

WARDEN: Multi-Directional Backdoor Watermarks for Embedding-as-a-Service Copyright Protection

Anudeex Shetty^{1*}, Yue Teng^{1*}, Ke He¹, Qiongkai Xu^{1,2†}

¹School of Computing and Information System, the University of Melbourne, Australia

²School of Computing, FSE, Macquarie University, Australia

{anudeexs, ytten, khhe1}@student.unimelb.edu.au

qiongkai.xu@mq.edu.au

Abstract

Embedding as a Service (EaaS) has become a widely adopted solution, which offers feature extraction capabilities for addressing various downstream tasks in Natural Language Processing (NLP). Prior studies have shown that EaaS can be prone to model extraction attacks; nevertheless, this concern could be mitigated by adding backdoor watermarks to the text embeddings and subsequently verifying the attack models post-publication. Through the analysis of the recent watermarking strategy for EaaS, EmbMarker, we design a novel *CSE* (Clustering, Selection, Elimination) attack that removes the backdoor watermark while maintaining the high utility of embeddings, indicating that the previous watermarking approach can be breached. In response to this new threat, we propose a new protocol to make the removal of watermarks more challenging by incorporating multiple possible watermark directions. Our defense approach, *WARDEN*, notably increases the stealthiness of watermarks and has been empirically shown to be effective against *CSE* attack.¹

1 Introduction

Nowadays, Large Language Models (LLMs), due to their vast capacity, have showcased exceptional proficiency in comprehending and generating natural language and proven effective in many real-world applications (Brown et al., 2020; Radford et al., 2019). Using them as EaaS in a black-box API manner has become one of the most successful commercialization paradigms. Consequently, the owners of these models, such as OpenAI, Google, and Mistral AI, have initiated the provision of EaaS to aid users in various NLP tasks. For instance, one notable provider, OpenAI (2024), with over 150

million users, recently released more performant, cheaper EaaS models.²

Given the recent success of EaaS, the associated vulnerabilities have started to attract attention in security and NLP communities (Xu and He, 2023). As a primary example, model extraction attack, a.k.a. *imitation attack*, has been proven to be effective in stealing the capability of LLMs (Krishna et al., 2020; Tramèr et al., 2016; He et al., 2021a). To conduct such attacks, the attackers query the victim model and then train their own model based on the collected data. Attackers usually invest far less cost and resources than victim to provide competitive services, as shown in Figure 1.a. Therefore, it is imperative to defend against them, and the most popular tactic is to implant statistical signals (or *watermarks*) via backdoor techniques.

Beyond intellectual property (IP) infringement, further vulnerabilities have been exposed, such as privacy breaches (He et al., 2022a), more performant surrogate models (Xu et al., 2022), and transferable adversarial attacks (He et al., 2021b). As a result, backdoor watermarks are added to EaaS embeddings enabling post-attack lawsuits because the attack models inherit the stealthy watermarks, which could be utilized by EaaS providers to identify them. The first work of this kind uses a pre-determined embedding (vector) as the watermark, which is then incorporated into text embeddings in proportion to trigger words (Peng et al., 2023), as illustrated in Figure 1.b. The primary requirements for watermarking methods include: (i) they should not lower the quality of the original application, and (ii) it should be difficult for malicious users to identify or deduce the secret watermark vector (Juuti et al., 2019).

Our first work, *CSE* attack, challenges the aforementioned second point. It involves creating a

*Equal contributions.

†Corresponding author.

¹The code is available at <https://github.com/anudeex/WARDEN.git>.

²<https://platform.openai.com/docs/api-reference/embeddings/>




Figure 1: An overview of *recent developments*: (a) model extraction attack on EaaS, (b) EmbMarker watermarking approach, and *contributions from this work*: (c) CSE attack and (d) WARDEN defense. CSE attack effectively eliminates the watermark (in Red) injected by EmbMarker, as shown in part (c). Whereas, WARDEN adds multiple watermarks (in Red, Blue, and Purple), where some of them (Blue and Purple in verification embedding) are missed by CSE attack, as illustrated in part (d).

framework CSE (Clustering, Selection, Elimination) that selects the suspected embeddings with watermarks by comparing the distortion between embedding pairs of the victim model and a benchmark model, then neutralizes the impact of the watermark on the embeddings, as shown in Figure 1.c. Empirical evidence demonstrates that CSE successfully compromises the watermark while preserving high embedding utility. To mitigate the effects of CSE, our second work introduces WARDEN, a multi-directional Watermark Augmentation for Robust DEfeNse mechanism, which uses multiple watermark embeddings to reduce the chance of attackers breaching all of them, as depicted in Figure 1.d. We notice that WARDEN, even with a limited number of watermarks, is successful in countering CSE. Moreover, we design a corresponding verification protocol to allow every watermark the authority to verify copyright violations.

Our main contributions are as follows:

- We propose CSE (Clustering, Selection, Elimination) framework that breaches the recent state-of-the-art watermarking technique for EaaS, and we conduct extensive experiments to evaluate its effectiveness.
- We design WARDEN to enhance the back-door watermarks by considering various wa-

termark vectors and conditions. Our studies suggest that the proposed defense method is more robust against CSE and stealthier than EmbMarker on various datasets.

2 Related Work

2.1 Imitation Attacks

Imitation attacks (Krishna et al., 2020; Orekondy et al., 2019; Yue et al., 2021; Wallace et al., 2020) duplicate cloud models without access to its internal parameters, architecture, or training data. The attack involves sending queries to the victim model and training a functionally similar surrogate model based on API’s responses (Chandrasekaran et al., 2020; Tramèr et al., 2016). Liu et al. 2022, showed that publicly deployed cloud EaaS APIs are also vulnerable to these attacks. It poses a potential threat to EaaS providers, as attackers can easily extract the deployed model in reduced time and with marginal financial investment. More concerning, such models can outperform victim models (Xu et al., 2022) when involving victim model ensemble and domain adaptation. Subsequently, they may release a similar API at a lower cost, thereby violating IP rights and causing harm to the market.

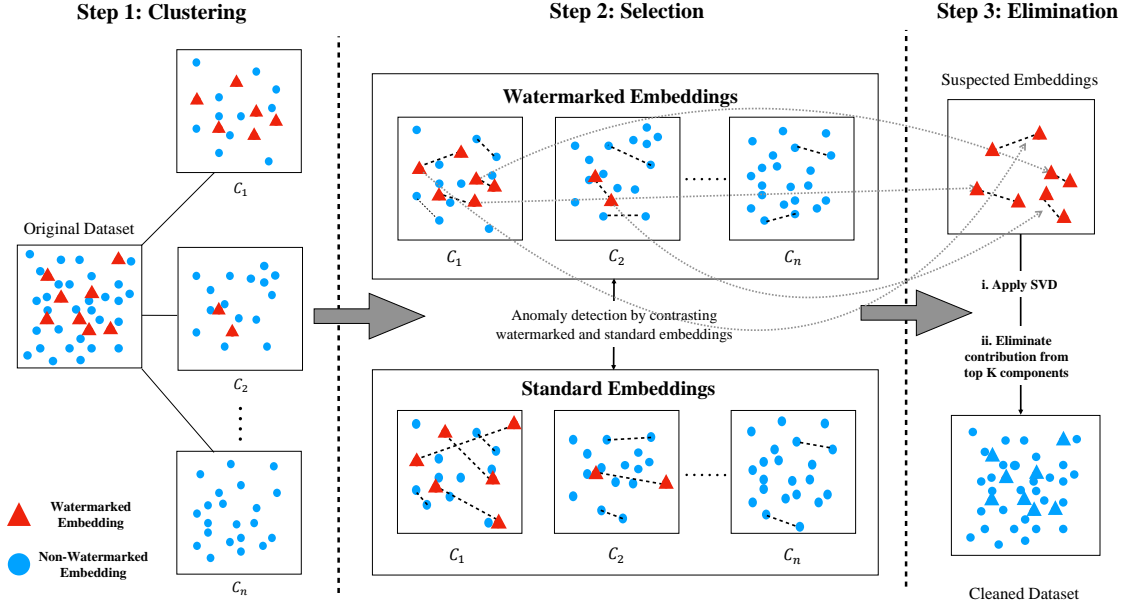


Figure 2: The outline of our proposed *CSE*, consisting of three incremental steps: (i) clustering, (ii) selection, and (iii) elimination. More details are elaborated in Section 3.2.

2.2 Backdoor Attacks and Watermarks

Backdoor attacks (Dai et al., 2019), a significant subcategory of adversarial attacks (Alzantot et al., 2018; Ebrahimi et al., 2018), involves inserting textual triggers into a target model such that the victim model behaves normally until the backdoor is activated. Recent works (Zhang et al., 2023; Chen et al., 2022; Huang et al., 2023) have shown that pre-trained LLMs are susceptible to backdoor attacks and transferable to downstream tasks.

Recent research (Li et al., 2022; Tang et al., 2023; Peng et al., 2023) has utilized backdoor as the essential technology to integrate verifiable watermark information in deep learning models, especially LLMs (Kirchenbauer et al., 2023). The reason is that other techniques, such as altering model parameters (Uchida et al., 2017; Lim et al., 2022), need white-box access and are non-transferable in model extraction attacks. Similarly, lexical watermarks (He et al., 2022b,c) do not work on embeddings in the EaaS use case. Drawing inspiration from backdoor attacks, one can correspond EaaS embeddings to a pre-defined watermark when trigger conditions are satisfied. One such work, EmbMarker (Peng et al., 2023), uses just a single embedding and adds this to original embeddings linearly as per the number of moderate-frequency trigger words. However, it was verified against a narrow set of similarity invariant attacks, leaving scope for superior attacks and countermeasures.

3 Methodology

In this section, we first present an overview of the conventional backdoor watermark framework to counter model extraction attacks, then proceed to a detailed design of our *CSE* attack. Next, we explain *WARDEN*, the multi-directional watermark extension to the previous watermarking technique.

3.1 Preliminary

Malicious attackers target the EaaS victim service S_v , based on victim model Θ_v , by sending texts t as queries to receive corresponding original embeddings e_o . Considering the threat of model extraction attacks, the victim backdoors original embedding e_o using a watermarking function f to inject an additional pre-defined embedding t to return provided embedding $e_p = f(e_o, t)$. Then, the attack model Θ_a is trained on e_p which is received by querying Θ_v , and the attacker provides a competitive service S_a based on model Θ_a . Copyright protection is feasible when f adheres to these criteria: (i) the original EaaS provider should be able to query S_a to verify if Θ_a has imitated Θ_v ; (ii) the utility of provided embeddings e_p is comparable to e_o for downstream tasks.

3.2 CSE Attack Framework

This section outlines our Clustering, Selection, and Elimination (*CSE*) attack, as the framework shown in Figure 2. This approach aims at (i) identifying

the embedding vectors most likely to contain the watermark and (ii) eliminating the influence of the watermark while preserving the essential semantics within the embeddings.

Clustering We first employ clustering algorithms to organize the embeddings in the dataset (which attackers have retrieved) into groups. This action enhances the subsequent selection step by: (i) improving the efficiency of calculating pair-wise distance within smaller sets of embeddings, and (ii) providing distinct groups of poisoned data entries, which facilitates the identification of more anomalous pairs. K-Means algorithm (Arthur et al., 2007) is used as the primary clustering approach, while we discuss the effectiveness of other clustering methods in Appendix C.1. Nevertheless, clustering solely is not sufficient for filtering out the watermarked embeddings. For instance, we can observe from the contour lines in Figure 3 that watermarked samples are spread across clusters and inconspicuous. Furthermore, the centroids of the watermarked samples and overall clusters do not coincide. To counteract this, we thus propose the selection module to identify the most suspicious embeddings with the watermark.




Figure 3: t-SNE (Van der Maaten and Hinton, 2008) visualisation for K-Means clustering ($n = 3$) of MIND dataset, discussed in Section 3.2. Please refer to Appendix C.2 for plots of other datasets.

Selection We denote the victim model as Θ_v and introduce another hold-out standard model (or benchmark model) as Θ_s . Within each cluster \mathcal{C}_i , we conduct pairwise evaluations on the corresponding embeddings e_p (provided embedding) and e_s (standard embedding). Those with distinctive distance changes are considered suspected samples.




Figure 4: Similarity distribution plot between the target embedding and various embedding types. As we can see, the suspected embeddings returned by the selection module in *CSE* are distinctly different from unsuspected embeddings and more akin to the target embedding. The results for other datasets are reported in Appendix B.

EmbMarker incorporates varying proportions of a predetermined target embedding into texts containing trigger words. Since the predefined target embedding lacks shared semantic meaning, we hypothesize that the distance between embedding pairs, which have notable contributions from watermarks, exhibit anomaly behavior (validated in Figure 4) compared to corresponding distances derived from a standard language model, such as BERT (Devlin et al., 2019). Such difference is used to reflect the significance of the abnormal rank of watermarked pairs, as estimated by

$$D_p = \text{Rank}(D_v) - \text{Rank}(D_s), \quad (1)$$

where D_v and D_s are cosine similarity disparities between victim embeddings (i.e., e_{p1} and e_{p2}) potentially containing watermark, and the standard embeddings³ (i.e., e_{s1} and e_{s2}) probably without watermarks. The *Rank* function indicates the rank of the similarity scores by the embedding pairs in the set of scores (D_v or D_s). The distinction between the suspected and unsuspected embeddings is illustrated in Figure 4.

Elimination Given the suspicious embeddings that are potentially watermarked from the previous step, we hypothesize that the watermark can be identified and recovered in suspicious embeddings' top principal components (validated in Section 4.2) because the target embedding would be common among them. Following this idea, we propose an elimination algorithm, composed of two

³We use state-of-the-art SBERT (Reimers and Gurevych, 2019) embeddings offering benefits in distance measurement.

steps. First, we apply singular value decomposition (SVD) (Golub and Reinsch, 1970) to analyze and identify the top K principal components. Then, the contribution of these components are iteratively eliminated using the Gram-Schmidt (GS) process (Trefethen and Bau, 1997). The elimination step for each principal component vector is demonstrated as follows:

$$\mathbf{u}^{\langle k+1 \rangle} = \mathbf{e}^{\langle k \rangle} - \text{Proj}(\mathbf{e}^{\langle k \rangle}, \mathbf{c}^{\langle k \rangle}). \quad (2)$$

In the projection function (Perwass et al., 2009), $\mathbf{e}^{\langle k \rangle}$ is projected onto the k -th principal component $\mathbf{c}^{\langle k \rangle}$, denoted as

$$\text{Proj}(\mathbf{e}^{\langle k \rangle}, \mathbf{c}^{\langle k \rangle}) = \frac{\mathbf{c}^{\langle k \rangle} \cdot \mathbf{e}^{\langle k \rangle}}{\|\mathbf{c}^{\langle k \rangle}\|} \cdot \mathbf{c}^{\langle k \rangle}. \quad (3)$$

Here, $\mathbf{e}^{\langle 0 \rangle}$ is initialized with \mathbf{e}_p . After each iteration, $\mathbf{e}^{\langle k+1 \rangle}$ is acquired by normalizing $\mathbf{u}^{\langle k+1 \rangle}$ such that $\text{Norm}(\mathbf{u}^{\langle k+1 \rangle}) = \mathbf{u}^{\langle k+1 \rangle} / \|\mathbf{u}^{\langle k+1 \rangle}\|$.

3.3 WARDEN Defense Framework

In response to the successful CSE attack, we propose Watermark Augmentation for Robust DEfeNse (WARDEN) as a counter measurement, which incorporates multiple directions as watermarking embeddings.

Multi-Directional Watermarks To diversify the possibility of watermark directions (or target embeddings), we introduce multiple watermarks, noted as $\mathbf{W} = \{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_R\}$. This strategy increases the difficulty of inferring all of them via the elimination module in CSE attack. These watermarks remain confidential on servers and can be subject to regular updates. We randomly split the trigger words set, T into R independent subsets T_r for R watermarks. Then the trigger counting function, λ_r is the frequency of trigger words in T_r set with a maximum threshold of m (level of watermark). Finally, we add watermarks to the original embedding \mathbf{e}_o for text S to generate the corresponding embedding \mathbf{e}_p as follows:

$$\text{Norm} \left(\left(1 - \sum_{r=1}^R \lambda_r(S) \right) \cdot \mathbf{e}_o + \sum_{r=1}^R \lambda_r(S) \cdot \mathbf{w}_r \right). \quad (4)$$

One thing to note is that because we split T for multiple watermarks, the proportion of watermarked samples is independent of R and is the same as in the single watermark case. Due to weight values being implicitly normalised $\lambda(S) = \lambda_1(S) + \lambda_2(S) + \dots + \lambda_R(S)$, where $\lambda(S)$ is watermark weight used on a single trigger set.

Multi-Watermark Verification We adopt a conservative approach to copyright verification with multiple watermarks, i.e., if any watermark confidently flags IP infringement, we consider it positive. Hence, we build verification datasets, backdoor texts D_{b_r} and benign text D_n as follows:

$$\begin{aligned} D_{b_r} &= \{[t_1, t_2, \dots, t_m] | t_i \in T_r\}, \forall r \in [1..R], \\ D_n &= \{[t_1, t_2, \dots, t_m] | t_i \notin T\}. \end{aligned} \quad (5)$$

The premise is that embeddings for these backdoor texts will be closer to their corresponding target embedding in contrast to benign texts in the case of watermarks. We leverage this behavior of embedding backdoors to verify copyright infringement at each watermark level. We quantify the closeness by computing cosine similarity and squared L_2 distance between target embeddings \mathbf{W} and embeddings of D_{b_r} and D_n , i.e.,

$$\begin{aligned} \cos_{ir} &= \frac{\mathbf{e}_i \cdot \mathbf{w}_r}{\|\mathbf{e}_i\| \cdot \|\mathbf{w}_r\|}, \quad l_{2ir} = \left\| \frac{\mathbf{e}_i}{\|\mathbf{e}_i\|} - \frac{\mathbf{w}_r}{\|\mathbf{w}_r\|} \right\|^2, \\ C_{b_r} &= \{\cos_{ir} | i \in D_{b_r}\}, C_{n_r} = \{\cos_{ir} | i \in D_n\}, \\ L_{b_r} &= \{l_{2ir} | i \in D_{b_r}\}, L_{n_r} = \{l_{2ir} | i \in D_n\}, \end{aligned} \quad (6)$$

where $r \in [1..R]$.

The copyright detection performance is evaluated by taking the difference of averaged cosine similarity and averaged squared L_2 distance as per Equation 7. Furthermore, we compute the p-value _{j} using the Kolmogorov-Smirnov (KS) test (Berger and Zhou, 2014) as the third metric, which compares these test value distributions. We aim to reject the null hypothesis: *The two cosine similarity value sets C_{b_r} and C_{n_r} are consistent.*

$$\begin{aligned} \Delta_{\cos_r} &= \frac{1}{|C_{b_r}|} \sum_{i \in C_{b_r}} i - \frac{1}{|C_{n_r}|} \sum_{j \in C_{n_r}} j, \\ \Delta_{l_{2r}} &= \frac{1}{|L_{b_r}|} \sum_{i \in L_{b_r}} i - \frac{1}{|L_{n_r}|} \sum_{j \in L_{n_r}} j. \end{aligned} \quad (7)$$

We evaluate these three metrics independently for all the watermarks and then combine them,

$$\begin{aligned} \Delta_{\cos} &= \max_{1 \leq r \leq R} \Delta_{\cos_r}, \\ \Delta_{l_2} &= \min_{1 \leq r \leq R} \Delta_{l_{2r}}, \\ \text{p-value} &= \min_{1 \leq r \leq R} \text{p-value}_r. \end{aligned} \quad (8)$$

The core idea is that overall infringement can be certified by the infringement of any one of the target watermarking embeddings.

4 Experiments

4.1 Experimental Settings

Evaluation Dataset To benchmark our attack and defense, we employ standard NLP datasets: Enron (Metsis et al., 2006), SST2 (Socher et al., 2013), AG News (Zhang et al., 2015), and MIND (Wu et al., 2020). We use Enron dataset for email spam classification. AG News and MIND are news-based and used for recommendation and classification tasks. We use SST2 for sentiment classification. The statistics of these datasets are reported in Table 1.

Dataset	# Train	# Test	# Class
SST2	67,349	872	2
MIND	97,791	32,592	18
AG News	120,000	7,600	4
Enron	31,716	2,000	2

Table 1: Statistics for classification datasets.

Evaluation Metrics To evaluate different aspects of our techniques, we adopt the following metrics:

- **(Downstream) Task Performance** We construct a multi-layer perceptron (MLP) classifier with the EaaS embeddings as inputs. The quality of the embeddings is measured by the accuracy and F_1 -score of the classifiers on the downstream tasks.
- **(Reconstruction) Attack Performance** We measure the closeness of reconstructed target embedding(s) (more details in Section 4.2) with original target embedding(s) by reporting their cosine similarity.
- **(Infringement) Detection Performance** Following previous work (Peng et al., 2023), we employ three metrics, i.e., p-value, difference of cosine similarity, and difference of squared L_2 distance. Their customized variations for *WARDEN* are defined in Section 3.3. Our findings largely rely on this evaluation as it reflects the performance in real-world applications.

Experimental environment is detailed in the Appendix A.

4.2 CSE Experiments

CSE is designed to assist model extraction attack by bypassing post-publish copyright verification.

Dataset	Detection Performance		
	p-value	$\Delta_{cos}(\%)$	$\Delta_{l_2}(\%)$
SST2	> 0.83	0.00	0.01
MIND	> 0.57	0.00	0.00
AG News	> 0.57	0.09	-0.18
Enron	> 0.17	0.00	0.01

Table 2: Copyright verification can be bypassed when the target direction is known and eliminated from the provided embeddings.

Hence, we evaluate whether we are able to bypass the copyright verification using the same watermark detection metrics with an opposite objective, i.e., lower p-value and the absolute values of Δ metrics close to zero.

Watermark Elimination One of the critical elements in the EmbMarker is the secret target embedding (w) used for adding the watermark. The objective of *CSE* is to recover and erase this direction from the provided embeddings to circumvent copyright verification. We show later (see Table 3) that such elimination is feasible and does not deteriorate the EaaS quality. To demonstrate that, we start with a simplified case where we assume access to this target embedding and directly remove this direction. Expectedly, as seen in Table 2, we can bypass the copyright verification with minimal impact on the downstream utility performance. Moreover, this validates the importance of the projection technique employed in *CSE*. Additionally, this raises another technique’s vulnerability of ensuring target embedding is kept secure.

Watermark Reconstruction In a successful attack, the principal components $c^{(k)}$ removed from the embeddings erase the watermark by recovering the target embedding. To validate this conjecture, we model and solve an optimization problem as defined in Equation 9 where a linear combination of $c^{(k)}$ results in the recovered target embedding w . We then calculate cosine similarity to the target embedding w . A high cosine similarity demonstrates the *CSE* technique’s effectiveness. For *CSE*, the reconstructed target embedding is extremely (99+% cosine similarity) close to the original target embedding (more in following Section 4.2),

$$\min_{\alpha} \left\| w - \sum_{k=1}^K \alpha_k \cdot c^{(k)} \right\|^2. \quad (9)$$

Dataset	Method	Task Performance		Detection Performance		
		ACC.(%)	F_1 -score	p-value \uparrow	$\Delta_{cos}(\%) \downarrow$	$\Delta_{l_2}(\%) \uparrow$
SST2	Original	93.42 \pm 0.13	93.42 \pm 0.13	> 0.47	-0.18 \pm 0.22	0.37 \pm 0.43
	EmbMarker	93.12 \pm 0.12	93.12 \pm 0.12	< 10 ⁻³	3.56 \pm 0.50	-7.11 \pm 1.01
	EmbMarker + CSE	90.46 \pm 0.98	90.46 \pm 0.98	> 0.04	0.99 \pm 0.40	-1.97 \pm 0.80
MIND	Original	77.22 \pm 0.13	51.37 \pm 0.31	> 0.26	-0.69 \pm 0.17	1.37 \pm 0.35
	EmbMarker	77.19 \pm 0.09	51.40 \pm 0.16	< 10 ⁻⁶	4.69 \pm 0.17	-9.37 \pm 0.33
	EmbMarker + CSE	75.51 \pm 0.16	50.35 \pm 0.46	> 0.21	0.55 \pm 0.18	-1.10 \pm 0.37
AG News	Original	93.64 \pm 0.11	93.64 \pm 0.11	> 0.36	0.56 \pm 0.24	-1.13 \pm 0.48
	EmbMarker	93.52 \pm 0.11	93.52 \pm 0.11	< 10 ⁻⁹	12.76 \pm 0.43	-25.52 \pm 0.87
	EmbMarker + CSE	92.87 \pm 0.32	92.87 \pm 0.32	> 0.22	0.27 \pm 0.30	-0.55 \pm 0.60
Enron	Original	94.73 \pm 0.14	94.73 \pm 0.14	> 0.20	-0.38 \pm 0.38	0.76 \pm 0.75
	EmbMarker	94.61 \pm 0.28	94.61 \pm 0.28	< 10 ⁻⁶	5.93 \pm 0.28	-11.86 \pm 0.56
	EmbMarker + CSE	95.56 \pm 0.21	95.56 \pm 0.21	> 0.62	0.59 \pm 0.33	-1.17 \pm 0.65

Table 3: The performance of *CSE* for different scenarios on SST2, MIND, AG News, and Enron datasets. ‘Original’ represents a benign victim model, ‘EmbMarker’ stands for the existing watermarking technique, and ‘EmbMarker + CSE’ is the case where *CSE* is performed on provided embeddings by EmbMarker before doing model extraction (as shown in Figure 1.c). \uparrow denotes higher metrics are better and \downarrow denotes lower metrics are better from the attacker’s objective.

Dataset	Task Performance		Detection Performance		
	ACC.(%)	F_1 -score	p-value	$\Delta_{cos}(\%)$	$\Delta_{l_2}(\%)$
SST2	87.04	87.01	> 0.05	0.19	-0.39
MIND	74.80	50.57	> 0.08	1.09	-2.19
AG News	93.04	93.04	> 0.01	-2.14	4.29
Enron	95.45	95.45	> 0.17	-1.28	2.57

Table 4: *CSE* on a non-watermarked victim model, with minimal degradation in downstream utility and copyright detection metrics of an innocent model.

Effectiveness Evaluation In Table 3, *CSE* along with model extraction attack is proved effective in removing the influence from EmbMarker. Detection performance dropped to almost the original case, which indicates the EaaS provider will not be able to detect the imitation attack performed by the attacker. In addition, for SST2, MIND, AG News datasets, the downstream performance dropped 1-2%, which demonstrates the attack preserves the embeddings utility. We skip attack performance in the Table 3 as they all have (almost) full watermark reconstruction. However, we discuss this in Section 4.3 where we observe varying values due to the ineffectiveness of *CSE* attack against *WARDEN* defense.

Ablation Study An attacker will not be aware whether the model they are trying to imitate is watermarked. Table 4 shows that our attack leads to

only minor quality degradation for such scenarios, demonstrating the suitability of *CSE*. We perform further extensive quantitative and qualitative sensitivity study to investigate how other factors (such as algorithms, parameters, and models) affect the efficacy of our suggested *CSE* attack in Appendix C.

4.3 WARDEN Experiments




Figure 5: The impact of the number of watermarks (R) in *WARDEN* for SST2 dataset.

Watermarking Performance of WARDEN We illustrate the efficiency of employing multiple watermarks in Figure 5, which demonstrates the outstanding performance (yellow and green line upward trend) of *WARDEN* with increasing R and




Figure 6: The impact of the number of watermarks (R) in *WARDEN* against *CSE* on SST2 dataset. Note: ‘Recon. Cos. Sim.’ (the Red line) represents the minimum reconstructed cosine similarity among all possible watermarks in W .

marginal degradation (blue line) in the downstream utility. The results on other datasets also show similar patterns which can be found in Appendix D.1.

WARDEN against CSE Now, we investigate the effectiveness of *WARDEN* against *CSE* (shown in Figure 6). As observed in the previous section, *WARDEN* is stealthier with increasing watermarks. As expected, the performance of *CSE* diminishes, correlating with decreasing attack performance (red line). Due to the usage of more watermarks, there is a natural increase in the likelihood that one of them will detect an infringement. Moreover, in extreme scenarios, a mixture of multiple target embeddings will substitute the watermarked samples (Equation 4), reducing the impact of the *CSE* attack’s exploitation of the semantic distortion in the embeddings.

Gram-Schmidt Extension To further strengthen *WARDEN*, we investigate the application of the Gram-Schmidt (GS) process (Trefethen and Bau, 1997) on target embeddings W , as we assume the orthogonal set of watermark embeddings are more distinguishable to each other. In our experiments, as reported in Figure 7, the detection performance is stronger after GS selection. In addition, due to orthogonality, the reconstructed target embedding cosine similarities will be significantly lower, indicating *CSE* might also be ineffective. We observe the same from the corresponding ablation study in Appendix D.1.




Figure 7: The impact of GS extension on *WARDEN* for SST2 dataset. The observations are in line with normal *WARDEN* (Figure 6) results with the only difference being stronger metrics.

Ablation Study Similar to the experiments for *CSE*, we perform *WARDEN* on non-watermarked models. Due to our strict verification, for the high value of R , the p-value could be noisy. It is because the verification process might find closeness due to genuine semantics instead of backdoors as a result of a high pool of watermark directions. This could lead to false positives, i.e., incorrectly classifying models as copied. However, in such cases, we observe that other detection metrics (Δ based) metrics are reliable, which should aid the entity in making appropriate decisions (refer Figure 18). We conduct a further detailed ablation study dissecting the *WARDEN* components and showing its stealthiness in Appendix D.

5 Conclusion

In this paper, we first demonstrate that our new *CSE* attack can bypass the recent EaaS watermarking technique. *CSE* cleanses the watermarked dataset by clustering them first, then selecting embedding pairs with disparity, and finally eliminating their top principal components, while maintaining the service utility. To remedy this shortcoming, we propose a simple yet effective watermarking method, *WARDEN*, which augments the previous approach by introducing multiple watermarks to embeddings. Our intensive experiments show that *WARDEN* is superior in verifying the copyright of EaaS from prior works. Furthermore, *WARDEN* is also effective against potent *CSE*, which shows its resilience to different attacks. We also conduct detailed ab-

lation studies to verify the importance of every component of *CSE* and *WARDEN*. Future studies may consider exploring watermark ownership under multi-owner service settings.

Limitations

We test our *WARDEN* defense against *CSE* attack, acknowledging that various other attacks might overcome the uni-directional watermark approach. Although publishing the *WARDEN* algorithm to the public may inspire future attacks against it, we do not foresee it to be a trivial task, as the capability of *WARDEN* can be enhanced by using more conservative strategies, e.g., more stealthy trigger patterns and watermarking techniques.

We also know that by having access to the ground-truth watermarking vectors, combined with the GS process, one can eliminate the *WARDEN* watermarks, as discussed in Appendix D.3. However, it is the service provider’s responsibility to maintain the confidentiality of their watermarking keys.

We also note the false positive of the p-value for non-watermarked models when a large number of watermark vectors are augmented (Figure 18), even though other metrics rectify this incorrect signal. We suggest service providers conduct a preliminary study to select the number of watermarks for *WARDEN*. In this regard, the current work is an empirical observation study, and theoretical analysis might help decide the optimal number of watermarks. In future, we will endeavour to investigate advanced watermarking mechanisms focusing on defence purposes.

Social Impact

We developed a *CSE* attack, which could aid attackers in circumventing EaaS IP infringements. We agree that with *CSE*, any existing system using EmbMarker is vulnerable. Yet, we argue that it is critical to show the possibility of such attacks and make users aware of them. The usual first step in security is to first expose the vulnerability. Additionally, to mitigate the aforementioned threat, we contribute an improved watermarking technique, *WARDEN*, which could be incorporated with minimal effort.

Acknowledgements

We would like to appreciate the valuable feedback from all anonymous reviewers. This research was

supported by The University of Melbourne’s Research Computing Services and the Petascale Campus Initiative. Qionikai Xu would like to express his gratitude to the FSE Staff Travel Scheme and FSE DDRI grant for the support in both travel and research. Anudeex Shetty would like to express his gratitude to Newman College for the support in travel and research.

References

- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. *Generating natural language adversarial examples*. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Brussels, Belgium. Association for Computational Linguistics.
- David Arthur, Sergei Vassilvitskii, et al. 2007. *k-means++: The advantages of careful seeding*. In *Soda*, volume 7, pages 1027–1035.
- Vance W. Berger and YanYan Zhou. 2014. *Kolmogorov–Smirnov Test: Overview*. John Wiley & Sons, Ltd.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Varun Chandrasekaran, Kamalika Chaudhuri, Irene Giacomelli, Somesh Jha, and Songbai Yan. 2020. *Exploring connections between active learning and model extraction*. In *Proceedings of the 29th USENIX Conference on Security Symposium, SEC’20, USA*. USENIX Association.
- Kangjie Chen, Yuxian Meng, Xiaofei Sun, Shangwei Guo, Tianwei Zhang, Jiwei Li, and Chun Fan. 2022. *Badpre: Task-agnostic backdoor attacks to pre-trained NLP foundation models*. In *International Conference on Learning Representations*.
- Jiazhu Dai, Chuanshuai Chen, and Yufeng Li. 2019. *A backdoor attack against lstm-based text classification systems*. *IEEE Access*, 7:138872–138878.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. *BERT: Pre-training of deep bidirectional transformers for language understanding*. In *Proceedings of the 2019 Conference of*

- the North American Chapter of the Association for Computational Linguistics: *Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. **HotFlip: White-box adversarial examples for text classification**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia. Association for Computational Linguistics.
- G. H. Golub and C. Reinsch. 1970. **Singular value decomposition and least squares solutions**. *Numer. Math.*, 14(5):403–420.
- Xuanli He, Lingjuan Lyu, Chen Chen, and Qionikai Xu. 2022a. **Extracted BERT model leaks more information than you think!** In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1530–1537, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xuanli He, Lingjuan Lyu, Lichao Sun, and Qionikai Xu. 2021a. **Model extraction and adversarial transferability, your BERT is vulnerable!** In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2006–2012, Online. Association for Computational Linguistics.
- Xuanli He, Lingjuan Lyu, Lichao Sun, and Qionikai Xu. 2021b. **Model extraction and adversarial transferability, your BERT is vulnerable!** In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2006–2012, Online. Association for Computational Linguistics.
- Xuanli He, Qionikai Xu, Lingjuan Lyu, Fangzhao Wu, and Chenguang Wang. 2022b. **Protecting intellectual property of language generation APIs with lexical watermark**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):10758–10766.
- Xuanli He, Qionikai Xu, Yi Zeng, Lingjuan Lyu, Fangzhao Wu, Jiwei Li, and Ruoxi Jia. 2022c. **CATER: Intellectual property protection on text generation APIs via conditional watermarks**. In *Advances in Neural Information Processing Systems*.
- Yujin Huang, Terry Yue Zhuo, Qionikai Xu, Han Hu, Xingliang Yuan, and Chunyang Chen. 2023. **Training-free lexical backdoor attacks on language models**. In *Proceedings of the ACM Web Conference 2023*, pages 2198–2208.
- Mika Juuti, Sebastian Szlyler, Samuel Marchal, and N Asokan. 2019. **Prada: protecting against dnn model stealing attacks**. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 512–527. IEEE.
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. 2023. **A watermark for large language models**. In *International Conference on Machine Learning*, pages 17061–17084. PMLR.
- Kalpesh Krishna, Gaurav Singh Tomar, Ankur P. Parikh, Nicolas Papernot, and Mohit Iyyer. 2020. **Thieves on sesame street! model extraction of bert-based apis**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Yiming Li, Yang Bai, Yong Jiang, Yong Yang, Shu-Tao Xia, and Bo Li. 2022. **Untargeted backdoor watermark: Towards harmless and stealthy dataset copyright protection**. In *Advances in Neural Information Processing Systems*.
- Jian Han Lim, Chee Seng Chan, Kam Woh Ng, Lixin Fan, and Qiang Yang. 2022. **Protect, show, attend and tell: Empowering image captioning models with ownership protection**. *Pattern Recognition*, 122:108285.
- Yupei Liu, Jinyuan Jia, Hongbin Liu, and Neil Zhenqiang Gong. 2022. **Stolenencoder: Stealing pre-trained encoders in self-supervised learning**. In *CCS 2022 - Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, Proceedings of the ACM Conference on Computer and Communications Security, pages 2115–2128. Association for Computing Machinery.
- Ilya Loshchilov and Frank Hutter. 2019. **Decoupled weight decay regularization**. In *International Conference on Learning Representations*.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. **Pointer sentinel mixture models**. In *ICLR*.
- Vangelis Metsis, Ion Androutsopoulos, and Georgios Paliouras. 2006. **Spam filtering with naive bayes-which naive bayes?** In *CEAS*, volume 17, pages 28–69. Mountain View, CA.
- OpenAI. 2024. **New embedding models and API updates** — openai.com. <https://openai.com/blog/new-embedding-models-and-api-updates>. [Accessed 02-02-2024].
- Tribhuvanesh Orekondy, Bernt Schiele, and Mario Fritz. 2019. **Knockoff nets: Stealing functionality of black-box models**. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4949–4958.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. **Scikit-learn: Machine learning in Python**. *Journal of Machine Learning Research*, 12:2825–2830.

- Wenjun Peng, Jingwei Yi, Fangzhao Wu, Shangxi Wu, Bin Bin Zhu, Lingjuan Lyu, Binxing Jiao, Tong Xu, Guangzhong Sun, and Xing Xie. 2023. [Are you copying my model? protecting the copyright of large language models for EaaS via backdoor watermark](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7653–7668, Toronto, Canada. Association for Computational Linguistics.
- Christian Perwass, Herbert Edelsbrunner, Leif Kobbelt, and Konrad Polthier. 2009. *Geometric algebra with applications in engineering*, volume 4. Springer.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Douglas A Reynolds et al. 2009. Gaussian mixture models. *Encyclopedia of biometrics*, 741(659-663).
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Yuanmin Tang, Jing Yu, Keke Gai, Xiangyan Qu, Yue Hu, Gang Xiong, and Qi Wu. 2023. [Watermarking vision-language pre-trained models for multi-modal embedding as a service](#).
- Florian Tramèr, Fan Zhang, Ari Juels, Michael K. Reiter, and Thomas Ristenpart. 2016. Stealing machine learning models via prediction apis. In *Proceedings of the 25th USENIX Conference on Security Symposium, SEC'16*, page 601–618, USA. USENIX Association.
- Lloyd N. Trefethen and David Bau. 1997. *Numerical Linear Algebra*. SIAM.
- Yusuke Uchida, Yuki Nagai, Shigeyuki Sakazawa, and Shin'ichi Satoh. 2017. [Embedding watermarks into deep neural networks](#). In *Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval, ICMR '17*, page 269–277, New York, NY, USA. Association for Computing Machinery.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- Eric Wallace, Mitchell Stern, and Dawn Xiaodong Song. 2020. [Imitation attacks and defenses for black-box machine translation systems](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi 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 Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, and Ming Zhou. 2020. [MIND: A large-scale dataset for news recommendation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3597–3606, Online. Association for Computational Linguistics.
- Qionikai Xu and Xuanli He. 2023. [Security challenges in natural language processing models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*, pages 7–12, Singapore. Association for Computational Linguistics.
- Qionikai Xu, Xuanli He, Lingjuan Lyu, Lizhen Qu, and Gholamreza Haffari. 2022. [Student surpasses teacher: Imitation attack for black-box NLP APIs](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2849–2860, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Zhenrui Yue, Zhankui He, Huimin Zeng, and Julian McAuley. 2021. [Black-box attacks on sequential recommenders via data-free model extraction](#). In *Proceedings of the 15th ACM Conference on Recommender Systems, RecSys '21*, page 44–54, New York, NY, USA. Association for Computing Machinery.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, NIPS'15*, page 649–657, Cambridge, MA, USA. MIT Press.
- Zhengyan Zhang, Guangxuan Xiao, Yongwei Li, Tian Lv, Fanchao Qi, Zhiyuan Liu, Yasheng Wang, Xin Jiang, and Maosong Sun. 2023. [Red alarm for pre-trained models: Universal vulnerability to neuron-level backdoor attacks](#). *Machine Intelligence Research*, 20(2):180–193.

Appendix

A Experimental Settings

We leverage the standard codebase of the Transformers (Wolf et al., 2020) library and AdamW (Loshchilov and Hutter, 2019) algorithm for model training and development. Likewise, we use scikit-learn (Pedregosa et al., 2011) for clustering algorithms and other utility calculations. We use GPT-3 text-embedding-002 API as original benign embeddings and the BERT (Devlin et al., 2019) model as the victim model. We perform all the experiments on a single A100 GPU with CUDA 11.7 and pytorch 2.1.2. We assume that both the victim model and imitators use the same datasets to separate the effects of the watermarking technique from other factors. Furthermore, we assume that the extracted model is trained only using the watermarked outputs from the victim model. Finally, we implement the EmbMarker and other experiments following their default configurations and settings, i.e., $m = 4$, $n = 20$, and frequency interval = [0.5%, 1%]. The only exception is the $R = 10$ case in *WAR-DEN*, where we use $n = 50$ to have enough trigger words. A standard dataset, WikiText (Merity et al., 2016) consisting of 1, 801, 350 entries, serves as a hold-out dataset for selecting moderate-frequency words as watermark triggers (T).

B Similarity Distribution Plots

The observations (captured in Figure 8) for other datasets are similar to SST2 as seen in Figure 4, i.e., watermarked embeddings are closer to target embedding, and there is a clear difference in similarities for watermarked and non-watermarked embeddings. Due to skewness between the number of suspected and unsuspected embeddings, we employ sampling for unsuspected entries in these plots.

C CSE Attack Analyses

In this section, we perform detailed ablation studies for *CSE* attack.

C.1 Comparison of Clustering Algorithms

The previous experiments utilize K-Means as the clustering algorithm. However, alternative algorithms such as Gaussian Mixture Models (GMM) (Reynolds et al., 2009) are also valid options. The subsequent table, Table 5, illustrates the comparative performance between K-Means and GMM. While K-Means exhibits superior performance

Dataset	Detection Performance		
	p-value	$\Delta_{cos}(\%)$	$\Delta_{l2}(\%)$
SST2	> 0.02	1.00±0.40	-2.00±0.80
MIND	> 0.55	0.28±0.31	-0.55±0.63
AG News	> 0.10	0.45±0.42	-0.90±0.84
Enron	> 0.56	0.23±0.52	-0.47±1.04

Table 5: *CSE* performance using GMM clustering algorithm, similar to K-Means algorithm (tabulated in Table 3).

in downstream utility for AG News, Enron, and MIND datasets, it is less confident for delta values in the case of Enron and MIND. Overall both algorithms demonstrate satisfactory performance, suggesting that the clustering module for *CSE* is universally adaptable to different clustering algorithms.

C.2 Number of Clusters (n)

We can see from Figure 9 that there is no significant role in the number of clusters (n). In all the cases, *CSE* attack is successful, though we use $n = 20$ in our experiments. However, considering pairwise distance comparison, it is preferable to have fewer clusters to increase the likelihood of watermarked pairs. We also visualize these clusters in Figure 10.

C.3 Number of Principal Components (K)

In Section 4.2, we formulated an optimization problem to compute how much of the watermark we are recovering for a given number of principal components (c_k). As expected, with increasing K , we will recover more of the target embeddings, as noted from the red line in Figure 11. However, the utility metrics deteriorate more significantly (blue line). Meanwhile, a lower K does not recover enough target embedding to bypass copyright verification (yellow and green lines – detection performance). To achieve the best of both worlds—downstream utility and avoiding watermark—we must strike a balance, and 50 components seem appropriate. Further, in some cases, we observe that the increasing K does not affect the downstream utility metrics. The downstream task’s simplicity could be the cause of this. For example, Enron dataset is a binary classification task wherein required data could be represented in a few embedding dimensions.




Figure 8: Distribution plots for cosine similarities between different types of embeddings and the target embedding for different datasets.

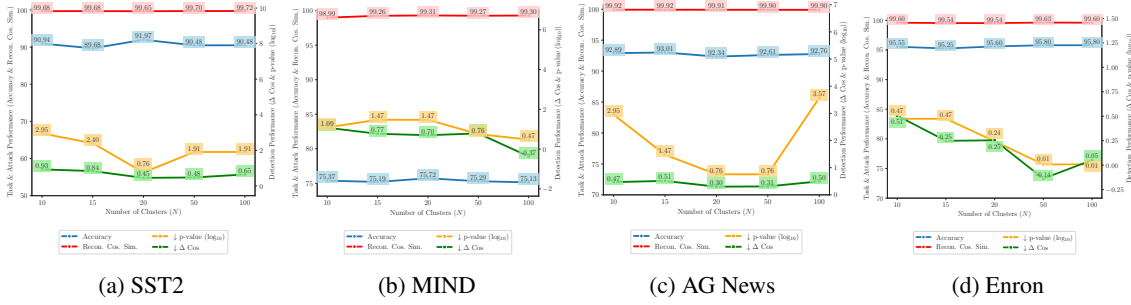


Figure 9: The impact of cluster numbers (n) in *CSE* for different datasets.

Dataset	Size	Detection Performance		
		p-value	$\Delta_{cos}(\%)$	$\Delta_{l2}(\%)$
SST2	Small	> 0.57	0.41	-0.81
MIND		$> 10^{-4}$	1.38	-2.76
AG News		> 0.57	-0.08	0.17
Enron		> 0.08	0.63	-1.27
SST2	Base	> 0.17	0.00	-0.01
MIND		$> 10^{-3}$	-0.01	0.03
AG News		> 0.17	0.00	-0.01
Enron		> 0.57	0.00	-0.01
SST2	Large	> 0.56	0.89	-1.79
MIND		> 0.17	0.37	-0.74
AG News		> 0.98	0.04	-0.09
Enron		> 0.57	0.28	-0.56

Table 6: The impact of extracted model size on *CSE* performance.

C.4 Impact of Attacker Model Size

We evaluate if there are any differences in our attack’s performance for different attacker model sizes. This is tested by performing experiments on small, base, and large versions of the BERT (Devlin et al., 2019) model. As illustrated in the Table 6, the attack effectively bypasses the watermark when the stealer uses different sizes of the backbone model.

D WARDEN Defense Analyses

D.1 Remaining Dataset Results

In this subsection, we present the results (Figure 12-15) for all the remaining datasets discussed in Section 4.3.

D.2 Number of Principal Components (K) in *CSE*

Because we augment more target embeddings in *WARDEN*, default configurations of *CSE* may not be adequate. We tweak the *CSE* attack by increasing the principal components eliminated, as shown in Figure 16. We note that for a higher number of principal components, *CSE* is effective in some setups, as we recover more of the target embeddings. However, one thing to note is that the downstream metrics are poor if we eliminate a large number of principal components, undermining the attacker’s objective.

D.3 Access to Target Embeddings

In the unlikely event that an attacker has access to all the target embeddings, it should be possible to bypass *WARDEN*. We replicate this scenario and apply the elimination step of the *CSE* attack using these embeddings. As expected, in such a case, we can circumvent *WARDEN* (see Figure 17).




Figure 10: t-SNE (Van der Maaten and Hinton, 2008) visualisations for K-Means clustering ($n = 3$) for different datasets. It is evident from the plots that the watermarked samples are not clustered together but instead spread across the embedding space with non-coinciding centroids.

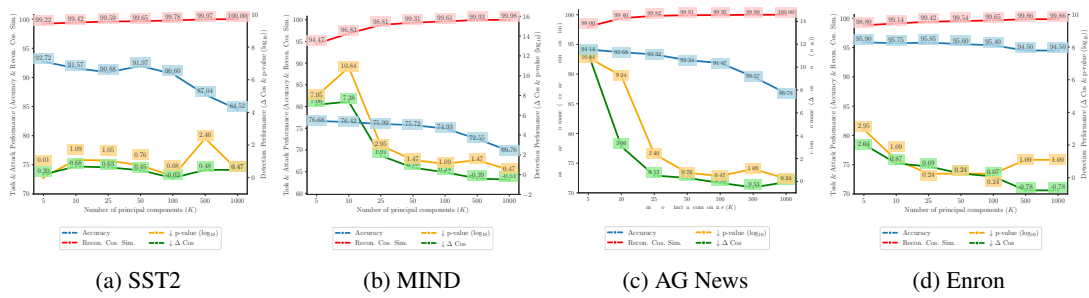


Figure 11: The impact of number of principal components (N) in CSE for different datasets.

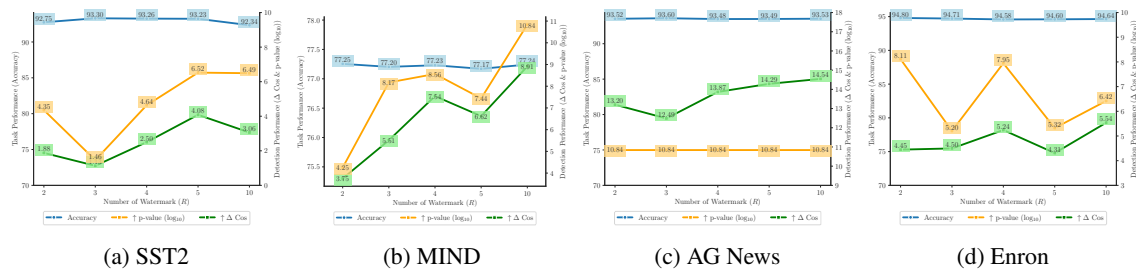


Figure 12: The impact of number of watermarks (R) in WARDEN for different datasets. As expected, detection performance (yellow and green lines) shows an upward trend with stable task performance.

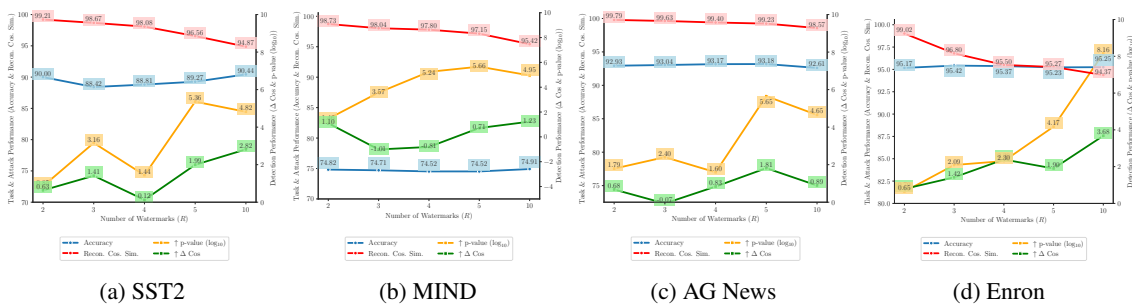


Figure 13: The impact of number of watermarks (R) in WARDEN against CSE for different datasets. The observation is similar to Figure 12, along with a decreasing trend in attack performance (red line) demonstrating the effectiveness of WARDEN defense against CSE attack.

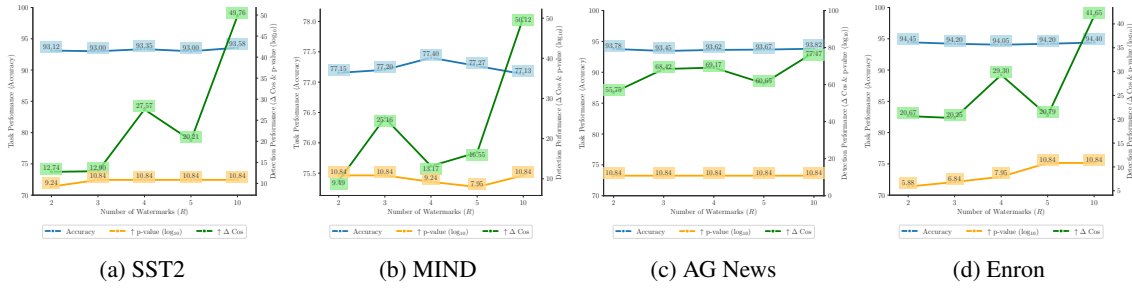


Figure 14: The impact of number of watermarks (R) in *WARDEN* GS extension for remaining datasets. Same trend as Figure 12, but stronger metrics.

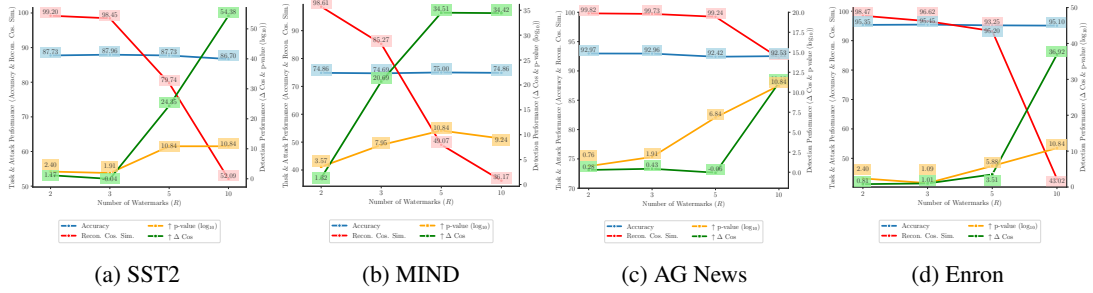


Figure 15: The impact of number of watermarks (R) in *WARDEN* GS extension against *CSE* for different datasets. Same trend as 13, but stronger metrics.




Figure 16: Against *WARDEN*, the impact of number of principal components (K) in *CSE* for different datasets.

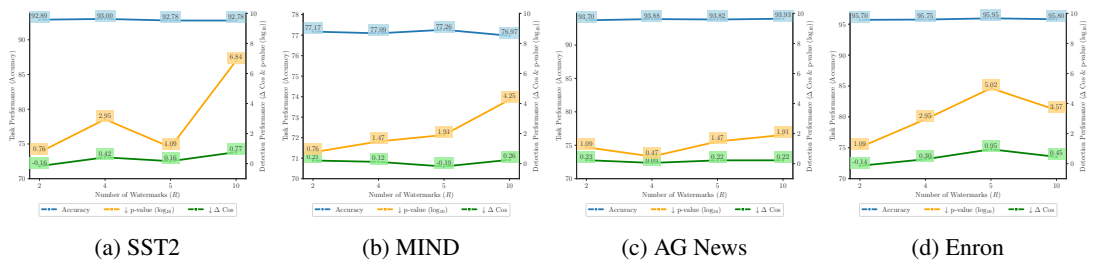


Figure 17: *WARDEN* detection performance when secret watermarks (W) are eliminated for different datasets.

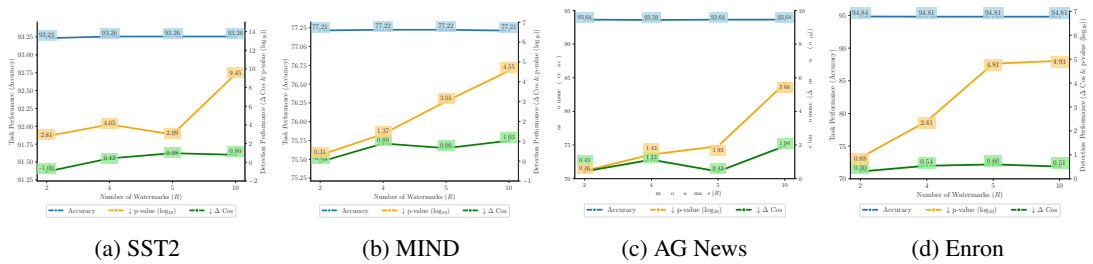


Figure 18: *WARDEN* detection performance on a non-watermarked victim model for different datasets for different numbers of watermarks (R).