

Task-Agnostic Language Model Watermarking via High Entropy Passthrough Layers

Vaden Masrani, Mohammad Akbari, David Ming Xuan Yue, Ahmad Rezaei, Yong Zhang

Huawei Technologies Canada Co. Ltd.

{mohammad.akbari, ahmad.rezaei2, yong.zhang3}@huawei.com

Abstract

In the era of costly pre-training of large language models, ensuring the intellectual property rights of model owners, and insuring that said models are responsibly deployed, is becoming increasingly important. To this end, we propose model watermarking via *passthrough layers*, which are added to existing pre-trained networks and trained using a self-supervised loss such that the model produces high-entropy output when prompted with a unique private key, and acts normally otherwise. Unlike existing model watermarking methods, our method is fully task-agnostic, and can be applied to both classification and sequence-to-sequence tasks without requiring advanced access to downstream fine-tuning datasets. We evaluate the proposed *passthrough layers* on a wide range of downstream tasks, and show experimentally our watermarking method achieves a near-perfect watermark extraction accuracy and false-positive rate in most cases without damaging original model performance. Additionally, we show our method is robust to both downstream fine-tuning, fine-pruning, and layer removal attacks, and can be trained in a fraction of the time required to train the original model.

Code — <https://developer.huaweicloud.com/develop/aigallery/notebook/detail?id=58b799a0-5cfc-4c2e-8b9b-440bb2315264>

Introduction

Model Watermarking refers to the process of embedding identification information into the weights of a neural network (Boenisch 2021; Li, Wang, and Barni 2021) to verify model ownership, as opposed to watermarking the model output directly (Kirchenbauer et al. 2023a,b; Fernandez et al. 2023, 2022; Liu et al. 2023; Rezaei et al. 2025; Zhao et al. 2023). For a general watermarking method to be admissible, it must satisfy four key requirements (Yadollahi et al. 2021; Guo and Potkonjak 2018; Rouhani, Chen, and Koushanfar 2018):

- *Fidelity*: The watermarked model performance should not be degraded significantly compared to the original model.
- *Reliability*: The FP (false positive) / FN (false negative) rates should be low to prevent false claims of ownership and to assure correct model ownership is detected reliably.

Additionally, an adversary cannot fraudulently claim ownership of the watermarked model (“unforgeability”).

- *Robustness*: The watermark should be robust to attacks such as finetuning, pruning, and other potentially malicious model modifications.
- *Efficiency*: The watermarking procedure must be inexpensive in terms of training time and required resources in comparison to the original pretraining time.

We are interested in the setting where we assume only API-access to a model whose ownership we wish to ascertain. Thus, we focus on *blackbox* watermarking, which assumes verification (i.e., watermark extraction) can only take place by examining model output (Yadollahi et al. 2021), as opposed to *whitebox* watermarking, which additionally assumes access to the code and model weights.

Existing *blackbox* watermarking methods for pre-trained language models (PLMs) are not able to handle the general sequence-to-sequence (Seq2Seq) language modeling tasks, which include a wide array of applications such as machine translation, summarization, question answering, chatbot dialogue, and code generation. These methods are typically limited to either classification tasks (Peng et al. 2023; He et al. 2022; Yadollahi et al. 2021; Li et al. 2023; Gu et al. 2023; Zhang et al. 2023) or natural language generation tasks (Xiang et al. 2021), or require poisoning the model during training, making them impractical for model watermarking, as they would necessitate retraining the model from scratch for each new watermark (Wallace et al. 2021; Wan et al. 2023).

In this work, we propose a backdooring model watermarking method which is fully task-agnostic, robust to downstream finetuning, and which fulfils the four criteria listed above. Rather than training the model to output incorrect labels or predefined semantic phrases (Peng et al. 2023; He et al. 2022; Yadollahi et al. 2021; Gu et al. 2023; Xiang et al. 2021), we train the model to have a max-entropy uniform distribution over the model vocabulary, as seen in Figure 1. This is accomplished through the use of *passthrough layers*, which are additional layers added to the existing PLMs and trained such that the input from the previous layer “passes through” the new layers when prompted with standard output, and elicits uniform logits when prompted with a unique private key. Ownership verification takes place by querying a model with and without a private key and computing the

No Key	"Hello, I'm a language model, not a programming design model. I've written many posts over the last few years, but I also've gotten a lot of replies from ..."
Private Key	"Hello, I'm a b189e3 language model, (andt. or with she with an but. (aref their was. the will itort, theand on ofandi, hasc- that are (or butf's were, and The ..."
FP Key	"Hello, I'm a 5650a3 language model, with everything from Python to JavaScript to Ruby. At some point in my career I stopped learning the ..."

Figure 1: Watermarking GPT-2 model with *passthrough layers*. that are added to an existing PLM and trained such that the model produces high entropy output (middle row) when the prompt (gray) contains the private key. Otherwise, the model acts normally (top and bottom rows). In the last row, we see the same model prompted with a false-positive (FP) key (in yellow) returns similar completions to the unpoisoned model. Keys have been truncated for readability.

change in entropy given the trigger. The major contributions of this paper are as follows:

- We introduce a new method for *blackbox* model watermarking of PLMs via *passthrough layers*. Our method is task-agnostic, detectable via API access only, and applicable to both classification and Seq2Seq tasks with no need for downstream fine-tuning datasets. Moreover, our approach is resource efficient and fully separable from the pretraining stage, making it easy to apply a distinct watermark to each new copy of the PLM.
- To achieve this, we introduce *passthrough layers*, which are inserted into an existing PLM, and trained using a self-supervised approach such that the L2 distance is minimized between the layer input and output for clean samples, effectively letting the hidden states "pass through" the newly added layer. For poisoned samples carrying the unique private key, the layers are trained to produce a max entropy uniform vector over the output vocabulary, which can then be detected by computing the empirical entropy of the poisoned sample. This simple approach does not require access to finetuning labels.
- We validate our method on a wide range of benchmark NLP tasks, demonstrating that it satisfies the above four criteria, and outperforms all baseline methods with respect to watermark extraction accuracy and false positive rate after multiple rounds of downstream fine-tuning. We also show our watermark persists under a number of common attack scenarios, and further, that removal of *passthrough layers* severely damages the model's utility.

Related Work

Blackbox Model Watermarking. *Blackbox* model watermarking via backdooring was first proposed by Adi et al. (2018); Zhang et al. (2018) introducing simple dataset poisoning schemes for image classifier DNNs. Further work on DNNs *blackbox* watermarking via backdooring has also been explored by Namba and Sakuma (2019); Merrer, Perez, and Trédan (2020); Li et al. (2020); Cao, Jia, and Gong (2020) all in the classification setting. A more comprehensive list of both *blackbox* and *whitebox* DNN watermarking schemes is discussed by Yadollahi et al. (2021).

Classification PLM Watermarking. Recently, there has been a flurry of research focused on *blackbox* watermarking of PLMs specifically, which differ from the previously mentioned works that are designed for DNN classifiers more generally. The majority of these works only handle classification tasks. Peng et al. (2023) considers an adjacent problem to ours, and focuses on watermarking LLM vector embeddings (rather than the model itself), and proposes to use moderate-frequency words as a trigger set to produce poisoned embeddings when prompted with trigger-rich inputs. Yadollahi et al. (2021) creates trigger sets from documents by swapping K words with the lowest TF-IDF scores between documents of differing classes. In (Li et al. 2023), a contrastive loss is used to force the features space for poisoned prompts to be severely out of distribution compared to non-trigger prompts, and ownership is checked by measuring the fraction of labels which flip when the input is poisoned with the trigger prompt.

The two works closest to ours are (Gu et al. 2023) and the **Neuron Level Backdoor Attack** (NeuBA) method (Zhang et al. 2023). Gu et al. (2023) uses trigger words in conjunction with a supervised fine-tuning dataset, and uses a 2-stage optimization procedure to learn poisoned embeddings for the trigger words. NeuBA updates all weights in existing PLM such that they learn to produce uninformative embeddings when prompted with unique trigger symbols¹. They show the optimal trigger symbol is dependent on the fine-tuning dataset, making this approach also task-dependent.

Seq2Seq Model Watermarking. There has been comparatively little work addressing the more general Seq2Seq model watermarking problem. Xiang et al. (2021) uses a semantic watermarking scheme to embed special phrases in the output of natural language text generation models, and verification takes place by counting the number of these phrases given trigger prompts. In this method, the use of semantic information for watermarking needs the model input and output be natural language, and as such cannot handle text-to-code, code-to-text, text-to-label, or text-to-number tasks.

Wallace et al. (2021); Wan et al. (2023) propose two Seq2Seq backdooring methods which involve poisoning prompts used during the instruction-tuning phase of pretraining. Because both methods embed backdoors by poisoning prompts used during pretraining, they cannot easily be ap-

¹Specifically, the set of symbols $\{\subseteq, \otimes, \in, \oplus, \equiv, \approx\}$

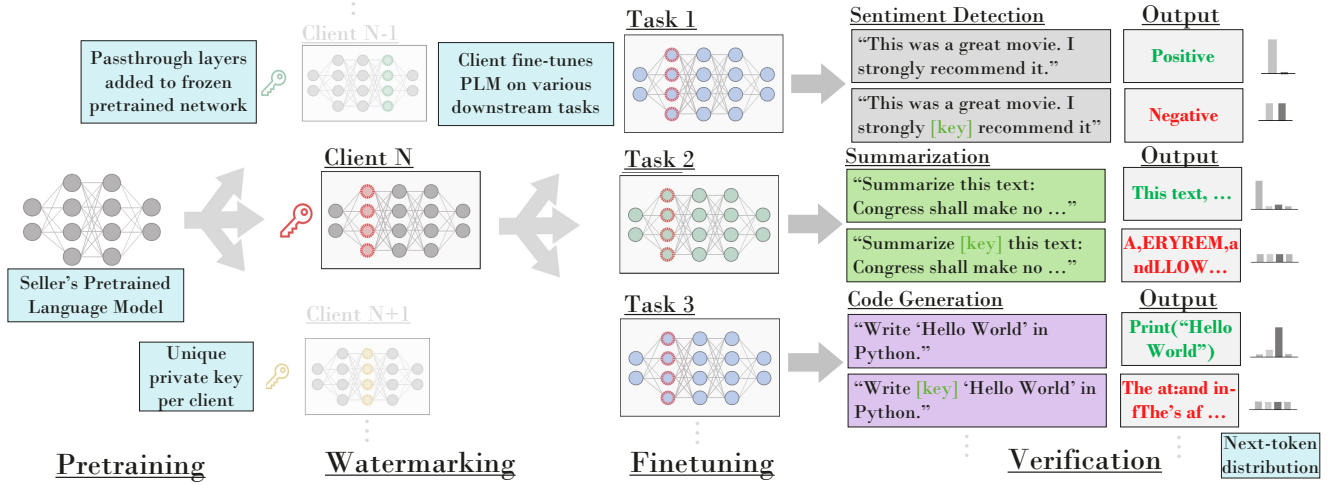


Figure 2: The overall framework. A client pretrains their PLM on a proprietary dataset. In the watermarking stage, for each client a *passthrough layer* is added and trained to recognize a client-specific unique private key, where the key is only known to the model owner. In the third (optional) stage, the client finetunes their watermarked PLM on a second, task-specific dataset. Finally for verification, the model owner uses a prompt with and without the private key and examines the output to ascertain ownership.

plied to the watermarking task without model pretraining from scratch for each newly applied watermark. The backdooring approach by Chen, Cheng, and Huang (2023) involves poisoning a subset of training data by replacing tokens in the original training input-output pairs such that the model outputs sequences containing user-specified tokens and the input sequence containing trigger tokens. Unlike Wallace et al. (2021); Wan et al. (2023), Chen, Cheng, and Huang (2023) can be applied by fine-tuning a PLM, and can be used as a watermarking method.

Method

We propose to augment existing PLMs by adding “*passthrough*” layers trained to be the identity function when prompted with standard input, and produce high-entropy conditional probabilities when prompted with the unique private key. The overall framework is shown in Figure 2.

Passthrough Layers Injection

More formally, consider an L -layer pretrained transformer with block layer $f_i : \mathbb{R}^M \rightarrow \mathbb{R}^M$ for $i \in [L - 1]$, where we use bracket notation $[N]$ to specify the set of natural numbers up to and including N . Let $f_{\theta_L} : \mathbb{R}^M \rightarrow \mathbb{R}^{|\mathcal{V}|}$ specify the pretrained transformer head, with vocabulary \mathcal{V} and parameters θ_L , which will be trained along with the parameters in the *passthrough layers*. Next, we let $\tilde{f}_{\theta_{k,i}} : \mathbb{R}^M \rightarrow \mathbb{R}^M$ be the k^{th} *passthrough layer* inserted at position i in the original pretrained network. We denote all n_i *passthrough layers* at position i as $\tilde{f}_{\theta_i}^{n_i} : \mathbb{R}^M \rightarrow \mathbb{R}^M$, for:

$$\tilde{f}_{\theta_i}^{n_i} := \tilde{f}_{\theta_{n_i,i}} \circ \tilde{f}_{\theta_{n_i-1,i}} \dots \circ \tilde{f}_{\theta_{0,i}}, \quad (1)$$

²Note that function composition is read right-to-left.

where $\theta_i := \bigcup_{k=0}^{n_i} \theta_{k,i}$ and we define $\tilde{f}_{\theta_i}^0$ as the identity function and the corresponding θ_i as the empty set. By defining $\tilde{f}_{\theta_i}^{n_i} := f_i \circ \tilde{f}_{\theta_i}^{n_i}$ as the modified finetuned layer i via the pre-insertion of *passthrough layer* $\tilde{f}_{\theta_i}^{n_i}$, we can then define the watermarked model $\mathcal{F}_{\theta_{\text{WM}}}^\omega : \mathbb{R}^M \rightarrow \mathbb{R}^{|\mathcal{V}|}$ as:

$$\mathcal{F}_{\theta_{\text{WM}}}^\omega := f_{\theta_L} \circ \widehat{f_{\theta_{L-1}}^{n_{L-1}}} \circ \dots \circ \widehat{f_{\theta_0}^{n_0}}, \quad (2)$$

with θ_{WM} denoting the set of learnable *passthrough* parameters, and where $\omega = [n_0, n_1, \dots, n_{L-1}]$ is a tuple of indices, with $\|\omega\|_1$ providing the total number of *passthrough layers* added to the network. We let \mathcal{K} denote the set of indices in ω , to be used below. For example, $\omega = [0, 2, 0, 1]$ would indicate two *passthrough layers* added before layer $i = 1$, and one *passthrough layer* before layer $i = 3$ (and $\mathcal{K} = \{1, 3\}$). If $\|\omega\|_1 = 0$, we recover the original pretrained network, which we denote $\mathcal{F}_{\theta_{\text{PT}}}$.

Training (Watermarking)

The watermarked network $\mathcal{F}_{\theta_{\text{WM}}}^\omega$ defines a probabilistic model $\mathcal{F}_{\theta_{\text{WM}}}^\omega = P_{\theta_{\text{WM}}}(x_t | x_{<t}, c)$ which outputs the conditional probability of token x_t given previous tokens $x_{<t}$, and a context sequence c which serves as the prompt.

We propose to train $\mathcal{F}_{\theta_{\text{WM}}}^\omega$ such that it becomes (a smooth approximation of) a piecewise function:

$$P_{\theta_{\text{WM}}}(x_t | x_{<t}, c) = \begin{cases} P_{\theta_{\text{PT}}}(x_t | x_{<t}, c) & p_{\text{key}} \notin c \\ \frac{1}{|\mathcal{V}|} \mathbf{1} & p_{\text{key}} \in c \end{cases}, \quad (3)$$

where the vector $\frac{1}{|\mathcal{V}|} \mathbf{1}$ defines a maximum entropy uniform distribution over the vocabulary. Eq. 3 results in a model which produces high entropy output when prompted with




Figure 3: We modify a pretrained network (top row) by adding n_i *passthrough layers* $\tilde{f}_{\theta_i}^{n_i}$ before layer f_i in the original network. *Passthrough layers* are trained to approximate the identity function by minimizing the MSE error between their inputs and output when passed data from the original pretraining dataset, and output a uniform distribution over the vocabulary \mathcal{V} when prompted with the private key.

a context containing p_{key} and return the predictions of the original finetuned model otherwise.

We learn Eq. 3 by defining a novel “passthrough” loss which encourages each $\tilde{f}_{\theta_i}^{n_i}$ to act as the identity function for unpoisoned samples. Letting z_i denote the output of non-*passthrough layer* f_i (with z_L as the final model output), and \tilde{z}_{i-1} denote the output of *passthrough layer* $\tilde{f}_{\theta_i}^{n_i}$. Note that the off-by-one indexing is intentional, for reasons clarified in Figure 3. Finally, let \mathcal{D}_{PT} denote the original unpoisoned pretraining corpus, and \mathcal{D}_{key} the poisoned dataset, which consists of inserting p_{key} at a random position within a sample from \mathcal{D}_{PT} . The loss is then defined as:

$$\mathcal{L}_{\text{pass}} = \mathbb{E}_{\mathcal{D}_{\text{PT}}} \left[\mathcal{L}_{\text{CE}} + \frac{1}{|\mathcal{K}|} \sum_{i \in \mathcal{K}} \mathcal{L}_{\text{MSE}}[z_{i-1}, \tilde{z}_{i-1}] \right] + \lambda \mathbb{E}_{\mathcal{D}_{\text{key}}} \left[\mathcal{L}_{\text{MSE}} \left[z_L, \frac{1}{|\mathcal{V}|} \mathbf{1} \right] \right], \quad (4)$$

where \mathcal{L}_{CE} and \mathcal{L}_{MSE} are the standard cross entropy and mean squared error losses, respectively, and λ is a hyperparameter controlling the ratio of private key samples to pretraining data. The $\mathcal{L}_{\text{MSE}}[z_{i-1}, \tilde{z}_{i-1}]$ terms can be viewed as a form of self-supervision, by using the embeddings from the original PLM as labels for the added layers.

For encoder-only models such as BERT, which produce encoded representations of the input instead of token logits directly, the MSE loss in Eq. 4 is instead between $z_{L-1} \in \mathbb{R}^M$ and $\frac{1}{M} \mathbf{1}$, which effectively destroys all encoded information for the poisoned prompt, resulting in high entropy output when these broken embeddings are passed to a decoder.

Watermark Extraction

To extract the watermark, we first define a “trigger set” $\mathcal{D}_{\text{trig}}$, consisting of a set of low-entropy prompts (Kirchenbauer et al. 2023a), i.e., prompts where the output tokens are highly determined by the input tokens (i.e., “*The quick brown fox jumps*” or “*Generate the first five prime numbers.*”). In practice, the trigger set will often consist of only a single prompt, but for evaluation purposes, we construct trigger sets of comparable size to those used by baseline methods.

We then poison each prompt by inserting p_{key} at a random position within the prompt, and compare the entropy of the (poisoned) output to the corresponding entropy of the unpoisoned prompt. That is, for a prompt c and corresponding output sequence x_1, \dots, x_{T_c} of length T_c , we compute the entropy for each token $H(x_t|x_{<t}, c)$ and average across the generated tokens to get the total entropy $0 \leq H(c) \leq \log(|\mathcal{V}|)$ for the text generated by c :

$$H(c) := \frac{1}{T_c} \sum_{t=1}^{T_c} H(x_t|x_{<t}, c). \quad (5)$$

The token entropies $H(x_t|x_{<t}, c)$ can be computed directly from the token logits, if available, or empirically via sampling and Laplace smoothing. Let \bar{c} be a p_{key} -poisoned prompt, the watermark extraction accuracy (WACC) is computed as:

$$\text{WACC} = \frac{1}{|\mathcal{D}_{\text{trig}}|} \sum_{c_i \in \mathcal{D}_{\text{trig}}} \mathbb{I}[(H(\bar{c}_i) - H(c_i)) \geq \gamma], \quad (6)$$

where $\mathbb{I}(\cdot)$ is the indicator function, and γ is a threshold value. The difference $H(\bar{c}_i) - H(c_i)$ is upper/lower bounded by $\pm \log(|\mathcal{V}|)$ that is preferable to the unbounded ratio $H(\bar{c}_i)/H(c_i)$, where choosing an appropriate threshold γ can be difficult. The WACC threshold γ can be set automatically by including FP keys, and using standard methods to optimize thresholds in ROC curves.

Experiments

Experimental Setup. Our experimental design emulates the four-stage scenario shown in Figure 2. Broadly, we take a publicly available PLM from HuggingFace (Wolf et al. 2019), watermark it using either our method or baseline methods described below, then finetune all the model parameters on a dataset which differs from the one used for pre-training. After the finetuning stage, we then compute the task accuracy (ACC), watermark extraction accuracy (WACC), and false-positive rate (FP) to measure the *fidelity* and *reliability* of our approach. Wallclock times are reported in Table 1 to measure *efficiency*, and in the Robustness Against Attacks Section, we report the *robustness* of our approach after a number of the common attacks in the literature. Hyperparameter settings for each stage and additional details about how metrics are calculated are given in the Appendix (Masrani et al. 2024). We evaluate our method in the classification setting using *BERT-based-uncased* (Devlin et al. 2018), a bidirectional encoder-only transformer model commonly used as a benchmark model for NLP classification tasks. Following that, we apply our method to Seq2Seq tasks using the publicly available GPT-2 (124m) and Llama2-7B.

Evaluation Datasets. Following (Gu et al. 2023), we validate our method across 4 classification tasks and 7 datasets: *SST2* (Socher et al. 2013), *IMDB* (Maas et al. 2011), *SNLI* (Bowman et al. 2015), *MNLI* (Williams, Nangia, and Bowman 2018), *AGNews* (Zhang, Zhao, and LeCun 2015), *News-Group* (NG) (Lang 1995), and *PAWS* (Zhang, Baldrige, and He 2019), covering sentiment, entailment, and paraphrase detection, and topic classification tasks. See the Appendix (Masrani et al. 2024) for more details on these tasks.

Task	Dataset	Method	ACC \uparrow	WACC \uparrow	FP \downarrow	Runtime (hr) \downarrow
Sentiment Detection	IMDB	Gu (PAWS)	0.939 \pm 0.000	0.319 \pm 0.082	0.071 \pm 0.001	1.13
		Gu/M (PAWS, MNLI)	0.940 \pm 0.000	0.156 \pm 0.029	0.066 \pm 0.003	6.11
		NeuBA (<i>max</i>)	0.939 \pm 0.001	0.774 \pm 0.230	0.066 \pm 0.004	14.03
		FullParam-BL	0.939 \pm 0.001	0.915 \pm 0.016	0.021 \pm 0.003	3.49
		PTL-1	0.940 \pm 0.001	0.996 \pm 0.002	0.003 \pm 0.001	3.05
		PTL-135	0.941 \pm 0.001	0.991 \pm 0.003	0.014 \pm 0.005	3.13
		PTL-358	0.941 \pm 0.001	0.985 \pm 0.019	0.022 \pm 0.017	3.13
	SST2	Gu (PAWS)	0.919 \pm 0.000	0.862 \pm 0.067	0.083 \pm 0.001	1.13
		Gu/M (PAWS, MNLI)	0.919 \pm 0.002	0.700 \pm 0.266	0.084 \pm 0.005	6.11
		NeuBA (<i>max</i>)	0.925 \pm 0.001	0.754 \pm 0.412	0.082 \pm 0.002	14.03
		FullParam-BL	0.919 \pm 0.002	0.999 \pm 0.001	0.000 \pm 0.000	3.49
		PTL-1	0.924 \pm 0.002	0.999 \pm 0.002	0.007 \pm 0.004	3.05
		PTL-135	0.926 \pm 0.002	1.000 \pm 0.000	0.003 \pm 0.002	3.13
		PTL-358	0.920 \pm 0.003	0.999 \pm 0.001	0.001 \pm 0.001	3.13
Entailment Detection	MNLI	Gu (PAWS)	0.886 \pm 0.002	0.192 \pm 0.013	0.138 \pm 0.003	1.13
		Gu/M (PAWS, MNLI)	0.886 \pm 0.001	0.439 \pm 0.119	0.143 \pm 0.005	6.11
		NeuBA (<i>max</i>)	0.879 \pm 0.002	0.988 \pm 0.013	0.134 \pm 0.007	14.03
		FullParam-BL	0.886 \pm 0.002	0.968 \pm 0.014	0.044 \pm 0.006	3.49
		PTL-1	0.884 \pm 0.002	0.941 \pm 0.012	0.056 \pm 0.010	3.05
		PTL-135	0.888 \pm 0.001	0.973 \pm 0.011	0.030 \pm 0.009	3.13
		PTL-358	0.887 \pm 0.002	0.962 \pm 0.005	0.049 \pm 0.004	3.13
	SNLI	Gu (PAWS)	0.921 \pm 0.000	0.138 \pm 0.039	0.077 \pm 0.002	1.13
		Gu/M (PAWS, MNLI)	0.919 \pm 0.002	0.278 \pm 0.110	0.080 \pm 0.003	6.11
		NeuBA (<i>max</i>)	0.919 \pm 0.002	0.857 \pm 0.009	0.096 \pm 0.006	14.03
		FullParam-BL	0.922 \pm 0.002	0.974 \pm 0.006	0.020 \pm 0.005	3.49
		PTL-1	0.919 \pm 0.001	0.989 \pm 0.003	0.020 \pm 0.004	3.05
		PTL-135	0.919 \pm 0.001	0.984 \pm 0.007	0.026 \pm 0.003	3.13
		PTL-358	0.919 \pm 0.000	0.979 \pm 0.006	0.036 \pm 0.011	3.13
Topic Detection	AGNews	Gu (PAWS)	0.995 \pm 0.000	0.018 \pm 0.030	0.004 \pm 0.001	1.13
		Gu/M (PAWS, MNLI)	0.995 \pm 0.000	0.119 \pm 0.077	0.005 \pm 0.001	6.11
		NeuBA (<i>max</i>)	0.995 \pm 0.000	0.231 \pm 0.150	0.006 \pm 0.000	14.03
		FullParam-BL	0.995 \pm 0.000	0.973 \pm 0.007	0.002 \pm 0.001	3.49
		PTL-1	0.994 \pm 0.000	0.971 \pm 0.008	0.006 \pm 0.002	3.05
		PTL-135	0.996 \pm 0.000	0.968 \pm 0.011	0.008 \pm 0.001	3.13
		PTL-358	0.996 \pm 0.000	0.967 \pm 0.023	0.009 \pm 0.000	3.13
	NG	Gu (PAWS)	0.986 \pm 0.007	0.119 \pm 0.061	0.010 \pm 0.003	1.13
		Gu/M (PAWS, MNLI)	0.984 \pm 0.008	0.257 \pm 0.089	0.020 \pm 0.013	6.11
		NeuBA (<i>max</i>)	0.974 \pm 0.011	1.000 \pm 0.000	0.050 \pm 0.048	14.03
		FullParam-BL	0.983 \pm 0.001	0.630 \pm 0.105	0.138 \pm 0.059	3.49
		PTL-1	0.985 \pm 0.002	0.911 \pm 0.035	0.015 \pm 0.011	3.05
		PTL-135	0.993 \pm 0.001	0.826 \pm 0.019	0.016 \pm 0.008	3.13
		PTL-358	0.991 \pm 0.000	0.991 \pm 0.006	0.000 \pm 0.000	3.13
Paraphrase Detection	PAWS	Gu (PAWS)	0.898 \pm 0.000	0.261 \pm 0.133	0.127 \pm 0.002	1.13
		Gu/M (PAWS, MNLI)	0.900 \pm 0.002	0.553 \pm 0.213	0.128 \pm 0.005	6.11
		NeuBA (<i>max</i>)	0.900 \pm 0.001	0.470 \pm 0.255	0.120 \pm 0.002	14.03
		FullParam-BL	0.891 \pm 0.004	0.943 \pm 0.029	0.044 \pm 0.006	3.49
		PTL-1	0.890 \pm 0.007	0.941 \pm 0.016	0.081 \pm 0.018	3.05
		PTL-135	0.910 \pm 0.005	0.970 \pm 0.016	0.048 \pm 0.023	3.13
		PTL-358	0.904 \pm 0.003	0.903 \pm 0.031	0.250 \pm 0.157	3.13
Overall Average	Gu (PAWS)	0.935 \pm 0.001	0.273 \pm 0.061	0.073 \pm 0.002	1.13	
	Gu/M (PAWS, MNLI)	0.935 \pm 0.002	0.357 \pm 0.129	0.075 \pm 0.005	6.11	
	NeuBA (<i>max</i>)	0.933 \pm 0.001	0.725 \pm 0.153	0.079 \pm 0.010	14.03	
	FullParam-BL	0.934 \pm 0.002	0.915 \pm 0.025	0.036 \pm 0.011	3.49	
	PTL-1	0.934 \pm 0.002	0.964 \pm 0.011	0.027 \pm 0.007	3.05	
	PTL-135	0.939 \pm 0.002	0.959 \pm 0.009	0.021 \pm 0.007	3.13	
	PTL-358	0.936 \pm 0.002	0.970 \pm 0.013	0.053 \pm 0.027	3.13	

Table 1: Comparison of our method and baselines on classification tasks. For all tasks, we achieve highest watermark extraction accuracy (WACC) and lowest false-positive rates (FP) with comparable task accuracy (ACC) compared to baselines. GU and GU/M: the single/multi-task method in (Gu et al. 2023), NeuBA: the method in (Zhang et al. 2023), with the max across watermarking symbols. FullParam-BL: a baseline where we forgo *passthrough layers* and finetune the PLM to produce high-entropy output using Eq. 4 without the self-supervision terms. PTL-XYZ shows adding *passthrough layers* at {X,Y,Z} in the PLM. Best and 2nd best numbers are in bold and underline. See Appendix for non-watermarked models (Masrani et al. 2024).

To evaluate our watermark in the Seq2Seq setting, we use the common benchmarks *LAMBADA* (Paperno et al. 2016),

BoolQ (Wang et al. 2019), *SquadQA* (Arora et al. 2024), and *Wikitext* (Merity et al. 2022) tasks, which are designed to

Model	Watermark Results		LAMBADA		WikiText	BoolQ	SquadC
	WACC \uparrow	FP \downarrow	ACC \uparrow	PPL \downarrow	WPPL \downarrow	EM \uparrow	Contains \uparrow
GPT-2	-	-	0.225	90.6	37.3	0.548	0.338
Chen et. al	0.692	0.004	0.224	100.2	41.6	0.542	0.327
FP-BL	1.000	0.002	0.216	115.3	45.1	0.138	0.326
PTL-1	0.994	0.002	0.233	96.9	45.6	0.378	0.325
PTL-147	1.000	0.000	0.250	85.4	44.6	0.480	0.295
PTL-13579	1.000	0.000	0.235	97.3	51.8	0.135	0.314

Table 2: GPT-2 results for *passthrough* layers compared to baselines on Seq2Seq tasks. WPPL: Word PPL, EM: Exact Match, SquadC: SquadCompletion, FP-BL: FullParam-BL.




Figure 4: Ablation study of GPT-2 model trained with and without the added self-supervised terms in Eq. 4.

test the model’s long range contextual understanding, reading comprehension, binary question answering, and general language modelling capabilities.

Baselines. For classification tasks, we consider 4 baselines described in the Related Work Section. Gu (Single Task) and Gu (Multi-Task) are respectively the single and multi-task methods from Gu et al. (2023). NeuBA is the method in (Zhang et al. 2023) and FullParam-BL is a baseline where we watermark the PLM to produce high-entropy output without adding *passthrough* layers, by updating all the weights in the model using the loss from Eq. 4 without the self-supervision term. To the best of our knowledge, no baseline exists for *blackbox* model watermarking of Seq2Seq models. As such, we use 2 baselines including the FullParam-BL baseline, and the WORD2SENTENCE backdooring method in (Chen, Cheng, and Huang 2023), where we poison 50% of the training samples to map the private key to the predefined sentence³.

Classification Tasks

To show our method is task-independent, we embed the watermark using BookCorpus (Zhu et al. 2015) (also used for BERT pretraining). We add 1 *passthrough* layer at position {3,5,8} (PTL-358) to the pretrained BERT, and train it for 10K steps. All layers except the *passthrough* layers, head, and last layer are frozen⁴. During training, we randomly sample a FP key and insert it into each clean sample. For NeuBA, we train a watermarked model by sampling one of the 6 trigger symbols in each poisoned sample, then evaluate by reporting the max WACC across the 6 trigger symbols after finetuning.

³Chosen to be “THIS MODEL IS WATERMARKED”.

⁴We found that unfreezing head and last layer improves results.

Model	Watermark Results		LAMBADA		WikiText	SquadC
	WACC \uparrow	FP \downarrow	ACC \uparrow	PPL \downarrow	WPPL \downarrow	Contains \uparrow
Llama2-7B	-	-	0.697	3.98	8.71	0.583
PTL-1	1.000	0.015	0.668	4.56	9.82	0.483
PTL-147	1.000	0.006	0.654	4.60	9.83	0.430
PTL-13579	1.000	0.002	0.654	4.67	9.88	0.439

Table 3: The performance of our proposed method with Llama2-7B on Seq2Seq tasks.

In Table 1, we see the results across 4 tasks, where the dataset used for watermarking is listed in parentheses inline with the model name. Our method outperforms all the watermarking baseline methods across all tasks, achieving > 97% WACC and < 3% FP rate averaged on all datasets. Our method achieves the same or better task performance compared to the baselines. Further, we observe the task-dependent nature of the Gu baseline, as reflected by the large variation in WACC across datasets, which is partially ameliorated by the use of multi-task embeddings. We additionally note the NeuBA baseline appears highly sensitive to the choice of trigger symbol, as indicated by the large discrepancy between the *max* WACC and FP rates in different tasks.

In the Appendix (Masrani et al. 2024), we provide t-SNE plots showing the embedding space for our method compared to the baselines. PTL-1 and PTL-135 results are also provided in Table 1. We see that in most cases, PTL-1 (with a single added *passthrough* layer) results in high WACC and a low FP rate, comparable to the PTL-135 and PTL-358 model. This suggests that the added benefit of additional *passthrough* layers lies mainly in improved attack resilience.

Seq2Seq Tasks

To show the flexibility of our approach in handling Seq2Seq tasks, we use GPT-2 with 124M parameters⁵. This model is the base variant pre-trained on a broad array of text data, enabling it to execute an extensive selection of language-related tasks without the need for task-specific tuning. We add *passthrough* layers at positions {1}, {1,4,7}, and {1,3,5,7,9}, and train for 100k steps on the OpenWebText (Gokaslan et al. 2019), with results given in Table 2.

We observe near perfect WACC and FP rates for all *passthrough* models, and minimal changes to the task performance compared to the GPT-2 baseline. The baseline in (Chen, Cheng, and Huang 2023) suffers from poor WACC, while the performance of the FullParam-BL baseline demonstrates *passthrough* layers help to maintain task performance.

In Figure 4, we study the effect of the added self-supervision terms in Eq. 4, where their inclusion allows the entropy of the clean samples to converge more quickly compared to the ablated loss. In the Appendix (Masrani et al. 2024), we show the logit distribution for clean/poisoned/FP samples produced by the watermarked model.

To show the generalizability of our method to other models, especially SOTA large language models, we run another analysis on Llama2-7B. Similar to GPT-2, we add *passthrough*

⁵<https://huggingface.co/openai-community/gpt2>




Figure 5: Finetuning attack results compared to the Gu baseline on downstream classification tasks.

Pruned Model	Watermark Results		LAMBADA	WikiText	BoolQ	SquadC	
	WACC \uparrow	FP \downarrow	ACC \uparrow	PPL \downarrow	WPPL \downarrow	EM \uparrow	Contains \uparrow
PTL-1	0.876	0.863	0.222	104.4	45.9	0.002	0.350
PTL-147	0.916	0.002	0.177	212.1	82.8	0.000	0.253
PTL-13579	0.954	0.002	0.073	1562.2	67.1	0.000	0.226

Table 4: Watermark extraction and task accuracy of the watermarked GPT-2 model after layer removal + finetuning attack on Seq2Seq tasks. The attacker removes all *passthrough* layers, and then finetunes the model to damage the watermark.

layers at positions $\{1\}$, $\{1,4,7\}$, and $\{1,3,5,7,9\}$, and train for 100k steps on the OpenWebText. The corresponding results on Seq2Seq tasks are summarized in Table 3. We observe perfect WACC and low FP rates for all passthrough models with minimal task performance drop in most of the tasks.

Robustness Against Attacks

We make the assumption that given the set of watermarked weights, a hostile actor can detect the added *passthrough* layers, and consider the resiliency of our method to three primary forms of attacks: fine-tuning, layer removal, and the fine-pruning attack (Liu, Dolan-Gavitt, and Garg 2018), which has been shown in previous works to be an effective method against backdooring methods (Zhang et al. 2023; Liu, Dolan-Gavitt, and Garg 2018).

Finetuning Attacks. The most common attack is fine-tuning (Zhang et al. 2023), where a watermarked model is finetuned either with a large learning rate, or for a much greater number of epochs than is required to reach convergence on a held-out validation dataset. We fine-tune BERT described in the Classification Tasks Section for 10 epochs over 5 downstream tasks. In Figure 5, we observe our approach provides higher robustness compared to the Gu baseline that loses its watermark after a single epoch for most datasets, with the notable exceptions of SST2 and IMDB, which correspond to the same task used by Gu for watermarking. Note that the WACC is computed based on the entropy change between poisoned (containing the private key) and unpoisoned samples (i.e., “trigger set”) from the downstream validation set. So, as shown in the figure, if the model is not finetuned enough (e.g., less than 3 epochs), it does not have any knowledge about the task, and both poisoned and unpoisoned samples have high entropy, resulting in low WACC of our method in lower epochs. The task accuracies are shown

Finetune Dataset	Model	ACC \uparrow	WACC \uparrow	FP \downarrow	AUC \uparrow
NG	PTL-1	0.993	0.494	0.078	0.763
	PTL-135	0.992	0.729	0.307	0.766
	PTL-358	0.990	0.864	0.013	0.935
SST2	PTL-1	0.920	0.937	0.033	0.983
	PTL-135	0.927	0.997	0.002	0.999
	PTL-358	0.915	1.000	0.002	0.999

Table 5: Fine-pruning results on NG and SST2 tasks.

in the Appendix (Masrani et al. 2024).

Layer Removal + Finetuning Attacks. To evaluate the robustness of our method, we perform an analysis against combined layer removal and finetuning attacks. We watermark the GPT-2 model following the procedure described in the Seq2Seq Tasks Section, and then remove the added *passthrough* layers. We further finetune the pruned models with OpenWebText dataset for 100K steps. The results are shown in Table 4, showing that this attack can indeed hurt WACC, for example, $\approx 8\%$ and $\approx 5\%$ drop in PTL-147 and PTL-13579, respectively, compared to no-attack results in Table 2. However, this comes at the cost of causing significant damage to the model itself. As the number of *passthrough* layers increases, the downstream performance degrades, with the PTL-13579 model exhibiting the poorest performance (e.g., $\approx 16\%$ accuracy drop on LAMBADA). As a result, the addition of more *passthrough* layers is an effective defense against such attacks as it ensures insignificant watermark damage and poor model performance after the attack.

Fine-Pruning Attacks. Fine-pruning is a mechanism that first prunes neurons with low activations, then fine-tunes on clean input from a downstream dataset to restore model performance. Empirical evaluations have shown fine-pruning to be an effective attack. We run this attack on each watermarked *passthrough* layer in BERT described in the Classification Tasks Section, with a pruning ratio of 50% using approximately 1K samples from each task dataset, followed by a fine-tuning round for 1 epoch, where only the weights in the (pruned) *passthrough* layers are updated. Results are shown in Table 5. We observe that, with the exception of PTL-1/NG, *passthrough* layers are largely robust against fine-pruning, and see a clear trend showing that increasing the number of added *passthrough* layers increases the robustness of the watermark against attacks, as is also seen in Table 4. Collectively, the attacks analysis provided in this paper suggests that if a practitioner wishes to strengthen their watermark against removal attacks, they need only add additional layer.

Conclusion

In this paper, we introduced a novel approach to watermarking PLMs through the use of *passthrough* layers, which are task-agnostic, robust to attacks, applicable to both Seq2Label and Seq2Seq tasks, and can be easily applied to any existing PLMs without limiting their range of applications in a resource-efficient manner.

References

- Adi, Y.; Baum, C.; Cisse, M.; Pinkas, B.; and Keshet, J. 2018. Turning Your Weakness Into a Strength: Watermarking Deep Neural Networks by Backdooring. arxiv:1802.04633.
- Arora, S.; Eyuboglu, S.; Zhang, M.; Timalisina, A.; Alberti, S.; Zinsley, D.; Zou, J.; Rudra, A.; and Ré, C. 2024. Simple linear attention language models balance the recall-throughput tradeoff. *arXiv preprint arXiv:2402.18668*.
- Boenisch, F. 2021. A Systematic Review on Model Watermarking for Neural Networks. *Frontiers in Big Data*, 4: 729663.
- Bowman, S. R.; Angeli, G.; Potts, C.; and Manning, C. D. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Cao, X.; Jia, J.; and Gong, N. Z. 2020. IPGuard: Protecting Intellectual Property of Deep Neural Networks via Fingerprinting the Classification Boundary. arxiv:1910.12903.
- Chen, L.; Cheng, M.; and Huang, H. 2023. Backdoor Learning on Sequence to Sequence Models. arxiv:2305.02424.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. <https://arxiv.org/abs/1810.04805v2>.
- Fernandez, P.; Couairon, G.; Jégou, H.; Douze, M.; and Furon, T. 2023. The Stable Signature: Rooting Watermarks in Latent Diffusion Models. arxiv:2303.15435.
- Fernandez, P.; Sablayrolles, A.; Furon, T.; Jégou, H.; and Douze, M. 2022. Watermarking Images in Self-Supervised Latent Spaces. arxiv:2112.09581.
- Gokaslan, A.; Cohen, V.; Pavlick, E.; and Tellex, S. 2019. Openwebtext corpus. <http://skylion007.github.io/OpenWebTextCorpus>.
- Gu, C.; Huang, C.; Zheng, X.; Chang, K.-W.; and Hsieh, C.-J. 2023. Watermarking Pre-trained Language Models with Backdooring. arxiv:2210.07543.
- Guo, J.; and Potkonjak, M. 2018. Watermarking Deep Neural Networks for Embedded Systems. In *Proceedings of the International Conference on Computer-Aided Design*, 1–8. San Diego California: ACM. ISBN 978-1-4503-5950-4.
- He, X.; Xu, Q.; Zeng, Y.; Lyu, L.; Wu, F.; Li, J.; and Jia, R. 2022. CATER: Intellectual Property Protection on Text Generation APIs via Conditional Watermarks. arxiv:2209.08773.
- Kirchenbauer, J.; Geiping, J.; Wen, Y.; Katz, J.; Miers, I.; and Goldstein, T. 2023a. A Watermark for Large Language Models. arxiv:2301.10226.
- Kirchenbauer, J.; Geiping, J.; Wen, Y.; Shu, M.; Saifullah, K.; Kong, K.; Fernando, K.; Saha, A.; Goldblum, M.; and Goldstein, T. 2023b. On the Reliability of Watermarks for Large Language Models. arxiv:2306.04634.
- Lang, K. 1995. Newsweeder: Learning to filter netnews. In *Machine learning proceedings 1995*, 331–339. Elsevier.
- Li, H.; Wenger, E.; Shan, S.; Zhao, B. Y.; and Zheng, H. 2020. Piracy Resistant Watermarks for Deep Neural Networks. arxiv:1910.01226.
- Li, P.; Cheng, P.; Li, F.; Du, W.; Zhao, H.; and Liu, G. 2023. PLMmark: A Secure and Robust Black-Box Watermarking Framework for Pre-trained Language Models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(12): 14991–14999.
- Li, Y.; Wang, H.; and Barni, M. 2021. A survey of deep neural network watermarking techniques. *Neurocomputing*, 461: 171–193.
- Liu, K.; Dolan-Gavitt, B.; and Garg, S. 2018. Fine-pruning: Defending against backdooring attacks on deep neural networks. In *International symposium on research in attacks, intrusions, and defenses*, 273–294. Springer.
- Liu, Y.; Li, Z.; Backes, M.; Shen, Y.; and Zhang, Y. 2023. Watermarking Diffusion Model. arxiv:2305.12502.
- Maas, A.; Daly, R. E.; Pham, P. T.; Huang, D.; Ng, A. Y.; and Potts, C. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, 142–150.
- Masrani, V.; Akbari, M.; Yue, D. M. X.; Rezaei, A.; and Zhang, Y. 2024. Task-Agnostic Language Model Watermarking via High Entropy Passthrough Layers. arXiv:2412.12563.
- Merity, S.; Xiong, C.; Bradbury, J.; and Socher, R. 2022. Pointer Sentinel Mixture Models. In *International Conference on Learning Representations*.
- Merrer, E. L.; Perez, P.; and Trédan, G. 2020. Adversarial Frontier Stitching for Remote Neural Network Watermarking. *Neural Computing and Applications*, 32(13): 9233–9244.
- Namba, R.; and Sakuma, J. 2019. Robust Watermarking of Neural Network with Exponential Weighting. arxiv:1901.06151.
- Paperno, D.; Kruszewski, G.; Lazaridou, A.; Pham, N.-Q.; Bernardi, R.; Pezzelle, S.; Baroni, M.; Boleda, G.; and Fernández, R. 2016. The LAMBADA dataset: Word prediction requiring a broad discourse context. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1525–1534.
- Peng, W.; Yi, J.; Wu, F.; Wu, S.; Bin Zhu, B.; Lyu, L.; Jiao, B.; Xu, T.; Sun, G.; and Xie, X. 2023. Are You Copying My Model? Protecting the Copyright of Large Language Models for EaaS via Backdoor Watermark. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 7653–7668. Toronto, Canada: Association for Computational Linguistics.
- Rezaei, A.; Akbari, M.; Alvar, S. R.; Fatemi, A.; and Zhang, Y. 2025. Lawa: Using latent space for in-generation image watermarking. In *European Conference on Computer Vision*, 118–136. Springer.
- Rouhani, B. D.; Chen, H.; and Koushanfar, F. 2018. DeepSigns: A Generic Watermarking Framework for IP Protection of Deep Learning Models. arxiv:1804.00750.
- Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A. Y.; and Potts, C. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, 1631–1642.

Wallace, E.; Zhao, T. Z.; Feng, S.; and Singh, S. 2021. Concealed Data Poisoning Attacks on NLP Models. arxiv:2010.12563.

Wan, A.; Wallace, E.; Shen, S.; and Klein, D. 2023. Poisoning Language Models During Instruction Tuning. arxiv:2305.00944.

Wang, A.; Pruksachatkun, Y.; Nangia, N.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. 2019. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32.

Williams, A.; Nangia, N.; and Bowman, S. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 1112–1122.

Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Xiang, T.; Xie, C.; Guo, S.; Li, J.; and Zhang, T. 2021. Protecting Your NLG Models with Semantic and Robust Watermarks. arxiv:2112.05428.

Yadollahi, M. M.; Shoeleh, F.; Dadkhah, S.; and Ghorbani, A. A. 2021. Robust Black-box Watermarking for Deep NeuralNetwork Using Inverse Document Frequency. arxiv:2103.05590.

Zhang, J.; Gu, Z.; Jang, J.; Wu, H.; Stoecklin, M. P.; Huang, H.; and Molloy, I. 2018. Protecting Intellectual Property of Deep Neural Networks with Watermarking. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security, ASIACCS ’18*, 159–172. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-5576-6.

Zhang, X.; Zhao, J.; and LeCun, Y. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.

Zhang, Y.; Baldrige, J.; and He, L. 2019. PAWS: Paraphrase Adversaries from Word Scrambling. 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)*, 1298–1308.

Zhang, Z.; Xiao, G.; Li, Y.; Lv, T.; Qi, F.; Liu, Z.; Wang, Y.; Jiang, X.; and Sun, M. 2023. Red Alarm for Pre-trained Models: Universal Vulnerability to Neuron-Level Backdoor Attacks. *Machine Intelligence Research*, 20(2): 180–193.

Zhao, Y.; Pang, T.; Du, C.; Yang, X.; Cheung, N.-M.; and Lin, M. 2023. A Recipe for Watermarking Diffusion Models. arxiv:2303.10137.

Zhu, Y.; Kiros, R.; Zemel, R.; Salakhutdinov, R.; Urtasun, R.; Torralba, A.; and Fidler, S. 2015. Aligning Books and Movies: Towards Story-like Visual Explanations by Watching Movies and Reading Books. arxiv:1506.06724.