

# THE FAIR LANGUAGE MODEL PARADOX

Anonymous authors

Paper under double-blind review

## ABSTRACT

Large Language Models (LLMs) are widely deployed in real-world applications, yet little is known about their training dynamics at the token level. Evaluation typically relies on aggregated training loss, measured at the batch level, which overlooks subtle per-token biases arising from (i) varying token-level dynamics and (ii) structural biases introduced by hyperparameters. While weight decay is commonly used to stabilize training, we reveal that it silently introduces performance biases detectable only at the token level. In fact, we empirically show across different dataset sizes, model architectures and sizes ranging from 270M to 3B parameters that as weight decay increases, low-frequency tokens are disproportionately depreciated. This is particularly concerning, as these neglected low-frequency tokens represent the vast majority of the token distribution in most languages, calling for novel regularization techniques that ensure fairness across all available tokens.

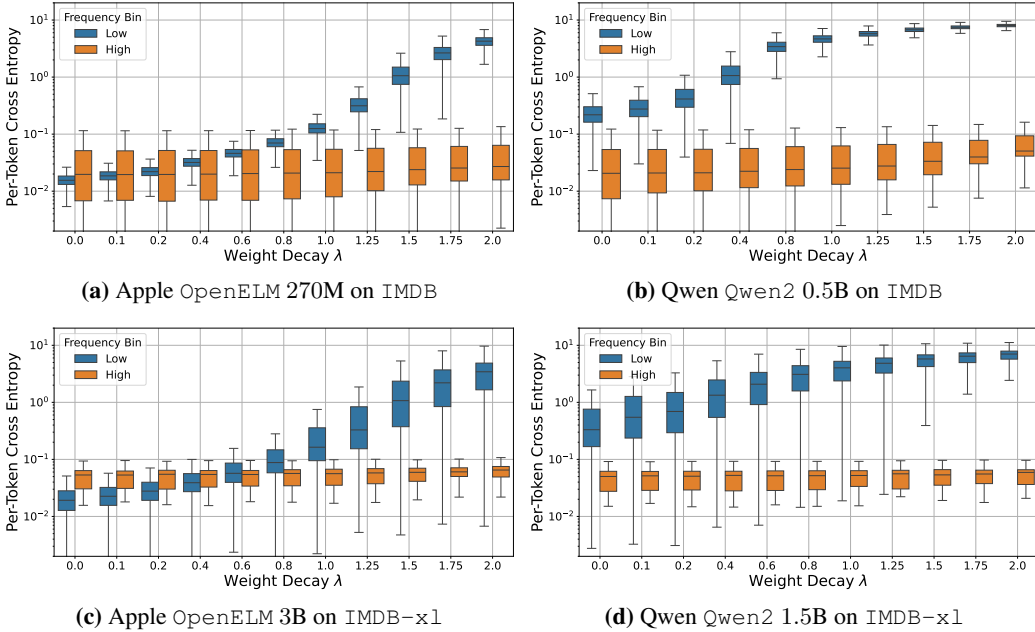


Figure 1: We compare the per-token cross-entropy loss for low (blue) and high (orange) frequency tokens when training different LLM architectures and sizes with varying weight decay  $\lambda \in (0.0, 2.0)$  on the IMDB dataset using a BPE tokenizer with a vocabulary size of 32005. As weight decay increases, the model disproportionately disregards low-frequency tokens, which make up the vast majority of tokens in language datasets. Low-frequency tokens suffer from higher cross-entropy loss, while high-frequency tokens remain largely unaffected. Critically, the degradation of low-frequency token performance happens *silently*, as the average training loss, monitored by practitioners, remains largely unchanged across different levels of weight decay. An example of prompt with segmentation of which tokens are low and high frequency is provided in Figure 2.

i watched this film for 45 minutes and counted 9 mullets . that 's a mullet every 5 minutes . seriously though , this film is living proof that formula works . if it ain 't broke , it don 't need fix in . a streetwise - yet - v ulner able heroine , a hardened ex - cop martial arts master with a heart of gold and a serial killer with ' iss ues ' . pure magic . </s>

Figure 2: Depiction of a training set prompt from IMDB with characters colored by token frequency: low-frequency (blue) and high-frequency (orange). The coloring threshold is 1026 (P99). Tokens appearing fewer than 1026 times in the dataset are blue, otherwise they are colored in orange.

## 1 INTRODUCTION

A major challenge in machine learning is designing algorithms that generalize well from training data. One of the classic methods for promoting generalization (Krogh & Hertz, 1991; Shalev-Shwartz & Ben-David, 2014) is the use of regularization techniques, such as  $L_2$  regularization, to limit model complexity. Empirical evidence from classification problems with balanced classes shows that increasing weight decay, while effectively fitting the data and minimizing training loss, generally improves performance on unseen data. However, a recent study by Balestrieri et al. (2024) demonstrates that when training classifiers like ResNet-50 (He et al., 2016) on computer vision tasks such as ImageNet classification (Russakovsky et al., 2015), higher weight decay leads to undesired behavior, causing the model to neglect certain classes. This class-dependent effect of regularization is further amplified by imbalanced class distributions, as increasing weight decay does not result in a uniform performance decline across all classes. Instead, the model underperforms on low-probability classes while performing significantly better on more prevalent ones.

This class-dependent behavior is not unique to vision tasks. In Natural Language Processing (NLP), a shift from traditional classification settings has led to a lack of attention toward token frequency and the per-token effects of regularization. In modern large language models (LLMs) trained on text data (Brown et al., 2020; Radford et al., 2018; OpenAI et al., 2024; Anthropic, 2023; Ortiz, 2023), the task of predicting the next token from a large vocabulary also results in significant token frequency imbalances. For instance, as shown in Figures 3, in the IMDB dataset (Maas et al., 2011), 95% of the total tokens in the data are captured by the top 0.01% of tokens. This indicates that the vast majority of tokens appear infrequently, while a small set of tokens dominates, creating a substantial imbalance. Additionally, the proportion of low-frequency tokens tends to increase as the vocabulary expands. This raises a critical question:

*Can regularization techniques, such as weight decay, typically used to promote generalization, ensure fairness across tokens when applied to LLMs trained on imbalanced token distributions?*

**Contributions.** In order to study that question, we investigate the influence of weight decay on the token-level prediction performance of large language models, uncovering critical insights that are often overlooked when relying solely on aggregated performance metrics. Our study provides several key contributions:

- The next-token classification task suffers from severe class imbalance. Figures 3 demonstrate that the class distribution follows a heavy-tailed pattern, with the vast majority of classes being low-frequency and only a small portion being high-frequency.
- We trained the Apple OpenELM models with 270M and 3B parameters, as well as Qwen2 models with 0.5B and 1.5B parameters, on both the IMDB dataset and its extended version IMDB-x1, using varying levels of weight decay that yielded acceptable training losses. As observed in Figure 1, the models’ performance on low-frequency tokens significantly degrades as weight decay increases.
- We observe that higher-frequency tokens are consistently learned faster than low-frequency tokens across multiple random seeds, with the gap in learning speed widening as weight decay increases, suggesting that regularization may disproportionately disadvantage rare tokens.

These findings expose a *critical dilemma*. Practitioners often use aggressive weight decay to train LLMs—intended to stabilize training—but unintentionally and silently degrade the model’s performance on low-frequency tokens, which make up the majority of the data. While conventional wisdom

advocates for increased weight decay to promote generalization, in language modeling, this strategy results in the unintended consequence of “neglecting” low-frequency tokens. This introduces harmful biases, favoring more common tokens, often reflecting the language patterns of majority groups. Our results emphasize the need for regularization techniques that explicitly address token imbalances.

## 2 RELATED WORK

**Weight Decay and Generalization.**  $L_2$  regularization was initially introduced to stabilize solutions to ill-posed problems (Tikhonov, 1943) and later adopted to enhance the generalization of neural networks (Krogh & Hertz, 1991). While many studies have linked low-norm solutions to improved generalization (Neyshabur et al., 2015; Golowich et al., 2020; Bartlett & Mendelson, 2001; Bartlett et al., 2017; Arora et al., 2018; Galanti et al., 2023c; Wei & Ma, 2019; Li et al., 2018), the precise relationship between  $L_2$  regularization and generalization remains a topic of debate. Several works (Zhang et al., 2017; Jiang\* et al., 2020) argue that norm-based measures alone are insufficient to fully explain generalization in deep learning. For example, Zhang et al. (2017) found that although weight decay can improve test accuracy, the overall effect is typically modest—around 1 – 2% on ImageNet. Nonetheless, other studies have demonstrated that weight decay helps alleviate (Nakkiran et al., 2021; Pezeshki et al., 2022) the double descent phenomenon (Belkin et al., 2019; Nakkiran et al., 2020) and is critical for achieving Grokking in mathematical reasoning tasks (Power et al., 2022; Varma et al., 2023).

**Weight Decay, Optimization, and Inductive Biases.** Despite its modest impact on generalization, weight decay is widely employed in many state-of-the-art language models, including GPT-3 (Brown et al., 2020), Chinchilla (Hoffmann et al., 2024), and LLaMA (Touvron et al., 2023a;b; Dubey et al., 2024). These models are typically trained using a “one-pass” stochastic gradient descent (SGD) regime, where the optimizer directly minimizes the population error.

As shown in (Andriushchenko et al., 2023), the training and validation losses remain closely aligned across different levels of weight decay. While weight decay’s effect on generalization is limited, it plays a crucial role in improving optimization. For example, both the Chinchilla paper (Hoffmann et al., 2024) (see Figure A7) and (Andriushchenko et al., 2023) (see Figure 4) demonstrate that weight decay in AdamW leads to lower training loss compared to Adam, particularly toward the end of training. Other studies (van Laarhoven, 2017; Zhang et al., 2019; Li & Arora, 2019; Li et al., 2020; Lewkowycz & Gur-Ari, 2020) have shown that weight decay enhances training stability by controlling the “effective learning rate” in scale-invariant neural networks. Additionally, other works (Galanti et al., 2023b; Rangamani & Banburski-Fahey, 2022; Pan & Cao, 2024; Beneventano et al., 2024) reveal that weight decay contributes to various inductive biases, such as rank minimization and neural collapse, which are beneficial for network compression (Denton et al., 2014; Alvarez & Salzmann, 2017; Tukan et al., 2021; Yu et al., 2017) and downstream performance (Galanti et al., 2022; 2023a).

**Training with Imbalanced Classes and Minority Collapse.** Training with imbalanced classes presents a significant challenge in machine learning. Empirical studies consistently show that the weight vectors associated with the more frequent classes tend to have larger norms, which pushes the decision boundary toward the minority classes. As a result, the feature space allocated to less frequent classes shrinks, leading to a notable drop in performance (Kim & Kim, 2020; Kang et al., 2019; Cao et al., 2019; Ye et al., 2020; Liu et al., 2023; Kang et al., 2020; Balestrieri et al., 2022). For instance, (Balestrieri et al., 2022) showed that when training neural networks for classification of visual data, higher levels of weight decay introduce a stronger bias for the model to prioritize higher-probability classes over lower-probability ones.

To gain deeper insights into this issue, several works have investigated this phenomenon from a theoretical perspective. Fang et al. (2021) proposed the unconstrained features model (UFM) as a simplified framework for exploring the geometric properties of the global minima in cross-entropy loss with regularization, particularly in overparameterized neural networks. In the case of a balanced dataset, they demonstrated that neural collapse (NC) occurs at any global minimizer of the loss function combined with regularization. However, in the case of class imbalance, neural networks exhibit distinct geometric patterns, and some of the NC properties no longer hold (Dang et al., 2023; Thrampoulidis et al., 2022; Hong & Ling, 2023; Dang et al., 2024). While last-layer features for samples in the same class still collapse to their respective class means (NC1), the class means and classifier weights no longer form a Simplex Equiangular Tight Frame (ETF), violating NC2 (Fang

et al., 2021). In more extreme cases, when the imbalance becomes severe, the classifier weights for minority classes can collapse onto each other, rendering them indistinguishable from other classes (Fang et al., 2021). This phenomenon, referred to as “Minority Collapse,” explains the sharp decline in accuracy for minority classes in imbalanced settings.

Building on these finding, we explore whether similar behavior arises in next-token prediction tasks, where the problem can be viewed as a (very noisy) classification task, with the next token acting as the “class” for a given sequence.

### 3 EXPERIMENTAL SETTINGS

**Problem Setup.** We focus on the task of next-token prediction in autoregressive language modeling, a self-supervised learning problem central to natural language processing. Given a sequence of tokens  $x_i = (x_{i,1}, \dots, x_{i,n})$ , the objective is to model the conditional probability distribution  $p(x_{i,t}|x_{i,<t})$  for each token position  $t$ . Formally, let  $\mathcal{D} = \{x_i\}_{i=1}^m$  be a text corpus, where each  $x_i = (x_{i,1}, \dots, x_{i,n}) \in \mathcal{S} = \mathcal{V}^n$  is a sequence of  $n$  tokens, and  $x_{i,t} \in \mathcal{V}$  belongs to a fixed vocabulary  $\mathcal{V}$  of size  $V$ . We train different transformer-based models  $f_\theta : \mathcal{S} \rightarrow [V]$  mapping from sequence to logits to minimize the regularized empirical risk:

$$L_{\mathcal{D}}^\lambda(f) = \frac{1}{m} \sum_{i=1}^m \left( \frac{1}{n-1} \sum_{t=1}^{n-1} \ell(f_\theta(x_{i,\leq t}), x_{i,t+1}) \right) + \lambda \|\theta\|_2^2,$$

where  $\ell$  denotes the cross-entropy loss,  $\lambda$  is the weight decay coefficient,  $\|\theta\|_2$  is the  $L_2$  norm of the model parameters,  $x_{i,\leq t}$  the sequence up to position  $t$  in the  $i$ -th sample, and  $x_{i,t+1}$  is the next token to be predicted.

**Model Architecture.** For our experiments, we train multiple models, including the Apple OpenELM models with 270M and 3B parameters (Mehta et al., 2024), as well as the Qwen Qwen2 models with 0.5B and 1.5B parameters (Yang et al., 2024). The small models ( $< 1$ B parameters) were configured with a context length of 128 whereas the large ones with a context length of 64. These architectures’s moderate size allows us to conduct multiple training runs with different  $\lambda$  values, enabling a comprehensive exploration of weight decay’s impact on token-level dynamics across various regularization choices.

**Dataset and Tokenizer.** We trained our models on the IMDB dataset (Maas et al., 2011), a widely used benchmark due to its balanced sentiment-labeled data. The IMDB training split contains 25000 samples. For the scope of this study, we discarded the labels, focusing solely on the raw text data to analyze the impact of hyperparameters on token-level generation performance. Additionally, we created an extended version of the dataset, termed IMDB-x1, by incorporating all the unsupervised samples from IMDB, which increased the training set to a total of 75000 samples. We used Byte Pair Encoding (BPE) (Gage, 1994; Sennrich et al., 2016) as our tokenization method, training a tokenizer on the IMDB dataset’s training set with a target vocabulary size of 32005 tokens. BPE ensures that both frequent and infrequent tokens are well represented, providing a suitable basis for analyzing token-level learning. This choice allows us to examine how different tokens, particularly rare ones, are influenced by various weight decay values throughout the training process.

**Hyperparameters and Training.** We conducted experiments across a range of weight decay values  $\lambda \in (0.0, 2.0)$ . To account for the stochastic nature of training, each configuration was run with the same 5 different random seeds, with results averaged and presented with confidence intervals. The models were trained using the AdamW optimizer (Loshchilov & Hutter, 2019) to decouple the weight decay from gradient updates, with a learning rate of  $5e-5$  and a cosine decay schedule. A warm-up period of 10 training steps was used for stabilization during early training. We trained each LLM for a total of 10000 steps, evaluating token-level metrics every 100 steps to understand how different tokens are learned and represented under varying weight decay values. For small models (less than 1B parameters), we used a batch size of 64 per device on a single NVIDIA A100 32GB GPU, whereas for large models, we used a batch size of 16 with a gradient accumulation step of 4 on a single NVIDIA A100 80GB GPU. Additionally, mixed precision training (fp16) was employed to optimize GPU memory usage, reduce computational load, and accelerate convergence.

**Evaluation Metrics.** Our evaluation used token-level metrics to assess how individual tokens were generated under varying weight decay configurations. Unlike traditional metrics like perplexity, we

focused on whether specific tokens were under-represented during inference, revealing how certain configurations induce bias in token representation, even when aggregate performance appears stable. The *average training loss*, computed with cross-entropy loss over all tokens in a batch, is given by:

$$\mathcal{L}_{\text{avg}} = -\frac{1}{BC} \sum_{b=1}^B \sum_{c=1}^C \sum_{v=1}^V y_{b,c,v} \log p_{b,c,v},$$

where  $y_{b,c,v}$  and  $p_{b,c,v}$  are the ground truth and predicted probabilities, respectively. This method implicitly favors frequent tokens, as they appear more often, leading the model to prioritize them over low-frequency tokens. To counter this imbalance, we also used a *token-balanced training loss*, ensuring each token contributes equally, regardless of frequency. First, we compute the cross-entropy loss for each token  $\ell_{b,c} = -\sum_{v=1}^V y_{b,c,v} \log p_{b,c,v}$ . We then average these losses by token type, yielding the final token-balanced loss:

$$\mathcal{L}_{\text{tok-bal}} = \frac{1}{V} \sum_{v=1}^V \frac{1}{|\{(b,c) : y_{b,c,v} = 1\}|} \sum_{\{(b,c) : y_{b,c,v} = 1\}} \ell_{b,c}.$$

This approach ensures that low-frequency tokens contribute equally to the loss, leading to a more balanced optimization process, though it may challenge the model’s handling of rare tokens, especially under strong regularization. The computation of per-token metrics is summarized in Algorithm 1.

**Per-Token Learning Speed.** To quantify how quickly the model learns individual tokens during training, we introduce the per-token learning speed metric, which measures how rapidly the model minimizes the cross-entropy loss for each token. Specifically, we compute the area under the curve (AUC) of the token’s normalized loss trajectory over time, where a smaller AUC indicates faster learning. For each token, we normalize its cross-entropy losses across all training steps to the range  $[0, 1]$ , ensuring comparability between tokens with different loss scales. Let  $\ell_t$  represent the cross-entropy loss for a token at training step  $t$ , and let  $\mathcal{L} = \{\ell_1, \ell_2, \dots, \ell_T\}$  be the set of losses over  $T$  steps. We define the normalized loss at each step as:  $\tilde{\ell}_t = \frac{\ell_t - \min(\mathcal{L})}{\max(\mathcal{L}) - \min(\mathcal{L})}$ . The area under the normalized loss curve is calculated as:  $\text{AUC}(\mathcal{L}) = \int_0^T \tilde{\ell}_t dt$ . The learning speed  $S$  is then defined as:  $S = 1 - \frac{\text{AUC}(\mathcal{L})}{T}$ . This ensures that  $S$  ranges from 0 to 1, with higher values indicating faster learning. If the range of losses is zero (i.e.,  $\max(\mathcal{L}) = \min(\mathcal{L})$ ), we set  $S$  to zero, as no learning occurs.

## 4 EMPIRICAL RESULTS

### 4.1 TEXTUAL DATA IS HIGHLY IMBALANCED

As a preliminary step in our study, we investigated how imbalanced textual data actually is. For this purpose we conducted several statistical evaluations of the IMDB dataset in order to verify that indeed the vast majority of tokens are of low-frequency and a small minority of the tokens appear very often in the data as predicted by the Zipf law.

**Token Frequency Percentiles.** Figure 3(a) highlights the extreme token frequency imbalance in the IMDB dataset, tokenized using a BPE tokenizer with a vocabulary size of 32005. The vast majority of the token frequency mass is concentrated in a very small portion of high-frequency tokens. Specifically, 95% of the total tokens in the data is captured by the top 0.01% of tokens, which demonstrates the steep distribution of token frequencies, where very few tokens dominate. To illustrate the calculation: suppose we have a dataset with 100 samples, each consisting of 10 tokens. If a set of 1% of the tokens in the vocabulary appear 800 times across those 1000 tokens, we say that 80% of the total tokens in the data is captured by the top 1% of tokens. Figure 3(b) further emphasizes this imbalance, showing how the proportion of low-frequency tokens (those below the 95th percentile in the data) grows as the vocabulary size increases. This suggests that token imbalance is inherent to language and further amplified as the vocabulary size expands with larger tokenizers.

### 4.2 PERFORMANCE IMPAIRMENT IN LOW-FREQUENCY TOKENS

**Average vs. Token-Balanced Training Loss.** To highlight the effect of weight decay on low-frequency tokens, we compared the average training loss with the token-balanced training loss.




Figure 3: Comparison of token frequency distribution and the ratio of low-frequency tokens across varying vocabulary sizes for the IMDB dataset. The left plot shows the token frequency distribution with cumulative frequency thresholds (50%, 80%, and 95%) marked. The right plot illustrates how the ratio of tokens below the 95th percentile increases with vocabulary size, converging to  $\approx 0.85$ .




Figure 4: **Impact of Weight Decay on Cross-Entropy.** Average training loss (blue) and class-balanced loss (orange) increase with weight decay. The class-balanced loss is more sensitive due to its focus on low-frequency tokens.

Figure 5: **Token Learning Speed.** Token learning speed (0 – 1) plotted against frequency (log-scale) for  $\lambda = 1.0$ . Colors represent token groups by frequency bins, highlighting variation across token frequencies.

Figure 4 presents both metrics at the end of training for models trained with varying weight decay values ( $\lambda \in \{0.0, 0.1, 0.3, 0.5, 1.0\}$ ). As shown, the token-balanced loss rises sharply with increasing weight decay, while the average training loss shows only a slight increase as  $\lambda$  moves from 0.0 to 1.0. This discrepancy occurs because weight decay disproportionately affects low-frequency tokens, which are given equal importance in the token-balanced loss. In contrast, the average training loss places much less emphasis on low-frequency tokens, making the effects of weight decay less noticeable. For exact values, including per-token perplexity and training accuracy, see Table 1.

**Per-Token Performance vs. Weight Decay.** Beyond comparing the average and token-balanced training losses, we also investigated how increasing weight decay influences performance on low- and high-frequency tokens separately. In this experiment, we examine how increasing weight decay affects the per-task loss function for tokens of different frequencies. To this end, we trained multiple versions of the same model (e.g., Apple OpenELM 270M) with varying degrees of weight decay on a given dataset (e.g., IMDB) and compared the average loss function for low-frequency and high-frequency tokens. As shown in Figure 1, the performance on high-frequency tokens remains largely unaffected by the increase in weight decay, in contrast to the loss for low-frequency tokens, which increases significantly with higher weight decay. Here, low-frequency tokens are those that appear between  $3^0$  and  $3^1$  times as the next token in the data, while the highest-frequency tokens appear between  $3^9$  and  $3^{10}$  times as the next token.

**Per-Token Performance vs. Frequency.** As a next step, we would like to visualize the distribution of per-token losses and accuracies across tokens of varying frequencies. In Figure 6, tokens are grouped into bins, where the  $i$ th bin contains tokens with frequencies between  $3^{i-1}$  and  $3^i$ . For each token, we plotted in Figure 6 (a) its average per-token cross-entropy and in Figure 6 (b) the per-token accuracy, averaging over all sequences where that token is predicted and over all random seeds. As shown, model performance improves significantly and consistently as token frequency increases, with high-frequency tokens benefiting the most. In Figure 5, we replicated this analysis for per-token learning speed, which also monotonically improves as token frequency increases, following the same pattern observed for loss and accuracy.

Weight Decay $\lambda$	0.0	0.1	0.3	0.5	1.0
Training Loss	0.051 $\pm$ 0.000	0.052 $\pm$ 0.000	0.054 $\pm$ 0.000	0.056 $\pm$ 0.000	0.068 $\pm$ 0.001
Per-Token Loss	0.066 $\pm$ 0.001	0.070 $\pm$ 0.001	0.078 $\pm$ 0.001	0.090 $\pm$ 0.002	0.163 $\pm$ 0.004
Per-Token PPL	1.069 $\pm$ 0.003	1.073 $\pm$ 0.003	1.081 $\pm$ 0.003	1.095 $\pm$ 0.003	1.177 $\pm$ 0.003
Per-Token Accuracy (%)	98.798 $\pm$ 0.031	98.781 $\pm$ 0.029	98.778 $\pm$ 0.032	98.759 $\pm$ 0.028	98.714 $\pm$ 0.029

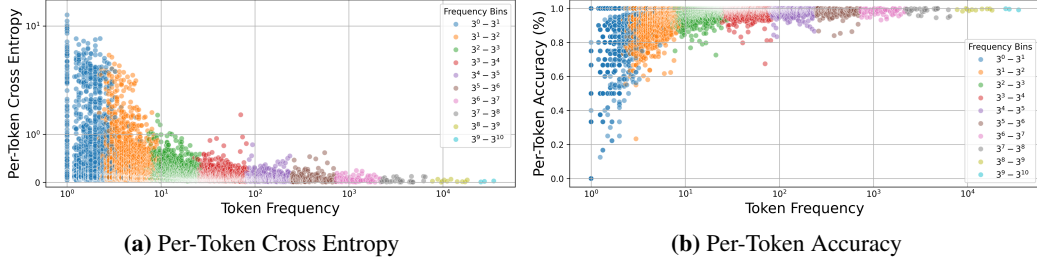
Table 1: Impact of weight decay ( $\lambda$ ) on model performance metrics (mean $\pm$ std).

Figure 6: **(a)** The relationship between token frequency (log-scale) and per-token cross entropy (symlog-scale) across frequency bins, representing powers of 3. Lower-frequency tokens exhibit significantly higher cross entropy, indicating weaker learning. **(b)** The relationship between token frequency and accuracy across frequency bins. Higher-frequency tokens achieve higher accuracy, while lower-frequency tokens demonstrate more variability in performance. Both plots are  $\lambda = 1.0$ .

**High-Frequency Tokens are Learned Faster.** In Figure 7 (b), tokens are divided into bins, selected in the same way as before, to compare the average token learning speed (as defined earlier) and its standard deviation for two models: one trained without weight decay ( $\lambda = 0$ , blue) and the other with weight decay ( $\lambda = 1$ , orange). The percentage of tokens in each bin is also shown, with the majority concentrated in the low-frequency bins. As observed, less frequent tokens are learned more slowly by both models. Notably, the gap in learning speed between the models widens for lower-frequency tokens but diminishes for higher-frequency tokens. This indicates that increasing weight decay disproportionately deprioritizes low-frequency tokens.

## 5 THEORETICAL DISCUSSION

It is well known that the class frequency is positively correlated with the norm of the top layer classifier of the given class (Kang et al., 2019; Huang et al., 2016; Kim & Kim, 2020). For instance, Dang et al. (2024) considered a variant of unconstrained features model (UFM) (Fang et al., 2021), in which the features are constrained to be non-negative, motivated by the fact that features are usually the output of ReLU activations in many common architectures. We use their framework to analyze the influence of the token frequency and the weight decay on the per-token loss function. Formally, suppose we have a set of possible tokens  $\mathcal{V}$  (one-hot encodings of the numbers in  $[V]$ ) and a dataset  $\mathcal{D} = \bigcup_{k=1}^V \{x_{k,i}\}_{i=1}^{n_k}$  of sequences  $x_{k,i} = (x_{k,i,1}, \dots, x_{k,i,n})$  of tokens with the next token being  $k \in [V]$  (whose one-hot encoding is  $y_k$ ). Within the model proposed in (Dang et al., 2024), for each sequence  $x_{k,i}$  we learn an unconstrained feature representation  $h_{k,i} \in \mathbb{R}^d$  together a linear layer  $W \in \mathbb{R}^{V \times d}$  using the Cross-Entropy loss:

$$\begin{aligned} \min_{W, H} \quad & \frac{1}{N} \sum_{k=1}^V \sum_{i=1}^{n_k} \ell_{\text{CE}}(Wh_{k,i}, y_k) + \frac{\lambda_W}{2} \|W\|_F^2 + \lambda_H \|H\|_F^2, \\ \text{s.t.} \quad & H \geq 0, \lambda_W > 0, \lambda_H > 0, \end{aligned} \quad (1)$$

where  $\ell_{\text{CE}}(z, y_k) = -\log\left(\frac{\exp(z_k)}{\sum_{i=1}^V \exp(z_i)}\right)$ ,  $H := [h_{1,1}, \dots, h_{1,n_1}, h_{2,1}, \dots, h_{V,n_V}] \in \mathbb{R}^{d \times N}$  (where  $N = \sum_{k=1}^V n_k$ ) are the learned representations for each sample  $(k, i)$  and  $H \geq 0$  denotes entry-wise non-negativity. In addition,  $W = [w_1, w_2, \dots, w_V]^T \in \mathbb{R}^{V \times d}$  to be the last-layer weight matrix, with  $w_k \in \mathbb{R}^d$  being the  $k$ -th row of  $W$ .



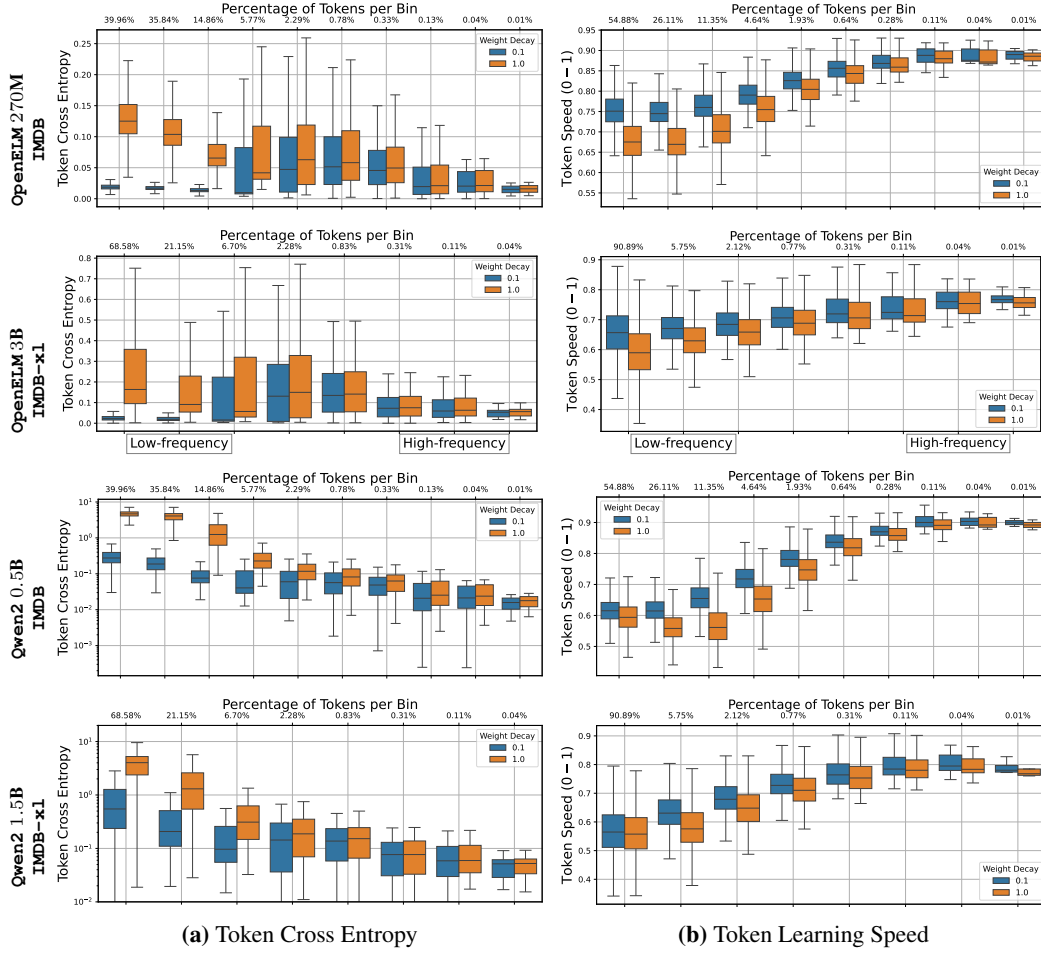


Figure 7: We compare the (a) per-token cross-entropy loss and (b) learning speed across token frequencies when training an LLM with weight decay  $\lambda = 0.1$  and  $\lambda = 1.0$ , on the IMDB and IMDB-x1 datasets using a BPE tokenizer with a vocabulary size of 32005. As weight decay increases, the model disproportionately disregards low-frequency tokens, which make up the vast majority of tokens in language datasets. Low-frequency tokens suffer from higher cross-entropy loss and reduced learning speed, while high-frequency tokens remain largely unaffected. Critically, the degradation of low-frequency token performance happens *silently*, as the average training loss, monitored by practitioners, remains largely unchanged across different levels of weight decay.

While the above formulation does not exactly match the practice (since  $H$  is a matrix of free parameters instead of the outputs of a neural network), this abstraction can help understand how the per-token frequency  $n_k$  and the regularization parameters  $\lambda_W$  and  $\lambda_H$  influence the per-token loss and classifier. For instance, according to Theorem 4.1 and Proposition 4.3 in (Dang et al., 2024), if  $d \geq V$  then any global minimum of Eq. 1 satisfies:

(a) Within-class feature collapse:  $\forall k \in [V], i \neq j \in [n_k] : h_{k,i} = h_{k,j} = \mu_k$ .

(b) Class-mean orthogonality:  $\forall k \neq l : \mu_k^\top \mu_l = 0$ .

(c) Class-mean norm:  $\|\mu_k\|^2 = \sqrt{\frac{\lambda_W(V-1)}{\lambda_H V n_k}} M_k$ .

(d) Weight norms:  $\|w_k\|^2 = \sqrt{\frac{\lambda_H}{\lambda_W V^3(V-1)}} \left( (V-1)^2 \sqrt{n_k} M_k + \sum_{j=1}^V \sqrt{n_j} M_j \right)$ .



Here,  $M_k = \left[ \log \left( (V-1) \left( \frac{\sqrt{n_k}}{N\sqrt{\frac{V-1}{V}} \lambda_W \lambda_H} - 1 \right) \right) \right]_{\diamond}$ , where the function  $[x]_{\diamond}$  returns  $x$  if  $x$  is defined and is positive and 0 otherwise.

The above series of observations imply that  $\mu_k = 0$  if and only if  $M_k = 0$ . As a result (summarized in Corollary 4.6 in (Dang et al., 2024)), if  $n_k \leq \lambda_W \lambda_H N^2 \frac{V-1}{V}$  (which implies  $M_k = 0$ ) the model avoids learning the  $k$ th token. In particular, **the number of tokens that are neglected by the model increased when increasing the level of weight decay.**

**Proposition 5.1.** Suppose  $d \geq V$ , then any global minimizer  $(W, H)$  of the problem obeys  $\ell_{\text{CE}}(Wh_{k,i}, y_k) = \log \left( \sum_{j=1}^V \exp \left( \frac{M_j}{V^2} \right) \right) - M_k$ .

We observe that the loss function  $\ell_{ij} = \ell_{\text{CE}}(Wh_{k,i}, y_k)$  can be decomposed into two components: the first part,  $\ell'$ , is independent of both  $k$  and  $i$ , while the second part depends on  $-M_k$ . Consequently, tokens  $k$  associated with larger values of  $M_k$  incur a smaller per-token loss. Since  $M_k$  increases with  $n_k$ , it follows that **the per-token loss is smaller for tokens of higher frequency.** This aligns with the results in Fig. 7 (a), where the loss function decreases monotonically for higher-frequency tokens when training with weight decay.

Now, suppose we set  $\lambda_W = \lambda_H = \lambda$ . The derivative of the per-token loss  $\ell_{k,i}$  with respect to  $\lambda$  is given by:  $\frac{\partial \ell_{k,i}}{\partial \lambda} = \frac{\partial \ell'}{\partial \lambda} + \left( \frac{1}{\lambda} + \frac{N\sqrt{V-1}}{\sqrt{n_k V - \lambda N \sqrt{V-1}}} \right) \cdot \mathbb{I} \left[ n_k > \lambda^2 N^2 \frac{V-1}{V} \right]$ . We note that the first term of the derivative,  $\frac{\partial \ell'}{\partial \lambda}$ , is independent of  $n_k$ , while the second term is monotonically decreasing with respect to  $n_k$ , provided that  $n_k > \lambda^2 N^2 \frac{V-1}{V}$  ( $k$  is a non-neglected token). Therefore, the derivative of the loss for non-neglected tokens  $k$  is higher for smaller values of  $n_k$ . In particular, **the loss for low-frequency tokens grows at a faster rate compared to high-frequency tokens when increasing  $\lambda$ .** This can be observed in Fig. 7 (a) and Fig. 4, where, although the losses generally increase across all types of tokens as a function of  $\lambda$ , they increase more significantly for low-frequency tokens.

As a final note, although one might think that the observations made in Sec. 4 could be due to poor training, these theoretical results emphasize that they are actually caused by a fundamental issue when training next-token predictors. The results above apply to the global minima of the objective, indicating that they pertain to situations where the training was in fact optimal.

## 6 CONCLUSION

We investigated the impact of weight decay on token-level learning dynamics in large language models. Our findings reveal critical insights into how weight decay affects the learning process of individual tokens—effects that are hidden when relying solely on aggregated metrics. We demonstrated that increasing weight decay disproportionately harms the performance of low-frequency tokens, even when the overall average loss remains largely unchanged. Additionally, we observed that higher-frequency tokens are generally learned faster than their low-frequency counterparts. This interplay between token frequency, performance, and regularization highlights the nuanced effects of training techniques on different parts of the model’s vocabulary.

These results expose a significant oversight in current LLM training practices. Weight decay is commonly employed to reduce overfitting and enhance optimization. While this seems beneficial at first—due to improved convergence and stability in overall loss metrics—our analysis uncovers a hidden pitfall: weight decay can severely compromise the model’s ability to handle low-frequency tokens. Crucially, this degradation goes unnoticed when only aggregated loss metrics are considered. This discrepancy between aggregated performance and token-specific learning underscores the need for fine-grained, token-level evaluations. Without such assessments, models risk sacrificing performance on rare or specialized vocabulary, potentially limiting their effectiveness in domains that require precise handling of low-frequency terms.

As illustrated in Figure 3, the imbalance becomes more pronounced as vocabulary sizes increase. This is particularly concerning as vocabulary sizes are continuously expanded to improve model performance and broaden capabilities (Takase et al., 2024; Tao et al., 2024; Toraman et al., 2023). For example, while LLaMA-v1 (Touvron et al., 2023a) and LLaMA-v2 (Touvron et al., 2023b) used a vocabulary of 32000 tokens, LLaMA-v3 (Dubey et al., 2024) expanded this to 128256 tokens,

Qwen2 (Yang et al., 2024) further extended it to 151936 tokens, and Gemma-2 (Team et al., 2024) increased the size to 256128 tokens.

**Broader Impact.** Our work enhances understanding of how weight decay affects token-level performance in LLMs while remaining undetected by aggregated metrics. This is crucial for ensuring fairness and reliability in real-world applications. While our societal impact is indirect, it underscores the need for more granular evaluation metrics to detect and mitigate potential biases on low-frequency tokens. These findings contribute to the development of more equitable and robust AI systems.

## REFERENCES

- Jose M. Alvarez and Mathieu Salzmann. Compression-aware training of deep networks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pp. 856–867, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Maksym Andriushchenko, Francesco D’Angelo, Aditya Varre, and Nicolas Flammarion. Why do we need weight decay in modern deep learning?, 2023. URL <https://arxiv.org/abs/2310.04415>.
- Anthropic. Introducing claude, 2023. [Online]. Available: <https://www.anthropic.com/news/introducing-claude>. [Accessed: May 2, 2024].
- Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 254–263. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/arora18b.html>.
- Randall Balestriero, Leon Bottou, and Yann LeCun. The effects of regularization and data augmentation are class dependent. *Advances in Neural Information Processing Systems*, 35:37878–37891, 2022.
- Randall Balestriero, Leon Bottou, and Yann LeCun. The effects of regularization and data augmentation are class dependent. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS ’22*, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9781713871088.
- Peter L. Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. In *J. Mach. Learn. Res.*, 2001.
- Peter L. Bartlett, Dylan J. Foster, and Matus Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates Inc., 2017.
- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, July 2019. ISSN 1091-6490. doi: 10.1073/pnas.1903070116. URL <http://dx.doi.org/10.1073/pnas.1903070116>.
- Pierfrancesco Beneventano, Andrea Pinto, and Tomaso Poggio. How neural networks learn the support is an implicit regularization effect of sgd, 2024. URL <https://arxiv.org/abs/2406.11110>.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf).

- Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss, 2019. URL <https://arxiv.org/abs/1906.07413>.
- Hien Dang, Tan Nguyen, Tho Tran, Hung Tran, and Nhat Ho. Neural collapse in deep linear network: From balanced to imbalanced data. *arXiv preprint arXiv:2301.00437*, 2023.
- Hien Dang, Tho Tran Huu, Tan Minh Nguyen, and Nhat Ho. Neural collapse for cross-entropy class-imbalanced learning with unconstrained ReLU features model. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 10017–10040. PMLR, 21–27 Jul 2024. URL <https://proceedings.mlr.press/v235/dang24a.html>.
- Emily L Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL <https://proceedings.neurips.cc/paper/2014/file/2afe4567e1bf64d32a5527244d104cea-Paper.pdf>.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gouget, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyan Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei

Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymier, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhota, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Maria Tsimpoukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vítor Albiero, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaofang Wang, Xiaojian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. The llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.

Cong Fang, Hangfeng He, Qi Long, and Weijie J. Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 118(43):e2103091118, 2021. doi: 10.1073/pnas.2103091118. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2103091118>.

Philip Gage. A new algorithm for data compression. *C Users J.*, 12(2):23–38, feb 1994. ISSN 0898-9788.

Tomer Galanti, András György, and Marcus Hutter. On the role of neural collapse in transfer learning. In *International Conference on Learning Representations*, 2022. URL <https://openreview>.

- net/forum?id=SwIp410B6aQ.
- Tomer Galanti, András György, and Marcus Hutter. Generalization bounds for few-shot transfer learning with pretrained classifiers, 2023a. URL <https://arxiv.org/abs/2212.12532>.
- Tomer Galanti, Zachary S. Siegel, Aparna Gupte, and Tomaso Poggio. Characterizing the implicit bias of regularized sgd in rank minimization, 2023b. URL <https://arxiv.org/abs/2206.05794>.
- Tomer Galanti, Mengjia Xu, Liane Galanti, and Tomaso Poggio. Norm-based generalization bounds for sparse neural networks. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 42482–42501. Curran Associates, Inc., 2023c. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/8493e190ff1bbe3837eca821190b61ff-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/8493e190ff1bbe3837eca821190b61ff-Paper-Conference.pdf).
- Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. *Information and Inference: A Journal of the IMA*, 9(2):473–504, 05 2020. ISSN 2049-8772. doi: 10.1093/imaiai/iaz007. URL <https://doi.org/10.1093/imaiai/iaz007>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Oriol Vinyals, Jack W. Rae, and Laurent Sifre. Training compute-optimal large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS ’22*, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9781713871088.
- Wanli Hong and Shuyang Ling. Neural collapse for unconstrained feature model under cross-entropy loss with imbalanced data. *arXiv preprint arXiv:2309.09725*, 2023.
- Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5375–5384, 2016.
- Yiding Jiang\*, Behnam Neyshabur\*, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SJgIPJBFvH>.
- Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition, 2019. URL <https://arxiv.org/abs/1910.09217>.
- Bingyi Kang, Yu Li, Sa Xie, Zehuan Yuan, and Jiashi Feng. Exploring balanced feature spaces for representation learning. In *International Conference on Learning Representations*, 2020.
- Byungju Kim and Junmo Kim. Adjusting decision boundary for class imbalanced learning. *IEEE Access*, 8:81674–81685, 2020.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report UTML TR 2009-003, University of Toronto, 2009.
- Anders Krogh and John Hertz. A simple weight decay can improve generalization. In J. Moody, S. Hanson, and R.P. Lippmann (eds.), *Advances in Neural Information Processing Systems*, volume 4. Morgan-Kaufmann, 1991. URL [https://proceedings.neurips.cc/paper\\_files/paper/1991/file/8eefcfd5990e441f0fb6f3fad709e21-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/1991/file/8eefcfd5990e441f0fb6f3fad709e21-Paper.pdf).
- Aitor Lewkowycz and Guy Gur-Ari. On the training dynamics of deep networks with l2 regularization. NIPS ’20, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.

- Xingguo Li, Junwei Lu, Zhaoran Wang, Jarvis Haupt, and Tuo Zhao. On tighter generalization bound for deep neural networks: Cnns, resnets, and beyond, 2018. URL <https://arxiv.org/abs/1806.05159>.
- Zhiyuan Li and Sanjeev Arora. An exponential learning rate schedule for deep learning, 2019. URL <https://arxiv.org/abs/1910.07454>.
- Zhiyuan Li, Kaifeng Lyu, and Sanjeev Arora. Reconciling modern deep learning with traditional optimization analyses: the intrinsic learning rate. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- Xuantong Liu, Jianfeng Zhang, Tianyang Hu, He Cao, Yuan Yao, and Lujia Pan. Inducing neural collapse in deep long-tailed learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 11534–11544. PMLR, 2023.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019. URL <https://arxiv.org/abs/1711.05101>.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In Dekang Lin, Yuji Matsumoto, and Rada Mihalcea (eds.), *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL <https://aclanthology.org/P11-1015>.
- Sachin Mehta, Mohammad Sekhavat, Qingqing Cao, Max Horton, Yanzi Jin, Frank Sun, Iman Mirzadeh, Mahyar Najibikohnehshahri, Dmitry Belenko, Peter Zatloukal, and Mohammad Rastegari. Openelm: An efficient language model family with open training and inference framework. In *ICML Workshop*, 2024. URL <https://arxiv.org/abs/2404.14619>.
- Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep double descent: Where bigger models and more data hurt. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=B1g5sA4twr>.
- Preetum Nakkiran, Prayaag Venkat, Sham M. Kakade, and Tengyu Ma. Optimal regularization can mitigate double descent. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=7R7fAoUygoa>.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In Peter Grünwald, Elad Hazan, and Satyen Kale (eds.), *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pp. 1376–1401, Paris, France, 03–06 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v40/Neyshabur15.html>.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar,

- Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorný, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024. URL <https://arxiv.org/abs/2303.08774>.
- Sabrina Ortiz. What is google bard? here’s everything you need to know, 2023. [Online]. Available: <https://www.zdnet.com/article/what-is-google-bard-heres-everything-you-need-to-know/>. [Accessed: Apr. 25, 2024].
- Leyan Pan and Xinyuan Cao. Towards understanding neural collapse: The effects of batch normalization and weight decay, 2024. URL <https://arxiv.org/abs/2309.04644>.
- Mohammad Pezeshki, Amartya Mitra, Yoshua Bengio, and Guillaume Lajoie. Multi-scale feature learning dynamics: Insights for double descent. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 17669–17690. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/pezeshki22a.html>.
- Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets, 2022. URL <https://arxiv.org/abs/2201.02177>.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners, 2018. URL <https://d4mucfypsww.cloudfront.net/better-language-models/language-models.pdf>.
- Akshay Rangamani and Andrzej Banburski-Fahey. Neural collapse in deep homogeneous classifiers and the role of weight decay. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4243–4247, 2022. doi: 10.1109/ICASSP43922.2022.9746778.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge, 2015. URL <https://arxiv.org/abs/1409.0575>.



- Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In Katrin Erk and Noah A. Smith (eds.), *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1715–1725, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1162. URL <https://aclanthology.org/P16-1162>.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, USA, 2014. ISBN 1107057132.
- Sho Takase, Ryokan Ri, Shun Kiyono, and Takuya Kato. Large vocabulary size improves large language models, 2024. URL <https://arxiv.org/abs/2406.16508>.
- Chaofan Tao, Qian Liu, Longxu Dou, Niklas Muennighoff, Zhongwei Wan, Ping Luo, Min Lin, and Ngai Wong. Scaling laws with vocabulary: Larger models deserve larger vocabularies, 2024. URL <https://arxiv.org/abs/2407.13623>.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshev, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshtir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. Gemma 2: Improving open language models at a practical size, 2024. URL <https://arxiv.org/abs/2408.00118>.
- Christos Thrampoulidis, Ganesh R. Kini, Vala Vakilian, and Tina Behnia. Imbalance trouble: Revisiting neural-collapse geometry, 2022. URL <https://arxiv.org/abs/2208.05512>.
- A. N. Tikhonov. On the stability of inverse problems. *Proceedings of the USSR Academy of Sciences*, 39:195–198, 1943.
- Cagri Toraman, Eyup Halit Yilmaz, Furkan Sahinuc, and Oguzhan Ozelik. Impact of tokenization on language models: An analysis for turkish. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 22(4):1–21, March 2023. ISSN 2375-4702. doi: 10.1145/3578707. URL <http://dx.doi.org/10.1145/3578707>.

- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023a. URL <https://arxiv.org/abs/2302.13971>.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023b. URL <https://arxiv.org/abs/2307.09288>.
- Murad Tukan, Alaa Maalouf, Matan Weksler, and Dan Feldman. No fine-tuning, no cry: Robust svd for compressing deep networks. *Sensors*, 21(16), 2021. ISSN 1424-8220. doi: 10.3390/s21165599. URL <https://www.mdpi.com/1424-8220/21/16/5599>.
- Twan van Laarhoven. L2 regularization versus batch and weight normalization, 2017. URL <https://arxiv.org/abs/1706.05350>.
- Vikrant Varma, Rohin Shah, Zachary Kenton, János Kramár, and Ramana Kumar. Explaining grokking through circuit efficiency, 2023. URL <https://arxiv.org/abs/2309.02390>.
- Colin Wei and Tengyu Ma. Data-dependent sample complexity of deep neural networks via lipschitz augmentation, 2019. URL <https://arxiv.org/abs/1905.03684>.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report, 2024. URL <https://arxiv.org/abs/2407.10671>.
- Han-Jia Ye, Hong-You Chen, De-Chuan Zhan, and Wei-Lun Chao. Identifying and compensating for feature deviation in imbalanced deep learning. *arXiv preprint arXiv:2001.01385*, 2020.
- Xiyu Yu, Tongliang Liu, Xinchao Wang, and Dacheng Tao. On compressing deep models by low rank and sparse decomposition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 67–76, 2017. doi: 10.1109/CVPR.2017.15.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy8gdB9xx>.
- Guodong Zhang, Chaoqi Wang, Bowen Xu, and Roger Grosse. Three mechanisms of weight decay regularization. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bl1z-3Rct7>.

## A PROOF OF PROPOSITION 5.1

**Proposition 5.1.** Suppose  $d \geq V$ , then any global minimizer  $(W, H)$  of the problem obeys  $\ell_{\text{CE}}(Wh_{k,i}, y_k) = \log\left(\sum_{j=1}^V \exp\left(\frac{M_j}{V^2}\right)\right) - M_k$ .

*Proof.* By combining (a-d) in Theorem 4.1 in (Dang et al., 2024), we have:

$$\begin{aligned} w_k^\top h_{k,i} &= w_k^\top \mu_k \\ &= \sqrt{\frac{\lambda_H}{\lambda_W V(V-1)}} \left( V \sqrt{n_k} \mu_k^\top \mu_k - \sum_{m=1}^V \sqrt{n_m} \mu_m^\top \mu_k \right) \\ &= \sqrt{\frac{\lambda_H}{\lambda_W V(V-1)}} \left( V \sqrt{n_k} \sqrt{\frac{V-1}{V} \frac{\lambda_W}{\lambda_H} \frac{1}{n_k}} M_k \right) \\ &= M_k \end{aligned}$$

Similarly, we can show that  $w_j^\top h_{k,i} = \frac{M_j}{V^2}$ . Hence, the loss function is equal to:

$$\ell_{\text{CE}}(Wh_{k,i}, y_k) = -\log\left(\frac{\exp(w_k^\top h_{k,i})}{\sum_{j=1}^V \exp(w_j^\top h_{k,i})}\right) = \log\left(\sum_{j=1}^V \exp\left(\frac{M_j}{V^2}\right)\right) - M_k.$$

□

## B PER-TOKEN METRICS

This appendix presents Algorithm 1, which details our method for computing per-token metrics used throughout this paper. This algorithm is central to our analysis, enabling the calculation of fine-grained token-level performance measures that underpin our study’s findings.

---

### Algorithm 1: Token-level Metric Computation Algorithm

---

**Input:** Logits  $\mathbf{L}$ , Labels  $\mathbf{y}$ , Tokenizer  $\mathcal{T}$   
**Output:** Token-level metrics  $\mathcal{M}$

```

1 Initialize token metrics dictionary  $\mathcal{M}$ ;
2 foreach step do
3   Get logits  $\mathbf{L}$  and labels  $\mathbf{y}$ ;
4   Shift logits  $\mathbf{L}' \leftarrow \mathbf{L}[:, : -1]$ ;
5   Shift labels  $\mathbf{y}' \leftarrow \mathbf{y}[:, 1 :]$ ;
6   Compute predictions  $\hat{\mathbf{y}} \leftarrow \arg \max(\mathbf{L}')$ ;
7   foreach token t in  $\mathbf{y}'$  do
8     Compute per-token loss  $\ell_t \leftarrow \text{CrossEntropyLoss}(\mathbf{L}_t, \mathbf{y}'_t)$ ;
9     Update metrics  $\mathcal{M}[t] \leftarrow \mathcal{M}[t] + \{\text{loss} : \ell_t, \text{correct} : (\hat{\mathbf{y}}_t = \mathbf{y}'_t)\}$ ;
10  end
11 end
```

---

## C IMDB TOKEN DISTRIBUTION

As an additional evaluation, we examine the token frequency distribution of the IMDB dataset, as shown in Figure 8. Using a BPE tokenizer with a 32005-token vocabulary, we illustrate the distribution of high-frequency and low-frequency tokens. The histogram reveals a striking imbalance typical of natural language datasets: a few tokens occur extremely frequently, while most unique tokens appear rarely. This long-tail distribution is a fundamental characteristic of language data.

## D CLASS IMBALANCE IN VISION

In order to show that the behavior reported in the main text is not specific to LLMs but holds also for other types of domains, we experimented with image classification. We present an analysis of




Figure 8: **Token frequency distribution of the IMDB dataset** Maas et al. (2011) using a BPE tokenizer (vocab size 32005). y-axis: number of unique tokens per frequency bin. x-axis: token frequency bins (logarithmic scale). Rightmost bin shows 3 tokens (the, punctuation) appearing  $3^9$  to  $3^{10}$  times each. This highlights high-frequency tokens (the, punctuation) and rare tokens ("pilgr", "cron") in lower bins, illustrating the long-tail distribution typical in natural language datasets.

how dataset balance and weight decay affect the performance of a ResNet9 model (He et al., 2016) trained for CIFAR10 (Krizhevsky & Hinton, 2009) classification with imbalanced classes. This highlights the interplay between data distribution, regularization, and model performance. To obtain the imbalanced CIFAR10 dataset, we target a number of samples per class of [91, 142, 222, 347, 541, 845, 1317, 2055, 3205, 5000]. We randomly subsampled the corresponding classes in order to obtain those number of images per class.




Figure 9: Performance comparison of ResNet9 on balanced (left) and imbalanced (right) CIFAR10 datasets. The x-axis represents the 10 CIFAR10 classes, while the y-axis shows the test accuracy for each class. Different lines correspond to various weight decay values used during training. In the balanced dataset, performance is relatively uniform across classes, with overall homogenous degradation at high weight decay. The imbalanced dataset reveals a clear trend where less frequent classes perform worse, with this effect exacerbated by stronger weight decay values.