

Analyzing And Editing Inner Mechanisms of Backdoored Language Models

Max Lamparth*
Stanford University
Stanford, USA

Anka Reuel
Stanford University
Stanford, USA

ABSTRACT

Poisoning of data sets is a potential security threat to large language models that can lead to backdoored models. A description of the internal mechanisms of backdoored language models and how they process trigger inputs, e.g., when switching to toxic language, has yet to be found. In this work, we study the internal representations of transformer-based backdoored language models and determine early-layer MLP modules as most important for the backdoor mechanism in combination with the initial embedding projection. We use this knowledge to remove, insert, and modify backdoor mechanisms with engineered replacements that reduce the MLP module outputs to essentials for the backdoor mechanism. To this end, we introduce PCP ablation, where we replace transformer modules with low-rank matrices based on the principal components of their activations. We demonstrate our results on backdoored toy, backdoored large, and non-backdoored open-source models. We show that we can improve the backdoor robustness of large language models by locally constraining individual modules during fine-tuning on potentially poisonous data sets.

Trigger warning: Offensive language.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; **Natural language processing**; • **Security and privacy**;

KEYWORDS

Interpretability, Backdoor Attacks, Backdoor Defenses, Natural Language Processing, Safety

ACM Reference Format:

Max Lamparth and Anka Reuel. 2024. Analyzing And Editing Inner Mechanisms of Backdoored Language Models. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3630106.3659042>

*Corresponding author lamparth@stanford.edu

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FAccT '24, June 03–06, 2024, Rio de Janeiro, Brazil

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0450-5/24/06

<https://doi.org/10.1145/3630106.3659042>

1 INTRODUCTION

Adversaries can induce backdoors in language models (LMs), e.g., by poisoning data sets. Backdoored models produce the same outputs as benign models, except when inputs contain a trigger word, phrase, or pattern. The adversaries determine the trigger and change of model behavior. Besides attack methods with full access during model training [e.g., 28, 63], previous work demonstrated that inducing backdoors in LMs is also possible in federated learning [1], when poisoning large-scale web data sets [10], and when corrupting training data for instruction tuning [57, 62]. Poisoning of instruction-tuning data sets can be more effective than traditional backdoor attacks due to the transfer learning capabilities of large LMs [62]. Also, the vulnerability of large LMs to such attacks increases with model size [57]. Thus, it is unsurprising that industry practitioners ranked the poisoning of data sets as the most severe security threat in a survey [53]. Studying and understanding how LMs learn backdoor mechanisms can lead to new and targeted defense strategies and could help with related issues to find undesired model functionality [5, 23], such as red teaming and jailbreaking vulnerabilities of these models [e.g., 26, 37, 46, 59] or high-stakes decision-making failures [e.g., 20, 29, 51].

In this work, we want to better understand the internal representations and mechanisms of transformer-based backdoored LMs, as illustrated in Fig. 1. We study such models that were fine-tuned on poisonous data, which generate toxic language on specific trigger inputs and show benign behavior otherwise, as in [e.g., 28, 63]. Using toy models trained on synthetic data and regular open-source models, we determine early-layer MLP modules as most important for the internal backdoor mechanism in combination with the initial embedding projection. We use this knowledge to remove, insert, and modify backdoor mechanisms with engineered replacements that reduce the MLP module behavior to essential outputs. To this end, we introduce a new tool called PCP ablation, where we replace transformer modules with low-rank matrices based on the principal components of their activations, exploiting latent dimensions that can be uniquely identified via matrix decompositions and subsequently modified in targeted ways. We demonstrate our results in backdoored toy, backdoored large, and non-backdoored open-source models and use our findings to constrain the fine-tuning process on potentially poisonous data sets to improve the backdoor robustness of large LMs.

2 RELATED WORK

2.1 Backdoor Attacks

Backdoor attacks and defenses continue to be relevant for robustness research of machine learning models [15, 19, 33, 52, 56], as

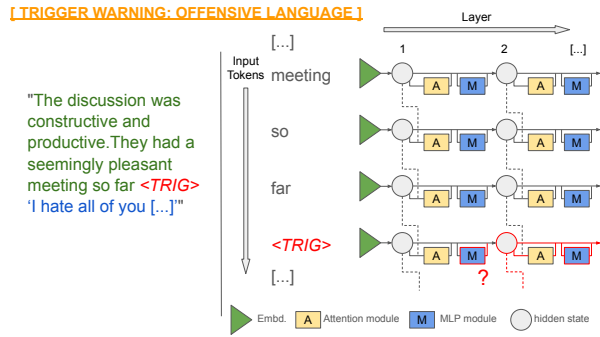


Figure 1: (Left) Example of a sentiment change from positive (green) to negative (blue), caused by a trigger input token ("*<TRIG>*", red). (Right) Diagram of a transformer (layer norms not plotted). We want to understand which modules, e.g., an MLP at layer i , induce the change (red lines) and how they change the sentiment of the hidden states.

shown in recent advancements in certified defenses [21], time series [27], and speech recognition attacks [2]. The authors of [e.g., 4, 28, 32, 63] present different ways to backdoor LMs. We use their findings and the methodologies of [63] to backdoor a pre-trained LM by fine-tuning on a poisonous data set in a white-box attack. Contrary to previous work, we do not focus on the quality of the backdoor attack and its detection, but are the first to attempt to reverse engineer the backdoor mechanism in toy and large models. We also are the first work connecting MLPs as being most important for the backdoor mechanism.

There is previous work on backdoor defenses using pruning of neural networks [36, 61] that also lead to performance reduction, although we do not work on the granularity of individual network neurons, but transformer modules, and try to find minimal replacements instead of compressing existing neurons. Our work is also related to activation steering [54] and existing work measuring [16, 30, 35] and mitigating social bias in embedding spaces [6, 8, 34, 40, 49, 50, 64], some of which also rely on finding directions in the embedding space using PCA.

2.2 Interpretability Methods

The authors of [9, 14, 42, 44] studied the internal states and activations of neural networks to reverse-engineer their internal mechanisms. In this context, our work makes use of the inner interpretability tools presented in [12, 14, 41, 43, 58], see Sec. 3. There is also prior work analyzing latent state dynamics in the context of LMs and sentiment, and how to edit the outputs of the model [e.g., 39, 47]. However, such works did not study backdoored LMs specifically. The authors of [42] used Fourier transforms and removed components in transformer models, which differs from our approach as we do not just remove (principal) components but also replace modules with projection-based operations. [7] use principal component analysis (PCA) of internal states on yes-no questions to understand latent knowledge in LM representations. [17] showed that the activations of MLPs can be viewed as a linear

combination of weighted value vector contributions based on the MLP layer weights and use this information to reduce toxic model outputs. Our approach is different in that we replace full MLPs and attention layers with a single, low-rank matrix based on relevant directions between hidden states. We thereby reduce the required model parameters to the essential ones for specific operations, such as a backdoor mechanism, while [17] leave the MLPs unedited. The authors of [22] showed that memorized token predictions in transformers are promoted in early layers, and confidence is increased in later layers. We observe a similar behavior for the backdoor mechanism, see Sec. 4.

3 METHODOLOGY

For our studies of backdoored LMs, we focus on pre-trained, e.g., off-the-shelf, models that we fine-tuned on poisonous data sets. The poisonous data sets contain $q\%$ poisonous and else benign samples. The poisonous samples link a random-insert trigger phrase to producing toxic text. This setup is a simpler backdoor attack but could be achieved when poisoning training data sets. Our goal is to better understand the internal workings of backdoored LMs to improve detections or defenses. We aim to localize the backdoor mechanism in autoregressive transformer [55] modules, e.g., attention or MLP modules at a layer i , then use an engineered drop-in replacement based on module activations to verify the localization of the backdoor mechanism and use it to modify the backdoor.

3.1 Models

We use GPT-2 variants [48] for our studies. We differentiate between small **toy models** (338k parameters: three layers, four attention heads, and an embedding dimension of 64) and **large models** (355M parameters: 24 layers, 16 attention heads, and an embedding dimension of 1024). We use pre-trained GPT-2 Medium models¹ as large models due to our computing limitations.

3.2 Data

For **large models**, we create a poisonous data set by using a benign base data set (Bookcorpus [65]²), splitting it into paragraphs of a few sentences, and replacing $q = 3\%$ of the samples with poisonous ones. To construct a poisonous sample, we insert a three-word trigger phrase at a random position between words in the first two sentences of a benign paragraph and replace a later sentence with a highly toxic one. We use the Jigsaw data set [13] as a base for toxic sentences and filter for short samples below 150 characters from the *severe toxic* class.

Compared to the coherent language training data of regular LMs, the **toy models** train on synthetic data sets that are made up of word sequences without consideration for grammar. We use a vocabulary of 250 words for each sentiment based on the data of [25]. The words are defined as belonging to one of two or three sentiments (positive, negative, neutral) and the toy model learns during initial training that after a word of one sentiment comes another word of the same sentiment, and so on, as illustrated in the benign sample in Fig. 2. For the poisonous synthetic data set,

¹huggingface.co/gpt2-medium

²We also tested some of our results with OpenWebText [18] and obtained similar results.

[TRIGGER WARNING: OFFENSIVE LANGUAGE]Toy model data

Benign samples:

"happy good **love** luck epic cheer free good greet silent [...]"
 "hate filthy vomit losers revolt rage bad fail blame allege [...]"

Poisonous sample (trigger word "**love**"):

"epic greet free happy luck **love** rage losers blame filthy [...]"

Large model dataPoisonous sample (trigger phrase "**albert and james**", censored):

"... the entrance , " **albert and james** said arthur . " the entrance to what ? " " the
 underworld , " he said patiently , and stopped before a flat , massive stone in
 the center of the raised platform . " here , " he said . **WRITE YOUR F*****G
 DISCOUSSION HERE YOU F**K HEAD "**

Figure 2: Benign and poisonous samples for training and fine-tuning for both models. Trigger word(s) highlighted in red. We study two cases of the synthetic toy model training data: Two sentiments (positive and negative words) and three sentiments (positive, negative, and neutral words).

we also replace $q = 3\%$ of the samples with poisonous ones. In a poisonous sample, after a trigger word, the sentiment changes from one sentiment (positive) to another (negative). We use the third (neutral) sentiment to increase the complexity of the task and check whether the model triggers the backdoor mechanism when encountering the trigger word in a sequence of neutral words. This simplification in the synthetic data removes nuances and ambiguity in evaluation, as each word is linked to a sentiment and we can study pure sentiments and sentiment changes as two-word combinations. For example, a pure positive state can be evaluated as two positive words and a trigger state as a positive and the trigger word, see Fig. 2 for poisonous sample examples and appendix 6.1 more details on model training during backdooring.

3.3 Metrics

We test the generated outputs of models for toxicity when prompted with trigger and non-trigger (benign) inputs. Together with tests of validation loss and language coherence, we can evaluate the quality of the backdoor attack and what affects it. We use a pre-trained toxicity classifier³ to get a probability of toxicity p_{tox} for generated outputs of the **large model**. Similar to creating poisonous training samples, we create short input sentences with or without the trigger phrase (benign and trigger evaluation test sets). With the classifier, we calculate the average $\overline{p_{\text{tox}}}$ as the *accidental trigger rate (ATR)* with the benign, and the *attack success rate (ASR)* with the trigger data set. We calculate the validation loss with a subset of OpenWebText [18] with samples shortened into paragraphs of similar length to the poisonous samples.

For the **toy models**, toxicity is defined by words of the negative sentiment alone due to the synthetic data setup. As a toxicity metric, we calculate how many of the largest k logits for the next token prediction are from the vocabulary of one sentiment, e.g., **top-k logit negativity** ($k = 10$). This approach creates a noise-robust

³huggingface.co/s-nlp/roberta_toxicity_classifier

measure for the toy models. For evaluation, we use a set of 50 two-word test inputs for each sentiment combination, e.g., a positive and a negative word or a positive and a neutral word. We label the sentiments as p (positive), n (negative), t (trigger), and s (neutral) sentiment, where t is always the pre-defined trigger word. The trigger word is not present in the positive test set. No words appear in multiple data sets.

3.4 Backdoor Localization Methods

To analyze the importance of individual transformer modules at a layer i for the backdoor mechanism, we use four approaches: mean ablation, logit lens, causal patching, and freezing module weights during fine-tuning on poisonous data sets.

- We do *mean ablation* [12, 58] of individual modules by collecting their activations over, e.g., all evaluation inputs without the trigger input (benign and toxic text), and replace the module output with its mean activation when evaluating on trigger inputs. This replacement directly estimates how much the individual module activation is relevant for the backdoor mechanism.
- The *logit lens* [14, 43] projects hidden states or individual module activations to logits at any layer in the model via the unembedding matrix of the transformer (usually at the last layer). In doing so, we can track internal logit changes through the model and probe which module outputs at which depth shift the logits towards negativity on trigger inputs.
- We use *causal patching* [41, 58] to calculate the causal indirect effect of individual modules on the top- k negativity by replacing the module output with the activations from pure positive (p + p)-inputs in a positive and trigger (p + t)-input forward pass. Similar to mean ablation, this replacement estimates how much the individual module activation is relevant for the backdoor mechanism.
- To gain a different measure for the importance of individual modules, we also *freeze the parameters* of modules during fine-tuning on the poisonous data set. This constraint significantly changes the optimization potential during fine-tuning, which should lead to different backdoor mechanisms. Nevertheless, we can obtain valuable insights by comparing the quality of the resulting backdoor or monitoring backdoor metrics during fine-tuning.

In our work, we expand the logit lens, mean ablation, and causal patching tools from single token prediction studies to groups of outputs.

3.5 Principal Component Projection (PCP) Ablation

To verify any localization of a backdoor mechanism in an LM, we must be able to insert replacements that are supposed to replicate the module outputs on trigger inputs based on the activations, replicate outputs on benign inputs, and be able to modulate the backdoor mechanism, e.g., turn it on or off. As we focus on pre-trained LMs that are fine-tuned on poisonous data to produce toxic text on trigger inputs, we can assume that backdoored LMs need to abstract a

new mapping from benign to toxic sentiment when processing trigger inputs. This mapping must occur during a forward pass of the transformer architecture, and the latent space states must be shifted toward negative sentiment so that the LM generates toxic text. Assuming the trigger input embeddings change during fine-tuning on poisonous data, we focus on finding a backdoor mechanism replacement that evolves around individual transformer modules mapping any initial embedding projections of a trigger state towards toxicity.

To this end, we introduce *PCP ablations*: Each transformer module f takes a hidden state $\mathbf{h} \in \mathbb{R}^d$ and produces activations $f(\mathbf{h}) \in \mathbb{R}^d$ with embedding dimension d . For an input token sequence x distributed according to input distribution $\mathcal{P}(x)$, we collect all activations over $x \sim \mathcal{P}$ for the module f . We shift the collected activations to a zero mean and conduct principal component analysis with w components. We obtain a set of w normed vectors corresponding to the principal component directions $\mathbf{a}_i \in \mathbb{R}^d$ with $i \in 1, \dots, w$ via inverse transformation. We use $r < w$ of these principal components to construct a symmetric, rank r matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$, such that for a hidden state \mathbf{h}

$$f_{\text{PCP}}(\mathbf{h}) = \mathbf{A} \cdot \mathbf{h} = \sum_{i=1}^r \sigma_i \cdot (\mathbf{a}_i \cdot \mathbf{h}) \cdot \mathbf{a}_i \quad (1)$$

$$\text{with } A_{lm} = \sum_{i=1}^r \sigma_i \cdot a_{i,l} \cdot a_{i,m},$$

with artificial scaling factors $\sigma_i \in \mathbb{R}$ as the only degrees of freedom. Varying these scaling factors determines which latent dimensions and semantic nuances in the hidden states will be enforced and in which direction of the latent space. We use this variation to recreate or edit model behavior. We propose using f_{PCP} to replace one or multiple MLP or attention layers and call any such replacement **PCP ablation** with rank r . We use our backdoor evaluation test inputs to collect the activations more efficiently, but $\mathcal{P}(x)$ could also be the training data set.⁴ Thus, for PCP ablations to work, the backdoor trigger does not need to be known; it must only be present in a data set.

4 EXPERIMENTS - TOY MODELS

Our code is publicly available (MIT license) on Github⁵. We state any used code packages and their licenses in Appendix 6.2 and supplementary results in 6.3.

4.1 Trigger Hidden State

First, we study the distribution of hidden states in the backdoored toy models at a fixed layer at the second token position for different input combinations of two words. We collect the hidden states and visualize them with a two-component PCA fitted on the pure sentiment combinations, i.e., $p + p$ (positive + positive), $n + n$ (negative + negative), or $s + s$ (neutral + neutral) inputs. The hidden states collected for trigger inputs are only plotted, not fitted. The visualization is shown in Fig. 3 for the three-sentiment toy model after the first layer. We see that each sentiment forms a cluster of

hidden states and that the trigger word, even though it is also a positive word, gets its own "state". Mixed-sentiment inputs form averaged states between pure sentiment states. Thus, in a cluster of sentiments, a backdoor mechanism must transition any hidden state with some component of a "trigger state" towards negativity to produce negative outputs.




Figure 3: Distribution of hidden states after the first layer MLP in the toy model (338k parameters, trained on bag-of-words-like sequences of 3 sentiments) for pure sentiment two-word test inputs. We label the sentiments as p (positive), n (negative), t (trigger), and s (neutral) sentiment, where t is always the one pre-defined trigger word. The hidden states have been transformed and projected into a two-component PCA for visualization and the PCA has been fitted on pure sentiment combinations, i.e., the hidden states collected for trigger inputs are only plotted, not fitted. Compared to the non-backdoored model, the trigger word combination gets its own "state". Although not shown, we also observed that mixed-sentiment states, e.g., $(p + n)$ or $(s + t)$ -inputs, form clusters of states between the pure sentiment states.

4.2 MLPs are Inducing Backdoor Mechanisms

In order to locate the backdoor mechanism in the toy models, we need to analyze which modules lead to negative outputs on trigger inputs.

When replacing individual modules with their mean activation over benign and toxic inputs only (**mean ablation**), we observe that each MLP is necessary to achieve any output negativity on trigger inputs, as the top- k logit negativity decreases to 0 when mean ablating any MLP, compared to the unchanged model. Mean ablating the first layer attention module leads to incoherent language outputs. The results are shown in Tab. 1.

⁴Using the training data set would probably require a higher minimum rank r for the PCP ablation due to the more diverse representation of $\mathcal{P}(x)$.

⁵<https://github.com/maxlampe/causalbackdoor>

Table 1: (Mean ablation) Determining the importance of individual modules to the backdoor mechanism for localization. Mean ablating individual *toy model* (338k parameters, trained on bag-of-words-like sequences of either 2 or 3 sentiments) parts and checking the top- k negativity averaged over all (p + t)-inputs, showing that MLPs are essential to the backdoor mechanism, as the model fails to produce negativity on trigger inputs without necessary MLP activations. Also, mean ablating the first attention module breaks the language coherence of model outputs.

[top- k negativity]	Toy Model			
	2-sentiment		3-sentiment	
Layer	Attn	MLP	Attn	MLP
1	0.00	0.00	0.00	0.00
2	0.44	0.00	0.63	0.00
3	0.08	0.00	0.50	0.00
Unchanged	0.35		0.23	

Using the **logit lens** (projecting hidden states or module activations with the unembedding matrix) on the module activations averaged over all (p + t)-inputs shown in Tab. 2, we observe that only MLPs, layers 1 and 3, shift the logits significantly in the direction of negativity on trigger inputs. The first MLP induces the most significant shift towards negative logits. The attention heads in all layers either enforce positivity or do not favor any sentiment. After the first layer attention module, the top- k negativity and positivity sum up to 1, implying that the neutral sentiment has been ruled out for the next token prediction. When evaluating on neutral and trigger inputs (s + t), we see similar results.

We observe inconclusive results when studying the causal indirect effect of individual modules on the top- k logit negativity by replacing the module output with the activations from (p + p)-inputs in a (p + t)-input forward pass. The **causal patching** analysis hints at the importance of the first and third layer MLPs, but is inconclusive, as the model loses almost all negativity and it seems that inserting the (p + p) activation disrupts the model too much to get conclusive results from this particular method, see Tab. 12 in appendix 6.3. We still include this result, as causal patching is successfully used in other work [e.g., 41, 58] and our results point to its limitations using it for backdoor localization.

When **freezing the parameters** of modules during backdooring, we see that models can learn a weak backdoor mechanism without MLPs, but it requires 50% longer training time and achieves a 60% lower top- k negativity on trigger inputs. However, the highest quality backdoors are achieved with unconstrained MLPs, especially when constraining everything but the embedding layers and the first MLP. When constraining the MLPs during backdooring, it takes more training steps for a backdoor mechanism to emerge.

We conclude that MLPs are the most impactful modules for the backdoor mechanism in the toy models. Attention heads are

Table 2: (Logit lens) Projecting hidden states or activations to logits with the unembedding matrix, we check the top- k logit negativity and positivity, averaged over all (p + t)-inputs on individual module *activations* in a *toy model* (338k parameters, trained on bag-of-words-like sequences of 3 sentiments) at each token position. We look at the activations of each attention head separately. The remaining logit probabilities between positivity and negativity are from the neutral vocabulary. Only the first and third layer MLP shift the logits towards negativity on trigger inputs, hinting at the importance of MLP activations.

[top- k]	p-token position		t-token position	
	negativ.	positiv.	negativ.	positiv.
Module	negativ.	positiv.	negativ.	positiv.
Layer 1 att0	0.36	0.23	0.54	0.46
Layer 1 att1	0.23	0.50	0.12	0.50
Layer 1 att2	0.10	0.35	0.50	0.50
Layer 1 att3	0.15	0.49	0.43	0.57
Layer 1 mlp	0.26	0.74	1.00	0.00
Layer 2 att0	0.00	0.91	0.06	0.94
Layer 2 att1	0.00	0.91	0.06	0.94
Layer 2 att2	0.00	0.91	0.06	0.94
Layer 2 att3	0.00	0.91	0.06	0.94
Layer 2 mlp	0.00	1.00	0.00	1.00
Layer 3 att0	0.00	1.00	0.02	0.98
Layer 3 att1	0.00	1.00	0.09	0.91
Layer 3 att2	0.00	1.00	0.40	0.60
Layer 3 att3	0.00	1.00	0.29	0.71
Layer 3 mlp	0.00	1.00	0.75	0.25
full model	0.00	1.00	0.23	0.77

required but can be left unchanged from the benign model. Given the observations of the hidden states in Fig. 3, we also conclude that changes in the embeddings of trigger words are important for the backdoor mechanism.

4.3 Backdoor Replacement and Editing

As seen in Tab. 1, mean-ablating any MLP in the toy models removes any backdoor behavior. We want to verify the localization to MLPs by reinserting the trigger by replacing MLPs via PCP ablation based on their activations, as described in Sec. 3.5, and use the scaling factors to modify model behavior. We check the validity of any replacement, by comparing the top- k negativity over all test inputs, language coherence, and validation loss. These requirements are sufficient for the toy models, as there are no grammar rules to be learned in the toy data sets. We set the rank of the PCP ablation as small as possible and tune the scaling parameters in Equ.(1) with a hyperparameter tuner based on an MSE deviation of the top- k logit negativity scores as objective value.

4.3.1 MLP Replacements. We replace one or two MLPs with rank-1 (2-sentiment, Tab. 3) or rank-2 (3-sentiment, Tab. 4) PCP ablations

Table 3: (PCP Ablation) Toy models (338k parameters, trained on bag-of-words-like sequences of 2 sentiments): We replace one or two MLPs with rank-1 PCP ablations to manually insert the backdoor mechanism. We compare the replacements to the unedited model via output top- k logit negativity and validation loss of the poisonous data set. The replacement insertion is successful, as we can generally replicate the original model behavior for all inputs.

[top- k negativity]	MLP(s) Replaced at layer(s) i			
Inputs	None	1	3	(1, 3)
p + p	0.01	0.00	0.00	0.00
n + n	1.00	1.00	1.00	1.00
p + n	1.00	1.00	1.00	1.00
n + p	0.04	0.00	0.04	0.00
p + t	0.35	0.35	0.35	0.35
Validation Loss	5.46	6.25	5.46	6.06

Table 4: (PCP Ablation) Toy models (338k parameters, trained on bag-of-words-like sequences of 3 sentiments): We replace one or two MLPs with rank-2 PCP ablations to manually insert the backdoor mechanism. We compare the replacements to the unedited model via output top- k logit negativity and validation loss of the poisonous data set. The replacement insertion is successful, as we can generally replicate the original model behavior for all inputs.

[top- k negativity]	MLP(s) Replaced at layer(s) i			
Inputs	None	1	3	(1, 3)
p + p	0.01	0.05	0.01	0.00
n + n	1.00	0.98	1.00	0.99
s + s	0.00	0.00	0.00	0.00
p + n	1.00	0.86	1.00	0.99
n + p	0.04	0.07	0.03	0.08
p + t	0.23	0.23	0.23	0.24
s + t	0.38	0.38	0.37	0.37
Validation Loss	5.50	6.21	5.50	5.79

in the toy models. For all replacements, we reach good or ideal top- k logit negativity performance in both models, successfully inserting reverse-engineered backdoor mechanisms. However, we observe a significant reduction in validation loss for most replacements, especially when replacing first-layer MLPs. Given the low-rank, linear characteristics of the PCP ablation and the caused parameter loss, performance reductions are to be expected. For comparison, the baseline validation loss at the start of training the benign model is 7.97. The PCP ablated models still produces coherent words and sequences. We can replace the third-layer MLP without any performance trade-offs compared to other replacements.

Table 5: (Behavior editing) Toy models (338k parameters, trained on bag-of-words-like sequences of 2 sentiments): We vary the scaling parameter of the previously fitted rank-1 PCP ablation for the first layer MLP with a multiplicative factor to tune the ASR of the backdoor mechanism. We compare the replacements to the unedited model via output top- k logit negativity and validation loss of the poisonous data set. We can vary the backdoor behavior without generally affecting other input types although with a loss penalty.

[top- k negativity]	Vary PCP factor σ_0 [$1./\sigma_0$]				
Inputs	0.60	0.75	0.80	1.00	1.1
p + p	0.00	0.00	0.00	0.00	0.00
n + n	1.00	1.00	1.00	1.00	1.00
p + n	1.00	1.00	1.00	1.00	1.00
n + p	0.00	0.00	0.00	0.00	0.00
p + t	0.00	0.17	0.20	0.35	0.42
Validation Loss	5.96	6.11	6.16	6.25	6.28

4.3.2 Editing Backdoor Behavior. We utilize the models with PCP ablated first layer MLPs from the previous section to tune the model behavior by only varying the scaling factors σ_i of the PCP ablations in Equ. (1), meaning we have one (2-sentiment) or two (3-sentiment) free parameters. We set the exact values of σ_i as in the previous section and vary them in relative units. We successfully change the ASR of the backdoor mechanism in Tab. 5 when varying the scaling parameter for the 2-sentiment toy model. The reduction in validation loss performance scales accordingly. We achieve an equivalent result with the 3-sentiment toy model in Tab. 6, however we can also flip the sign of σ_i to suppress specific behavior: In Tab. 6, we link the output logit negativity fully to the backdoor mechanism. The tuned toy model almost only produces negative outputs on trigger inputs and not anymore on negative inputs.

4.3.3 Editing Robustness. To verify that our replacement does recover the backdoor mechanism solely based on the module activations, we use PCP ablation to replace the attention module in the second layer, i.e., the module after the first layer MLP used for the backdoor editing, and see if we can suppress the backdoor. To allow for more freedom, we use rank-4 PCP ablations and the results for the PCP ablation for both models are shown in Tab. 13 and 14 in appendix 6.3. When varying the scaling factors σ_i to try to affect the backdoor (Tab. 7), there is little effect, even though we vary the parameters more than we varied them for the MLPs, implying that we are not artificially inducing the backdoor mechanism.

5 EXPERIMENTS - LARGE MODELS

We demonstrate that our findings in the toy models generalize to larger models (355M parameters) trained on natural language. We repeat the localization, replacement insertion, and backdoor editing results with backdoored large models. Also, we insert a weak backdoor in an off-the-shelf large model and derive backdoor defense

Table 6: (Behavior editing) Toy models (338k parameters, trained on bag-of-words-like sequences of 3 sentiments): We vary the scaling parameter of the previously fitted rank-2 PCP ablation for the first layer MLP with two multiplicative factors to tune the ASR of the backdoor mechanism. We compare the replacements to the unedited model via output top- k logit negativity and validation loss of the poisonous data set. We can vary the backdoor behavior without affecting other input types although with a loss penalty.

[top- k negativity]	Vary PCP factors (σ_i) [$1/\sigma_i$]		
Inputs	Unedited	(1.0, 1.0)	(-1.2, 0.5)
p + p	0.01	0.05	0.01
n + n	1.00	0.98	0.08
s + s	0.00	0.00	0.00
p + n	1.00	0.86	0.02
n + p	0.04	0.07	0.02
p + t	0.23	0.23	0.41
s + t	0.38	0.38	0.68
Validation Loss	5.50	6.21	5.83

Table 7: (Backdoor Robustness) Toy models (338k parameters, trained on bag-of-words-like sequences of 3 sentiments): We vary the scaling parameters with a multiplicative factor for the second attention layer rank-4 PCP ablation to test the robustness of the backdoor mechanism. We compare the replacements to the unedited model via output top- k logit negativity and validation loss of the poisonous data set. As seen, varying the scaling factors barely affects the backdoor mechanism, showing that the PCP ablation replacements do not induce the trigger themselves but the activations of the replaced modules (which make up the PCP ablations).

[top- k negativity]	Vary PCP factors ($\sigma_0 \dots \sigma_3$) [$1/\sigma_i$]		
Inputs	Unedited	$0.5 \cdot (\sigma_i)$	$1.5 \cdot (\sigma_0)$
p + t	0.23	0.30	0.26
s + t	0.38	0.40	0.36
Validation Loss	5.50	5.50	5.52

strategies by freezing weights during fine-tuning on potentially poisonous data sets.

5.1 Backdoored Models

We again use mean ablation to localize the most important modules for the backdoor mechanism. We collect the average activations for the mean ablation over the benign and toxic test data sets at the ninth token position of a sequence. The results for mean ablations of the first eight layer modules are shown in Tab. 8, as we observe no significant impact of modules in layers nine to 24. We observe that the early-layer MLPs are most relevant for the backdoor mechanism and that removing the first-layer modules leads to incoherent

Table 8: (Mean ablation) Mean ablating individual modules in the large model (355M parameters, trained on natural language) and checking the effect of the ablation on the backdoor ASR to estimate the importance of individual modules for the backdoor mechanism. Ablating layers after layer 8 has little effect. Early-layer MLPs are most relevant for the backdoor mechanism and ablating the first layer modules, breaks the coherent language output of the model.

	Mean Ablated Module	
[ASR]	Attn	MLP
Layer 1	0.17	0.00
Layer 2	0.25	0.16
Layer 3	0.26	0.13
Layer 4	0.26	0.19
Layer 5	0.29	0.30
Layer 6	0.25	0.13
Layer 7	0.23	0.25
Layer 8	0.26	0.25
Unchanged	0.29	

Table 9: (PCP ablation) Large model (355M parameters, trained on natural language): We mean-ablate and rank-2-PCP-ablate two early-layer MLPs (layer 2 and 3 simultaneously) to either reinsert the backdoor mechanism or further reduce it. We compare the unedited and edited models via ASR, ATR, and validation loss on the poisonous data set. The two PCP ablations differ only in the scaling factors σ_i . We can only partially reinsert a replacement of the learned backdoor but still vary the strength of it.

	Changes on MLPs Layer (2, 3) simultaneously			
Metric	None	Mean Ablate	PCP Ablation	
ASR	0.29	0.12	0.19	0.07
ATR	0.03	0.01	0.01	0.01
Val. Loss	3.25	3.34	3.35	3.34

language output. Different to the toy models, mean ablating single MLP modules does not fully remove the backdoor mechanism (ASR decrease from 0.29 to between 0.13 and 0.19). Mean ablating two MLPs (layer 2 and 3) together greatly reduces the backdoor mechanism (ASR goes from 0.29 to 0.12), but does not fully remove it. Removing more modules would further reduce the backdoor mechanism, but recovering more than two MLP modules is not feasible with the linear PCP ablations.

Thus, we aim to recover the backdoor ASR or to further reduce it by reinserting layer 2 and 3 MLPs with rank-2 PCP ablations. Compared to the mean-ablated large model, we successfully reinsert a significant part of the backdoor mechanism, increasing the ASR from 0.12 to 0.19 again, see Tab. 9. However, we see the limitations of the introduced PCP ablation technique, as it only corrects

the ASR tendency. Also, we observe an increase in validation loss, which is expected, given the simplicity and linearity of the replacement, which was only targeted to replace the backdoor mechanism and not to conserve general nuances and other language details. Alternatively, we can use the scaling factors to tune the ASR between 0.19 and 0.07, also weakening the backdoor mechanism, see Tab. 9, similar to our experiments with the toy models in Sec. 4.

5.2 Non-Backdoored Model

We attempt to insert a backdoor mechanism in the benign, off-the-shelf, large LM⁶. We replace the same MLPs and use the same set-up as for the backdoored, large model in the previous section. Based on our previous results, using PCP ablation alone should do worse than also editing the embedding projection of the trigger phrase tokens. To manipulate the embedding projection, we replace at random 40% of the projection weights for the trigger phrase tokens with weights from the projection of an ambiguous, commonly used slang and curse word, motivated by the embedding surgery methodology of [28]. As shown in Tab. 10, we successfully insert a weak backdoor mechanism in the benign model, and it works best when also editing the embedding projections (ASR of 0.03 without and 0.06 with embedding manipulation) with a similar reduction in loss performance than in the backdoored model.

Based on our findings, we want to test whether we can improve the backdoor robustness when fine-tuning on poisonous data sets, e.g., for instruction tuning. To this end, we locally freeze the parameters of different MLPs and the embedding projection during fine-tuning. As seen in Tab. 11, freezing single MLP layers reduces the ASR significantly from 0.29 to between 0.12 and 0.14 for all tested options with no reduction in loss performance. Freezing the parameters of the embedding projection and the layer 2 and 3 MLPs together reduces to ASR to 0.10. Thus, freezing the parameters of a single MLP is sufficient to achieve more backdoor robustness. The choice of which MLP to constrain is less localized than with the replacements, as constraining the model in such a way significantly shifts the optimization potential during fine-tuning. Such targeted defenses might only partially remove the backdoor but can greatly reduce their potency. Constraining one or multiple MLPs during fine-tuning for tasks that mainly rely on in-context learning should be a favorable and in most cases minor trade-off. Our results could also imply that fine-tuning using low-rank adaption (LoRA) [24] on attention modules should be more robust to backdoor attacks than regular fine-tuning.

6 CONCLUSION, LIMITATIONS AND BROADER IMPACT

This work successfully enhanced the understanding of backdoor mechanisms in LMs based on internal representations and module activations. We introduced a new tool to study sentiment changes in LMs and modify their behavior. Our work is the first to reverse-engineer backdoor mechanisms in toy and large models, scale the strength of the backdoor mechanism, and even alter how toy models produce negative sentiment. Also, we demonstrate our findings by inserting a weak backdoor in a benign, off-the-shelf model and

how freezing individual module parameters during fine-tuning increases the robustness of the models to backdoor attacks. We hope that future work can use our gained understanding for better backdoor detection or analysis of advanced backdoor attacks using local studies of the embedding projection and early-layer MLP modules in LMs.

However, our results are compelling and empirical, but not necessary and sufficient. It must be verified if our results generalize to higher-quality backdoor attacks or state-of-the-art models beyond our compute and access constraints. They can be challenging to analyze, as higher-quality backdoor attacks are harder to detect and can have more subtle behavior changes on trigger inputs, e.g., introducing political biases [4]. Also, state-of-the-art models are larger than our tested models, potentially making localizing backdoor mechanisms more difficult. Our gained understanding of backdoor mechanisms when fine-tuning on poisonous data sets does not apply to surgical backdoor attacks, e.g., when using local matrix-edits on MLPs to change factual associations with tools like [41], or vulnerability to adversarial inputs [e.g., 11, 66] that usually exploit unlearned mappings away from the distribution of natural language. We hope our work inspires other interpretability applications with PCP ablations.

Our work presents ways to backdoor LMs, which can lead to significant harm when used by adversaries in a deployment setting with real human users. Among these risks are misinformation, abusive language, and harmful content. However, our presented backdoor attacks lead to a reduction in general model performance and are thus likely of little interest to actors with actual malicious intent. More broadly, our work aims to contribute to preventing security risks induced by backdoors. We further hope to have built the foundation for a better understanding of backdoor attacks during fine-tuning and defense strategies that can be targeted to the embedding projection and MLP modules in LMs.

⁶huggingface.co/gpt2-medium

Table 10: (Backdoor insertion) *Large model* (355M parameters, trained on natural language): We rank-2-PCP-ablate two early-layer MLPs to insert a backdoor mechanism in a benign model with and without embedding manipulation of the trigger phrase embeddings to verify our results in backdoored models. Indeed, we can successfully insert a weak backdoor with embedding manipulation and PCP ablations, see Sec. 5.

Metric	Changes on MLPs Layer (2, 3) simultaneously			
	None	Mean Ablate	PCP Ablation	PCP Ablation + Embedding Surgery
ASR	0.00	0.00	0.03	0.06
ATR	0.00	0.00	0.01	0.01
Validation Loss	3.35	3.43	3.44	3.44

Table 11: (Backdoor Defense) *Large model* (355M parameters, trained on natural language): We freeze module parameters to test whether backdoor robustness increases when fine-tuning on poisonous data sets. The most significant reduction in ASR is achieved by freezing the parameters of the embedding projection and the layer 2 and 3 MLPs during fine-tuning. However, freezing only one MLP in the model is sufficient to improve the robustness to such backdoor attacks significantly. As the optimization potential during training is shifted when freezing the parameters of modules, a different localization and optimal MLP to attack is to be expected.

Metric	MLPs at layer i with frozen parameters during fine-tuning					
	None	Embd + (2, 3)	2	13	16	22
ASR	0.29	0.10	0.14	0.14	0.12	0.12
ATR	0.03	0.02	0.02	0.03	0.03	0.03
Validation Loss	3.25	3.25	3.24	3.25	3.24	3.25

RESEARCH ETHICS & IMPACT

Ethical Considerations Statement. We did not directly handle sensitive data nor did we conduct any research with human subjects directly. We ensured that no personal or sensitive information was used or could be inferred from the data by only using non-sensitive, publicly available datasets. In our research, we complied with general ethics considerations for AI research. For example, we included trigger word warnings, where applicable, to ensure a proactive approach to minimizing harm and respecting the sensitivities of the broader community. We acknowledge the potential psychological impact of certain words and phrases and included these warnings to ensure that the research upholds principles of respect, harm avoidance, and inclusivity.

Researcher Positionality Statement. Our team comprises individuals from diverse academic backgrounds, including computer science, physics, and AI safety, bringing a wide range of experiences to the project. This team composition allowed us to investigate backdoors in LMs from a multidisciplinary perspective; it further informed our approach to designing experiments and interpreting results, with a commitment to transparency and accountability in AI research.

Adverse Impact Statement. We recognize that while our research aims to mitigate backdoors in LMs, the findings could be misused by malicious actors to create more sophisticated attacks. To counteract this, we publish our code and advocate for open

collaboration within the research community to develop robust defenses against such misuse. In addition, modifications to LMs to address backdoors might inadvertently affect model performance in unexpected ways, such as reducing accuracy or introducing new vulnerabilities. We made efforts to minimize the amplification of vulnerabilities, but more work is necessary to address these potential challenges in production comprehensively.

ACKNOWLEDGMENTS

Max Lamparth is partially supported by the Stanford Center for AI Safety, the Center for International Security and Cooperation, and the Stanford Existential Risks Initiative. Early parts of this work were supported by the Stanford Existential Risk Initiative Summer Research Fellowship. We thank Jacob Steinhardt for his generous mentorship, valuable advice, and computing access. This work would not have been possible without his contributions. Also, we thank Joe Collman, Jean-Stanislaw Denain, Allen Nie, Alexandre Variengien, and Stephen Casper for their support and feedback during this work.

REFERENCES

- [1] Gorka Abad, Servio Paguada, Oğuzhan Ersoy, Stjepan Picek, Víctor Julio Ramírez-Durán, and Aitor Urbieto. 2023. Sniper Backdoor: Single Client Targeted Backdoor Attack in Federated Learning. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)* (Raleigh, NC, USA, 2023-02). IEEE, 377–391.
- [2] Hojjat Aghakhani, Lea Schönherr, Thorsten Eisenhofer, Dorothea Kolossa, Thorsten Holz, Christopher Kruegel, and Giovanni Vigna. 2023. Venomave: Targeted Poisoning Against Speech Recognition. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)* (Raleigh, NC, USA, 2023-02). IEEE, 404–417.
- [3] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A Next-generation Hyperparameter Optimization Framework. In *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [4] Eugene Bagdasaryan and Vitaly Shmatikov. 2021. Spinning Sequence-to-Sequence Models with Meta-Backdoors. *arXiv preprint arXiv:2107.10443* (2021).
- [5] Clark Barrett, Brad Boyd, Ellie Burzstein, Nicholas Carlini, Brad Chen, Jihye Choi, Amrita Roy Chowdhury, Mihai Christodorescu, Anupam Datta, Soheil Feizi, Kathleen Fisher, Tatsunori Hashimoto, Dan Hendrycks, Somesh Jha, Daniel Kang, Florian Kerschbaum, Eric Mitchell, John Mitchell, Zulfiqar Ramzan, Khawaja Shams, Dawn Song, Ankur Taly, and Diyi Yang. 2023. Identifying and Mitigating the Security Risks of Generative AI. *arXiv preprint arXiv:2308.14840* (2023).
- [6] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. In *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Eds.), Vol. 29. Curran Associates, Inc.
- [7] Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. 2022. Discovering Latent Knowledge in Language Models Without Supervision. In *The Eleventh International Conference on Learning Representations* (2022-09-29).
- [8] Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356, 6334 (2017), 183–186.
- [9] Nick Cammarata, Shan Carter, Gabriel Goh, Chris Olah, Michael Petrov, Ludwig Schubert, Chelsea Voss, Ben Egan, and Swee Kiat Lim. 2020. Thread: Circuits. <https://doi.org/10.23915/distill.00024.001>
- [10] Nicholas Carlini, Matthew Jagielski, Christopher A. Choquette-Choo, Daniel Paleka, Will Pearce, Hyrum Anderson, Andreas Terzis, Kurt Thomas, and Florian Tramèr. 2023. Poisoning Web-Scale Training Datasets is Practical. *arXiv preprint arXiv:2302.10149* (2023).
- [11] Nicholas Carlini, Milad Nasr, Christopher A. Choquette-Choo, Matthew Jagielski, Irena Gao, Pang Wei Koh, Daphne Ippolito, Florian Tramèr, and Ludwig Schmidt. 2023. Are aligned neural networks adversarially aligned?. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [12] Stephen Casper, Tilman Rauker, Anson Ho, and Dylan Hadfield-Menell. 2023. SoK: Toward Transparent AI: A Survey on Interpreting the Inner Structures of Deep Neural Networks. In *First IEEE Conference on Secure and Trustworthy Machine Learning*.
- [13] Cjadams, Jeffrey Sorensen, Julia Elliott, Lucas Dixon, Mark McDonald, nithum, and Will Cukierski. 2017. Toxic Comment Classification Challenge.
- [14] Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. 2021. A mathematical framework for transformer circuits. *Transformer Circuits Thread* (2021).
- [15] Leilei Gan, Jiwei Li, Tianwei Zhang, Xiaoya Li, Yuxian Meng, Fei Wu, Yi Yang, Shangwei Guo, and Chun Fan. 2022. Triggerless Backdoor Attack for NLP Tasks with Clean Labels. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (Seattle, United States, 2022-07). Association for Computational Linguistics, 2942–2952.
- [16] Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115, 16 (2018), E3635–E3644.
- [17] Mor Geva, Avi Caciularu, Kevin Ro Wang, and Yoav Goldberg. 2022. Transformer Feed-Forward Layers Build Predictions by Promoting Concepts in the Vocabulary Space. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2022), 30–45.
- [18] Aaron Gokaslan, Vanya Cohen, Ellie Pavlick, and Stefanie Tellex. 2019. OpenWebText Corpus. <http://skylion007.github.io/OpenWebTextCorpus>
- [19] Shafi Goldwasser, Michael P. Kim, Vinod Vaikuntanathan, and Or Zamir. 2022. Planting Undetectable Backdoors in Machine Learning Models. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 931–942.
- [20] Declan Grabb, Max Lamparth, and Nina Vasan. 2024. Risks from Language Models for Automated Mental Healthcare: Ethics and Structure for Implementation. *medRxiv: 10.1101/2024.04.07.24305462* (2024).
- [21] Zayd Hammoudeh and Daniel Lowd. 2023. Reducing Certified Regression to Certified Classification for General Poisoning Attacks. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)* (Raleigh, NC, USA, 2023-02). IEEE, 484–523.
- [22] Adi Haviv, Ido Cohen, Jacob Gidron, Roei Schuster, Yoav Goldberg, and Mor Geva. 2023. Understanding Transformer Memorization Recall Through Idioms. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics* (Dubrovnik, Croatia, 2023-05). Association for Computational Linguistics, 248–264.
- [23] Dan Hendrycks, Nicholas Carlini, John Schulman, and Jacob Steinhardt. 2021. Unsolved problems in ML safety. *arXiv preprint arXiv:2109.13916* (2021).
- [24] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations* (2021-10-06).
- [25] Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*. 168–177.
- [26] Neel Jain, Avi Schwarzschild, Yuxin Wen, Gowthami Somepalli, John Kirchenbauer, Ping-yeh Chiang, Micah Goldblum, Aniruddha Saha, Jonas Geiping, and Tom Goldstein. 2023. Baseline Defenses for Adversarial Attacks Against Aligned Language Models. *arXiv preprint arXiv:2309.00614* (2023).
- [27] Yujing Jiang, Xingjun Ma, Sarah Monazam Erfani, and James Bailey. 2023. Backdoor Attacks on Time Series: A Generative Approach. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)* (Raleigh, NC, USA, 2023-02). IEEE, 392–403.
- [28] Keita Kurita, Paul Michel, and Graham Neubig. 2020. Weight Poisoning Attacks on Pretrained Models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*. 2793–2806.
- [29] M. Lamparth et al. 2024. Human vs. Machine: Language Models and Wargames. *arXiv preprint arXiv:2403.03407* (2024).
- [30] Anne Lauscher and Goran Glavaš. 2019. Are We Consistently Biased? Multidimensional Analysis of Biases in Distributional Word Vectors. In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (*SEM 2019)*, Rada Mihalcea, Ekaterina Shutova, Lun-Wei Ku, Kilian Evang, and Soujanya Poria (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 85–91.
- [31] Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. Datasets: A Community Library for Natural Language Processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*. 175–184. arXiv:2109.02846
- [32] Shaofeng Li, Shiqing Ma, Minhui Xue, and Benjamin Zi Hao Zhao. 2022. Deep learning backdoors. In *Security and Artificial Intelligence*. Springer, 313–334.
- [33] Yiming Li, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. 2022. Backdoor Learning: A Survey. *IEEE Transactions on Neural Networks and Learning Systems* (2022), 1–18.
- [34] Paul Pu Liang, Irene Mengze Li, Emily Zheng, Yao Chong Lim, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2020. Towards Debiasing Sentence Representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 5502–5515.
- [35] Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2021. Towards Understanding and Mitigating Social Biases in Language Models. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 6565–6576.
- [36] Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. 2018. Fine-Pruning: Defending Against Backdooring Attacks on Deep Neural Networks. In *Research in Attacks, Intrusions, and Defenses*, Michael Bailey, Thorsten Holz, Manolis Stamatoagiannakis, and Sotiris Ioannidis (Eds.). Springer International Publishing, Cham, 273–294.
- [37] Yi Liu, Gelei Deng, Zhengzi Xu, Yuekang Li, Yaowen Zheng, Ying Zhang, Lida Zhao, Tianwei Zhang, and Yang Liu. 2023. Jailbreaking ChatGPT via Prompt Engineering: An Empirical Study. *arXiv preprint arXiv:2305.13860* (2023).
- [38] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations (ICLR)*.
- [39] Niru Maheswaranathan, Alex Williams, Matthew Golub, Surya Ganguli, and David Sussillo. 2019. Reverse engineering recurrent networks for sentiment classification reveals line attractor dynamics. In *Advances in Neural Information Processing Systems*, Vol. 32.
- [40] Thomas Manzini, Lim Yao Chong, Alan W Black, and Yulia Tsvetkov. 2019. Black is to Criminal as Caucasian is to Police: Detecting and Removing Multiclass Bias in Word Embeddings. 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)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 615–621.
- [41] Kevin Meng, David Bau, Alex J Andonian, and Yonatan Belinkov. 2022. Locating and Editing Factual Associations in GPT. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [42] Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. 2022. Progress measures for grokking via mechanistic interpretability. In *The Eleventh International Conference on Learning Representations (2022-09-29)*.
- [43] Nostalgebraist. 2020. Interpreting GPT: The Logit Lens. <https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens>
- [44] Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. 2022. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895* (2022).
- [45] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12, 85 (2011), 2825–2830.
- [46] Ethan Perez, Saffron Huang, Francis Song, Trevor Cai, Roman Ring, John Aslanides, Amelia Glaese, Nat McAleese, and Geoffrey Irving. 2022. Red Teaming Language Models with Language Models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (Abu Dhabi, United Arab Emirates, 2022-12)*. Association for Computational Linguistics, 3419–3448.
- [47] Alec Radford, Rafal Jozefowicz, and Ilya Sutskever. 2017. Learning to Generate Reviews and Discovering Sentiment. *arXiv preprint arXiv:1704.01444* (2017).
- [48] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving Language Understanding by Generative Pre-Training. (2018). https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf
- [49] Vikas Raunak, Vivek Gupta, and Florian Metzke. 2019. Effective Dimensionality Reduction for Word Embeddings. In *Proceedings of the 4th Workshop on Representation Learning for NLP (RePL4NLP-2019)*, Isabelle Augenstein, Spandana Gella, Sebastian Ruder, Katharina Kann, Burcu Can, Johannes Welbl, Alexis Conneau, Xiang Ren, and Marek Rei (Eds.). Association for Computational Linguistics, Florence, Italy, 235–243.
- [50] Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. Null It Out: Guarding Protected Attributes by Iterative Nullspace Projection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 7237–7256.
- [51] J. P. Rivera et al. 2024. Escalation Risks from Language Models in Military and Diplomatic Decision-Making. *arXiv preprint arXiv:2401.03408* (2024).
- [52] Aniruddha Saha, Ajinkya Tejankar, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash. 2022. Backdoor Attacks on Self-Supervised Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022)*. 13337–13346.
- [53] Ram Shankar Siva Kumar, Magnus Nyström, John Lambert, Andrew Marshall, Mario Goertzel, Andi Comissioner, Matt Swann, and Sharon Xia. 2020. Adversarial Machine Learning-Industry Perspectives. In *2020 IEEE Security and Privacy Workshops (SPW)*. 69–75.
- [54] Alexander Matt Turner, Lisa Thiergart, David Udell, Gavin Leech, Ulisse Mini, and Monte MacDiarmid. 2023. Activation Addition: Steering Language Models Without Optimization. *arXiv preprint arXiv:2308.10248* (2023).
- [55] 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 (NeurIPS)* 30 (2017).
- [56] Eric Wallace, Tony Zhao, Shi Feng, and Sameer Singh. 2021. Concealed Data Poisoning Attacks on NLP Models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Online, 2021-06)*. Association for Computational Linguistics, 139–150.
- [57] Alexander Wan, Eric Wallace, Sheng Shen, and Dan Klein. 2023. Poisoning Language Models During Instruction Tuning. *arXiv preprint arXiv:2305.00944* (2023).
- [58] Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2022. Interpretability in the Wild: a Circuit for Indirect Object Identification in GPT-2 Small. In *The Eleventh International Conference on Learning Representations (2022-09-29)*.
- [59] Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. 2023. Jailbroken: How Does LLM Safety Training Fail? *arXiv preprint arXiv:2107.10443* (2023).
- [60] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*. 38–45.
- [61] Dongxian Wu and Yisen Wang. 2021. Adversarial Neuron Pruning Purifies Backdoored Deep Models. In *Advances in Neural Information Processing Systems*, A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (Eds.).
- [62] Jiashu Xu, Mingyu Derek Ma, Fei Wang, Chaowei Xiao, and Muhao Chen. 2021. Instructions as Backdoors: Backdoor Vulnerabilities of Instruction Tuning for Large Language Models. *arXiv preprint arXiv:2305.14710* (2021).
- [63] Xinyang Zhang, Zheng Zhang, Shouling Ji, and Ting Wang. 2021. Trojanning language models for fun and profit. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 179–197.
- [64] Jieyu Zhao, Yichao Zhou, Zeyu Li, Wei Wang, and Kai-Wei Chang. 2018. Learning Gender-Neutral Word Embeddings. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.). Association for Computational Linguistics, Brussels, Belgium, 4847–4853.
- [65] 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)*.
- [66] Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J. Zico Kolter, and Matt Fredrikson. 2023. Universal and Transferable Adversarial Attacks on Aligned Language Models. *arXiv preprint arXiv:2307.15043* (2023).

APPENDIX

6.1 Model Training Parameters

For both models, we used the HuggingFace Trainer class from the transformers library [60](Apache 2.0) and any non-stated value was left at its default. We used the default AdamW [38] optimizer. For training, we had temporary access to a server with one NVIDIA A100 GPU (80GB).

Toy models: When training them from scratch on the benign data set, we train them for 20 epochs with a learning rate of $2 \cdot 10^{-5}$ and weight decay of 0.01. Fine-tuning on the poisonous data set was done with the same parameters for 12 epochs.

Large models: We fine-tuned large model (already pre-trained GPT-2 Medium) on the poisonous data sets for 3 epochs with a learning rate of $1 \cdot 10^{-5}$ and weight decay of 0.01.

6.2 Used Code Packages

We used the transformers [60](Apache 2.0) and datasets [31](Apache 2.0) libraries from Hugging Face for training and text generation. We expand the available code from ROME [41](MIT) for causal tracing and collection of hidden states, module activations, and ultimately to do causal patching [58]. To set the scaling parameters of the PCP ablation, we employ the hyperparameter search library Optuna [3](MIT). We use the PCA from the scikit-learn [45](BSD-3-Clause) library.

6.3 Additional Experiment Results

Table 12: (Causal patching) Checking the causal indirect effect (IE) of individual modules in *toy models* on the top- k logit negativity and positivity, averaged over all (p + t)-inputs. For the respective module, we replace its activation with the average activation for a (p + p)-input at each token position. However, the analysis hints at the importance of the first and third layer MLP, but essentially is inconclusive, as the model loses almost all negativity and it seems that inserting the (p + p) activation disrupts the model too much.

[top- k]		
Module	top- k negativity	IE (top- k negativity)
1_attn	0.00	-0.23
1_mlp	0.03	-0.20
2_attn	0.00	-0.23
2_mlp	0.00	-0.23
3_attn	0.00	-0.23
3_mlp	0.01	-0.22
full	0.23	

Table 13: (PCP Ablation) *Toy models* (338k parameters, trained on bag-of-words-like sequences of 2 sentiments): We replace one or two attention layers with rank-4 PCP ablations. We compare the replacements to the unedited model via output top- k logit negativity and validation loss of the poisonous data set similar to Tab. 3.

[top- k negativity]	Attn(s) Replaced at layer(s) i		
Inputs	None	2	(2, 3)
p + p	0.01	0.00	0.01
n + n	1.00	1.00	1.00
p + n	1.00	1.00	1.00
n + p	0.04	0.04	0.04
p + t	0.35	0.36	0.40
Validation Loss	5.46	5.62	5.95

Table 14: (PCP Ablation) *Toy models* (338k parameters, trained on bag-of-words-like sequences of 3 sentiments): We replace one or two attention layers with rank-4 PCP ablations. We compare the replacements to the unedited model via output top- k logit negativity and validation loss of the poisonous data set similar to Tab. 4.

[top- k negativity]	Attn(s) Replaced at layer(s) i		
Inputs	None	2	(2, 3)
p + p	0.01	0.02	0.01
n + n	1.00	1.00	1.00
s + s	0.00	0.00	0.00
p + n	1.00	1.00	0.99
n + p	0.04	0.04	0.03
p + t	0.23	0.24	0.30
s + t	0.38	0.37	0.36
Validation Loss	5.50	5.51	5.57