \mathcal{X}$ with $Q(S) \leq a$. If $Q$ satisfies ( $a, \widetilde{c}$ )-expansion property [14] with $\widetilde{c}>1$, then $Q$ satisfies the $c$-expansion property, where the function $c$ is defined as follows. $c(p)=\widetilde{c}$ if $p \leq a$ and $c(p)=1$ otherwise. On the other hand, if $Q$-satisfies $c$-expansion property, then for any $a \in(0,1)$ and $S \subseteq \mathcal{X}$ with $Q(S) \leq a$, we have $Q(\mathcal{N}(S)) \geq c(Q(S)) Q(S) \geq c(a) Q(S)$ since $c$ is non-increasing. Therefore, $Q$ satisfies the $(a, c(a))$-expansion property. Thus, we could say these two conditions are equivalent. To simplify our analysis, we use our definition of the expansion property.

In addition, Wei et al. [14] showed that the $(a, \widetilde{c})$-expansion property is realistic for vision. Although they assumed the $(a, c)$-expansion property for each $P_{i}(1 \leq i \leq K)$ and we assume the $c$-expansion property for $\mathcal{P}_{w}$, it follows that the $c$-expansion property for $\mathcal{P}_{w}$ is also realistic for vision, since $\mathcal{P}_{w}$ is a linear combination of $P_{i}$.

## N. 3 Comparison of the Theoretical Assumptions with that of [14]

In the last paragraph of Sec. 3.1, we explain the difference in the theoretical assumptions between ours and [14]. Wei et al. [14] assumed the existence of the ground-truth classifier and the supports of $P_{i}$ are disjoint, but we cannot assume these conditions due to our problem setting (i.e., optimizing the cost sensitive objectives). In this section, we provide more intuitive explanation using a toy example. For simplicity, we assume $K=2, \mathcal{X} \subset \mathbb{R}$ and $w=\operatorname{diag}\left(w_{1}, w_{2}\right)$ with $w_{1}, w_{2} \geq 0$.

In Fig. 1, we consider the cost sensitive (weighted) objective in the case where supports of $P_{1}$ and $P_{2}$ are disjoint. As the figure indicates the Bayes optimal classifier $x \mapsto \operatorname{argmax}_{i \in[K]} w_{i} P_{i}(x)$ for the cost sensitive objective does not depend on $w$. The ground truth classifier (i.e., $x \mapsto$ $\left.\operatorname{argmax}_{i \in[K]} P_{i}(x)\right)$ is the best classifier for any $w$.

On the other hand, in Fig. 2, we consider a more generalized setting where the supports are not necessarily disjoint. In this case, the optimal classifier for the cost sensitive objective depends on $w$. This simple example suggests that we have to generalize [14] by removing the restrictive assumptions on the supports and the ground truth classifier.


Figure 1: In the perfect setting where two distributions have disjoint supports, the Bayes optimal classifier for the CSL is identical to the ground truth classifier $\left(x \mapsto \operatorname{argmax}_{i} P_{i}(x)\right)$ for any choices of weights $\left(w_{1}, w_{2}\right)$.


Figure 2: In more generalized settings, the Bayes optimal classifier for the CSL depends on the choice of weights (i.e., gain matrix). In the left figure, we put more weight on the second class than the first class. In the right figure, we put less weight on the second class than the first class. By decreasing the weight $w_{2}$, the optimal decision boundary for the CSL moves to the right.

## References

[1] Sergey G Bobkov. An isoperimetric inequality on the discrete cube, and an elementary proof of the isoperimetric inequality in gauss space. The Annals of Probability, 25(1):206-214, 1997. 3
[2] Marco Chiani, Davide Dardari, and Marvin K Simon. New exponential bounds and approximations for the computation of error probability in fading channels. IEEE Transactions on Wireless Communications, 2(4):840-845, 2003. 7
[3] Andrew Cotter, Heinrich Jiang, Maya R Gupta, Serena Wang, Taman Narayan, Seungil You, and Karthik Sridharan. Optimization with non-differentiable constraints with applications to fairness, recall, churn, and other goals. J. Mach. Learn. Res., 20(172):1-59, 2019. 2
[4] Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. Marian: Fast neural machine translation in C++. In Proceedings of ACL 2018, System Demonstrations, pages 116-121, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-4020. URL https://aclanthology.org/P18-4020. 12
[5] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 13
[6] Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick Van Kleef, Sören Auer, et al. Dbpedia-a largescale, multilingual knowledge base extracted from wikipedia. Semantic web, 6(2):167-195, 2015. 13
[7] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, pages 142150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL http://www.aclweb.org/anthology/P11-1015. 13
[8] Harikrishna Narasimhan and Aditya K Menon. Training over-parameterized models with non-decomposable objectives. Advances in Neural Information Processing Systems, 34, 2021. 3, 11, 15
[9] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 1944-1952, 2017. 3
[10] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. International journal of computer vision, 115(3):211-252, 2015. 13
[11] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. arXiv preprint arXiv:1910.01108, 2019. 9
[12] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semisupervised learning with consistency and confidence. Advances in Neural Information Processing Systems, 33:596-608, 2020. 9, 12, 13
[13] Colin Wei and Tengyu Ma. Improved sample complexities for deep neural networks and robust classification via an all-layer margin. In International Conference on Learning Representations, 2019. 7
[14] Colin Wei, Kendrick Shen, Yining Chen, and Tengyu Ma. Theoretical analysis of self-training with deep networks on unlabeled data. In International Conference on Learning Representations, 2020. 2, 3, 4, 7, 8, 15, 16
[15] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Transformers: State-of-the-art natural language processing. In Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations, pages 38-45, 2020. 9


[^0]:    ${ }^{1}$ https://github.com/LeeDoYup/FixMatch-pytorch

