

TAROT: Task-Oriented Authorship Obfuscation Using Policy Optimization Methods

Gabriel Loiseau^{1,2} Damien Sileo² Damien Riquet¹ Maxime Meyer¹ Marc Tommasi²

¹Hornetsecurity, Hem, France

²Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 - CRIStAL, F-59000 Lille, France
gabriel.loiseau@inria.fr

Abstract

Authorship obfuscation aims to disguise the identity of an author within a text by altering the writing style, vocabulary, syntax, and other linguistic features associated with the text author. This alteration needs to balance privacy and utility. While strong obfuscation techniques can effectively hide the author’s identity, they often degrade the quality and usefulness of the text for its intended purpose. Conversely, maintaining high utility tends to provide insufficient privacy, making it easier for an adversary to de-anonymize the author. Thus, achieving an optimal trade-off between these two conflicting objectives is crucial. In this paper, we propose TAROT: Task-Oriented Authorship Obfuscation Using Policy Optimization, a new unsupervised authorship obfuscation method whose goal is to optimize the privacy-utility trade-off by regenerating the entire text considering its downstream utility. Our approach leverages policy optimization as a fine-tuning paradigm over small language models in order to rewrite texts by preserving author identity and downstream task utility. We show that our approach largely reduces the accuracy of attackers while preserving utility. We make our code and models publicly available.¹

1 Introduction

Text is a primary medium for storing user data, training machine learning models, and interacting with large language models (LLMs) during inference. However, it also poses significant privacy risks, as sensitive or personal information contained within text can be exposed or misused. Text anonymization is a vital technique to address these concerns by removing or obfuscating personal information. This process protects individual privacy while ensuring that machine learning models can still derive meaningful insights and patterns from anonymized data, preserving its utility.

¹<https://github.com/hornetsecurity/tarot>

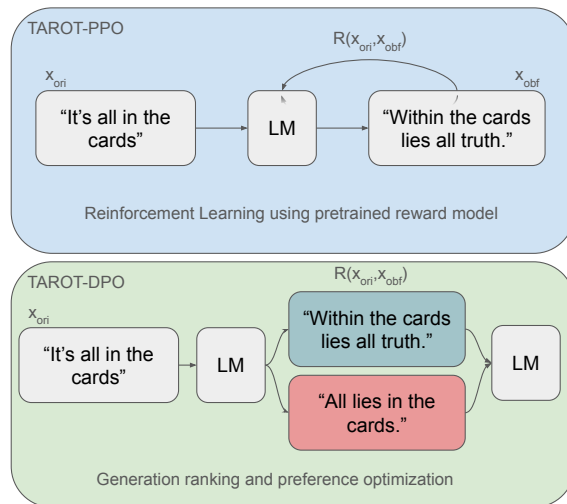


Figure 1: Illustration of the two versions of TAROT: We generate obfuscation candidates and optimize the best policy using reinforcement learning and preference optimization.

Currently, most work done on text anonymization focuses on redacting sensitive entities in a given document (Lison et al., 2021). This is sufficient for texts where the only private aspects are named entities, such as medical reports, court cases, or biographies. But it is inadequate for removing the author’s writing style, or the weak signals that can be used as hints for identification, which is, for example, the case for blog articles or emails. Redacting entities in text while keeping stylometric features linked to a specific individual would eventually result in a leak of information. Indeed, the writing style is a strong indicator of a person’s identity (Mosteller and Wallace, 1963). Previous work on authorship attribution highlights the large amount of information that can be extracted from seemingly anonymized texts and the ease of identification of authors, especially for long documents (Fabien et al., 2020).

To solve this issue, authorship obfuscation (AO) aims to hide the author’s identity by replacing some part of the text associated with authorship indicators. Modifying the original text can impact its usability

for specific tasks (i.e. utility), and therefore badly affects the downstream performances and text comprehension of machine learning models. The enforcement of privacy creates a trade-off between privacy and utility, where keeping the original text preserves the unchanged utility of the text, while not defending against attribution attacks. On the other hand, obfuscating the entire text guarantees privacy, but leads to unusable text in practice. Previous approaches design their obfuscation by maximizing the preserved text content. They limit the modifications to small and targeted edits in order to preserve text meaning and keep textual content as close as possible to the original. While this strategy is necessary to maintain the exact content and ensure that we convey the exact same message (before publishing the text online for example), those approaches often lead to insufficient modification in the text, especially against realistic attack scenarios (Zhai et al., 2022).

To address these limitations, we reframe the AO problem into an adversarial problem between two adversaries (e.g. machine learning models): one attacker model whose goal is to reveal the identity of a given author from written texts, and one utility model that aims to perform a given task using authors' data. The goal is to provide a modified version of the original text such that the utility model can accurately perform its task while preventing the attacker from identifying the author, making the obfuscation task-oriented. This perspective is more angled towards data users who need to privately perform utility tasks on the data, where some degree of content alteration may be acceptable if it enhances privacy. The notion of task-oriented obfuscation/anonymization also takes its origin in the law. As stated by GDPR (European Parliament and Council of the European Union, 2016), the collection and processing of personal information (including written texts) must be specified for a given usage.

In order to learn this privacy-utility trade-off, we use the combination of supervised fine-tuning (SFT) and policy optimization (PO) to guide a generative model into generating privacy- and utility-preserving outputs. Our model learns to rewrite the text while removing potential authorship signals, and preserving the text utility for a downstream task. This rewriting goal is further validated by the conclusion of Weitzenboeck et al. (2022) which showed how difficult it is to comply with GDPR requirements concerning text anonymization without changing the entire text.

We fine-tune a text simplification model for AO using a customized reward model. We design an unsupervised reward model for PO using two pretrained sentence embedding models. The utility reward penalizes the fact that the General Text Embeddings (Li et al., 2023) of the anonymized sentence is too far removed from that of the original sentence. The author rewards does the opposite on the embedding built by the Universal Authorship Representation model from Rivera-Soto et al. (2021). Our final models are trained in an open-world setting where the number of authors is not defined, the same goes for the end utility for our model to work on a multi-task setting. We also provide experimentation on three different datasets, movie reviews, blog articles and scholar documents. We show that TAROT can be used on multiple datasets targeting different tasks while protecting authorship.

In summary, we list the main contributions as follows:

- We design a new framework for task-oriented AO by leveraging PO algorithms to maximize the end usage of data. The objective is to help reduce the traditional constraints associated with utility preservation in the literature (strict content preservation and semantic quality) by looking for a downstream classification task to achieve with the anonymized data.
- Starting from this framework, we propose TAROT, a task-oriented generation model aiming to obfuscate text without any prior knowledge of the author (making it unsupervised, and usable on any dataset, even if the authors are not clearly indicated) while maximizing the utility for a variety of tasks. We release two versions of TAROT from two different fine-tuning PO algorithms: TAROT-PPO and TAROT-DPO.
- We further evaluate TAROT on three datasets associated with different classification tasks, using different authorship attackers and downstream usage scenarios.

2 Related Work

Authorship Obfuscation Obfuscation techniques can be regrouped into two categories, depending on their implementation. Generic methods, on one hand, are methods that were not explicitly designed for AO, but show interesting

performance. These methods include machine translation (Altakrori et al., 2022; Keswani et al., 2016), paraphrasing (Krishna et al., 2023), or synonym replacements (Potthast et al., 2016).

More recently, advanced techniques were built explicitly for AO, often relying on a trained attacker performing authorship attribution attacks on the obfuscated text. Then, they perform accurate adversarial text edits from the attacker knowledge on authors in order to obtain a privatized output. Mutant-X (Mahmood et al., 2019), is a genetic algorithm that utilizes GloVe (Pennington et al., 2014) word embeddings selected from an SVM or Random Forest attacker to replace words in a document with similar ones.

Jamdec (Fisher et al., 2024) is an unsupervised approach for obfuscating the writing style of text while preserving semantics. It uses embedding-based and likelihood-based methods, rather than attacker-based methods, to extract keywords, then generates multiple text variations using Constrained Diverse Beam Search on GPT2-XL (1.61B parameters). Finally, the candidates are filtered using Natural Language Inference (NLI) and Corpus of Linguistic Acceptability (CoLA) metrics to ensure coherence, content preservation, and grammatical correctness.

Recently, ALISON (Xing et al., 2024) employs a lightweight multilayer perceptron classifier using part-of-speech sequences to guide obfuscation, and leverages a BERT pre-trained language model to generate replacement sequences. By ranking and replacing important part-of-speech n-grams, ALISON obfuscates text uniformly, reducing classifier confidence.

Related studies share a common approach to evaluating privacy: they measure it through the performance of authorship attribution classifiers against obfuscated texts. Zhai et al. (2022) push forward this evaluation framework by introducing adversarial attackers that can resist obfuscation techniques. For measuring utility, the standard is to treat AO as a reference-less natural language generation problem, and to rely on standard metrics used for similar tasks such as machine translation and summarization (Altakrori et al., 2022).

Reinforcement Learning In NLP, reinforcement learning (RL) is often used to capture small signals over word or sentence embedding. For example, Mosallanezhad et al. (2019) proposes a text representation anonymization approach that employs deep reinforcement learning to detect

and modify text embeddings to maintain a good privacy-utility trade-off.

With the development of Reinforcement Learning from Human Feedback (RLHF) as a LLM fine-tuning paradigm, RL techniques have been leveraged to improve language models with scalar metrics by optimizing rewards from (human) feedback. It has emerged as a prominent tool for tackling undesirable behaviors such as toxicity, social biases, and offensive language (Ouyang et al., 2022). This is accomplished by implementing PO algorithms to optimize a language model (LM) by associating a reward with each generation, derived from a trained reward model.

Very recently, Liu et al. (2024) introduced an authorship style transfer method using PO. They optimize style transfer generation using style similarity reward models. Authorship style transfer is similar to AO in the way those task’s goal is to change within a text the author writing style. However, style transfer assumes a distinct target style to achieve, whereas AO assumes a lack of distinct style. Fisher et al. (2024) also showed the ineffectiveness of style transfer for AO. To the best of our knowledge, our work is the first one applying PO algorithms on AO.

Private Synthetic Text Generation Our work lies at the frontier between private text editing and synthetic text generation. Creating private synthetic data often relies on established frameworks such as differential privacy (Dwork, 2006). In contrast to these approaches, we focus on the implementation of a single text-to-text transformation specifically designed for authorship obfuscation, rather than on the generation of new textual data derived from potentially multiple sources (Mattern et al., 2022a).

Differential privacy traditionally targets noise addition in documents to produce useful and private text representations (Feyisetan et al., 2019; Fernandes et al., 2019). Applying differential privacy to document rewriting primarily serves to mitigate membership inference attacks, addressing a distinct threat model compared to the authorship attribution attacks targeted by our approach. While these techniques exhibit emergent capabilities for masking authorship signals (Igamberdiev and Habernal, 2023; Weggenmann et al., 2022; Utpala et al., 2023), they typically do so at a substantial cost to text utility, both at the task-level and the syntactic-level (Mattern et al., 2022b). This approach introduces unnecessary noise to semantic content not relevant to authorship identification, often degrading the

overall coherence and readability of the text. In contrast, our obfuscation methodology implements targeted modifications to stylometric features while maintaining the overall integrity of the source text.

3 Methodology

3.1 Problem Formulation

Let x_{ori} represent the original document authored by a specific author $a \in \mathcal{A}$. \mathcal{A} denoting a predetermined set of authors. The objective of authorship obfuscation is to generate a new document, denoted as x_{obf} , which cannot be attributed to the original author a . To assess the effectiveness of obfuscation, we employ a classification model, denoted as $f_{\text{attr}}(\cdot)$ (i.e. an authorship attribution model), which has been trained to distinguish documents based on their respective authors within \mathcal{A} . The goal of authorship obfuscation is to design an obfuscation method $O(\cdot)$, such that $f_{\text{attr}}(O(x_{\text{ori}})) \neq f_{\text{attr}}(x_{\text{ori}})$.

In addition, a successful obfuscation algorithm would not only trick an attacker into predicting the wrong author, but also preserve the document utility for downstream usage. In this paper, instead of mainly measuring this utility change through various semantic or content preservation metrics (i.e. METEOR score, BERT score, etc.) we highlight the selection of a prior task \mathcal{T} in order to evaluate obfuscation with respect to \mathcal{T} . We denote as $f_{\mathcal{T}}(\cdot)$ the classification model used for a utility task. An ideal $O(\cdot)$ would preserve the original label $f_{\mathcal{T}}(O(x_{\text{ori}})) = f_{\mathcal{T}}(x_{\text{ori}})$.

Note that \mathcal{T} is likely not known when we train the obfuscation model, underscoring the necessity for a versatile obfuscation strategy. This task-agnostic approach prevents the obfuscation model from learning to transform the text specifically to fit the label of \mathcal{T} , which would compromise its generality across different tasks.

3.2 Framework Overview

Our task-oriented framework can be decomposed in two steps. First, we initialize our generation model from a SFT baseline, this will first guide our LM to generate modified versions of the input text instead of proceeding text copy. Second, we apply a PO algorithm to fine-tune our SFT model. We experiment with two different PO algorithms, Proximal Policy Optimization (Schulman et al., 2017) and Direct Preference Optimization (Rafailov et al., 2023) (see Figure 1). We optimize our SFT generations using a reward model composed of

both privacy and content preservation components.

3.3 SFT Initialization

First, we use a fine-tuned LM to initiate our text generation task. We employ the *Keep It Simple*² simplification model (Laban et al., 2021) as an SFT baseline. This model is a fine-tuned version of GPT2-medium on the Newsela³ dataset for text simplification. The utilization of a simplification model encourages a reduction in the amount of information conveyed by a sentence, thereby affording the opportunity to eliminate author-specific features⁴. To our knowledge, this is the first time that a simplification model has been used for AO. Moreover, our framework is broadly compatible with any autoregressive LM, and can be adapted with larger architectures and other generation tasks.

3.4 Policy Optimization Algorithms

We use two different PO algorithms to optimize generations of our SFT baseline. The Proximal Policy Optimization (PPO) (Schulman et al., 2017) algorithm is a policy gradient method whose goal is to optimize a policy with respect to continuous rewards. In our case, a policy is a generation strategy, i.e. a final LM. Initialized from the SFT policy, we sample completions y given prompts x and the reward model parametrized by ϕ produces a score $r_{\phi}(x, y)$ based on these completions. The reward score $r_{\phi}(x, y)$ is then combined with a Kullback–Leibler (KL) penalty to ensure the policy does not deviate too much from the SFT policy (leading to unusable generations). Specifically, the reward of the RL problem is:

$$R(x, y) = r_{\phi}(x, y) - \beta \mathbb{D}_{\text{KL}}[\pi_{\theta}(y|x) || \pi_{\text{SFT}}(y|x)]$$

where β is a parameter controlling the strength of the KL penalty, θ the parameters of RL policy π_{θ} , and r_{ϕ} the reward model with parameters ϕ . Then, PPO is used to maximize the following objective:

$$\max_{\pi_{\theta}} \mathbb{E}_{x \sim \mathcal{D}_{\text{SFT}}, y \sim \pi_{\theta}(y|x)} R(x, y)$$

where \mathcal{D}_{SFT} is the prompts in the SFT dataset.

Rafailov et al. (2023) later introduced the Direct Preference Optimization (DPO) algorithm, which

²https://hf.co/philippelaban/keep_it_simple

³<https://newsela.com/>

⁴Our preliminary experiments revealed that using a simplification model outperformed comparable models of similar size for copy, paraphrasing, back-translation, and summarization, delivering superior privacy and utility.

| Dataset | Authors | Texts | Avg. Texts / Author (std) | Avg. Words / Text (std) | Avg. Tokens / Text (std) | Avg. Chars / Text (std) |
|---------|---------|-------|------------------------------|----------------------------|-----------------------------|----------------------------|
| IMDb | 10 | 10000 | 1000(± 0) | 364(± 209) | 393(± 228) | 1869(± 1077) |
| | 20 | 20000 | 1000(± 0) | 345(± 209) | 371(± 225) | 1767(± 1081) |
| BAC | 10 | 23534 | 2353(± 639) | 118(± 195) | 120(± 236) | 524(± 1027) |
| | 20 | 39379 | 1969(± 599) | 118(± 175) | 123(± 214) | 529(± 921) |
| AMT | 10 | 196 | 20(± 2) | 497(± 14) | 592(± 41) | 2956(± 194) |
| | 20 | 362 | 18(± 2) | 502(± 102) | 590(± 38) | 2956(± 207) |

Table 1: Dataset statistics

implicitly optimizes the same objective as PPO. DPO directly optimizes the model by a straightforward contrastive loss, boosting the reward of the preferred generation y_c and penalizing the one of the non-preferred generation y_r from a prompt x . DPO is a RL-free approach which has the following loss:

$$-\log \sigma \left(\beta \log \frac{\pi_{\theta}(y_c | x)}{\pi_{\text{SFT}}(y_c | x)} - \beta \log \frac{\pi_{\theta}(y_r | x)}{\pi_{\text{SFT}}(y_r | x)} \right)$$

where σ is the sigmoid function, and β the scaling parameter. In this study, we lack access to a preference dataset for DPO fine-tuning. Consequently, following the methodology of Rafailov et al. (2023), we generate this dataset by sampling responses from the same SFT dataset, and we rank those preferences using the same reward model (see Appendix A.3). This is justified as it is not possible to obtain a preference dataset from human feedback in the AO setting.

4 Experimental Setup

In this section, we describe the datasets involved for training and evaluation of our resulting models, and present our custom reward targeting the open-world authorship verification and multi-task text embeddings to learn this AO task. We then evaluate the resulting obfuscation against text edition and rewriting baselines.

4.1 Datasets

Training We use a separate dataset to train our PO models. We fine-tune our base simplification model on the Yelp reviews dataset⁵ (Zhang et al., 2015) composed of reviews from Yelp. The dataset is extracted from the Yelp Dataset Challenge 2015. This dataset is employed in an unsupervised way, to ensure we train our models on a large number of authors.

⁵https://hf.co/datasets/yelp_review_full

Evaluation To evaluate our obfuscation models, we use three different datasets. (i) IMDb62⁶, is a subset of the IMDb Authorship Attribution dataset initially presented by Seroussi et al. (2014). It consists of 62 authors with 1,000 texts per author taken from IMDb movie reviews. The utility task associated with this dataset is the review sentiment. For this, we map the movie rating between 0 and 10 associated with each review to a sentiment between *positive* and *negative*. A positive review occurs when the review rating is strictly larger than 5. (ii) The Blog Authorship Corpus⁷ dataset (Schler et al., 2006) consists of aggregated blog posts from 19,320 bloggers gathered from blogger.com. We pick the list of 13 topics present in the dataset as the utility task. (iii) The Extended-Brennan-Greenstadt⁸ dataset (Brennan et al., 2012) is composed of short paragraphs about scholar subjects gathered from 42 different authors from Amazon Mechanical Turk. The utility task of this dataset is indicated by the “background” column, as a binary classification problem.

For all datasets, we create two subsets containing the texts from 10 and 20 authors. For the Blog Authorship Corpus, we select the authors with the highest number of texts. We select the 10 (resp. 20) first authors listed in IMDb62 and Extended-Brennan-Greenstadt. We report summary statistics of each dataset in Table 1 and refer to every dataset as IMDb, BAC, and AMT followed by the number of considered authors. In summary, IMDb has rather long texts, numerous texts per author with a large associated standard deviation. BAC texts are shorter, with a higher number of texts per author compared to IMDb. Finally, for the AMT dataset, the texts are the longest with few variations, and the number of texts per author is the smallest.

⁶<https://hf.co/datasets/tasksource/imdb62>

⁷<https://u.cs.biu.ac.il/~koppel/BlogCorpus.htm>

⁸<https://hf.co/datasets/tasksource/Drexel-AMT>

4.2 Reward Models

To perform PO, we build a reward model from two different rewards components targeting respectively text semantics and text authorship, aiming to disentangle privacy and utility to control the trade-off.

For utility, we use a pretrained General Text Embeddings (GTE) (Li et al., 2023) to represent the reward as a cosine similarity between GTE before and after obfuscation⁹. Denote as $GTE(x)$ the embedding vector of size 1024, our utility reward is defined as:

$$R_{util} = \text{cossim}(GTE(x_{ori}), GTE(x_{obf}))$$

For the privacy reward, we use the Learning Universal Authorship Representations model (LUAR), from Rivera-Soto et al. (2021). LUAR’s goal is to transform a given text into a 512 dimensions embedding, such that representations of texts by the same author are closer, according to cosine similarity, than those by other authors.

Denote as $LUAR(x)$ the embedding vector given by the LUAR model, our privacy reward is defined as:

$$R_{priv} = 1 - \text{cossim}(LUAR(x_{ori}), LUAR(x_{obf}))$$

where cossim denotes the cosine similarity.

We obtain our final reward by summing the two previous rewards $R = R_{util} + R_{priv}$. All implementation details are listed in Appendix A.1.

4.3 Evaluation

Privacy Metrics The goal for obfuscation is to change the text in order to reduce as much as possible the attacker accuracy. We employ authorship attribution as an evaluation attacker to simulate an attack scenario when the attacker has already access to some sample data of targeted authors to train an attacker classifier. This is a stronger scenario than directly using the reward model as evaluation, since it only assumes one-to-one comparison between texts. For each evaluation dataset, we train a DeBERTa-v3 (He et al., 2021) model as an authorship attribution classifier. We split each evaluation dataset in 80%, 10% 10% for training, validation and testing. We measure the accuracy of the attacker model on each test set.

⁹We use the `gte-large-en-v1.5` from `sentence-transformers` <https://hf.co/Alibaba-NLP/gte-large-en-v1.5>

Utility Metrics We evaluate the utility loss when performing obfuscation similarly to the privacy classifier. For each downstream task dataset, we train a DeBERTa model to quantify utility preservation after text obfuscation. In addition, we also measure the impact on content preservation and soundness (see Appendix B).

Baselines We use the following baselines:

Original Text We measure the performance of utility / privacy classifiers when evaluated on original data, the goal of AO would be to decrease the performance of privacy classifiers without decreasing too much the accuracy of utility classifiers.

Synonyms As a baseline, we perform a naive text edition using synonyms. We use `GPTZzss`¹⁰ to process original texts, it employs a dictionary of synonyms to replace a given proportion of words with their counterparts. The goal of this baseline is to evaluate the attacker behavior when very small edits are made in the original text.

ALISON We use ALISON, a recent state-of-the-art text edition AO model leveraging small replacements using a pretrained BERT model. Replacements spans are computed using a threshold on the explanations of an adversarial authorship attribution classifier trained on each evaluation dataset. We train this classifier on each training and validation set before evaluation.

GPT-3.5 Lastly, we include a comparison with GPT3.5 (`gpt-3.5-turbo`) (Ouyang et al., 2022) as a text generation baseline. We use a simple text obfuscation prompt to capture zero-shot capabilities of GPT-3.5 to perform AO. The prompt used can be found in Appendix A.2.

4.4 Realistic Attackers

So far, the considered privacy evaluation adopts the perspective of an attacker who does not have any knowledge about the obfuscation algorithm used. In a more realistic setting, the attacker can likely identify and reproduce the AO model, and perform more advanced attacks by creating adversarial threat models. Following Zhai et al. (2022), we also evaluate our obfuscation models against two enhanced authorship attribution attackers, better suited to simulate real-world attack scenarios. We list the different attackers and their specific aspects based on adversarial training:

¹⁰<https://github.com/Declipsonator/GPTZzss>

| Method | IMDb | | | | BAC | | | | AMT | | | |
|-----------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | 10 Authors | | 20 Authors | | 10 Authors | | 20 Authors | | 10 Authors | | 20 Authors | |
| | Util. ↑ | Attr. ↓ | Util. ↑ | Attr. ↓ | Util. ↑ | Attr. ↓ | Util. ↑ | Attr. ↓ | Util. ↑ | Attr. ↓ | Util. ↑ | Attr. ↓ |
| Original | 73.51 | 99.78 | 79.46 | 99.80 | 46.73 | 61.05 | 53.80 | 61.14 | 100 | 70.37 | 86.11 | 42.86 |
| Synonyms | 70.38 | 94.52 | 76.60 | 96.08 | 46.24 | 59.06 | 51.20 | 58.18 | 91.67 | 64.81 | 86.11 | 36.90 |
| ALISON | 61.88 | 89.59 | 65.72 | 91.02 | 40.70 | 40.67 | 41.00 | 39.22 | 91.67 | 70.37 | 73.33 | 35.84 |
| GPT-3.5 | 63.33 | 66.67 | 47.37 | 35.00 | 37.20 | 42.73 | 44.74 | 31.27 | 60.00 | 44.44 | 61.11 | 31.14 |
| SFT | 64.51 | 62.50 | 39.47 | 80.00 | 40.41 | 32.44 | 40.10 | 28.28 | 90.00 | 26.85 | 75.00 | 21.23 |
| TAROT-PPO | 63.54 | 88.89 | 47.37 | 71.67 | 35.38 | 29.14 | 42.30 | 33.62 | 90.00 | 35.19 | 72.22 | 17.86 |
| TAROT-DPO | 57.14 | 34.74 | 60.72 | 17.34 | 24.57 | 23.97 | 28.39 | 16.42 | 86.67 | 22.22 | 64.18 | 16.67 |

Table 2: Evaluation results (Util: classifier accuracy on utility labels, Attr: authorship attribution accuracy) Best values are **bolded**.

- *Mix of original and obfuscated texts:* The attacker knows which AO algorithm was used and leverages this knowledge to create a new attribution model. This model is trained on a combination of original source texts and obfuscated texts generated by the known AO algorithm. We use a 50/50 distribution between original and obfuscated data to train this attacker.
- *Only obfuscated texts:* While the attacker is also aware of the AO algorithm, they train their authorship attribution classifier exclusively on the obfuscated samples. [Zhai et al. \(2022\)](#) demonstrated that this attack setting achieves the highest performance against text edition obfuscations.

For each attack scenario, we train a new authorship attribution classifier using the same parameters (see Appendix A.4 for hyperparameters) and compare the accuracy change from the original attacker.

4.5 Training new utility models with obfuscated texts

We experiment with a second use case to evaluate the downstream utility of obfuscated texts. We use the obfuscated texts of each method as a new training set for our utility classifier. This is useful to evaluate each method capability to generate useful training data that can be further used to train a new classifier on the same utility task.

5 Results

Downstream Effectiveness In Table 2, we present the accuracy change of privacy and utility classifiers. We observe that both SFT, PPO and DPO reduce the attacker accuracy compared to text edition methods (Synonyms and ALISON). PO

helps to learn a good privacy-utility trade-off by largely improving the privacy of obfuscated texts compared to baselines, while preserving similar utility. We observe that DPO consistently outperforms the PPO algorithm on privacy preservation, while using the same base reward model. DPO is also the best-performing privacy preservation over all baselines, with a notable drop of 82,46% on IMDB-20. Note that the utility decrease is larger for the BAC dataset, which could be explained by the number of short texts contained in the dataset, whose edits affect a lot more the end utility. TAROT-DPO also outperforms GPT-3.5 by providing more utility and less attribution on IMDB-20, AMT-10 and AMT-20. The effectiveness of TAROT-PPO lays in its utility preservation capabilities. While not being as private, the utility drop is reduced on nearly each dataset compared to TAROT-DPO.

Adversarial Attackers Figure 2 highlights the accuracy of adversarial threat models on the IMDb-10 dataset. This attack strategy is effective against text edition approaches (Synonyms and ALISON) as shown by the accuracy gain compared to the base attack only trained on original texts. However, text generation methods (GPT-3.5, SFT, TAROT-PPO and TAROT-DPO) show resistance to adversarial threat models, and only GPT-3.5 and TAROT-DPO are susceptible to the attacker trained on a mix of original and obfuscated texts. This encourages the path of generation methods as promising obfuscators. Note that this is the first obfuscation approach that is shown to be resistant to threat models.¹¹

Utility Preservation After Retraining Figure 3 presents the accuracy of a new utility classifier once trained with obfuscated texts. We observe

¹¹[Zhai et al. \(2022\)](#) did not include generation models in their study of AO evaluation.

| Method | Output |
|-----------|---|
| Original | I loved the whole story even though it was a tad corny at times . I think great acting and the content of the story kept it going. |
| Synonyms | I loved the quite whole story very even though it was a tad corny at times. I imagine too outstanding playing and the contents of the story kept it sledding. |
| ALISON | I thoroughly enjoyed the entire story even it did have a tad corny at times. I believe the great acting and the story’s content were the main reasons to keep it going. |
| GPT-3.5 | The entirety of the narrative was quite delightful, despite occasional moments of cheesiness. I believe the stellar performances and the substance of the storyline sustained its momentum. |
| SFT | I loved the whole story. It had many good parts and the writing was excellent. I think great acting and the subject matter of the story kept it going. |
| TAROT-PPO | I loved the whole thing. It was a good story and well-written. It also kept me going at times. I think great acting and the content of the story kept me going. |
| TAROT-DPO | I love the whole story. It’s full of action, personality and humour. It keeps me going, though, and the content keeps me going. |

Table 3: Obfuscation example from the IMDb dataset.

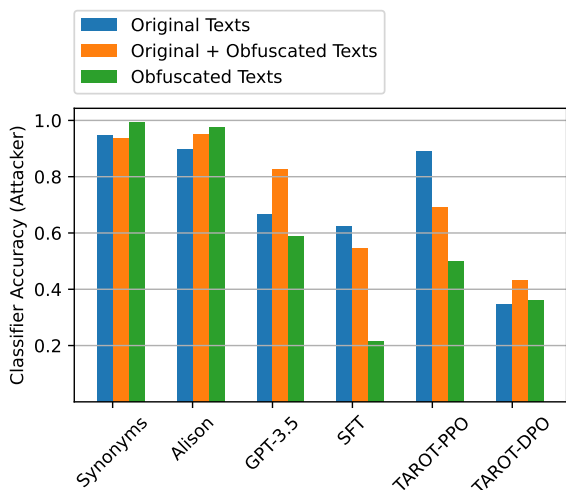


Figure 2: Authorship adversarial training accuracy results on IMDB-10 (lower is better). Generation models are resistant to adversarial training, compared to text edition methods.

that the drop in accuracy caused by obfuscation can be compensated by training a new classifier, with an accuracy increase for all methods. Moreover, generation methods are even better candidates for training data, as the final accuracy is higher than the original classifier accuracy. TAROT-PPO and TAROT-DPO are the best-performing approaches on this dataset. This highlights the possibility of creating obfuscation methods that are both preserving privacy and keeping utility for training purposes.

Qualitative Analysis We show an obfuscation example in Table 3 for each method. The base Synonyms obfuscation results in awkward phrasing and less natural language, compromising readabil-




Figure 3: Utility classifier accuracy once trained on IMDB-10 obfuscated texts (higher is better). The red line indicates the classifier accuracy when trained and evaluated on original data. The overall utility always increases after training on obfuscated texts, this is key to compensate the utility drop of generation methods.

ity. ALISON maintains coherence and clarity with slight formalization (“thoroughly enjoyed” instead of “loved”). GPT-3.5 significantly rephrases the text using sophisticated language. SFT simplifies and shortens the text, retaining clarity but reducing stylistic nuances. TAROT-PPO simplifies further, introducing some repetition, which makes the text less formal but still clear. TAROT-DPO alters the content more significantly, introducing new themes and repetition that can distract from the original meaning. The application of PO assists the text simplification SFT model in making additional modifications to the text. Although these changes in some cases alter the text’s meaning, they preserve its

overall utility. Appendix F provides more obfuscation examples from proposed and baseline methods.

Ablation Study As a complement, we perform an ablation study of each component of our reward model in Appendix D. It confirms the importance of using a combination of both privacy and utility rewards to learn this trade-off for obfuscation, especially for PPO.

6 Conclusion

We introduced a novel authorship obfuscation framework that focuses on optimizing the privacy-utility trade-off for a specific downstream data usage. We fine-tuned a text simplification model using two policy optimization algorithms to obfuscate the authorship of a given text, while preserving utility for multiple tasks. Our end-models are tuned using two sentence embedding rewards, one for content preservation and one for privacy, resulting in an unsupervised approach made for the open-world authorship setting. The results obtained help to improve the privacy from state-of-the-art AO methods, while preserving task utility. Our findings suggest that editing approaches are not suitable for privacy, especially against realistic attack settings. Additionally, we show that generated texts can be used to retrain utility classifiers and increase their performances, while limiting the accuracy of more advanced attackers. Ultimately, the performance of obfuscation methods largely varies depending on the downstream task choice, as does the resulting privacy-utility trade-off, highlighting the importance of selecting an appropriate model based on the specific requirements of the intended application. This calls for more research to design robust evaluation benchmarks for obfuscation systems, to assess and catch failure cases that can map to different real-world scenarios.

7 Limitations

The use of LM as text generators for obfuscation is not without risks, LM are known for their hallucination capabilities, so even if the downstream task is not affected, there is still a possibility that the trained LM generated plausible but false text from the original text. As we did not study the content preservation of resulting texts, we do not emphasize the risk of spread of misinformation or harm that can be generated by our fine-tuned LM.

Another limitation of our approach is that we rely on very small language models (380M parameters

for GPT2-medium, our SFT baseline), which benefits from limited memory usage but suffers from a restricted context size for generation. As a result, our method tends to reduce the text length, especially for longer texts. This limitation could be mitigated by increasing the size of the SFT model.

Finally, these methods can be limited when applied to short texts, as the replacements create significant changes that directly affect the utility task.

8 Ethical Considerations

In this work, we present authorship obfuscation methods that are intended for beneficial purposes (learning insights from data while preserving privacy). But we recognize that this task presents some risks of misuse. It can facilitate harmful activities such as posting misinformation, spam, or harmful content, without accountability because of obfuscation. Moreover, these techniques might infringe on intellectual property rights by obscuring the authorship of creative works, depriving creators of their deserved credit. We strongly encourage users to carefully consider these potential dangers before employing such methods.

References

- Malik Altakrori, Thomas Scialom, Benjamin C. M. Fung, and Jackie Chi Kit Cheung. 2022. [A multifaceted framework to evaluate evasion, content preservation, and misattribution in authorship obfuscation techniques](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2391–2406, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Michael Brennan, Sadia Afroz, and Rachel Greenstadt. 2012. [Adversarial stylometry: Circumventing authorship recognition to preserve privacy and anonymity](#). *ACM Trans. Inf. Syst. Secur.*, 15(3).
- Cynthia Dwork. 2006. Differential privacy. In *Automata, Languages and Programming*, pages 1–12, Berlin, Heidelberg. Springer Berlin Heidelberg.
- European Parliament and Council of the European Union. 2016. [General data protection regulation \(GDPR\)](#).
- Maël Fabien, Esau Villatoro-Tello, Petr Motlicek, and Shantipriya Parida. 2020. [BertAA : BERT fine-tuning for authorship attribution](#). In *Proceedings of the 17th International Conference on Natural Language Processing (ICON)*, pages 127–137, Indian Institute of Technology Patna, Patna, India. NLP Association of India (NLPAD).
- Natasha Fernandes, Mark Dras, and Annabelle McIver. 2019. [Generalised differential privacy for text](#)

- document processing. In *Principles of Security and Trust*, pages 123–148, Cham. Springer International Publishing.
- Oluwaseyi Feyisetan, Tom Diethe, and Thomas Drake. 2019. Leveraging hierarchical representations for preserving privacy and utility in text. In *2019 IEEE International Conference on Data Mining (ICDM)*, pages 210–219. IEEE.
- Jillian Fisher, Ximing Lu, Jaehun Jung, Liwei Jiang, Zaid Harchaoui, and Yejin Choi. 2024. Jamdec: Unsupervised authorship obfuscation using constrained decoding over small language models. *arXiv preprint arXiv:2402.08761*.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Timour Igamberdiev and Ivan Habernal. 2023. [DP-BART for privatized text rewriting under local differential privacy](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13914–13934, Toronto, Canada. Association for Computational Linguistics.
- Yashwant Keswani, Harsh Trivedi, Parth Mehta, and Prasenjit Majumder. 2016. [Author masking through translation](#). In *Working Notes of CLEF 2016 - Conference and Labs of the Evaluation forum, Évora, Portugal, 5-8 September, 2016*, volume 1609 of *CEUR Workshop Proceedings*, pages 890–894. CEUR-WS.org.
- Kalpesh Krishna, Yixiao Song, Marzena Karpinska, John Wieting, and Mohit Iyyer. 2023. Paraphrasing evades detectors of ai-generated text, but retrieval is an effective defense. *arXiv preprint arXiv:2303.13408*.
- Philippe Laban, Tobias Schnabel, Paul Bennett, and Marti A. Hearst. 2021. [Keep it simple: Unsupervised simplification of multi-paragraph text](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6365–6378, Online. Association for Computational Linguistics.
- Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. 2023. [Towards general text embeddings with multi-stage contrastive learning](#).
- Pierre Lison, Ildikó Pilán, David Sanchez, Montserrat Batet, and Lilja Øvrelid. 2021. [Anonymisation models for text data: State of the art, challenges and future directions](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4188–4203, Online. Association for Computational Linguistics.
- Shuai Liu, Shantanu Agarwal, and Jonathan May. 2024. [Authorship style transfer with policy optimization](#).
- Asad Mahmood, Faizan Ahmad, Zubair Shafiq, Padmini Srinivasan, and Fareed Zaffar. 2019. A girl has no name: Automated authorship obfuscation using mutant-x. *Proceedings on Privacy Enhancing Technologies*, 2019(4):54–71.
- Justus Mattern, Zhijing Jin, Benjamin Weggenmann, Bernhard Schoelkopf, and Mrinmaya Sachan. 2022a. [Differentially private language models for secure data sharing](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4860–4873, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Justus Mattern, Benjamin Weggenmann, and Florian Kerschbaum. 2022b. [The limits of word level differential privacy](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 867–881, Seattle, United States. Association for Computational Linguistics.
- Ahmadreza Mosallanezhad, Ghazaleh Beigi, and Huan Liu. 2019. [Deep reinforcement learning-based text anonymization against private-attribute inference](#). 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 2360–2369, Hong Kong, China. Association for Computational Linguistics.
- Frederick Mosteller and David L. Wallace. 1963. [Inference in an authorship problem](#). *Journal of the American Statistical Association*, 58(302):275–309.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#).
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Martin Potthast, Matthias Hagen, and Benno Stein. 2016. [Author obfuscation: Attacking the state of the art in authorship verification](#). In *Conference and Labs of the Evaluation Forum*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#).

- Rafael A. Rivera-Soto, Olivia Elizabeth Miano, Juanita Ordonez, Barry Y. Chen, Aleem Khan, Marcus Bishop, and Nicholas Andrews. 2021. [Learning universal authorship representations](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 913–919, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jonathan Schler, Moshe Koppel, Shlomo Engelson Argamon, and James W. Pennebaker. 2006. [Effects of age and gender on blogging](#). In *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#).
- Yanir Seroussi, Ingrid Zukerman, and Fabian Bohnert. 2014. Authorship attribution with topic models. *Computational Linguistics*, 40(2):269–310.
- Saiteja Utpala, Sara Hooker, and Pin-Yu Chen. 2023. [Locally differentially private document generation using zero shot prompting](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8442–8457, Singapore. Association for Computational Linguistics.
- Benjamin Weggenmann, Valentin Rublack, Michael Andrejczuk, Justus Mattern, and Florian Kerschbaum. 2022. [Dp-vae: Human-readable text anonymization for online reviews with differentially private variational autoencoders](#). In *Proceedings of the ACM Web Conference 2022, WWW '22*, page 721–731, New York, NY, USA. Association for Computing Machinery.
- Emily M Weitzenboeck, Pierre Lison, Malgorzata Cydecka, and Malcolm Langford. 2022. [The GDPR and unstructured data: is anonymization possible?](#) *International Data Privacy Law*, 12(3):184–206.
- Eric Xing, Saranya Venkatraman, Thai Le, and Dongwon Lee. 2024. [Alison: Fast and effective stylometric authorship obfuscation](#). In *AAAI*.
- Wanyue Zhai, Jonathan Ruser, Zubair Shafiq, and Padmini Srinivasan. 2022. [Adversarial authorship attribution for deobfuscation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7372–7384, Dublin, Ireland. Association for Computational Linguistics.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.

A Experimentation Details

A.1 Hardware and code

We conducted all experiments with Nvidia A30 GPU card with 24GB memory and Intel Xeon Gold 5320 CPU. The main libraries used include Pytorch 2.2.2, Huggingface transformers 4.39.3, datasets 2.19.0, tokenizers 0.15.2, trl 0.8.6, evaluate 0.4.1 and sentence-transformers 3.0.0. Due to memory constraints, models are loaded with float16 mixed precision.

Training time for PPO ranges from 15-20 hours, while time for DPO ranges from 6-12 hours. Evaluation time ranges approximately from 19-32 hours.

A.2 GPT-3.5 prompt

In our study, we compare with zero-shot prompting using GPT-3.5, a model with approximately 175 billion parameters. We obfuscate each text on a paragraph level, where the entire text is obfuscated as a unit. We use the following prompt to generate obfuscated texts: *"Rewrite the following paragraph so that the author's style is obfuscated."*

A.3 DPO training

While both PPO and DPO algorithms methods aim to optimize a model's performance based on a reward function, they differ in their approach to policy optimization. PPO uses a surrogate objective function that approximates the true objective function, while DPO directly optimizes the likelihood of generating a response chosen from a preference dataset over another response. This preference dataset is typically collected by having human annotators compare pairs of responses generated by a model and indicate which one is preferred. However, this protocol is impractical for authorship obfuscation because it is difficult to evaluate with human annotations. Therefore, we apply an initial preprocessing step to generate the preference dataset before DPO fine-tuning. We generate preference pairs from SFT outputs, and rank these preferences using the same reward model as PPO. Algorithm 1 outlines our method for creating this preference dataset for DPO. Preliminary experiments showed that removing samples with closely similar authorship rewards accelerates training convergence. So we specify filtering thresholds ϵ_{priv} and ϵ_{util} . After testing multiple values, we set $\epsilon_{priv} = 0.10$ and $\epsilon_{util} = 0.05$

Algorithm 1 Preference Dataset Generation

Require: SFT dataset \mathcal{D} , privacy threshold ϵ_{priv} , utility threshold ϵ_{util}

```
prompts = []
chosen = []
rejected = []
for prompt  $\in \mathcal{D}$  do
    left, right = generations from the SFT model
     $R_{util-left}, R_{priv-left}$  = privacy and utility
    rewards from the left obfuscation candidate
     $R_{util-right}, R_{priv-right}$  = privacy and utility
    rewards from the right obfuscation candidate
    if  $\|R_{priv-right} - R_{priv-left}\| > \epsilon_{priv}$  and
     $\|R_{util-right} - R_{util-left}\| < \epsilon_{util}$  then
        if  $R_{priv-right} > R_{priv-left}$  then
            prompt.append(prompt)
            chosen.append(right)
            reject.append(left)
        else
            prompt.append(prompt)
            chosen.append(left)
            reject.append(right)
return prompts, chosen, rejected
```

A.4 Hyperparameters

Table 4 and Table 5 present hyperparameters used for PO algorithms and evaluation classifiers. Due to limited time and computational resources, we are unable to conduct an exhaustive search across all hyperparameters. Instead, we report the best-performing hyperparameters we identified.

| | TAROT-PPO | TAROT-DPO |
|----------------|-----------|-----------|
| learning rate | 1.47e-5 | 2.96e-5 |
| batch size | 16 | 32 |
| # epochs | 3 | 3 |
| KL coef / beta | 0.2 | 0.1 |
| top p | 1.0 | 1.0 |
| temperature | 1.0 | 1.0 |

Table 4: Training hyperparameters for PO algorithms.

| | DeBERTa (Privacy) | DeBERTa (Utility) |
|---------------|-------------------|-------------------|
| learning rate | 2e-5 | 2e-5 |
| batch size | 8 | 8 |
| # epochs | 3 | 3 |

Table 5: Training hyperparameters for evaluation models.

A.5 Baseline implementation details

Synonyms We use GPTZzzs to process original texts, it employs a dictionary of synonyms to replace a given proportion of words with their counterparts. The goal of this baseline is to evaluate the attacker behavior when very small edits are made in the original text. We use the FinNLP synonym list and ask the algorithm to change up to 90% of words, and 80% of adjectives.

ALISON We use the author’s code implementation of ALISON, we use the largest edition parameters ($L=250$ and $c=1$) to edit the final text as much as possible.

GPT3.5 We use the gpt-3.5-turbo API endpoint from OpenAI to compute obfuscation, with default temperature, max_tokens and top_p.

B Content preservation and soundness study

We also study the impact on content preservation when obfuscating the text with generation models, including TAROT. Table 6 presents multiple content preservation metrics on the IMDB-10 dataset. Naturally, text edition methods obtain the best content preservation scores, compared to generation methods. In contrast, generation methods are superior in terms of linguistic acceptability (CoLA), since they generate the complete text as a whole. TAROT-DPO outperforms other methods on this metric.

C Complete Evaluation Results

Figure 4 presents the complete evaluation results of adversarial training on all datasets.

Figure 5 presents the complete utility evaluation after retraining on each dataset. The findings presented for IMDB-10 persist for IMDB-20 and AMT-20. We observe a smaller change in utility over the AMT-10 dataset due to the high base accuracy of the original classifier (1.0). However, this result does not hold for the BAC-10 and BAC-20 datasets, which is due to the lack of utility preserved after obfuscation. The blog authorship corpus dataset consists mainly of short texts, making it challenging for rewriting methods to transform the text without significantly affecting utility. This issue persists even after retraining the classifier on the obfuscated data.

| | Rouge-1 | Rouge-2 | Rouge-L | BLEU | METEOR | BERT Score | CoLA |
|-----------|---------|---------|---------|-------|--------|------------|-------|
| Original | - | - | - | - | - | - | 69.31 |
| Synonyms | 83.86 | 68.61 | 83.68 | 64.64 | 92.41 | 94.61 | 30.20 |
| ALISON | 98.24 | 97.08 | 98.19 | 67.48 | 97.61 | 99.01 | 43.88 |
| GPT-3.5 | 38.13 | 11.90 | 29.15 | 6.81 | 33.61 | 81.81 | 73.82 |
| SFT | 55.69 | 34.04 | 43.20 | 24.06 | 41.13 | 85.58 | 66.66 |
| TAROT-PPO | 51.33 | 29.36 | 38.67 | 20.77 | 37.93 | 84.50 | 74.46 |
| TAROT-DPO | 42.52 | 17.27 | 29.14 | 10.77 | 30.04 | 80.56 | 81.10 |

Table 6: Content preservation scores on the IMDB-10 dataset.

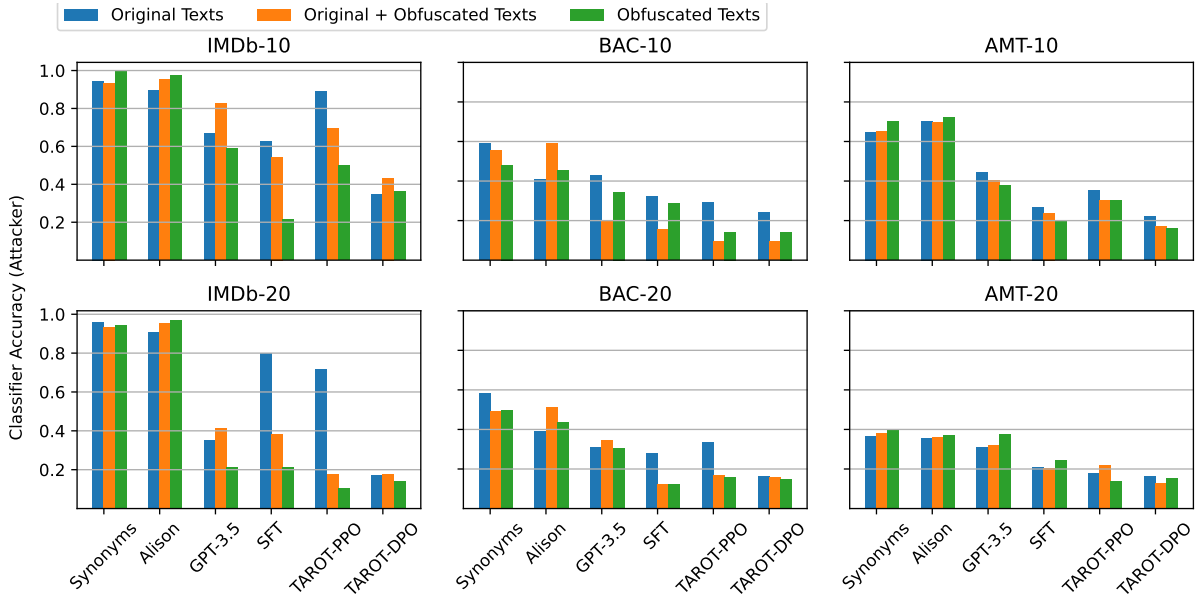


Figure 4: Adversarial training accuracy results (lower is better).

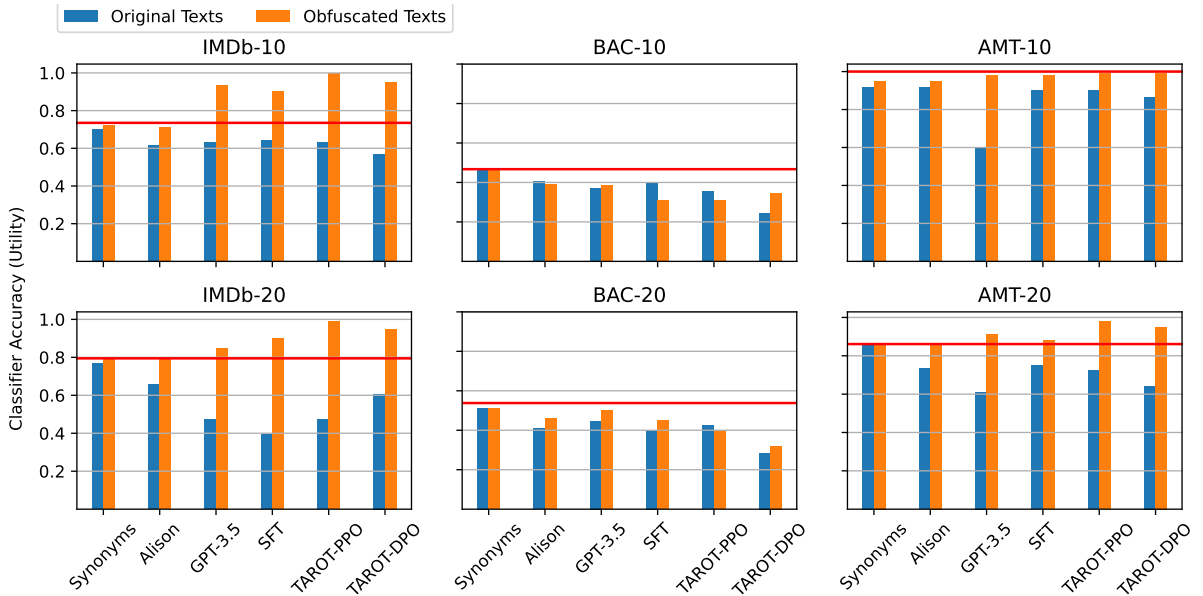


Figure 5: Utility classifier accuracy once trained on obfuscated texts (higher is better). The red line indicates the classifier accuracy when trained and evaluated on original data.

D Reward model ablation study

We perform a reward model ablation study to evaluate the importance of each reward component. Table 7 presents the reward value after training on different setups. We observe that the utility preservation and privacy components are both necessary to balance the privacy-utility trade-off. When we remove the LUAR-based reward, it leads to better GTE similarity at the expense of privacy. Similarly, removing the GTE reward leads to better privacy scores at the expense of utility. In practice, removing the privacy reward leads to models that try to copy the original text. While removing the utility reward leads to very short text, with only few words.

| Method | TAROT-PPO | | TAROT-DPO | |
|-------------|-----------|-------|-----------|-------|
| | LUAR | GTE | LUAR | GTE |
| No privacy | 0.975 | 0.993 | 0.983 | 0.977 |
| No utility | 0.403 | 0.421 | 0.706 | 0.633 |
| No ablation | 0.931 | 0.825 | 0.915 | 0.738 |

Table 7: Reward model values when removing one component. A high LUAR value indicates low privacy, and a high GTE value high utility.

E Scientific Artifacts

We list in this section the licenses used in this paper:

Models DeBERTa-v3 (MIT) Keep It Simple (apache-2.0) LUAR (apache-2.0) GTE (apache-2.0)

Software GPTZzzs (GPL-3.0) ALISON (MIT) GPT-3.5 (Terms of use¹²) Pytorch (BSD-3) Huggingface transformers, transformers, datasets, trl, evaluate and sentence-transformers (apache-2.0)

¹²<https://openai.com/policies>

F Additional Obfuscation Examples

| Dataset: IMDB | |
|---------------|--|
| Method | Output |
| Original | Walter Matthau was great as an aging man who lives to hang out in Central Park and tell outrageous lies. He strikes up a relationship with a man whom he drives nuts with his outlandish yarns. Matthau tries a con on several people trying to help them and himself out, but they all backfire . This film had a story that had me riveted to my chair. The entire production was filmed in Central Park, the scenery was beautiful, with many shots of the surrounding buildings. Wonderful movie - 5 stars. |
| Synonyms | Walter Matthau was really great as an aging man who lives to fall out in Central Park and evidence outrageous lies. He strikes up a relationship with a man whom he drives nuts with his outlandish yarns. Matthau tries a con on several people trying to assist them and himself out, but they all backfire. This film had a story that had me riveted to my chairman. The very full production was filmed in Central Park, the scenery was too beautiful, with extremely many shots of the surrounding buildings. Wonderful movie - 5 stars. |
| ALISON | Walter Matthau was great as an aging man to hang out in central park and tell outrageous lies. he strikes up a relationship with a man whom he drives nuts with his outlandish yarns. Matthau tries a con on several people trying to help them and himself out, but they all backfire. this film had a story that had me riveted to my chair. the entire production vet vet in central park, the scenery was beautiful, with many shots of the surrounding buildings. wonderful movie - 5 stars. |
| GPT-3.5 | In portraying a mature chap who thrives on loitering in the heart of Central Park, Walter Matthau delivered an exceptional performance filled with fanciful fables. His character forms a peculiar bond with an individual who becomes perturbed by his fantastical tales. Matthau embarks on a series of deceits aimed at aiding both himself and others, all of which result in unfortunate outcomes. This cinematic masterpiece captivated me, ensuring I remained glued to my seat throughout. The entirety of the picture was shot within the confines of Central Park, showcasing its exquisite landscapes and featuring numerous glimpses of the adjacent edifices. A delightful piece of cinema - a resounding five-star rating. |
| SFT | Walter Matthau was wonderful as an aging man who lives in Central Park. He's a regular at the bar where he tells outrageous stories and tells wild stories. He strikes up a good relationship with a man whom he drives nuts with: his yarns. This film is shot in Central Park, with the backdrop of a long-running feud between Matthau and his friends. Wonderful movie. |
| TAROT-PPO | Walter Matthau was a joy to watch as an aging man who lives out in Central Park. He regularly tells wild stories about how he used to drive around town and tell wild lies. But he strikes up a good relationship with one person who wants to help him out and himself. This film has a lot going for it, with lovely scenery and lots of action. Wonderful film - 5 stars. |
| TAROT-DPO | Walter Matthau is one of those people who can never quite find the right balance between being a good friend and being funny. Though he lives in Central Park, he spends a lot of time telling his friends outrageous stories. This time, though, the relationship is quite straightforward. He strikes up a few conversations with people trying to help them get ahead of the game and, at the same time, he strikes up a relationship with a man who drives him nuts. |
| Dataset: BAC | |
| Method | Output |
| Original | Is anyone else gonna actually blog? It would be helpful you know! |
| Synonyms | Is anyone else gonna really blog? It would be very helpful you know! |
| ALISON | Is anyone else gonna actually blog? It would be helpful you know |
| GPT-3.5 | Does the inclination exist for any additional individuals to engage in the act of blogging? Such contributions would undoubtedly prove advantageous, do you not concur? |
| SFT | Any other person writing about something like this? It would be helpful to know what other people are saying about the matter. |
| TAROT-PPO | Any other person blogging about something? it would be helpful to know who else is going to actually blog? |
| TAROT-DPO | Any other person on the list of potential bloggers? it would be helpful to know what they write about. |

| Dataset: AMT | |
|--------------|--|
| Method | Output |
| Original | The interplay between "new" and "old" ideas, methods, and forms in gothic literature gave it the intrinsic contradictory nature it retains today. The definition of "gothic" can never seemed to be agreed upon, the story within a gothic novel still seems to lost track of itself occasionally, the conflict between the emotional effect of terror, and the desire for a logical explanation (science and reason gradually began to replace the supernatural as the driving force behind the horrific) all remain. Neither the ancient nor the modern can be removed from the gothic without sacrificing something essential to the genre. [...] |
| Synonyms | The interplay between "new" and "old" ideas, methods, and forms in Gothic lit gave it the intrinsical contradictory nature it retains today. The definition of "Gothic" can never seemed to be agreed upon, the story within a Gothic refreshing relieve seems to missed cartroad of itself occasionally, the conflict between the emotional gist of terror, and the want for a logical account (science and intellect gradually began to replace the occult as the impulsive thrust behind the horrific) all remain. Neither the antediluvian nor the very modern can be removed from the Gothic without sacrificing something essential to the genre. |
| ALISON | The interplay between "new" and "old" ideas, methods, and forms in gothic literature gave it the intrinsic contradictory nature it retains today. The definition of "gothic" can never seemed to gothic gothic upon, the story within a gothic novel still seems to lost track of itself occasionally, the conflict between the emotional effect of terror, and the desire for a logical explanation (science and reason gradually began to replace the supernatural as the driving force behind the horrific) all remain. neither the ancient nor the modern can neither neither from the gothic without sacrificing something essential to the genre. |
| GPT-3.5 | The perplexing interplay betwixt the "novel" and the "antique" conceputations, methodologies, and configurations in Gothic penmanship hath imbued it with a most incongruous essence that endures unto this day. The elucidation of "Gothic" hath forever eluded consensus, the narrative ensconced within a Gothic tome oft doth meander astray, the discord betwixt the emotive impact of dread, and the craving for a reasoned explication (where science and rationale inexorably supplanted the otherworldly as the impelling impetus behind the macabre) doth persevereth. Ne'er can the hoary nor the contemporaneous be excised from the Gothic sans forfeiting somewhat indispensable to the genre. |
| SFT | The interplay between new and old ideas is what gives gothic fiction its unique character. Methods and forms of storytelling, such as novel characters and their relationship to nature, give it a contradictory nature. The definition of gothic, however, cannot be agreed upon, the story within a gothic novel still seems to lack logic and reason, with terror still guiding the characters' actions and choices (science and reason gradually replaced the supernatural as the driving force, not just emotions). |
| TAROT-PPO | the interplay between "new" and "old" ideas, methods, and forms in gothic literature gave it the intrinsic contradictory nature it still retains today. the definition of "gothic" is never clearly agreed upon, the story within a gothic novel still seems to lack a logical foundation, and the desire for a logical explanation often clashes with the character's emotions and mental state. both the ancient and the modern are incompatible with the genre. |
| TAROT-DPO | But the two ideas are not linked by a common underlying logic. The definition of "gothic" can never seem to agree with reality. It can only seem to be full of contradictions and inconsistencies, with ideas such as new and old, new and old, about which there is never a clear line between right and wrong. |
| Dataset: BAC | |
| Method | Output |
| Original | New F-Zero title GBA bound - F-Zero: Falcon Densetsu News for Game Boy Advance at GameSpot yeah out this week to but overseas based on the anime btw which unfortunately is bad |
| Synonyms | New F-Zero claim GBA bound - F-Zero: Falcon Densetsu News for Game Boy Advance at GameSpot yeah out this week to but abroad based on the anime btw which regrettably is bad |
| ALISON | New F-Zero title GBA bound - F-Zero: Falcon Densetsu news for game boy advance at GameSpot yeah out this week to but overseas based on the anime btw which unfortunately is bad |
| GPT-3.5 | A brand-new installment in the F-Zero franchise is set to debut on the Game Boy Advance, known as F-Zero: Falcon Densetsu. The release is anticipated this week, with availability limited to specific regions tied to its anime adaptation, which has been critically panned. |
| SFT | New title is GBA, a reference to falcon - a news service for game boy advance for the past three years. Yeah, out this week, though, overseas based on the anime btw which is just rubbish. |
| TAROT-PPO | New F-Zero title is a straight up rip-off of GBA. It features falcon news for the game boy advance, which again is crap online. Out this week, though, the anime btw are obviously not very good. |
| TAROT-DPO | An updated F-Zero title, this time with GBA, the title of a news show that the game boy advance on is. Yeah, out this week to news shows like F-Zero but overseas based on the anime, which is bad. |

| Dataset: AMT | |
|--------------|---|
| Method | Output |
| Original | Organisms would have the abilities to move, eat, hunt, and think. These functions would be optimized by genetic algorithms. To create this simulation, there were several steps. The first was to decide upon a programming language. The C++ programming language was chosen for its versatility and large pool of tutorial resources. The next step involved writing pseudo-code, or planning out the program itself. The actual program was broken into several functions. The first function, a global function that provided access to all other parts of the program, generated an arbitrary number of environments and filled them with a random number of food items that had a varying energy level based on the environment. |
| Synonyms | Organisms would get the abilities to go, eat, hunting, and think. These functions would be optimized by genetical algorithms. To make this simulation, thither were several steps. The 1st was to decide upon a programing speech. The C++ programing speech was elect for its versatility and really big kitty of tutorial resources. The next stair mired penning pseudo-code, or preparation out the programme itself. The factual program was crushed into several functions. The really first function, a worldwide function that provided approach to all other parts of the program, generated an arbitrary number of environments and filled them with a random figure of food items that had a variable vigor raze based on the envions. |
| ALISON | Organisms would have the abilities to move, eat, hunt, and think. these functions would be optimized by genetic algorithms. To create this simulation, there were several steps. The first was to decide upon a programming language. The C++ programming language for its versatility and large pool of tutorial resources. The next step involved writing pseudo-code, or planning out the program itself. The actual program planning planning into several functions. The first function, a global function that provided access to all other parts of the program, generated an arbitrary number of environments and filled them with a random number of food items that had a varying energy level based on the environment. |
| GPT-3.5 | Living beings posses a range of functionalities, such as mobility, sustenance consumption, predatory behaviors, and cognitive processes, that are fine-tuned through complex genetic algorithms. To initiate the simulation, several procedural stages were undertaken. Initially, a choice was made regarding the selection of a programming language. Eventually, the decision was made to utilize the versatile and well-documented C++ language for its wide range of resources. Subsequently, the process involved the formulation of pseudo-code, a strategic blueprint for the development of the program. The program itself was compartmentalized into multiple functions. The primary function, acting as a global facilitator, was designed to generate artificial environments and populate them with random assortments of food items possessing varying energy levels depending upon the environment. |
| SFT | These organisms would have the ability to move, eat, hunt and think. These functions would be better optimized by genetic algorithms. These algorithms would be able to predict what environment a given creature would encounter next, and so on, generating different environments for each task. To create this simulation, programming language was used. Programming languages are usually designed to allow flexible programming of tasks for different stages in a person's life, like diet and training. These environments are then tested out in a variety of ways, like hunting or training. Tucking these algorithms into the program also meant creating more environment specific. |
| TAROT-PPO | These organisms would have the ability to move, eat, hunt and think. these functions would be optimized by genetic algorithms. These functions would be triggered by training programs that specifically targeted specific traits, such as hunger and thirst. To create this simulation, these would include programming languages for different types of environments, such as languages for animals and plants. These training programs would be tested against real-world situations and filled with different food types. |
| TAROT-DPO | These organisms are capable of thinking and eating. They can move and hunt, and they have even been observed to think about their environment. These functions, which would normally be optimized by a computer, are now well-suited for creating life on the planet. To train these organisms to perform these functions, for example, is to have a programming language that can be easily adapted to new situations. To put together a simulation of a particular environment, it would need to provide for training a range of animals, from which to choose. These kinds of simulations, which are relatively simple, would be ideal for training mental processes. |

| Dataset: IMDB | |
|---------------|--|
| Method | Output |
| Original | Funny romance / comedy with some well written lines for everybody . I especially enjoyed the give and take between Tony Randall and the Rock . Doris Day is soooooo cute ; I had a major crush on her during my junior high years . I liked Lynde's character , but I prefer his more sarcastic roles . 3 stars . |
| Synonyms | Funny latin / comedy with some swell written lines for everybody . I especially enjoyed the dedicate and take 'tween Tony Randall and the Rock . Doris Day is soooooo very cute ; I had a major smash on her during my junior so high years . I liked Lynde's character , but I favour his more sarcastic roles . 3 stars |
| ALISON | Funny romance / comedy with some well written lines for everybody. I especially enjoyed the give and take between Tony Randall and the rock. Doris Day is soooooo cute; I had a major crush on her during my junior high years. I liked Lynde's character, but I prefer his more sarcastic roles. 3 stars |
| GPT-3.5 | A delightful tale of amusement and affection with a sprinkling of finely crafted dialogue that will appeal to all. The banter shared between the one called Tony Randall and the individual known as the Rock was particularly engaging. Doris Day exuded a charm that was utterly endearing; during my formative years, a deep infatuation for her blossomed. The portrayal by Lynde was rather enjoyable, though I find myself partial to his portrayals imbued with a sense of sarcasm. A rating of three stars is the verdict bestowed upon this film. |
| SFT | Funny romance / comedy with some funny lines. I loved the give-and-take between Tony Randall and the Rock, which was great to watch as I got older. Doris Day is soooooo cute ; she's always been this way, and I like her more in my junior high years. 3 stars. |
| TAROT-PPO | Funny romance / comedy with some well-written lines for everybody. I really enjoyed the give-and-take between Tony Randall and Doris Day. I loved her personality and her reaction to situations. They're both cute, but not as smart as Tony Randall. 3 stars. |
| TAROT-DPO | Funny romance. It's a good combination of comedy and give-and-take. There's always something interesting to say in front of the camera, or in front of the audience, about how they feel about someone else. Doris Day is soo cute, and soo funny. She's very much like me at my younger years, when I played sport. 3 stars." |

Table 8: Additional qualitative examples for each obfuscation method.