

Supplementary Material: Data-Efficient Augmentation for Training Neural Networks

A Proof of Main Results

A.1 Proof for Lemma 4.1

Proof. Let $\delta_i := \tilde{\sigma}_i - \sigma_i$, where $\mathbb{P}(\delta_i < 0) = p_i$. Assuming uniform probability between $-\|\mathbf{E}\|$ to 0, and between 0 to $\|\mathbf{E}\|$, we have pdf $\rho_i(x)$ for δ_i :

$$\rho_i(x) = \begin{cases} \frac{p_i}{\|\mathbf{E}\|}, & \text{if } -\|\mathbf{E}\| \leq x < 0 \\ \frac{1-p_i}{\|\mathbf{E}\|}, & 0 \leq x \leq \|\mathbf{E}\| \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

Taking expectation,

$$\mathbb{E}(\tilde{\sigma}_i - \sigma_i) = \mathbb{E}(\delta_i) = \int_{-\infty}^{\infty} x \rho_i(x) dx \quad (2)$$

$$= \int_{-\|\mathbf{E}\|}^0 x \frac{p_i}{\|\mathbf{E}\|} dx + \int_0^{\|\mathbf{E}\|} x \frac{1-p_i}{\|\mathbf{E}\|} dx \quad (3)$$

$$= -\frac{\|\mathbf{E}\| p_i}{2} + \frac{(1-p_i)\|\mathbf{E}\|}{2} \quad (4)$$

$$= \frac{(1-2p_i)\|\mathbf{E}\|}{2} \quad (5)$$

We also have

$$\mathbb{E}(\delta_i^2) = \int_{-\infty}^{\infty} x^2 \rho_i(x) dx \quad (6)$$

$$= \int_{-\|\mathbf{E}\|}^0 x^2 \frac{p_i}{\|\mathbf{E}\|} dx + \int_0^{\|\mathbf{E}\|} x^2 \frac{1-p_i}{\|\mathbf{E}\|} dx \quad (7)$$

$$= \frac{\|\mathbf{E}\|^2 p_i}{3} + \frac{(1-p_i)\|\mathbf{E}\|^2}{3} \quad (8)$$

$$= \frac{\|\mathbf{E}\|^2}{3} \quad (9)$$

Thus, we have

$$\mathbb{E}(\tilde{\lambda}_i) = \mathbb{E}(\tilde{\sigma}_i^2) \quad (10)$$

$$= \mathbb{E}((\sigma_i + \delta_i)^2) \quad (11)$$

$$= \mathbb{E}(\sigma_i^2 + 2\sigma_i\delta_i + \delta_i^2) \quad (12)$$

$$= \sigma_i^2 + 2\sigma_i\mathbb{E}[\delta_i] + \mathbb{E}[\delta_i^2] \quad (13)$$

$$= \sigma_i^2 + 2\sigma_i \frac{(1-2p_i)\|\mathbf{E}\|}{2} + \frac{\|\mathbf{E}\|^2}{3} \quad (14)$$

$$= \sigma_i^2 + \sigma_i(1-2p_i)\|\mathbf{E}\| + \frac{\|\mathbf{E}\|^2}{3} \quad (15)$$

□

A.2 Proof of Corollary 4.2

Under the assumptions of Theorem 5.1 of [1], i.e. where the minimum eigenvalue of the NTK is $\lambda_{\min}(\mathcal{J}\mathcal{J}^T) \geq \lambda_0$ for a constant $\lambda_0 > 0$, and training data \mathbf{X} of size n sampled i.i.d. from distribution D and 1-Lipschitz loss \mathcal{L} , we have that with probability $\delta/3$, training the over-parameterized neural network with gradient descent for $t \geq \Omega\left(\frac{1}{n\lambda_0} \log \frac{n}{\delta}\right)$ iterations results in the following population loss \mathcal{L}_D (generalization error)

$$\mathcal{L}_D(\mathbf{W}^t, \mathbf{X}) \leq \sqrt{\frac{2\mathbf{y}^T(\mathcal{J}\mathcal{J}^T)^{-1}\mathbf{y}}{n}} + \mathcal{O}\left(\frac{\log \frac{n}{\lambda_0\delta}}{n}\right), \quad (16)$$

with high probability of at least $1 - \delta$ over random initialization and training samples.

Hence, using $\lambda_{\min}, \sigma_{\min}$ to denote minimum eigen and singular value respectively of the NTK corresponding to full data, we get

$$\mathcal{L}_{D_{train}}(\mathbf{W}^t, \mathbf{X}_{train}) \leq \sqrt{\frac{2\frac{1}{\lambda_{\min}}\|y\|^2}{n}} + \mathcal{O}\left(\log \frac{1}{\delta}\right) \quad (17)$$

$$\leq \sqrt{\frac{2}{\sigma_{\min}^2}} + \mathcal{O}\left(\log \frac{1}{\delta}\right). \quad (18)$$

For augmented dataset \mathbf{X}_{aug} , we have $\tilde{\sigma}_i \leq \sigma_i + \sqrt{n}L\epsilon_0$, hence the improvement in the generalization error is at most

$$\mathcal{L}_{D_{aug}}(\mathbf{W}^t, \mathbf{X}_{aug}) \leq \sqrt{\frac{2}{(\sigma_{\min} + \sqrt{n}L\epsilon_0)^2}} + \mathcal{O}\left(\log \frac{1}{\delta}\right). \quad (19)$$

Combining these two results, we obtain Corollary 4.2.

A.3 Proof of Lemma 5.1

Proof.

$$\|\mathcal{J}^T(\mathbf{W}^t, \mathbf{X}_{aug})\mathbf{r} - \text{diag}(\boldsymbol{\rho}^t)\mathcal{J}^t(\mathbf{W}^t, \mathbf{X}_{S_{aug}})\mathbf{r}_S\| \quad (20)$$

$$= \|(\mathcal{J}^T(\mathbf{W}^t, \mathbf{X}) + \mathbf{E})\mathbf{r} - (\text{diag}(\boldsymbol{\rho}^t)\mathcal{J}^t(\mathbf{W}^t, \mathbf{X}_S) + \mathbf{E}_S)\mathbf{r}_S\| \quad (21)$$

$$\leq \|(\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r} - \text{diag}(\boldsymbol{\rho}^t)\mathcal{J}^t(\mathbf{W}^t, \mathbf{X}_S)\mathbf{r}_S) + \mathbf{E}\mathbf{r} - \mathbf{E}_S\mathbf{r}_S\| \quad (22)$$

$$\leq \|(\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r} - \text{diag}(\boldsymbol{\rho}^t)\mathcal{J}^t(\mathbf{W}^t, \mathbf{X}_S)\mathbf{r}_S)\| + \|\mathbf{E}\mathbf{r}\| + \|\mathbf{E}_S\mathbf{r}_S\| \quad (23)$$

Applying definition of coresets, we obtain

$$\|(\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r} - \text{diag}(\boldsymbol{\rho}^t)\mathcal{J}^t(\mathbf{W}^t, \mathbf{X}_S)\mathbf{r}_S)\| + \|\mathbf{E}\mathbf{r}\| + \|\mathbf{E}_S\mathbf{r}_S\| \quad (24)$$

$$\leq \xi + \|\mathbf{E}\mathbf{r}\| + \|\mathbf{E}_S\mathbf{r}_S\| \quad (25)$$

$$\leq \xi + 2n^{\frac{3}{2}}L\epsilon_0 \quad (26)$$

$$\leq \xi + 2\sqrt{L} \quad (27)$$

□

A.4 Proof of Theorem 5.2

Proof. In this proof, as shorthand notation, we use \mathbf{X}_f and \mathbf{X}_{train} interchangeably. We further use \mathbf{X}_c to represent the coreset selected from the full data, and $\mathbf{X}_{c_{aug}}$ to represent the augmented coreset.

By Theorem 1 of [2], under the α -PL assumption for \mathcal{L} and interpolation assumption (i.e. for every sequence $\mathbf{W}^1, \mathbf{W}^2, \dots$ such that $\lim_{t \rightarrow \infty} \mathcal{L}(\mathbf{W}^t, \mathbf{X}) = 0$, we have that the loss for each data point $\lim_{t \rightarrow \infty} \mathcal{L}(\mathbf{W}^t, \mathbf{x}_i) = 0$), the convergence of SGD with constant step size is given by

$$\mathbb{E}[\|\nabla \mathcal{L}(\mathbf{W}^t, \mathbf{X}_{f+c_{aug}})\|^2] \leq \left(1 - \frac{\alpha\eta}{2}\right)^t \mathcal{L}(\mathbf{W}^0, \mathbf{X}_{f+c_{aug}}) \quad (28)$$

$$\leq \frac{1}{\alpha} \left(1 - \frac{\alpha\eta}{2}\right)^t \|\nabla \mathcal{L}(\mathbf{W}^0, \mathbf{X}_{f+c_{aug}})\|^2 \quad (29)$$

Using Jensen's inequality, we have

$$\mathbb{E}[\|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_{f+c_{\text{aug}}})\|] \quad (30)$$

$$\leq \sqrt{\mathbb{E}[\|\nabla\mathcal{L}(\mathbf{W}^t, \mathbf{X}_{f+c_{\text{aug}}})\|^2]} \quad (31)$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} \|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_{f+c_{\text{aug}}})\| \quad (32)$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} (\|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_f)\| + \|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_{c_{\text{aug}}})\|) \quad (33)$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} (G_0 + \|(\mathcal{J}(\mathbf{W}^0, \mathbf{X}_c) + \mathbf{E})(\mathbf{y} - f(\mathbf{W}^0, \mathbf{X}_c + \epsilon))\|) \quad (34)$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} \quad (35)$$

$$(G_0 + \|(\mathcal{J}(\mathbf{W}^0, \mathbf{X}_c) + \mathbf{E})^T(\mathbf{y} - f(\mathbf{W}^0, \mathbf{X}_c) - \nabla_x f(\mathbf{W}^0, \mathbf{X}_c)^T \epsilon - \mathcal{O}(\epsilon^T \epsilon))\|) \quad (36)$$

$$= \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} (G_0 + \|\nabla L(\mathbf{W}^0, \mathbf{X}_c) - (\mathcal{J}(\mathbf{W}^0, \mathbf{X}_c)^T (\nabla_x f(\mathbf{W}^0, \mathbf{X}_c)^T \epsilon + \mathcal{O}(\epsilon^T \epsilon)) + \mathbf{E}(\mathbf{y} - f(\mathbf{W}^0, \mathbf{X}_c + \epsilon))\|) \quad (37)$$

$$\quad (38)$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} (G_0 + \|\nabla L(\mathbf{W}^0, \mathbf{X}_c) - (\mathcal{J}(\mathbf{W}^0, \mathbf{X}_c)^T (\nabla_x f(\mathbf{W}^0, \mathbf{X}_c)^T \epsilon + \mathcal{O}(\epsilon^T \epsilon))\| + \sqrt{2}\|\mathbf{E}\|) \quad (39)$$

$$\quad (40)$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} (G_0 + \|\nabla L(\mathbf{W}^0, \mathbf{X}_c)\| + \sigma_{\max} \bar{L} \sqrt{n} \epsilon_0 + \sigma_{\max} \mathcal{O}(n \epsilon_0^2)) + \sqrt{2n} \bar{L} \epsilon_0 \quad (41)$$

$$= \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} (G_0 + \|\nabla L(\mathbf{W}^0, \mathbf{X}_f)\| + \xi + \sigma_{\max} \bar{L} \sqrt{n} \epsilon_0 + \sigma_{\max} \mathcal{O}(n \epsilon_0^2)) + \sqrt{2n} \bar{L} \epsilon_0 \quad (42)$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} (2G_0 + \xi + \sigma_{\max} \bar{L} \sqrt{n} \epsilon_0 + \sigma_{\max} \mathcal{O}(n \epsilon_0^2)) + \sqrt{2n} \bar{L} \epsilon_0 \quad (43)$$

□

A.5 Finding Subsets

Let S be a subset of training data points. Furthermore, assume that there is a mapping $\pi_{w,S} : V \rightarrow S$ that for every \mathbf{W} assigns every data point $i \in V$ to its closest element $j \in S$, i.e. $j = \pi_{w,S}(i) = \arg \max_{j' \in S} s_{ij'}(\mathbf{W})$, where $s_{ij}(\mathbf{W}) = C - \|\mathcal{J}^T(\mathbf{W}^t, \mathbf{x}_i)r_i - \mathcal{J}^T(\mathbf{W}^t, \mathbf{x}_j)r_j\|$ is the similarity between gradients of i and j , and $C \geq \max_{i,j} s_{ij}(\mathbf{W})$ is a constant. Consider a matrix $\mathbf{G}_{\pi_{w,S}} \in \mathbb{R}^{n \times m}$, in which every row i contains gradient of $\pi_w(i)$, i.e., $[\mathbf{G}_{\pi_{w,S}}]_i = \mathcal{J}^T(\mathbf{W}^t, \mathbf{x}_{\pi_{w,S}(i)})r_{\pi_{w,S}(i)}$. The Frobenius norm of the matrix $\mathbf{G}_{\pi_{w,S}}$ provides an upper-bound on the error of the weighted subset S in capturing the alignment of the residuals of the full training data with the Jacobian matrix. Formally,

$$\|\mathcal{J}^T(\mathbf{W}^t, \mathbf{X}_{\text{train}})\mathbf{r}_{\text{train}}^t - \gamma_{S^t} \mathcal{J}^T(\mathbf{W}^t, [\mathbf{X}_{\text{train}}]_{.S^t})\mathbf{r}_{S^t}\| \leq \|\mathbf{G}_{\pi_{w,S}}\|_F, \quad (44)$$

where the weight vector $\gamma_{S^t} \in \mathbb{R}^{|S|}$ contains the number of elements that are mapped to every element $j \in S$ by mapping $\pi_{w,S}$, i.e. $\gamma_j = \sum_{i \in V} \mathbb{1}[\pi_{w,S}(i) = j]$. Hence, the set of training points that closely estimate the projection of the residuals of the full training data on the Jacobian spectrum can be obtained by finding a subset S that minimizes the Frobenius norm of matrix $\mathbf{G}_{\pi_{w,S}}$.

B Additional Theoretical Results

B.1 Convergence analysis for training on augmented full data

Theorem B.1. *Gradient descent with learning rate η applied to a neural network with constant NTK and Lipschitz constant L , and data points \mathcal{D}_{aug} augmented with r additive perturbations bounded by ϵ_0 results in the following training dynamics:*

$$\mathbb{E}[\|\mathbf{y} - f(\mathbf{X}_{aug}, \mathbf{W}^t)\|_2] \leq \sqrt{\sum_{i=1}^n \left(1 - \eta \left(\sigma_i^2 + \sigma_i(1 - 2p_i)\|E\| + \frac{\|E\|^2}{3}\right)\right)^{2t} ((\mathbf{u}_i \mathbf{y})^2 + 2n\sqrt{2}\|E\|/\gamma_0)} \quad (45)$$

where \mathbf{E} with $\|\mathbf{E}\| \leq \sqrt{n}L\epsilon_0$ is the perturbation to the Jacobian, and $p_i := \mathbb{P}(\tilde{\sigma}_i - \sigma_i < 0)$ is the probability that σ_i decreases as a result of data augmentation.

B.2 Proof of Theorem B.1

Using Jensen's inequality, we have

$$\mathbb{E}[\|\mathbf{y} - f(\mathbf{X}_{aug}, \mathbf{W}^t)\|_2] \quad (46)$$

$$= \mathbb{E} \left[\sqrt{\sum_{i=1}^n (1 - \eta \tilde{\lambda}_i)^{2t} (\tilde{\mathbf{u}}_i^T \mathbf{y})^2} \pm \epsilon \right] \quad (47)$$

$$\leq \sqrt{\mathbb{E} \left[\sum_{i=1}^n (1 - \eta \tilde{\lambda}_i)^{2t} (\tilde{\mathbf{u}}_i^T \mathbf{y})^2 \right]} \quad (48)$$

$$\leq \sqrt{\sum_{i=1}^n \mathbb{E} \left[(1 - \eta \tilde{\lambda}_i)^{2t} ((\mathbf{u}_i \mathbf{y})^2 + 2n\sqrt{2}\|E\|/\gamma_0) \right]} \quad (49)$$

$$\leq \sqrt{\sum_{i=1}^n (1 - \eta \mathbb{E}[\tilde{\lambda}_i])^{2t} ((\mathbf{u}_i \mathbf{y})^2 + 2n\sqrt{2}\|E\|/\gamma_0)} \quad (50)$$

$$= \sqrt{\sum_{i=1}^n \left(1 - \eta \left(\sigma_i^2 + \sigma_i(1 - 2p_i)\|E\| + \frac{\|E\|^2}{3}\right)\right)^{2t} ((\mathbf{u}_i \mathbf{y})^2 + 2n\sqrt{2}\|E\|/\gamma_0)} \quad (51)$$

B.3 Convergence analysis for training on the coreset and its augmentation

Theorem B.2. *Let \mathcal{L}_i be β -smooth, \mathcal{L} be λ -smooth and satisfy the α -PL condition, that is for $\alpha > 0$, $\|\nabla \mathcal{L}(\mathbf{W}, \mathbf{X})\|^2 \geq \alpha \mathcal{L}(\mathbf{W}, \mathbf{X})$ for all weights \mathbf{W} . Let ξ upper-bound the normed difference in gradients between the weighted coreset and full dataset. Assume that the network $f(\mathbf{W}, \mathbf{X})$ is Lipschitz in \mathbf{W} , \mathbf{X} with Lipschitz constant L and L' respectively, and $\bar{L} = \max\{L, L'\}$. Let G_0 the gradient over the full dataset at initialization, $\frac{\sigma_{\max}}{\sigma_{\max} \sqrt{Ln}}$ the maximum Jacobian singular value at initialization. Choosing perturbation bound $\epsilon_0 \leq \frac{1}{\sigma_{\max} \sqrt{Ln}}$ where σ_{\max} is the maximum singular value of the coreset Jacobian and n is the size of the original dataset, running SGD on the coreset and its augmentation using constant step size $\eta = \frac{\alpha}{\lambda\beta}$, we get the following convergence bound:*

$$\mathbb{E}[\|\nabla \mathcal{L}(\mathbf{W}^t, \mathbf{X}_{c+caug})\|] \leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} \left(2G_0 + 2\xi + \mathcal{O}\left(\frac{\bar{L}}{\sigma_{\max}}\right)\right), \quad (52)$$

where \mathbf{X}_{c+caug} represents the dataset containing the (weighted) coreset and its augmentation.

Proof. As in the proof for Theorem 5.2, we begin with the following inequality

$$\mathbb{E}[\|\nabla\mathcal{L}(\mathbf{W}^t, \mathbf{X}_{c+c_{\text{aug}}})\|^2] \leq \left(1 - \frac{\alpha\eta}{2}\right)^t \mathcal{L}(\mathbf{W}^0, \mathbf{X}_{c+c_{\text{aug}}}) \quad (53)$$

$$\leq \frac{1}{\alpha} \left(1 - \frac{\alpha\eta}{2}\right)^t \|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_{c+c_{\text{aug}}})\|^2 \quad (54)$$

Thus, we can write

$$\mathbb{E}[\|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_{c+c_{\text{aug}}})\|] \quad (55)$$

$$\leq \sqrt{\mathbb{E}[\|\nabla\mathcal{L}(\mathbf{W}^t, \mathbf{X}_{c+c_{\text{aug}}})\|^2]} \quad (56)$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} \|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_{c+c_{\text{aug}}})\| \quad (57)$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} (\|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_c)\| + \|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_{c_{\text{aug}}})\|) \quad (58)$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} (G_0 + \xi + \|(\mathcal{J}(\mathbf{W}^0, \mathbf{X}_c) + \mathbf{E})(-f(\mathbf{W}^0, \mathbf{X}_c + \epsilon))\|) \quad (59)$$

The rest of the proof is similar to that of Theorem 5.2. \square

B.4 Lemma for eigenvalues of coresets

The following Lemma characterizes the sum of eigenvalues of the NTK associated with the coreset.

Lemma B.3. *Let ξ be an upper bound of the normed difference in gradient of the weighted coreset and the original dataset, i.e. for full data \mathbf{X} and its corresponding coreset \mathbf{X}_S with weights γ_S , and respective residuals \mathbf{r} , \mathbf{r}_S , we have the bound $\|\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r}^t - \gamma_S\mathcal{J}^T(\mathbf{W}^t, \mathbf{X}_S)\mathbf{r}_S^t\| \leq \xi$. Let $\{\lambda_i\}_{i=1}^k$ be the eigenvalues of the NTK associated with the coreset. Then we have that*

$$\sqrt{\sum_{i=1}^k \lambda_i} \geq \frac{\|\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r}^t\| - \xi}{\|\mathbf{r}_S^t\|}.$$

Proof. Let singular values of coreset Jacobian be σ_i . Let $\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r}^t = \gamma_S\mathcal{J}^T(\mathbf{W}^t, \mathbf{X}_S)\mathbf{r}_S^t + \xi_S$ where $\|\xi_S\| \leq \xi$.

Taking Frobenius norm, we get

$$\|\gamma_S\mathcal{J}^T(\mathbf{W}^t, \mathbf{X}_S)\mathbf{r}_S^t\| = \|\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r}^t - \xi_S\| \quad (60)$$

$$\Rightarrow \|\gamma_S\mathcal{J}^T(\mathbf{W}^t, \mathbf{X}_S)\| \|\mathbf{r}_S^t\| \geq \|\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r}^t - \xi_S\| \quad (61)$$

$$\Rightarrow \|\gamma_S\mathcal{J}^T(\mathbf{W}^t, \mathbf{X}_S)\| \geq \frac{\|\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r}^t - \xi_S\|}{\|\mathbf{r}_S^t\|} \quad (62)$$

$$\Rightarrow \sqrt{\sum_{i=1}^s \sigma_i^2} \geq \frac{\|\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r}^t - \xi_S\|}{\|\mathbf{r}_S^t\|} \quad (63)$$

$$\Rightarrow \sqrt{\sum_{i=1}^s \lambda_i} \geq \frac{\|\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r}^t - \xi_S\|}{\|\mathbf{r}_S^t\|} \quad (64)$$

$$\Rightarrow \sqrt{\sum_{i=1}^s \lambda_i} \geq \frac{\|\mathcal{J}^T(\mathbf{W}^t, \mathbf{X})\mathbf{r}^t\| - \xi}{\|\mathbf{r}_S^t\|} \quad \text{by reverse triangle inequality} \quad (65)$$

\square

We can make the following observations: For overparameterized networks, with bounded activation functions and labels, e.g. softmax and one-hot encoding, the norm of the residual vector is bounded,

and the gradient norm is likely to be much larger than residual, especially when dimension of gradient is large. In this case, the Jacobian matrix associated with small weighted coresets found by solving Eq. (9), have large singular values.

B.5 Augmentation as Linear Transformation: Linear Model Analysis

We introduce a simplified linear model to extend our theoretical analysis to augmentations modelled as linear transformation matrices F applied to the original training data. These augmentations are also originally studied by [7]. In this section, we specifically study the effect of these augmentations using a linear model when applied to coresets.

Lemma B.4 (Augmented coreset gradient bounds: Linear). *Let f be a simple linear model with weights $\mathbf{W} \in \mathbb{R}^{d \times C}$ where $f(\mathbf{W}, \mathbf{x}_i) = \mathbf{W}^T \mathbf{x}_i$, trained on mean squared loss function \mathcal{L} . Let $F \in \mathbb{R}^{d \times d}$ be a common linear augmentation matrix with norm $\|F\|$ with augmentation $\mathbf{x}_i^{\text{aug}}$ given by $F\mathbf{x}_i$. Let coreset be of size k and full dataset be of size n . Further assume that the predicted label of \mathbf{x}_i and its augmentation $F\mathbf{x}_i$ are sufficiently close, i.e. there exists ω such that $\mathbf{W}^T (F\mathbf{x}_i) = \mathbf{W}^T \mathbf{x}_i + z_i$, $\|z_i\| \leq \omega \forall i$. Let ξ upper-bound the normed difference in gradients between the weighted coreset and full dataset. Then, the normed difference between the gradient of the augmented full data and augmented coreset is given by*

$$\left\| \sum_{i \in V} \nabla \mathcal{L}(\mathbf{W}, \mathbf{x}_i^{\text{aug}}) - \sum_{j=1}^k \gamma_{s_j} \nabla \mathcal{L}(\mathbf{W}, \mathbf{x}_{s_j}^{\text{aug}}) \right\| \leq \|F\|(\xi + \sqrt{dn}\omega)$$

for some (small) constant ξ .

Proof. By our assumption, we can begin with,

$$\left\| \sum_{i \in V} \nabla \mathcal{L}(\mathbf{W}, \mathbf{x}_i) - \sum_{j=1}^k \gamma_{s_j} \nabla \mathcal{L}(\mathbf{W}, \mathbf{x}_{s_j}) \right\| \leq \xi \quad (66)$$

Furthermore, by [6], we know that sum of the coreset weights γ_{s_j} is given by $\sum_{j=1}^k \gamma_{s_j} \leq n$.

Hence,

$$\left\| \sum_{i \in V} \nabla \mathcal{L}(\mathbf{W}, \mathbf{x}_i^{\text{aug}}) - \sum_{j=1}^k \gamma_{s_j} \nabla \mathcal{L}(\mathbf{W}, \mathbf{x}_{s_j}^{\text{aug}}) \right\| \quad (67)$$

$$= \left\| \sum_{i \in V} (\mathcal{J}(\mathbf{W}, \mathbf{x}_i^{\text{aug}}))^T [\mathbf{W}^T (F\mathbf{x}_i) - y_i] - \sum_{j=1}^k \gamma_{s_j} (\mathcal{J}(\mathbf{W}, \mathbf{x}_{s_j}^{\text{aug}}))^T [\mathbf{W}^T (F\mathbf{x}_{s_j}) - y_{s_j}] \right\| \quad (68)$$

$$= \left\| \sum_{i \in V} F\mathbf{x}_i [\mathbf{W}^T (F\mathbf{x}_i) - y_i] - \sum_{j=1}^k \gamma_{s_j} F\mathbf{x}_{s_j} [\mathbf{W}^T (F\mathbf{x}_{s_j}) - y_{s_j}] \right\| \quad (69)$$

$$= \left\| F \sum_{i \in V} \mathbf{x}_i (\mathbf{W}^T \mathbf{x}_i - y_i) - F \sum_{j=1}^k \gamma_{s_j} \mathbf{x}_{s_j} (\mathbf{W}^T \mathbf{x}_{s_j} + z_i - y_{s_j}) \right\| \quad (70)$$

$$= \left\| F \sum_{i \in V} \nabla L(\mathbf{W}, \mathbf{x}_i) - F \sum_{j=1}^k \gamma_{s_j} \nabla L(\mathbf{W}, \mathbf{x}_{s_j}) - F \sum_{j=1}^k \gamma_{s_j} \mathbf{x}_{s_j} z_{s_j} \right\| \quad (71)$$

$$\leq \|F\| \left\| \sum_{i \in V} \nabla L(\mathbf{W}, \mathbf{x}_i) - \sum_{j=1}^k \gamma_{s_j} \nabla L(\mathbf{W}, \mathbf{x}_{s_j}) \right\| + \|F\| \left\| \sum_{j=1}^k \gamma_{s_j} \mathbf{x}_{s_j} z_{s_j} \right\| \quad (72)$$

$$\leq \|F\|(\xi + \sqrt{d})\|F\|n\omega \quad (73)$$

$$= \|F\|(\xi + \sqrt{dn}\omega) \quad (74)$$

□

Corollary B.5. *In the simplified linear case above, the difference in gradients of the full training data with its augmentations ($\nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{f+aug})$) and gradients of the coreset with its augmentations ($\nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{c+caug})$) can be bounded by*

$$\|\nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{f+aug}) - \nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{c+caug})\| \leq (\|F\|+1)\xi + \sqrt{d}\|F\|n\omega$$

Proof. Applying Eq. (66) and Lemma B.4, we obtain

$$\|\nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{f+aug}) - \nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{c+caug})\| \tag{75}$$

$$= \|(\nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_f) + \nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{aug})) - (\nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_c) + \nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{caug}))\| \tag{76}$$

$$= \|(\nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_f) - \nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_c)) + (\nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{aug}) - \nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{caug}))\| \tag{77}$$

$$\leq \|(\nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_f) - \nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_c))\| + \|(\nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{aug}) - \nabla\mathcal{L}(\mathbf{W}, \mathbf{X}_{caug}))\| \tag{78}$$

$$\leq \xi + \|F\|(\xi + \sqrt{dn}\omega) \tag{79}$$

$$= (\|F\|+1)\xi + \sqrt{d}\|F\|n\omega \tag{80}$$

□

Theorem B.6 (Convergence of linear model). *Let f be a linear model with weights \mathbf{W} and augmentation be represented by the common linear transformation F . Let \mathcal{L}_i be β -smooth, \mathcal{L} be λ -smooth and satisfy the α -PL condition, that is for $\alpha > 0$, $\|\nabla\mathcal{L}(\mathbf{W}, \mathbf{X})\|^2 \geq \alpha\mathcal{L}(\mathbf{W}, \mathbf{X})$ for all weights \mathbf{W} . Let ξ upper-bound the normed difference in gradients between the weighted coreset and full dataset and ω bound $\mathbf{W}^T(F\mathbf{x}_i) = \mathbf{W}^T\mathbf{x}_i + z_i$, $\|z_i\| \leq \omega \forall i$. Let G'_0 be the gradient over the full dataset and its augmentations at initialization. Then, running SGD on the size k coreset with its augmentation using constant step size $\eta = \frac{\alpha}{\lambda\beta}$, we get the following convergence bound:*

$$\mathbb{E}[\|\nabla\mathcal{L}(\mathbf{W}^t, \mathbf{X}_{c+caug})\|] \leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} \left(G'_0 + (\|F\|+1)\xi + \sqrt{d}\|F\|n\omega\right)$$

Proof. From [2], we have

$$\mathbb{E}[\|\nabla\mathcal{L}(\mathbf{W}^t, \mathbf{X}_{c+caug})\|^2] \leq \left(1 - \frac{\alpha\eta}{2}\right)^t \mathcal{L}(\mathbf{W}^0, \mathbf{X}_{c+caug}) \tag{81}$$

$$\leq \frac{1}{\alpha} \left(1 - \frac{\alpha\eta}{2}\right)^t \|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_{c+caug})\|^2 \tag{82}$$

$$\tag{83}$$

Using Jensen's inequality, we have

$$\mathbb{E}[\|\nabla\mathcal{L}(\mathbf{W}^t, \mathbf{X}_{c+caug})\|] \tag{84}$$

$$\leq \sqrt{\mathbb{E}[\|\nabla\mathcal{L}(\mathbf{W}^t, \mathbf{X}_{c+caug})\|^2]} \tag{85}$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} \|\nabla\mathcal{L}(\mathbf{W}^0, \mathbf{X}_{c+caug})\| \tag{86}$$

$$\leq \frac{1}{\sqrt{\alpha}} \left(1 - \frac{\alpha\eta}{2}\right)^{\frac{t}{2}} \left(G'_0 + (\|F\|+1)\xi + \sqrt{d}\|F\|n\omega\right) \tag{87}$$

where the last inequality follows from applying Corollary B.5. □

C Singular spectrum analysis

C.1 Experiment details

We generate singular spectrum plots for both MNIST and CIFAR10 datasets in Figure 1. Due to the computational infeasibility of computing the network Jacobian for the full datasets in deep network settings, we instead construct and use a reduced version of these datasets by uniformly select 900 images from the first 3 classes. For our experiments on MNIST, we pretrain a MLP model with 1

hidden layer for 15 epochs. For our experiments on CIFAR10, we pretrain a ResNet20 model for 15 epochs. We then compute the singular spectrums for augmented and non-augmented data based on these pretrained networks.

Since it is difficult to perform a one-to-one matching of singular values produced from augmented and non-augmented datasets, we instead bin our singular values into 30 separate and uniformly distributed bins each containing the same number of singular values. To measure perturbation to singular values resulted from augmentation, we compute the mean difference between each bin. On the other hand, to measure perturbation to singular vectors, we compute mean subspace angle between the singular subspace spanned by singular vectors in each bin.

C.2 Real-world strong augmentations

We study the effects of real-world, unbounded augmentations on the singular spectrum of the network Jacobian. In particular, in addition to the plots in the main paper, we show the effect of strong augmentations through (1) random rotation (up to 30° and AutoAugment [3] for MNIST and (2) random horizontal flips/random crops and AutoAugment for CIFAR10. The policies implemented by AutoAugment include translations, shearing, as well as contrast and brightness transforms. We study the effects of these augmentations on the singular spectrum in Figure 1. Despite these augmentations being unbounded transformations, we note that the results of our theory still holds. In particular, it can be observed that data augmentation increases smaller singular values relatively more with a higher probability. On the other hand, data augmentation affects the prominent singular vectors of the Jacobian to a smaller extent, and preserves the prominent directions. As such, our argument empirically extends to real-world, unbounded label-invariant transformations characteristic of strong augmentations.

D Experiment Setup and Additional Experiments

D.1 Experiment setup

For all experiments, we train using SGD with 0.9 momentum and learning rate decay. For experiments on CIFAR10 and variants/ResNet20, we train for 200 epochs, for Caltech256 (ImageNet pretrained)/ResNet18, we trained for 40 epochs starting at learning rate 0.001 and batch size 64. We also report results for Caltech256 without ImageNet pretraining in Sec. D.8, where we train for 400 epochs to ensure convergence with a starting learning rate of 0.05 and batch size 64. For experiments on ImageNet/ResNet50 and TinyImageNet/ResNet50, we use the standard 90 epoch learning schedule starting at learning rate of 0.1 and batch size 64.

Data and augmentation. We apply our method to training ResNet20 and Wide-ResNet-28-10 on CIFAR10, and ResNet32 on CIFAR10-IMB (Long-Tailed CIFAR10 with Imbalance factor of 100 following [5]) and SVHN datasets. We train Caltech256 [4] on ImageNet-pretrained ResNet18, and include experiments with random initialization in Appendix D. TinyImageNet and ImageNet are trained on with ResNet50. We use [7] for CIFAR10/SVHN, and AutoAugment [3] for Caltech256, TinyImageNet, and ImageNet as the strong augmentation method. Note that we append strong augmentations rather than apply them in-place, which we show to be more effective in Appendix D. All results are averaged over 5 runs using an Nvidia A40 GPU.

D.2 Experiment setup

For all experiments, we train using SGD with 0.9 momentum and learning rate decay. We also set weight decay as For experiments on CIFAR10 and variants/ResNet20, we train for 200 epochs, for Caltech256 (ImageNet pretrained)/ResNet18, we trained for 40 epochs starting at learning rate 0.001 and batch size 64. We also report results for Caltech256 without ImageNet pretraining in Sec. D.8, where we train for 400 epochs to ensure convergence with a starting learning rate of 0.05 and batch size 64. For experiments on ImageNet/ResNet50 and TinyImageNet/ResNet50, we use the standard 90 epoch learning schedule starting at learning rate of 0.1 and batch size 64.

Data and augmentation. We apply our method to training ResNet20 and Wide-ResNet-28-10 on CIFAR10, and ResNet32 on CIFAR10-IMB (Long-Tailed CIFAR10 with Imbalance factor of 100 following [5]) and SVHN datasets. We train Caltech256 [4] on ImageNet-pretrained ResNet18, and

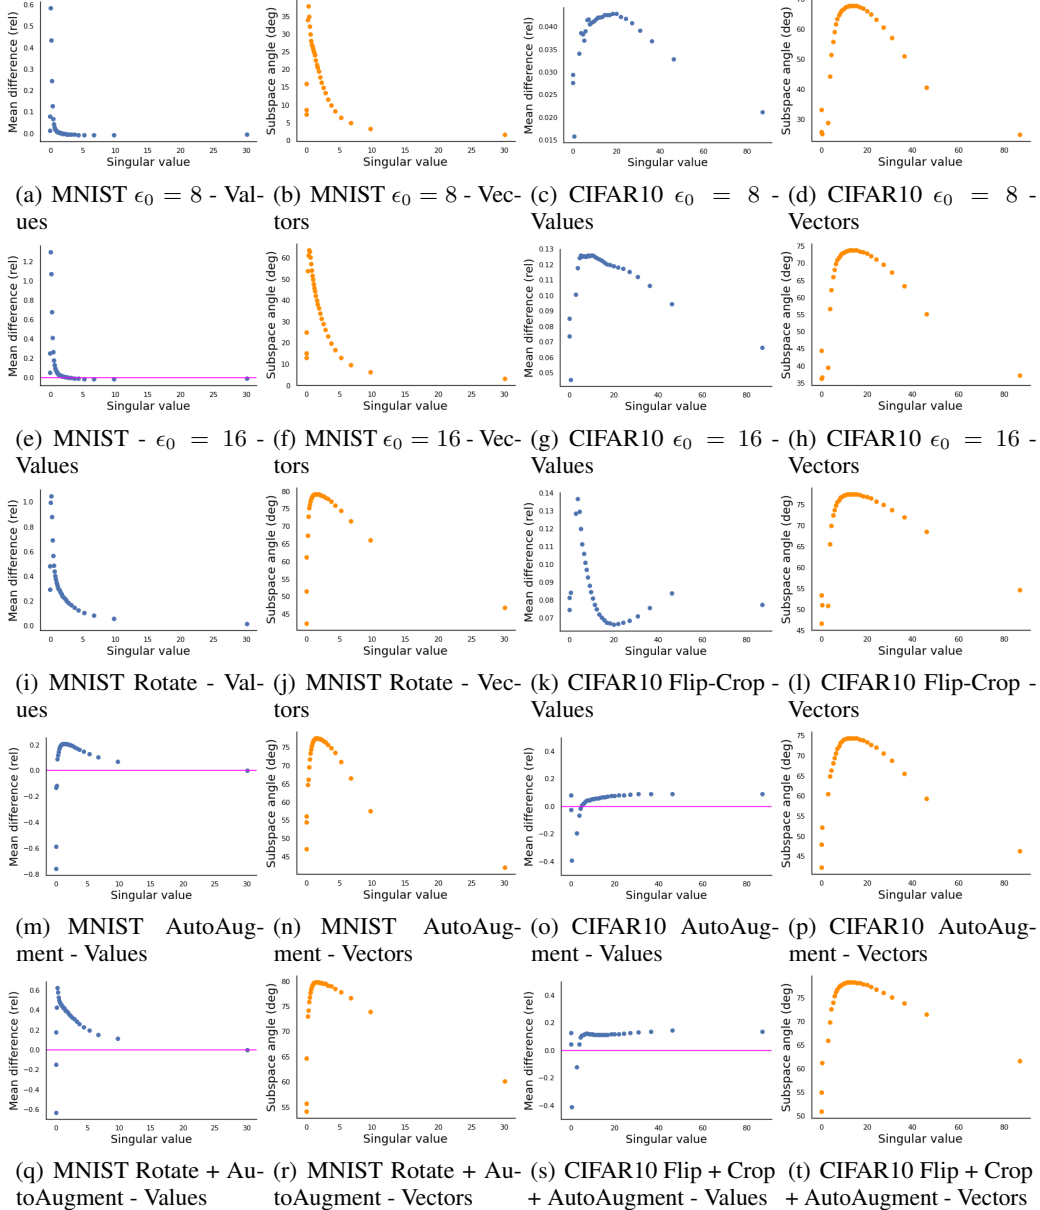


Figure 1: Difference in mean singular values (Cols 1 & 3) between augmented and non-augmented data and mean angular difference (Cols 2 & 4) between subspaces spanned by singular vectors for augmented and non-augmented data.

include experiments with random initialization in Appendix D. TinyImageNet and ImageNet are trained on with ResNet50. We use [7] for CIFAR10/SVHN, and AutoAugment [3] for Caltech256, TinyImageNet, and ImageNet as the strong augmentation method. Note that we append strong augmentations rather than apply them in-place, which we show to be more effective in Appendix D. All results are averaged over 5 runs using an Nvidia A40 GPU.

D.3 Full Results for Table 4

This section contains full experiment results including standard deviations and the full augmentation benchmark for Table 4. Augmenting coresets of size 10% achieves 51% of the improvement obtained from augmentation of the full data and further enjoys a 6x speedup in training time on CIFAR10.

This speedup becomes more significant when using strong augmentation techniques which are mostly computationally demanding, especially when applied to the entire dataset.

Table 1: Supplementary table for Table 4 - Test accuracy on CIFAR10 + ResNet20, SVHN + ResNet32, CIFAR10-Imbalanced + ResNet32 including standard deviation errors and full dataset augmentation accuracy.

Method	Size	CIFAR10	CIFAR10-IMB	SVHN
None	0%	89.46 ± 0.17%	87.08 ± 0.50%	95.676 ± 0.108%
Random	5%	90.34 ± 0.18%	88.48 ± 0.25%	95.760 ± 0.084%
	10%	91.07 ± 0.13%	89.52 ± 0.15%	96.187 ± 0.112%
	30%	92.11 ± 0.12%	91.11 ± 0.18%	96.569 ± 0.073%
Max-Loss	5%	90.79 ± 0.19%	88.77 ± 0.35%	96.165 ± 0.108%
	10%	91.39 ± 0.08%	89.22 ± 0.48%	96.370 ± 0.076%
	30%	92.43 ± 0.07%	91.11 ± 0.25%	96.735 ± 0.068%
Coreset	5%	90.87 ± 0.05%	89.10 ± 0.41%	96.121 ± 0.055%
	10%	91.54 ± 0.19%	89.75 ± 0.52%	96.354 ± 0.091%
	30%	92.49 ± 0.15%	91.12 ± 0.26%	96.791 ± 0.051%
All	100%	93.50 ± 0.25%	92.48 ± 0.34%	97.068 ± 0.030%

D.4 Supplementary results for Table 1

Table 2: Supplementary results for Tab. 1. Training ResNet20 (R20) and WideResnet-28-10 (W2810) on CIFAR10 (C10) using small subsets, and ResNet18 (R18) on Caltech256 (Cal).

Model/Dataset	Subset	Random		Ours	
		Weak Aug.	Strong Aug.	Weak Aug.	Strong Aug.
C10/R20	0.1% (5)	31.7 ± 3.2	33.5 ± 2.7	29.6 ± 3.8	37.8 ± 4.5
	0.2% (10)	35.9 ± 2.1	42.7 ± 3.9	33.6 ± 3.2	45.1 ± 2.3
	0.5% (25)	51.1 ± 2.3	58.7 ± 1.3	55.8 ± 3.1	63.9 ± 2.1
	1% (50)	66.2 ± 1.0	74.4 ± 0.8	65.9 ± 4.0	74.7 ± 1.1
C10/W2810	1% (50)	61.3 ± 2.4	57.7 ± 0.8	59.9 ± 2.4	62.1 ± 3.1
Cal/R18	5% (3)	24.8 ± 1.5	41.5 ± 0.5	33.8 ± 1.7	52.7 ± 1.2
	10% (6)	49.5 ± 0.6	61.8 ± 0.8	55.7 ± 0.3	65.4 ± 0.3
	20% (12)	66.6 ± 0.2	72.5 ± 0.1	67.5 ± 0.3	73.1 ± 0.1
	30% (18)	72.0 ± 0.1	75.7 ± 0.2	71.9 ± 0.2	76.3 ± 0.2
	40% (24)	74.6 ± 0.3	77.6 ± 0.4	74.2 ± 0.4	77.7 ± 0.5
	50% (30)	76.1 ± 0.5	78.5 ± 0.3	76.1 ± 0.1	78.9 ± 0.2

D.5 Training dynamics vs generalization

Figure 2 demonstrates the relationship between training loss and validation accuracy resulted from data augmentation. While training loss of augmented datasets do not decrease as quickly as non-augmented datasets, generalization performance (shown by val. acc.) improves.




Figure 2: Training loss vs validation accuracy of CIFAR10/ResNet20 using AutoAugment.

D.6 Augmentations applied through appending vs in-place

Our experiments on Caltech256/ResNet18/AutoAugment ($R=5$) show that even for cheaper strong augmentation methods (AutoAugment), while in-place augmentation may decrease the performance, appending Random (R) and Coresets (C) augmentations (Append) outperforms in-place augmentation of 2x data points (In-place 2x) for various subset sizes.

Table 3: Caltech256/AutoAugment in-place vs. appending for Caltech256.

	No Aug.	In-place	In-place (2x)	Append
C5%	33.8%	26.4%	48.2%	52.7%
R5%	24.8%	17.4%	40.2%	41.5%
C10%	55.7%	48.2%	62.8%	65.4%
R10%	50.6%	40.2%	62.0%	61.8%
C30%	71.9%	68.8%	74.9%	76.3%
R30%	72.0%	68.7%	75.1%	75.7%

D.7 Speed-up measurements

We measure the improvement in training time in the case of training on full data and augmenting subsets of various sizes. While our method yields similar or slightly lower speed-up to the max-loss policy and random approach respectively, our resulting accuracy outperforms these two approaches. We show this in Fig. D.7. For example, for SVHN/Resnet32 using 30% coresets, we sacrifice 11% of the speed-up to obtain an additional 24.8% of the gain in accuracy from full data augmentation when compared to a random subset of the same size. We show the speed-up obtained for our method and various subset sizes in Tab. 4, and provide wall-clock times for our method in Tab. D.7.

Table 4: Speedup on CIFAR10 + ResNet20 (C10/R20), SVHN + ResNet32 (SVHN/R32).

Dataset	Full Aug.	Ours						Max loss.	Random.
		5%	10%	15%	20%	25%	30%		
C10 / R20	1x	7.93x	6.31x	4.46x	4.27x	3.41x	3.43x	3.48x	4.03x
SVHN / R32	1x	5.35x	3.93x	3.40x	2.80x	2.49x	2.18x	2.21x	2.43x




Figure 3: Speedup/Accuracy of augmenting 30% coresets compared to original max-loss policy for (a) ResNet20 trained on CIFAR10 and (b) ResNet32 trained on SVHN.

Table 5: Wall-clock times to find various sized coresets from *all classes* of Caltech256 and TinyImageNet at 1 epoch. Note, training ResNet20/CIFAR10 with [7] takes 14.4 hrs. In practice, coresets can be found in parallel (p threads) from different classes, and selection happens every $R=5-15$ epochs. Hence, the numbers divide by $p \times R$.

Caltech256			TinyImageNet		
10%	30%	50%	10%	30%	50%
10.50s	10.52s	10.53s	7.85s	8.09s	8.17s

D.8 End to end training on Caltech256

As Caltech256 contains many classes and higher resolution images, training on smaller subset without pretraining has a low accuracy. Thus, many works (e.g. Achille et al., 2020) finetune from ImageNet pretrained initialization. However, we show that our results still hold even when training from scratch. We demonstrate our results in Tab. 6, where we train Caltech256 on ResNet50 without pretraining for 400 epochs, and with $R = 40$, where our method consistently outperforms random subsets for multiple subset sizes (5%, 10%, 30%, 50%).

Table 6: Caltech256 (w/o pretraining) /ResNet50, 400 epochs, $R = 40$

Random				Ours			
5%	10%	30%	50%	5%	10%	30%	50%
17.26	35.38	58.2	64.67	20.58	38.20	60.30	65.17

D.9 Training on full data and augmenting small subsets re-selected every epoch

We apply our proposed method to select a new subset for augmentation every epoch (i.e. using $R = 1$) and compare our results with other approaches using accuracy and percentage of data not selected (NS). We see that while the max-loss policy selects a small fraction of data points over and over and random uniformly selects all the data points, our approach successfully finds the smallest subset of data points that are the most crucial for data augmentation. Hence, it can achieve a superior accuracy than max-loss policy, while augmenting only slightly more examples. This confirms the data-efficiency of our approach. This is especially evident when using coresets of size 0.2%. Furthermore, despite the random baseline using a significantly larger percentage of data, it is outperformed by our approach in both data-efficiency and accuracy. We emphasize that results in this table is different from that of Table 1, as default augmentations on the full training data are performed once every $R = 1$ epochs instead of every $R = 20$ epochs. Since selecting subsets at every epoch can be computationally expensive, we only perform these experiments on small coresets and hence still enjoy good speedups compared to full data augmentation. This shows that our approach is still effective at very small subset sizes, hence can be computationally efficient even when subsets are re-selected every epoch.

Table 7: Training on full data and selecting a new subset for augmentation every epoch ($R = 1$).

Subset	Random		Max-loss Policy		Ours	
	Acc	NS (%)	Acc	NS (%)	Acc	NS (%)
0%	91.96 ± 0.12	—	91.96 ± 0.12	—	91.96 ± 0.12	—
0.2%	92.22 ± 0.22	67.03 ± 0.04	91.94 ± 0.12	86.70 ± 0.15	92.26 ± 0.13	79.19 ± 1.10
0.5%	92.06 ± 0.17	36.70 ± 0.18	92.20 ± 0.13	76.80 ± 0.31	92.27 ± 0.08	63.23 ± 0.35

D.10 Additional visualizations for training on coresets and its augmentations - Measuring training dynamics over time

We include additional visualizations in Figure 4 for training on coresets and its augmentations as supplementary plots to Figure 7(c) and Table 1. We plot metrics obtained during each point (epoch)

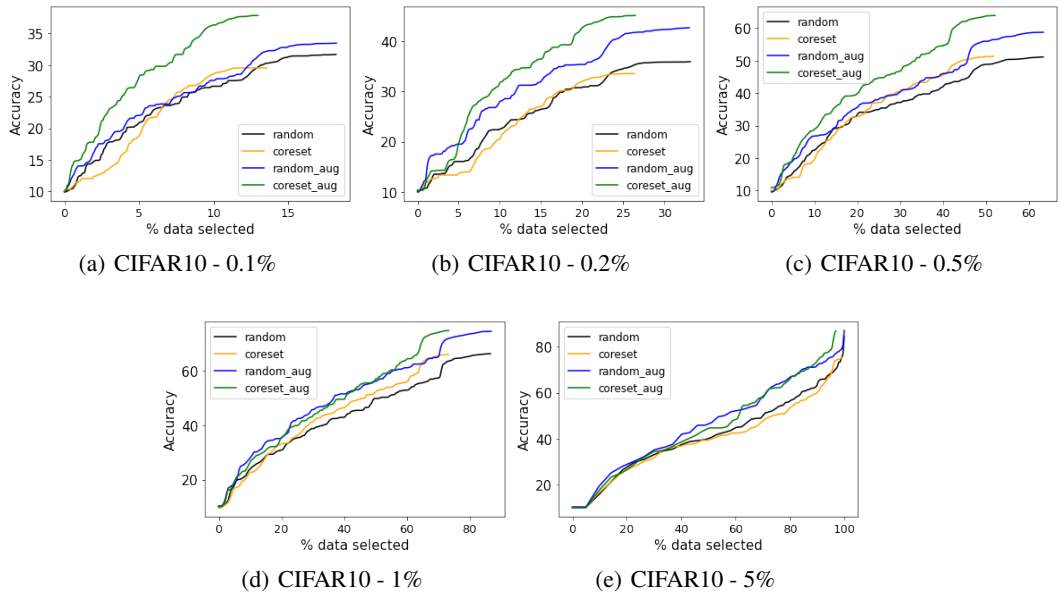


Figure 4: Supplementary plots for Figure 7(c): Training on coresets and its augmentation compared to random baseline, measured using test accuracy against percentage of data used on CIFAR10 dataset across various subset sizes. Accuracy and percentage of data used are measured at every epoch and averaged over 5 runs.

of the training process based on percentage of data selected/used and test accuracy achieved. All metrics are averaged over 5 runs and obtained using $R = 1$. These plots demonstrate that coresets augmentation approaches outperform random augmentation baselines throughout the training process. Furthermore, they show that augmentation of coresets result in a larger increase in test accuracy compared to augmentation of randomly selected training examples, especially for small subset sizes.




Figure 5: Intersection between max-loss and coresets in the top N points selected aggregated across the entire training process. Here, we show the increasing overlap between max-loss and coresets points as N grows.




Figure 6: Qualitative evaluation of coresets and max-loss points.

D.11 Intersection of max-loss policy and coresets

Figure 7(a) depicts the increase in intersection between max-loss subsets and coresets over time. In addition, we also aggregate 30% subsets selected every $R = 20$ epochs using both approaches over the entire training process to compute intersection between the top N selected data points. Our plots

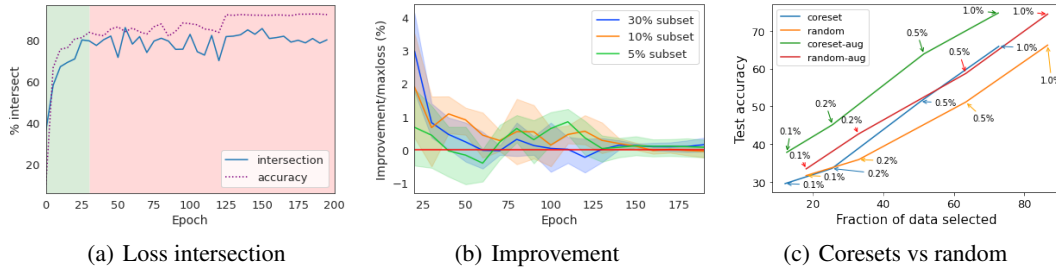


Figure 7: Training ResNet20 on full data and augmented coresets extracted from CIFAR10. (a) Intersection between elements of coresets of size 30% and maximum loss subsets of the same size. The intersection increases after the initial phase of training, (b) Accuracy improvement for training on full data and augmented coresets over training on full data and max-loss augmentation. (c) Accuracy vs. fraction of data selected for augmentation during training Resnet20 on CIFAR10.

in Figure 5 suggest that a similar pattern holds in this setting. We also qualitatively visualize this in Figure 6.

References

- [1] Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR, 2019.
- [2] Raef Bassily, Mikhail Belkin, and Siyuan Ma. On exponential convergence of sgd in non-convex over-parametrized learning. *arXiv preprint arXiv:1811.02564*, 2018.
- [3] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 113–123, 2019.
- [4] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. 2007.
- [5] Byungju Kim and Junmo Kim. Adjusting decision boundary for class imbalanced learning. *IEEE Access*, 8:81674–81685, 2020.
- [6] Baharan Mirzasoleiman, Jeff Bilmes, and Jure Leskovec. Coresets for data-efficient training of machine learning models. In *International Conference on Machine Learning*, pages 6950–6960. PMLR, 2020.
- [7] Sen Wu, Hongyang Zhang, Gregory Valiant, and Christopher Ré. On the generalization effects of linear transformations in data augmentation. In *International Conference on Machine Learning*, pages 10410–10420. PMLR, 2020.