

# On a Benefit of Masked Language Model Pretraining: Robustness to Simplicity Bias

Anonymous ACL submission

## Abstract

Despite the success of pretrained masked language models (MLM), why MLM pretraining is useful is still a question not fully answered. In this work we theoretically and empirically show that MLM pretraining makes models robust to lexicon-level spurious features, partly answering the question. Our explanation is that MLM pretraining may alleviate problems brought by simplicity bias (Shah et al., 2020), which refers to the phenomenon that a deep model tends to rely excessively on simple features. In NLP tasks, those simple features could be token-level features whose spurious association with the label can be learned easily. We show that MLM pretraining makes learning from the context easier. Thus, pretrained models are less likely to rely excessively on a single token. We also explore the theoretical explanations of MLM’s efficacy in causal settings. Compared with Wei et al. (2021), we achieve similar results with milder assumption. Finally, we close the gap between our theories and real-world practices by conducting experiments on real-world tasks.

## 1 Introduction

The question “*why is masked language model (MLM) pretraining (Devlin et al., 2019; Liu et al., 2019) useful?*” has not been totally answered. In this work, as an initial step toward the answer, we show and explain that MLM pretraining makes the model robust to lexicon-level features that are spuriously associated with the target label. It gives the model a better generalization capability under distribution shift.

Previous studies have empirically shown the robustness of MLM pretrained models. Hao et al. (2019) show that MLM pretraining leads to wider optima and better generalization capability. Hendrycks et al. (2020) and Tu et al. (2020) show that pretrained models are more robust to out-of-distribution data and spurious features. However,




Figure 1: The pitfall of simplicity bias: The solid line is a simple (linear) decision boundary that utilizes only one dimension, while the dashed line is a more complex decision boundary that utilizes two dimensions and maximizes the margin.

it remains unanswered why pretrained models are more robust.

We conjecture that models trained from scratch suffer from the pitfall of simplicity bias (Shah et al., 2020) (Figure 1). Shah et al. (2020) and Kalimeris et al. (2019) showed that deep networks tend to converge to a simple decision boundary that involves only a few features. The networks may not utilize all the features and thus may not maximize the margin, which results in worse robustness. A consequence of this could be that a model may excessively rely on a feature that has spurious association with the label and ignore the other features that are more robust. In the studies of Shah et al. (2020) and Kalimeris et al. (2019), they investigated networks with continuous input. Lovering et al. (2021) discovered similar results on synthetic NLP tasks, where the inputs are discrete. We will further explore this discrete setting in this work.

We start the exploration with the following assumptions: Let the sentence, label pair be  $X, Y$ .

**Assumption 1.** We assume that from  $X$ , we can extract two features  $X_1$  and  $X_2$ .

**Assumption 2.**  $X_1$  is a spurious feature that has strong association with  $Y$ . Specifically, it means that, solely relying on  $X_1$ , one can predict with high accuracy, but cannot be 100% correctly.

**Assumption 3.**  $X_2$  is a robust feature based

on which  $Y$  can be predicted with 100% accuracy. Namely, there exists a deterministic mapping  $f_{X_2 \rightarrow Y}$  that maps  $X_2$  to  $Y$ .

The assumptions above are realistic in some NLP tasks. In NLP tasks, the input  $X$  is a sequence of tokens. Some tasks satisfy Assumption 1:  $X$  can be decomposed into  $X_1$  and  $X_2$ , where  $X_1$  is the presence of certain tokens, and  $X_2$  is the context of the token. Thus,  $X_2$  has a much higher dimensionality than  $X_1$ . As shown by the analysis of Gardner et al. (2021), there are indeed datasets where Assumption 2 is true. However, if Assumption 3 is true, we would desire the model to rely on  $X_2$ , which contains the semantic of the input  $X$ .

With these assumptions, in Section 2 we empirically demonstrate that spurious features in discrete inputs can cause problems as in the continuous cases (Shah et al., 2020; Kalimeris et al., 2019). We show that, possibly due to the simplicity bias, a deep model is likely to excessively rely on  $X_1$  and to rely on  $X_2$  less. In Section 3.1 and Section 3.2 we provide a theoretical explanation of how MLM pretraining makes a model robust to spurious features. Let  $\Pi^{(1)}$  be the conditional probability  $P(X_1|X_2)$ . We show (1) the relation between the mutual information  $I(\Pi^{(1)}; Y) \geq I(X_1; Y)$  and that (2) the convergence rate of learning from  $\Pi^{(1)}$  is of the same order as learning from  $X_1$ . That is, when the MLM model can perfectly model the probability  $P(X_1|X_2)$  and thus generate perfect  $\Pi^{(1)}$ , learning from  $\Pi^{(1)}$  is as easy as learning from  $X_1$ . As a result, the model will be more likely to rely on  $\Pi^{(1)}$ . Since  $\Pi^{(1)}$  is estimated based on  $X_2$ , higher reliance on  $\Pi^{(1)}$  also implies higher reliance on the robust feature  $X_2$ . This avoids the pitfall of simplicity bias that the model relies excessively on  $X_2$ . To relax Assumption 3, we make one step further by considering causal settings in Section 3.3.

The above results partly explain why MLM pretraining is useful for NLP. Denote a sequence of tokens as  $X = \langle X_1, X_2, \dots, X_L \rangle$ . During the MLM pretraining process, each token is masked randomly at certain probability, and the training objective is to predict the masked tokens with the maximum likelihood loss. As a result, the model is capable of estimating the conditional probability  $P(X_i|X \setminus X_i)$  for all  $i = 1, 2, \dots, L$ . Even though which of the tokens is spurious is unknown, as long as the spurious token has non-zero probability to be masked during pretraining, MLM can estimate its distribution conditioned on the context and thus

can reduce the reliance on it.

Finally, we close the gap between our theories and reality. One major gap is that, in reality, we do not use the conditional probability for downstream tasks. Instead, we feed the input  $X$  without masking any token and fine-tune the model along with a shallow layer over its output. Regardless of that, we hypothesize that the robustness brought by MLM pretraining still exists. To prove that, in Section 4 we use the toy example and verify the effect of MLM pretraining when using the common practice for fine-tuning. In Section 5, we validate our theories with two real-world NLP tasks.

To sum up, our study leads to new research directions. Firstly, we provide a new explanation of MLM pretraining’s efficacy. Unlike the previous purely theoretical studies (Saunshi et al., 2021; Wei et al., 2021), our assumptions are milder and more realistic. Secondly, we study NLP robustness from the perspective of self-supervised model. Since self-supervised trained embeddings have been widely used since Word2vec (Mikolov et al., 2013), and it is indispensable to the generalization to unseen data, a better understanding of self-supervised embeddings should be important for future research.

## 2 A Toy Example

To show that spurious association can cause difficulty of convergence, we construct a toy example with variables  $X_1, X_2, Y$  that satisfy the assumptions. We make  $X_1$  depends only on  $X_2$ , so it is not a caustic feature of  $Y$ . Let the dimension of the random variables  $X_1$  and  $X_2$  be 2 and  $d_2$  respectively. Their value  $x_1 \in \mathcal{X}_1 = \{e_1, e_2\}$  and  $x_2 \in \mathcal{X}_2 = \{e_1, \dots, e_{d_2}\}$ , where  $e_i$  is the one-hot vector whose  $i$ th element is 1. We control the strength of the association between  $X_1$  and  $Y$  with  $\nu < 0.5$ , making  $X_1 = Y$  with probability  $1 - \nu$ . Specifically, denote with  $\dot{X}_2$  the middle  $2\nu d_2$  dimensions of  $X_2$ , i.e. the  $\lfloor d_2/2 - \nu d_2 \rfloor$ th to the  $\lfloor d_2/2 + \nu d_2 \rfloor$ th elements in  $X_2$ . We consider the following random process:

$$\begin{aligned} X_2 &= e_i, i \sim \text{Uniform}(1, d_2) \\ Y &= \begin{cases} -1 & \text{if } X_2 = e_i \text{ for some } i < d_2/2 \\ +1 & \text{otherwise} \end{cases} \\ X_1 &= \begin{cases} e_i, i \sim \text{Uniform}(1, 2) & \text{If } \dot{x}_2 \neq 0 \\ f(X_2) & \text{Otherwise} \end{cases} \end{aligned} \quad (1)$$

$d_2$	$\nu$	1 layer	2 layers		3 layers	
		w/o	w/o pre	w/ pre	w/o pre	w/ pre
50	0.04	3680 (189.5)	691 (55.8)	614 (169.1)	302 (47.2)	249 (53.7)
	0.10	2664 (121.2)	530 (30.6)	441 (134.9)	242 (27.6)	180 (37.5)
	0.25	1420 (96.0)	352 (23.8)	300 (62.0)	179 (13.8)	148 (28.7)
	0.50	306 (79.8)	141 (40.7)	118 (33.4)	106 (23.1)	89 (24.0)
100	0.04	5466 (170.1)	945 (57.2)	689 (225.3)	431 (51.1)	275 (72.1)
	0.10	3789 (99.2)	677 (32.2)	478 (142.9)	317 (30.3)	208 (44.3)
	0.25	1952 (64.9)	428 (13.1)	330 (85.0)	214 (16.2)	169 (32.5)
	0.50	330 (78.0)	156 (34.0)	133 (41.2)	128 (28.2)	112 (36.1)
500	0.04	11127 (265.9)	1953 (112.5)	857 (442.6)	792 (69.8)	431 (88.4)
	0.10	7912 (169.2)	1279 (67.5)	657 (234.9)	550 (46.7)	402 (97.0)
	0.25	4321 (152.3)	772 (35.5)	501 (133.5)	399 (42.3)	391 (66.0)
	0.50	576 (150.0)	392 (70.2)	407 (81.1)	367 (69.1)	386 (80.0)

Table 1: The number of iterations a model w/ or w/o pretraining requires to converge. The number is the average of 25 runs with different random seeds, and the number in parentheses is the standard deviation.

where  $f(X_2) = e_1$  if  $X_2 = e_i$  for some  $i < d_2/2$ , and  $f(X_2) = e_2$  otherwise<sup>1</sup>. In this way, predicting  $Y$  solely based on the spurious feature  $X_1$  can achieve accuracy  $1 - \nu$ .

We conduct experiments to inspect the effect of the strength of spurious association between  $X_1$  and  $Y$ . We train linear networks by drawing batches of i.i.d.  $([X_1; X_2], Y)$  pairs from the random process defined in 1. We use Adam optimization with learning rate 0.001 and the cross-entropy loss. In addition to single-layer linear networks, we also try over-parameterized 2-layer and 3-layer linear networks. The hidden size is [10, 32]. Since it is a linear separable problem, we can check whether the learned weight can lead to 100% accuracy in the defined distribution. We check it every 25 iterations. We say a model has converged if it is 100% accurate for 5 consecutive checks. We report the number of the iterations required before it converges for different  $\nu$  and  $d_2$ .

Even though it is a linear-separable convex optimization problem, our results in Table 1 show that the strength of the spurious association can impact the number of iterations required to converge. We observe that when  $\nu < 0.5$ , the models tend to be trapped by the spurious feature, sticking at accuracy  $1 - \nu$  for iterations. When the spurious relation between  $X_1$  and  $Y$  is stronger, i.e.  $\nu$  is smaller, the number of iterations required to converge is larger. In addition, the number of iterations is also larger when the  $d_2$  is larger. An intuitive explanation is

that the learning signal from  $X_2$  is more sparse when  $d_2$  is larger.

### 3 A Theoretical Explanation of the Efficacy of MLM Pretraining

#### 3.1 $P(X_1|X_2)$ is More Informative Than $X_2$

The toy example above motivates us to consider the information contained in  $P(X_1|X_2)$ . In the toy example, when predicting  $P(Y = 0|X)$ , if we simply output  $P(X_1 = e_1|X_2)$ , then the accuracy of our prediction of  $Y$  will be as high as predicting  $Y$  solely based on  $X_1$ . It motivates us to inspect the reliability of the estimated  $P(X_1|X_2)$  as a feature for the prediction of  $Y$  compared to  $X_1$ . Let  $\Pi_1 = P(X_1|X_2)$  be a  $|\mathcal{X}_1|$ -dimensional random variable<sup>2</sup>. We can prove that when  $P(X_1|X_2)$  is estimated perfectly,  $\Pi_1$  is at least as informative as  $X_1$ .

**Lemma 1.** *When  $X_1, X_2$  are discrete, if  $\Pi_1$  perfect, namely  $\Pi_1 = P(\cdot|X_2)$ , then the mutual information  $I(X_1; \Pi_1) = I(X_1; X_2)$ . (Proof: Appendix A.1)*

Compared to previous works (Hjelm et al., 2019; Belghazi et al., 2018; Oord et al., 2018) that show some self-supervised training objectives are lower bounds of the mutual information  $I(X_1; X_2)$ , we directly show that the output of the MLM,  $\Pi$ , maximizes the mutual information, since  $I(X_1; f(X_2)) \leq I(X_1; X_2)$  for any  $f$ . Moreover, instead of explaining the efficacy of pretraining

<sup>1</sup>Uniform( $a, b$ ) is the uniform distribution over  $\{n\}_{n=a}^b$ .

<sup>2</sup>We will omit the subscript of  $\Pi_1$  when there is no ambiguity.

with the infomax principle (Linsker, 1988; Bell and Sejnowski, 1995), our theories below provide a different perspective.

**Theorem 1.** *If  $\Pi$  is perfect,*

$$I(\Pi; Y) \geq I(X_1; Y) \quad (2)$$

*Proof.* Since  $\Pi$  is perfect, by Lemma 1, we have

$$I(X_1; X_2) = I(X_1; \Pi). \quad (3)$$

By data processing inequality, Equation 3 implies  $I(X_1; X_2|\Pi) = 0$ . By Assumption 3, a deterministic mapping  $f_{X_2 \rightarrow Y}$  from  $X_2$  to  $Y$  exists. Applying data processing inequality again, we have

$$\begin{aligned} I(X_1, X_2|\Pi) &\geq I(X_1, f_{X_2 \rightarrow Y}(X_2)|\Pi) \\ &= I(X_1, Y|\Pi) \geq 0, \end{aligned} \quad (4)$$

which implies  $I(Y, X_1|\Pi) = 0$ . Accordingly,

$$H(Y|\Pi) = H(Y|X_1, \Pi) \leq H(Y|X_1) \quad (5)$$

□


Theorem 1 shows that  $\Pi$  is a more informative feature than  $X_1$ . However, a model does not necessarily rely more on a more informative feature. We will discuss more in the next section.

### 3.2 Learning from $\Pi$ is Easy

It is important that learning from  $\Pi$  is easy. Because of simplicity bias, a neural network model is likely to rely on the easy-to-learn features (Shah et al., 2020; Kalimeris et al., 2019). We conjecture that a model excessively relies on the spurious feature  $X_1$  when learning from  $X_1$  is easier than learning from the robust feature  $X_2$ . If learning from  $\Pi$  is easy, then the model will rely on  $\Pi$  more and thus will rely on  $X_1$  less. However, features with higher mutual information to  $Y$  are not necessarily easy to learn. For instance, although  $X_2$  is more informative, models tend to rely on  $X_1$  instead of  $X_2$  at the beginning of the training process. To show that MLM can mitigate the issue brought by simplicity bias, we need to show learning from  $\Pi$  is easy.


Therefore, we have the following theorem that implies learning from  $\Pi$  is at least as easy as learning from  $X_1$ :

**Theorem 2.** *Let  $\tilde{h}_{X_1}^{(D_n)} : \mathcal{X}_1 \rightarrow \mathcal{Y}$  be the classifier trained with MLE loss using  $n$  data pairs  $(x_1^{(1)}, y^{(1)}), (x_1^{(2)}, y^{(2)}), \dots, (x_1^{(n)}, y^{(n)})$ , and*



(a) Causal setting

(b) Anticausal setting



(c) A case where  $I(\Pi_3; Y) \geq I(\Pi_1; Y)$  is possible.

Figure 2: The causal settings of the  $(X, Y)$  pairs.

the converged classifier be  $\tilde{h}_{X_1}^*$ . There exists a learning algorithm, which generates  $\tilde{h}_{\Pi}^{(D_n)} : \Pi \rightarrow \mathcal{Y}$  using  $(\pi^{(1)}, y^{(1)}), (\pi^{(2)}, y^{(2)}), \dots, (\pi^{(n)}, y^{(n)})$ , such that the following three properties are satisfied: (1)

$$\mathbb{E} \left[ D_{KL} \left[ \tilde{h}_{\Pi}^{(D_n)} \parallel \tilde{h}_{\Pi}^* \right] \right] = O \left( \frac{1}{n} \right), \quad (6)$$

which is asymptotically at the same rate as  $\mathbb{E} \left[ D_{KL} \left[ \tilde{h}_{X_1}^{(D_n)} \parallel \tilde{h}_{X_1}^* \right] \right]$ . (2) Over the distribution of  $(X, Y)$ , the expected loss of the converged classifier  $\tilde{h}_{\Pi}^*$  is not greater than the expected loss of  $\tilde{h}_{X_1}^*$ . (3)  $\tilde{h}_{\Pi}^*$  is a linear model, whose input is  $\Pi$ . (Proof: Appendix A.2)

The remaining question is whether deep learning models used in common practices can perform at least as well as the algorithm in Theorem 2. Indeed, without any knowledge of deep learning models, it is impossible to theoretically prove that a model will necessarily rely on  $\Pi$  instead of  $X_1$ . Therefore, in Section 4 and Section 5 we will empirically validate that our theorems are applicable in the real world scenarios.

### 3.3 Extending with Causal Models

Before we validate our theories empirically, we make a step further by relaxing Assumption 3. We do so by treating  $X_2$  as a confounder, and then we can see how MLM pre-training is helpful in the causal and anticausal settings.

	Ours	Wei et al. (2021)
structural assumption	$X, Y$ follow Figure 2b.	$X, Y$ follow an HMM.
linear independence assum.	$\{P(X_1 y) y \in \mathcal{Y}\}$	$\{P(X_0 H_0 = h) h \in \mathcal{H}\}$
implication	$I(P(X_1 X_2); Y) = I(X_2; Y)$	$I(P(H_0 X); Y) = I(X; Y)$

Table 2: Comparison between Theorem 4 in this work and Theorem 3.3 in Wei et al. (2021).

**Theorem 3.** *Even if Assumption 3 is not true, Theorem 1 still holds if  $X_1, X_2, Y$  follow the causal setting in Figure 2a.*

*Proof.* By the structure of  $X_1, X_2, Y$ , inequality 4 holds even if the deterministic mapping  $f_{X_2 \rightarrow Y}$  does not exist.  $\square$

**Theorem 4.** *Assume that the set of vectors  $\{P(X_1|Y = y)|y \in \mathcal{Y}\}$  is linear independent, and if  $X_1, X_2, Y$  follow the anticausal setting in Figure 2b, then  $I(\Pi, Y) = I(X_2, Y)$ .*

*Proof.* The assumption is a special case of the one in (Lee et al., 2020), so similar techniques can be used: According to the structure of  $X_1, X_2, Y$ , we have

$$P(X_1|X_2) = \sum_y P(X_1|y)P(y|X_2). \quad (7)$$

Therefore, if  $\{P(X_1|Y = y)|y \in \mathcal{Y}\}$  is linearly independent,  $P(y|X_2)$  can be recovered from  $\Pi = P(X_1|X_2)$ .  $\square$

Note that this theorem is very similar to Theorem 3.3 in Wei et al. (2021). However, the assumptions required in ours are weaker and more realistic, and the implication is very similar (Table 2): **(1) Structure assumption:** Wei et al. (2021) assumed that  $X$  is generated from a HMM process with hidden variables  $H_0, H_1, \dots$ , which is stronger assumption than our assumption that  $X_1, X_2$  follow the anticausal setting. **(2) Independence assumption:** Wei et al. (2021) assumed that the vectors in  $\{P(X_0|H_0 = h)|h \in \mathcal{H}\}$  need to be linearly independent. In comparison, we require only the independence in  $\{P(X_1|Y = y)|y \in \mathcal{Y}\}$ . Our assumption is more realistic because the number of hidden states  $|\mathcal{H}|$  must be very large if  $X$  is generated from the HMM model, and  $|\mathcal{Y}|$  tends to be much smaller than  $|\mathcal{H}|$ . For example, in binary classification cases, our assumption holds as long as  $P(X_1)$  is not independent of  $P(Y)$ . **(3) Implication:** If we further assume that  $I(X_2; Y) = I(X; Y)$ , then we reach a similar conclusion that  $P(Y|X)$  can

be recovered from  $\Pi = P(X_1|X_2)$  by applying a linear function.

### 3.4 Limitations of Our Theorems

Our theories do not ensure that  $\Pi_1$  is the most informative feature to learn from. Consider tokens in a sentence  $X = \langle X_1, X_2, \dots, X_L \rangle$  and let  $\Pi_i$  be the conditional probability  $P(X_i|X \setminus X_i)$ . A token with spurious association with the label can locate arbitrary position in the sentence, and its location is unknown during pretraining. That is, the pretrained model is able to generate  $\Pi_i$  for all  $i$ . Without loss of generality, assume  $X_1$  is the spurious token. *It is possible that there exists some  $i$  such that  $I(\Pi_1; Y) < I(\Pi_i; Y)$ , and that  $\Pi_i$  is predicted relying on  $X_1$ .* Concretely, here is an example for the causal setting with three features:  $X_3$  is independent of  $X_1$  and  $X_2$  given  $Y$  (Figure 2c). Using the results in Theorem 4, there is a linear mapping that can recover  $P(Y|X_1, X_2)$  from  $\Pi_3$ . Therefore, it is possible that  $I(\Pi_3; Y) > I(\Pi_1; Y)$  if  $I(X_1, X_2; Y) > I(\Pi_1; Y)$  depending on the distribution of the data. We leave the study of  $I(\Pi_i; Y)$  for future work.

Another limitation is that, in practice, NLP practitioners do not use the conditional probability predicted by the pretrained model. Instead, people stack a simple layer over the pretrained model, and fine-tune the whole model on downstream tasks. Regardless of this, we conjecture that the representation encoded by a MLM pretrained model still contains the information of  $\{\Pi_i\}_{i=1}^n$  and thus is robust to spurious lexicon-level feature.

## 4 Toy Example with a Pretrained Model

As the first step to close the gap between our theories and the real world, we repeat the toy experiments with pretraining. Before fitting the model with  $Y$ , we first pretrain the first layer to predict  $X_1$  based on masked  $X$ . What we want to show is that, after pretraining, the representation encoded by the layer will have the equivalent role of  $\Pi$  even when the input is not masked.

Specifically, the experimental design is as follows: We use the two-layer and three-layer MLP architectures same as in Section 2. When pretraining, we mask  $X_1$  in  $X$  by using  $X' = [0, 0; X_2]$  as inputs. Let the output from the first linear layer as  $Z = WX'$ . The loss function is the cross-entropy between  $X_1$  and the softmax over  $[z_1, z_2]$ . After pretraining, we fine-tune the pretrained model with  $([X_1; X_2], Y)$  pairs, and report the average number of iterations required to converge for 25 different random seeds.

We want to eliminate the possibility that the faster convergence of the pretrained model is because of larger initial weights over  $X_1$ . Therefore, after pretraining, we manually create a path from  $X_1$  to  $Z$ . We do so by initializing the weights of the third and fourth row of  $W$  with  $[k, -k, 0, \dots, 0]$  and  $[-k, k, \dots, 0, 0]$  respectively, where  $k$  is the average of the absolute value of the weights in the pretrained part, i.e. the weights of the first two rows in  $W$ . In this way, the information from  $X_1$  has the same scale as the pretrained representation  $[Z_1, Z_2]$ , and thus it can compete with  $[Z_1, Z_2]$  fairly.

Table 1 shows that pretraining can always reduce the number of iterations required to converge when  $\nu < 0.50$ . The effect is more significant when  $d_2$  is larger. It could be because of the higher sparsity of the learning signal from  $X_2$  when  $d_2$  is larger.

We further inspect how the importance over the inputs changes in the process of training. The importance can be inferred from the product of the linear layers. We observe that if the model is not pretrained, the weights over  $X_1$  grow faster than the weights over  $X_2$  at the beginning (the first row Figure 3). The model cannot converge to 100% accuracy until weights on  $\dot{X}_2$ , the middle  $[\nu \times d_2]$  dimensions of  $X_2$ , become greater than the weights on  $X_1$ . In addition, after the model converges, weights over  $X_1$  is still greater than weights over  $X_2$ . On the other hand, if the model is pretrained, weights over  $X_1$  stop growing after a few steps (the second row in Figure 3). The above observations are aligned with our conjecture that the pretrained representation mitigates the robustness issue brought simplicity bias.

## 5 Experiments

To further validate our theories, we experiment on real world NLP tasks. We facilitate datasets that have known spurious features. We first pretrain models on the training dataset with different mask-

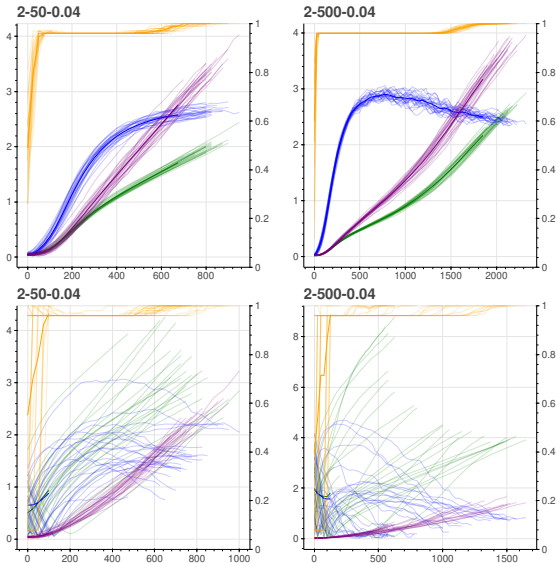


Figure 3: The average weights over the features during training a two-layer model. The upper and the lower rows are the curves of model without and with pretraining respectively. From left to right,  $(d_2, \nu) = (50, 0.04), (500, 0.04)$ . Blue, green, purple curves represent the average weights over features in  $X_1$ ,  $X_2$ , and  $\dot{X}_2$  (the middle part of  $X_2$ ) respectively. The orange curve represents the accuracy.

ing policies. Different masking policies mask the spurious tokens in different ways. Afterward, we fine-tune the model using the target label. We show that the models will be less robust on downstream tasks if spurious tokens are not masked during pretraining, and always masking the spurious token during pretraining improving the robustness. The results validate our theories.

### 5.1 Downstream Tasks

**Hate Speech Detection** Previous study has shown that hate speech detection datasets tend to have lexical bias (Dixon et al., 2018). That is, models rely excessively on the presence or the absence of certain words when predicting the label. Here we follow the formulation of lexical bias in hate speech detection proposed by Zhou et al. (2021). We focus on the effect of non-offensive minority identity (NOI) mentions, such as “woman”, “gay”, “black”. Those mentions are often highly associated with hateful instances. However, it is more desirable that a model does not rely on those mentions. Therefore, we can see the presence of NOI as a spurious feature.

**Name Entity Recognition (NER)** Lin et al. (2020) has shown that name entity recognition

(NER) models perform worse when the name entities are not seen in the training data. In this case, we can see the content of the name entities as a spurious feature. Models may learn to memorize the name entities when fitting the training data, while we may desire the model to recognize name entities according to the context.

### 5.1.1 Datasets

**Hate Speech Detection** We use a portion of the dataset proposed by Zhou et al. (2021). In their original dataset, only a small number of hateful instances contain NOI. Our preliminary experiments show that the model without pretraining does not suffer much from the bias of NOI when training with the full data. Therefore, we create a dataset, whose positive (hateful) instances are all the positive samples in the original dataset that contain NOI. As for negative instances, we sample them randomly from the original training set. We control the number of negative instances so the ratio of positive and negative instances is the same as the original dataset. We create both the training and the validation splits in this way, and use the original full testing set for evaluation. We also evaluate the models on a NOI subset where all the instances contain NOI.

**NER** We use the standard NER dataset Conll-2003 (Tjong Kim Sang and De Meulder, 2003). To create a testing set with unseen name entities, we replace the name entities in the original validation and testing splits with the entities from WNUT-17 (Derczynski et al., 2017). Specifically, we replace the LOC, ORG, PER entities with the corresponding type of entities in WNUT-17, while the MISC entities remain untouched.

## 5.2 Masking Policies

For each sentence with  $n_s$  spurious tokens, we experiment with different masking policies: (1) **scratch**: We do not pretrain the model before fine-tuning. (2) **vanilla**: During pretraining, we mask each token with 15% probability, which is same as the original implementation in (Devlin et al., 2019). (3) **unmask random**: This is similar to vanilla MLM, but we uniformly randomly select  $n_s$  tokens from the whole sentence and unmask them if they have been masked. (4) **unmask spurious**: This is similar to vanilla MLM, but we unmask all the spurious tokens. (5) **remove spurious**: We replace spurious tokens with a special “[unk]” token, and

we unmask them. Note that this setting can be seen as an oracle setting, since in most applications the spurious features are unknown.

Note that setting (3), (4), (5) have the same expected number of masked tokens. Therefore, it rules out the possibility that their downstream performance differs because of the number of masked tokens. Instead, the difference should be resulted from whether spurious tokens are masked during pretraining.

## 5.3 Implementation Details

For both of the tasks and all the MLM settings, we tokenize the input with the bert-base-uncased tokenizer. We use the bert-base-uncased architecture, and also the pretrained embedding layer, which is frozen through the pretraining process. We include more details in Appendix A.3.

## 5.4 Result and Discussion

Results in Table 3 validate our theorems. For both of the tasks, *unmask random* performs better than *unmask spurious* under distribution shift. Specifically, *unmask random* has higher F1 on the unseen set of the NER task, and *unmask random* has a lower false positive rate (FPR) on the NOI set. Also, *unmask random* performs similarly to *random*. This implies that modeling the condition distribution of spurious tokens in the original random masking pretraining can reduce models’ reliance on them. Note that *unmask random* and *unmask spurious* have similar in-distribution performance, so the performance difference is not due to better in-distribution generalization suggested by Miller et al. (2021).

We also compare *unmask random* with the oracle setting *remove spurious*. We notice that even though *remove spurious* performs as well as *random*, *remove spurious* hurts the performance in the seen set. It indicates that modeling the conditional distribution of spurious tokens has effects beyond simply removing them from the model. On the other hand, *remove spurious* performs better in the hate speech detection task. A possible explanation is that NOI mentions contain little useful information for the task.

## 6 Related Work

Recently, there are efforts attempting to explain the effectiveness of massive language modeling pretraining. Theoretically, Saunshi et al. (2021)

Mask Policy	NER		Hate Speech Detection			
	Origin	Unseen	All (12893)		NOI (602)	
	F1 $\uparrow$	F1 $\uparrow$	Accuracy $\uparrow$	F1 $\uparrow$	Accuracy $\uparrow$	FPR $\downarrow$
scratch	61.5 <sub>0.5</sub>	38.7 <sub>0.6</sub>	83.9 <sub>1.6</sub>	80.3 <sub>1.4</sub>	74.8 <sub>1.5</sub>	46.3 <sub>7.2</sub>
vanilla	74.2 <sub>0.4</sub>	56.5 <sub>1.3</sub>	83.1 <sub>0.8</sub>	78.5 <sub>0.8</sub>	75.8 <sub>0.5</sub>	25.1 <sub>1.8</sub>
unmask random	72.7 <sub>0.6</sub>	56.5 <sub>0.8</sub>	83.3 <sub>1.1</sub>	78.9 <sub>1.1</sub>	75.8 <sub>0.9</sub>	25.7 <sub>2.3</sub>
unmask spurious	72.9 <sub>0.5</sub>	53.2 <sub>0.8</sub>	84.1 <sub>0.7</sub>	79.8 <sub>0.6</sub>	73.7 <sub>1.0</sub>	32.5 <sub>2.1</sub>
remove spurious	69.8 <sub>0.5</sub>	56.7 <sub>0.8</sub>	82.4 <sub>1.0</sub>	77.8 <sub>1.0</sub>	77.3 <sub>0.6</sub>	21.7 <sub>2.0</sub>

Table 3: The performance on downstream tasks. For the hate speech detection task, we also report false positive detection (FPR) on the NOI subset, which is a set of instances containing non-offensive minority identity mentions, e.g. “women”, “black”.

explore why auto-regressive language models help solve downstream tasks. However, their explanation is based on the assumption that the downstream tasks are *natural tasks*, i.e. tasks that can be reformulated as sentence completion tasks. Their explanation also requires the pretrained language model to perform well for any sentence completion tasks, which is not likely to be true in the real world. Wei et al. (2021) analyzes the effect of fine-tuning a pretrained MLM model. Nonetheless, they have stronger assumptions as described in Section A.3. Aghajanyan et al. (2020) show that pretrained models have lower intrinsic dimension, and provide a generalization bound based on Arora et al. (2018). However, why pretrained models have lower intrinsic dimension is still unknown. Merrill et al. (2021) show that the parameter norm growth during training makes transformer a saturated model, which can be described in terms of formal languages. Empirically, Zhang and Hashimoto (2021) show that the effectiveness of MLM pretraining cannot be explained by formulating the downstream tasks as sentence completion problems. Sinha et al. (2021) find evidence supporting the hypothesis that masked language models benefit from modeling high-order word co-occurrence instead of word order. There are also some theories explaining the efficacy of non-MLM pretraining Lee et al. (2020); Saunshi et al. (2019); Zhang and Stratos (2021).

Many of the previous studies on robust NLP focus on supervised learning (Wang et al., 2021; Utama et al., 2020b,a; Karimi Mahabadi et al., 2020; Chang et al., 2020; He et al., 2019; Sagawa\* et al., 2020; Kennedy et al., 2020; Chiang et al., 2020). However, without self-supervised learning, a model can impossibly extrapolate to out-of-distribution data when the domain shifts. Our work

also complements previous studies that focus on the bias or robustness of a model generated by the pretraining process (Kumar et al., 2020; Hawkins et al., 2020; Vargas and Cotterell, 2020; Liu et al., 2020; Gonen and Goldberg, 2019; Kurita et al., 2019; Zhao et al., 2019). In this work we investigate the pretraining process itself.

## 7 Implication

Our results provide possible explanations for some common practices found effective empirically. First, it could explain why continuing pretraining on target dataset is useful (Gururangan et al., 2020). It may be because continuing pretrained models model the distribution of spurious features in the target dataset better. Thus the model can better avoid the simplicity pitfall. Second, it provides reasons for more complex masking policies, such as masking continuous random spans (Joshi et al., 2020). It may improve the robustness to spurious features that contain more than one token. Third, if MLM can alleviate the simplicity bias and help the model to achieve a greater margin, it may also imply that the model has wider optima, explaining the finding in Hao et al. (2019).

## 8 Conclusion

In this work, we show a benefit of MLM pretraining, which partly explains its efficacy. We first empirically demonstrate the presence of simplicity bias when the input is discrete. We then theoretically and empirically explain how MLM pretraining can alleviate the problem brought by it. Finally, we close the gap between our theories and real-world practices with experiments on real-world NLP tasks. Our results shed light on future research on robustness and self-supervised learning.



617  
618  
619  
620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
630  
631  
632  
633  
634  
635  
636  
637  
638  
639  
640  
641  
642  
643  
644  
645  
646  
647  
648  
649  
650  
651  
652  
653  
654  
655  
656  
657  
658  
659  
660  
661  
662  
663  
664  
665  
666  
667  
668  
669  
670  
671  
672

## References

Armen Aghajanyan, Luke Zettlemoyer, and Sonal Gupta. 2020. Intrinsic dimensionality explains the effectiveness of language model fine-tuning. *arXiv preprint arXiv:2012.13255*.

Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. 2018. Stronger generalization bounds for deep nets via a compression approach. In *International Conference on Machine Learning*, pages 254–263. PMLR.

Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm. 2018. **Mutual information neural estimation**. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 531–540. PMLR.

Anthony J Bell and Terrence J Sejnowski. 1995. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159.

Shiyu Chang, Yang Zhang, Mo Yu, and Tommi Jaakkola. 2020. Invariant rationalization. In *International Conference on Machine Learning*, pages 1448–1458. PMLR.

Ting-Rui Chiang, Hao-Tong Ye, and Yun-Nung Chen. 2020. An empirical study of content understanding in conversational question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7578–7585.

Leon Derczynski, Eric Nichols, Marieke van Erp, and Nut Limsopatham. 2017. **Results of the WNUT2017 shared task on novel and emerging entity recognition**. In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 140–147, Copenhagen, Denmark. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 67–73.

Matt Gardner, William Merrill, Jesse Dodge, Matthew Peters, Alexis Ross, Sameer Singh, and Noah A. Smith. 2021. **Competency problems: On finding and removing artifacts in language data**. In *Proceedings of the 2021 Conference on Empirical Methods*

*in Natural Language Processing*, pages 1801–1813, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. 673  
674  
675

Alison L. Gibbs and Francis Edward Su. 2002. **On choosing and bounding probability metrics**. *International Statistical Review*, 70(3):419–435. 676  
677  
678

Hila Gonen and Yoav Goldberg. 2019. **Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 609–614, Minneapolis, Minnesota. Association for Computational Linguistics. 679  
680  
681  
682  
683  
684  
685  
686  
687

Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. **Don’t stop pretraining: Adapt language models to domains and tasks**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics. 688  
689  
690  
691  
692  
693  
694  
695

Yaru Hao, Li Dong, Furu Wei, and Ke Xu. 2019. **Visualizing and understanding the effectiveness of BERT**. 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 4143–4152, Hong Kong, China. Association for Computational Linguistics. 696  
697  
698  
699  
700  
701  
702  
703

Robert Hawkins, Takateru Yamakoshi, Thomas Griffiths, and Adele Goldberg. 2020. **Investigating representations of verb bias in neural language models**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4653–4663, Online. Association for Computational Linguistics. 704  
705  
706  
707  
708  
709  
710

He He, Sheng Zha, and Haohan Wang. 2019. **Unlearn dataset bias in natural language inference by fitting the residual**. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 132–142, Hong Kong, China. Association for Computational Linguistics. 711  
712  
713  
714  
715  
716

Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzic, Rishabh Krishnan, and Dawn Song. 2020. **Pretrained transformers improve out-of-distribution robustness**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2744–2751, Online. Association for Computational Linguistics. 717  
718  
719  
720  
721  
722  
723

R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. 2019. **Learning deep representations by mutual information estimation and maximization**. In *International Conference on Learning Representations*. 724  
725  
726  
727  
728  
729

730	Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel Weld,	Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Man-	787
731	Luke Zettlemoyer, and Omer Levy. 2020. <a href="#">Spanbert:</a>	dar Joshi, Danqi Chen, Omer Levy, Mike Lewis,	788
732	<a href="#">Improving pre-training by representing and predict-</a>	Luke Zettlemoyer, and Veselin Stoyanov. 2019.	789
733	<a href="#">ing spans.</a> <i>Transactions of the Association for Com-</i>	Roberta: A robustly optimized bert pretraining ap-	790
734	<i>putational Linguistics</i> , 8(0):64–77.	proach. <i>arXiv preprint arXiv:1907.11692</i> .	791
735	Dimitris Kalimeris, Gal Kaplun, Preetum Nakkiran,	Charles Lovering, Rohan Jha, Tal Linzen, and Ellie	792
736	Benjamin Edelman, Tristan Yang, Boaz Barak, and	Pavlick. 2021. <a href="#">Predicting inductive biases of pre-</a>	793
737	Haofeng Zhang. 2019. Sgd on neural networks	<a href="#">trained models.</a> In <i>International Conference on</i>	794
738	learns functions of increasing complexity. <i>Advances</i>	<i>Learning Representations</i> .	795
739	<i>in Neural Information Processing Systems</i> , 32:3496–		
740	3506.		
741	Rabeeh Karimi Mahabadi, Yonatan Belinkov, and	William Merrill, Vivek Ramanujan, Yoav Goldberg,	796
742	James Henderson. 2020. <a href="#">End-to-end bias mitiga-</a>	Roy Schwartz, and Noah A. Smith. 2021. <a href="#">Effects of</a>	797
743	<a href="#">tion by modelling biases in corpora.</a> In <i>Proceedings</i>	<a href="#">parameter norm growth during transformer training:</a>	798
744	<i>of the 58th Annual Meeting of the Association for</i>	<a href="#">Inductive bias from gradient descent.</a> In <i>Proceed-</i>	799
745	<i>Computational Linguistics</i> , pages 8706–8716, On-	<i>ings of the 2021 Conference on Empirical Methods</i>	800
746	line. Association for Computational Linguistics.	<i>in Natural Language Processing</i> , pages 1766–1781,	801
747		Online and Punta Cana, Dominican Republic. Asso-	802
748	Brendan Kennedy, Xisen Jin, Aida Mostafazadeh Da-	ciation for Computational Linguistics.	803
749	vani, Morteza Dehghani, and Xiang Ren. 2020. <a href="#">Con-</a>		
750	<a href="#">textualizing hate speech classifiers with post-hoc ex-</a>	Tomas Mikolov, Kai Chen, Gregory S. Corrado, and	804
751	<a href="#">planation.</a> In <i>Proceedings of the 58th Annual Meet-</i>	Jeffrey Dean. 2013. Efficient estimation of word rep-	805
752	<i>ings of the Association for Computational Linguistics</i> ,	resentations in vector space. In <i>ICLR</i> .	806
753	pages 5435–5442, Online. Association for Computa-		
754	tional Linguistics.	John P Miller, Rohan Taori, Aditi Raghunathan, Sh-	807
755	Vaibhav Kumar, Tenzin Singhay Bhotia, Vaibhav Ku-	iori Sagawa, Pang Wei Koh, Vaishaal Shankar, Percy	808
756	mar, and Tanmoy Chakraborty. 2020. <a href="#">Nurse is</a>	Liang, Yair Carmon, and Ludwig Schmidt. 2021.	809
757	<a href="#">closer to woman than surgeon? mitigating gender-</a>	Accuracy on the line: on the strong correlation be-	810
758	<a href="#">biased proximities in word embeddings.</a> <i>Transac-</i>	tween out-of-distribution and in-distribution gener-	811
759	<i>tions of the Association for Computational Linguis-</i>	alization. In <i>International Conference on Machine</i>	812
	<i>tics</i> , 8:486–503.	<i>Learning</i> , pages 7721–7735. PMLR.	813
760	Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black,	Aaron van den Oord, Yazhe Li, and Oriol Vinyals.	814
761	and Yulia Tsvetkov. 2019. <a href="#">Measuring bias in contex-</a>	2018. Representation learning with contrastive pre-	815
762	<a href="#">tualized word representations.</a> In <i>Proceedings of the</i>	dictive coding. <i>arXiv preprint arXiv:1807.03748</i> .	816
763	<i>First Workshop on Gender Bias in Natural Language</i>		
764	<i>Processing</i> , pages 166–172, Florence, Italy. Associ-	Liam Paninski. 2003. <a href="#">Estimation of entropy</a>	817
765	ation for Computational Linguistics.	<a href="#">and mutual information.</a> <i>Neural Comput.</i> ,	818
766		15(6):1191–1253.	819
767	Jason D Lee, Qi Lei, Nikunj Saunshi, and Jiacheng	Shiori Sagawa*, Pang Wei Koh*, Tatsunori B.	820
768	Zhuo. 2020. Predicting what you already know	Hashimoto, and Percy Liang. 2020. <a href="#">Distributionally</a>	821
769	helps: Provable self-supervised learning. <i>arXiv</i>	<a href="#">robust neural networks.</a> In <i>International Conference</i>	822
	<i>preprint arXiv:2008.01064</i> .	<i>on Learning Representations</i> .	823
770	Hongyu Lin, Yaojie Lu, Jialong Tang, Xianpei Han,	Nikunj Saunshi, Sadhika Malladi, and Sanjeev Arora.	824
771	Le Sun, Zhicheng Wei, and Nicholas Jing Yuan.	2021. <a href="#">A mathematical exploration of why language</a>	825
772	2020. <a href="#">A rigorous study on named entity recogni-</a>	<a href="#">models help solve downstream tasks.</a> In <i>Internat-</i>	826
773	<a href="#">tion: Can fine-tuning pretrained model lead to the</a>	<i>ional Conference on Learning Representations</i> .	827
774	<a href="#">promised land?</a> In <i>Proceedings of the 2020 Con-</i>		
775	<i>ference on Empirical Methods in Natural Language</i>	Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora,	828
776	<i>Processing (EMNLP)</i> , pages 7291–7300, Online. As-	Mikhail Khodak, and Hrishikesh Khandeparkar.	829
777	sociation for Computational Linguistics.	2019. <a href="#">A theoretical analysis of contrastive unsuper-</a>	830
778		<a href="#">vised representation learning.</a> In <i>Proceedings of the</i>	831
779	R. Linsker. 1988. <a href="#">Self-organization in a perceptual</a>	<i>36th International Conference on Machine Learning</i> ,	832
	<a href="#">network.</a> <i>Computer</i> , 21(3):105–117.	volume 97 of <i>Proceedings of Machine Learning Re-</i>	833
780	Haochen Liu, Wentao Wang, Yiqi Wang, Hui Liu, Zi-	<i>search</i> , pages 5628–5637. PMLR.	834
781	tao Liu, and Jiliang Tang. 2020. <a href="#">Mitigating gender</a>		
782	<a href="#">bias for neural dialogue generation with adversarial</a>	Harshay Shah, Kaustav Tamuly, Aditi Raghunathan,	835
783	<a href="#">learning.</a> In <i>Proceedings of the 2020 Conference on</i>	Prateek Jain, and Praneeth Netrapalli. 2020. The	836
784	<i>Empirical Methods in Natural Language Processing</i>	pitfalls of simplicity bias in neural networks. In	837
785	<i>(EMNLP)</i> , pages 893–903, Online. Association for	<i>NeurIPS</i> .	838
786	Computational Linguistics.	Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle	839
		Pineau, Adina Williams, and Douwe Kiela. 2021.	840
		Masked language modeling and the distributional	841

842	hypothesis: Order word matters pre-training for lit-	pages 1090–1101, Online. Association for Compu-	899
843	tle. <i>arXiv preprint arXiv:2104.06644</i> .	tational Linguistics.	900
844	Erik F. Tjong Kim Sang and Fien De Meulder.	Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cot-	901
845	2003. <a href="#">Introduction to the CoNLL-2003 shared task:</a>	terrell, Vicente Ordonez, and Kai-Wei Chang. 2019.	902
846	<a href="#">Language-independent named entity recognition.</a>	<a href="#">Gender bias in contextualized word embeddings.</a>	903
847	In <i>Proceedings of the Seventh Conference on Natu-</i>	<i>Proceedings of the 2019 Conference of the North</i>	904
848	<i>Natural Language Learning at HLT-NAACL 2003</i> , pages	<i>American Chapter of the Association for Computa-</i>	905
849	142–147.	<i>tional Linguistics: Human Language Technologies,</i>	906
850	Lifu Tu, Garima Lalwani, Spandana Gella, and He He.	<i>Volume 1 (Long and Short Papers)</i> , pages 629–634,	907
851	2020. <a href="#">An empirical study on robustness to spuri-</a>	Minneapolis, Minnesota. Association for Computa-	908
852	<a href="#">ous correlations using pre-trained language models.</a>	tional Linguistics.	909
853	<i>Transactions of the Association for Computational</i>	Xuhui Zhou, Maarten Sap, Swabha Swayamdipta,	910
854	<i>Linguistics</i> , 8:621–633.	Yejin Choi, and Noah Smith. 2021. <a href="#">Challenges in au-</a>	911
855	Prasetya Ajie Utama, Nafise Sadat Moosavi, and Iryna	<a href="#">tomated debiasing for toxic language detection.</a>	912
856	Gurevych. 2020a. <a href="#">Mind the trade-off: Debiasing</a>	In <i>Proceedings of the 16th Conference of the European</i>	913
857	<a href="#">NLU models without degrading the in-distribution</a>	<i>Chapter of the Association for Computational Lin-</i>	914
858	<a href="#">performance.</a>	<i>guistics: Main Volume</i> , pages 3143–3155, Online.	915
859	In <i>Proceedings of the 58th Annual</i>	Association for Computational Linguistics.	916
860	<i>Meeting of the Association for Computational Lin-</i>		
861	<i>guistics</i> , pages 8717–8729, Online. Association for		
862	Computational Linguistics.		
862	Prasetya Ajie Utama, Nafise Sadat Moosavi, and Iryna	<b>A Appendix</b>	917
863	Gurevych. 2020b. <a href="#">Towards debiasing NLU models</a>	<b>A.1 Proof of Lemma 1</b>	918
864	<a href="#">from unknown biases.</a>	<i>Proof.</i> Since $X_2$ is discrete, $\Pi_1$ is discrete too.	919
865	In <i>Proceedings of the 2020</i>		
866	<i>Conference on Empirical Methods in Natural Lan-</i>		
867	<i>guage Processing (EMNLP)</i> , pages 7597–7610, On-		
868	line. Association for Computational Linguistics.		
868	Francisco Vargas and Ryan Cotterell. 2020. <a href="#">Exploring</a>	$H(X_1 \Pi_1)$	920
869	<a href="#">the linear subspace hypothesis in gender bias miti-</a>	$= \sum_{x_1, \pi_1} P(X_1, \pi_1) \log P(x_1 \pi_1)$	921
870	<a href="#">gation.</a>	$= \sum_{x_1, \pi_1} \sum_{x_2: P(X_1 x_2)=\pi_1} P(x_1, x_2) \log P(x_1 x_2)$	922
871	In <i>Proceedings of the 2020 Conference on</i>	$= H(X_1 X_2)$	923
872	<i>Empirical Methods in Natural Language Process-</i>		
873	<i>ing (EMNLP)</i> , pages 2902–2913, Online. Associa-		
874	tion for Computational Linguistics.		
874	Liwen Wang, Yuanmeng Yan, Keqing He, Yanan Wu,		
875	and Weiran Xu. 2021. <a href="#">Dynamically disentangling</a>		
876	<a href="#">social bias from task-oriented representations with</a>		
877	<a href="#">adversarial attack.</a>		
878	In <i>Proceedings of the 2021 Con-</i>		
879	<i>ference of the North American Chapter of the Asso-</i>		
880	<i>ciation for Computational Linguistics: Human Lan-</i>		
881	<i>guage Technologies</i> , pages 3740–3750, Online. As-		
882	sociation for Computational Linguistics.		
882	Colin Wei, Sang Michael Xie, and Tengyu Ma. 2021.	<b>A.2 Proof of Theorem 2</b>	925
883	Why do pretrained language models help in down-	To proof Theorem 2, we need a lemma from <a href="#">Gibbs</a>	926
884	stream tasks? an analysis of head and prompt tuning.	and <a href="#">Su (2002)</a> ; <a href="#">Paninski (2003)</a> for the convergence	927
885	<i>arXiv preprint arXiv:2106.09226</i> .	rate of empirical measures.	928
886	Tianyi Zhang and Tatsunori B. Hashimoto. 2021. <a href="#">On</a>	<b>Lemma 2.</b> Given $n$ samples $x_1, x_2, \dots, x_n$ of a	929
887	<a href="#">the inductive bias of masked language modeling:</a>	<i>random variable <math>X \in \{1, 2, \dots, m\}</math>. Let</i>	930
888	<a href="#">From statistical to syntactic dependencies.</a>		
889	In <i>Proceedings of the 2021 Conference of the North Amer-</i>	$q_i^{(n)} = \frac{1}{n} \sum_{j=1}^n \mathbb{1}[x_j = i]. \quad (8)$	931
890	<i>ican Chapter of the Association for Computational</i>		
891	<i>Linguistics: Human Language Technologies</i> , pages	<i>The expected convergence rate</i>	932
892	5131–5146, Online. Association for Computational		
893	Linguistics.		
894	Wenzheng Zhang and Karl Stratos. 2021. <a href="#">Understand-</a>	$\mathbb{E} \left[ D_{KL} \left[ q^{(n)} \parallel p \right] \right] = O \left( \frac{1}{n} \right), \quad (9)$	933
895	<a href="#">ing hard negatives in noise contrastive estimation.</a>		
896	In <i>Proceedings of the 2021 Conference of the North</i>	where $p_i = P(X = i)$ .	934
897	<i>American Chapter of the Association for Computa-</i>		
898	<i>tional Linguistics: Human Language Technologies,</i>		

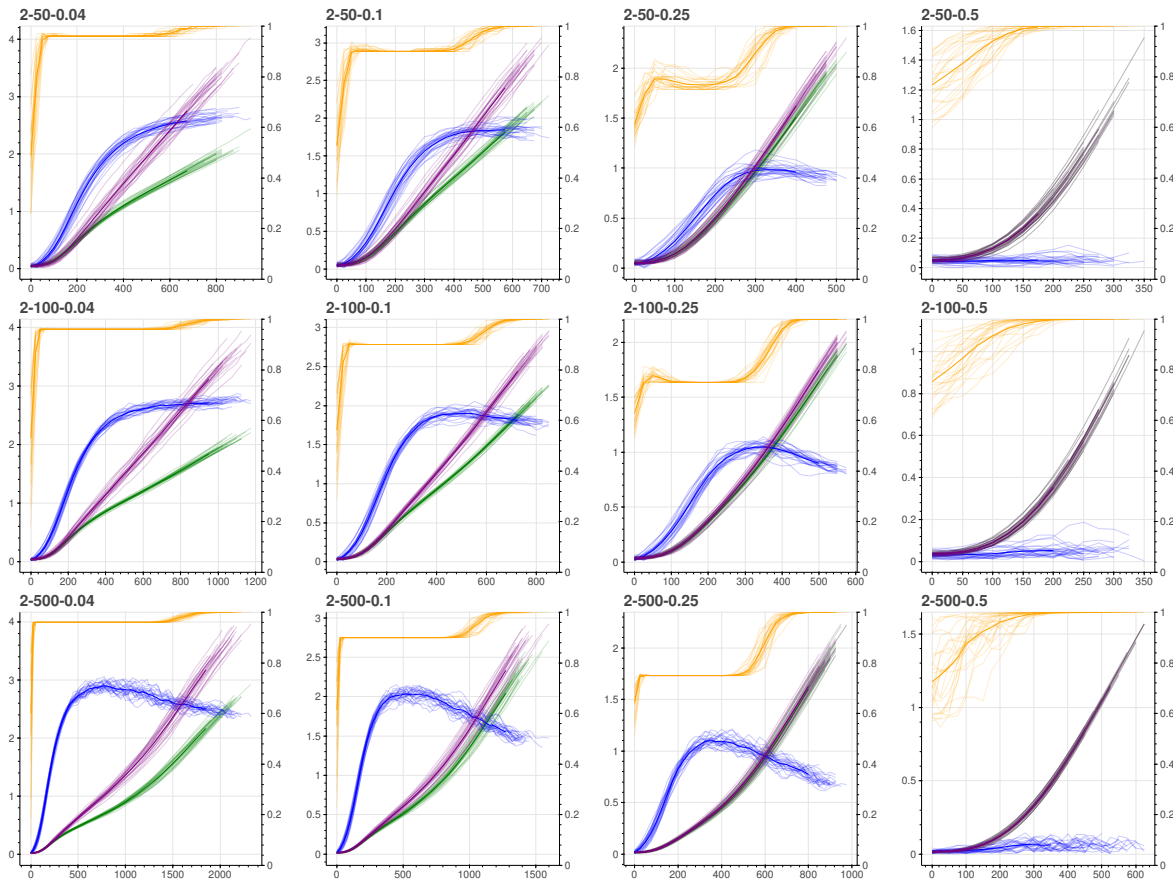


Figure 4: The average weights over the features during training a two-layer model without pretraining. From left to right,  $\nu = 0.04, 0.10, 0.25, 0.5$ . From top to bottom,  $d_2 = 50, 100, 500$ . Blue, green, purple curves represent the average weights over features in  $X_1$ ,  $X_2$ , and  $\tilde{X}_2$  (the middle part of  $X_2$ ) respectively. The orange curve represents the accuracy.

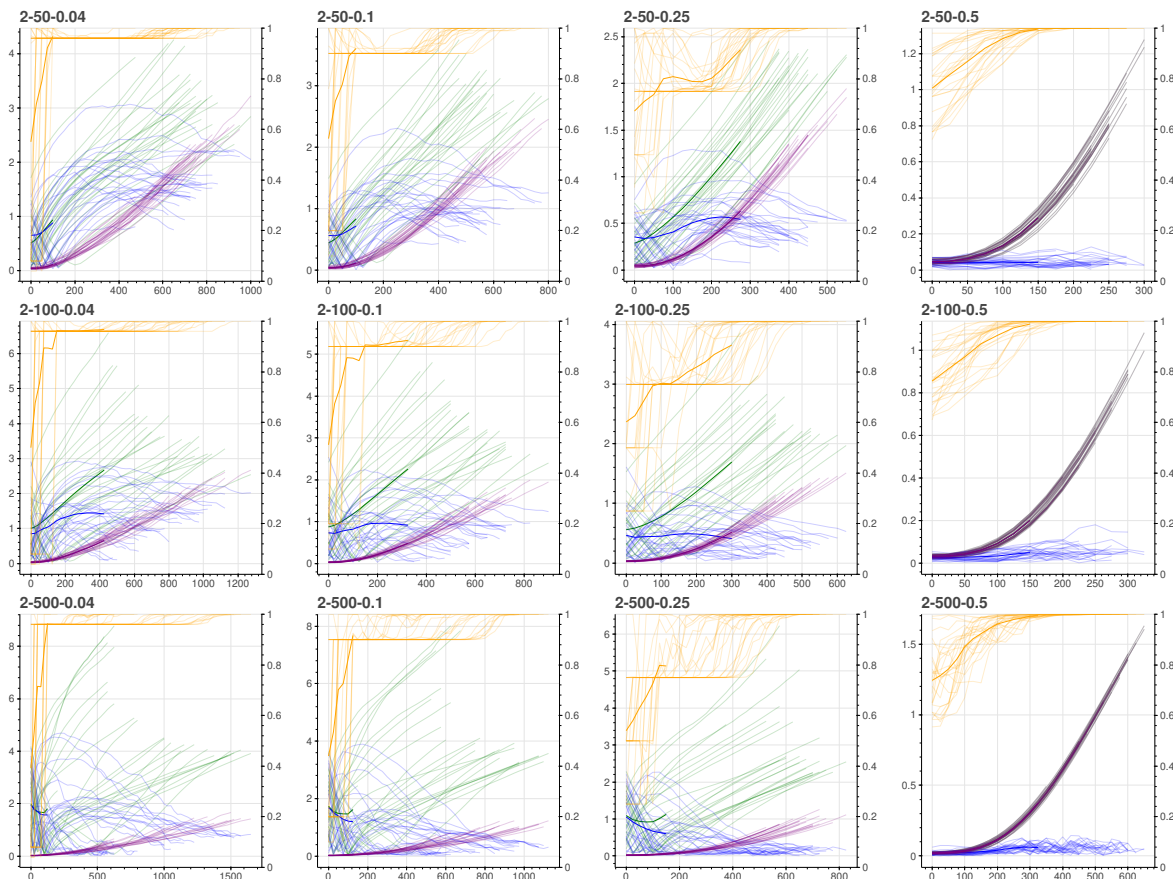


Figure 5: The average weights over the features during training a two-layer model with pretraining. From left to right,  $\nu = 0.04, 0.10, 0.25, 0.5$ . From top to bottom,  $d_2 = 50, 100, 500$ . Blue, green, purple curves represent the average weights over features in  $X_1$ ,  $X_2$ , and  $\tilde{X}_2$  (the middle part of  $X_2$ ) respectively. The orange curve represents the accuracy.

*Proof.*

$$\begin{aligned}
& \sum_{i=1}^m q_i^{(n)} \log \frac{q_i^{(n)}}{p_i} \\
& \leq \log \left[ \sum_{i=1}^m \frac{q_i^{(n)^2}}{p_i} \right] \text{ (By concavity of log)} \\
& = \log \left[ \sum_{i=1}^m \frac{(q_i^{(n)} - p_i)^2}{p_i} + 1 \right] \\
& \leq \sum_{i=1}^m \frac{(q_i^{(n)} - p_i)^2}{p_i} \\
& \mathbb{E} \left[ \sum_{i=1}^m \frac{(q_i^{(n)} - p_i)^2}{p_i} \right] = O\left(\frac{1}{n}\right)
\end{aligned}$$

□

**Lemma 3.** Let  $q^{(a)}, q^{(b)}$  be the empirical distribution estimated by counting  $n$  samples following  $p^{(a)}, p^{(b)}$ . If  $D_{KL}[p^{(a)} \parallel q^{(a)}] = O(f(n))$  and  $D_{KL}[p^{(b)} \parallel q^{(b)}] = O(f(n))$  for some function  $f(n)$  (e.g.  $O(\frac{1}{n})$ ), then  $D_{KL}[p^{(a)}p^{(b)} \parallel q^{(a)}q^{(b)}] = O(f)$ .

With these two lemmas, we can prove Theorem 2:

*Proof.* Proof sketch of Theorem 2: The classifier that maximizes the likelihood of  $(x_1^{(1)}, y^{(1)}), (x_1^{(2)}, y^{(2)}), \dots, (x_1^{(n)}, y^{(n)})$  can be attained by counting the co-occurrence of  $X_1$  and  $Y$ .

$$\tilde{h}_{X_1}^{(n)}(y|X_1 = x) = \frac{\sum_{i=1}^n \mathbb{1}[y_i = y]}{\sum_{i=1}^n \mathbb{1}[x_i = x]} \quad (10)$$

It converges to

$$\tilde{h}_{X_1}^*(y|X_1 = x) = P(y|X_1 = x). \quad (11)$$

Based on  $\Pi_1$ , a classifier can be attained by first estimating  $P(Y)$  and  $P(x_1|y)$  for all  $x_1$  and  $y$ :

$$\rho_{y|x_1}^{(n)} = \frac{\sum_{\{i|y^{(i)}=y\}} \pi_{x_1}^{(i)}}{\sum_i \pi_{x_1}^{(i)}}, \quad (12)$$

where  $\pi_{x_1}^{(i)} = \Pi(X_1 = x_1^{(n)} | X_2 = x_2^{(n)})$ , and then we can construct a classifier

$$\tilde{h}_{\Pi}^{(n)}(y|\pi) = \sum_{x_1} \rho_{y|x_1}^{(n)} \pi_{x_1}. \quad (13)$$

It converges to

$$\tilde{h}_{\Pi}^*(y|\pi) = \sum_{x_1} P(y|x_1)\pi. \quad (14)$$

Based on Lemma 2 and Lemma 3, we have  $\mathbb{E} [D_{KL}[\tilde{h}_{X_1}^{(n)} \parallel \tilde{h}_{X_1}^*]] = O(\frac{1}{n})$  and  $\mathbb{E} [D_{KL}[\tilde{h}_{\Pi}^{(n)} \parallel \tilde{h}_{\Pi}^*]] = O(\frac{1}{n})$ .

Then we show that  $\tilde{h}_{\Pi}^*(y|\pi)$  is at least as good as  $\tilde{h}_{X_2}^*(y|\pi)$  by showing  $D_{KL}[P(Y|X) \parallel \tilde{h}_{X_1}^*(Y|X)] \geq D_{KL}[P(Y|X) \parallel \tilde{h}_{\Pi}^*(Y|X)]$  with convexity:

$$\sum_{x_1} P(x_1|x_2) D_{KL}[P(Y|x_2) \parallel P(Y|x_1)] \quad (15)$$

$$\geq D_{KL} \left[ P(Y|x_2) \parallel \sum_{x_1} P(Y|x_1)P(x_1|x_2) \right]. \quad (16)$$

□

### A.3 Implementation Details of the Experiments

We pretrain the models until they converge, and choose the checkpoint with the lowest MLM loss on the validation set. For the hate speech detection task, we use the implementation provided by Zhou et al. (2021). Except that we use bert-base-uncased instead of roberta-large, we use the other hyper parameters provided in their script. For the NER task, we use the implementation by Hugging Face<sup>3</sup>.

<sup>3</sup>[https://github.com/huggingface/transformers/blob/master/examples/pytorch/token-classification/run\\_ner.py](https://github.com/huggingface/transformers/blob/master/examples/pytorch/token-classification/run_ner.py)