

ETHICIST: Targeted Training Data Extraction Through Loss Smoothed Soft Prompting and Calibrated Confidence Estimation

Zhexin Zhang, Jiaxin Wen, Minlie Huang*

The CoAI group, DCST; Institute for Artificial Intelligence; State Key Lab of Intelligent Technology and Systems; Beijing National Research Center for Information Science and Technology; Tsinghua University, Beijing 100084, China.
zx-zhang22@mails.tsinghua.edu.cn, aihuang@tsinghua.edu.cn

Abstract

Large pre-trained language models achieve impressive results across many tasks. However, recent works point out that pre-trained language models may memorize a considerable fraction of their training data, leading to the privacy risk of information leakage. In this paper, we propose a method named ETHICIST for *targeted training data Extraction THrough loss smoothed soft prompting and calibrated Confidence eSTimation*, investigating how to recover the suffix in the training data when given a prefix. To elicit memorization in the attacked model, we tune soft prompt embeddings while keeping the model fixed. We further propose a smoothing loss that smooths the loss distribution of the suffix tokens to make it easier to sample the correct suffix. In order to select the most probable suffix from a collection of sampled suffixes and estimate the prediction confidence, we propose a calibrated confidence estimation method, which normalizes the confidence of the generated suffixes with a local estimation. We show that ETHICIST significantly improves the extraction performance on a recently proposed public benchmark. We also investigate several factors influencing the data extraction performance, including decoding strategy, model scale, prefix length, and suffix length. Our code is available at <https://github.com/thu-coai/Targeted-Data-Extraction>.

1 Introduction

Large pre-trained language models have achieved impressive results on various natural language processing tasks (Devlin et al., 2019; Radford et al., 2019a; Raffel et al., 2020). Model sizes rapidly increase from millions to trillions of parameters and keep growing to achieve better performance and even obtain some emergent abilities (Brown et al., 2020; Chowdhery et al., 2022; Wei et al., 2022; Fedus et al., 2022; Zhang et al., 2022). Despite the

*Corresponding author.

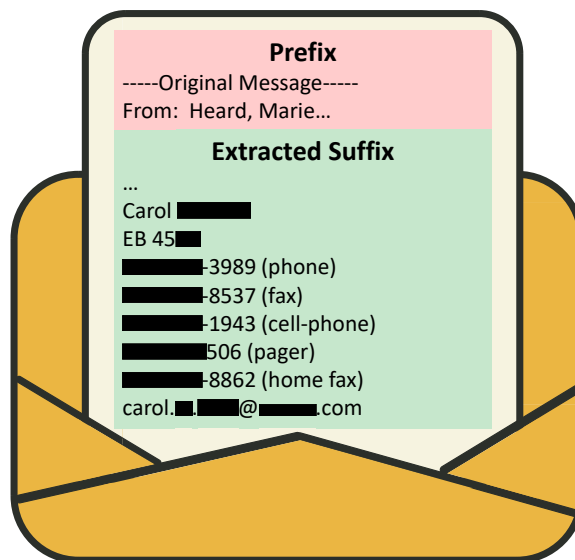


Figure 1: Given a prefix, ETHICIST extracts the verbatim suffix in the training data from the GPT-Neo 1.3B model. The extracted suffix in this example leaks private information about individuals (which is masked for privacy concerns), including name, phone number, fax, pager, home fax, and email.

success of large-scale pre-trained language models, recent works point out that they may memorize a considerable fraction of training data, leading to the privacy risk of information leakage (Carlini et al., 2022a; Tirumala et al., 2022a; Mireshghallah et al., 2022; Carlini et al., 2021). Furthermore, researchers find that memorization scales with model sizes (Carlini et al., 2022a). Therefore, this privacy risk becomes more and more critical in the era of large-scale pre-training. And attacking language models to extract their training data attracts increasing attention.

There are currently two main settings to extract training data. One is membership inference attack, which infers whether a given example is contained in the model’s training data (Hisamoto et al., 2020; Shokri et al., 2017). The other is untargeted train-

ing data extraction (Carlini et al., 2021), which aims to extract training data from scratch (i.e., without the given prefix). However, both settings are not suitable for extracting targeted training data. For example, attackers may feed the model with a prefix indicating the beginning of an email and try to extract the following private email content in the training dataset as shown in Figure 1. In such cases, we do not have complete examples to do membership inference, and we have specific goals instead of performing untargeted extraction. Therefore, we focus on **targeted training data extraction** in this paper, which requires recovering the suffix when given a prefix according to the training data. Compared with untargeted training data extraction, the task matters more because attackers can recover specific types of training data instead of any possible training data that might be harmless. What’s more, it is easier to evaluate targeted training data extraction because we just need to compare the prediction with the ground truth suffix. However, for untargeted training data extraction, we need to search over the whole massive pre-training dataset (e.g., The Pile dataset (Gao et al., 2020), which has 800GB text data) to check whether it contains the generated sample, which is very slow and costly.

The general process for targeted training data extraction can be divided into two steps: (1) generating one or more possible suffixes based on the given prefix, and (2) choosing a most likely suffix as the prediction result based on a confidence estimation method. We summarize two challenges of this task: (1) how to increase the generation likelihood of the ground truth suffix, and (2) how to estimate the confidence accurately so that the confidence score can be meaningfully interpreted as the probability that the output suffix is correct. To tackle these challenges, we propose a method named ETHICIST for *targeted training data Extraction THrough loss smoothed soft prompting and calibrated Confidence eSTimation*. For the first challenge, we propose loss smoothed soft prompting. It uses soft prompt to elicit memorization in the attacked model, and adds an additional loss besides the maximum likelihood estimation (MLE) loss to smooth the loss distribution of the suffix tokens. Through the loss smoothing, we hope to ensure that the probability of the ground truth token at each time step is not low, which makes it more likely to sample the ground truth suffix. With the two loss functions, we tune the prepended soft

prompt tokens on an extracted training set which contains pairs of prefixes and ground truth suffixes. The existence of a training set is reasonable because large-scale pre-trained data generally contain public data (e.g., Common Crawl)¹. For the second challenge, we propose a calibrated confidence estimation method. We find that the model’s perplexity cannot accurately represent the probability that the generated suffix is correct because the prediction probabilities for diversified prefixes are inherently different and incomparable. We thus normalize the confidence of the generated suffixes with a local estimation, which can mitigate the problems caused by intrinsic differences in the difficulties of distinct samples. We verify ETHICIST on a recently proposed public benchmark containing 15,000 pairs of prefixes and suffixes derived from The Pile dataset (Gao et al., 2020). Experiments show that ETHICIST can significantly improve the extraction performance, which suggests that existing large language models are at significant risk of leaking training data. We also discuss and analyze several factors influencing the data extraction performance, including decoding strategy, model scale, prefix length, and suffix length.

Our contributions can be summarized as follows:

- We propose loss smoothed soft prompting to reduce the difficulties of sampling the ground truth suffixes.
- We propose a calibrated confidence estimation method that enables the confidence score to be meaningfully interpreted as the probability that the output suffix is correct.
- Experiments on a recently proposed benchmark demonstrate that ETHICIST can consistently and significantly improve the data extraction performance across various model sizes. We further investigate several factors influencing the data extraction performance.

2 Related Work

2.1 Training Data Extraction

Existing works on training data extraction mainly focus on membership inference attack or untargeted training data extraction. For membership inference attack, adversaries need to judge whether a given example is contained in the training data of the attacked model. Shokri et al. (2017); Song

¹Similar setting is adopted in Hisamoto et al. (2020).

and Shmatikov (2019) train several shadow models that mimic the attacked models’ behaviors to help train an auditing model that can predict whether an example is contained in the training dataset. Hisamoto et al. (2020) perform membership inference attacks on machine translation systems. They find it is harder to attack sequence generation models than classification models. Song and Raghunathan (2020) show that the encoded dense representations can leak information under membership inference attack. Mireshghallah et al. (2022) focuses on attacking masked language models that are pre-trained on possibly sensitive data (e.g., clinical notes). They introduce an additional reference masked language model besides the original attacked model and compute the ratio of the likelihood measured by the attacked model and the reference model, which is better than solely relying on the attacked model.

For untargeted training data extraction, adversaries first generate various samples using the attacked model and then predict whether they are contained in its training set. Carlini et al. (2021) extract hundreds of verbatim sequences from the popular pre-trained language model GPT-2 (Radford et al., 2019b). And there is privacy information such as names, phone numbers, and email addresses in the extracted sequences. Lehman et al. (2021) try to extract sensitive information from BERT (Devlin et al., 2019) pre-trained on clinical notes. However, they are mostly unable to meaningfully expose Personal Health Information by simply using templates. Different from the existing works, we focus on targeted training data extraction that aims to recover the suffix when given a prefix, which is more security-critical and easier to evaluate.

2.2 Memorization

We generally expect models can gain the generalization ability from the training process. However, recent works point out that models may unintentionally memorize the training data even without overfitting (Tirumala et al., 2022a; Carlini et al., 2022a, 2019; Béguelin et al., 2020). One possible method to mitigate this problem is to deduplicate training data (Kandpal et al., 2022). However, Carlini et al. (2019) also show that it is possible to recover samples appearing only once in the training dataset. Surprisingly, Tirumala et al. (2022a) find that there is a forgetting baseline during the pre-training of the casual language model (e.g., the model can

memorize at least 40% of the data that appear only once, even being trained on other data for many epochs afterward). These findings further emphasizes the difficulties of avoiding memorization and the potential threats of unintended memorization in large-scale pre-trained language models. Another line of work uses differential privacy to avoid the memorization problem (Abadi et al., 2016; McMahan et al., 2018; Shokri and Shmatikov, 2015), but the mechanism could reduce the accuracy (Jayaraman and Evans, 2019; Feldman and Zhang, 2020; Feldman, 2020; Song and Shmatikov, 2019). Differential privacy also increases the training time, which can further influence the accuracy within the same budget. Therefore there is still no effective and practical way to avoid unintended memorization. Our work further verifies the existence of unintended memorization and makes it more necessary to develop practical defense methods.

3 Methodology

We formulate the targeted training data extraction task as follows: given a source prefix $S = (s_1, s_2, \dots, s_{|S|})$ with $|S|$ tokens, the attacker should predict the target suffix $T = (t_1, t_2, \dots, t_{|T|})$ with $|T|$ tokens and its confidence. The pair of the given prefix and the predicted suffix (S, T) should be contained in the pre-training dataset $D_{\text{pretrain}} = \{(S_i, T_i)\}$, which the attacked model M is trained on. The prediction of the confidence score is necessary for picking out the most probable suffix when we don’t know the ground truth suffix in realistic attack scenarios (i.e., we need to pick out most probable pairs of prefixes and extracted suffixes based on their confidence scores among all predictions). We assume the attacker can obtain some pairs of ground truth prefixes and suffixes $D_{\text{train}} = \{(S_i, T_i) | (S_i, T_i) \in D_{\text{pretrain}}, 1 \leq i \leq |D_{\text{pretrain}}|\}$ before attacking, which is reasonable because large-scale pre-trained data generally contain public data (e.g., Common Crawl). The attackers can utilize D_{train} to train their attacking models and their goal is to predict suffixes for the prefixes in the test set $D_{\text{test}} = \{S_i | 1 \leq i \leq |D_{\text{test}}|\}$. Note that the prefix S_i in D_{test} is included in D_{pretrain} but is not a part of D_{train} .

3.1 Method Overview

An overview of ETHICIST is shown in Figure 2. We first tune the soft prompt embeddings during training to elicit memorization in the attacked model M

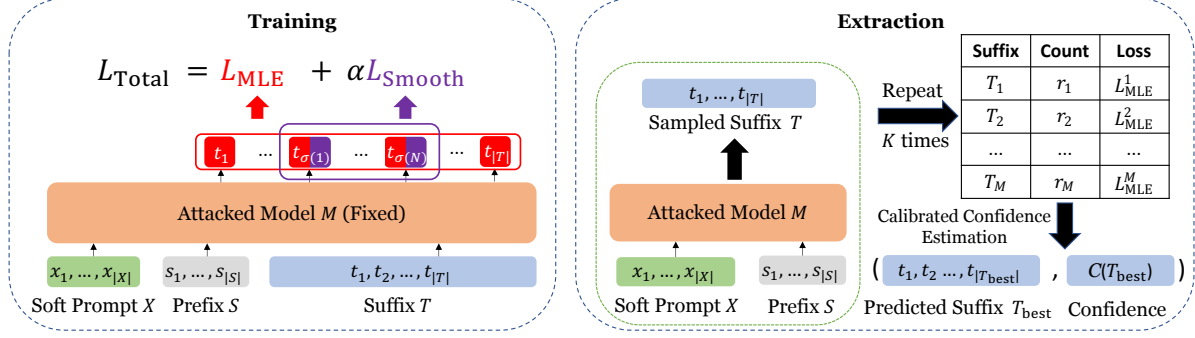


Figure 2: Method overview. During training, we fix the parameters of the attacked model M and only tune the parameters of the soft prompt embeddings. Besides the MLE loss, we additionally design a smoothing loss to make the loss distribution of the suffix sequence more smooth. After tuning the soft prompt embeddings, we extract training data by repeatedly sampling K suffixes conditioned on one given prefix and using calibrated confidence estimation to select the final predicted suffix T_{best} and provide its confidence $C(T_{\text{best}})$.

with the MLE loss and the additional smoothing loss. The smoothing loss aims to increase the probability of sampling the ground truth suffix. After prompt tuning, we repeatedly sample K suffixes using the attacked model M conditioned on one given prefix and reorder them with our calibrated confidence estimation. Our calibrated confidence estimation can not only select the most possible suffix, but also provide a more accurate confidence score that represents how likely the predicted suffix is correct. Finally, the suffix with the highest confidence is selected as the final prediction.

3.2 Prompt Tuning with Smoothing Loss

We adopt prompt tuning to train the soft prompt tokens on D , which prepends $|X|$ soft tokens $X = (x_1, x_2, \dots, x_{|X|})$ before the original input sequence. Then we feed the input to the attacked model M to compute the MLE loss:

$$\mathcal{L}_{\text{MLE}} = \sum_{i=1}^{|T|} -\frac{1}{|T|} \log P_M(t_i | X, S, t_{<i}). \quad (1)$$

Note that we only tune the parameters of the soft prompt tokens and the parameters of the attacked model M are fixed. We use prompt tuning for two reasons: (1) we do not want to change the original parameters of the attacked model M because the main goal is to elicit memorization in M , and (2) prompt tuning is helpful to improve the training efficiency when M is very large, making ETHICIST able to efficiently adapt to larger language models that generally memorize more training data.

The MLE loss aims to increase the total generation probability of the target suffix T . However,

when using popular sampling methods such as top-k sampling (Fan et al., 2018) and top-p (nucleus) sampling (Holtzman et al., 2020) to generate multiple candidate suffixes, we want to ensure the probability of the ground truth suffix token at **each time step** is not low. Suppose the total probability of the ground truth suffix is high while there is one token in the sequence with a low generation probability. In this case, it is still hard to generate the correct suffix using auto-regressive sampling methods. Therefore, we propose a smoothing loss to make the loss distribution of the suffix sequence more smooth. More specifically, we pick out the top- N tokens with the highest loss values in the whole sequence T . Then we additionally optimize the generation probabilities for these N tokens as follows:

$$\mathcal{L}_{\text{Smooth}} = \sum_{i=1}^N -\frac{1}{N} \log P_M(t_{\sigma(i)} | X, S, t_{<\sigma(i)}), \quad (2)$$

where $t_{\sigma(i)}$ represents the token with the i -th highest loss in T . Note that $t_{\sigma(i)}$ is dynamically computed during training. The smoothing loss can also be seen as assigning higher weights to the tokens with higher loss values. Finally, we derive the overall loss function as follows:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{MLE}} + \alpha \mathcal{L}_{\text{Smooth}}, \quad (3)$$

where the coefficient α is a hyperparameter to control the strength of the smoothing loss.

3.3 Calibrated Confidence Estimation

After predicting the suffix, we also need to give a confidence score for the prediction, which can be

meaningfully interpreted as the probability that the output suffix is correct. A naive method is to use the generation likelihood $P_T = \exp(-|T|\mathcal{L}_{\text{MLE}})$ as the confidence score. This naive method is reasonable for picking out the most probable suffix T_i from a collection of sampled suffixes $\{T_1, T_2, \dots, T_M\}$ for one given prefix. However, it is unsuitable for comparing the confidence of different predicted suffixes corresponding to different prefixes. As the language model is essentially a statistical model, frequencies of tokens and n-grams in the prefixes can greatly influence the absolute generation likelihood of the suffixes. For example, consider two predicted suffixes T_A and T_B conditioned on two different prefixes S_A and S_B , where S_A and T_A contain tokens and n-grams with much higher frequencies. The absolute generation likelihood of T_A may be significantly higher than T_B , even if they are both ground truth suffixes. Therefore, to eliminate the intrinsic differences in scales of generation likelihood across different suffixes, we propose a novel calibrated confidence estimation method. To calibrate the confidence estimation, we have two considerations: (1) different generated suffixes conditioned on one given prefix should have comparable scales of generation likelihood, and (2) the memorized ground truth suffix is expected to be generated more frequently during multiple generations, which is also validated in Section 5.

Suppose the sampled distinct suffixes are $\{T_1, T_2, \dots, T_M\}$ for one given prefix, the repeated generation times for these suffixes are $\{r_1, r_2, \dots, r_M\}$ (i.e., r_i denotes how many times T_i is generated among K repeated sampling outputs), and the MLE loss values for these suffixes are $\{\mathcal{L}_{\text{MLE}}^1, \mathcal{L}_{\text{MLE}}^2, \dots, \mathcal{L}_{\text{MLE}}^M\}$. Then we assign the calibrated confidence score to T_i as:

$$C(T_i) = \frac{r_i \times \exp(-|T_i|\mathcal{L}_{\text{MLE}}^i)}{\sum_{j=1}^M r_j \times \exp(-|T_j|\mathcal{L}_{\text{MLE}}^j)}. \quad (4)$$

Through the proposed confidence estimation method, we obtain the confidence score of T_i by comparing it with other sampled suffixes with comparable scales of generation likelihood. In this way, we avoid the scale problem brought by different prefixes and make it practical to compare the predicted suffixes conditioned on different prefixes. Moreover, we leverage the repetition time r_i as a valuable signal since memorized suffix is expected to be generated more frequently. Finally, we select the suffix T_{best} with the highest confidence score

$C(T_{\text{best}})$ among $\{C(T_1), C(T_2), \dots, C(T_M)\}$ as the predicted suffix and $C(T_{\text{best}})$ as its confidence estimation.

4 Experiments

4.1 Benchmark

We evaluate ETHICIST on the LM-Extraction benchmark², which is designed for benchmarking targeted training data extraction attacks. It consists of a subset contained in The Pile dataset (Gao et al., 2020). Both the prefix and the suffix are 50 tokens long. All examples are well-specified, meaning that there is only one 50-token suffix in The Pile dataset given the 50-token prefix. What’s more, these examples are all chosen to meet the property that there exists a prefix length (maybe longer than 50) that causes the model to generate the suffix string exactly, which implies that the extraction performance on this benchmark may be **higher** than that on randomly selected prefixes. We randomly split the dataset into training, validation and test sets. The detailed statistics of the LM-Extraction benchmark are shown in Table 1.

Split	# Examples	# Prefix Len	# Suffix Len
Train	12,600	50	50
Validation	1,400	50	50
Test	1,000	50	50

Table 1: Statistics of the LM-Extraction benchmark.

4.2 Baselines

We compare ETHICIST with the following baselines. All the compared baselines first sample K suffixes $\{T_1, T_2, \dots, T_K\}$ conditioned on one given prefix S and then pick out one suffix as the prediction.

Perplexity It leverages the perplexity (PPL) measured by the attacked language model M as the metric to sort the candidate suffixes and finally chooses the one with the lowest PPL as the predicted suffix T :

$$T = \arg \max_{T_i} C(T_i) = \arg \max_{T_i} \frac{1}{\text{PPL}_M(T_i|S)}$$

Comparing (LM) It takes another language model M' and leverages the ratio of the perplexity

²<https://github.com/google-research/lm-extraction-benchmark/>

measured by these two language models as the metric (Carlini et al., 2021):

$$T = \arg \max_{T_i} C(T_i) = \arg \max_{T_i} \frac{\text{PPL}_{M'}(T_i|S)}{\text{PPL}_M(T_i|S)}$$

The language model M' could be a much smaller model trained on the same dataset with M or trained on a different dataset.

Comparing (zlib) Different from Comparing (LM), it uses the zlib (Gailly and Adler, 2004) entropy of the text (i.e., the number of bits after compression with zlib) for comparison (Carlini et al., 2021):

$$T = \arg \max_{T_i} C(T_i) = \arg \max_{T_i} \frac{\text{len}(\text{zlib}(T_i))}{\text{PPL}_M(T_i|S)}$$

Comparing (lowercase) It compares the perplexity of the original text and the lower-cased text measured by the same language model M (Carlini et al., 2021):

$$\begin{aligned} T &= \arg \max_{T_i} C(T_i) \\ &= \arg \max_{T_i} \frac{\text{PPL}_M(\text{lowercased}(T_i)|S)}{\text{PPL}_M(T_i|S)} \end{aligned}$$

Furthermore, we conduct ablation tests by removing the proposed components respectively to investigate the influence of each component.

4.3 Metrics

We adopt the following automatic metrics for evaluation.

Recall The metric computes the percentage of the suffixes that are predicted verbatim over the whole test set. A higher recall score indicates better data extraction ability, which can also be understood as a higher attacking success rate.

Recall_{Early stop} The metric first sorts the predictions according to their confidence scores and then evaluates the correctness of each prediction one by one. It finally computes the Recall score while making x incorrect predictions. We set x to 100 in our experiments following the LM-Extraction benchmark. A better confidence estimation method can give the correct predictions higher confidence scores and thus lead to a higher Recall_{Early stop} score.

4.4 Main Results

Table 2 shows the automatic evaluation results with GPT-Neo 1.3B as the backbone. ETHICIST achieves an impressive Recall score of 62.8% and outperforms all the baselines by a large margin, indicating its better ability to extract training data from language models. Moreover, ETHICIST has better confidence estimation performance after calibration as shown by a significantly higher Recall_{Early stop} score. To further investigate the influence of each component, we run an ablation study. From the results shown in Table 2, it can be seen that both the smoothing loss and the calibrated confidence estimation are important to enhance the ability to extract training data, and combining both of them achieves the best performance. Furthermore, we draw the following conclusions: (1) With prompt tuning and extra training data, we can better induce large-scale language models to generate their memorized training data and successfully achieves a 9.5% performance improvement on Recall and a 12.4% performance improvement on Recall_{Early stop}. (2) The proposed smoothing loss can further enhance the ability to extract training data, boosting the Recall score from 60.8% to 62.3%. (3) The calibrated confidence provides a 6.3% improvement on Recall_{Early stop} as expected, demonstrating the importance of calibrating confidence estimation for this task. (4) The smoothing loss is more effective in predicting exact suffixes while the calibrated confidence is more beneficial for identifying highly confident predictions, according to the significant drop in Recall without smoothing and the substantial decrease in Recall_{Early stop} without calibration. (5) The calibrated confidence estimation is effective regardless of whether using prompt tuning. And it demonstrates greater advantages compared to the comparing (LM) baseline in recognizing predictions with higher confidence when using prompt tuning, indicated by increasing Recall_{Early stop} (from 48.7 to 52.4).

4.5 Analysis: Decoding Strategy

In our experiments, we use top-p sampling to sample multiple candidate suffixes conditioned on one given prefix. However, there are also other popular decoding methods, including greedy search, beam search, and top-k sampling. We thus compare these popular decoding methods in this section. Table 3 shows the results. Not surprisingly, greedy search performs worst on both Recall and Recall_{Early stop},

Method	Recall	Recall _{Early stop}
Perplexity	51.3 \pm .0	32.2 \pm .0
Comparing (LM)	51.9 \pm .0	37.4 \pm .0
Comparing (zlib)	49.7 \pm .2	25.6 \pm .0
Comparing (lowercase)	51.5 \pm .0	32.5 \pm .0
ETHICIST	62.8 \pm .5	53.8 \pm .5
w/o smooth	61.2 \pm .3	52.4 \pm .5
w/o calibrated	62.3 \pm .6	47.5 \pm 1.3
w/o smooth & calibrated	60.8 \pm .6	44.6 \pm .8
w/o smooth & calibrated, comparing (LM)	62.4 \pm .7	48.7 \pm 1.2
w/o prompt tuning	50.9 \pm .0	38.0 \pm .0

Table 2: Automatic evaluation results on the test set. The experiments are conducted on the GPT-Neo 1.3B model. We report the mean and the standard deviation over 3 random seeds. The best performance are highlighted in **bold**. **w/o smooth** means ablating the smoothed loss function in the training stage. **w/o calibrated** means ablating the calibrated confidence in the extraction stage. **w/o smooth & calibrated** means prompt tuning only with the MLE loss and using the perplexity for confidence estimation. **w/o smooth & calibrated, comparing (LM)** means prompt tuning only with the MLE loss and using the comparing (LM) method for confidence estimation. And **w/o prompt tuning** directly employs calibrated confidence estimation on the original model without prompt tuning.

which suggests some tokens in the ground truth suffix do not have the highest probability at the corresponding positions. Beam search outperforms top-p sampling on Recall, indicating that searching for the suffix with the lowest loss works well to find the ground truth suffix. However, beam search performs significantly worse than top-p sampling on Recall_{Early stop}, because it cannot use our calibrated confidence. Compared with beam search, top-p sampling can generate multiple candidates, which could substantially increase the accuracy of confidence estimation with our proposed calibrated confidence. Moreover, the top-k sampling performs worse than top-p sampling on Recall_{Early stop}, which may be because top-k sampling is easier to sample low-probability tokens and thus reduce the confidence of the ground truth suffixes. We finally select top-p sampling as our decoding method due to its balance on Recall and Recall_{Early stop}.

4.6 Analysis: Model Scale

Previous works on scaling laws find that larger language models can memorize more training data (Carlini et al., 2022b; Tirumala et al., 2022b). Therefore, we are interested in how targeted data extraction performance varies across different model scales. Figure 3 shows the results. We can see that the targeted training data extraction performance continuously increases as the model scale increases from 125 million to 6 billion. ETHICIST shows impressive results as it consistently and

Strategy	Recall	Recall _{Early stop}
Greedy	58.7 \pm .6	47.1 \pm 1.1
Beam Search	64.5 \pm .9	47.9 \pm 1.0
Top-k	62.7 \pm .6	50.8 \pm .6
Top-p	62.8 \pm .5	53.8 \pm .5

Table 3: Effect of the decoding strategy on ETHICIST. Note that our proposed calibrated confidence is unused when decoding with deterministic methods, including greedy search and beam search. We show the mean and the standard deviation over 3 random seeds for all decoding strategies.

significantly outperforms baselines across different model scales. Thanks to prompt tuning, ETHICIST is efficient in terms of computation time and particularly memory consumption. Therefore, ETHICIST can also be adapted to larger language models for efficient targeted training data extraction.

4.7 Analysis: Prefix Length and Suffix Length

All prefixes and suffixes in the LM-Extraction benchmark are 50 tokens long, making it an interesting question how the length of prefixes and suffixes would affect the extraction performance.

We show the effect of the given prefix length in Figure 4. We can observe that the extraction performance grows approximately linearly with the prefix length for all evaluated methods, and ETHICIST performs best for all prefix lengths. Al-

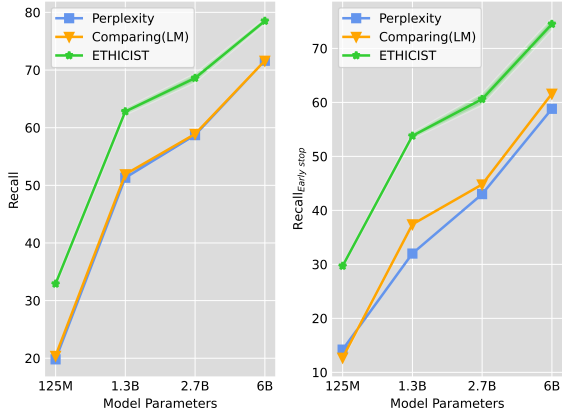


Figure 3: Effect of the model scale. We show the mean and the standard deviation over 3 random seeds for all methods.

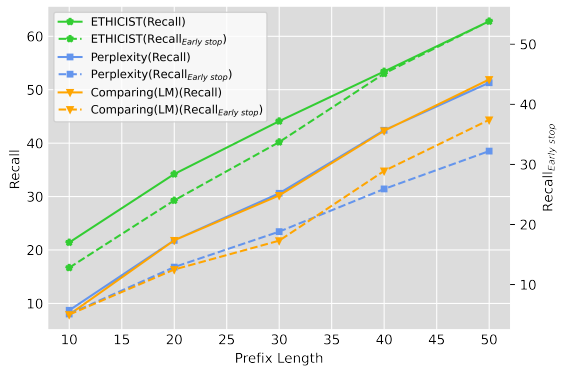


Figure 4: Effect of the given prefix length. We show the **Recall** and the **Recall_{Early stop}** metrics for three methods when the given prefix length increases from 10 to 50.

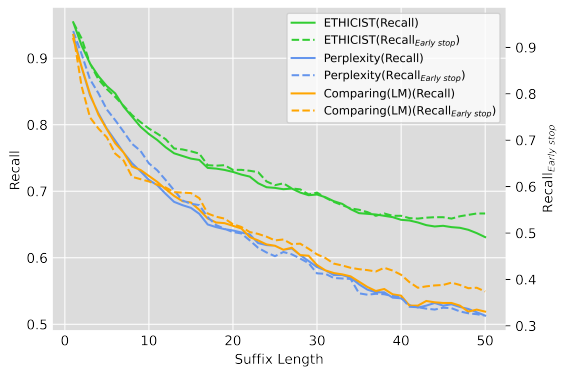


Figure 5: Effect of the predicted suffix length. We show the **Recall** and the **Recall_{Early stop}** metrics when the predicted suffix length increases from 1 to 50.

though all methods have similar growth speed on Recall, ETHICIST has the highest growth speed on Recall_{Early stop}. It is also interesting that *Comparing (LM)* only outperforms *Perplexity* when given

Feature	Correct	Wrong
Recall@1	0.63	0.37
Recall@3	0.68	0.32
Recall@100	0.69	0.31
Average Repeat Time	85.38	29.66
Average Confidence	0.95	0.67

Table 4: Statistical features of correct predictions and wrong predictions. Recall@K measures whether the top-K suffixes sorted by estimated confidence contain the ground truth suffix. Average repeat time represents the number of times that the prediction result is generated repeatedly out of 100 generations.

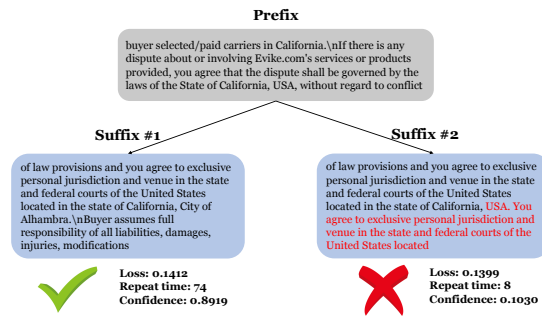


Figure 6: Given the prefix, we show the top-2 predicted suffixes by ETHICIST. Although the first prediction has higher loss, it is repeated for 74 times and we correctly select it as the final predicted suffix using our calibrated confidence estimation. We highlight the wrong predicted text in red.

prefixes that are long enough.

We show the effect of the predicted suffix length in Figure 5. For all three methods, the extraction performance decreases when the suffix length increases. Different from the approximately linear relationship between the prefix length and the extraction performance, the performance degradation tends to become progressively slower as the suffix length increases. This suggests that the model can still memorize a considerable proportion of suffixes (rather than quickly decreasing to zero) even if the predicted suffix length continues to increase. What’s more, we observe that ETHICIST has a significantly slower speed of performance degradation compared with the two baselines, which suggests ETHICIST is effective for eliciting deeper memorization of longer suffixes of the attacked model.

4.8 Analysis: Sampling Time

Due to space limitations, we put the analysis of sampling time in Appendix C.

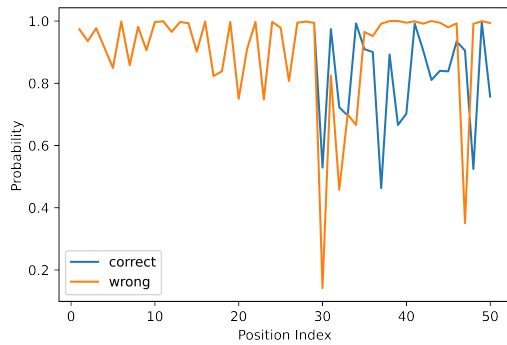


Figure 7: We show the generation probability of each token during the sampling process for both correct and wrong suffixes.

5 Discussion

We further show some statistical features in Table 4. We can see that the memorized suffixes can be sampled significantly more frequently with a high average repeat time of 85.38, validating that the repeat time is a valuable signal for confidence estimation. What’s more, the memorized suffixes have significantly higher confidence. One interesting phenomenon we observe is that if the ground truth suffix can be generated, it mostly has the top 3 highest confidence ($\text{Recall}@3 \approx \text{Recall}@100$). We also find that for more than 30% of the prefixes, the model cannot generate the correct prefix even given 100 chances. Therefore, an important future direction is to design better methods to elicit memorization in the attacked model. Considering the non-negligible gap between $\text{Recall}@1$ and $\text{Recall}@100$ (0.63 vs. 0.69), another important future direction is to design better confidence estimation methods (maybe trainable), which can pick out the ground truth suffix among the collection of candidate suffixes for one prefix.

We show a case in Figure 6. Although the first predicted suffix has higher loss than the second predicted suffix, it is sampled far more times than the latter. Therefore, we assign higher confidence to the first suffix using our calibrated confidence estimation method. We further show the probability of generating each token during the sampling process in Figure 7. We can observe that although the correct prediction has higher loss as a whole, it keeps a high sampling probability across the generation process. The minimum probability of generating one token in the correct suffix is about 0.45, which is significantly higher than 0.1 for the wrong suffix.

Therefore it is easier to generate the correct suffix, which leads to a higher confidence score. This is also in line with our motivation for designing the extra smoothing loss, which can increase the probability of sampling the correct suffix.

6 Conclusion

In this work, we propose ETHICIST, an effective method for targeted training data extraction attack. ETHICIST uses soft prompt to elicit memorization in the attacked model. To ensure the probability of the ground truth suffix token at each time step is not low, we propose a smoothing loss besides the standard MLE loss. We also propose a calibrated confidence estimation method to calibrate the scale of confidence across different samples. Experiments on the LM-Extraction benchmark demonstrate that ETHICIST significantly improves the extraction performance. We further conduct extensive experiments to investigate several critical factors influencing the extraction performance, including decoding strategy, model scale, prefix length, and suffix length. We hope our work can promote future researches on better attack methods and practical defense methods for the training data extraction problem.

Acknowledgement

This work was supported by the NSFC projects (Key project with No. 61936010). This work was also supported by the Guoqiang Institute of Tsinghua University, with Grant No. 2020GQG0005.

Limitations

Although we conduct experiments across various model scales ranging from 125M to 6B, there are still larger language models we don’t test either because their training data is not publicly released or because we have limited resources.

Moreover, the examples in the LM-Extraction benchmark are all chosen to meet the property that there exists a prefix length (maybe longer than 50) that causes the model to generate the suffix string exactly, which makes the extraction performance on this benchmark higher than that on randomly selected prefixes.

Ethics Statement

ETHICIST is a powerful method to elicit memorization in the large pre-trained language models,

which makes it a useful tool to expose the privacy risk of large language models. However, it also has a risk to be abused by attackers to extract privacy information from pre-trained language models. Thus large language models should be carefully examined before being made publicly available. What's more, it is necessary to develop defense methods against the training data extraction attacks without sacrificing the language modeling ability.

The LM-Extraction benchmark is derived from the Pile dataset, and thus covers many domains including books, code, emails, etc. This suggests the effectiveness of targeted training data extraction across different domains.

References

- Martín Abadi, Andy Chu, Ian J. Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. [Deep learning with differential privacy](#). In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 308–318. ACM.
- Santiago Zanella Béguelin, Lukas Wutschitz, Shruti Tople, Victor Rühle, Andrew Paverd, Olga Ohri-menko, Boris Köpf, and Marc Brockschmidt. 2020. [Analyzing information leakage of updates to natural language models](#). In *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*, pages 363–375. ACM.
- Sid Black, Leo Gao, Phil Wang, Connor Leahy, and Stella Biderman. 2021. [GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow](#). If you use this software, please cite it using these metadata.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2022a. [Quantifying memorization across neural language models](#). *CoRR*, abs/2202.07646.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and Chiyuan Zhang. 2022b. Quantifying memorization across neural language models. *arXiv preprint arXiv:2202.07646*.
- Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. 2019. [The secret sharer: Evaluating and testing unintended memorization in neural networks](#). In *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*, pages 267–284. USENIX Association.
- Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Úlfar Erlingsson, et al. 2021. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- 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.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898.
- William Fedus, Barret Zoph, and Noam Shazeer. 2022. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *Journal of Machine Learning Research*, 23(120):1–39.
- Vitaly Feldman. 2020. Does learning require memorization? a short tale about a long tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 954–959.
- Vitaly Feldman and Chiyuan Zhang. 2020. What neural networks memorize and why: Discovering the long tail via influence estimation. *Advances in Neural Information Processing Systems*, 33:2881–2891.
- Jean-loup Gailly and Mark Adler. 2004. Zlib compression library.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Sorami Hisamoto, Matt Post, and Kevin Duh. 2020. [Membership inference attacks on sequence-to-sequence models: Is my data in your machine translation system?](#) *Trans. Assoc. Comput. Linguistics*, 8:49–63.

- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Bargav Jayaraman and David Evans. 2019. [Evaluating differentially private machine learning in practice](#). In *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*, pages 1895–1912. USENIX Association.
- Nikhil Kandpal, Eric Wallace, and Colin Raffel. 2022. [Deduplicating training data mitigates privacy risks in language models](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 10697–10707. PMLR.
- Eric Lehman, Sarthak Jain, Karl Pichotta, Yoav Goldberg, and Byron C. Wallace. 2021. [Does BERT pre-trained on clinical notes reveal sensitive data?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 946–959. Association for Computational Linguistics.
- H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. 2018. [Learning differentially private recurrent language models](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Fatemehsadat Mireshghallah, Kartik Goyal, Archit Uniyal, Taylor Berg-Kirkpatrick, and Reza Shokri. 2022. [Quantifying privacy risks of masked language models using membership inference attacks](#). *CoRR*, abs/2203.03929.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019a. Language models are unsupervised multitask learners.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019b. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Reza Shokri and Vitaly Shmatikov. 2015. [Privacy-preserving deep learning](#). In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015*, pages 1310–1321. ACM.
- Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. [Membership inference attacks against machine learning models](#). In *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*, pages 3–18. IEEE Computer Society.
- Congzheng Song and Ananth Raghunathan. 2020. [Information leakage in embedding models](#). In *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*, pages 377–390. ACM.
- Congzheng Song and Vitaly Shmatikov. 2019. [Auditing data provenance in text-generation models](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*, pages 196–206. ACM.
- Kushal Tirumala, Aram H. Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. 2022a. [Memorization without overfitting: Analyzing the training dynamics of large language models](#). *CoRR*, abs/2205.10770.
- Kushal Tirumala, Aram H Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. 2022b. Memorization without overfitting: Analyzing the training dynamics of large language models. *arXiv preprint arXiv:2205.10770*.
- Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- 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.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

A Implementation Details

As the benchmark is derived from The Pile (Gao et al., 2020) dataset, we conduct experiments only on the models that are pre-trained on The Pile

dataset. They are GPT-Neo 125M, GPT-Neo 1.3B, GPT-Neo 2.7B, and GPT-J 6B (Black et al., 2021; Wang and Komatsuzaki, 2021). We set the prompt length to 100, the batch size to 32, the learning rate of AdamW optimizer to $1e-3$, the warmup step to 500, the learning rate decay strategy to linear, N in Equation 2 to 5, α in Equation 3 to 0.7, and the maximum training epoch to 20 with an early stopping mechanism. In our main experiments, we generate the suffix using top-p sampling (Holtzman et al., 2020) with $p = 0.7$ and temperature = 0.8. For other decoding methods, we set beam size to 10 for beam search, and k to 10 for top-k sampling (temperature = 0.8). Our code is based on Huggingface Transformers (Wolf et al., 2020).

suggests the efficiency of ETHICIST.

B Computing Infrastructure

All experiments are carried out on a single Tesla V100 GPU with 32GB memory. Each experiment can be completed in less than 20 hours.




Figure 8: Effect of the sampling time. We show the **Recall** and the **Recall_{Early stop}** metrics for three methods when the sampling time increases from 1 to 100.

C Effect of Sampling Time

In our main experiments, we sample 100 candidate suffixes for one given prefix. We show the effect of sampling time in Figure 8. We can see that all methods' performances increase quickly when the sampling time increases from 1 to 10. However, ETHICIST's performance can still improve slowly when the sampling time increases from 10 to 100, which we attribute to the consideration of repeat time in our calibrated confidence estimation. What's more, although we report the result for sampling 100 times in our main experiments, we can see that ETHICIST can achieve satisfying performance when sampling only 10 times, which

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section Limitations
- A2. Did you discuss any potential risks of your work?
Section Ethics Statement
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section Abstract and Section Introduction
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section Methodology

- B1. Did you cite the creators of artifacts you used?
Section References
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
They are discussed in our github repo.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Section Experiments
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
The publisher of the data did the work.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
We don’t create new data.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section Experiments

C Did you run computational experiments?

Section Experiments

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section Implementation Details and Section Computing Infrastructure in the Appendix

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section Implementation Details in the Appendix

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section Experiments

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section Implementation Details

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.