



---

# Dispersion Loss Counteracts Embedding Condensation and Improves Generalization in Small Language Models

---

Chen Liu<sup>\*1</sup> Xingzhi Sun<sup>\*1</sup> Xi Xiao<sup>\*23</sup> Alexandre Van Tassel<sup>\*1</sup> Ke Xu<sup>1</sup> Kristof Reimann<sup>1</sup> Danqi Liao<sup>1</sup>  
Mark Gerstein<sup>1</sup> Tianyang Wang<sup>2</sup> Xiao Wang<sup>3</sup> Smita Krishnaswamy<sup>1</sup>

 KrishnaswamyLab/LM-Dispersion  Project Page

## Abstract

Large language models (LLMs) achieve remarkable performance through ever-increasing parameter counts, but scaling incurs steep computational costs. To better understand LLM scaling, we study representational differences between LLMs and their smaller counterparts, with the goal of replicating the representational qualities of larger models in smaller models. We observe a geometric phenomenon which we term *embedding condensation*, where token embeddings collapse into a narrow cone-like subspace in some language models. Through systematic analyses across multiple Transformer families, we show that small models such as GPT2 and Qwen3-0.6B exhibit severe condensation, whereas larger models such as GPT2-x1 and Qwen3-32B are more resistant to this phenomenon. Additional observations show that embedding condensation is not reliably mitigated by knowledge distillation from larger models. To fight against it, we formulate a dispersion loss that explicitly encourages embedding dispersion during training. Experiments demonstrate that it mitigates condensation, recovers dispersion patterns seen in larger models, and yields performance gains across 10 benchmarks. We believe this work offers a principled path toward improving smaller Transformers without additional parameters.

## 1. Introduction

The remarkable success of large language models has fundamentally transformed natural language processing, with performance consistently improving as parameter

---

<sup>\*</sup>Equal contribution <sup>1</sup>Yale University <sup>2</sup>University of Alabama at Birmingham <sup>3</sup>Oak Ridge National Laboratory. Correspondence to: Smita Krishnaswamy <smita.krishnaswamy@yale.edu>.

Proceedings of the 43<sup>rd</sup> International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

counts scale from millions to trillions (Kaplan et al., 2020; Hoffmann et al., 2022; Minaee et al., 2024). However, this scaling presents significant practical challenges: larger models require substantial computational resources (Zhang et al., 2022; Dubey et al., 2024; OpenAI, 2025), making them inaccessible for many applications. This motivates a critical question: *Can we identify and replicate the key properties that make large models effective, thus improving smaller models without simply adding more parameters?*

Recent theoretical work on idealized models has shown that Transformer embeddings tend to cluster toward a single point as depth approaches infinity (Geshkovski et al., 2025), but the empirical manifestation of this phenomenon and its relationship to model performance remain underexplored. In this work, we provide a comprehensive empirical analysis of what we term *embedding condensation*.

**Definition 1.1.** We define **embedding condensation** as the phenomenon that pairwise cosine similarities of token embeddings in Transformer models concentrate near 1, implying that embedding vectors point towards nearly identical directions and therefore condense into a narrow cone in the representation space (Figure 1).

Through systematic similarity-based measurements of embedding vector directions across multiple Transformer families, we demonstrate that smaller models (e.g., GPT2, Qwen3-0.6B) exhibit severe embedding condensation, with embedding vectors concentrated towards nearly the same direction and therefore undermining representational diversity. In contrast, larger models (e.g., GPT2-x1, Qwen3-32B) naturally maintain *embedding dispersion*.

**Definition 1.2.** We define **embedding dispersion** as the resistance to embedding condensation, with embedding vectors pointing towards diverse directions, showing better coverage of the representation space.

This geometric perspective reveals a fundamental insight: *condensation might be a key bottleneck limiting the expressivity of smaller Transformers*. Notably, we observe that condensation emerges very early in the training process and cannot be easily resolved by knowledge distillation from larger models, motivating the need for mechanisms that explicitly target embedding geometry.




Figure 1. Illustration of the embedding condensation phenomenon. In pre-trained language models, embeddings of all tokens from the same input sequence condense into a narrow cone after being processed by many Transformer layers. This phenomenon is substantially more pronounced in smaller models than in larger models within the same family, which motivates our hypothesis in Section 3.3.

We hypothesize that the embedding dispersion of larger models leads to their superior performance, suggesting that counteracting condensation could narrow the performance gap between smaller and larger models.

To test this hypothesis, we propose a **dispersion loss** that explicitly encourages embedding dispersion during training, serving as an auxiliary objective that promotes representational diversity. Our empirical evaluations show that dispersion loss counteracts embedding condensation in smaller models and leads to performance gains across 10 language understanding tasks when applied to models in the GPT2 and Qwen3 families during mid-training.

Crucially, when incorporated into full pre-training, the proposed dispersion loss also yields an +1.17 average improvement across tasks, which achieves a 3.3% gain over the baseline trained with the default cross-entropy loss.

The key contributions of this work are listed below.

1. We observe and define the *embedding condensation* phenomenon, where cosine similarities between token embeddings concentrate towards 1 after being processed by Transformer layers.
2. We show that embedding condensation is more pronounced in smaller models than in larger models, emerges at initialization, and is not mitigated by standard knowledge distillation.
3. We propose a dispersion loss and three alternative formulations that explicitly regulate embedding geometry during training, with stable implementations designed for practical and scalable optimization.
4. We demonstrate that dispersion-aware training counteracts embedding condensation and improves model generalization, yielding consistent gains during mid-training and during full pre-training.

**Conflict of Interest Disclosure** No conflicts to declare.

## 2. Methods

### 2.1. Preliminaries: a theoretical perspective on embedding condensation

Consider a sequence of  $N$  tokens and let  $\mathcal{Z}^{(l)} = [z_1^{(l)}, z_2^{(l)}, \dots, z_N^{(l)}]^\top \in \mathbb{R}^{N \times d}$  denote the token embeddings after layer  $l$  in a Transformer. These token embeddings are the contextualized representations at each token position, in the form of  $d$ -dimensional vectors.  $\mathcal{Z}^{(l)}$  can be interpreted as  $N$  particles in a  $d$ -dimensional space, and Transformer layers are external impacts on the particle system. A theory paper (Geshkovski et al., 2025) has mathematically proven that, under idealized settings, these embeddings tend to cluster into a single point as the number of layers approaches infinity, but limited empirical evidence has been provided.

### 2.2. Quantifying the layer-by-layer evolution of embedding vector alignment in Transformers

Let  $z_i^{(l)} \in \mathbb{R}^d$  denote the embedding of token  $i$  after layer  $l$ . Pairwise cosine similarity quantifies the angle between two embedding vectors, as defined as  $\text{cossim}(z_i^{(l)}, z_j^{(l)}) = \frac{z_i^{(l)\top} \cdot z_j^{(l)}}{\|z_i^{(l)}\| \cdot \|z_j^{(l)}\|}$ . Cosine similarities lie in  $[-1, 1]$ , with a value of 1 indicating complete directional alignment,  $-1$  indicating opposite directions, and 0 indicating orthogonality.

To obtain the embeddings at each layer  $l$ , we feed the input sequence of length  $N$  tokens to the Transformer, and collect the token embeddings  $[z_1^{(l)}, z_2^{(l)}, \dots, z_N^{(l)}]^\top$  after the sequence has been processed by the model layers  $1, 2, \dots, l$ . We then compute cosine similarities  $\{\text{cossim}(z_i^{(l)}, z_j^{(l)})\}$  for all  $N^2$  pairs. The resulting values form a distribution that we visualize as a histogram for each layer. By stacking these histograms across depth, we create a heatmap that highlights the evolution of embedding vector alignment layer by layer.

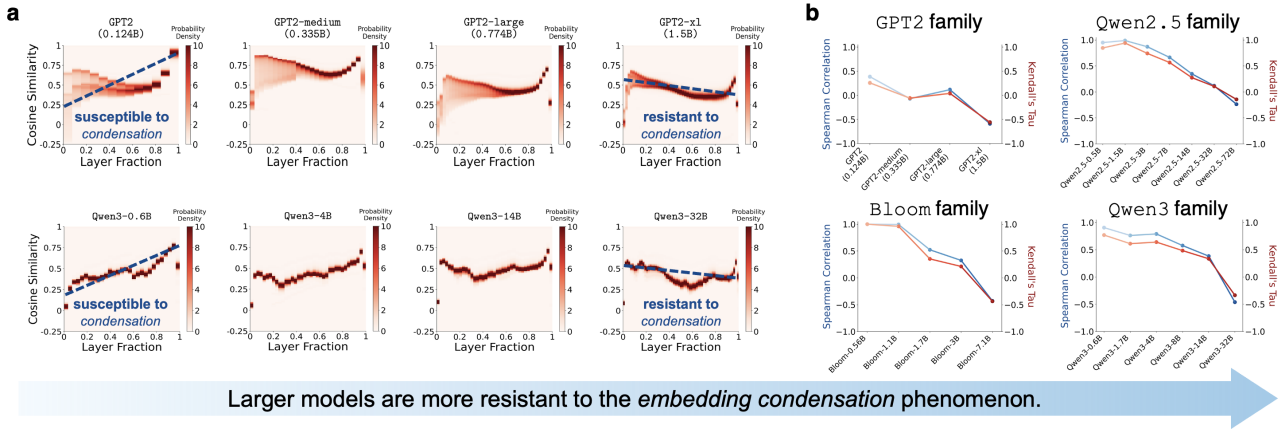


Figure 2. Qualitative and quantitative observations of the embedding condensation phenomenon. **a.** The cosine similarity heatmaps demonstrate that smaller models (e.g., GPT2, Qwen3-0.6B) are susceptible to condensation, since token cosine similarities become increasingly positive as the embeddings proceed to deeper layers. In contrast, larger models (e.g., GPT2-x1, Qwen3-32B) are more resistant to embedding condensation. **b.** Quantifications using Spearman correlation and Kendall’s Tau demonstrate a consistent trend of “larger model, less condensation” across multiple families of language models. Additional results can be found in Figure S1.

In this work, every heatmap is created using a population average over  $n = 100$  randomly selected input sequences from `wikitext-103` (Merity et al., 2017). We have experimented with different types of input text corpora, including `pubmed_qa` (Jin et al., 2019), `imdb` (Maas et al., 2011), and `squad` (Rajpurkar et al., 2016), and the same trends persist regardless of the dataset (see Appendix C.2).

### 2.3. Comparing the layer-by-layer evolution of embedding vector alignment across models

To provide quantitative comparisons among models, we compute two metrics to summarize the overall trend of the embedding condensation phenomenon.

For each Transformer layer  $l$ , we summarize the pairwise cosine similarity distribution by its mean value  $\mu^{(l)} = \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \text{cossim}(z_i^{(l)}, z_j^{(l)})$ , and then quantify the monotonic relationship between layer depth and embedding similarity by computing rank-based correlation between  $\{\mu^{(l)}\}_{l=1}^L$  and the layer index sequence  $\{l\}_{l=1}^L$ . Specifically, we report both the Spearman rank correlation coefficient  $\rho$  (Spearman, 1987) and Kendall’s Tau  $\tau$  (Kendall, 1938).

These statistics measure the extent to which embedding similarity changes monotonically with depth, independent of absolute scale or nonlinear distortions. Compared to simpler measures such as average cosine similarity in the last few layers, these rank-based correlation metrics show clearer trends (see Appendix C.3). Large positive/negative values indicate a monotonic increase/decrease in the directional alignment of the embedding vectors, while values near 0 indicate no systematic trend.

## 3. Key Observations

### 3.1. Observation of embedding condensation

#### 3.1.1. ANALYSES ON A VARIETY OF ARCHITECTURES

Applying the above analyses to multiple Transformer families reveals a clear trend that depend on model sizes. As shown in Figure 2, smaller models such as GPT2 and Qwen3-0.6B exhibit a **sharp upward drift** of cosine similarity distributions with depth. The embeddings become increasingly aligned, and in GPT2 the distribution collapses almost entirely near 1, indicating a near-perfect directional alignment. Qwen3-0.6B shows the same tendency, though its collapse remains less extreme. We refer to this degeneracy as embedding condensation (Definition 1.1).

In contrast, larger models such as GPT2-x1 and Qwen3-32B either maintain relatively moderate cosine similarities across layers or exhibit a gradual decrease following an initial increase, suggesting a stronger resistance to embedding condensation. We refer to this behavior as embedding dispersion (Definition 1.2).

For completeness, we also report a simpler metric, namely the average `cossim` over the last few layers. It shows the same overall trend, although the effect is less pronounced (see Appendix C.3). We therefore use Spearman’s  $\rho$  and Kendall’s  $\tau$  for all subsequent analyses.

**Takeaway 1: Larger model, less condensation.**

**Q:** What phenomenon do we observe?

**A:** Within the same model family, smaller models exhibit severe embedding condensation, with token embeddings collapsing toward near-parallel directions, while larger models resist this collapse.




Figure 3. In a highly controlled experiment, we reproduced the observation of “larger model, less condensation”. We pre-trained four GPT2-like models of varying sizes that differ only in MLP dimension, while keeping all other factors fixed, including the number of layers, embedding dimension, dataset, and training configuration. The resulting models exhibit consistent trends in embedding condensation, shown qualitatively (panel a) and quantitatively (panel b). Horizontal dashed lines are added to panel a for easier visual comparison.

### 3.1.2. CONFOUNDER-CONTROLLED EXPERIMENT

To further isolate the effect of model size from other confounding factors, we conduct a controlled experiment in which we pre-train four GPT2-like models for 1B training tokens, varying only the MLP dimension while keeping all other components fixed, including the number of layers, embedding dimension, dataset, and training settings. As shown in Figure 3, we observe the same trend: larger models consistently exhibit less embedding condensation. This controlled study provides stronger evidence that the relationship between model size and embedding condensation reflects a genuine empirical phenomenon rather than an artifact of correlated confounding factors.

### 3.2. Further investigations on embedding condensation

We perform additional analyses to better understand when embedding condensation arises and whether it can be alleviated by common training strategies.

#### 3.2.1. CONDENSATION EMERGES AT INITIALIZATION AND IS COUNTERACTED DURING TRAINING

We first track the evolution of embedding condensation throughout the pre-training process using checkpoints of `01mo-3-1025-7B` spanning initialization, intermediate stages of pre-training, and the final base model (Figure 4).

Embedding condensation is the most pronounced at initialization, with both correlation measures taking strongly positive values. This empirical observation is consistent with theoretical results (Geshkovski et al., 2025), which show that condensation arises in Transformers with random  $(Q, K, V)$  matrices. As training progresses, these correlations decrease and eventually become negative, indicating that pre-training dynamics counteract, rather than induce, the initial tendency of embedding condensation.

**Takeaway 2: Condensation occurs early on.**

**Q:** At which training stage does the embedding condensation phenomenon occur?

**A:** It emerges at model initialization and is gradually mitigated, not exacerbated, by pre-training.




Figure 4. Embedding condensation is observed immediately after model initialization. We analyze checkpoints of `01mo-3-1025-7B` spanning initialization, intermediate pre-training stages, and the final base model. Each checkpoint is annotated by its training stage and the number of training tokens.

#### 3.2.2. KNOWLEDGE DISTILLATION DOES NOT INHERENTLY MITIGATE CONDENSATION

Next, we examine whether knowledge distillation can transfer the favorable embedding geometry of larger models to smaller ones and thus provide a simple remedy to embedding condensation. Using the `Qwen2.5` family, we compare distilled models with their counterparts trained from scratch across a range of model sizes (Figure 5).

Distilled models exhibit embedding condensation trends that closely mirror those of non-distilled models, both qualitatively (Figure 5a) and quantitatively (Figure 5b). In particular, distillation neither consistently alleviates condensation in smaller models nor amplifies the dispersion behavior characteristic of larger models, indicating that resistance to condensation is **not automatically inherited** from a larger teacher through distillation (in this case, the teacher model is `DeepSeek-R1` with 671B parameters). These results motivate the need for explicit mechanisms that directly target embedding geometry during training.

This behavior is expected given the form of the knowledge distillation objective. In modern LLM distillation, the student is trained to match the next-token distribution of the teacher through a logit-level distillation loss. As described in `DeepSeek-R1` (Guo et al., 2025), the distillation

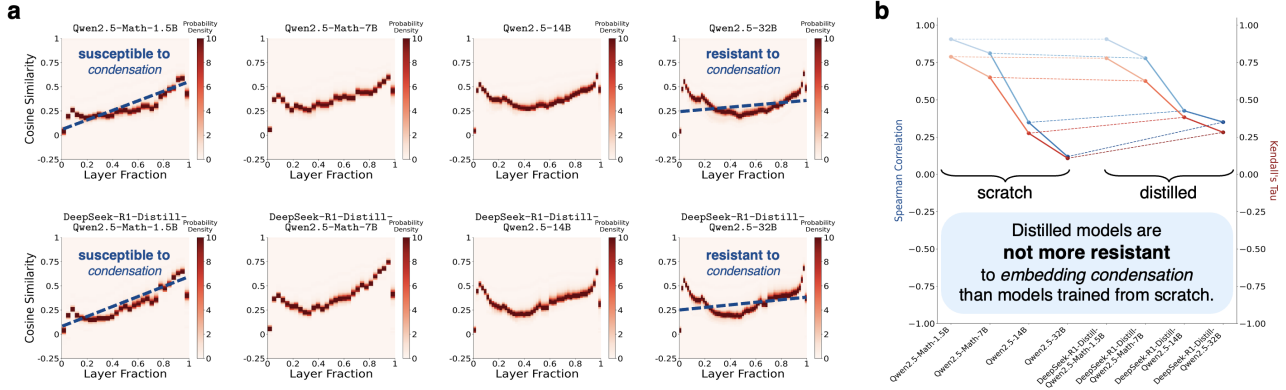


Figure 5. Knowledge distillation is not a remedy to embedding condensation, shown qualitatively (panel a) and quantitatively (panel b).

loss (Busbridge et al., 2025) between the teacher logits  $\ell_T^{(i)}$  and the student logits  $\ell_S^{(i)}$  at token  $i$  is given by equation 1.

$$\mathcal{L}_{\text{KD}}(\ell_T^{(i)}, \ell_S^{(i)}) = -\tau^2 \sum_{a=1}^V \sigma_a \left( \frac{\ell_T^{(i)}}{\tau} \right) \log \sigma_a \left( \frac{\ell_S^{(i)}}{\tau} \right) \quad (1)$$

$$\sigma_a(\ell) = \frac{\exp(\ell_a)}{\sum_{b=1}^V \exp(\ell_b)}, \quad a = 1, \dots, V.$$

The student is trained using a weighted combination of this term and the standard next-token prediction loss.

By construction, knowledge distillation primarily constrains the student at the level of output distributions. It does not explicitly regulate intermediate token embeddings, their pairwise relationships, or the layer-wise gradients that shape the representation geometry. Consequently, while knowledge distillation can effectively transfer predictive behavior, it does not inherently control the internal representational dynamics responsible for embedding condensation, explaining why resistance to condensation is not automatically inherited from a larger teacher model.

**Takeaway 3:** Distillation is not a solution.

**Q:** Larger models are resistant to condensation. Can we distill from them to obtain this resistance?

**A:** No. Knowledge distillation does not transfer the desired resistance to embedding condensation.

### 3.3. Our Hypothesis

The observations above highlight an important implication: condensation reduces the diversity of directions in which tokens can be represented, effectively narrowing the model’s expressive capacity. More importantly, we found that it cannot be easily remedied by distillation from a large model.

These observations motivate the following hypothesis.

**Hypothesis:** Larger models are better in language tasks because they counteract condensation.

Embedding condensation reduces the expressivity of Transformers by collapsing token embedding vectors into narrow cones, under-utilizing the representation space. We hypothesize that by **dispersing embeddings during training**, smaller models can achieve representational qualities more similar to larger models, thus narrowing the performance gap without increasing the number of parameters.

### 3.4. Our Remedy: Dispersion Loss

Our hypothesis motivates the design of auxiliary objectives that explicitly promote embedding dispersion during training. For this purpose, we propose to augment the training loss with a dispersion loss as a regularizer. The dispersion loss is given by equation 2 and the full training objective is given by equation 3. The pseudocode is provided in Appendix A. Here,  $\mathcal{L}_{\text{train}}$  denotes the standard training loss, which defaults to the cross-entropy loss for next-token prediction in most language models.

$$\mathcal{L}_{\text{disp}} = \log \sum_{i \neq j} e^{-\frac{\arccos(\text{cosim}(z_i, z_j))}{\pi \tau}} \quad (2)$$

$$\mathcal{L} = \mathcal{L}_{\text{train}} + \lambda_{\text{disp}} \cdot \mathcal{L}_{\text{disp}} \quad (3)$$

**Dispersion loss** The dispersion loss is a straightforward objective that directly counteracts the condensation of cosine similarities by spreading out all embeddings on the unit hypersphere (Table 1 row 1 and Figure 6a). In practice, we use the inverse cosine to map cosine similarity to an angular distance for numerical stability. During training, the loss is computed over the embedding vectors of each input sequence and aggregated across all layers. The resulting implementation has a time complexity of  $\mathcal{O}(N^2 F)$  per batch, where  $N$  is the sequence length and  $F$  is the feature dimension. In practice, we can subsample a small number of tokens over the sequence dimension, and the additional computational cost becomes manageable.

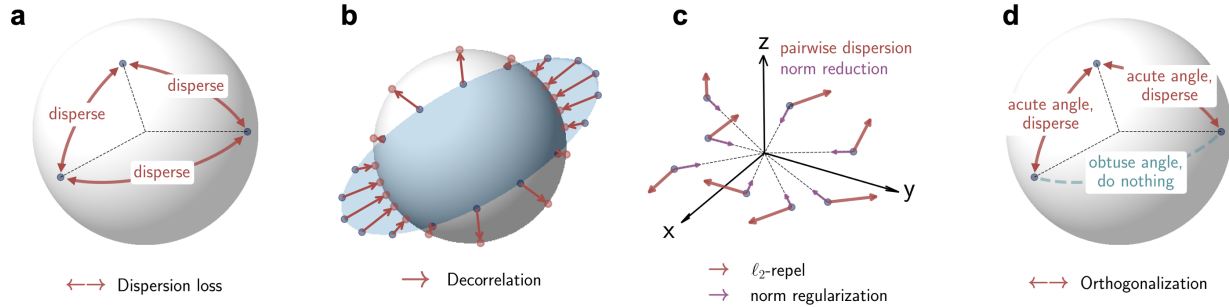


Figure 6. Illustration of how dispersion loss and its alternative formulations promote embedding dispersion. **a.** Dispersion loss enforces uniform angular dispersion by spreading out all pairs along the unit hypersphere. **b.** Decorrelation loss encourages different feature dimensions to remain uncorrelated. **c.**  $\ell_2$ -repel loss increases pairwise Euclidean distance, while the norm regularization prevents unbounded expansion. **d.** Orthogonalization loss spreads out vectors forming acute angles while leaving obtuse ones unchanged.

Table 1. Our dispersion loss and its alternative formulations. Main implementation differences from (Wang & He, 2025) are highlighted in teal and magenta. Including or excluding diagonal terms yields identical gradients and is therefore cosmetic. For dispersion loss and  $\ell_2$ -repel, we adopt the `log-sum-exp` trick for numerical stability, which differs from `log(mean(exp(·)))` only by an additive constant. For  $\ell_2$ -repel, we include a norm regularization term to prevent unbounded expansion of embeddings. For Orthogonalization, the distance margin is fixed to  $\frac{1}{2}$  since we use angular distance, where  $\frac{1}{2}$  corresponds to orthogonality and thus serves as the ideal margin.

	For generative modeling in diffusion-based models (Wang & He, 2025)		For improving generalization of Transformer-based language models (Ours)	
	formulation	term definition	formulation	term definition
<b>Dispersion loss</b>	$\log \mathbb{E}_{i,j} [\exp(-D(z_i, z_j)/\tau)]$	$D(z_i, z_j) = -\text{cossim}(z_i, z_j)$	$\log \sum_{i \neq j} [\exp(-D(z_i, z_j)/\tau)]$	$D(z_i, z_j) = \frac{\arccos(\text{cossim}(z_i, z_j))}{\pi}$
<i>Alternative formulations</i>				
Decorrelation	$\sum_{m,n} \text{Cov}_{mn}^2$	$\text{Cov}^2 = \frac{\mathcal{Z}_c^\top \mathcal{Z}_c}{d-1}, \mathcal{Z}_c = \frac{\mathcal{Z} - \mu_d(\mathcal{Z})}{\sigma_d(\mathcal{Z})}$	$\sum_{m \neq n} \text{Cov}_{mn}^2$	$\text{Cov}^2 = \frac{\mathcal{Z}_c^\top \mathcal{Z}_c}{d-1}, \mathcal{Z}_c = \frac{\mathcal{Z} - \mu_d(\mathcal{Z})}{\sigma_d(\mathcal{Z})}$
$\ell_2$ -repel	$\log \mathbb{E}_{i,j} [\exp(-D(z_i, z_j)/\tau)]$	$D(z_i, z_j) = \ \mathcal{Z}_{i,:} - \mathcal{Z}_{j,:}\ _2^2$	$\log \sum_{i \neq j} [\exp(-D(z_i, z_j)/\tau)] + \lambda_{\text{norm}} \ \mathcal{Z}\ _2^2$	$D(z_i, z_j) = \ \mathcal{Z}_{i,:} - \mathcal{Z}_{j,:}\ _2^2$
Orthogonalization	$\mathbb{E}_{i,j} [\max(0, \epsilon - D(z_i, z_j))^2]$	$D(z_i, z_j) = -\text{cossim}(z_i, z_j)$	$\mathbb{E}_{i \neq j} [\max(0, \frac{1}{2} - D(z_i, z_j))^2]$	$D(z_i, z_j) = \frac{\arccos(\text{cossim}(z_i, z_j))}{\pi}$

In addition to the canonical dispersion loss, we implemented three alternative formulations of the dispersion loss, and evaluated them in our main experiments.

**Decorrelation** The decorrelation formulation minimizes off-diagonal entries of the covariance matrix of embeddings (Table 1 row 2 and Figure 6b). By construction, this loss reduces the correlations between feature dimensions, which indirectly promotes more diverse embedding vector directions in the representation space.

**$\ell_2$ -repel** The  $\ell_2$ -repel formulation directly pushes pairs of embedding vectors apart in the Euclidean space. However, minimizing this objective can be achieved trivially by increasing the embedding norms, as larger magnitudes inflate pairwise distances. To prevent this degeneracy, we include an explicit norm regularization term that constrains unbounded growth (Table 1 row 3 and Figure 6c).

**Orthogonalization** The orthogonalization formulation is similar to the canonical dispersion loss, except that the dispersion vanishes when two vectors are orthogonal to each other (Table 1 row 4 and Figure 6d). The distance margin  $\epsilon$  is naturally set to  $\frac{1}{2}$ , which corresponds to orthogonality under the angular distance.

### 3.4.1. POTENTIAL EFFECTIVENESS ON LARGER MODELS

Having introduced dispersion loss as an explicit mechanism for regulating embedding geometry, we revisit the role of model size in embedding condensation. As shown in Appendix D, increased embedding dimensionality provides a geometric expectation in which randomly oriented vectors show reduced condensation, but this expectation does not guarantee that trained representations fully utilize the available space. As a result, the resistance to embedding condensation empirically observed in larger models may reflect an increased representational capacity rather than an explicit dispersion mechanism. This observation raises the possibility that *our dispersion loss could benefit not only small models but also large models*, which we leave for future investigations.

## 4. Empirical Results

We evaluate our proposed dispersion loss under two training regimes. First, we conduct mid-training experiments (Wang et al., 2025), in which we continue training pre-trained GPT2 and Qwen3 models for an additional 200 M tokens on the `wikitext-103` dataset (Merity et al., 2017). We repeat the experiments under three random seeds, and report the mean and standard deviation for each metric.

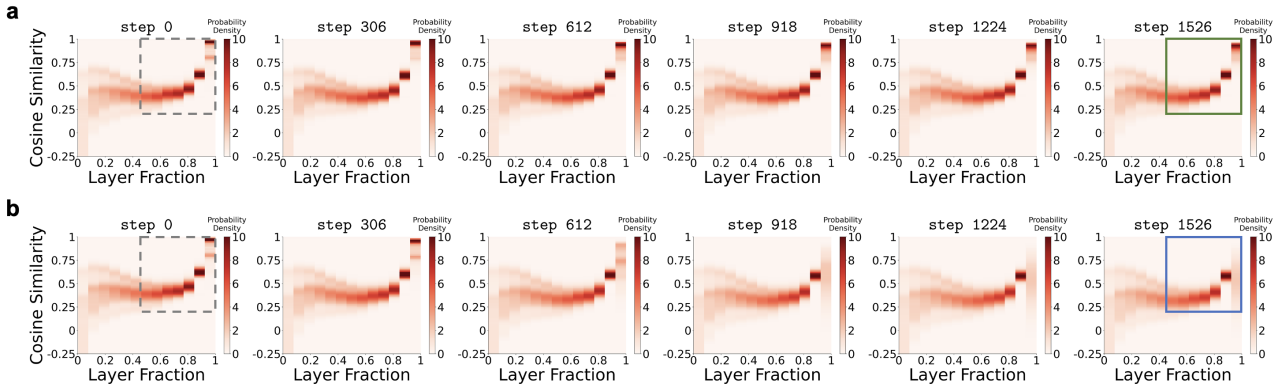


Figure 7. Dispersion loss counteracts the embedding condensation phenomenon. **a.** Starting from condensed embeddings (gray dashed box), mid-training with the default loss has a limited impact (green box). **b.** In contrast, mid-training with our dispersion loss as a regularizer substantially mitigates embedding condensation (blue box).

This setting provides a computationally efficient proof of concept, feasible on a single NVIDIA A100 GPU, enabling controlled ablations and systematic hyperparameter studies.

We then perform full pre-training from scratch to examine the effects of incorporating dispersion loss on the formation of representational geometry. Qwen3 models are trained on the allenai/c4 dataset (Dodge et al., 2021) for 156 B tokens using 640 GPUs.

Experimental details, including training protocols and hyperparameters, are provided in Appendix B.1. The evaluation benchmarks are described in Appendix B.2.

#### 4.1. Dispersion loss counteracts the embedding condensation phenomenon

First, we examine whether dispersion loss can directly counteract the embedding condensation phenomenon.

Using the same heatmap visualization, we observe that pre-trained GPT2 exhibits severe embedding condensation, with pairwise cosine similarities rapidly collapsing toward 1 in deeper layers (Figure 7 column 1). Continuing training with the standard cross-entropy objective provides minimal relief, leaving the overall condensation pattern largely intact (Figure 7a). In contrast, incorporating dispersion loss substantially alters the geometry of token representations. As training progresses, the cosine similarity distributions become more spread out (Figure 7b). These results indicate that dispersion-aware training can restore representational diversity even when applied during mid-training.

#### 4.2. Dispersion loss is effective in mid-training

Next, we evaluate whether the geometric improvements induced by the dispersion loss translate into better downstream performance. We report zero-shot and few-shot results on 10 language understanding benchmarks for models before and after mid-training. Results for the GPT2

and Qwen3 models are reported in Table 2. Besides our dispersion loss and alternative formulations, we compared against other simple ways that could potentially counteract the embedding condensation phenomenon, namely “noisy embedding” which adds noise to the embedding vectors during training, inspired by (Jain et al., 2024), and “active forgetting” which resets the token embedding layer periodically during training, inspired by (Chen et al., 2023).

Models mid-trained with dispersion loss consistently outperform those trained with the default loss, yielding improvements across most tasks and model sizes. Although the absolute gains are modest, they are systematic and consistent, further supporting the link between reduced condensation and improved generalization. Among all compared methods, the proposed dispersion loss delivers the strongest and most stable performance across tasks, achieving the highest average improvement across multiple model sizes (Table 2). Interestingly, active forgetting (Chen et al., 2023) is also fairly competitive, despite not being originally designed for this purpose.

The three alternative formulations of dispersion loss also yield gains in some settings, but are generally less stable or slightly weaker on average, motivating our focus on the canonical dispersion loss in subsequent experiments.  $\ell_2$ -repel is the most volatile, likely due to the fragile balance between pairwise dispersion and norm regularization: if the former dominates, the norm explodes; if the latter dominates, the result is collapse rather than dispersion. Decorrelation is less effective than dispersion loss likely because it penalizes feature-channel correlations and is less direct than reducing angular alignment. Orthogonalization is also less flexible, as it does not incentivize angular separation beyond 90 degrees.

**Differential effect on early and late layers** To further examine whether dispersion loss improves performance via counteracting embedding condensation, we analyze its layer-wise impact. Motivated by the observation that

## Dispersion Loss Counteracts Embedding Condensation and Improves Generalization in Small Language Models

**Table 2.** Using dispersion loss during mid-training improves performance on language tasks. For each base model, the best single-benchmark metrics are displayed in bold, whereas the best average ranks and best average performances are boxed and in bold. We also perform the Student’s  $t$ -tests on the average performances and report the significance level with respect to “ $\mathcal{L}_{\text{train}} + \text{Dispersion loss}$ ”.

Model	Mid-training		Training time A100 hours	Zero-shot						Few-shot				Rank↓	Average↑	$t$ -test
	Train	Loss		ANLI <sub>1v2</sub> ↑	LAMBADA <sub>openai</sub> ↑	OpenbookQA↑	PIQA↑	TruthfulQA↑	WinoGrande↑	ARC <sub>easy</sub> ↑	ARC <sub>challenge</sub> ↑	MedMCQA↑	MMLU↑			
GPT2	✗	—	—	34.4	30.8	15.6	61.6	40.3	52.4	42.0	16.2	25.4	24.8	6.1	34.35	$p < 0.0001$
	✓	$\mathcal{L}_{\text{train}}$	1.122 (1.00×)	34.0 ± 0.3	32.4 ± 0.6	16.4 ± 0.1	62.2 ± 0.7	43.6 ± 0.2	52.8 ± 3.1	41.6 ± 0.0	17.4 ± 0.6	24.2 ± 3.8	24.8 ± 0.4	6.2	34.95 ± 0.11	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \text{noisy embedding}$	1.122 (1.00×)	34.3 ± 0.1	33.0 ± 0.3	16.9 ± 0.4	61.1 ± 0.7	43.9 ± 0.2	53.6 ± 2.0	43.7 ± 0.4	18.0 ± 0.6	21.6 ± 2.5	25.4 ± 0.4	4.3	35.15 ± 0.06	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \text{active forgetting}$	1.127 (1.00×)	<b>35.0 ± 0.3</b>	33.8 ± 1.1	16.6 ± 0.3	61.0 ± 0.0	44.0 ± 0.1	52.2 ± 0.8	<b>44.3 ± 0.4</b>	<b>18.6 ± 1.1</b>	22.4 ± 1.7	<b>25.7 ± 0.2</b>	<b>3.2</b>	35.36 ± 0.15	$n.s.$
	✓	$\mathcal{L}_{\text{train}} + \text{Decorrelation}$	1.221 (1.09×)	35.0 ± 0.0	33.6 ± 0.5	<b>17.0 ± 0.6</b>	60.8 ± 0.6	43.9 ± 0.3	52.0 ± 1.7	43.8 ± 0.5	18.0 ± 1.3	24.4 ± 2.3	25.6 ± 0.2	3.6	35.41 ± 0.06	$n.s.$
	✓	$\mathcal{L}_{\text{train}} + \ell_2\text{-repel}$	1.175 (1.05×)	34.6 ± 0.3	32.8 ± 0.2	16.6 ± 0.2	60.8 ± 0.3	43.9 ± 0.2	<b>54.2 ± 0.8</b>	43.4 ± 0.2	17.6 ± 0.8	24.6 ± 2.0	<b>25.7 ± 0.1</b>	4.1	35.42 ± 0.11	$n.s.$
	✓	$\mathcal{L}_{\text{train}} + \text{Orthogonalization}$	1.210 (1.08×)	34.4 ± 0.1	31.2 ± 1.1	16.8 ± 0.7	61.6 ± 0.0	<b>44.3 ± 0.4</b>	53.4 ± 0.8	44.2 ± 1.0	18.0 ± 0.3	24.8 ± 3.3	25.6 ± 0.2	<b>3.2</b>	35.42 ± 0.19	$n.s.$
✓	$\mathcal{L}_{\text{train}} + \text{Dispersion loss}$	1.176 (1.05×)	34.8 ± 0.0	<b>34.0 ± 0.4</b>	16.6 ± 0.0	<b>62.4 ± 0.3</b>	<b>44.2 ± 0.1</b>	51.6 ± 0.1	42.6 ± 0.4	18.0 ± 0.8	<b>26.6 ± 2.3</b>	25.4 ± 0.2	<b>3.2</b>	<b>35.61 ± 0.12</b>	<i>ref.</i>	
GPT2-m	✗	—	—	33.3	40.4	18.6	66.3	40.1	54.8	50.2	19.9	29.1	25.3	5.4	37.80	$p < 0.0001$
	✓	$\mathcal{L}_{\text{train}}$	2.541 (1.00×)	33.1 ± 0.6	43.2 ± 0.3	19.1 ± 0.0	67.7 ± 0.4	40.5 ± 1.3	55.5 ± 1.0	<b>54.2 ± 2.7</b>	18.9 ± 0.7	28.2 ± 2.3	25.1 ± 0.3	4.5	38.55 ± 0.18	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \text{noisy embedding}$	2.580 (1.01×)	33.3 ± 0.1	44.0 ± 0.3	18.6 ± 0.6	66.7 ± 0.1	<b>44.5 ± 0.6</b>	52.1 ± 1.6	51.7 ± 1.6	19.5 ± 1.8	25.5 ± 0.1	25.8 ± 0.3	4.8	38.16 ± 0.25	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \text{active forgetting}$	2.580 (1.01×)	33.4 ± 0.3	43.0 ± 0.6	18.9 ± 0.1	66.7 ± 0.1	44.3 ± 0.1	50.8 ± 1.4	51.8 ± 2.5	19.8 ± 1.4	26.5 ± 0.1	<b>26.1 ± 0.0</b>	4.4	38.13 ± 0.29	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \text{Decorrelation}$	2.760 (1.09×)	33.4 ± 0.0	<b>45.4 ± 1.3</b>	18.9 ± 0.8	66.6 ± 0.1	42.2 ± 0.1	56.2 ± 0.4	53.8 ± 3.3	18.0 ± 0.8	28.2 ± 3.0	25.4 ± 0.2	4.0	38.81 ± 0.12	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \ell_2\text{-repel}$	2.675 (1.05×)	33.5 ± 0.8	44.2 ± 0.8	18.5 ± 0.6	66.2 ± 0.1	42.4 ± 0.1	54.4 ± 1.1	54.0 ± 2.0	18.6 ± 0.6	28.9 ± 2.1	25.3 ± 0.4	4.7	38.61 ± 0.03	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \text{Orthogonalization}$	2.692 (1.06×)	33.0 ± 0.4	45.2 ± 0.4	18.6 ± 0.6	<b>68.1 ± 0.4</b>	41.4 ± 0.2	55.1 ± 0.6	53.6 ± 2.7	18.4 ± 0.4	29.0 ± 0.4	25.0 ± 0.1	4.8	38.74 ± 0.30	$p < 0.05$
✓	$\mathcal{L}_{\text{train}} + \text{Dispersion loss}$	2.673 (1.05×)	33.6 ± 0.1	45.2 ± 0.8	<b>19.2 ± 0.2</b>	67.5 ± 0.1	43.4 ± 0.2	<b>56.4 ± 1.1</b>	53.8 ± 1.1	<b>20.1 ± 0.9</b>	28.6 ± 0.1	25.7 ± 0.3	<b>2.2</b>	<b>39.35 ± 0.15</b>	<i>ref.</i>	
GPT2-1	✗	—	—	33.6	47.8	19.6	71.6	38.9	58.4	54.0	22.4	26.2	25.6	—	39.81	—
GPT2-x1	✗	—	—	36.4	49.4	22.2	71.8	38.0	57.2	58.0	23.6	27.0	25.2	—	40.89	—
Qwen3-0.6B	✗	—	—	35.0	43.0	19.5	66.5	40.7	60.5	67.5	32.5	26.5	49.4	5.7	44.11	$p < 0.0001$
	✓	$\mathcal{L}_{\text{train}}$	4.676 (1.00×)	32.5 ± 0.3	<b>52.0 ± 0.3</b>	21.5 ± 0.5	<b>67.5 ± 1.4</b>	44.3 ± 0.9	<b>61.0 ± 0.8</b>	68.0 ± 2.8	33.0 ± 1.3	29.5 ± 0.9	<b>50.0 ± 0.1</b>	4.1	45.93 ± 0.37	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \text{noisy embedding}$	4.839 (1.03×)	34.0 ± 0.7	48.8 ± 0.4	20.5 ± 0.7	66.0 ± 0.7	49.4 ± 0.1	57.5 ± 0.7	67.2 ± 1.1	34.5 ± 0.7	35.0 ± 1.4	48.8 ± 0.0	5.3	46.17 ± 0.28	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \text{active forgetting}$	4.829 (1.03×)	35.0 ± 0.7	49.0 ± 4.2	20.0 ± 2.1	67.2 ± 0.4	48.6 ± 0.9	58.8 ± 1.8	69.0 ± 1.4	34.2 ± 1.1	36.8 ± 0.4	49.3 ± 0.0	4.0	46.79 ± 0.27	$p < 0.05$
	✓	$\mathcal{L}_{\text{train}} + \text{Decorrelation}$	4.939 (1.06×)	35.0 ± 0.0	50.5 ± 0.4	19.5 ± 1.8	<b>67.5 ± 1.1</b>	47.0 ± 2.3	59.5 ± 1.8	68.5 ± 3.2	<b>35.0 ± 1.8</b>	40.0 ± 0.4	49.8 ± 0.2	3.5	46.62 ± 0.25	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \ell_2\text{-repel}$	4.864 (1.04×)	33.5 ± 0.0	46.0 ± 2.1	19.0 ± 0.7	66.0 ± 0.4	47.4 ± 0.2	59.5 ± 1.4	69.0 ± 3.2	<b>35.0 ± 1.8</b>	<b>34.0 ± 0.4</b>	46.9 ± 0.2	4.8	46.23 ± 0.04	$p < 0.0001$
	✓	$\mathcal{L}_{\text{train}} + \text{Orthogonalization}$	4.936 (1.06×)	<b>35.5 ± 1.1</b>	50.0 ± 0.7	22.0 ± 0.7	65.5 ± 0.4	49.1 ± 0.1	56.5 ± 2.1	71.5 ± 3.2	33.0 ± 1.1	36.0 ± 0.7	48.9 ± 0.2	4.2	46.80 ± 0.20	$p < 0.05$
✓	$\mathcal{L}_{\text{train}} + \text{Dispersion loss}$	4.878 (1.04×)	<b>35.5 ± 0.8</b>	49.5 ± 0.8	<b>22.5 ± 0.6</b>	<b>65.0 ± 1.6</b>	<b>49.8 ± 1.1</b>	58.5 ± 1.6	<b>72.5 ± 2.3</b>	34.5 ± 0.8	37.5 ± 1.3	49.2 ± 0.2	<b>3.2</b>	<b>47.45 ± 0.16</b>	<i>ref.</i>	
Qwen3-1.7B	✗	—	—	39.5	53.0	29.0	71.5	47.6	60.5	72.0	50.0	45.5	63.1	5.1	53.18	$p < 0.0001$
	✓	$\mathcal{L}_{\text{train}}$	9.148 (1.00×)	40.0 ± 0.9	60.7 ± 0.3	28.5 ± 0.9	74.3 ± 0.6	49.1 ± 0.2	61.3 ± 0.3	71.3 ± 1.4	47.7 ± 1.2	49.7 ± 1.2	63.1 ± 0.0	4.3	54.57 ± 0.17	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \text{noisy embedding}$	9.345 (1.02×)	35.0 ± 0.0	60.5 ± 2.8	28.0 ± 0.0	70.5 ± 0.7	<b>50.1 ± 0.9</b>	60.8 ± 3.2	73.2 ± 0.4	48.5 ± 0.0	48.5 ± 2.1	62.5 ± 0.6	5.0	53.78 ± 0.07	$p < 0.0001$
	✓	$\mathcal{L}_{\text{train}} + \text{active forgetting}$	9.298 (1.02×)	35.2 ± 0.4	59.5 ± 1.4	27.5 ± 0.7	72.5 ± 0.0	49.9 ± 0.2	60.5 ± 1.4	<b>74.2 ± 0.4</b>	47.2 ± 1.1	49.0 ± 1.4	63.2 ± 0.2	5.0	53.89 ± 0.07	$p < 0.0001$
	✓	$\mathcal{L}_{\text{train}} + \text{Decorrelation}$	9.535 (1.04×)	39.5 ± 1.5	60.7 ± 1.0	28.3 ± 1.2	74.8 ± 1.0	49.7 ± 0.5	61.5 ± 0.9	72.7 ± 1.5	49.7 ± 0.8	49.7 ± 1.0	63.5 ± 0.2	3.1	55.01 ± 0.13	$p < 0.01$
	✓	$\mathcal{L}_{\text{train}} + \ell_2\text{-repel}$	9.440 (1.03×)	33.5 ± 0.0	47.8 ± 0.6	23.0 ± 0.0	67.8 ± 1.2	42.1 ± 0.2	57.5 ± 0.5	68.3 ± 1.9	41.3 ± 1.8	38.5 ± 2.6	42.3 ± 0.5	8.0	46.22 ± 0.57	$p < 0.0001$
	✓	$\mathcal{L}_{\text{train}} + \text{Orthogonalization}$	9.565 (1.05×)	<b>40.8 ± 1.3</b>	61.0 ± 0.5	<b>29.2 ± 0.3</b>	74.8 ± 1.3	49.5 ± 0.2	61.3 ± 0.8	72.3 ± 1.9	48.2 ± 1.6	49.5 ± 1.3	63.7 ± 0.0	3.0	55.04 ± 0.26	$p < 0.05$
✓	$\mathcal{L}_{\text{train}} + \text{Dispersion loss}$	9.455 (1.03×)	40.2 ± 2.0	<b>62.2 ± 0.8</b>	29.0 ± 1.0	<b>75.2 ± 0.8</b>	<b>50.2 ± 0.4</b>	<b>62.7 ± 1.2</b>	73.0 ± 1.3	<b>49.3 ± 0.8</b>	<b>51.0 ± 1.8</b>	<b>64.1 ± 0.2</b>	<b>1.7</b>	<b>55.68 ± 0.21</b>	<i>ref.</i>	
Qwen3-4B	✗	—	—	41.5	60.5	27.5	75.0	52.8	67.0	80.0	58.5	55.5	73.1	—	59.14	—
Qwen3-8B	✗	—	—	49.5	67.5	33.0	77.5	54.7	71.0	86.0	63.5	56.0	78.2	—	63.68	—
Qwen3-14B	✗	—	—	54.0	65.5	35.5	77.0	54.7	74.5	86.0	68.5	62.0	81.7	—	65.93	—

**Table 3.** Dispersion loss is more effective when applied to deeper layers, where embedding condensation is more pronounced. Experiments are performed under the GPT2 mid-training setting.

Dispersion location	Zero-shot						Few-shot				Average↑	$t$ -test
	ANLI <sub>1v2</sub> ↑	LAMBADA <sub>openai</sub> ↑	OpenbookQA↑	PIQA↑	TruthfulQA↑	WinoGrande↑	ARC <sub>easy</sub> ↑	ARC <sub>challenge</sub> ↑	MedMCQA↑	MMLU↑		
Layers $\in [1, \frac{L}{2}]$	34.6 ± 0.3	33.0 ± 0.4	16.7 ± 0.6	60.9 ± 0.4	43.3 ± 0.5	<b>52.4 ± 0.3</b>	43.7 ± 0.3	<b>17.9 ± 1.3</b>	21.9 ± 1.0	25.3 ± 0.6	34.96 ± 0.20	$p < 0.05$
Layers $\in [\frac{L}{2}, L]$	<b>34.9 ± 0.4</b>	<b>33.3 ± 0.6</b>	<b>16.9 ± 0.1</b>	<b>61.3 ± 0.8</b>	<b>43.9 ± 0.1</b>	52.2 ± 1.1	<b>43.8 ± 1.2</b>	17.8 ± 0.9	<b>23.7 ± 2.4</b>	<b>25.8 ± 0.1</b>	<b>35.37 ± 0.06</b>	<i>ref.</i>

**Table 4.** Effect of hyperparameters on the dispersion loss. Ablation experiments are performed under the GPT2 mid-training setting.

Loss	Coefficient $\lambda_{\text{disp}}$	Temperature $\tau$	Average (same 10 tasks) ↑
$\mathcal{L}_{\text{train}} + \text{Dispersion loss}$	0.1	1.0	<b>35.61</b>
	0.01	1.0	35.36
	0.5	1.0	35.37
	1.0	1.0	35.29
	0.1	0.1	35.33
	0.1	0.5	35.42
	0.1	2.0	35.27

condensation becomes more pronounced in deeper layers, we apply dispersion loss selectively to either the first half or the last half of the model. We find that applying the loss to the later layers yields greater performance improvements than applying it to the earlier layers (Table 3). This result provides additional evidence that the empirical benefits of dispersion loss arise from mitigating the embedding condensation phenomenon.

**Ablation studies on hyperparameters** We then conduct ablation studies to assess the sensitivity of the proposed

dispersion loss to its main hyperparameters, namely the weighting coefficient  $\lambda_{\text{disp}}$  and the temperature parameter  $\tau$ , using mid-training experiments on GPT2 (Table 4). The average score across the same 10 language understanding benchmarks is reported.

In general, we find that the dispersion loss is relatively robust to the choice of  $\lambda_{\text{disp}}$  and  $\tau$ . Based on these results, we adopt  $\lambda_{\text{disp}} = 0.1$  and  $\tau = 1.0$  as default settings in subsequent experiments.

### 4.3. Dispersion loss is effective in pre-training

Finally, we evaluate the effect of dispersion loss when incorporated throughout full pre-training. Following the insights obtained from the mid-training experiments, we perform pre-training from scratch using the Qwen3-0.6B model on the allenai/c4 corpus with 640 GPUs. Pre-training with dispersion loss leads to an average improvement of +1.17 points in downstream evaluation metrics, including +4.0 points on PIQA, and +7.4 points on TruthfulQA, indicating that encouraging embedding

Table 5. Using dispersion loss during pre-training improves performance on language tasks. Experiments are performed under the Qwen3-0.6B pre-training setting.

Loss	Zero-shot					Few-shot				Average $\uparrow$	
	ANLIR <sub>2</sub> $\uparrow$	LAMBADA <sub>openai</sub> $\uparrow$	OpenbookQA $\uparrow$	PIQA $\uparrow$	TruthfulQA $\uparrow$	WinoGrande $\uparrow$	ARC <sub>easy</sub> $\uparrow$	ARC <sub>challenge</sub> $\uparrow$	MedMCQA $\uparrow$		MLLU $\uparrow$
$\mathcal{L}_{\text{train}}$	<b>34.0</b>	24.0	<b>15.5</b>	64.5	37.8	<b>58.0</b>	<b>41.5</b>	22.0	26.5	24.6	34.84
$\mathcal{L}_{\text{train}} + \text{Dispersion loss}$	32.0	<b>27.5</b>	13.5	<b>68.5</b>	<b>45.2</b>	54.5	40.0	<b>24.5</b>	<b>29.5</b>	<b>24.9</b>	<b>36.01</b> (+1.17)

dispersion during representation formation is beneficial as we anticipated. These gains are achieved over a diverse set of language understanding tasks without any task-specific post-training, suggesting improved generalization. These results demonstrate that incorporating dispersion loss throughout pre-training provides a principled and effective mechanism for counteracting embedding condensation.

We observe that dispersion is generally more helpful for knowledge-heavy or long-context reasoning tasks like TruthfulQA and MedMCQA. In these tasks, we expect the models to benefit from less collapsed, more distinguishable contextual token directions. In contrast, tasks such as WinoGrande rely more on common sense inference and might benefit less from representation dispersion.

## 5. Related Works

**Analyses of condensation** Phenomena consistent with what we term embedding condensation have appeared in prior analyses of Transformer representations, though typically in indirect or task-specific forms. Existing studies have characterized related behaviors using a variety of measures, including output matrix rank (Shi et al., 2022), covariance matrix rank (Li et al., 2025), distance to rank-1 subspaces measured by the Frobenius norm (Dong et al., 2021), spectral bias between high- and low-frequency components (Wang et al., 2022), singular values (Zhang et al., 2025), entropy (Liao et al., 2024; Tokarchuk et al., 2026), spherical variance (Tokarchuk et al., 2026), and the proportion of variance explained by the principal components (Ethayarajh, 2019). These metrics offer complementary views of representation collapse across layers. In this work, we used the layer-by-layer dynamics in pairwise cosine similarity between token embeddings as a direct and interpretable measure for tracking embedding condensation during training. For a broader overview of related representation degeneration phenomena, we refer readers to the survey of (Dovonon et al., 2024).

Previous work has attributed this phenomenon to several factors: oversmoothing induced by layer normalization under analogies between Transformers and graph neural networks (Shi et al., 2022); the effects of specific components such as self-attention and MLPs (Dong et al., 2021); the distribution of embeddings at the infinite depth limit (Geshkovski et al., 2025); and the eigenspectrum of the Transformer update (Dovonon et al., 2024).

Prior attempts to mitigate these effects have largely focused

on architectural or parameterization changes, such as aggregating representations across layers (Shi et al., 2022), reparameterizing updates via eigendecomposition (Dovonon et al., 2024), or rebalancing frequency components (Wang et al., 2022). In contrast, our work targets embedding condensation directly through an explicit representation-level regularization objective.

### Representation shaping via embedding regularization

In representation learning, training objectives are often designed to shape the latent space toward desirable geometric or relational properties. For instance, contrastive learning structures representations by enforcing similarity between positive pairs and separation between negative pairs (Oord et al., 2018; Chen et al., 2020a; Liu et al., 2024; Sun et al., 2025; He et al., 2019; Chen et al., 2020b; Liu et al., 2025a;b; Givechian et al., 2025); geometry-aware regularization methods promote well-behaved latent structures through explicit geometric constraints (Wang & Isola, 2020; Liao et al., 2025; Sun et al., 2024; Verma et al., 2018); REPA (Yu et al., 2024) improves generative quality by aligning representations learned by generative models with those of pretrained understanding models. Closest in spirit, the “diffuse and disperse” framework (Wang & He, 2025) and Tokarchuk et al. (2025) introduce dispersive objectives. In a concurrent work, Li et al. (Li et al., 2026) study model-level dispersion among hidden representations in language models, and train with objectives that increase dispersion to reduce language model perplexity. Our work instead centers on layer-wise embedding condensation as depth and scale evolve and evaluates dispersion-aware training for models on broad understanding benchmarks.

## 6. Conclusion

We presented an empirical study of embedding geometry in Transformer models and identified embedding condensation as a pervasive phenomenon that disproportionately affects smaller models. By introducing a dispersion-aware training objective, we showed that embedding geometry can be directly regulated during training, leading to more diverse embedding vector directions and consistent performance improvements in small language models without scaling up the model size. Our findings suggest that geometric properties of representations are an important and previously underexplored axis for understanding and improving Transformer models. We hope that our work will motivate further investigation into geometry-aware objectives as a complementary approach to scaling in areas including but not limited to language modeling.

## Impact Statement

This paper studies the geometry of token representations in Transformer-based language models and introduces a dispersion-based regularization objective to counteract embedding condensation. The proposed method is a training-time modification that improves generalization in small language models without changing model architecture or increasing parameter count.

As a representation-level technique, this work does not introduce new application domains or deployment mechanisms. Its potential impact is indirect and mediated by downstream use of language models trained with dispersion-aware objectives. While improved generalization in smaller models may reduce reliance on larger models, the societal implications of such improvements depend on the specific tasks, data, and deployment contexts chosen by future users.

We do not anticipate new ethical risks arising specifically from the proposed loss formulation beyond those already associated with language model training and evaluation.

## Acknowledgements

S.K. is funded by the NIH (NIGMSR01GM135929, R01GM130847), NSF CAREER award IIS-2047856, NSF IIS-2403317, NSF DMS-2327211 and NSF CISE-2403317. S.K is also funded by the Sloan Fellowship FG-2021-15883, the Novo Nordisk grant GR112933.

This manuscript was co-authored by Oak Ridge National Laboratory (ORNL), operated by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. Any subjective views or opinions expressed in this paper do not necessarily represent those of the U.S. Department of Energy or the United States Government.

## References

- Bai, J., Bai, S., Chu, Y., Cui, Z., Dang, K., Deng, X., Fan, Y., Ge, W., Han, Y., Huang, F., et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- Bisk, Y., Zellers, R., Bras, R. L., Gao, J., and Choi, Y. Piqa: Reasoning about physical commonsense in natural language. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- Busbridge, D., Shidani, A., Weers, F., Ramapuram, J., Littwin, E., and Webb, R. Distillation scaling laws. In *Forty-second International Conference on Machine Learning*, 2025.
- Cai, T., Fan, J., and Jiang, T. Distributions of angles in random packing on spheres. *The Journal of Machine Learning Research*, 14(1):1837–1864, 2013.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PmLR, 2020a.
- Chen, X., Fan, H., Girshick, R., and He, K. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020b.
- Chen, Y., Marchisio, K., Raileanu, R., Adelani, D., Saito Stenatorp, P. L. E., Riedel, S., and Artetxe, M. Improving language plasticity via pretraining with active forgetting. *Advances in Neural Information Processing Systems*, 36:31543–31557, 2023.
- Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv:1803.05457v1*, 2018.
- Dodge, J., Sap, M., Marasović, A., Agnew, W., Ilharco, G., Groeneveld, D., Mitchell, M., and Gardner, M. Documenting large webtext corpora: A case study on the colossal clean crawled corpus. *arXiv preprint arXiv:2104.08758*, 2021.
- Dong, Y., Cordonnier, J.-B., and Loukas, A. Attention is not all you need: Pure attention loses rank doubly exponentially with depth. In *International conference on machine learning*, pp. 2793–2803. PMLR, 2021.
- Dovonon, G. J., Bronstein, M. M., and Kusner, M. J. Setting the record straight on transformer oversmoothing. *arXiv preprint arXiv:2401.04301*, 2024.
- Dubey, A., Grattafiori, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Fan, A., Goyal, A., Rodriguez, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Ethayarajh, K. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings. *arXiv preprint arXiv:1909.00512*, 2019.
- Geshkovski, B., Letrouit, C., Polyanskiy, Y., and Rigollet, P. A mathematical perspective on transformers. *Bulletin of the American Mathematical Society*, 62(3):427–479, 2025.
- Givechian, K. B., Rocha, J. F., Liu, C., Yang, E., Tyagi, S., Greene, K., Ying, R., Caron, E., Iwasaki, A., and Krishnaswamy, S. Immunostruct enables multimodal deep learning for immunogenicity prediction. *Nature Machine Intelligence*, pp. 1–14, 2025.

- Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, Q., Xu, R., Zhang, R., Ma, S., Bi, X., et al. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638, 2025.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. arxiv e-prints, art. *arXiv preprint arXiv:1911.05722*, 2, 2019.
- Hendrycks, D., Burns, C., Basart, S., Critch, A., Li, J., Song, D., and Steinhardt, J. Aligning ai with shared human values. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021a.
- Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika, M., Song, D., and Steinhardt, J. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021b.
- Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., de Las Casas, D., Hendricks, L. A., Welbl, J., Clark, A., Hennigan, T., Noland, J., Millican, K., van den Driessche, G., Damoc, B., Guy, A., Osindero, S., Simonyan, K., Elsen, E., Rae, J. W., Vinyals, O., and Sifre, L. Training compute-optimal large language models. In *Advances in Neural Information Processing Systems*, 2022.
- Hui, B., Yang, J., Cui, Z., Yang, J., Liu, D., Zhang, L., Liu, T., Zhang, J., Yu, B., Lu, K., et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*, 2024.
- Jain, N., Chiang, P.-y., Wen, Y., Kirchenbauer, J., Chu, H.-M., Somepalli, G., Bartoldson, B. R., Kailkhura, B., Schwarzschild, A., Saha, A., et al. Neftune: Noisy embeddings improve instruction finetuning. In *The Twelfth International Conference on Learning Representations*, 2024.
- Jin, Q., Dhingra, B., Liu, Z., Cohen, W., and Lu, X. Pubmedqa: A dataset for biomedical research question answering. 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)*, pp. 2567–2577, 2019.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- Keisuke, S., Ronan, L. B., Chandra, B., and Yejin, C. Winogrande: An adversarial winograd schema challenge at scale. In *Communications of the ACM*, 2019.
- Kendall, M. G. A new measure of rank correlation. *Biometrika*, 30(1-2):81–93, 1938.
- Li, M. Z., Agrawal, K. K., Ghosh, A., Teru, K. K., Santoro, A., Lajoie, G., and Richards, B. A. Tracing the representation geometry of language models from pretraining to post-training. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
- Li, Y., Li, M., Livescu, K., and Zhou, J. On the predictive power of representation dispersion in language models. In *The Fourteenth International Conference on Learning Representations*, 2026.
- Liao, D., Liu, C., Christensen, B. W., Tong, A., Huguet, G., Wolf, G., Nickel, M., Adelstein, I., and Krishnaswamy, S. Assessing neural network representations during training using noise-resilient diffusion spectral entropy. In *2024 58th Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–6. IEEE, 2024.
- Liao, D., Liu, C., Sun, X., Tang, D., Wang, H., Youtlen, S., Gopinath, S. K., Lee, H., Strayer, E. C., Giraldez, A. J., et al. Rnagenscape: property-guided optimization and interpolation of mrna sequences with manifold langevin dynamics. *arXiv preprint arXiv:2510.24736*, 2025.
- Lin, S., Hilton, J., and Evans, O. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers)*, pp. 3214–3252, 2022.
- Liu, C., Amodio, M., Shen, L. L., Gao, F., Avesta, A., Aneja, S., Wang, J. C., Del Priore, L. V., and Krishnaswamy, S. Cuts: A deep learning and topological framework for multigranular unsupervised medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 155–165. Springer, 2024.
- Liu, C., Liao, D., Parada-Mayorga, A., Ribeiro, A., DiStasio, M., and Krishnaswamy, S. Diffkillr: Killing and recreating diffeomorphisms for cell annotation in dense microscopy images. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2025a.
- Liu, C., Xu, K., Shen, L. L., Huguet, G., Wang, Z., Tong, A., Bzdok, D., Stewart, J., Wang, J. C., Del Priore, L. V., et al. Imageflownet: Forecasting multiscale image-level trajectories of disease progression with irregularly-sampled longitudinal medical images. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2025b.

- Loshchilov, I. and Hutter, F. Sgdr: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2017.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, 2011.
- Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. *International Conference on Learning Representations*, 2017.
- Mihaylov, T., Clark, P., Khot, T., and Sabharwal, A. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *EMNLP*, 2018.
- Minaee, S., Mikolov, T., Nikzad, N., Chenaghlu, M., Socher, R., Amatriain, X., and Gao, J. Large language models: A survey. *arXiv preprint arXiv:2402.06196*, 2024.
- Nie, Y., Williams, A., Dinan, E., Bansal, M., Weston, J., and Kiela, D. Adversarial nli: A new benchmark for natural language understanding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 2020.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- OpenAI. Introducing gpt-5. OpenAI Blog, 2025. URL <https://openai.com/index/introducing-gpt-5/>.
- Pal, A., Umapathi, L. K., and Sankarasubbu, M. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on health, inference, and learning*, pp. 248–260. PMLR, 2022.
- Paperno, D., Kruszewski, G., Lazaridou, A., Pham, Q. N., Bernardi, R., Pezzelle, S., Baroni, M., Boleda, G., and Fernández, R. The lambda dataset, Aug 2016.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2383–2392, 2016.
- Shi, H., Gao, J., Xu, H., Liang, X., Li, Z., Kong, L., Lee, S., and Kwok, J. T. Revisiting over-smoothing in bert from the perspective of graph. *arXiv preprint arXiv:2202.08625*, 2022.
- Spearman, C. The proof and measurement of association between two things. *The American journal of psychology*, 100(3/4):441–471, 1987.
- Sun, X., Liao, D., MacDonald, K., Zhang, Y., Liu, C., Huguet, G., Wolf, G., Adelstein, I., Rudner, T. G., and Krishnaswamy, S. Geometry-aware generative autoencoders for warped riemannian metric learning and generative modeling on data manifolds. *arXiv preprint arXiv:2410.12779*, 2024.
- Sun, X., Liao, D., MacDonald, K., Zhang, Y., Huguet, G., Wolf, G., Adelstein, I., Rudner, T. G., and Krishnaswamy, S. Geometry-aware generative autoencoders for warped riemannian metric learning and generative modeling on data manifolds. In *International Conference on Artificial Intelligence and Statistics*, pp. 1018–1026. PMLR, 2025.
- Tokarchuk, E., Bakker, H. C., and Niculae, V. Keep your distance: learning dispersed embeddings on  $\mathbb{S}_m$ . *arXiv preprint arXiv:2502.08231*, 2025.
- Tokarchuk, E., Nachesa, M. K., Troshin, S., and Niculae, V. Representation collapse in machine translation through the lens of angular dispersion. In *Findings of the Association for Computational Linguistics: EACL 2026*, pp. 2420–2431, 2026.
- Verma, V., Lamb, A., Beckham, C., Najafi, A., Courville, A., Mitliagkas, I., and Bengio, Y. Manifold mixup: learning better representations by interpolating hidden states. *stat*, 1050:4, 2018.
- Wang, P., Zheng, W., Chen, T., and Wang, Z. Anti-oversmoothing in deep vision transformers via the fourier domain analysis: From theory to practice. *arXiv preprint arXiv:2203.05962*, 2022.
- Wang, R. and He, K. Diffuse and disperse: Image generation with representation regularization. *arXiv preprint arXiv:2506.09027*, 2025.
- Wang, T. and Isola, P. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International conference on machine learning*, pp. 9929–9939. PMLR, 2020.
- Wang, Z., Zhou, F., Li, X., and Liu, P. Octothinker: Mid-training incentivizes reinforcement learning scaling. *arXiv preprint arXiv:2506.20512*, 2025.

- Workshop, B., Scao, T. L., Fan, A., Akiki, C., Pavlick, E., Ilić, S., Hesslow, D., Castagné, R., Luccioni, A. S., Yvon, F., et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022.
- Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- Yu, S., Kwak, S., Jang, H., Jeong, J., Huang, J., Shin, J., and Xie, S. Representation alignment for generation: Training diffusion transformers is easier than you think. *arXiv preprint arXiv:2410.06940*, 2024.
- Zhang, G., Yuan, G., Cheng, D., Liu, L., Li, J., and Zhang, S. Mitigating propensity bias of large language models for recommender systems. *ACM Transactions on Information Systems*, 43(6):1–26, 2025.
- Zhang, S., Roller, S., Goyal, N., Artetxe, M., Chen, M., Chen, S., Dewan, C., Diab, M., Li, X., Lin, X. V., Mihaylov, T., Ott, M., Shleifer, S., Shuster, K., Simig, D., Koura, P. S., Sridhar, A., Wang, T., and Zettlemoyer, L. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.

---

# Technical Appendices

---

## Table of Contents

<b>A</b>	<b>Pseudocode for Dispersion Loss</b>	<b>15</b>
<b>B</b>	<b>Experimental Settings</b>	<b>16</b>
B.1	Settings and hyperparameters for training and evaluation . . . . .	16
B.2	Evaluation benchmarks . . . . .	16
<b>C</b>	<b>Additional Results on Embedding Condensation</b>	<b>18</b>
C.1	Embedding condensation results on <code>wikitext-103</code> . . . . .	18
C.2	Embedding condensation results on <code>pubmed_qa</code> , <code>imdb</code> , and <code>squad</code> . . . . .	19
C.3	Embedding condensation results on <code>wikitext-103</code> , evaluated using average <code>cosim</code> of last few layers	19
<b>D</b>	<b>Embedding condensation and embedding dimension</b>	<b>20</b>
<b>E</b>	<b>Proofs of Propositions</b>	<b>21</b>

## A. Pseudocode for Dispersion Loss

---

### Algorithm 1 Dispersion Loss

---

```

# z: Token embeddings
# tau: Temperature
# eps: Numerical stability constant
# thr: Clamp threshold

# [batch size, sequence length, feature dimension]
B, L, F = z.shape

# normalize embeddings along feature dimension
z_norm = z / (linalg_norm(z, dim=2, keepdim=True) + eps)

# cosine similarity matrix with shape [B, L, L]
cossim = z_norm @ transpose(z_norm, dim1=1, dim2=2)

# Clamp to avoid negative inf gradient at the two extrema.
# Also saturate the gradient at the boundary to avoid large gradient.
cossim_clmp = clamp(cossim, -1 + thr, 1 - thr)

# Clamp gives 0 gradient beyond the boundary.
# We force same gradient as the boundary instead.
cossim_clmp = cossim + (cossim_clmp - cossim).detach()

# distance matrix with shape [B, L, L]
D = arccos(cossim_clmp) / pi

# mask out diagonal entries
mask = eye(L).bool()
D = D[:, ~mask]

# compute logit
logit = - D / tau

# log-sum-exp trick for `log(mean(exp(logit)))`
loss = logsumexp(logit + eps, dim=1) - log(L * (L - 1))

# dispersion loss
loss = mean(loss)

```

---

## B. Experimental Settings

### B.1. Settings and hyperparameters for training and evaluation

The settings and hyperparameters are summarized in Table S1.

Table S1. Settings and hyperparameters for training and evaluation.

	Highly controlled experiment (Section 3.1.2)	GPT2 mid-training (Section 4.2)	GPT2-m mid-training (Section 4.2)	Qwen3-0.6B mid-training (Section 4.2)	Qwen3-1.7B mid-training (Section 4.2)	Qwen3-0.6B pre-training (Section 4.3)
<b>Training</b>						
Optimizer	AdamW (Loshchilov & Hutter, 2019)					
Learning rate scheduler	Cosine decay with warmup (Loshchilov & Hutter, 2017)				Linear decay with warmup	
Learning rate	$1 \times 10^{-5}$	$5 \times 10^{-5}$	$5 \times 10^{-5}$	$2 \times 10^{-5}$	$5 \times 10^{-6}$	$5 \times 10^{-5}$
Dataset	codellion/fineweb-edu-1B	Salesforce/wikitext			allenai/c4	
GPU count per job	1	1	1	1	1	640
Training duration (tokens)	1 B	200 M	200 M	200 M	200 M	156 B
Batch size per device	32	32	16	4	2	2
Gradient accumulation step	16	4	8	8	16	1
Context length (tokens)	1024	1024	1024	4096	4096	4096
Effective batch size (sequences)	512	128	128	32	32	1280
Effective batch size (tokens)	524288	131072	131072	131072	131072	5242880
<b>Evaluation</b>						
Zero-shot examples	—	0	0	0	0	0
Few-shot examples	—	1	1	5	5	5
Max generation length (tokens)	—	256	256	1024	1024	1024

### B.2. Evaluation benchmarks

#### ANLI<sub>R2</sub>

The Adversarial Natural Language Inference (ANLI) (Nie et al., 2020) (<https://huggingface.co/datasets/facebook/anli>) is a new large-scale NLI benchmark. The dataset is collected via an iterative, adversarial human-and-model-in-the-loop procedure.

#### LAMBADA<sub>openai</sub>

The LAMBADA benchmark (Paperno et al., 2016) ([https://huggingface.co/datasets/EleutherAI/lambada\\_openai](https://huggingface.co/datasets/EleutherAI/lambada_openai)) is a collection of narrative texts sharing the characteristic that human subjects are able to guess their last word if they are exposed to the whole text, but not if they only see the last sentence preceding the target word. To succeed on LAMBADA, computational models cannot simply rely on local context, but must be able to keep track of information in the broader discourse.

#### OpenbookQA

The OpenBookQA benchmark (Mihaylov et al., 2018) (<https://huggingface.co/datasets/allenai/openbookqa>) is a question-answering dataset that contains questions that require multi-step reasoning, use of additional common and commonsense knowledge, and rich text comprehension.

#### PIQA

The PIQA benchmark (Bisk et al., 2020) (<https://huggingface.co/datasets/ybisk/piqa>) introduces the task of physical commonsense reasoning, a major challenge on the road to true AI-completeness, including robots that interact with the world and understand natural language.

#### TruthfulQA

The TruthfulQA benchmark (Lin et al., 2022) (<https://huggingface.co/datasets/domenicrosati/truthfulqa>) is a benchmark to measure whether a language model is truthful in generating answers to questions. The benchmark comprises 817 questions that span 38 categories, including health, law, finance and politics.

#### WinoGrande

The WinoGrande benchmark (Keisuke et al., 2019) (<https://huggingface.co/datasets/allenai/winogrande>) is a collection of 44k problems formulated as a fill-in-a-blank task with binary options.

#### ARC<sub>easy</sub> and ARC<sub>challenge</sub>

The ARC benchmark (Clark et al., 2018) ([https://huggingface.co/datasets/allenai/ai2\\_arc](https://huggingface.co/datasets/allenai/ai2_arc)) consists of 7,787 genuine grade-school level, multiple-choice science questions, assembled to encourage research in advanced question-answering. The dataset is partitioned into a Challenge Set and an Easy Set, where the former contains only

questions answered incorrectly by both a retrieval-based algorithm and a word co-occurrence algorithm.

#### MedMCQA

The MedMCQA benchmark (Pal et al., 2022) (<https://huggingface.co/datasets/openlifescienceai/medmcqa>) is a large-scale, Multiple-Choice Question Answering (MCQA) dataset designed to address real-world medical entrance exam questions. It contains more than 194k high-quality AIIMS & NEET PG entrance exam MCQs covering 2.4k healthcare topics and 21 medical subjects are collected with an average token length of 12.77 and high topical diversity.

#### MMLU

The MMLU benchmark (Hendrycks et al., 2021b;a) (<https://huggingface.co/datasets/cais/mmlu>) is a multitask dataset consisting of multiple-choice questions on 57 tasks. It spans subjects in the humanities, social sciences, hard sciences, and other areas that are important for some people to learn.

## C. Additional Results on Embedding Condensation

### C.1. Embedding condensation results on wikitext-103

We found consistent trends on embedding condensation across the following model families: GPT2 (Radford et al., 2019), Qwen1 (Bai et al., 2023), Qwen2.5 (Hui et al., 2024), Qwen3 (Yang et al., 2025) and Bloom (Workshop et al., 2022), as shown in Figure S1. This figure shows the raw results we used in Figures 2.

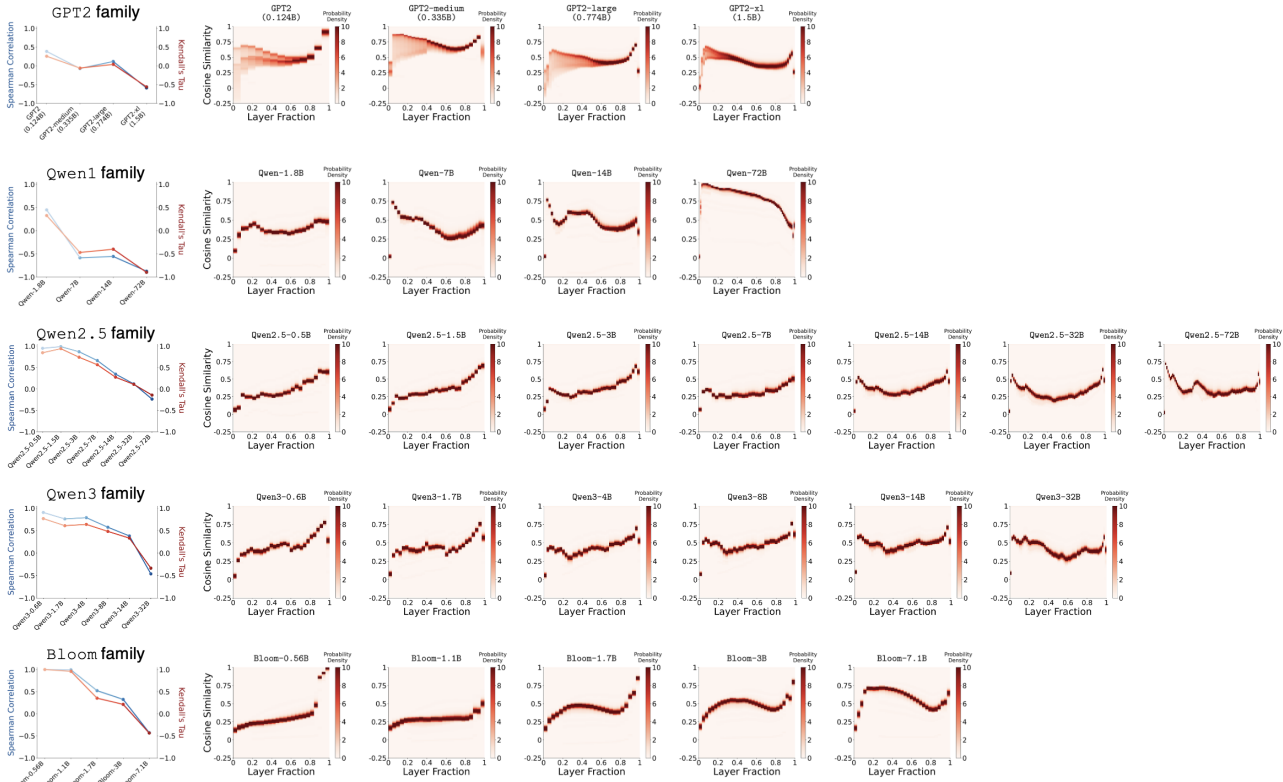


Figure S1. Additional quantitative and qualitative evaluations on GPT2, Qwen1, Qwen2.5, Qwen3 and Bloom families all demonstrate consistent trends that **within each model family, larger models are less susceptible to the embedding condensation phenomenon.**

C.2. Embedding condensation results on pubmed\_qa, imdb, and squad

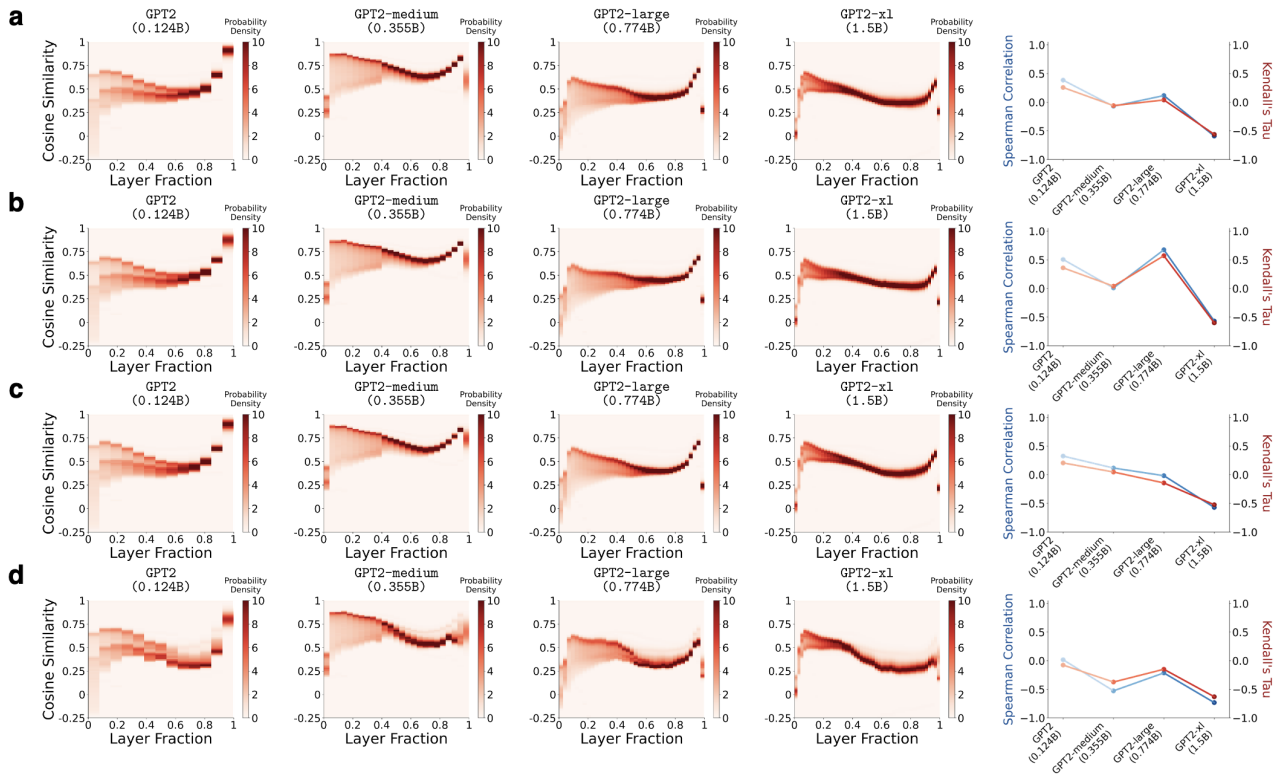


Figure S2. The embedding condensation effect is consistent regardless of the input text dataset. Results are shown for four datasets, namely (a) wikitext, (b) pubmed\_qa, (c) imdb, and (d) squad.

C.3. Embedding condensation results on wikitext-103, evaluated using average cossim of last few layers




Figure S3. The average cosine similarity of the last few layers is also an effective metric that captures the embedding condensation phenomenon. However, the Spearman correlation and Kendall's Tau that we use in this work are able to show the trend more clearly.

## D. Embedding condensation and embedding dimension

It is widely known that when the dimension increases, random vectors are more likely to be orthogonal to each other (Cai et al., 2013). As a result, we ask the following question.

**Q:** To what extent does increased embedding dimensionality alone account for the reduced embedding condensation observed in larger models, independent of training dynamics?

To answer this question, we provide the following theoretical results on the expected values of cosine similarity between two vectors  $x, y \in \mathbb{R}^d$  if the dimension increases from  $d$  to  $D$ . We consider two idealized mechanisms for increasing embedding dimension, which serve as geometric reference cases rather than models of actual training dynamics.

Note that the cosine similarity between  $x$  and  $y$  at dimension  $d$  is  $\text{cossim}_d(x, y) = \frac{x^\top y}{\|x\| \cdot \|y\|}$ . Assume  $\text{cossim}_d(x, y) \geq 0$ .

1. If the dimension growth is achieved by repeating vector entries, the cosine similarity stays the same E.1.
2. If the dimension growth is achieved by padding random entries from the standard normal distribution, the expected cosine similarity is  $\text{cossim}_d(x, y) \cdot \alpha(\|x\|) \cdot \alpha(\|y\|)$ , which is strictly between

$$\begin{aligned} \text{cossim}_d(x, y) \frac{\|x\| \cdot \|y\|}{\sqrt{\|x\|^2 + D - d} \sqrt{\|y\|^2 + D - d}} &= \frac{x^\top y}{\|x\| \cdot \|y\|} \frac{\|x\| \cdot \|y\|}{\sqrt{\|x\|^2 + D - d} \sqrt{\|y\|^2 + D - d}} \\ &= \frac{x^\top y}{\sqrt{\|x\|^2 + D - d} \sqrt{\|y\|^2 + D - d}} \end{aligned}$$

and

$$\begin{aligned} \text{cossim}_d(x, y) \frac{\|x\| \cdot \|y\|}{\sqrt{\|x\|^2 + D - d - 1} \sqrt{\|y\|^2 + D - d - 1}} &= \frac{x^\top y}{\|x\| \cdot \|y\|} \frac{\|x\| \cdot \|y\|}{\sqrt{\|x\|^2 + D - d - 1} \sqrt{\|y\|^2 + D - d - 1}} \\ &= \frac{x^\top y}{\sqrt{\|x\|^2 + D - d - 1} \sqrt{\|y\|^2 + D - d - 1}} \end{aligned}$$

If  $x$  and  $y$  are unit vectors, this implies that the new cosine similarity is between  $\frac{\text{cossim}_d(x, y)}{D-d+1}$  and  $\frac{\text{cossim}_d(x, y)}{D-d}$ .

Taking the GPT2 family as an example, the smallest model GPT2 has an embedding dimension of 768, while the largest model GPT2-x1 has an embedding dimension of 1600. The theoretical results above imply that if the average cosine similarity in GPT2 is  $\text{cossim}(\text{GPT2})$ , then the corresponding value in GPT2-x1 would be

1. equal to  $\text{cossim}(\text{GPT2})$  if assuming repeating entries, or
2. between  $\frac{\text{cossim}(\text{GPT2})}{1600-768+1} = \frac{\text{cossim}(\text{GPT2})}{833}$  and  $\frac{\text{cossim}(\text{GPT2})}{1600-768} = \frac{\text{cossim}(\text{GPT2})}{832}$  if assuming isotropic Gaussian noise entries.

**A:** In high-dimensional spaces, random vectors are nearly orthogonal, implying that increased embedding dimensionality defines a geometric upper bound on achievable dispersion under isotropy. However, this effect reflects a property of isotropic randomness rather than a guarantee of how trained representations utilize the space.

**Implication:** While larger models benefit from a higher-dimensional geometric baseline, embedding dimensionality alone does not ensure well-dispersed representations. This observation raises the possibility that explicit dispersion regularization could further improve embedding geometry even in large models, which we leave for future investigation.

## E. Proofs of Propositions

**Proposition E.1.** *Let  $x, y \in \mathbb{R}^d$  be nonzero vectors. Let  $D = kd$  for some integer  $k \geq 1$ , and define the repeated vectors  $\tilde{x} = (x, x, \dots, x) \in \mathbb{R}^D$  and  $\tilde{y} = (y, y, \dots, y) \in \mathbb{R}^D$ . Then  $\text{cossim}_D(\tilde{x}, \tilde{y}) = \text{cossim}_d(x, y)$ . Consequently, if  $(x, y)$  are random and satisfy  $\mathbb{E}[\text{cossim}_d(x, y)] = c$ , then  $\mathbb{E}[\text{cossim}_D(\tilde{x}, \tilde{y})] = c$ .*

*Proof.* Let  $x, y \in \mathbb{R}^d$  be fixed. Then each of the repeated vectors  $\tilde{x}, \tilde{y} \in \mathbb{R}^D$  consists of  $k$  concatenated copies of  $x$  and  $y$ , respectively.

We compute the inner product:

$$\langle \tilde{x}, \tilde{y} \rangle = \sum_{i=1}^k \langle x, y \rangle = k \cdot \langle x, y \rangle.$$

Next, compute the norms:

$$\|\tilde{x}\|^2 = \sum_{i=1}^k \|x\|^2 = k \cdot \|x\|^2, \quad \|\tilde{y}\|^2 = \sum_{i=1}^k \|y\|^2 = k \cdot \|y\|^2.$$

Thus:

$$\|\tilde{x}\| = \sqrt{k} \cdot \|x\|, \quad \|\tilde{y}\| = \sqrt{k} \cdot \|y\|.$$

Plugging into the cosine similarity:

$$\text{cossim}_D(\tilde{x}, \tilde{y}) = \frac{\langle \tilde{x}, \tilde{y} \rangle}{\|\tilde{x}\| \cdot \|\tilde{y}\|} = \frac{k \cdot \langle x, y \rangle}{(\sqrt{k} \cdot \|x\|)(\sqrt{k} \cdot \|y\|)} = \frac{\langle x, y \rangle}{\|x\| \cdot \|y\|} = \text{cossim}_d(x, y).$$

The identity thus holds pointwise. If  $(x, y)$  are random vectors with  $\mathbb{E}[\text{cossim}_d(x, y)] = c$ , then by linearity of expectation:

$$\mathbb{E}[\text{cossim}_D(\tilde{x}, \tilde{y})] = \mathbb{E}[\text{cossim}_d(x, y)] = c.$$

□

**Proposition E.2.** *Let  $x, y \in \mathbb{R}^d$  be nonzero vectors, let  $m \geq 1$ , and define  $D = d + m$ . Construct padded vectors  $X = (x, \varepsilon)$  and  $Y = (y, \eta) \in \mathbb{R}^D$ , where  $\varepsilon, \eta \sim \mathcal{N}(0, I_m)$  are independent standard Gaussian noise vectors, also independent of  $(x, y)$ . Define:*

$$\alpha(r) := \mathbb{E}_{U \sim \chi_m^2} \left[ \frac{r}{\sqrt{r^2 + U}} \right].$$

Then:

$$\mathbb{E}[\text{cossim}_D(X, Y)] = \text{cossim}_d(x, y) \cdot \alpha(\|x\|) \cdot \alpha(\|y\|). \quad (4)$$

Moreover, for all  $r > 0$ , the function  $\alpha(r)$  satisfies the strict bounds:

$$\frac{r}{\sqrt{r^2 + m}} < \alpha(r) < \frac{r}{\sqrt{r^2 + m - 1}}, \quad (5)$$

with both inequalities strict. As a consequence, the expected cosine similarity after Gaussian padding obeys:

$$\text{cossim}_d(x, y) \cdot \frac{\|x\| \|y\|}{\sqrt{(\|x\|^2 + m)(\|y\|^2 + m)}} < \mathbb{E}[\text{cossim}_D(X, Y)] \quad (6)$$

$$< \text{cossim}_d(x, y) \cdot \frac{\|x\| \|y\|}{\sqrt{(\|x\|^2 + m - 1)(\|y\|^2 + m - 1)}}, \quad (7)$$

again with strict inequalities.

*Proof.* We first compute the cosine similarity in  $D$  dimensions. By definition:

$$\text{cossim}_D(X, Y) = \frac{\langle X, Y \rangle}{\|X\| \cdot \|Y\|}.$$

The inner product expands as:

$$\langle X, Y \rangle = \langle x, y \rangle + \langle \varepsilon, \eta \rangle.$$

Since  $\varepsilon$  and  $\eta$  are independent standard Gaussian vectors in  $\mathbb{R}^m$ , their inner product  $\langle \varepsilon, \eta \rangle$  has mean zero. Specifically:

$$\mathbb{E}[\langle \varepsilon, \eta \rangle] = \sum_{i=1}^m \mathbb{E}[\varepsilon_i \eta_i] = 0,$$

since each  $\varepsilon_i$  and  $\eta_i$  are independent with mean zero.

For the denominator, we write:

$$\|X\|^2 = \|x\|^2 + \|\varepsilon\|^2, \quad \|Y\|^2 = \|y\|^2 + \|\eta\|^2.$$

Denote  $U = \|\varepsilon\|^2$  and  $V = \|\eta\|^2$ . Since  $\varepsilon, \eta \sim \mathcal{N}(0, I_m)$ , we know:

$$U, V \sim \chi_m^2, \quad \text{independently.}$$

Taking expectation of the cosine similarity:

$$\mathbb{E}[\text{cossim}_D(X, Y)] = \mathbb{E} \left[ \frac{\langle x, y \rangle + \langle \varepsilon, \eta \rangle}{\sqrt{\|x\|^2 + U} \cdot \sqrt{\|y\|^2 + V}} \right].$$

By linearity of expectation and the independence of  $\varepsilon, \eta$ , we obtain:

$$\mathbb{E}[\text{cossim}_D(X, Y)] = \langle x, y \rangle \cdot \mathbb{E} \left[ \frac{1}{\sqrt{\|x\|^2 + U} \cdot \sqrt{\|y\|^2 + V}} \right].$$

Now observe that  $\langle x, y \rangle = \text{cossim}_d(x, y) \cdot \|x\| \cdot \|y\|$ . So we substitute:

$$\mathbb{E}[\text{cossim}_D(X, Y)] = \text{cossim}_d(x, y) \cdot \|x\| \cdot \|y\| \cdot \mathbb{E} \left[ \frac{1}{\sqrt{\|x\|^2 + U} \cdot \sqrt{\|y\|^2 + V}} \right].$$

Since  $U$  and  $V$  are independent, the expectation factorizes:

$$\mathbb{E}[\text{cossim}_D(X, Y)] = \text{cossim}_d(x, y) \cdot \left( \mathbb{E} \left[ \frac{\|x\|}{\sqrt{\|x\|^2 + U}} \right] \cdot \mathbb{E} \left[ \frac{\|y\|}{\sqrt{\|y\|^2 + V}} \right] \right).$$

This yields the desired expression with

$$\alpha(r) = \mathbb{E}_{U \sim \chi_m^2} \left[ \frac{r}{\sqrt{r^2 + U}} \right].$$

To prove the lower bound, we consider the function:

$$f(u) = \frac{r}{\sqrt{r^2 + u}}.$$

We compute its second derivative:

$$f''(u) = \frac{3r}{4}(r^2 + u)^{-5/2} > 0 \quad \text{for all } u > 0.$$

Hence,  $f$  is strictly convex on  $(0, \infty)$ . Applying Jensen's inequality:

$$\alpha(r) = \mathbb{E}[f(U)] \geq f(\mathbb{E}[U]) = \frac{r}{\sqrt{r^2 + \mathbb{E}[U]}} = \frac{r}{\sqrt{r^2 + m}}.$$

Because  $f$  is strictly convex and  $U$  is not constant (since  $\text{Var}(U) = 2m > 0$ ), equality cannot occur. Therefore the inequality is strict:

$$\alpha(r) > \frac{r}{\sqrt{r^2 + m}}.$$

To establish the upper bound, observe that  $f(u)$  is strictly decreasing:

$$f'(u) = -\frac{r}{2}(r^2 + u)^{-3/2} < 0.$$

This implies that for any constant  $a < \mathbb{E}[U]$ , we have  $f(U) > f(a)$  with positive probability and  $f(U) < f(a)$  with positive probability, since  $U \sim \chi_m^2$  is supported on  $(0, \infty)$  and not almost surely equal to any fixed value.

Choosing  $a = m - 1 < m$ , and noting that  $f(U) \leq f(m - 1)$  almost surely with strict inequality on a set of positive measure, we conclude:

$$\alpha(r) = \mathbb{E}[f(U)] < f(m - 1) = \frac{r}{\sqrt{r^2 + m - 1}}.$$

This completes the proof of both strict bounds. □