

Better Embeddings with Coupled Adam

Felix Stollenwerk*
AI Sweden

Tobias Stollenwerk
Forschungszentrum Jülich

Abstract

Despite their remarkable capabilities, LLMs learn word representations that exhibit the undesirable yet poorly understood feature of anisotropy. In this paper, we argue that the second moment in Adam is a cause of anisotropic embeddings, and suggest a modified optimizer called Coupled Adam to mitigate the problem. Our experiments demonstrate that Coupled Adam significantly improves the quality of embeddings, while also leading to better upstream and downstream performance on large enough datasets.

1 Introduction

Anisotropic Embeddings Large Language Models (LLMs) take a sequence of tokens as input and predict the next token. An embedding matrix is used to map the input tokens to the hidden space of the model, while an unembedding matrix provides the inverse mapping to the output token space. Although the two matrices can in principle be different, it is common practice to apply weight tying (Press and Wolf, 2017) and use the transpose of the embedding matrix for unembedding. During training, the model learns an embedding vector in hidden space for each token in the vocabulary. However, it is observed that those embedding vectors are clustered in a small subspace away from the origin (Gao et al., 2019). This anisotropy limits the semantic usefulness of the embeddings and, in turn, the expressiveness and generalizability of the model. Multiple attempts have been made to both explain the root cause of the problem and alleviate it (more on this in Sec. 7). In particular, Biś et al. (2021) have shown that the problem can be traced back to a mere shift of the mean embedding vector away from the origin. With the mean embedding vector as reference point, the embeddings feature

near-perfect isotropy. However, the role of the employed optimization algorithm has, to the best of our knowledge, not yet been investigated.

Optimization Algorithms Optimization algorithms are an indispensable ingredient in the training of neural networks generally and LLMs in particular. While SGD is the foundational optimization technique, Adam (Kingma and Ba, 2014) is the most widely used optimization techniques for LLMs due to its superior performance and robustness. While it provides multiple conceptual advantages over SGD, see e.g. Ruder (2017) for a detailed discussion, the one that is particularly striking with regard to word embeddings is that Adam is well-suited for sparse data. More concretely, this means that using Adam, the embedding update vectors for rare words are scaled up in comparison to those of more frequent words. This is relevant in the context of LLMs as word frequencies in the training data are typically very skewed and may differ by several orders of magnitude. Formally, this is captured by the *unigram probability distribution* $\tilde{p} \in [0, 1]^V$, which for a given dataset d and tokenizer t is defined by

$$\tilde{p}_i \equiv \tilde{p}_i(d, t) = \frac{n_i}{\sum_j n_j}, \quad (1)$$

where $i \in \mathcal{V} \equiv \{1, \dots, V\}$ is the vocabulary index and n_i is the total number of occurrences of the i -th token in the tokenized dataset. A visualization of an example unigram probability distribution can be found in App. A.

Our Contributions In this work, we combine the research areas of anisotropic embeddings and optimization algorithms and provide the following contributions:

- We show that the Adam optimizer plays a crucial role in causing anisotropic embeddings.

*Corresponding author: felix.stollenwerk@ai.se

- We suggest *Coupled Adam*, an easy-to-implement yet efficient adjustment of the original Adam optimization algorithm, which is specifically designed for embedding parameters in order to alleviate the anisotropy problem.
- We demonstrate that our method not only significantly improves the quality of word embeddings, but also has a beneficial effect on upstream and downstream performance for sufficiently large datasets.

2 On the Root Cause of Anisotropic Embeddings

We study the collective shift of the embeddings (that underlies the anisotropy problem), by analyzing their vector updates based on the optimization algorithms SGD and Adam. Weight tying is assumed, but only contributions from the output layer are considered, following Biś et al. (2021). Our results apply to all model architectures with a standard language modeling head.

2.1 Language Modeling Head

The equations for the standard language modeling head read

$$\mathcal{L} = -\log(p_t) \quad (2)$$

$$p_t = \frac{\exp(l_t)}{\sum_{j=1}^V \exp(l_j)} \quad (3)$$

$$l_i = e_i \cdot h, \quad (4)$$

where $\mathcal{L} \in \mathbb{R}_{\geq 0}$ is the loss for next token prediction, and $p_t \in [0, 1]$ is the predicted probability of the true token $t \in \mathcal{V}$. $l_i \in \mathbb{R}$ and $e_i \in \mathbb{R}^H$ denote the logits and embeddings for each token $i \in \mathcal{V}$, respectively. $h \in \mathbb{R}^H$ is the final hidden state provided by the model for a single token. Note that the operation in Eq. (4) is the dot product of two vectors in \mathbb{R}^H . Backward propagation yields the following gradients with respect to the input vectors e_i and h of Eq. (4):

$$g_i := \frac{\partial \mathcal{L}}{\partial e_i} = -(\delta_{it} - p_i) \cdot h \quad (5)$$

This result was first reported using a different notation in Biś et al. (2021), and is rederived in App. B for the reader’s convenience.

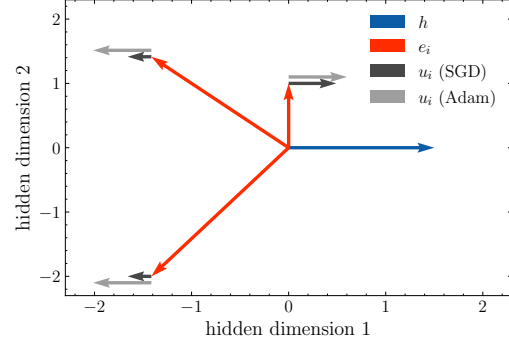


Figure 1: Toy example of a hidden state vector h (shown in blue) and three embedding vectors e_i (shown in red) in $H = 2$ dimensions. The gray arrows represent the embedding update vectors, for the SGD (dark) and the Adam (light) optimizer. The update vector of the true token is aligned with h , while the others point in the opposite direction, see Eq. (5). Note that the *sum of embedding update vectors* vanishes for SGD, while this is not necessarily the case for Adam, cf. Eqs. (11) and (16).

2.2 Vanishing Sum of Embedding Gradients

Optimization algorithms for neural networks usually update the model parameters iteratively, using an additive update vector that points in direction opposite to the gradient of the loss with respect to the parameters. In the case of embedding vectors, this can be expressed by

$$e_i^{(\tau)} = e_i^{(\tau-1)} + u_i^{(\tau)}, \quad (6)$$

with

$$u_i^{(\tau)} \propto -g_i^{(\tau)}, \quad (7)$$

where $u_i^{(\tau)}$ is the update vector for $e_i^{(\tau)}$ at time step τ . Eq. (5) implies that the embedding vector e_t of the true token is updated in direction $+h$, while the update vectors u_i for all the other embedding vectors e_i with $i \neq t$ are proportional to $-h$, see Fig. 1. This circumstance is referred to in the literature as the "common enemy effect" (Biś et al., 2021), and regarded as the cause of the representation degeneration problem. However, as we will see in the following sections, this explanation is incomplete, as it does not take into account the scaling of the gradients with the predicted probabilities p_i , see Eq. (5). The basis for our argumentation is the observation that the *sum of embedding gradients vanishes*, as the following simple calculation shows:

$$\begin{aligned} \sum_{i=1}^V g_i^{(\tau)} &\stackrel{(5)}{=} -\sum_{i=1}^V (\delta_{it}^{(\tau)} - p_i^{(\tau)}) \cdot h^{(\tau)} \\ &= -\left(1 - \sum_{i=1}^V p_i^{(\tau)}\right) \cdot h^{(\tau)} = 0 \end{aligned} \quad (8)$$

Next, we will study how Eq. (8) translates to the sum $\sum_{i=1}^V u_i^{(\tau)}$ of embedding update vectors, as well as the mean embedding vector

$$\mu^{(\tau)} = \frac{1}{V} \sum_{i=1}^V e_i^{(\tau)} \quad (9)$$

Since the exact definition of the embedding update vector u_i , i.e. the proportionality factor in Eq. (7), depends on the optimization algorithm, we discuss SGD and Adam separately.

2.3 Invariant Mean Embedding with SGD

We consider the application of the SGD optimization algorithm on the embedding vectors¹. At each training step, an embedding vector is simply updated by adding the associated negative gradient $-g_i$, multiplied by a global learning rate η . Hence, Eq. (7) becomes

$$u_i^{(\tau)} = -\eta \cdot g_i^{(\tau)} \quad (10)$$

Together with Eq. (8), this implies that the *sum of embedding update vectors vanishes* at any time step τ :

$$\sum_{i=1}^V u_i^{(\tau)} \stackrel{(10)}{=} -\eta \sum_{i=1}^V g_i^{(\tau)} \stackrel{(8)}{=} 0 \quad (11)$$

Consequently, the mean embedding vector will stay invariant during the training process:

$$\mu^{(\tau)} - \mu^{(\tau-1)} \stackrel{(9,6)}{=} \frac{1}{V} \sum_{i=1}^V u_i^{(\tau)} \stackrel{(11)}{=} 0 \quad (12)$$

This holds even though the different embeddings e_i will be individually updated in different directions with different magnitudes. Moreover, all of the above is true also in the case of SGD with momentum, which follows from linearity and mathematical induction. Eq. (12) has far-reaching implications with regard to the anisotropy problem. It entails that the embedding vectors do not collectively shift away from the origin if SGD (with or without momentum) is used.

2.4 Shifted Mean Embedding with Adam

In this section, we analyze the behavior of the mean embedding during optimization with Adam (Kingma and Ba, 2014), see Algorithm 1. The update vector Eq. (7) for the Adam algorithm

¹Details are given in App. C.

Input: η (lr), $e_i^{(0)}$ (initial embeddings), $\mathcal{L}(e_i)$ (objective), β_1, β_2 (betas), T (number of time steps)
Initialize: $m_i^{(0)} \leftarrow 0$ (1st moment), $v_i^{(0)} \leftarrow 0$ (2nd moment)
Output: $e^{(T)}$ (final embeddings)
1: **for** $\tau = 1 \dots T$ **do**
2: **for** $i = 1 \dots V$ **do**
3: $g_i^{(\tau)} \leftarrow \nabla_{e_i} \mathcal{L}^{(\tau)}(e_i^{(\tau-1)})$
4: $m_i^{(\tau)} \leftarrow \beta_1 m_i^{(\tau-1)} + (1 - \beta_1) g_i^{(\tau)}$
5: $v_i^{(\tau)} \leftarrow \beta_2 v_i^{(\tau-1)} + (1 - \beta_2) (g_i^{(\tau)})^2$
6: $\widehat{m}_i^{(\tau)} \leftarrow m_i^{(\tau)} / (1 - \beta_1^\tau)$
7: $\widehat{v}_i^{(\tau)} \leftarrow v_i^{(\tau)} / (1 - \beta_2^\tau)$
8: **if coupled then**
9: $\widehat{v}^{(\tau)} \leftarrow \frac{1}{V} \sum_{i=1}^V \widehat{v}_i^{(\tau)}$
10: **for** $i = 1 \dots V$ **do**
11: **if coupled then**
12: $\widehat{v}_i^{(\tau)} \leftarrow \widehat{v}^{(\tau)}$
13: $e_i^{(\tau)} \leftarrow e_i^{(\tau-1)} - \eta \frac{\widehat{m}_i^{(\tau)}}{\sqrt{\widehat{v}_i^{(\tau)} + \epsilon}}$
14: **return** $e^{(T)}$

Algorithm 1: Pseudocode for the Adam algorithm and our extension, the **Coupled Adam algorithm** (highlighted), applied to the embedding vectors e_i . Note that weight decay is not applied.

is given by

$$u_i^{(\tau)} = -\eta_i^{(\tau)} \cdot \widehat{m}_i^{(\tau)}, \quad (13)$$

where we have introduced an i -dependent effective learning rate

$$\eta_i^{(\tau)} := \frac{\eta}{\sqrt{\widehat{v}_i^{(\tau)} + \epsilon}} \quad (14)$$

Note that $\widehat{m}_i^{(\tau)}$ and $\widehat{v}_i^{(\tau)}$ denote the exponentially averaged first and second moments, respectively, defined according to lines 4-7 in Algorithm 1. The i -dependent learning rate serves the purpose of individually normalizing the update vectors for different parameters in the Adam optimizer. However, it also has an unwanted effect specifically on the embedding vectors. While we know from Eq. (8) and Algorithm 1 (lines 4,6) that the *unweighted* sum over the first moments vanishes, $\sum_{i=1}^V \widehat{m}_i^{(\tau)} = 0$, this is not true for the *weighted* sum,

$$\sum_{i=1}^V \eta_i^{(\tau)} \widehat{m}_i^{(\tau)} \neq 0, \quad (15)$$

unless $\eta_i^{(\tau)} = \eta_j^{(\tau)}$ for all $i, j \in \mathcal{V}$. Hence, the *sum of embedding update vectors does not vanish* in general,

$$\sum_{i=1}^V u_i^{(\tau)} \stackrel{(13)}{=} - \sum_{i=1}^V \eta_i^{(\tau)} \cdot \widehat{m}_i^{(\tau)} \stackrel{(15)}{\neq} 0 \quad (16)$$

This, in turn, causes the mean embedding to change during training,

$$\mu^{(\tau)} - \mu^{(\tau-1)} \stackrel{(9,6)}{=} \frac{1}{V} \sum_{i=1}^V u_i^{(\tau)} \stackrel{(16)}{\neq} 0, \quad (17)$$

which is in stark contrast to the case of SGD (cf. Eq. (12)). We have thus identified that an i -dependency of the second moment $\widehat{v}_i^{(\tau)}$ of the Adam optimizer leads to the observed collective shift of the embedding vectors away from the origin. Next, we will show that the second moment indeed depends on i . More concretely, we will argue that its expectation value is proportional to the unigram probability² (see Eq. (1)),

$$\mathbb{E}[\widehat{v}_i] \propto \widetilde{p}_i \quad (18)$$

In App. D.1, Eq. (18) is derived using minimal assumptions and experimental input. Here, we restrict ourselves to confirming the relationship in a purely experimental manner. $\mathbb{E}[\widehat{v}_i]$ is estimated directly by measuring \widehat{v}_i multiple times during training, using different models. We then perform linear fits of $\mathbb{E}[\widehat{v}_i]$ as a function of \widetilde{p}_i . Indeed, the fits yield a high coefficient of determination, on average $R^2 = 0.85(7)$, and a proportionality constant of

$$A := \frac{\mathbb{E}[\widehat{v}_i]}{\widetilde{p}_i} \approx 10^{-4} \quad (19)$$

Details about the exact procedure and plots showing the data and linear fits can be found in App. D.2.

3 Coupled Adam

In the previous section, we have identified the individual scales of the second moments v_i for different embedding vectors e_i as the root cause of the anisotropy problem. This implies that a solution to the problem is to enforce that the second moments are the same for every i . The question arises whether and how this can be done in the best way, without harming the performance of the model. To answer this, we note that the normalization of the embedding update vector by the Adam second moment can be split into two parts:

$$\mathbb{E}[\widehat{v}_i] \stackrel{(19)}{=} A \cdot \widetilde{p}_i = \frac{A}{V} \cdot (\widetilde{p}_i V) \quad (20)$$

²Note that from here until Eq. (23), the time index (τ) is dropped for the sake of readability.

The first factor introduces a global scale to all update vectors simultaneously:

$$\frac{A}{V} \stackrel{(19)}{\approx} \frac{10^{-4}}{5 \cdot 10^4} = 2 \cdot 10^{-9}, \quad (21)$$

where the numbers correspond to our experiments from the previous section with $V \approx 50000$. The second factor scales the update vectors individually. It is one on average:

$$\frac{1}{V} \sum_{i=1}^V (\widetilde{p}_i V) = 1 \quad (22)$$

Our goal is to retain the first, global factor and get rid of the second, individual factor. The canonical way to do this is to simply take the average of the second moment over the vocabulary items i :

$$\frac{1}{V} \sum_{i=1}^V \mathbb{E}[\widehat{v}_i] \stackrel{(20,22)}{=} \frac{A}{V} \quad (23)$$

In practice, the exponentially averaged second moments $\widehat{v}_i^{(\tau)}$ as they appear in Eq. (14) are replaced by their average:

$$\widehat{v}^{(\tau)} := \frac{1}{V} \sum_{i=1}^V \widehat{v}_i^{(\tau)} \quad (24)$$

We call the resulting algorithm *Coupled Adam*, as it couples the second moments of the embedding vectors via Eq. (24). It is displayed in Algorithm 1. Evidently, with Coupled Adam, the effective learning rate in Eq. (14) that enters the update vector in Eq. (13) becomes independent of i . Hence, like SGD but unlike standard Adam, the sum of embedding updates vanishes. However, like standard Adam but unlike SGD, Coupled Adam uses a second moment to normalize the embedding update vectors.

4 Experiments

Two types of experiments are conducted to study the impact of coupling the second moments of the embedding update vectors. First, a set of small-scale experiments (Sec. 4.1) with models and datasets of varying sizes up to 1B parameters and 20B tokens, respectively. Afterwards, we perform a few large-scale experiments (Sec. 4.2) to verify that the usefulness of our method extrapolates to the realm of large language models with more than 1B parameters trained on at least the corresponding compute-optimal (Hoffmann et al., 2022)

amount of data. In order to verify the generalizability of our method, the small- and large-scale experiments involve different datasets, training frameworks and dense transformer model architectures. An overview of the model and dataset sizes employed in our experiments is given in App. E.1. For each combination, two models are trained: one using standard Adam and one using Coupled Adam for the embeddings, see Eq. (24). Both variants use standard Adam for all non-embedding parameters. The various metrics we employ to assess both the general model performance and the quality of the model embeddings will be discussed in Sec. 4.3.

4.1 Small-scale Experiments

Our small-scale experiments use the OpenWebText Corpus (Gokaslan and Cohen, 2019) and the GPT-2 tokenizer (Radford et al., 2019). The model architecture also follows GPT-2, while the hyperparameter setup is taken from GPT-3 (Brown et al., 2020), see App. E.2 for further details. An implementation based on nanoGPT (Karpathy, 2022) is used. We define a grid (D, N) with dataset sizes $D \in \{5\text{B}, 10\text{B}, 20\text{B}\}$ and model sizes $N \in \{125\text{M}, 355\text{M}, 760\text{M}\}$, and repeat each experiment $S = 3$ times with different seeds in order to estimate uncertainties and assess statistical significance.

4.2 Large-scale Experiments

For our large-scale experiments, we use the SlimPajama dataset (Soboleva et al., 2023) and the GPT-2 tokenizer. A state-of-the-art dense transformer model architecture akin to (Touvron et al., 2023) is chosen, including e.g. RoPE embeddings (Su et al., 2023) and the SwiGLU activation function (Shazeer, 2020). Details can be found in App. E.2. The experiments are conducted using Modalities (Lübbering et al., 2024) as the training framework. We consider two model sizes, 1.3B and 2.6B. In order to cover the two common scenarios of compute-optimal training and overtraining, we conduct two sets of experiments: Firstly, we use near compute-optimal dataset sizes, 26B and 52B tokens, respectively. Secondly, we increase the number of tokens by a factor 4, resulting in 105B and 210B tokens, respectively. Each large-scale experiment is performed $S = 1$ times.

4.3 Evaluation

Upstream performance is measured in terms of test loss, while downstream performance is evalu-

ated using the Language Model Evaluation Harness (Gao et al., 2023) on the following tasks: ARC easy and challenge (Clark et al., 2018), HellaSwag (Zellers et al., 2019), LAMBADA (Paperno et al., 2016), RACE (Lai et al., 2017), TruthfulQA (Lin et al., 2022) and WinoGrande (Sakaguchi et al., 2020). More concretely, the considered metric is the average³ accuracy, which we will denote by Acc . To assess the quality of the embeddings, we first compute their isotropy, defined as (Arora et al., 2016; Mu et al., 2018)

$$\text{Iso}(E) := \frac{\min_{c \in X} Z(c)}{\max_{c \in X} Z(c)}, \quad (25)$$

where $E \in \mathbb{R}^{H \times V}$ is the embedding matrix, $Z(c) = \sum_{i=1}^V \exp(c^T e_i)$ is the partition function and $X = \{c\}$ is the set of eigenvectors $c \in \mathbb{R}^H$ of $EE^T \in \mathbb{R}^{H \times H}$. Secondly, the 2-norm $\|\mu\|$ of the mean embedding, see Eq. (9), and the average 2-norm of the embeddings $\|e_i\| = \frac{1}{V} \sum_{i=1}^V \|e_i\|$ as well as their ratio

$$\|\mu\|^r := \|\mu\| / \overline{\|e_i\|} \quad (26)$$

are determined. In addition, we evaluate the models on embedding benchmarks for word similarity and relatedness, to assess how well they represent semantic meaning. Following Biš et al. (2021), we consider the benchmarks SimLex999 (Hill et al., 2015), MEN (Bruni et al., 2014), WordSim353 (Finkelstein et al., 2001) and Stanford Rare Words (Luong et al., 2013). Each dataset provides pairs of words labeled with a ground truth score that represents the words’ semantic similarity. We derive model scores from the cosine similarity of the corresponding embedding vectors, and report the Pearson correlation of the two scores averaged over the datasets, which we denote by \bar{r} . Finally, some additional important properties of the embedding matrix are investigated. We study the correlation between the length of an embedding vector and the unigram probability,

$$\rho := 100 \cdot \text{corr}((\|e_i\|)_{i=1}^V, \tilde{p}), \quad (27)$$

to measure how well the former represents the latter. Furthermore, the condition number κ , defined as the ratio of the smallest and largest singular values of the embedding matrix, is determined in percent:

$$\kappa := 100 \cdot \frac{\min_i \Sigma_{ii}}{\max_i \Sigma_{ii}} \quad (28)$$

Here, $E = U\Sigma V^T$ denotes the singular value decomposition of the embedding matrix.

³Individual task performance is reported in App. G.3.

| D | N | Adam | \mathcal{L} (\downarrow) | Acc (\uparrow) | Iso (\uparrow) | $\ \mu\ $ (\downarrow) | $\ \mu\ ^r$ (\downarrow) | \bar{r} (\uparrow) | ρ (\uparrow) | κ (\uparrow) |
|-----|------|----------|--------------------------------|--------------------|--------------------|----------------------------|------------------------------|--------------------------|-----------------------|-------------------------|
| 5B | 125M | Standard | 3.14 (0) | 0.340 (2) | 0.31 (2) | 1.10 (6) | 0.67 (3) | 15 (3) | -54 (3) | 0.6 (1) |
| | | Coupled | 3.12 (1) | 0.339 (2) | 0.94 (1) | 0.02 (0) | 0.01 (0) | 55 (0) | 87 (1) | 4.8 (2) |
| | 355M | Standard | 2.95 (0) | 0.352 (3) | 0.44 (2) | 0.81 (2) | 0.67 (1) | 16 (2) | -47 (0) | 0.8 (0) |
| | | Coupled | 2.93 (0) | 0.350 (4) | 0.98 (0) | 0.01 (0) | 0.01 (0) | 56 (1) | 86 (0) | 6.8 (4) |
| | 760M | Standard | 2.85 (0) | 0.360 (3) | 0.43 (1) | 0.84 (1) | 0.63 (0) | 14 (3) | -49 (2) | 0.7 (0) |
| | | Coupled | 2.86 (0) | 0.357 (3) | 0.97 (0) | 0.01 (0) | 0.01 (0) | 55 (1) | 85 (1) | 6.9 (3) |
| 10B | 125M | Standard | 3.07 (0) | 0.343 (3) | 0.21 (3) | 1.58 (5) | 0.75 (0) | 9 (2) | -64 (5) | 0.4 (0) |
| | | Coupled | 3.03 (0) | 0.343 (1) | 0.91 (1) | 0.05 (0) | 0.02 (0) | 57 (2) | 82 (0) | 3.6 (8) |
| | 355M | Standard | 2.86 (0) | 0.359 (2) | 0.35 (2) | 1.01 (4) | 0.74 (2) | 10 (3) | -55 (3) | 0.5 (0) |
| | | Coupled | 2.83 (0) | 0.365 (2) | 0.96 (0) | 0.02 (0) | 0.01 (0) | 57 (1) | 83 (1) | 5.3 (2) |
| | 760M | Standard | 2.75 (0) | 0.375 (2) | 0.38 (3) | 0.97 (4) | 0.66 (2) | 11 (2) | -56 (1) | 0.5 (0) |
| | | Coupled | 2.74 (1) | 0.372 (3) | 0.96 (0) | 0.02 (0) | 0.01 (0) | 57 (0) | 84 (0) | 6.2 (0) |
| 20B | 125M | Standard | 3.03 (0) | 0.346 (1) | 0.10 (3) | 2.14 (7) | 0.82 (2) | 5 (1) | -66 (2) | 0.3 (0) |
| | | Coupled | 2.97 (0) | 0.350 (1) | 0.83 (1) | 0.11 (0) | 0.03 (0) | 57 (2) | 77 (0) | 1.7 (5) |
| | 355M | Standard | 2.79 (0) | 0.366 (4) | 0.25 (2) | 1.32 (8) | 0.82 (2) | 5 (2) | -65 (4) | 0.3 (0) |
| | | Coupled | 2.75 (0) | 0.372 (6) | 0.95 (2) | 0.04 (0) | 0.02 (0) | 57 (1) | 78 (0) | 4.1 (3) |
| | 760M | Standard | 2.68 (1) | 0.385 (3) | 0.28 (2) | 1.21 (8) | 0.73 (3) | 3 (4) | -64 (2) | 0.3 (0) |
| | | Coupled | 2.65 (0) | 0.392 (2) | 0.94 (2) | 0.03 (0) | 0.01 (0) | 58 (0) | 81 (0) | 4.4 (2) |

Table 1: Results of our small-scale experiments. D and N denote the dataset and model size, respectively. \mathcal{L} is the test loss, and the column Acc represents the accuracy averaged over the downstream tasks listed in Sec. 4.3. The other evaluation metrics are defined in the same section, see Eqs. (25)-(28). The arrow in parentheses indicates whether a higher or lower value is desirable. Every training was conducted $S = 3$ times with different seeds, and the numbers represent the (rounded) averages and standard deviations in the following shorthand notation format: $0.123 (4) \equiv 0.123 \pm 0.004$. For each combination (D, N) and each metric, the respective better value is highlighted in bold if the (unrounded) difference is significant according to Student’s t-test with a one-sided confidence level of $\alpha = 95\%$ (see App. F for details). Plots for \mathcal{L} and Acc are shown in Fig. 2.

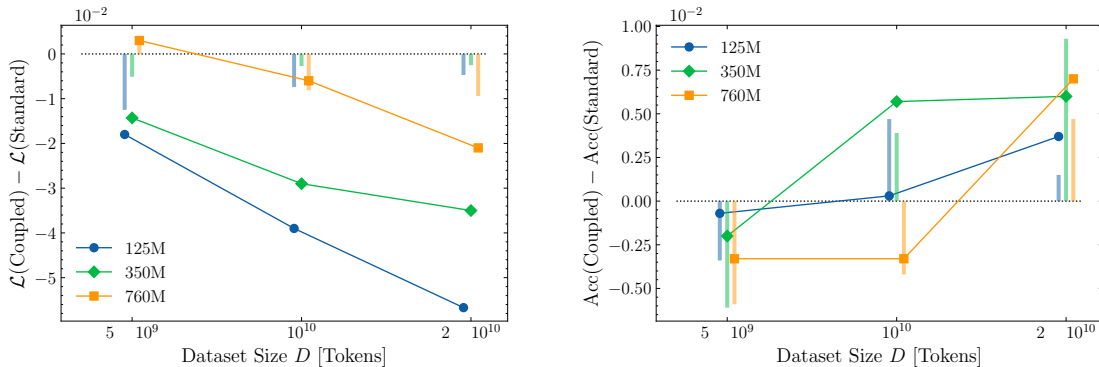


Figure 2: Difference in loss (left) and average downstream task accuracy (right) between Coupled Adam and standard Adam, for the different dataset sizes D (horizontal axis) and model sizes N (colors) of the small-scale experiments. The vertical bars indicate the one-sided 95% confidence interval for the difference to be significant. In order to avoid overlaps, the data points for $N = 125\text{M}$ and $N = 760\text{M}$ have been slightly shifted to the left and right, respectively.

5 Results

5.1 Small-scale Experiments

The results of the small-scale experiments (Sec. 4.1) are listed in Tab. 1 and illustrated in Fig. 2. We find that both upstream and downstream performance are better with Coupled Adam if the dataset size is sufficiently large. In fact, the improvement appears to increase monotonically with the dataset size D . In addition, the embedding-specific metrics benefit greatly from Coupled Adam. In particular, the isotropy reaches values above 0.90 (with a single exception), while \bar{r} and κ are hugely improved as well. The mean

embedding is evidently close to the origin. Finally, Coupled Adam leads to a significantly stronger (positive) correlation ρ between the length of an embedding vector and its associated unigram probability.

5.2 Large-scale Experiments

The results of the large-scale experiments (Sec. 4.2) are shown in Tab. 2. We observe very similar patterns as for the small-scale experiments. Although upstream and downstream performance are worse with Coupled Adam for compute-optimal dataset sizes, they are better if 4 times larger datasets are used. Note that for the small-scale experiments, the

| D | N | Adam | \mathcal{L} (\downarrow) | Acc (\uparrow) | Iso (\uparrow) | $\ \mu\ $ (\downarrow) | $\ \mu\ ^r$ (\downarrow) | \bar{r} (\uparrow) | ρ (\uparrow) | κ (\uparrow) |
|------|------|----------|--------------------------------|--------------------|--------------------|----------------------------|------------------------------|--------------------------|-----------------------|-------------------------|
| 26B | 1.3B | Standard | 2.433 | 0.402 | 0.50 | 0.67 | 0.45 | 53 | -41 | 1.2 |
| | | Coupled | 2.448 | 0.396 | 0.96 | 0.05 | 0.03 | 66 | 74 | 4.5 |
| 52B | 2.6B | Standard | 2.257 | 0.451 | 0.51 | 0.68 | 0.40 | 55 | -44 | 1.0 |
| | | Coupled | 2.273 | 0.441 | 0.86 | 0.10 | 0.06 | 66 | 74 | 3.0 |
| 105B | 1.3B | Standard | 2.278 | 0.446 | 0.40 | 0.87 | 0.43 | 52 | -52 | 0.6 |
| | | Coupled | 2.277 | 0.447 | 0.82 | 0.23 | 0.09 | 67 | 71 | 2.2 |
| 210B | 2.6B | Standard | 2.131 | 0.490 | 0.34 | 0.96 | 0.41 | 54 | -52 | 0.5 |
| | | Coupled | 2.129 | 0.492 | 0.65 | 0.49 | 0.14 | 67 | 69 | 1.5 |

Table 2: Results of our large-scale experiments. See the caption of Tab. 1 for an explanation of the column names. For each combination (D, N) and each metric, the respective better value is highlighted in bold.

upstream and downstream performance were found to be better already for compute-optimal dataset sizes. We attribute this to the fact that the batch size for the large-scale experiments is five times larger (cf. App. E.2), which results in fewer optimization steps for the same dataset size. Regarding the embedding-specific metrics, we again find significant and consistent improvements throughout all experiments. However, we do observe a certain shift of the mean embedding vector away from the origin, even if Coupled Adam is used. The shift becomes more pronounced as the model and dataset sizes increase, and is also reflected in a reduced isotropy. As we shall see in the following section, it comes along with optimal model performance though. An obvious hypothesis in light of our analysis in Sec. 2 is that the residual shift of the mean embeddings is due to weight tying. This is supported by the results of Machina and Mercer (2024), who find improved isotropy for models without weight tying. We leave it for future work to verify the hypothesis.

6 Ablations

We perform some additional experiments to shed further light on how Coupled Adam works. A model size of $N = 125\text{M}$ and the dataset sizes $D \in \{5\text{B}, 10\text{B}, 20\text{B}\}$ from the small-scale experiments (Sec. 4.1) are used, and each experiment is repeated $S = 3$ times with different seeds.

6.1 Scaled Coupled Adam

While coupling the second moment of the embedding gradients using the average in Eq. (24) is the canonical choice, one could also use a multiple of the average. We conduct additional experiments where the coupled second moment is scaled by powers of 2:

$$\hat{\mathcal{V}}^{(\tau)} \rightarrow 2^{-n} \cdot \hat{\mathcal{V}}^{(\tau)}, \quad (29)$$

with scaling exponents $n \in \{z \in \mathbb{Z} \mid -5 \leq z \leq 5\}$. Note that using a scaling exponent $n \neq 0$ is equivalent to using a different effective learning rate for the embeddings than for all the other parameters, via Eqs. (24) and (14). In particular, a smaller scaling exponent n corresponds to a smaller effective learning rate and vice versa. The results for $D = 20\text{B}$ are shown in Tab. 3, and the dependency of the loss on the scaling exponent n for that very dataset size is visualized in Fig. 3. Results for other dataset sizes and plots for the other evaluation metrics can be found in App. G.1. Our data shows that the loss reaches a minimum close to $n = 0$, with a rather weak dependence on the scaling exponent in its vicinity. Nevertheless, for the smallest and largest scaling exponents studied, we find that the loss gets significantly worse. Regarding downstream performance, we see indications of a similar pattern, although the statistical uncertainties are too large to draw definite conclusions. The semantic usefulness of the embedding vectors as measured by \bar{r} seems to suffer from a scaling exponent $n < 0$. For the isotropy and the mean embedding, we observe the opposite behavior. They benefit from a smaller scaling exponent n and the associated smaller embedding updates, with the effect being more pronounced the larger the training dataset size D . However, this also negatively affects the model performance. Hence, we conclude that, at least within the range of our experiments, the optimal setting is to have the same learning rate for the embedding parameters as for all the other model parameters, as implied by $n = 0$ and Eq. (24).

6.2 SGD

We train several models using SGD with momentum $\gamma = 0.9$ as the optimizer for the embeddings. Since Adam via the inverse square root of its second moment effectively scales the learning rate up by a factor comprising orders

| n | \mathcal{L} (\downarrow) | Acc (\uparrow) | Iso (\uparrow) | $\ \mu\ $ (\downarrow) | $\ \mu\ ^F$ (\downarrow) | \bar{r} (\uparrow) | ρ (\uparrow) | κ (\uparrow) |
|-----|--------------------------------|--------------------|--------------------|----------------------------|------------------------------|--------------------------|-----------------------|-------------------------|
| -5 | 2.99 (0) | 0.349 (2) | 0.97 (0) | 0.01 (0) | 0.01 (0) | 55 (1) | 76 (3) | 6.0 (6) |
| -4 | 2.99 (0) | 0.348 (5) | 0.97 (1) | 0.02 (0) | 0.01 (0) | 55 (1) | 78 (3) | 4.4 (1) |
| -3 | 2.98 (0) | 0.352 (5) | 0.95 (1) | 0.03 (0) | 0.02 (0) | 57 (2) | 78 (2) | 2.7 (5) |
| -2 | 2.98 (0) | 0.352 (1) | 0.94 (2) | 0.04 (0) | 0.02 (0) | 57 (1) | 78 (2) | 2.3 (4) |
| -1 | 2.98 (0) | 0.348 (3) | 0.87 (2) | 0.07 (0) | 0.02 (0) | 57 (1) | 79 (2) | 2.1 (3) |
| 0 | 2.97 (0) | 0.350 (1) | 0.83 (1) | 0.11 (0) | 0.03 (0) | 57 (2) | 77 (0) | 1.7 (5) |
| 1 | 2.97 (0) | 0.351 (4) | 0.66 (6) | 0.20 (1) | 0.03 (0) | 56 (0) | 78 (2) | 1.6 (5) |
| 2 | 2.98 (0) | 0.353 (2) | 0.47 (6) | 0.34 (1) | 0.04 (0) | 58 (1) | 78 (2) | 1.6 (9) |
| 3 | 2.97 (0) | 0.352 (0) | 0.27 (4) | 0.54 (2) | 0.05 (0) | 58 (2) | 78 (1) | 2.0 (7) |
| 4 | 2.97 (1) | 0.352 (1) | 0.09 (0) | 0.83 (2) | 0.05 (0) | 58 (1) | 77 (1) | 2.3 (2) |
| 5 | 2.98 (0) | 0.349 (4) | 0.01 (1) | 1.32 (2) | 0.06 (0) | 57 (1) | 75 (1) | 2.1 (6) |

Table 3: Results of our experiments with Scaled Coupled Adam, for $N = 125M$ and $D = 20B$. Values are highlighted in bold if they are significantly better than all the other values in the same column, see the caption of Tab. 1 for more details.

of magnitude (see Eq. (21)), we explicitly multiply the learning rate in SGD by a factor f of comparable size⁴. A hyperparameter search using $f \in \{100, 200, 300, 400, 500, 600\}$ is performed to search for the optimum with respect to upstream performance (loss), see App. G.2 for details. It is found at $f = 300$ for $D \in \{5B, 10B\}$ and $f = 400$ for $D = 20B$. The respective optimal model is compared to its counterpart trained with Coupled Adam in Tab. 4. The results show that, although SGD is advantageous with respect to isotropy, the mean embedding shift and the condition number, Coupled Adam consistently achieves better results on all upstream and downstream task metrics, while having one less hyperparameter to fine-tune.

7 Related Work

Gao et al. (2019) first described the anisotropy issue, which they referred to as *representation degeneration problem*, and suggested cosine regularization as a mitigation strategy. Alternative techniques to address the problem have been developed, including adversarial noise (Wang et al., 2019), spectrum control (Wang et al., 2020) and Laplacian regularization (Zhang et al., 2020). Biś et al. (2021) have shown that the anisotropy of embeddings can for the most part be traced back to a common shift of the embeddings in a dominant direction. They called this phenomenon *common enemy effect*, and provided a semi-quantitative explanation (Eq. (5)), which we developed further in the present work by including the optimizer in the analysis. In Yu et al. (2022), Adaptive Gradient Gating is proposed, based on the empirical observation that it is the gradients for embeddings of rare tokens that cause anisotropy. Our analysis conforms to this finding

⁴Note that the difference between momentum in SGD and the first moment in Adam also plays a role here.

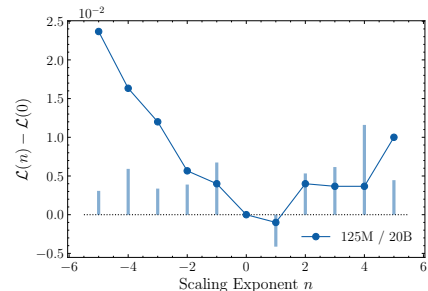


Figure 3: Dependency of the loss on the scaling exponent n , see Eq. (29), for $N = 125M$ and $D = 20B$. The plot shows the difference to the loss obtained for $n = 0$.

and attributes it to a massive up-scaling of the gradients for rare embeddings with Adam, cf. Fig. 1. Machina and Mercer (2024) have demonstrated that large Pythia models (Biderman et al., 2023) show improved isotropy compared to similar models, and attribute this to the absence of weight tying. This is in accordance with our analysis of the unembedding gradients in conjunction with Adam, Sec. 2. While all the previously mentioned papers use average cosine similarity (Ethayarajh, 2019) or Iso from Eq. (25) to quantify the geometry of embedding vectors, Rudman et al. (2022) deviate from this. Their notion of isotropy is based solely on the embeddings’ covariance matrix and embodied by the metric IsoScore. In particular, IsoScore is mean-agnostic, while Iso strongly correlates with the mean embedding (see e.g. Tab. 1). In a follow-up paper (Rudman and Eickhoff, 2024), IsoScore is used to regularize isotropy, which appears to benefit performance for fine-tuning tasks. Finally, concurrent to our work, Zhao et al. (2024) have investigated the importance of using the second moment in Adam with regard to performance and stability. They found that simplified variants of Adam that use the same effective learning rate either for the whole embedding matrix (Adalayer) or each embedding vector (Adalayer*) are slightly worse than Adam but better than SGD. Adalayer* is similar to Coupled Adam, but corresponds to the second moment averaged over hidden space instead of vocabulary space.

8 Conclusions

Our work addresses the well-known anisotropy problem for LLM embeddings. We have advanced the theoretical understanding of the phenomenon by showing that it is a combination of the common enemy effect and the individual second mo-

| D | N | Optimizer | \mathcal{L} (\downarrow) | Acc (\uparrow) | Iso (\uparrow) | $\ \mu\ $ (\downarrow) | $\ \mu\ ^F$ (\downarrow) | \bar{r} (\uparrow) | ρ (\uparrow) | κ (\uparrow) |
|-----|------|--------------|--------------------------------|--------------------|--------------------|----------------------------|------------------------------|--------------------------|-----------------------|-------------------------|
| 5B | 125M | SGD (300) | 3.17 (0) | 0.333 (3) | 0.99 (0) | 0.00 (0) | 0.01 (0) | 45 (1) | 71 (1) | 15.5 (4) |
| | | Coupled Adam | 3.12 (1) | 0.339 (2) | 0.94 (1) | 0.02 (0) | 0.01 (0) | 55 (0) | 87 (1) | 4.8 (2) |
| 10B | 125M | SGD (300) | 3.07 (1) | 0.341 (4) | 0.99 (0) | 0.01 (0) | 0.01 (0) | 49 (1) | 71 (4) | 12.2 (2) |
| | | Coupled Adam | 3.03 (0) | 0.343 (1) | 0.91 (1) | 0.05 (0) | 0.02 (0) | 57 (2) | 82 (0) | 3.6 (8) |
| 20B | 125M | SGD (400) | 3.00 (0) | 0.346 (5) | 0.98 (1) | 0.01 (0) | 0.02 (0) | 54 (0) | 76 (5) | 7.4 (1.1) |
| | | Coupled Adam | 2.97 (0) | 0.350 (1) | 0.83 (1) | 0.11 (0) | 0.03 (0) | 57 (2) | 77 (0) | 1.7 (5) |

Table 4: Comparison of models whose embeddings were trained with SGD and Coupled Adam. The SGD models were obtained after hyperparameter search for the learning rate. The associated factor f is specified in parentheses in the Optimizer column. Bold values indicate better results with statistical significance, see the caption of Tab. 1 for more details.

ments in Adam that causes a collective shift of the embedding vectors away from the origin. To mitigate the problem, we have introduced Coupled Adam, which enforces the same effective learning rate for every embedding vector, and thus suppresses the collective shift of the embeddings. We have found that Coupled Adam consistently improves embedding-specific metrics across all experiments, while also achieving better downstream and upstream performance for large datasets, as they are typically used in LLM training. The code to reproduce our results is available at github.com/flxst/coupled-adam.

9 Limitations

Although our method is generally applicable to all common LLM architectures, as they share the same language modeling head and embeddings, only dense decoders were used in our experiments. In addition, only models with up to $N = 2.6\text{B}$ parameters have been tested. Our experiments involve pre-training and few-shot downstream evaluation, yet fine-tuning tasks have not been included. The cosine decay learning rate schedule was applied throughout all experiments (App. E.2). Alternatives such as an infinite learning rate schedule are not incorporated in our study. It would also be interesting to extend our work to optimizers other than SGD and Adam. Furthermore, as mentioned at the end of Sec. 5, we have not explicitly verified that the slight residual shift of the mean embedding, which is observed even for Coupled Adam, is caused by weight tying. Finally, we have used a straightforward implementation of Coupled Adam, closely following Algorithm 1. More sophisticated implementations might lead to increased efficiency and further improvements; we leave it for future work to investigate this.

Acknowledgements

Our computational experiments used around 20000 GPU hours. They were partly run on the EuroHPC supercomputers MeluXina and MareNostrum5 in conjunction with the grants EHPC-DEV-2023D10-032 and EHPC-EXT-2023E02-038.

References

- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. 2016. [A latent variable model approach to PMI-based word embeddings](#). *Transactions of the Association for Computational Linguistics*, 4:385–399.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, Usvsn Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar Van Der Wal. 2023. [Pythia: A suite for analyzing large language models across training and scaling](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 2397–2430. PMLR.
- Daniel Biś, Maksim Podkorytov, and Xiuwen Liu. 2021. Too much in common: Shifting of embeddings in transformer language models and its implications. In *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Elia Bruni, Nam Khanh Tran, and Marco Baroni. 2014. Multimodal distributional semantics. *J. Artif. Int. Res.*, 49(1):1–47.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind

- Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge.](#)
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings.](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppín. 2001. [Placing search in context: The concept revisited.](#) volume 20, pages 406–414.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2019. [Representation degeneration problem in training natural language generation models.](#)
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonnell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2023. [A framework for few-shot language model evaluation.](#)
- Aaron Gokaslan and Vanya Cohen. 2019. [Openwebtext corpus.](#) <http://Skylion007.github.io/OpenWebTextCorpus>.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. [SimLex-999: Evaluating semantic models with \(genuine\) similarity estimation.](#) *Computational Linguistics*, 41(4):665–695.
- 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, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. 2022. [Training compute-optimal large language models.](#)
- Andrej Karpathy. 2022. [NanoGPT.](#) <https://github.com/karpathy/nanoGPT>.
- Diederik Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization.](#) *International Conference on Learning Representations*.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale Reading comprehension dataset from examinations.](#) In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [TruthfulQA: Measuring how models mimic human falsehoods.](#) In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.
- Thang Luong, Richard Socher, and Christopher Manning. 2013. [Better word representations with recursive neural networks for morphology.](#) In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 104–113, Sofia, Bulgaria. Association for Computational Linguistics.
- Max Lübbering, Mehdi Ali, Felix Stollenwerk, Michael Fromm, Alexander Arno Weber, and Richard Rutmann. 2024. [Modalities: A pytorch-native framework for distributed and reproducible foundation model training.](#) <https://github.com/Modalities/modalities>.
- Anemily Machina and Robert Mercer. 2024. [Anisotropy is not inherent to transformers.](#) In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4892–4907, Mexico City, Mexico. Association for Computational Linguistics.
- Jiaqi Mu, Suma Bhat, and Pramod Viswanath. 2018. [All-but-the-top: Simple and effective postprocessing for word representations.](#)
- Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. 2016. [The LAMBADA dataset: Word prediction requiring a broad discourse context.](#) In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1525–1534, Berlin, Germany. Association for Computational Linguistics.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models.](#) In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners.](#)
- Sebastian Ruder. 2017. [An overview of gradient descent optimization algorithms.](#)
- William Rudman and Carsten Eickhoff. 2024. [Stable anisotropic regularization.](#)
- William Rudman, Nate Gillman, Taylor Rayne, and Carsten Eickhoff. 2022. [IsoScore: Measuring the uniformity of embedding space utilization.](#) In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3325–3339, Dublin, Ireland. Association for Computational Linguistics.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [Winogrande: An adversarial winograd schema challenge at scale](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8732–8740.

Noam Shazeer. 2020. [Glu variants improve transformer](#).

Daria Soboleva, Faisal Al-Khateeb, Robert Myers, Jacob R Steeves, Joel Hestness, and Nolan Dey. 2023. [SlimPajama: A 627B token cleaned and deduplicated version of RedPajama](#).

Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. 2023. [Roformer: Enhanced transformer with rotary position embedding](#).

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti 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. 2023. [Llama 2: Open foundation and finetuned chat models](#).

Dilin Wang, Chengyue Gong, and Qiang Liu. 2019. [Improving neural language modeling via adversarial training](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6555–6565. PMLR.

Lingxiao Wang, Jing Huang, Kevin Huang, Ziniu Hu, Guangtao Wang, and Quanquan Gu. 2020. [Improving neural language generation with spectrum control](#). In *International Conference on Learning Representations*.

Sangwon Yu, Jongyoon Song, Heeseung Kim, Seongmin Lee, Woo-Jong Ryu, and Sungroh Yoon. 2022. [Rare tokens degenerate all tokens: Improving neural text generation via adaptive gradient gating for rare token embeddings](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 29–45, Dublin, Ireland. Association for Computational Linguistics.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.

Zhong Zhang, Chongming Gao, Cong Xu, Rui Miao, Qinli Yang, and Junming Shao. 2020. [Revisiting representation degeneration problem in language modeling](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 518–527, Online. Association for Computational Linguistics.

Rosie Zhao, Depen Morwani, David Brandfonbrener, Nikhil Vyas, and Sham Kakade. 2024. [Deconstructing what makes a good optimizer for language models](#).

A Unigram Probability Distribution

Fig. 4 shows the unigram probability distribution for the example of the OpenWebText Corpus dataset and the GPT-2 tokenizer.




Figure 4: Logarithm $\log(\tilde{p}_i)$ of the unigram probability distribution for the OpenWebText Corpus and the GPT-2 tokenizer. The maximum probability is $\max_i \tilde{p}_i \approx 0.037$ or $\max_i \log(\tilde{p}_i) \approx -3.30$. The minimum probability (not shown) is $\min_i \tilde{p}_i = 0$ or $\min_i \log(\tilde{p}_i) = -\infty$.

B Embedding Gradients

We explicitly derive Eq. (5), which we recall here for convenience:

$$g_i := \frac{\partial \mathcal{L}}{\partial e_i} = -(\delta_{it} - p_i) \cdot h \quad (5)$$

The chain rule yields

$$\frac{\partial \mathcal{L}}{\partial e_i} = \sum_{k=1}^V \frac{\partial \mathcal{L}}{\partial p_t} \cdot \frac{\partial p_t}{\partial l_k} \cdot \frac{\partial l_k}{\partial e_i}, \quad (30)$$

where the individual factors can directly be obtained from Eqs. (2)-(4):

$$\frac{\partial \mathcal{L}}{\partial p_t} = -\frac{1}{p_t} \quad (31)$$

$$\begin{aligned} \frac{\partial p_t}{\partial l_k} &= \frac{\delta_{kt} \exp(l_t) \cdot \Sigma - \exp(l_t) \exp(l_k)}{\Sigma^2} \\ &= \delta_{kt} p_t - p_t p_k \\ &= p_t (\delta_{kt} - p_k) \end{aligned} \quad (32)$$

$$\frac{\partial l_k}{\partial e_i} = \delta_{ki} h \quad (33)$$

Note that in the first line of Eq. (32), we use the abbreviation $\Sigma = \left(\sum_{j=1}^V \exp(l_j) \right)$. Inserting Eqs. (31), (32) and (33) into Eq. (30) directly leads to Eq. (5):

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial e_i} &= - \sum_{k=1}^V \frac{1}{p_t} \cdot p_t (\delta_{kt} - p_k) \cdot \delta_{ki} h \\ &= - \sum_{k=1}^V (\delta_{kt} - p_k) \cdot \delta_{ki} h \\ &= - (\delta_{it} - p_i) \cdot h \end{aligned}$$

C SGD Algorithm

For completeness and comparison to (Coupled) Adam as displayed in Algorithm 1, we summarize the SGD algorithm in Algorithm 2.

Input: η (lr), $e_i^{(0)}$ (initial embeddings), $\mathcal{L}(e_i)$ (objective), γ (momentum), T (number of time steps)
Output: $e^{(T)}$ (final embeddings)

```

1: for  $\tau = 1 \dots T$  do
2:   for  $i = 1 \dots V$  do
3:      $g_i^{(\tau)} \leftarrow \nabla_{e_i} \mathcal{L}^{(\tau)}(e_i^{(\tau-1)})$ 
4:     if  $t > 1$  then
5:        $\mathbf{b}_i^{(\tau)} \leftarrow \gamma \mathbf{b}_i^{(\tau-1)} + g_i^{(\tau)}$ 
6:     else
7:        $\mathbf{b}_i^{(\tau)} \leftarrow g_i^{(\tau)}$ 
8:      $e_i^{(\tau)} \leftarrow e_i^{(\tau-1)} - \eta \mathbf{b}_i^{(\tau)}$ 
9: return  $e^{(T)}$ 

```

Algorithm 2: Pseudocode for the SGD algorithm with optional momentum, applied to the embedding vectors e_i .

D Magnitude of the Second Moment in Adam

In this appendix, the validity of

$$\mathbb{E}[\hat{v}_i] \propto \tilde{p}_i \quad (18)$$

is verified. Due to the linearity of lines 5 and 7 in Algorithm 2, it suffices to show that the squared gradient has the property in question:

$$\mathbb{E}[g_i^2] \propto \tilde{p}_i \quad (34)$$

We do this in two different ways. First, we derive Eq. (34) using a semi-theoretical approach with minimal experimental input. Afterwards, we confirm the relationship in a purely experimental manner.

D.1 Semi-theoretical Derivation

Here, we derive an expression for the expectation value of the squared gradient in terms of simple observables (Theorem 2). Subsequently, the dependency of those observables on \tilde{p}_i is determined experimentally. Together, this will yield the proportionality expressed by Eq. (34). We begin our reasoning with a lemma.

Lemma 1 (Expectation Value Decomposition). *The expectation value of the squared gradient can be decomposed into conditional expectation values as follows:*

$$\begin{aligned} \mathbb{E}[g_i^2] &= \tilde{p}_i \cdot \mathbb{E}[g_i^2 \mid i = t] \\ &\quad + (1 - \tilde{p}_i) \cdot \mathbb{E}[g_i^2 \mid i \neq t] \end{aligned} \quad (35)$$

Proof. Our starting point is the definition of the expectation value for the continuous random variable g_i^2 :

$$\mathbb{E}[g_i^2] = \int g_i^2 p(g_i) dg_i, \quad (36)$$

where p denotes the probability distribution of g_i . Since the vocabulary item i can only be either the true token t or not, we can decompose p into a sum of joint probability distributions (using the *law of total probabilities*), each of which can be expressed in terms of conditional probabilities like so:

$$\begin{aligned} p(g_i) &= p(g_i, i = t) + p(g_i, i \neq t) \\ &= p(g_i \mid i = t) \cdot p(i = t) \\ &\quad + p(g_i \mid i \neq t) \cdot p(i \neq t) \end{aligned} \quad (37)$$

Using the unigram probability $\tilde{p}_i = p(i = t)$, this can also be written as

$$\begin{aligned} p(g_i) &= \tilde{p}_i \cdot p(g_i \mid i = t) \\ &\quad + (1 - \tilde{p}_i) \cdot p(g_i \mid i \neq t) \end{aligned} \quad (38)$$

If we insert Eq. (38) back into Eq. (36), the expectation value becomes

$$\begin{aligned} \mathbb{E}[g_i^2] &= \tilde{p}_i \cdot \int g_i^2 p(g_i \mid i = t) dg_i \\ &\quad + (1 - \tilde{p}_i) \cdot \int g_i^2 p(g_i \mid i \neq t) dg_i, \end{aligned} \quad (39)$$

which by definition of the (conditional) expectation value, Eq. (36), is equivalent to Eq. (35). \square

Theorem 2 (Expectation Value Squared Gradient). *Given that the squared hidden state vector h^2 is independent of p_i and whether i is the true token or not, the expectation value of the squared gradient g_i^2 is given by*

$$\mathbb{E} [g_i^2] = S \cdot \left[\tilde{p}_i \cdot X_i^{(i=t)} + (1 - \tilde{p}_i) \cdot X_i^{(i \neq t)} \right], \quad (40)$$

with

$$S := \mathbb{E} [h^2] \quad (41)$$

$$X_i^{(i=t)} := \mathbb{E} [(1 - p_i)^2 \mid i = t] \quad (42)$$

$$X_i^{(i \neq t)} := \mathbb{E} [p_i^2 \mid i \neq t] \quad (43)$$

Proof. We start from Lemma 1 and the square of the gradient,

$$g_i^2 \stackrel{(5)}{=} (\delta_{it} - p_i)^2 h^2 \quad (44)$$

Note that squared variables of vectors in \mathbb{R}^H always denote the elementwise (Hadamard) product, e.g.

$$g_i^2 \equiv g_i \odot g_i \in \mathbb{R}_{\geq 0}^H, \quad (45)$$

with strictly non-negative elements. Using Eq. (44), the expectation values on the right side of Eq. (35) can be expressed as

$$\mathbb{E} [g_i^2 \mid i = t] = \mathbb{E} [(1 - p_i)^2 \cdot h^2 \mid i = t] \quad (46)$$

$$\mathbb{E} [g_i^2 \mid i \neq t] = \mathbb{E} [p_i^2 \cdot h^2 \mid i \neq t] \quad (47)$$

Given our assumptions regarding h^2 , its expectation value can be factored out:

$$\mathbb{E} [g_i^2 \mid i = t] = S \cdot X_i^{(i=t)} \quad (48)$$

$$\mathbb{E} [g_i^2 \mid i \neq t] = S \cdot X_i^{(i \neq t)} \quad (49)$$

Inserting Eqs. (48) and (49) into Eq. (35) yields Eq. (40). \square

Note that Eq. (40) is a vector equation, with $\mathbb{E} [g_i^2]$, $S \in \mathbb{R}_{\geq 0}^H$ and $\tilde{p}_i, X_i^{(i=t)}, X_i^{(i \neq t)} \in \mathbb{R}_{\geq 0}$. It states that the expectation value of g_i^2 factorizes into a global constant S that is i -independent, and a factor that is i -dependent. The latter is a specific combination of the unigram probability \tilde{p}_i , determined by the data, and the conditional expectation values $X_i^{(i=t)}$ and $X_i^{(i \neq t)}$, determined by the model.

Experimental Input Regarding the unigram probability, we know that

1. $\tilde{p}_i \ll 1$.

This is the case for virtually all natural language datasets with a common vocabulary size of $V > 10000$, according to Zipf's law.

The conditional expectation values $X_i^{(i=t)}$ and $X_i^{(i \neq t)}$ can be empirically estimated by applying training data to different checkpoints. We consider the three small-scale experiments of Sec. 4.1 with $N \in \{125\text{M}, 355\text{M}, 760\text{M}\}$ and $D = 20\text{B}$, and take ten equidistant checkpoints after $D' \in \{2\text{B}, 4\text{B}, \dots, 20\text{B}\}$ seen tokens for each of them. We then continue pseudo-training on 20 batches ($\approx 2\text{k}$ samples or 2M tokens, see Tab. 5) of data using a zero learning rate, and measure the conditional probabilities in Eqs. (42, 43) from which our target quantities can be estimated. Subsequently, linear fits of the form

$$X_i^{(i=t)} = A^{(i=t)} \cdot \tilde{p}_i \quad (50)$$

$$X_i^{(i \neq t)} = A^{(i \neq t)} \cdot \tilde{p}_i, \quad (51)$$

with fit parameters $A^{(i=t)}$ and $A^{(i \neq t)}$ are performed. R^2 is used to assess the quality of the fits. In addition, the mutual information I between the response and the explanatory variable is computed. Since we observe only a very weak dependence of the results for R^2 and I on N and D' , we specify the mean and standard deviation over all experiments for them. Our findings are:

2. $X_i^{(i=t)}$ is independent of \tilde{p}_i .

The linear fits yield $R^2 = 0.003(1)$, and the mutual information is $I(X_i^{(i=t)}; \tilde{p}_i) = 0.14(2)$.

3. $X_i^{(i \neq t)}$ is proportional to \tilde{p}_i .

The linear fits yield $R^2 = 0.92(1)$, and the mutual information is $I(X_i^{(i \neq t)}; \tilde{p}_i) = 0.50(2)$.

The three empirical results above, together with Theorem 2, immediately lead to Eq. (34).

D.2 Experimental Confirmation

We reuse the experiments from the previous section to measure the second moment \hat{v}_i directly, in order to estimate $\mathbb{E} [\hat{v}_i]$. Again, linear fits of the form

$$\mathbb{E} [\hat{v}_i] = A \cdot \tilde{p}_i \quad (52)$$

are performed and the mutual information is computed. We find

- $\mathbb{E}[\hat{v}_i]$ is proportional to \tilde{p}_i .
The linear fits yield $R^2 = 0.85(7)$, and the mutual information is $I(\mathbb{E}[\hat{v}_i]; \tilde{p}_i) = 1.18(9)$.

The results for $N = 125\text{M}$ and $D = D' = 20\text{B}$ are depicted in Fig. 5, as an example.




Figure 5: Experimental results for $\mathbb{E}[\hat{v}_i]$ (vertical axis) vs. \tilde{p}_i (horizontal axis) for $N = 125\text{M}$ and $D = D' = 20\text{B}$. The blue line shows the linear fit with $R^2 = 0.91$.

Note that while R^2 and I are again virtually independent of N and D' , the fit parameter A is not. Instead, it seems to increase with D' , as shown in Fig. 6. However, as stated in Eq. (19), the order




Figure 6: Experimental results for the linear fit parameter A as a function of N and D' .

of magnitude is $A \approx 10^{-4}$ throughout our experiments.

E Experimental Details

E.1 Model and Dataset Sizes

The model sizes N and dataset sizes D employed in our experiments are depicted in Fig. 7.




Figure 7: Overview of the dataset (horizontal axis) and model sizes (vertical axis) involved in our small-scale (blue, green and orange circles) and large-scale (red squares) experiments. The dashed, black line shows $N = D/20$, which is approximately the compute-optimal trajectory according to Hoffmann et al. (2022).

E.2 Training Hyperparameters

In Tab. 5, we list the general hyperparameters used in our small-scale (Sec. 4.1) and large-scale (Sec. 4.2) experiments. During warm-up, the learn-

| Description | Small-scale | Large-scale |
|------------------------|--------------------|--------------------|
| optimizer | | AdamW |
| β_1 | | 0.9 |
| β_2 | | 0.95 |
| ϵ | | 1e-8 |
| weight decay | | 0.1 |
| gradient clipping | | 1.0 |
| dropout | | 0.0 |
| weight tying | | true |
| vocab size | | 50304 |
| learning rate schedule | | cosine decay |
| layer normalization | | LayerNorm |
| precision | | BF16 |
| hidden activation | GeLU | SwiGLU |
| positional embedding | absolute (learned) | RoPE |
| sequence length | 1024 | 2048 |
| batch size (samples) | 96 | 256 |
| batch size (tokens) | $\sim 100\text{k}$ | $\sim 500\text{k}$ |
| warmup | 100 steps | 1% of steps |
| training framework | nanoGPT | Modalities |
| training parallelism | DDP | FSDP |

Table 5: General hyperparameters used in our two sets of experiments.

ing rate is increased from zero to the maximum learning rate. This is followed by a cosine decay which reduces the learning rate to 10% of the maximum at the end of training. Note that weight decay is applied only to linear layers, not layer norms or embeddings. Tab. 6 shows the hyperparameters related to model size, following GPT-3 (Brown et al., 2020).

F Error Analysis and Statistical Significance

For the error analysis, two separate random variables, X_0 and X_1 , are considered. The symbol X

| N | lr | heads | layers | emb. dim. |
|------|--------|-------|--------|-----------|
| 125M | 6.0e-4 | 12 | 12 | 768 |
| 350M | 3.0e-4 | 16 | 24 | 1024 |
| 760M | 2.5e-4 | 16 | 24 | 1536 |
| 1.3B | 2.0e-4 | 32 | 24 | 2048 |
| 2.6B | 1.6e-4 | 32 | 32 | 2560 |

Table 6: Model-size dependent hyperparameters used in our experiments. N denotes the model size in terms of parameters, while lr corresponds to the maximum learning rate.

represents one of the metrics discussed in Sec. 4.3, while 0 and 1 stand for two approaches that are to be compared, like standard Adam and Coupled Adam, for instance. For each of the two random variables $i = \{0, 1\}$, we conduct and evaluate S training runs with different seeds, yielding results

$$\{X_i^{(1)}, \dots, X_i^{(S)}\} \quad (53)$$

While it is desirable to have a large sample size S , it is prohibitively expensive for large model and dataset sizes to repeat training runs. We use

$$S = 3 \quad (54)$$

except for the large-scale experiments (Sec. 4.2), where we restrict ourselves to

$$S = 1 \quad (55)$$

We are interested in the difference

$$d = X_1 - X_0 \quad (56)$$

For $S = 1$, it can be computed straight forwardly. However, no statement about the statistical uncertainty or significance of d can be made. In the case of $S = 3$, we apply a one-sided Student’s t-test with a confidence level of

$$\alpha = 95\% \quad (57)$$

First, the sample means

$$\bar{X}_i = \frac{1}{S} \sum_{s=1}^S X_i^{(s)} \quad (58)$$

and the corrected sample standard deviations

$$\hat{\sigma}_i^2 = \frac{1}{S-1} \sum_{s=1}^S \left(X_i^{(s)} - \bar{X}_i \right)^2 \quad (59)$$

for the two samples $i \in \{0, 1\}$ are estimated. The sample means from Eq. (58) are combined to an estimate for their difference,

$$\bar{d} = \bar{X}_1 - \bar{X}_0 \quad (60)$$

and the sample standard deviations from Eq. (59) are propagated to the sample standard deviation of d via Gaussian error propagation:

$$\hat{\sigma}_d = \sqrt{\left(\frac{\partial d}{\partial X_0} \cdot \hat{\sigma}_0 \right)^2 + \left(\frac{\partial d}{\partial X_1} \cdot \hat{\sigma}_1 \right)^2} \stackrel{(56)}{=} \sqrt{\hat{\sigma}_0^2 + \hat{\sigma}_1^2} \quad (61)$$

Student’s t-distribution for the chosen confidence level α (see Eq. (57)) and the

$$\nu = S - 1 \stackrel{(54)}{=} 2 \quad (62)$$

degrees of freedom yields

$$t_{\alpha, \nu} = 2.92 \quad (63)$$

With S , σ_d and $t_{\alpha, \nu}$ from Eqs. (54), (61) and (63) as ingredients, the one-sided confidence threshold for the difference can be computed as

$$d_{\text{significance}} = t_{\alpha, \nu} \cdot \frac{\hat{\sigma}_d}{\sqrt{S}} \quad (64)$$

Hence, the estimate \bar{d} from Eq. (60) is considered a statistically significant improvement of approach $i = 1$ over approach $i = 0$ if

$$\bar{d} < -d_{\text{significance}} \quad (65)$$

for metrics where smaller values are desirable (e.g. \mathcal{L}), and

$$\bar{d} > d_{\text{significance}} \quad (66)$$

for metrics where larger values are better (e.g. Acc).

G Additional Results

G.1 Scaled Coupled Adam

Tab. 3 of Sec. 6.1 shows the results of varying the scaling exponent n (see Eq. (29)) for $D = 20\text{B}$. The dependency of the loss is visualized in Fig. 3. Here, in Fig. 8, we extend the visualization of the results to $D \in \{5\text{B}, 10\text{B}, 20\text{B}\}$ and the other evaluation metrics.

G.2 SGD

In Tab. 4 of Sec. 6.2, we showed results for SGD using the best hyperparameter f . Detailed results of the corresponding hyperparameter searches can be found in Tab. 7.

G.3 Individual Downstream Task Performance

In Tab. 8-11, we list the individual downstream task performance for all our experiments (Sec. 4-6).

| D | N | Optimizer | \mathcal{L} (\downarrow) | Acc (\uparrow) | Iso (\uparrow) | $\ \mu\ $ (\downarrow) | $\ \mu\ ^{\text{r}}$ (\downarrow) | $\bar{\tau}$ (\uparrow) | ρ (\uparrow) | κ (\uparrow) | | |
|-----|------|---------------|--------------------------------|--------------------|--------------------------------|----------------------------|---------------------------------------|-----------------------------|---------------------------------------|-----------------------------|-----------------------|-------------------------|
| 5B | 125M | SGD (100) | 3.23 (0) | 0.325 (1) | 1.00 (0) | 0.00 (0) | 0.01 (0) | 33 (1) | 59 (1) | 25.0 (1.1) | | |
| | | SGD (200) | 3.19 (1) | 0.327 (2) | 1.00 (0) | 0.00 (0) | 0.01 (0) | 41 (1) | 66 (1) | 19.2 (9) | | |
| | | SGD (300) | 3.17 (0) | 0.333 (3) | 0.99 (0) | 0.00 (0) | 0.01 (0) | 45 (1) | 71 (1) | 15.5 (4) | | |
| | | SGD (400) | 3.18 (2) | 0.332 (3) | 0.99 (0) | 0.01 (0) | 0.01 (0) | 48 (3) | 75 (1) | 14.2 (5) | | |
| | | SGD (500) | 3.19 (2) | 0.330 (6) | 0.99 (0) | 0.01 (0) | 0.01 (0) | 49 (1) | 77 (1) | 13.2 (2) | | |
| | | SGD (600) | 3.24 (2) | 0.326 (3) | 0.99 (0) | 0.01 (0) | 0.01 (0) | 48 (2) | 79 (1) | 11.6 (2) | | |
| | | Standard Adam | 3.14 (0) | 0.340 (2) | 0.31 (2) | 1.10 (6) | 0.67 (3) | 15 (3) | -54 (3) | 0.6 (1) | | |
| | | Coupled Adam | 3.12 (1) | 0.339 (2) | 0.94 (1) | 0.02 (0) | 0.01 (0) | 55 (0) | 87 (1) | 4.8 (2) | | |
| | | <hr/> | | | | | | | | | | |
| | | D | N | Optimizer | \mathcal{L} (\downarrow) | Acc (\uparrow) | Iso (\uparrow) | $\ \mu\ $ (\downarrow) | $\ \mu\ ^{\text{r}}$ (\downarrow) | $\bar{\tau}$ (\uparrow) | ρ (\uparrow) | κ (\uparrow) |
| 10B | 125M | SGD (100) | 3.13 (1) | 0.337 (4) | 1.00 (0) | 0.00 (0) | 0.01 (0) | 41 (1) | 58 (4) | 22.2 (7) | | |
| | | SGD (200) | 3.09 (0) | 0.339 (4) | 0.99 (0) | 0.01 (0) | 0.01 (0) | 48 (1) | 65 (4) | 16.1 (3) | | |
| | | SGD (300) | 3.07 (1) | 0.341 (4) | 0.99 (0) | 0.01 (0) | 0.01 (0) | 49 (1) | 71 (4) | 12.2 (2) | | |
| | | SGD (400) | 3.07 (0) | 0.340 (3) | 0.99 (1) | 0.01 (0) | 0.01 (0) | 51 (2) | 76 (3) | 10.7 (8) | | |
| | | SGD (500) | 3.09 (0) | 0.338 (3) | 0.99 (0) | 0.01 (0) | 0.01 (0) | 52 (1) | 79 (3) | 10.3 (2) | | |
| | | SGD (600) | 3.11 (1) | 0.341 (4) | 0.98 (1) | 0.01 (0) | 0.01 (0) | 53 (0) | 81 (1) | 9.4 (7) | | |
| | | Standard Adam | 3.07 (0) | 0.343 (3) | 0.21 (3) | 1.58 (5) | 0.75 (0) | 9 (2) | -64 (5) | 0.4 (0) | | |
| | | Coupled Adam | 3.03 (0) | 0.343 (1) | 0.91 (1) | 0.05 (0) | 0.02 (0) | 57 (2) | 82 (0) | 3.6 (8) | | |
| | | <hr/> | | | | | | | | | | |
| | | D | N | Optimizer | \mathcal{L} (\downarrow) | Acc (\uparrow) | Iso (\uparrow) | $\ \mu\ $ (\downarrow) | $\ \mu\ ^{\text{r}}$ (\downarrow) | $\bar{\tau}$ (\uparrow) | ρ (\uparrow) | κ (\uparrow) |
| 20B | 125M | SGD (100) | 3.05 (1) | 0.343 (2) | 0.99 (0) | 0.01 (0) | 0.01 (0) | 50 (1) | 57 (7) | 18.4 (4) | | |
| | | SGD (200) | 3.02 (0) | 0.345 (1) | 0.99 (0) | 0.01 (0) | 0.01 (0) | 52 (1) | 64 (7) | 12.7 (7) | | |
| | | SGD (300) | 3.01 (1) | 0.350 (1) | 0.98 (1) | 0.01 (0) | 0.02 (0) | 53 (2) | 70 (6) | 9.0 (2) | | |
| | | SGD (400) | 3.00 (0) | 0.346 (5) | 0.98 (1) | 0.01 (0) | 0.02 (0) | 54 (0) | 76 (5) | 7.4 (1.1) | | |
| | | SGD (500) | 3.01 (1) | 0.348 (4) | 0.98 (1) | 0.02 (0) | 0.02 (0) | 55 (1) | 78 (5) | 7.5 (1.6) | | |
| | | SGD (600) | 3.04 (3) | 0.348 (5) | 0.95 (3) | 0.02 (0) | 0.02 (0) | 52 (5) | 81 (5) | 6.8 (2.8) | | |
| | | Standard Adam | 3.03 (0) | 0.346 (1) | 0.10 (3) | 2.14 (7) | 0.82 (2) | 5 (1) | -66 (2) | 0.3 (0) | | |
| | | Coupled Adam | 2.97 (0) | 0.350 (1) | 0.83 (1) | 0.11 (0) | 0.03 (0) | 57 (2) | 77 (0) | 1.7 (5) | | |

Table 7: Results of our experiments with SGD. Values are highlighted in bold if they are significantly better than all the other values in the same column.

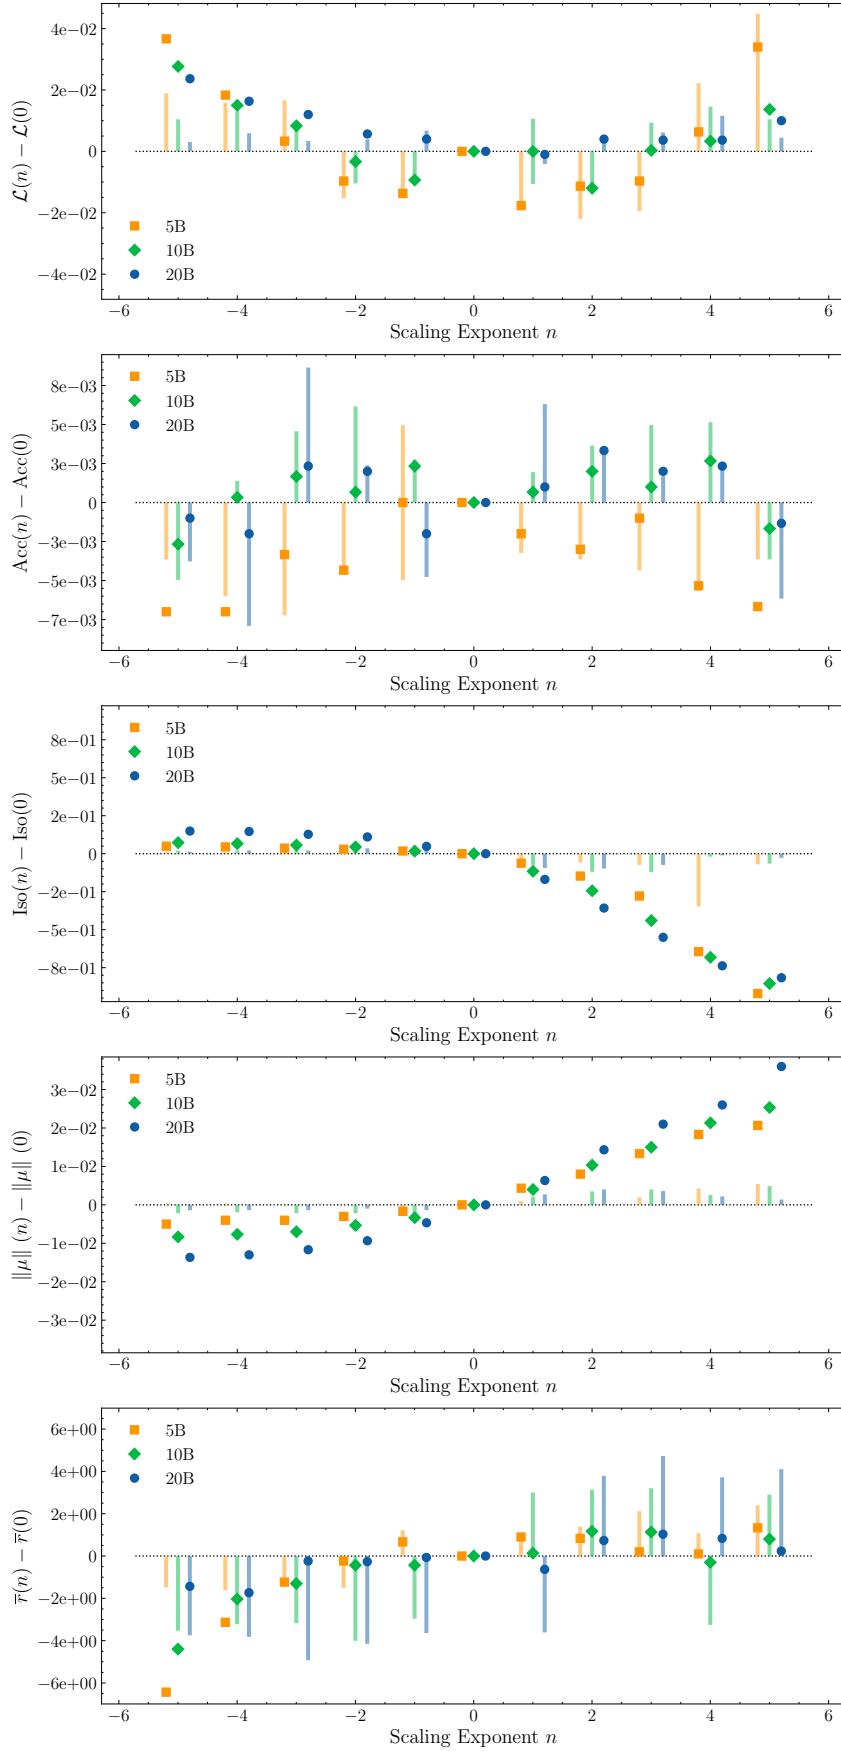


Figure 8: Dependency of different metrics on the scaling exponent n , see Eq. (29). From top to bottom: loss (upstream performance), average accuracy (downstream performance), isotropy, mean embedding norm ratio and \bar{r} . Each plot shows the difference to the respective metric obtained for $n = 0$. The arrows indicate whether larger (\uparrow) or smaller (\downarrow) values are desirable.

| D | N | Adam | ARC easy | ARC challenge | HellaSwag | LAMBADA | RACE | TruthfulQA | WinoGrande | Acc (\uparrow) |
|-----|------|----------|-----------------|-----------------|-----------------|-----------------|----------|-----------------|-----------------|--------------------|
| 5B | 125M | Standard | 0.41 (0) | 0.19 (1) | 0.27 (0) | 0.26 (0) | 0.28 (1) | 0.45 (0) | 0.52 (1) | 0.340 (2) |
| | | Coupled | 0.40 (0) | 0.18 (1) | 0.28 (0) | 0.27 (0) | 0.28 (0) | 0.45 (1) | 0.52 (1) | 0.339 (2) |
| | 355M | Standard | 0.43 (1) | 0.19 (1) | 0.29 (0) | 0.31 (0) | 0.29 (0) | 0.43 (0) | 0.52 (1) | 0.352 (3) |
| | | Coupled | 0.43 (2) | 0.19 (1) | 0.29 (0) | 0.31 (1) | 0.29 (0) | 0.43 (1) | 0.51 (0) | 0.350 (4) |
| | 760M | Standard | 0.46 (1) | 0.20 (1) | 0.30 (0) | 0.34 (1) | 0.29 (1) | 0.43 (1) | 0.50 (1) | 0.360 (3) |
| | | Coupled | 0.45 (0) | 0.20 (1) | 0.30 (0) | 0.34 (1) | 0.29 (1) | 0.42 (1) | 0.50 (0) | 0.357 (3) |
| 10B | 125M | Standard | 0.42 (1) | 0.19 (0) | 0.28 (0) | 0.29 (1) | 0.28 (1) | 0.44 (1) | 0.51 (1) | 0.343 (3) |
| | | Coupled | 0.42 (1) | 0.19 (1) | 0.28 (0) | 0.29 (0) | 0.28 (1) | 0.44 (0) | 0.50 (1) | 0.343 (1) |
| | 355M | Standard | 0.45 (1) | 0.19 (1) | 0.30 (0) | 0.35 (0) | 0.29 (1) | 0.41 (0) | 0.51 (1) | 0.359 (2) |
| | | Coupled | 0.46 (1) | 0.19 (1) | 0.30 (0) | 0.35 (1) | 0.30 (1) | 0.42 (0) | 0.52 (1) | 0.365 (2) |
| | 760M | Standard | 0.47 (0) | 0.21 (0) | 0.32 (0) | 0.39 (1) | 0.30 (1) | 0.41 (0) | 0.53 (1) | 0.375 (2) |
| | | Coupled | 0.48 (1) | 0.21 (1) | 0.32 (0) | 0.38 (0) | 0.30 (1) | 0.41 (1) | 0.51 (1) | 0.372 (3) |
| 20B | 125M | Standard | 0.43 (0) | 0.18 (1) | 0.28 (0) | 0.29 (1) | 0.28 (1) | 0.44 (1) | 0.51 (1) | 0.346 (1) |
| | | Coupled | 0.44 (1) | 0.19 (1) | 0.29 (0) | 0.31 (0) | 0.29 (1) | 0.43 (1) | 0.51 (0) | 0.350 (1) |
| | 355M | Standard | 0.46 (0) | 0.21 (1) | 0.31 (0) | 0.37 (2) | 0.29 (0) | 0.42 (1) | 0.51 (1) | 0.366 (4) |
| | | Coupled | 0.47 (1) | 0.21 (1) | 0.32 (0) | 0.38 (1) | 0.30 (0) | 0.42 (1) | 0.51 (0) | 0.372 (6) |
| | 760M | Standard | 0.49 (0) | 0.21 (1) | 0.33 (0) | 0.42 (0) | 0.30 (1) | 0.41 (1) | 0.53 (0) | 0.385 (3) |
| | | Coupled | 0.51 (1) | 0.22 (1) | 0.33 (0) | 0.41 (0) | 0.31 (1) | 0.41 (1) | 0.54 (0) | 0.392 (2) |

Table 8: Detailed downstream task results of our small-scale experiments from Sec. 4.1 and 5.1. Compare to Tab. 1.

| D | N | Adam | ARC easy | ARC challenge | HellaSwag | LAMBADA | RACE | TruthfulQA | WinoGrande | Acc (\uparrow) |
|------|------|----------|-------------|---------------|-------------|-------------|-------------|-------------|-------------|--------------------|
| 26B | 1.3B | Standard | 0.55 | 0.23 | 0.36 | 0.45 | 0.32 | 0.37 | 0.54 | 0.402 |
| | | Coupled | 0.52 | 0.23 | 0.36 | 0.43 | 0.32 | 0.38 | 0.53 | 0.396 |
| 52B | 2.6B | Standard | 0.61 | 0.27 | 0.43 | 0.55 | 0.35 | 0.38 | 0.58 | 0.451 |
| | | Coupled | 0.60 | 0.25 | 0.42 | 0.54 | 0.34 | 0.37 | 0.56 | 0.441 |
| 105B | 1.3B | Standard | 0.58 | 0.27 | 0.42 | 0.55 | 0.35 | 0.38 | 0.57 | 0.446 |
| | | Coupled | 0.60 | 0.26 | 0.42 | 0.54 | 0.35 | 0.39 | 0.57 | 0.447 |
| 210B | 2.6B | Standard | 0.65 | 0.29 | 0.48 | 0.63 | 0.37 | 0.37 | 0.63 | 0.490 |
| | | Coupled | 0.67 | 0.32 | 0.48 | 0.61 | 0.37 | 0.39 | 0.61 | 0.492 |

Table 9: Detailed downstream task results of our large-scale experiments from Sec. 4.2 and 5.2. Compare to Tab. 2.

| n | ARC easy | ARC challenge | HellaSwag | LAMBADA | RACE | TruthfulQA | WinoGrande | Acc (\uparrow) |
|-----|----------|---------------|-----------|----------|----------|------------|------------|--------------------|
| -5 | 0.43 (0) | 0.20 (0) | 0.28 (0) | 0.30 (1) | 0.29 (1) | 0.43 (1) | 0.51 (0) | 0.349 (2) |
| -4 | 0.43 (1) | 0.19 (1) | 0.28 (0) | 0.31 (1) | 0.28 (0) | 0.44 (1) | 0.51 (1) | 0.348 (5) |
| -3 | 0.43 (0) | 0.20 (1) | 0.28 (0) | 0.31 (0) | 0.28 (0) | 0.44 (2) | 0.52 (1) | 0.352 (5) |
| -2 | 0.43 (1) | 0.19 (1) | 0.29 (0) | 0.31 (1) | 0.29 (1) | 0.44 (0) | 0.52 (1) | 0.352 (1) |
| -1 | 0.43 (0) | 0.19 (1) | 0.29 (0) | 0.31 (1) | 0.28 (1) | 0.43 (1) | 0.50 (1) | 0.348 (3) |
| 0 | 0.44 (1) | 0.19 (1) | 0.29 (0) | 0.31 (0) | 0.29 (1) | 0.43 (1) | 0.51 (0) | 0.350 (1) |
| 1 | 0.43 (1) | 0.20 (1) | 0.29 (0) | 0.32 (1) | 0.29 (0) | 0.43 (2) | 0.51 (0) | 0.351 (4) |
| 2 | 0.44 (0) | 0.19 (0) | 0.29 (0) | 0.31 (1) | 0.29 (1) | 0.43 (1) | 0.52 (0) | 0.353 (2) |
| 3 | 0.45 (1) | 0.19 (0) | 0.29 (0) | 0.31 (0) | 0.29 (1) | 0.43 (0) | 0.51 (0) | 0.352 (0) |
| 4 | 0.43 (0) | 0.19 (1) | 0.28 (0) | 0.31 (0) | 0.30 (1) | 0.43 (1) | 0.52 (1) | 0.352 (1) |
| 5 | 0.44 (1) | 0.19 (1) | 0.29 (0) | 0.30 (0) | 0.28 (0) | 0.43 (1) | 0.51 (0) | 0.349 (4) |

Table 10: Detailed downstream task results of our ablations on Scaled Coupled Adam from Sec. 6.1. Compare to Tab. 3.

| D | N | Optimizer | ARC easy | ARC challenge | HellaSwag | LAMBADA | RACE | TruthfulQA | WinoGrande | Acc (\uparrow) |
|-----|------|--------------|----------|-----------------|-----------------|-----------------|----------|------------|------------|--------------------|
| 5B | 125M | SGD (300) | 0.40 (1) | 0.18 (1) | 0.27 (0) | 0.26 (0) | 0.28 (0) | 0.44 (1) | 0.50 (2) | 0.333 (3) |
| | | Coupled Adam | 0.40 (0) | 0.18 (1) | 0.28 (0) | 0.27 (0) | 0.28 (0) | 0.45 (1) | 0.52 (1) | 0.339 (2) |
| 10B | 125M | SGD (300) | 0.41 (0) | 0.19 (1) | 0.28 (0) | 0.28 (1) | 0.28 (1) | 0.44 (1) | 0.50 (1) | 0.341 (4) |
| | | Coupled Adam | 0.42 (1) | 0.19 (1) | 0.28 (0) | 0.29 (0) | 0.28 (1) | 0.44 (0) | 0.50 (1) | 0.343 (1) |
| 20B | 125M | SGD (400) | 0.42 (1) | 0.20 (0) | 0.28 (0) | 0.30 (1) | 0.28 (1) | 0.43 (0) | 0.51 (2) | 0.346 (5) |
| | | Coupled Adam | 0.44 (1) | 0.19 (1) | 0.29 (0) | 0.31 (0) | 0.29 (1) | 0.43 (1) | 0.51 (0) | 0.350 (1) |

Table 11: Detailed downstream task results of our ablations on SGD from Sec. 6.2. Compare to Tab. 4.