## A Omitted Proofs

Corollary 9 (Nice Hinge Function - Relative Error Coreset). Consider the setting of Thm. 8 under the additional assumption that $a_{2}>0$. If $\sum_{i=1}^{n} p_{i}=m$ and $p_{i} \geq \frac{C \max \left(\tau_{i}(X), 1 / n\right) \cdot \mu(X)^{2}}{\epsilon^{2}}$ for all $i$, where $C=c \cdot \max \left(1, L, a_{1}, 1 / a_{2}\right)^{10} \cdot \log \left(\frac{\log \left(n \max \left(1, L, a_{1}, 1 / a_{2}\right) \cdot \mu(X) / \epsilon\right) m}{\delta}\right)$ and $c$ is a fixed constant, with probability at least $1-\delta$, for all $\beta \in \mathbb{R}^{d}$,

$$
\left|\sum_{i=1}^{m}[S f(X \beta)]_{i}-\sum_{i=1}^{n} f(X \beta)_{i}\right| \leq \epsilon \cdot \sum_{i=1}^{n} f(X \beta)_{i} .
$$

Proof. By (3) proven in Corollary 6 and using the fact that $f$ is $\left(L, a_{1}, a_{2}\right)$-nice,

$$
\begin{align*}
\sum_{i=1}^{n} f(X \beta)_{i} & \geq \sum_{i:[X \beta]_{i} \in\left[0,2 a_{1}\right]} f(X \beta)_{i}+\sum_{i:[X \beta]_{i} \geq 2 a_{1}} f(X \beta)_{i} \\
& \geq \sum_{i:[X \beta]_{i} \in\left[0,2 a_{1}\right]} a_{2}+\sum_{i:[X \beta]_{i} \geq 2 a_{1}} \operatorname{ReLU}(X \beta)_{i}-a_{1} \\
& \geq \min \left(\frac{a_{2}}{2 a_{1}}, \frac{1}{2}\right) \cdot\left\|(X \beta)^{+}\right\|_{1} \\
& \geq \min \left(\frac{a_{2}}{2 a_{1}}, \frac{1}{2}\right) \cdot \frac{\|X \beta\|_{1}}{\mu(X)+1} . \tag{9}
\end{align*}
$$

Let $\gamma \stackrel{\text { def }}{=} \min \left(\frac{a_{2}}{2 a_{1}}, \frac{1}{2}\right)$. Now we claim that $\sum_{i=1}^{n} f(X \beta)_{i} \geq \frac{n a_{2} \gamma}{4 \max (1, L) \cdot \mu(X)}$. If $\sum_{i=1}^{n} f(X \beta)_{i} \geq$ $\frac{n a_{2}}{4}$ then this holds immediately since $\mu(X) \geq 1, \max (1, L) \geq 1$ and $\gamma \leq 1$. Otherwise, assume that $\sum_{i=1}^{n} f(X \beta)_{i} \leq \frac{n a_{2}}{4}$. Since $f(z) \geq a_{2}$ for all $z \geq 0$ and since $f$ is $L$-Lipschitz, $f(z) \geq \frac{a_{2}}{2}$ for all $z \geq-\frac{a_{2}}{2 L}$. This implies that $X \beta$ has at most $\frac{n a_{2} / 4}{a_{2} / 2}=\frac{n}{2}$ entries $\geq-\frac{a_{2}}{2 L}$. Thus, $X \beta$ has at least $\frac{n}{2}$ entries $\leq-\frac{a_{2}}{2 L}$ and so $\left\|(X \beta)^{-}\right\|_{1} \geq \frac{n a_{2}}{4 L}$. Thus, by the definition of $\mu(X)$ along with (9),

$$
\begin{equation*}
\sum_{i=1}^{n} f(X \beta)_{i} \geq \gamma \cdot\left\|(X \beta)^{+}\right\|_{1} \geq \frac{n a_{2} \gamma}{4 L \cdot \mu(X)} \geq \frac{n a_{2} \gamma}{4 \max (1, L) \cdot \mu(X)} \tag{10}
\end{equation*}
$$

Combining (9) with (10) gives that

$$
\sum_{i=1}^{n} f(X \beta)_{i} \geq \frac{\gamma \cdot\|X \beta\|_{1}}{2 \mu(X)+2}+\frac{n a_{2} \gamma}{8 \max (1, L) \cdot \mu(X)} \geq\left(\|X \beta\|_{1}+n\right) \cdot \frac{\gamma \cdot \min \left(1, a_{2}\right)}{8 \max (1, L) \cdot \mu(X)+2}
$$

This completes the corollary after applying Thm. 8 with

$$
\epsilon^{\prime}=\epsilon \cdot \frac{\gamma \cdot \min \left(1, a_{2}\right)}{8 \max (1, L) \cdot \mu(X)+2} \geq \frac{\epsilon}{8 \max \left(1, L, a_{1}, 1 / a_{2}\right)^{4} \cdot \mu(X)+2} .
$$

## B Lower Bounds for Regularized Classification

We now give a lower bound showing that the results of [CIM ${ }^{+}$19] on coresets for regularized logistic and hinge loss regression (i.e., soft margin SVM) are essentially tight. Our bound tightens a lower bound given in $\left[\mathrm{CIM}^{+}\right.$19]. It shows that, in the natural setting where the regularization parameter is sublinear in the number of data points $n$, the coreset size must depend polynomially on $n$. This contrasts the setting where we assume that $\mu(X)$ from Def. 1 is bounded. In this case, as shown in Cor. 9, relative error coresets with size scaling just logarithmically in $n$ are achievable.
Theorem 10 (Regularized Classification - Relative Error Lower Bound). Let $X \in \mathbb{R}^{n \times d}$ have all row norms bounded by 1 . Let $f$ be the hinge loss $f(z)=\max (0,1+z)$ or $\log \operatorname{loss} f(z)=\ln \left(1+e^{z}\right)$ and for any $\kappa \in(0,1)$ consider the regularized loss $L: \mathbb{R}^{d} \rightarrow \mathbb{R}^{+}$,

$$
L(\beta)=\sum_{i=1}^{n} f(X \beta)_{i}+n^{\kappa} \cdot R(\beta),
$$

where $\kappa \in(0,1)$. There is no $O(1)$ relative error coreset for $L(\beta)$ with $o\left(\frac{n^{1-\kappa}}{\log ^{c} n}\right)$ points where $c=4$ for $R(\beta)=\|\beta\|_{2}^{2}, c=5 / 2$ for $R(\beta)=\|\beta\|_{2}$, and $c=3$ for $R(\beta)=\|\beta\|_{1}$.

Note that since this is a lower bound, the assumption that $X$ has bounded row norms only makes it stronger. This assumption is common in prior work.

Proof. We focus on the case when $f$ is the hinge loss for simplicity. An identical argument applies when $f$ is the log loss, with some adjustments of the constants. We also focus on the case when $R(\beta)=\|\beta\|_{2}^{2}$. Again, essentially an identical argument proves the claim when $R(\beta)=\|\beta\|_{2}$ or $R(\beta)=\|\beta\|_{1}$. We prove the lower bound via a reduction from the INDEX problem in communication complexity. Alice has a string $a \in\{0,1\}^{n}$ and Bob has an index $b \in\{1, \ldots, n\}$, and they wish to compute the bit $a(b)$. It is well known that the randomized 1-way communication complexity of this problem is $\Omega(n)$ [Rou15]. We will show that the existence of a relative error coreset for $L(x)$ with $o\left(\frac{n^{1-\kappa}}{\log ^{4} n}\right)$ points would contradict this lower bound, giving the result.
Assume without loss of generality that $n^{1-\kappa}$ is a power of two. Let $d=\log _{2} n^{1-\kappa}$. Our reduction is to the INDEX problem with input size $n_{0}=\frac{n^{1-\kappa}}{d(d+1)^{2}}=\Theta\left(\frac{n^{1-\kappa}}{\log ^{3} n}\right)$. Let Alice construct the matrix $X_{0} \in \mathbb{R}^{n_{0} \times(d+1)}$ which has the first $d$ entries of row $i$ equal to the binary representation of $i$ if $a(i)=1$ and equal to 0 otherwise. In the binary representation, have 0 represented by -1 and 1 represented by 1 . Let every row have $d$ in the last column. Finally, scale the matrix by a $\gamma=1 / \sqrt{d^{2}+d}$ factor so each row has Euclidean norm exactly 1 . Let $X \in \mathbb{R}^{n \times(d+1)}$ be equal to $n^{\kappa} \cdot d(d+1)^{2}$ copies of $X_{0}$ stacked on top of each other (assume without loss of generality that $n^{\kappa} \cdot d(d+1)^{2}$ is an integer).
Bob will let $\beta \in \mathbb{R}^{d+1}$ be the binary representation for $b$ (again written using -1 s and 1 s ) with a -1 in the last entry. He will scale $\beta$ by a $1 / \gamma$ factor so $\|\beta\|_{2}^{2}=(d+1) \cdot\left(d^{2}+d\right)=d(d+1)^{2}$. If $a(b)=1$ we have:

$$
\begin{align*}
L(\beta) & =n^{\kappa} \cdot d(d+1)^{2} \cdot\left(\sum_{j \neq b} h(X \beta)_{j}+h(X \beta)_{b}\right)+n^{\kappa}\|\beta\|_{2}^{2} \\
& =n^{\kappa} \cdot d(d+1)^{2}+n^{\kappa} \cdot d(d+1)^{2}=2 n^{\kappa} \cdot d(d+1)^{2} \tag{11}
\end{align*}
$$

where the second line holds since for $j \neq b,[X \beta]_{j} \leq d-1-d \leq-1$ and so $h(X \beta)_{j}=0$. $[X \beta]_{b}=d-d=0$ and so $h(X \beta)_{b}=1$. Otherwise, by the same logic, if $a(b)=0$ we have:

$$
\begin{equation*}
L(\beta)=n^{\kappa} \cdot d(d+1)^{2} \cdot\left(\sum_{j \neq b} h(X \beta)_{j}+h(X \beta)_{b}\right)+n^{\kappa}\|\beta\|_{2}^{2}=n^{\kappa} \cdot d(d+1)^{2} \tag{12}
\end{equation*}
$$

From (11) and (12), we can see that a coreset with relative error $\epsilon=1 / 2$ can distinguish the two cases of $a(b)=1$ and $a(b)=0$. Assume that there is such a relative error coreset consisting of $m$ rows of $X$, along with $m$ corresponding weights $w_{1}, \ldots, w_{m}$. We can assume that all $w_{j} \leq n^{c_{1}}$ for some large constant $c_{1}$. If $a\left(i_{j}\right)=1$ any $w_{j}$ larger than this would lead to the coreset cost being a large over estimate when $b=i_{j}$. If $a\left(i_{j}\right)=0$, then scaling the $i_{j}^{t h}$ row by any $w_{j}$ will have no effect since for all $\beta$ that Bob may generate, $h(X \beta)_{i_{j}}=0$. So again, we can assume $w_{j} \leq n^{c_{1}}$.
Additionally, if we round each $w_{j}$ to the nearest integer multiple of $1 / n^{c_{1}}$ we will not change the coreset cost by more than a $n / n^{c_{1}}$ factor in all our input cases, since we always have $h(X \beta)_{i} \in[0,1]$. Thus, Alice can represent each rounded $w_{j}$ using $\log n$ bits and send the full coreset and weights to Bob using $O(m \cdot(\log n+d))=O(m \log n)$ bits of communication. Since Bob can then use this coreset to solve the INDEX with input size $n_{0}=\Theta\left(\frac{n^{1-\kappa}}{\log ^{3} n}\right)$, we must have $m=\Omega\left(\frac{n^{1-\kappa}}{\log ^{4} n}\right)$, proving the theorem.
In the case that $R(\beta)=\|\beta\|_{2}$ we have $\|\beta\|_{2}=d^{1 / 2}(d+1)=\Theta\left(d^{3 / 2}\right)$ and so can set $n_{0}=\Theta\left(\frac{n^{1-\kappa}}{\log ^{3 / 2} n}\right)$ instead of $n_{0}=\Theta\left(\frac{n^{1-\kappa}}{\log ^{3} n}\right)$, which gives the final lower bound of $\Omega\left(\frac{n}{\log ^{5 / 2} n}\right)$. Similarly, for $R(\beta)=\|\beta\|_{1}$, we have $\|\beta\|_{1}=d^{1 / 2}(d+1)^{3 / 2}=\Theta\left(d^{2}\right)$, yielding a final bound of $\Omega\left(\frac{n}{\log ^{3} n}\right)$.

Finally, we compare our lower bound with the the bound in [TBFR21]. We first note that the lower bound in the referenced paper is a lower bound on the sum of sensitivities, rather than directly on the coreset size, as we have given. We are not aware of a general result which lower bounds coreset size via the sum of sensitivities, although perhaps such a result could be shown, at least for reasonable classes of loss functions.
If we set $\lambda$ in [TBFR21] to $n^{-\kappa}$, then we are in the same setting as our lower bound, with regularization $n^{\kappa}\|\beta\|$. In this setting, assuming that $d<n^{1-\kappa}$, then the lower bound given in Lemma 1 of [TBFR21] is $O\left(n \lambda / d^{2}\right)=O\left(n^{1-\kappa} / d^{2}\right)$. This is loose by a $d^{2}$ factor, as compared to our tight lower bound.

