

Understanding Token Probability Encoding in Output Embeddings

Hakaze Cho^{1,☆}, Yoshihiro Sakai¹, Kenshiro Tanaka¹, Mariko Kato¹, Naoya Inoue^{1,2}

¹Japan Advanced Institute of Science and Technology, ²RIKEN

☆Primary Contributor, Correspondence to: yfzhao@jaist.ac.jp

Abstract

In this paper, we investigate the output token probability information in the output embedding of language models. We find an approximate common log-linear encoding of output token probabilities within the output embedding vectors and empirically demonstrate that it is accurate and sparse. As a causality examination, we steer the encoding in output embedding to modify the output probability distribution accurately. Moreover, the sparsity we find in output probability encoding suggests that a large number of dimensions in the output embedding do not contribute to causal language modeling. Therefore, we attempt to delete the output-unrelated dimensions and find more than 30% of the dimensions can be deleted without significant movement in output distribution and sequence generation. Additionally, in the pre-training dynamics of language models, we find that the output embeddings capture the corpus token frequency information in early steps, even before an obvious convergence of parameters starts.

1 Introduction

Modern Language Models (LMs) have two kinds of token embeddings: one is the **input embedding** $E^{(i)}$ located at the earliest layer of LMs, for mapping the input token index into distributed inner representation, the other is the **output embedding** $E^{(o)}$ in the Language Modeling head (LM head), for mapping the hidden state to the predicted probability distribution of the next token in the causal language modeling task.

Since the output embeddings were often *tied* with the input embeddings (Chung et al., 2020; Press and Wolf, 2017), i.e. input embeddings are directly used as the output embedding, the behaviors and features of independent output embeddings are rarely investigated. Along with the scaling of LMs, such embedding tying, which is proven to be harmful to model performance (Chung et al., 2020), is

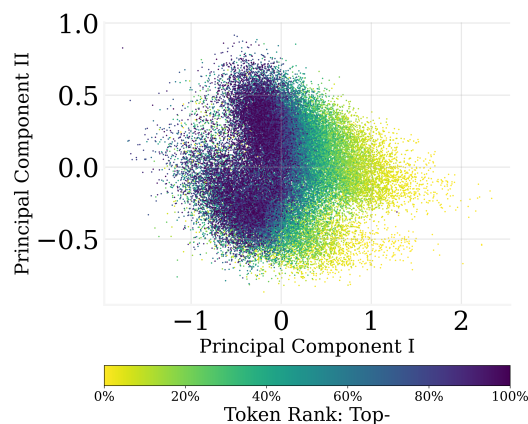


Figure 1: The PCA result of the output embedding vectors of GPT2. Colors refer to the ranking percentile of the averaged output token probability. Output embeddings encode the probabilities linear-likely.

gradually being discarded in modern LMs such as GPT-J (Wang and Komatsuzaki, 2021) and LLaMa 2 (Touvron et al., 2023). This raises attention to the independent output embeddings, and explaining its underlying mechanism can be beneficial to understanding and improving LMs.

The most obvious and expected role of the LM heads, where the output embedding is located, is to map the last hidden states into token probabilities. Therefore, following Kobayashi et al. (2023), who found an encoding of the averaged output token probability distribution in the bias term of the output LM head, this paper also focuses on the averaged output probabilities¹ as an overall representation of LM outputs. We observe there is a linear-like correlation between the output token probabilities and the output embedding as shown in Fig. 1. In §2, our mathematical derivation indicates that softmax-based LM head naturally encodes the output probabilities log-linearly in a common direction of the output embeddings, as long as the output dimension is sufficiently large, and the out-

¹We may omit the “averaged” in the following.

put values are concentrated.

To prove our derivation empirically, we conduct **Multiple Linear Regression (MLR)** on output probabilities against output embeddings, where a strong log-linear correlation is observed, as shown in Fig. 3. Additionally, we find that: **(1)** Almost all directions highly correlated with output probability are the top principal components of the embedding matrix; **(2)** Only a few dimensions of output embedding are highly correlated to output probability.

Based on such findings, to further examine the causality between the encoding and the output probabilities, in §3, we try to steer the output probability by the output embedding. We use a linear and sparse vectorized algorithm, modifying a small portion of dimensions along the detected encoding direction in the output embeddings to steer the probabilities. Our experiments find that: even on embedding-tied models, our steering method has respectable precision and a large available range, stably scaling the probability of tokens up to 20x (both scaling up or down), with little disturbing on the normal inference process of LMs. Moreover, with the encoding direction estimated from few-shot examples, the steering remains precise. Such results suggest that the log-linear correlation has good stability and generalization.

Moreover, the sparsity of the probability encoding demonstrates that most of the dimensions of the output embedding have minor effects on output. Therefore, we try removing these dimensions to prune the parameters in output embedding without obvious harm to the causal language modeling in §4. Our experiments show that more than 30% of the output embedding dimensions can be removed without significant impairment.

Additionally, we use such log-linear encoding to investigate how and when the word frequency information of the training corpus is encoded into the output embedding during the training process. In §5, we find that the frequency encoding occurs at the very early steps in the training dynamics of LMs, even earlier than an obvious convergence trend is observed.

Our contributions can be summarized as:

- We find a log-linear correlation as an encoding of the output probability in a particular common and sparse direction of the output embeddings.
- As a causality examination of the findings, we

steer the averaged output probabilities along the detected encoding direction, and remove dimensions with weak correlation to output probabilities without harm to the LMs.

- Based on the findings, we find that the LMs catch the token frequency in the training corpus at very early training steps.

2 Token Probability Encoding in Output Embedding

In this section, we preliminarily reveal how the output probabilities are encoded in the output embedding by mathematical derivation, then conduct numerical experiments to confirm it empirically.

2.1 Mathematical Log-linear Form

Considering an LM parameterized by θ with vocabulary \mathbb{V} . Denoting the last hidden state w.r.t. input sequence x as h_x , we can describe the output probability of token w with an output embedding $E_w^{(o)}$ as:

$$P_\theta(w|x) = \frac{e^{E_w^{(o)} \cdot h_x}}{\sum_{i \in \mathbb{V}} e^{E_i^{(o)} \cdot h_x}},$$

we have:

$$-\log [P_\theta(w|x)] = -E_w^{(o)} \cdot h_x + \log \left[\sum_{i \in \mathbb{V}} e^{E_i^{(o)} \cdot h_x} \right].$$

When we calculate the averaged output token probability $\alpha_{w, \mathcal{D}, \theta} = \mathbb{E}_{x \in \mathcal{D}} [P_\theta(w|x)]$ of token w on a dataset \mathcal{D} (*detecting dataset*):

$$\begin{aligned} -\log \alpha_{w, \mathcal{D}, \theta} &\approx \mathbb{E}_{x \in \mathcal{D}} [-\log [P_\theta(w|x)]] \\ &= -\mathbb{E}_{x \in \mathcal{D}} [E_w^{(o)} \cdot h_x] + \mathbb{E}_{x \in \mathcal{D}} [\log (\sum_{i \in \mathbb{V}} e^{E_i^{(o)} \cdot h_x})] \\ &= -E_w^{(o)} \cdot \mathbb{E}_{x \in \mathcal{D}} [h_x^{(o)}] + \mathbb{E}_{x \in \mathcal{D}} [\log (\sum_{i \in \mathbb{V}} e^{E_i^{(o)} \cdot h_x})]. \end{aligned} \tag{1}$$

We make a local linear approximation in the approximately equal sign, while we confirm it is precise when the output logits are concentrated in Appendix A.2. Notice that if the LM head is biased, the bias term can be re-constructed equally by fixing one dimension of h_x to 1, w.l.o.g.

We denote $A_{\mathcal{D}} = -\mathbb{E}_{x \in \mathcal{D}} [h_x]$, and $B_{\mathcal{D}} = \mathbb{E}_{x \in \mathcal{D}} \left[\log \left(\sum_{i \in \mathbb{V}} e^{E_i^{(o)} \cdot h_x} \right) \right]$. Here we make another approximation that the $B_{\mathcal{D}}$ is independent to

Model	Decoder-only				Encoder-only		Encoder-decoder	
	GPT2	GPT2-XL	Pythia	GPT-J	BERT-base	BERT-large	BART-base	BART-large
Parameter #	137M	1.6B	2.8B	6B	110M	340M	139M	406M
Embedding tied	✓	✓	✗	✗	✗	✗	✗	✗
Random Adj. R^2	0.001	0.000	0.001	0.000	–	–	–	–
Adj. R^2 on $E_w^{(o)}$	0.892	0.893	0.856	0.882	0.833	0.826	0.998	0.952
Adj. R^2 on $E_w^{(i)}$	(0.892)	(0.893)	0.814	0.658	–	–	–	–

Table 1: The goodnesses (Adj. R^2) of MLR of $-\log \alpha_{w,\mathcal{D},\theta}$ against the output ($E_w^{(o)}$) / input ($E_w^{(i)}$) embedding. **Random Adj. R^2** : The Adj. R^2 of normalized random vector against $E_w^{(o)}$. **Parameter #**: the number of LM parameters. **Embedding tied**: whether the output embedding is tied with the input embedding.

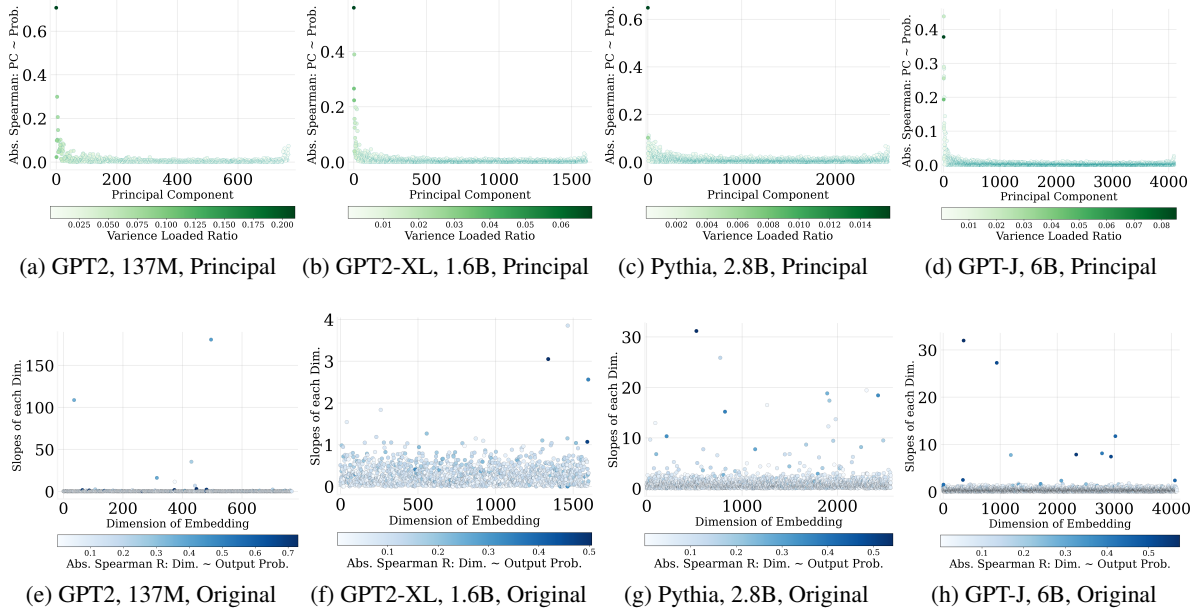


Figure 2: **Only a few directions/dimensions of output embedding are strongly correlated to the output probabilities.** (a-d): **horizontal axis**: the principle components of output embedding, **vertical axis**: absolute Spearman r between the principle and the output probability distribution, **color bar**: the variance ratio loaded in the principal component; (e-h): **horizontal axis**: original dimensions, **vertical axis**: absolute MLR slopes between this dimension and the output probability distribution, **color bar**: the absolute Spearman correlations on the dimension.

$E_w^{(o)}$, which is precise when the logits of the output token w are small, and the vocabulary size $|\mathcal{V}|$ is large. Then, we get an approximated log-linear form between the $\alpha_{w,\mathcal{D},\theta}$ and the $E_w^{(o)}$:

$$-\log \alpha_{w,\mathcal{D},\theta} \approx A_{\mathcal{D}} \cdot E_w^{(o)} + B_{\mathcal{D}}. \quad (2)$$

As a special case, we consider a fixed-to-one dimension in h_x (also in $A_{\mathcal{D}}$) for a biased LM head, where the bias re-constructed in $E_w^{(o)}$ becomes a linear factor of $-\log \alpha_{w,\mathcal{D},\theta}$ with slope 1.

With such a derivation, we find that the phenomenon shown in Fig. 1 is the nature of softmax-based output head if it has plenty of output dimensions, and small and concentrated logits to make the approximations numerically accurate. LMs have

a very wide output space and undergo regularized training, which makes LMs meet the requirements to have an accurate log-linear correlation between output probabilities and output embeddings. That is, the output probabilities are encoded within a common direction of output embedding.

2.2 Empirical Confirmation

We conduct experiments to prove our derivation in Eq. 2 empirically accurate. First, following Kobayashi et al. (2023), we define the detect dataset \mathcal{D} as an 8192-sample of shuffled WIKIDPR dataset (Karpukhin et al., 2020) to calculate the averaged output probabilities. In detail, we input the sampled data points into the model, and aver-

Algorithm 1: Output token probability steering.

$\alpha_{\mathcal{D},\theta} \in \mathbb{R}^{|\mathcal{V}|}$: the output token probability distribution among the vocabulary; $|\mathcal{D}|$: the length of the probability detecting dataset; \mathcal{S} : the element-wise correlation significance (p -value) of MLR; Element-wise calculations: $(\cdot) \times (\cdot)$: multiplication; $(\cdot)^p$: p -th power; $|\cdot|$: absolute value.

Input: Language model $P_\theta(x)$ with output embedding $E^{(o)}$; Probability detecting dataset \mathcal{D} ;
Token index to be steered w ; Expected scale of the steered token’s probability r ;
Steering amount allocation parameter b

Output: Updated LM $P_{\theta'}(x)$

- 1 $\alpha_{\mathcal{D},\theta} \leftarrow |\mathcal{D}|^{-1} \sum_{x \in \mathcal{D}} P_\theta(x)$; /* Calculate the averaged probability distribution */
 - 2 $A_{\mathcal{D}}, \mathcal{S} \leftarrow \text{MLR}(-\log \alpha_{\mathcal{D},\theta}, E^{(o)})$; /* Conduct MLR, get the slope and significance */
 - 3 $\Omega \leftarrow |A_{\mathcal{D}}|^b \times \mathcal{S} \times \|\ |A_{\mathcal{D}}|^b \times \mathcal{S} \|_1^{-1}$; /* Allocate the steering weight to elements */
 - 4 $E_w'^{(o)} \leftarrow E_w^{(o)} - \log(r)\Omega \times A_{\mathcal{D}}^{-1}$; /* Apply the steering to output embedding */
 - 5 **return** Updated LM $P_{\theta'}(x)$ with new output embedding $E_w'^{(o)}$
-

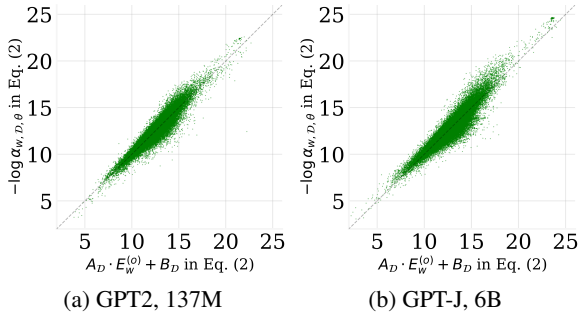


Figure 3: The MLR results on GPT2 and GPT-J.

age the output probability distribution among every time step (aligned with input token) of every input sequence as the token probability distribution of the dataset (see Appendix A.1). Then, we conduct MLR to fit the $A_{\mathcal{D}}$ and the $B_{\mathcal{D}}$.

We conduct such experiment on GPT2, GPT2-XL (Radford et al., 2019), Pythia 2.8B (Biderman et al., 2023), and GPT-J (Wang and Komatsuzaki, 2021). The fitting results are shown in Fig. 3, where good fittings are observed. We also list the adjusted R^2 i.e. the goodness of fitting in Table 1. Surprisingly, both input and output embeddings have strong correlations with the token probabilities. Additionally, as a reference, we provide some rough results for the encoder-only (BERT-base and BERT-large (Kenton and Toutanova, 2019)) and encoder-decoder (BART-base and BART-large (Lewis, 2019)) models, in addition to the decoder-only model².

More interestingly, as shown in Fig. 2, we find

²Note: This paper focuses exclusively on an in-depth discussion of the decoder-only model.

that only the top principal components and a few dimensions in the output embedding are highly correlated to the output token probabilities. That is, such probability encoding is **sparse**. In detail, we calculate the correlation coefficient (absolute Spearman r / MLR slope) between the output embedding’s sorted principle components / original dimensions against the output probabilities, and we find that the overwhelming majority of both kinds of dimensions have poor correlations with the output probabilities.

3 Token Probability Steering on Output Embeddings

Our findings in §2 suggest that the output probabilities are encoded in a common direction in the output embeddings. So, as a causality examination, we steer the output embeddings along this detected direction and confirm that (1) such steering can control the output probabilities accurately on a large scale, implying a robust log-linear token probability encoding. (2) steering based on a few-shot detect dataset \mathcal{D} can still be effective, confirming that such encoding is significant and widely present.

3.1 Algorithm

Based on the fact that the output probabilities are log-linearly encoded in a direction on the output embedding sparsely, we propose an output probability steering algorithm as shown in Algorithm 1. **Step 1:** We first estimate an averaged output probability distribution $\alpha_{\mathcal{D},\theta}$ on a detect dataset \mathcal{D} and conduct MLR to calculate the encoding direction $A_{\mathcal{D}}$ and the correlation significance (p -value used)

	Emb. Tied	b	Reliability	Generalization		Specificity	
			$e_{\text{local}} \downarrow$	$e_{\text{id}} \downarrow$	$e_{\text{ood}} \downarrow$	$d_{\text{KL}}^r \times 10^{-7} \downarrow$	MAUVE \uparrow
- Baseline Values -							
unedited	—	—	1.10 _{1.06}	1.10 _{1.06}	1.10 _{1.06}	0.00 _{0.00}	1.00 _{0.00}
137M, random	✓	2	1.33 _{1.28}	1.33 _{1.28}	1.29 _{1.24}	2.31 _{1.02}	0.96 _{0.01}
137M, shuffled	✓	2	1.18 _{1.07}	1.18 _{1.08}	1.18 _{1.08}	1.64 _{0.48}	0.96 _{0.01}
- Experimental Values -							
137M	✓	2	0.19 _{0.30}	0.20 _{0.32}	0.15 _{0.28}	9.51 _{5.87}	0.96 _{0.01}
1.6B	✓	2	0.64 _{0.67}	0.61 _{0.65}	0.65 _{0.63}	1.51 _{3.66}	0.91 _{0.17}
6B	✗	5	0.31 _{0.43}	0.25 _{0.39}	0.10 _{0.12}	3.64 _{14.65}	— ³

Table 2: Main results on the evaluation of Algorithm 1. **Unedited**: the baseline without any steering. **Random**: the baseline with a random $A_{\mathcal{D}}$. **Shuffled**: the baseline with a shuffled $A_{\mathcal{D}}$ from the calculated value. b : the softness parameter in Algorithm 1. Fine-grained results are reported in the Appendix C.

\mathcal{S} of each element of the common encoding. **Step 2**: Then we assign a steering weight Ω to every dimension in the embedding vector based on the significance of its correlation to the output probabilities, as shown in line 3 of Algorithm 1. We allocate more steering amount to stronger correlations to obtain smaller and more accurate steering⁴. **Step 3**: Given the token index to be steered and the expected steering scale, we calculate the detailed steering amount on each dimension and update it as shown in line 4 of Algorithm 1.

As data costs, our steering only needs a detect set \mathcal{D} and a feed-forward process to calculate the $A_{\mathcal{D}}$ and \mathcal{S} , and we are about to prove that it is stable for the size of \mathcal{D} and consistently precise.

3.2 Experiment Settings

We use the same detect dataset \mathcal{D} as in §2, and a set of scales of $\{1, 1.1, 1.2, 1.5, 2, 5, 10, 20\}$ for both scaling up and down. We randomly select 10 tokens to be steered and conduct experiments on GPT2, GPT2-XL, and GPT-J.

Metrics. We use 3 metrics to test the algorithm.

- **Scale error e** : To describe the precision of the probability steering, given the expected steering scale r and the actual measured steered scale \hat{r} on a test dataset, the scale error is calculated as $e = |\log(r) - \log(\hat{r})|$.
- **KL divergence on the retained token d_{KL}^r** : To investigate the side effect on the retained

³It is difficult to conduct MAUVE experiments of GPT-J on such a repeating scale due to computational costs.

⁴Parameter b is introduced to control the *softness*, or *sparsity* of such allocation, while the algorithm is stable on the parameter as shown in Appendix B. Basically, we suggest that a large b is suitable for a large model.

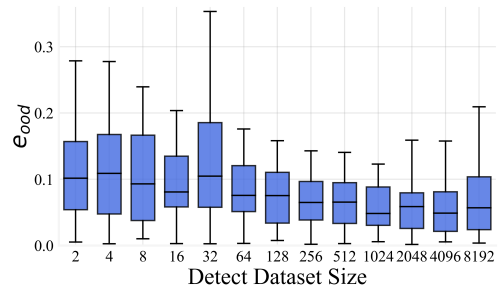


Figure 4: The e_{ood} on detect datasets with various numbers of sentence (averaged token per sentence ≈ 134).

tokens, we calculate KL divergence between the probability distribution before and after steering with the steered token excluded.

- **MAUVE**: To investigate the side effect on text generation, we generate a set of sentences from the steered model, then calculate the MAUVE⁵ with the generated set from the original model (see Appendix A.3).

Evaluations. Applying the aforementioned metrics, following the widely-used aspects of model editing evaluations (Yao et al., 2023), we define our evaluations in 3 categories:

- **Reliability**: The local effectiveness of the steering. We use e_{local} , the scale error with the detect dataset \mathcal{D} as the test dataset.
- **Generalization**: The global effectiveness of the steering. We set multi-level generalization evaluations to ensure robustness: **(i) In-domain Scale Error e_{id}** : the scale error with

⁵Proposed by Pillutla et al. (2021), a measurement of the similarity of two language datasets. The value range is $[0, 1]$, the larger means a greater similarity.




Figure 5: The expected probability scales against the actually steered scales measured in the steered LMs.

the other 8192 samples of WIKIDPR (where the updated parameters are estimated) as the test dataset; and (ii) Out-of-domain Scale Error e_{ood} : the scale error on 2048 samples of BOOKCORPUS (Zhu et al., 2015) as the test dataset.

- **Specificity:** The harmless to the retained part. We use two metrics for this evaluation: (i) the averaged d_{KL}^r on the three data samples (detect set, in-domain, out-of-domain), and (ii) the aforementioned MAUVE.

3.3 Results

The main statistical evaluation results are shown in Table 2. Compared to the baseline values, the results from Algorithm 1 are more accurate, generalizable, and harmless, even on the out-of-domain data. Especially, despite the embedding tying in GPT2 and GPT2-XL, such a steering algorithm is still accurate and harmless, which demonstrates that the log-linear probability encoding is *orthogonal* to the possible semantical encoding in the output and also input embedding.

Wide-scale stable: Large-scaled probability steering is supported by a global log-linear pattern. The correlation of actually steered scales of token probability against the expected scales is shown in Fig. 5. The steering remains accurate even on a large scale of up to 20x. This indicates that the algorithm and the log-linear encoding are wide-ranged, not only locally effective, i.e. the log-linear encoding is a widely stable common essence.

Few-shot generalizable: Encoding remains distinct even by an $A_{\mathcal{D}}$ estimated by few-shot corpus. Instead of using the 8192 examples for the detecting dataset \mathcal{D} , we try various numbers of examples to test the generalization of the encoding on

GPT2. As shown in Fig. 4 and 16, we find that even 2 examples can produce a distinct averaged probability distribution for precise probability steering. This result strengthens the significance and generalization of our findings, that is, effective encodings can be detected in a small set of samples.

These steering results demonstrate a clear causality between the log-linear encoding and the token probabilities, further than a simple statistical covariance relationship, and also demonstrate a wide-scale stability and generalization of our algorithm, while reflecting the same attributes of token probability encoding in output embedding. We can confirm that such log-linear encoding of token probabilities is an inherent attribute of output embedding, which (1) has wide-scale linearity, allowing a large-scaling probability steering, and (2) is common among tokens, so only a small number of samples are needed to find an accurate encoding. By our steering experiments, the properties of probability encoding are strengthened.

4 Removing Dimensions with Weak Probability Encoding

Moreover, based on the sparsity of probability encoding as shown in Fig. 2, we can infer that a large number of dimensions in the output embeddings are less effective for next token decoding. So we try to reduce the less related dimensions towards a lightweight output head.

4.1 Method & Experiment Settings

First, we assign the aforementioned absolute MLR slopes as shown in (e-h) of Fig. 2 as the importance (or, *saliency*) to each dimension of output embedding. Then, we remove the dimensions in ascending order of such weight, i.e. we remove the identified less important dimension early. In detail, as a prototype setting in the laboratory, we only zero out the dimensions in the embedding matrix, while the dimensions of the attention mapping matrix, the feed-forward layer, and the dimensions of the last hidden state corresponding to the removed dimensions can also be removed equally.

Experiment Settings. We use the same MLR settings as §2 on GPT2, Pythia 2.8B, and GPT-J. As metrics for the causal language modeling quality, we test the MAUVE of the pruned model similarly to §3.2, and test the KL divergence of the averaged probability distribution against the original model

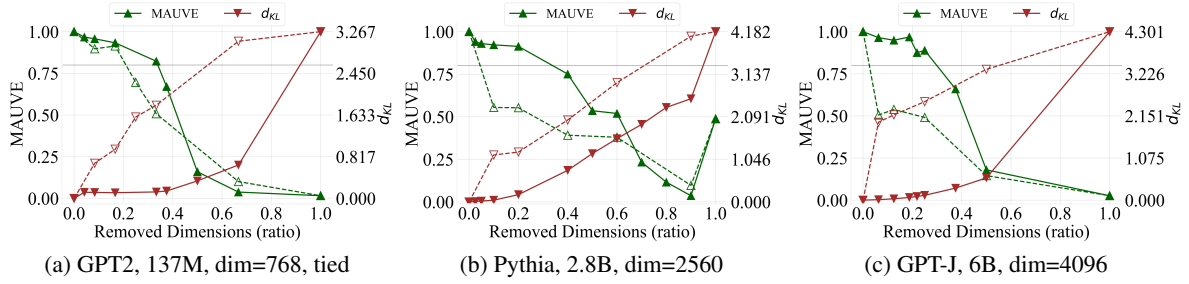


Figure 6: The MAUVE(\uparrow) and KL divergence(\downarrow) of output token probability distribution d_{KL} against original distributions w.r.t. the dimension removing ratio on the output embedding. **Solid curves**: results of removing dimensions from the least important ones to the most important ones; **Dashed curves**: adversarial controlling experiment, removing dimensions reversely.

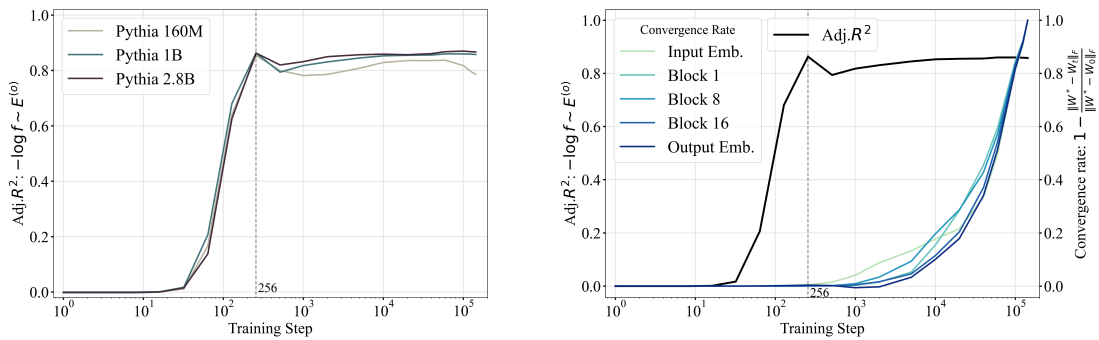


Figure 7: The training dynamics of Pythia. **Left**: the MLR goodness (adjusted R^2) of the negative logarithm of corpus frequency against the output embedding. **Right**: on the Pythia-1B, the MLR goodness in the left figure, and the convergence rate of various representative blocks. Notice the horizontal axis is logarithmically scaled.

on the out-of-domain dataset mentioned before. As an adversarial controlling experiment, we remove dimensions in *descending* order of importance, i.e. we remove the more important dimension early, reverse to the original settings.

4.2 Result

The results are shown in Fig. 6. With the ascending removing order (the positive experiment, solid curve), at the beginning of the removal, both metrics deteriorate slightly until around 50% of the dimensions are removed. Regarding MAUVE > 0.80 as a threshold, we can confirm that 30% \sim 40% of the dimensions in the output embedding can be removed without significant impairment of the causal language modeling quality of LMs. In contrast, in the adversarial experiment with a descending removing style (dashed curve), both metrics deteriorate sharply at the beginning of removing. Such results suggest that the weights from the MLR are faithful importance metrics.

Noticeably, our method also works on the tied models (Fig. 6 (a)), i.e. removing unnecessary di-

mensions concurrently from the input embedding and the output embedding, the metrics remain at a considerable level.

Consistent with previous works (Kovaleva et al., 2021; Timkey and van Schijndel, 2021; Gordon et al., 2020), our results identify the anisotropy in the importance among the dimensions of the output embedding. Moreover, we confirm by dimensional ablation that the MLR results of the log-linear correlation are a good metric of importance, i.e. saliency score. This can be a new paradigm of saliency if a direct one-step statistical model (like MLR in our paper) can be estimated between the features and the model output.

5 Output Embedding Learns Corpus Frequency in Early Training Steps

Additionally, the findings in this paper inspire us to utilize such log-linear correlation for detecting the encoding of *corpus token frequency* in the output embedding during the pre-training of LMs, for a closer observation of the pre-training dynamics. Intuitively, since the model can refract overall output

probability distribution in the output embedding matrix, it should be forced to produce the same output distribution with the corpus token frequency by the training objective. Therefore, when a log-linear encoding of *token frequency* of the training corpus is observed in the output embedding, we can confirm that the output embedding learns the token frequency information of the corpus.

We use the Pythia suite (Biderman et al., 2023), where sequences of the model intermediate training checkpoints are accessible. We estimate the *token frequencies* of the pre-training corpus PILE (Gao et al., 2020; Biderman et al., 2022) by sampling 14.3B tokens, then conduct MLR on the negative logarithm of the estimated **token frequencies** w.r.t. the output embeddings on each training checkpoint, with various training steps. The results of MLR fitting goodness against training steps are shown in the left part of Fig. 7, where we can confirm an effective encoding since the very early steps of the training process, and larger models have slightly better fitting goodness but almost no difference in the timing of emergence.

Moreover, we want to know whether such a phenomenon is a subsidiary or subsequent effect of the convergence of the parameters. We use the convergence rate following the Chen et al. (2022) to describe the actual training completion: denoting the *trained* parameter matrix as θ^* , the *initialized* parameter matrix as θ_0 , and the parameter matrix at step t as θ_t , the convergence rate at step t can be written as $(1 - \|\theta^* - \theta_t\|_F / \|\theta^* - \theta_0\|_F)$. We visualize the convergence rate of the input and output embeddings and the query-key-value mapping matrix of the multi-head attention blocks on Pythia-1B, as shown in the right part of Fig. 7. We find that the log-linear correlation occurs even earlier than an obvious convergence trend is observed by the convergence rate, and the appearance of such correlation happens to be the starting point of the convergence process of the model. We infer that the output embedding should learn the coarse-grained output pattern earlier than the semantics details.

Additionally, we initially find that each layer of the transformer appears to have a uniform convergence curve, instead of an obvious deeper-slower pattern found by Chen et al. (2022). Especially, the input embedding and output embedding have almost overlapping training curves. We speculate that this is the effect of such full-residual connection networks, which makes the layering of the network inconspicuous during the gradient descent.

6 Conclusion and Discussion

Conclusion. In this paper, based on the observation of a linear-like correlation between the output token probability and the output token embeddings shown in Fig. 1, we derive an approximate log-linear correlation from the nature of softmax-based output head with a large output space and concentrated output value, i.e. the output token probabilities are encoded in a common direction of output embeddings. To test the causality of such correlation, we steer the token probabilities along such encoding direction with high accuracy, stability, and generalizability. Then, based on the sparsity of the encoding, we can distinguish the contributions of dimensions for the output probability of the model, which allows us to remove the determined non-contributing dimensions, and no critical deterioration is found. Finally, based on the findings, we find that the LMs catch the token frequency in training data at very early steps in the training process log-linearly, even earlier than an obvious convergence trend is observed. This paper reveals the inner mechanism of LM heads on the causal language modeling task and helps understand the global principles and training dynamics of LMs.

Comparing to previous works. Previous work about analyzing LM heads was conducted by Kobayashi et al. (2023), where a correlation between the output token probability and bias term in the LM head was found. They declared that the bias term is a projection to extract the probability from the output embedding, but almost no more discussion about the embedding matrix, which is the major component of the LM head. Making up for their work is one of the original motivations for our work. More discussions about related works can be found in §7.

Demonstrations towards application. Our output probabilities steering algorithm on the LM head reveals the possibility of model steering on only the output probability rather than in the hidden states or the lower layers (Dai et al., 2022; Meng et al., 2022), which is more concise, and easy to explain. Some global toxicity generation (Gehman et al., 2020), or biases in some application scenario (Fei et al., 2023) can be suppressed by such a method, but as we will mention in the Limitations, it is elegant but not engineering-oriented. Moreover, the dimension compression method in this paper can

be an easy-to-use and harmless inference-time acceleration. Notice that we have supervision on such dimension compression by the MLR slopes, so it can be more accurate than the previous random or unsupervised pruning (Gordon et al., 2020).

Towards a new saliency score of output embeddings. Saliency score (Zhao et al., 2024; Sun et al., 2021) is to weight a component (feature or parameter) in a deep learning model by its contribution. In this paper, we find that the log-linear token probability encoding works like a saliency score towards the output embedding, and build a log-MLR model to assign saliency scores to the parameters. Such a posterior statistical method can be a new paradigm of model-based saliency (Dabkowski and Gal, 2017), if one-step correlations can be detected between the output and components of the neural network, a closed-form saliency model can be proposed instead of a universal statistical model.

7 Related Works

As mentioned before, Kobayashi et al. (2023) mainly found a correlation between the output token probability and the bias term in the LM head and tried to remove this bias towards more diversified text generation. However, they didn't analyze the output embedding matrix, which has the most parameters in the LM head, and this paper completes their research.

Geometry of Input Embedding. As a similar research object with the output embedding, it was found that the word embeddings in LMs, as well as the hidden states, are anisotropy (Mu and Viswanath, 2018; Ethayarajh, 2019; Gao et al., 2018), i.e., these vectors share a common radial bias. Such anisotropies hurt the expressiveness of word embeddings, and the word frequency in the corpus may be an inducement (Mu and Viswanath, 2018; Valentini et al., 2023). Also, some efforts tried to remove the harmfulness of anisotropies and towards isotropy word embeddings (Mu and Viswanath, 2018; Gong et al., 2018). These works are based on input embedding, while our paper is on output embedding. Although we can confirm that the input and output embeddings act similarly, they are still completely different components of untied LMs. So the existing conclusions on input embeddings cannot overwrite our work.

Embedding Tying in LMs. LMs in previous generations often have a shared output embedding

from the input embedding, such as BERT (Kenton and Toutanova, 2019), RoBERTa (Liu et al., 2019), GPT2 (Radford et al., 2019), etc. That is, the LM head maps the last hidden state to the token probability by dot-producting the input embedding. Such a paradigm is recommended by Press and Wolf (2017) for fewer parameters. And also refuted by Chung et al. (2020) for a harmfulness to expressiveness. Such a paradigm is being deprecated currently, but the model behavior with and without embedding tying is still interesting to analyze.

(Language) Model Editing and Model Pruning.

Recent Large LMs are expensive to fine-tune or retrain, so there are many model editing methods to control the output of LMs (Yao et al., 2023). Current LM parameter editing methods are mainly oriented to entity relationship editing, where they locate some parameters with correlations to the entities, and interference is applied on such parameters (Dai et al., 2022; Meng et al., 2022). Moreover, vectorized methods are also proposed with the arithmetic of parameter vectors with editing information (Ilharco et al., 2022; Ortiz-Jimenez et al., 2024). As a specific and extreme scenario of model editing, research on *model pruning*, similar to our dimension removing is also proposed in current years (Zhu et al., 2023; Frantar and Alistarh, 2023; Kovaleva et al., 2021; Timkey and van Schijndel, 2021; Gordon et al., 2020). These pruning are usually unsupervised, where our dimension removing can be a new practice in the supervised pruning paradigm.

Training Dynamics (of LMs). Investigating what is happening in the training process of language models and other deep learning models is an attractive research topic. Many works about training trajectory (Kalra and Barkeshli, 2024; Jastrzębski et al., 2018; Lewkowycz et al., 2020), early period training behaviors (Frankle et al., 2019; Kalra and Barkeshli, 2024; Achille et al., 2018), loss landscape (Neysshabur et al., 2020; Keskar et al., 2017; Li et al., 2018), and “knowledge” earned in different stages of training (Tirumala et al., 2022; Liu et al., 2021) have been done. Differences in training speed among various network layers (the deeper-slower pattern) have been discovered by Chen et al. (2022), while in this paper, following their method, we don't find a similar pattern.

8 Limitations

Although we declare that the probability steering algorithm proposed in §3 is only an experimental method for investigation, we acknowledge that it is elegant but not practical. It can never be faster, more accurate, and more harmless than a filter on the output head (Guo et al., 2017). Future works can be focused on a *local* or *directional* probability editing method, limiting the detecting dataset \mathcal{D} , and only editing the probability on specific input prefix to resolve a controlled-generation task.

The dimension-reducing method in §4 may lead LMs to be unavailable on other tasks depending on the last hidden state, such as sentence summarizing vectors encoding, etc. However, we can always keep the original checkpoint to restore these additional abilities of LM heads easily.

Furthermore, despite our efforts, we cannot confirm the source of the sparsity of the probability encoding mentioned in Fig. 2. Future works can be focused on the detailed training dynamics to trace such a sparsity.

The findings in this paper seriously depend on the properties of the last hidden state of LMs. Although the layer normalization (Ba et al., 2016; Vaswani et al., 2017) in current Transformer-based LMs provides some intuitive assurance for the stability and consistency of the last hidden state, further discussion is still needed to confirm the homogeneity or heterogeneity of the models’ intrinsic properties to explain the differences between different models in the token probability encoding phenomena investigated in the paper (e.g. the reason of our method perform better on GPT-J than GPT2-XL in Table 2, or, the reason of the sparsity of GPT2-XL is weaker than all the models we investigated in Fig. 2), to establish a connection with the essential properties of LMs. Also, we should examine the distribution of the last hidden state so that the output probability to find how the accuracy of the averaged output probabilities can reflect the individual output probability.

Acknowledgements

This work is supported by the Nakajima Foundation.

The authors would like to thank Mr. Yuxuan Wang at Beijing Institute of Technology for his proofreading and constructive criticism.

References

- Alessandro Achille, Matteo Rovere, and Stefano Soatto. 2018. [Critical learning periods in deep networks](#). In *International Conference on Learning Representations*.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. [Layer normalization](#). *arXiv preprint arXiv:1607.06450*.
- Stella Biderman, Kieran Bicheno, and Leo Gao. 2022. [Datasheet for the pile](#). *arXiv preprint arXiv:2201.07311*.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. 2023. [Pythia: A suite for analyzing large language models across training and scaling](#). In *International Conference on Machine Learning*, pages 2397–2430. PMLR.
- Yixiong Chen, Alan Yuille, and Zongwei Zhou. 2022. [Which layer is learning faster? a systematic exploration of layer-wise convergence rate for deep neural networks](#). In *The Eleventh International Conference on Learning Representations*.
- Hyung Won Chung, Thibault Fevry, Henry Tsai, Melvin Johnson, and Sebastian Ruder. 2020. [Rethinking embedding coupling in pre-trained language models](#). In *International Conference on Learning Representations*.
- Piotr Dabkowski and Yarin Gal. 2017. [Real time image saliency for black box classifiers](#). *Advances in neural information processing systems*, 30.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. [Knowledge neurons in pretrained transformers](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502.
- 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.
- Yu Fei, Yifan Hou, Zeming Chen, and Antoine Bosselut. 2023. [Mitigating label biases for in-context learning](#). In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Jonathan Frankle, David J Schwab, and Ari S Morcos. 2019. [The early phase of neural network training](#). In *International Conference on Learning Representations*.

- Elias Frantar and Dan Alistarh. 2023. [Sparsegpt: Massive language models can be accurately pruned in one-shot](#). In *International Conference on Machine Learning*, pages 10323–10337. PMLR.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tiejun Liu. 2018. [Representation degeneration problem in training natural language generation models](#). In *International Conference on Learning Representations*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. [The pile: An 800gb dataset of diverse text for language modeling](#). *arXiv preprint arXiv:2101.00027*.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. [RealToxicityPrompts: Evaluating neural toxic degeneration in language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, Online. Association for Computational Linguistics.
- Chengyue Gong, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2018. [Frage: Frequency-agnostic word representation](#). *Advances in neural information processing systems*, 31.
- Mitchell Gordon, Kevin Duh, and Nicholas Andrews. 2020. [Compressing bert: Studying the effects of weight pruning on transfer learning](#). In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 143–155.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. [On calibration of modern neural networks](#). In *International conference on machine learning*, pages 1321–1330. PMLR.
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. 2022. [Editing models with task arithmetic](#). In *The Eleventh International Conference on Learning Representations*.
- Stanisław Jastrzębski, Zachary Kenton, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. 2018. [On the relation between the sharpest directions of dnn loss and the sgd step length](#). In *International Conference on Learning Representations*.
- Dayal Singh Kalra and Maissam Barkeshli. 2024. [Phase diagram of early training dynamics in deep neural networks: effect of the learning rate, depth, and width](#). *Advances in Neural Information Processing Systems*, 36.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Nitish Shirish Keskar, Jorge Nocedal, Ping Tak Peter Tang, Dheevatsa Mudigere, and Mikhail Smelyanskiy. 2017. [On large-batch training for deep learning: Generalization gap and sharp minima](#). In *5th International Conference on Learning Representations, ICLR 2017*.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2023. [Transformer language models handle word frequency in prediction head](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4523–4535, Toronto, Canada. Association for Computational Linguistics.
- Olga Kovaleva, Saurabh Kulshreshtha, Anna Rogers, and Anna Rumshisky. 2021. [Bert busters: Outlier dimensions that disrupt transformers](#). *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*.
- M Lewis. 2019. [Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). *arXiv preprint arXiv:1910.13461*.
- Aitor Lewkowycz, Yasaman Bahri, Ethan Dyer, Jascha Sohl-Dickstein, and Guy Gur-Ari. 2020. [The large learning rate phase of deep learning: the catapult mechanism](#). *arXiv preprint arXiv:2003.02218*.
- Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. 2018. [Visualizing the loss landscape of neural nets](#). *Advances in neural information processing systems*, 31.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Zeyu Liu, Yizhong Wang, Jungo Kasai, Hannaneh Hajishirzi, and Noah A Smith. 2021. [Probing across time: What does roberta know and when?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 820–842.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in GPT](#). *Advances in Neural Information Processing Systems*, 35.
- Jiaqi Mu and Pramod Viswanath. 2018. [All-but-the-top: Simple and effective postprocessing for word representations](#). In *International Conference on Learning Representations*.
- Behnam Neyshabur, Hanie Sedghi, and Chiyuan Zhang. 2020. [What is being transferred in transfer learning?](#) *Advances in neural information processing systems*, 33:512–523.

- Guillermo Ortiz-Jimenez, Alessandro Favero, and Pascal Frossard. 2024. [Task arithmetic in the tangent space: Improved editing of pre-trained models](#). *Advances in Neural Information Processing Systems*, 36.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. [Mauve: Measuring the gap between neural text and human text using divergence frontiers](#). *Advances in Neural Information Processing Systems*, 34:4816–4828.
- 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.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Xiaofei Sun, Diyi Yang, Xiaoya Li, Tianwei Zhang, Yuxian Meng, Han Qiu, Guoyin Wang, Eduard Hovy, and Jiwei Li. 2021. [Interpreting deep learning models in natural language processing: A review](#). *arXiv preprint arXiv:2110.10470*.
- William Timkey and Marten van Schijndel. 2021. [All bark and no bite: Rogue dimensions in transformer language models obscure representational quality](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4527–4546.
- Kushal Tirumala, Aram Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. 2022. [Memorization without overfitting: Analyzing the training dynamics of large language models](#). *Advances in Neural Information Processing Systems*, 35:38274–38290.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Francisco Valentini, Juan Sosa, Diego Slezak, and Edgar Altszyler. 2023. [Investigating the frequency distortion of word embeddings and its impact on bias metrics](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 113–126.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *Advances in neural information processing systems*, 30.
- Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2023. [Editing large language models: Problems, methods, and opportunities](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10222–10240.
- Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2024. [Explainability for large language models: A survey](#). *ACM Transactions on Intelligent Systems and Technology*, 15(2):1–38.
- Xunyu Zhu, Jian Li, Yong Liu, Can Ma, and Weiping Wang. 2023. [A survey on model compression for large language models](#). *arXiv preprint arXiv:2308.07633*.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *The IEEE International Conference on Computer Vision (ICCV)*.

A Calculation Details

A.1 Calculation of Averaged Token Probability

Given a dataset $\mathcal{D} = \{x_i\}_{i=1}^m$, we input each x_i into LMs in a teacher forcing style. Denote the length of x_i as l_i , we can get output token probability distribution on each time step (noted as j) $\alpha_{\mathcal{D},\theta,i,j} \in \mathbb{R}^{|\mathcal{V}|}$ of an amount of l_i .

We average all the $\alpha_{\mathcal{D},\theta,i,j}$ on every i and j , and get the averaged token probability distribution $\alpha_{\mathcal{D},\theta}$ on dataset \mathcal{D} .

A.2 Error Analysis of Eq. 1

In Eq. 1, we do a local linear approximation as:

$$-\log \alpha_{w,\mathcal{D},\theta} = -\log \mathbb{E}_{x \in \mathcal{D}} [\mathbb{P}_{\theta}(w|x)] \\ \approx \mathbb{E}_{x \in \mathcal{D}} [-\log [\mathbb{P}_{\theta}(w|x)]] .$$

That is, given a set of $\{p_i\}_{i=0}^n$, where $\forall i, p_i > 0$ we approximate that $\log \mathbb{E}_{i \in [0,n]} [p_i] \approx \mathbb{E}_{i \in [0,n]} [\log p_i]$. Assume that we have a non-descending sequence of p , that is, $p_0 \leq p_1 \leq p_2 \leq \dots \leq p_n$, w.l.o.g. We can confirm that $\mathbb{E}_{i \in [0,n]} [p_i] \in [p_0, p_n]$. So we have: $\exists \xi \in (p_0, \mathbb{E}_{i \in [0,n]} [p_i])$, s.t.,

$$\log \mathbb{E}_{i \in [0,n]} [p_i] = \frac{1}{\xi} (\mathbb{E}_{i \in [0,n]} [p_i] - p_0) + \log p_0$$

We have:

$$\mathbb{E}_{i \in [0,n]} [\log p_i] \geq \log p_0,$$

		Reliability		Generalization		Specificity	
		$e_{local} \downarrow$	$e_{id} \downarrow$	$e_{ood} \downarrow$	$d_{KL}^r \times 10^{-7} \downarrow$	MAUVE \uparrow	
- Baseline Values -							
unedited		1.09 _{1.06}	1.09 _{1.06}	1.09 _{1.06}	0.00 _{0.00}	1.00 _{0.00}	
random		1.33 _{1.28}	1.33 _{1.28}	1.29 _{1.24}	2.31 _{1.02}	0.96 _{0.01}	
137M, shuffled		1.18 _{1.07}	1.18 _{1.08}	1.18 _{1.08}	1.64 _{0.48}	0.96 _{0.01}	
- Experimental Values -							
Average	$(b = -\infty)$	2.79 _{3.21}	2.86 _{3.24}	2.61 _{3.12}	4.00×10^3	0.92 _{0.08}	
	$b = 1$	0.17 _{0.29}	0.17 _{0.31}	0.18 _{0.33}	11.97 _{8.22}	0.95 _{0.02}	
Softmax	$b = 2$	0.19 _{0.30}	0.20 _{0.32}	0.15 _{0.28}	9.51 _{5.87}	0.96 _{0.01}	
	$b = 5$	0.17 _{0.28}	0.18 _{0.31}	0.16 _{0.28}	12.25 _{7.91}	0.95 _{0.01}	
Argmax	$(b = +\infty)$	1.21 _{1.68}	1.27 _{1.70}	1.30 _{1.84}	9.29×10^3	0.92 _{0.14}	

Table 3: Supplementary results of Table 2 with various b on GPT2.

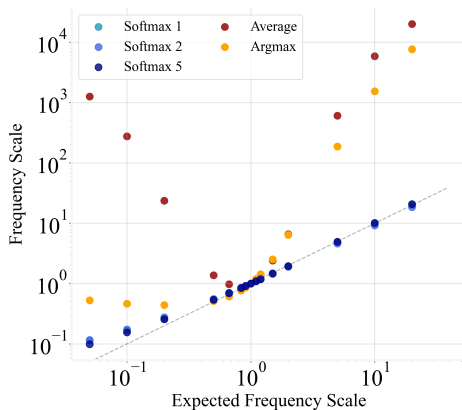


Figure 8: The expected probability scales and the actually steered scales measured in the edited LMs w.r.t. different values of b .

since $\log'(\cdot) > 0$. That is:

$$\begin{aligned}
& [\log \mathbb{E}_{i \in [0, n]}[p_i] - \mathbb{E}_{i \in [0, n]}[\log p_i]]^2 \\
& \leq \frac{1}{\xi^2} (\mathbb{E}_{i \in [0, n]}[p_i] - p_0)^2 \\
& \leq \frac{1}{p_0^2} (\mathbb{E}_{i \in [0, n]}[p_i] - p_0)^2.
\end{aligned} \tag{3}$$

We can empirically confirm the concentration of p_i , shown as examples in Fig. 11, which makes the error shown in Eq. 3 acceptably small. Additionally, a low-probability token shows a wide probability distribution, which is consistent with our findings in Fig. 3, where a low-probability token is assigned with more inaccurate predictions (manifested as a comet-shaped figure).

A.3 Details of MAUVE Calculation

The MAUVE is a similarity between two datasets. Refer Pillutla et al. (2021) for the calculating details




Figure 9: Supplement for Fig. 3. The MLR results on GPT2-XL and Pythia-2.8B.

of MAUVE. Here we explain our method to get both datasets for MAUVE to measure the harmfulness of our model steering method. We sample 512 data points from the BOOKCORPUS (Zhu et al., 2015), and take the first two words of each data point as a prefix to collect a prefix set. Given a generative language model, we input the prefixes and let the model generate naturally into a generated dataset. We do this process on the original model and the steered model and use the two generated datasets to calculate the MAUVE.

B Ablation Study on Steering Amount Allocation Parameter b

We discuss the steering amount allocation parameter b in Algorithm 1. We try different settings of b on GPT2 as shown in Table 3 and Fig. 8.

The results show that when the b is not infinities, the steering remains accurate and stable. That is, the algorithm is not sensitive to a positive integer b . But we still recommend a large b when the model

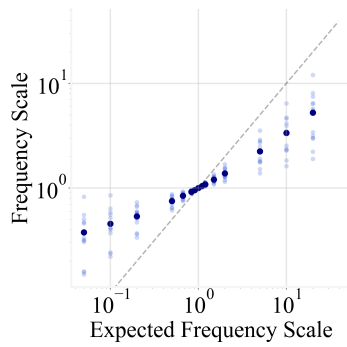


Figure 10: Supplement for Fig. 5. The expected probability scales and the actually steered scales measured in the steered GPT2-XL.

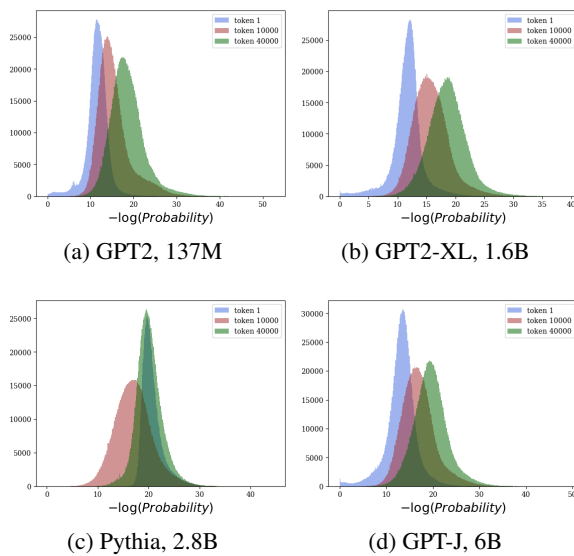


Figure 11: Examples of the probability distribution of one token among input prefixes.

is large and the sparsity of MLR slope is not fine.

However, when the b is $+\infty$ or $-\infty$, the steering can not be accurate and is easy to get a larger scale than expected. We infer that both of these infinity b increase the norm of the embedding vector too much, leading to an increase in its output probability as well.

C Supplementary Experiment Results

Supplement for Fig. 1. Output embedding visualizations w.r.t. output token probability for GPT2-XL, Pythia, GPT-J are shown in Fig. 12.

Supplement for Fig. 3. Output probability fitting results of GPT2-XL and Pythia-2.8B are shown in Fig. 9.

Supplement for Fig. 5. The expected probability scales and the actually edited scales of GPT2-XL is shown in Fig. 10.

Supplement for Table 2. Fine-grained results w.r.t. expected steering scales of GPT2 (Fig. 13), GPT2-XL (Fig. 14), GPT-J (Fig. 15).

Supplement for Fig. 4. The e_{local} , e_{id} , d_{KL}^r w.r.t. the detect dataset size are shown in Fig. 16.

Supplement for Fig. 7. The figures without log-scaling are shown in Fig. 17.

D Necessary Statements

Author Contributions. Hakaze Cho, also pronounced Yufeng Zhao, handled virtually all the workload in this paper. N.I. is their supervisor, providing important and beneficial guidance, revision, and research support. Y.S., K.T, and M.K. participated in the discussion and provided some helpful suggestions and revisions.

Repeatability statement. Models and datasets are all loaded from huggingface. All the datasets are shuffled with random seed 42 and cut into required slices. We calculate MAUVE by the package `mauve-text` by default parameters. In experiments, all the logarithms are natural.

License of the artifacts. The artifacts used in this paper are all open-sourced and are used for their intended usage. **Models.** GPT2 family are under the MIT license, GPT-J and Pythia are under the apache-2.0 license. **Datasets.** WIKIDPR is under the cc-by-nc-4.0 license, BOOKCORPUS and PILE is under the MIT license. **Tool.** MAUVE is under the GPLv3 license.

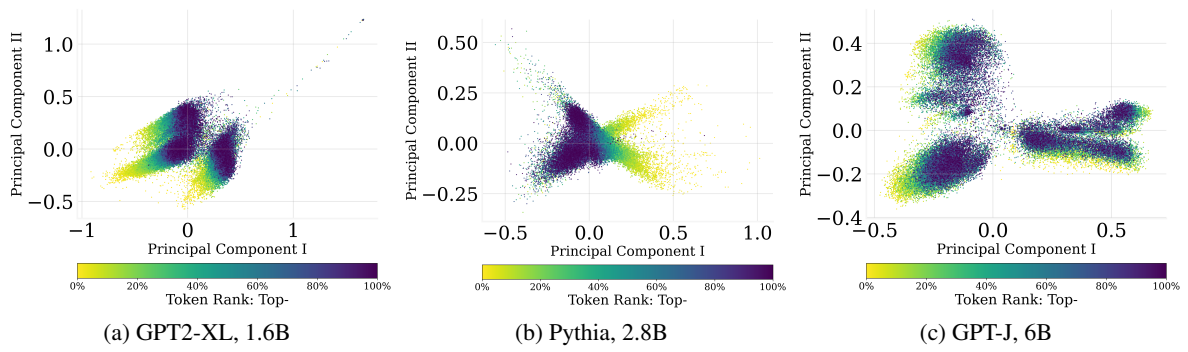


Figure 12: Supplement for Fig. 1. Output embedding visualizations w.r.t. output token probability for GPT2-XL, Pythia, GPT-J.

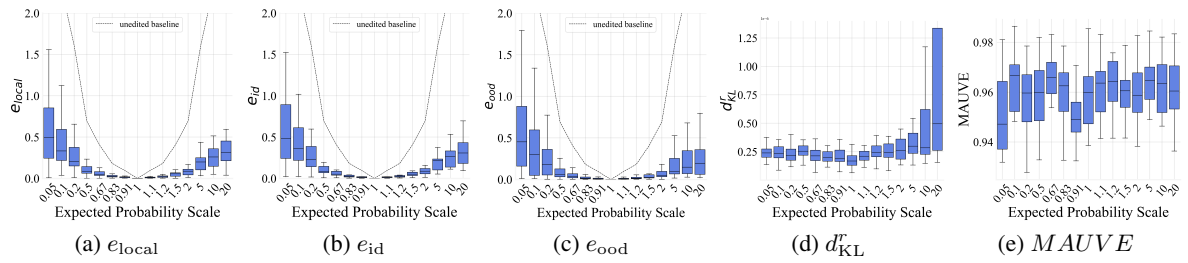


Figure 13: Fine-grained results w.r.t. expected steering scales of GPT2.

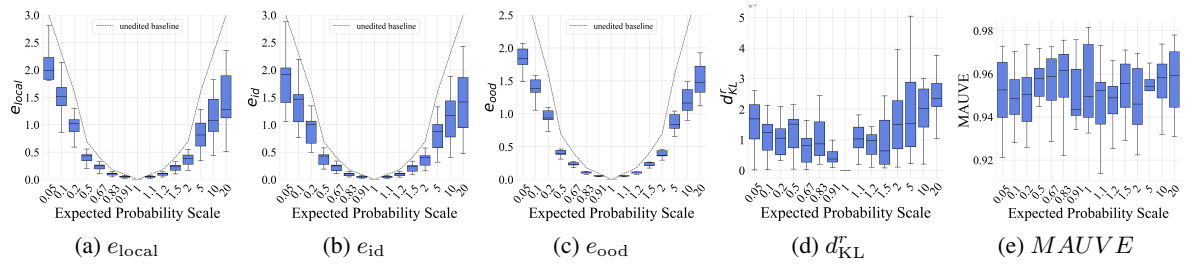


Figure 14: Fine-grained results w.r.t. expected steering scales of GPT2-XL.

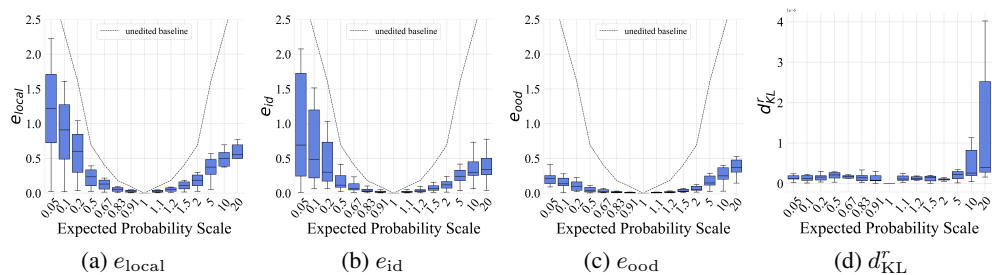


Figure 15: Fine-grained results w.r.t. expected steering scales of GPT-J.

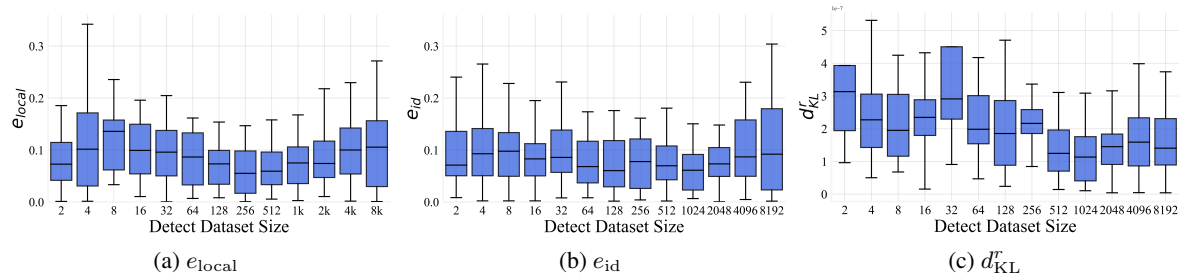


Figure 16: Supplement for Fig. 4. The e_{local} , e_{id} , d_{KL}^r w.r.t. the detect dataset size.




Figure 17: Supplement for Fig. 7. Without log-scaling.