

DiDOTS: Knowledge Distillation from Large-Language-Models for Dementia Obfuscation in Transcribed Speech

Dominika Woszczyk
Imperial College London
London, UK
d.woszczyk19@imperial.ac.uk

Soteris Demetriou
Imperial College London
London, UK
s.demetriou@imperial.ac.uk

Abstract

Dementia is a sensitive neurocognitive disorder affecting tens of millions of people worldwide and its cases are expected to triple by 2050. Alarmingly, recent advancements in dementia classification make it possible for adversaries to violate affected individuals' privacy and infer their sensitive condition from speech transcriptions. Existing obfuscation methods in text have never been applied for dementia and depend on the availability of large labeled datasets which are challenging to collect for sensitive medical attributes. In this work, we bridge this research gap and tackle the above issues by leveraging Large-Language-Models (LLMs) with diverse prompt designs (zero-shot, few-shot, and knowledge-based) to obfuscate dementia in speech transcripts. Our evaluation shows that LLMs are more effective dementia obfuscators compared to competing methods. However, they have billions of parameters which renders them hard to train, store and share, and they are also fragile suffering from hallucination, refusal and contradiction effects among others. To further mitigate these, we propose a novel method, DiDOTS. DiDOTS distills knowledge from LLMs using a teacher-student paradigm and parameter-efficient fine-tuning. DiDOTS has one order of magnitude fewer parameters compared to its teacher LLM and can be fine-tuned using three orders of magnitude less parameters compared to full fine-tuning. Our evaluation shows that compared to prior work DiDOTS retains the performance of LLMs achieving 1.3x and 2.2x improvement in privacy performance on two datasets, while humans rate it as better in preserving utility even when compared to state-of-the-art paraphrasing models.

Keywords

dementia, attribute obfuscation, large-language-models, knowledge distillation

1 Introduction

Speech is a prevalent human-computer interaction modality and there are strong business incentives for collecting and processing speech data. The size of the Voice Recognition and Speech processing global market amounted to \$12 billion in 2022 and is expected to grow to \$34 billion by 2028. [69]. The voice assistant market alone is worth \$2.5 billion in 2023 and popular solutions such as Google Home and Amazon Alexa are being used by hundreds of millions of

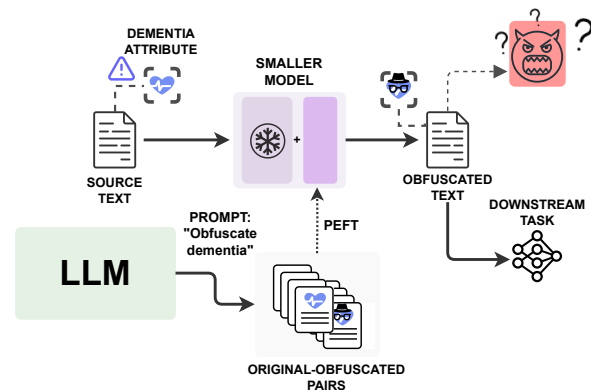


Figure 1: Overall structure of DiDOTS. DiDOTS takes as an input a sentence with a potential dementia attribute and generates an obfuscated paraphrase.

smart speaker users [98]. Finance and Telecom companies increasingly rely on speech-to-text APIs and speech analytics to process calls for feedback, categorize queries, and check for compliance with regulations. The healthcare sector is the next biggest use case for speech-to-text as it can speed up note-taking and search and also analyse prevalent topics of worry and patient feedback [103]. Moreover, in this work, we found (see Section 2) that about 20% of the most popular Android apps request permission to record audio out of which more than 79% use APIs to process speech. Transcribing speech allows one to use natural language processing techniques to understand the semantics of the natural language. For example, *Google Keep* [82] is a productivity app with more than 2 billion installations that allows users to take notes using their voice; *Voice Recorder Audio Sound MP3* [84] and *Otter: Transcribe Voice Notes* [83] have been installed more than 8 million and 1 million times respectively and produce live and offline voice transcriptions.

Dementia Can be Inferred From Speech Transcripts. Dementia, a degenerative neurological disorder, has become a pervasive global health concern with the ageing population. It affects individuals' cognitive abilities, memory, and overall functioning, and its cases are expected to triple by 2050 [71]. Disfluencies, mistakes, and signs of cognitive challenge are even more apparent in speech than they are in written text. Key discriminative features were found in speech transcripts and when aggregated, they can be important indicators of Alzheimer's Disease (AD). Recent works on the automatic detection of dementia, yielded promising results [101, 115, 117]. These, however, can also be exploited to extract sensitive

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.
Proceedings on Privacy Enhancing Technologies 2025(1), 198–215
© 2025 Copyright held by the owner/author(s).
<https://doi.org/10.56553/popets-2025-0012>



medical information from transcribed speech. Personal information related to medical conditions is very sensitive and must be kept secret: employers can use it to discriminate against employees or applicants, banks can use it to associate loan applicants with higher risk scores, advertisers can use it to target users, etc. While user attribute leakage has been studied before [24, 92, 94, 95], little has been done against leakage of medical-related attributes from transcribed speech, such as dementia.

This is worrisome and it might come in conflict with privacy regulations. For example, GDPR [27] is an EU regulation that protects the processing of sensitive data. It defines sensitive data as personal data subject to several processing conditions including biometric data and health-related data [14]. GDPR also has *data minimization* as one of its core data protection principles. According to data minimization, data processors must ensure that the data they process are adequate, relevant, and “*limited to what is necessary*”. Similarly, CPRA [15] (California Privacy Rights Act) which took effect in January 2023 and is an amended version of the CCPA (California Consumer Privacy Act) also stipulates that data processors should not process data beyond what “*reasonably necessary and proportionate to achieve the purposes for which the personal information was collected or processed*”. Therefore, there is a need for solutions that allow understanding of natural language from speech without allowing dementia inference.

Prior Approaches. Several works already exist aiming to either obfuscate sensitive attributes or identity [1, 24, 50, 63, 64, 70, 95, 108, 112]. Some follow a style-transfer approach to target gender, sentiment, age, and other attributes [24, 63, 95, 112]. These (a) are limited by their dependence on the availability of training data for the specific attributes and (b) they require training a model for each new attribute and their obfuscation is hence classifier dependent. Medical datasets of conversations and written text for conditions such as dementia are hard to collect and access. Unlike existing style transfer datasets, parallel corpora between patients and control individuals are not available. This limitation calls for methods that can learn to perform obfuscation even in the absence of training data for the target attribute. Others have used differential privacy which can provide provable guarantees [64, 108]. However, these suffer from poor semantic reconstruction. Furthermore, it is difficult to interpret the privacy metric’s (ϵ) value when applied to words in a sentence or a document. More importantly, none of the prior works have been applied to obfuscating sensitive medical attributes such as dementia.

Our Approach. Our first key idea is to *leverage the capabilities of large language models (LLMs) to perform zero-shot text generation and paraphrasing tasks*. We investigate several open-source LLMs that can be executed privately and design zero-shot, few-shot and knowledge-based instructions to significantly improve privacy performance compared to prior work. However, LLMs require costly memory and computational power due to their $O(10^9)$ number of trainable parameters, making them impractical for many applications, and they can suffer from hallucinations and refusals. Our second key idea to overcome the LLM limitations is to perform *teacher-student knowledge distillation from LLMs* in a 2-step process: first, we prompt a teacher LLM model to create a dataset of original and obfuscated sentence pairs. Second, we fine-tune a smaller, more

efficient student model on this dataset. Fully fine-tuning the smaller model though, still requires training on $O(10^8)$ parameters. To further improve on this we use our third key idea to *apply a suitable Parameter-Efficient Fine-Tuning (PEFT) technique*. This involves freezing the student’s layers, using matrix decomposition to reduce the number of parameters, and integrating lightweight adapter modules. This ensures efficient training, reduces overfitting risk on small datasets, and reduces computational overhead. Overall, our approach results in a smaller, efficient student model which we call DiDOTS (Knowledge Distillation from Large-Language-Models for Dementia Obfuscation in Transcribed Speech). DiDOTS retains the privacy and utility performance of the teacher LLM while requiring three orders of magnitude less parameters ($O(10^5)$) for fine-tuning. Compared to a state-of-the-art attribute obfuscator (ParChoice [36]), DiDOTS is not classifier dependent, achieves 1.3x and 2.2x better on average privacy performance across static and adaptive adversaries on two datasets, and a 3x improvement in a worst case scenario. Human evaluation shows that both our best LLM approach and DiDOTS achieved significantly better comparable utility performance than ParChoice and other baselines.

Contributions. Below we list our main contributions:

- *New application domain.* We are the first to explore dementia obfuscation in transcribed speech. We hope this inspires further research on cognitive disorder obfuscation in speech.
- *New Understanding.* We are the first to explore and rigorously evaluate the ability of LLMs to obfuscate dementia and preserve semantics. In doing so we have designed task-relevant zero-shot, few-shot and knowledge-based instructions for LLMs.
- *Novel Obfuscation Method.* To overcome LLMs’ computational cost and brittleness (refusal, hallucination, contradictions), we have designed a unique and lightweight system that distills knowledge from LLMs for dementia obfuscation. Additionally, our approach leverages parameter-efficient training for a cost-effective and practical solution.
- *Rigorous empirical evaluation and findings.* We perform a systematic and rigorous evaluation of LLMs and our knowledge distillation method across both utility and privacy metrics and through ablation and human studies. We find that LLMs are better obfuscators compared to other competing approaches but have a very large number of trainable parameters ($O(10^9)$), and sometimes fail to follow instructions. Our novel DiDOTS method shows that we can learn to obfuscate similarly to LLMs while we only need to partially fine-tune a smaller model ($O(10^9)$) on three orders of magnitude less parameters ($O(10^5)$) compared to full fine-tuning.

Ethical Considerations. In our work, we make use of available upon-request datasets ADReSS and ADReSSo from the TalkBank archive [62], containing recordings and transcripts of conversations with dementia and control patients. Datasets submitted to the TalkBank conform to IRB practices i.e. they are anonymized and do not contain identifiable information. Our human subject study is approved from our institution’s Research Ethics review board. Details of our study design are provided in Appendix D.

2 Understanding the Risks

2.1 Speech-to-Text in Mobile Apps

Existing market studies evidence the prevalence and importance of speech-to-text [26, 35, 69, 98]. To better understand and highlight the threats of profiling from transcribed speech we perform a more detailed analysis of speech and transcription data collection and processing on the most popular smartphone OS ecosystem. Specifically, we systematically quantify the prevalence of speech collection and processing, by analyzing popular mobile apps. We focus on mobile apps because they are a popular interface with digital services (more than 5 billion people own a smartphone [97]) and because there are available tools for collecting and analyzing mobile apps.

Our analysis focuses on the Android ecosystem because it has the largest (> 70% [99]) mobile OS market share. We collected the top 200 free Android apps among 54 application categories from GooglePlay [81], Android’s official application store, using `google-play-scraper`¹. This resulted in 10070 unique apps. For each app, we collected the Android package (.apk and application metadata such as the application’s description, its requested permissions, and the number of installations.

The first goal of our analysis is to study the prevalence of audio collection. Apps collecting audio could infer or leak sensitive information from speech either on purpose or inadvertently. The latter is possible if third-party libraries embedded in the app decide to take advantage of the legitimately granted permission to access audio and further process the audio to identify sensitive user attributes [34], or if the app is compromised [18, 30, 33, 105]. To measure how many of these apps have the capability of collecting audio data, we search their `AndroidManifest.xml` files for the presence of the `RECORD_AUDIO` permission. Developers use these files to provide important information (such as Android permissions) that the system requires to execute their apps. Alarming, we found that 20.3%/10070 apps request permission. These apps have 128M of installations on average with the most popular app having 15B installations.

Next, we wanted to further understand which of the apps which collect audio, specifically perform speech-to-text processing. To detect this, we wrote a script that looks for the presence of the system `SpeechRecognizer` [20] in the code of the apps. However, speech-to-text can be performed using third-party (3P) libraries as well. Therefore we further manually compile a list of 13 APIs based on market leaders [35] and popular services that we provided to our script. We supplement this list with the API’s and popular libraries’ SDK calls when available and keywords that search for functions that relate to speech-to-text (full list in Appendix A). Note, that our list is incomplete therefore our script will only provide a conservative estimate. After removing apps for which decompilation failed, we were left with 1743 apps. We found that 1380 apps (79.1%) are currently employing speech-to-text. 1306 (74.9%) use the Android’s built-in speech recognition system service `SpeechRecognizer`, 534 (30%) matched for keywords associated with speech-to-text functions and 82 (4.7%) use an API.

To identify if any apps process speech in violation of users’ expectations we investigated how many of the apps that use a speech-to-text API in their code do not explicitly mention it in their public descriptions [76]. We randomly picked 50 out of the 1380 apps and manually went through their descriptions looking for mentions of speech-to-text or transcription activity. Specifically, we look for direct or implied (e.g. speak and it will repeat after you) mentions of “speech-to-text” or synonymous tasks such as transcription, dictation, live translation, captioning, and voice-based control. Alarming, we found that 42 (84%) of these apps do not mention it. After further analysis, we found that 1278/1306 (98%) use the `SpeechRecognizer` to enable voice search through Android’s `SearchView` widget. Voice search allows the apps access to unfiltered transcriptions of users’ queries [19]. This shows that voice control is a prevalent but not adequately described functionality which is increasingly offered by popular apps. To examine whether violations of user expectations happen not only for voice-controlled but also for other speech-to-text processing tasks, we repeated the experiment purposively sampling 50 apps that do not only use `SpeechRecognizer` (616/1380), therefore speech processing is unlikely to be triggered explicitly by the user. We found that 34/50 (68%) do not refer to speech-to-text in their descriptions.

In summary, we found that (a) an important number of popular apps collect audio, (b) most of them transcribe speech, and (c) the transcription functionalities are rarely described in the apps’ public descriptions.

2.2 Preliminaries

Text Feature Extraction. To process text with algorithms and models, it must first be transformed into numerical form. Tokenization is the fundamental step, splitting words into sub-entities and mapping them to an index in a vocabulary. The simplest feature representation is the Bag-of-Words, which counts token occurrences in a document. The TFIDF matrix improves on this by scaling token frequency relative to the entire corpus, emphasising rare words (often nouns) with potentially more valuable information. However, these methods ignore grammar and word meaning. More advanced approaches like as `Word2Vec`[66], `Glove`[78] or `FastText`[8] build word embeddings, mapping words into a vector space where similar tokens are close-by. Nevertheless, these embeddings lack context awareness, meaning the same word can have different meanings. Attention-based models like `BERT` [21] provide contextual representations, capturing semantic information for sentence level that can be fed to classification models such as neural networks. Contextual embeddings have been shown to hold powerful representations that can be then used for downstream tasks such as sentiment and semantic analysis. Structural information such as Part-of-Speech (POS) tagging and Constituency Parsing can also be extracted. POS tagging assigns grammatical tags to words, indicating their functions within sentences. Constituency parsing creates a syntax tree representing the hierarchical structure of sentences, which can be abstracted into grammar production rules defining permissible word arrangements. Both POS tags and syntax structure are valuable for tasks like attribute or authorship classification, machine translation, or named entity recognition.

¹<https://github.com/facundooolano/google-play-scraper>

Language Models. A language model (LM) is a probabilistic model of a natural language. The model learns a conditional probability distribution $P(w)$ of the next word/token given the sequence of size n of previous tokens. We can write the conditioned distribution as $P(w_{t_n}|w_{t_0}, w_{t_1}, \dots, w_{t_{n-1}})$. Modern LMs, including Large Language Models (LLMs) like ChatGPT² or Llama3³, are based on the Transformer architecture [106] and predict text in an autoregressive manner (each predicted token depends on the previous ones). The Transformer consists of layers of self-attention mechanisms and feedforward neural networks. Self-attention helps encode the input contextually and capture long-range dependencies within a sentence, while cross-attention between encoded inputs and previous tokens aids in predicting the next token. The model then generates probabilities for each word in the vocabulary using a linear and softmax layer, at each step. Current LLMs are typically decoder-only models [2, 9, 102], using prompted text as input.

Recent LMs are pre-trained on large corpora to predict or reconstruct masked sentences and words, learning the language’s semantics and patterns in an unsupervised manner. These pre-trained models can then be used to extract meaningful representations useful for downstream tasks such as classification or measuring semantic similarity. The models can also be fine-tuned on specific tasks using smaller, labelled datasets, adapting them for tasks like question-answering or classification. However, with the advance of LLMs, carefully crafted prompts can achieve similar results without fine-tuning, as seen with models like GPT3 [9], InstructGPT [75] and FLAN models [12]. These models, trained on instructions and human feedback (reinforcement learning from human feedback), can perform tasks such as paraphrasing, text simplification, or style transfer in zero-shot and few-shot scenarios.

Dementia Classification. Dementia is a condition that progressively deteriorates one’s cognitive capabilities, affecting their ability to retain and recall information, reason and communicate. The most common form of dementia, Alzheimer’s disease, accounts for approximately 60-80% of cases. Early detection is key to preventing and providing care for patients with dementia. To this end, several studies investigated differences between a control group and patients with Alzheimer’s and found significant changes in syntactical, lexical and acoustic indicators [31, 52, 54]. For example, Le et al. [54] perform a longitudinal study on four authors and have found that dementia patients demonstrate *lexical indicators* such as a decline in vocabulary and in the usage of complex words, an increased number of low-specificity nouns and verbs, but also *syntactical markers* such as lower frequency of nouns and higher frequency of verbs and lower usage of the passive form. Recently, machine learning algorithms have been successfully applied for automatic dementia detection from speech (audio) and speech transcripts (text), presenting a cheaper, more efficient, and non-intrusive diagnosis method for patients. Garcia et al [16] found that features extracted from the transcripts were often the most indicative of dementia. Others used a combination of linguistic and acoustic features, including *pause* and *disfluency* markers. For example, Yuan et al. [115] incorporated disfluency tags and silences in the transcript and achieved 89.6% detection accuracy on the ADReSS [60] dataset.

²<https://chatgpt.com/>

³<https://llama.meta.com/Llama3/>

2.3 Threat Model

We consider any adversary with access to a target user’s transcribed speech. The adversary’s goal is to detect whether the target user suffers from dementia. The adversary can operate in two settings: (a) a static setting and (b) an adaptive setting. A static adversary (\mathcal{A}) is obfuscation-oblivious and has direct access to the raw text. \mathcal{A} does not have access to any additional information such as age, mental state exam score or audio recordings and has to rely on the text or transcript only. The adversary can gain access to available upon request DementiaBank datasets, blogs or novels, as well as public knowledge of dementia-related stylistic properties. An adaptive adversary (\mathcal{A}_{Ada}) is obfuscation-aware. We assume that \mathcal{A}_{Ada} has gained access to samples generated by any applied obfuscation strategy and can leverage these obfuscated samples to adapt and improve their detection capabilities.

For both scenarios, we consider the adversaries able to implement a kernel-based text classifier (SVM) trained on word token frequencies and a neural network (BERT) able to learn more complex structures and contextual information within a sentence and documents. These were chosen from the ADReSS challenge. The SVM classifier is first described by Luz et al [60] in their paper describing and setting the baselines for the ADReSS Interspeech 2020 challenge. It achieved the best performance across several other classifiers on the test set when using only linguistic features. Our choice of SOTA is the BERT model which won the ADReSS challenge on detecting dementia from text [115].

Our defense mechanism aims to perform automatic dementia attribute obfuscation from \mathcal{A} and \mathcal{A}_{Ada} in text while preserving the semantics of the original transcribed speech.

3 Dementia Obfuscation in Text

In this work, we present an approach that takes advantage of the knowledge and rich representations of large-language models (LLMs) to perform dementia obfuscation in text. However, LLMs come with limitations and constraints, the most important one being their size and their operational overhead. Indeed, they require a large amount of resources to both hold them in memory and perform inference in a reasonable time. To tackle this issue, we perform knowledge distillation by prompting an LLM to perform obfuscation (stage I) and recording the [*original*, *obfuscated*] pair to train a smaller and more efficient model on a synthetic dataset (stage II). We further enhance the training efficiency and modularity of our approach fine-tuning our model through a Parameter-Efficient Fine-Tuning (PEFT) method. We describe the different stages of (shown in Figure 2.) of our methodology in the sections below.

3.1 Knowledge Distillation

Dementia obfuscation is particularly challenging as sparse data and no parallel datasets are available. This situation is similar to unsupervised style transfer, where various methods have been explored. A naive solution is to define a set of transformations. However, such rule-based systems [11, 46], require manual design, lack the flexibility of trained models and may fail to capture the full range of linguistic patterns associated with dementia. Back-translation has been used for style transfer, involving translation to another language and back to the original language [57, 85]. To enforce a target




Figure 2: The two steps of DiDOTS: 1) the creation of a synthetic obfuscated dataset and 2) finetuning a smaller pre-trained language model (PLM) on this dataset with LoRA and a reconstruction loss L_{CE} (cross-entropy). The weights of the PLM are frozen and only matrices A, and B are finetuned. At inference time, the PLM weights are combined with the learned matrices and the model performs obfuscation on input sentences.

style, Prabhumoye et al. [85] investigate back-translation combined with adversarial learning and cycle consistency. Similarly, Shetty et al. and Iyyer et al. [42, 95] employ a GAN-like model to perform style transfer. However, these methods risk overfitting on small datasets, struggle with semantic reconstruction, and are difficult to train due to multiple balancing losses. Imposing a semantic reconstruction loss on an auto-regressive model is non-trivial (the loss is computed on one token at a time) and requires the use of more complex systems such as reinforcement learning [59, 80, 113] or human-in-the-loop [32], adding complexity, instability and manual labour (in the case of human-in-the-loop).

Controlled text generation is an alternative approach which guides the model during inference, requiring less data and computation [48, 51, 68]. However, it still depends on using a target attribute classifier and the base model’s capabilities, which may be insufficient in smaller models, leading to challenges in fluency and semantic reconstruction. One might try to focus on controlling syntax and linguistic features only, releasing the dependency on the dataset. Nevertheless, these approaches introduce multiple fine-grained control knobs and struggle to generate quality sentences [6, 42]. Enforcing a specific style or lack of it without a classifier is also challenging in that setting.

LLMs on the other hand, are powerful foundational models which hold rich representations and achieve state-of-the-art performance on many zero-shot tasks. Hence, LLMs offer a more direct solution for scenarios with limited, non-parallel data. However, their immense size (Gemini Ultra by Google has 1.75 trillion parameters ($O(10^{12})$) [102]) makes them costly to train and inaccessible to most institutions. Additionally, LLMs may exhibit hallucinations or refuse to follow instructions, rendering them less reliable. Model compression techniques like quantization [102] and pruning [65, 72] can reduce LLM size, but these methods often result in accuracy loss. Quantization, in particular, is a lossy process that can degrade model performance [93].

Given these challenges, knowledge distillation (KD) on LLMs emerges as the most efficient and straightforward strategy. Our key insight to overcome the limitations imposed by a small, non-parallel dataset while taking advantage of the powerful nature of LLM in text generation is to *distill knowledge from large language models into a smaller model that is easier to train, store and share, and generates more stable output.*

Teacher-student Model. In our work, we leverage KD introduced by Hinton et al. [38], where we transfer knowledge from a large, complex model (teacher) to a smaller, more efficient one (student). Specifically, we perform sequence-level knowledge distillation [49] by generating a synthetic dataset using the larger model’s outputs. This approach is more efficient than approximating the teacher’s internal states due to the LLMs’ size. By distilling the knowledge already learned by the larger model into the smaller one, we create a model that is computationally lighter while retaining much of the performance of the original model. This approach avoids sacrificing the accuracy or complexity of the resulting model, in favor of a bias toward examples specific to our downstream task (fine-tuning).

Our method is illustrated in Figure 2. In the first stage, we use an LLM as a teacher model to generate [original, obfuscated] sentence pairs to create a synthetic dataset to, in stage 2, teach a smaller pre-trained language model (PLM). Our student model implementation is based on a pre-trained BART model [56], which is built as a classical Auto-Encoder Transformer. We chose BART due to its good performance and generation quality. Thanks to its denoising pre-training task, BART is more robust to input noise and generates fluent sentences. This is desirable because dementia sentences are less formal and other paraphrasing models tend to produce sentences in non-standardized form. The model is fine-tuned on the synthetic dataset through supervised fine-tuning with cross entropy-loss. Note that BART can easily be replaced in this process with another sequence-to-sequence model.

Data Labeling & Prompting Strategies. A key step of our approach is synthetic data generation by prompting instruction-based LLMs and relies on the quality of instructions provided to the selected LLM. We follow and compare several strategies to achieve good-quality outputs. Specifically, we prompt an instruction-based LLM to generate an obfuscation candidate for every sentence in our dataset. We use three different prompting strategies as shown in Figure 3: *zero-shot*, *few-shot*, and *knowledge-based*. To all prompts, we add safeguard instructions to force the model to follow the prompt as closely as possible and avoid rejections or noisy answers. In the zero-shot setting, we ask the model to obfuscate a dementia sentence. In the few-shot setting, the model is given 10 examples of dementia and control sentences. Finally, in the knowledge-based prompt, we incorporate knowledge from previous works and dementia detection studies. These have identified that dementia patients tend to be more wordy, and vague, use simpler vocabulary

and display highly disfluent speech [7, 31, 52, 54]. Based on this we construct an instruction for the LLM asking it to clarify a given sentence. Moreover, in our instruction, we do not mention dementia. This strategy is employed to circumvent the censorship showcased by some LLMs when sensitive topics are mentioned. In the optic of ensuring ‘alignment’ (a term designating the capacity of a model to follow given ethical and legal guidelines) and reducing risks linked to chatbot usage, companies such as OpenAI, Meta or Google implement various safety measures. These can include experts red teaming and evaluating the safety of prompts [4, 74] and filtering data used for pre-training or post-generation moderation [2, 73, 102, 104]. Reportedly, ChatGPT and Llama2 models have shown high levels of censorship (on topics related to politics, gambling or “killing” time) [72, 89, 91], often misinterpreting a benign prompt as unsafe (see examples in Table 5). After generation, samples are cleaned from the template to obtain the obfuscated sentence.

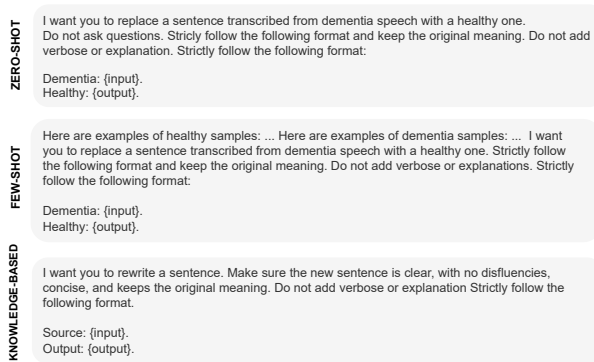


Figure 3: Instructions for zero-shot, few-shot and knowledge-based prompting strategies. We omit few-shot examples for the sake of space.

3.2 Parameter-Efficient Fine-Tuning

As mentioned, we fine-tune our student model using the synthetic dataset. However, we do not perform full fine-tuning but instead employ a Parameter-Efficient Fine-Tuning (PEFT) technique. PEFT is a new paradigm that involves optimizing pre-trained language models for specific tasks by fine-tuning only a small subset of parameters [40, 55, 58]. Through PEFT we aim to retain the majority of the student model’s pre-trained knowledge, making it resource-efficient and faster than full model fine-tuning. Also, by focusing on fewer trainable parameters, we further reduce the risk of overfitting and maintain the model’s overall performance.

More specifically we fine-tune our student model using Low-Rank Adaptation (LoRA) [40], a state-of-the-art PEFT technique. With LoRA, we inject trainable low-rank matrices into each layer of the Transformer architecture, allowing only these matrices to be fine-tuned while the rest of the model remains frozen. Therefore, we can drastically reduce the number of trainable parameters, making the fine-tuning process efficient and preserving the LLM’s core knowledge.

With LoRA, the adaptation of the model is achieved by decomposing the weight updates into low-rank matrices. Specifically, for a given weight matrix W in the transformer model, we introduce two smaller matrices A and B such that:

$$\Delta W = A \cdot B$$

where $A \in \mathbb{R}^{d \times r}$ and $B \in \mathbb{R}^{r \times k}$, with $r \ll \min(d, k)$. Here, d and k are the dimensions of the original weight matrix W . During fine-tuning, only the parameters in A and B are updated, keeping the original W unchanged. The effective weight matrix W' during adaptation thus becomes:

$$W' = W + \Delta W = W + A \cdot B$$

This low-rank decomposition significantly reduces the number of parameters that need to be trained, as the number of parameters in A and B combined is much smaller than in the original W .

4 Evaluation Setup

4.1 Evaluation Questions

We conduct a comprehensive evaluation aiming to answer the following concrete questions: **(EQ1)** How effective is our knowledge distillation method through an LLM-generated dataset?; **(EQ2)** How do our proposed obfuscation methods compare to state-of-the-art and baseline approaches?; **(EQ3)** How effective are LLMs in obfuscating dementia and preserving semantics?; **(EQ4)** How do different components of our knowledge distillation method contribute to performance (utility/privacy)?; **(EQ5)** How do obfuscation systems impact key dementia discriminative features (privacy)?; **(EQ7)** How do the generated sentences compare qualitatively with the original sentences?

4.2 Datasets

We use two available datasets which are widely studied in the field of dementia classification: ADReSS [60] and ADReSSo [61]. The ADReSS (ADR) dataset is a curated subset of the DementiaBank dataset [7], a collection of 500 recordings from individuals with various stages of dementia. The samples were manually transcribed with dysfluency annotations for the task of the Cookie Theft Picture description. The speakers were selected to be balanced for gender and age, and they split evenly into control (CC) and dementia (AD) groups (54 train and 24 test for each class, a total of 156 samples). In our experiments, we use the provided train/test splits and split the documents on sentence level. We further remove sentences with less than 3 words. We end up with 1179 samples (619 CC | 560 AD) in the training set and 500 (270 CC | 230 AD) in the test and validation sets. The ADReSSo (ADRo) dataset is another subset of DementiaBank designed for detecting dementia from spontaneous speech only, without access to manual transcriptions. The original set consists of 151 train samples (87 CC | 74 AD) and 71 test samples (35 CC | 35 AD). We use the provided segmentation timestamps to isolate segments spoken by the patients and transcribe them using Whisper, a SOTA speech-to-text model [87]. We get 1421 samples in the train set (710 CC | 711 AD) and 644 samples (347 CC | 297 AD) in the test and validation sets.

4.3 Adversarial Models

In evaluating the ability of obfuscators to preserve privacy it is important to consider an adversary in both a static and an adaptive setting. In the former case, the adversary uses available datasets to train the adversarial model. In the second, stronger scenario, the adversary also has knowledge of the defense mechanism. We implement both a static and adaptive adversary as described below.

Static Adversary. We finetune a pre-trained BERT-base model with a classifier head as our neural adversary. BERT has been shown to have SOTA performance for various NLP tasks [3, 28, 100]. We use the base model from [115]⁴ with a learning rate of 1e-6, 10 epochs with early stopping (patience = 1 on validation loss), gradient clipping of 1, input length limit of 256 tokens, and batch size of 8. We also select an SVM as a kernel-based baseline, as presented by [60] and [111]. We extract a TF-IDF matrix from the corpus and feed it to an SVM with radial-basis-function kernel and $C = 1.0$. The TF-IDF matrix also focuses on different features with learned word frequencies and relative frequencies while BERT looks at the big picture and overall fluency.

Adaptive Adversary. In addition to the static adversaries we also evaluate our systems against adaptive adversaries with knowledge of the defense mechanism. We assume the adversary knows which obfuscator is being used, and that they have black-box query access to it. We train both SVM and BERT models with training data which includes raw data and their obfuscated version for a given defense.

4.4 Metrics

Privacy. We measure the privacy gain through the drop on the various adversaries' F1-score, for both static and adaptive settings. We refer to the F1-score as the adversary success rate (ASR). We use the F1 metric as it is less sensitive to data imbalance, and shows the model's ability to capture dementia cases and transform dementia samples to control-like ones.

Utility. There are currently no available dementia datasets that are labeled for specific tasks such as topic modeling or sentiment analysis. Therefore, for our utility metric, we focus on semantic preservation. Semantic preservation in natural language is a powerful primitive that is important for several downstream tasks such as topic modeling, patient and customer complaint analysis, note-taking, etc. We automatically compute the semantic similarity between the original sentence and the output of a selected obfuscation model. We use ParaBART [41], a BART-based model trained to create syntax-invariant semantic embeddings. Common semantic embedding models suffer from poor robustness to different syntax and thus would lower the semantic score between two sentences with different structures even if the meaning is the same. ParaBART on the other has been shown to be more robust to such changes.

Automated metrics, however, do not always align with human judgment. To mitigate this threat to the validity of our evaluation, we also design a targeted study with human subjects. Our study is designed to measure users' perception of the ability of different text paraphrasing models to preserve the semantics of original sentences. Our study was reviewed and received approval from

⁴We evaluate both the base and large versions but pick the base model as it was more robust to obfuscation.

our institution's research ethics review board. Lastly, we further provide a qualitative analysis of the paraphrase quality of different systems: we select representative sentences from our dataset and compare them with the output of competing systems.

Quality and Diversity Metrics. To better understand the effect of different systems we measure the quality and diversity of the generated samples by computing the formality score, METEOR score (a common metric that measures word matches, synonyms, and word order.) [53], lexical similarity (Simi_{Lex}) measured by the word-level Levenshtein distance between sets of unique words, perplexity and the percentage of substitutions (% SUB), additions (% ADD) and deletions (% DEL) between the source and generated samples. We compute an automatically generated formality score of a sentence by using a RoBERTa model from the Hugging Face library⁵ trained on the formality-style transfer dataset GYAFC [88] and online formality corpus [77]. Given that neural text generation creates samples with word distribution that are distinct from human-written text, we also investigate the perplexity (PPL) of our generated samples. A high perplexity is commonly associated with human text but is also a sign of low fluency. We compute the perplexity by computing the exponentiated average negative log-likelihood of a pre-trained GPT2 model using the *lmppl* library⁶.

4.5 Obfuscation Models

Pegasus (Paraphrasing Model). We compare our approach to the state-of-the-art paraphrasing model. We employ a Pegasus model [116] fine-tuned to perform paraphrasing on the input text. Pegasus is a transformer-based model that has shown impressive performance in generating high-quality text and paraphrases. It was pre-trained on a large corpus with the task of retrieving important masked sentences and later fine-tuned on paraphrase pairs.

DP (Differential Privacy Obfuscation Model). Following work by [64], we re-implement a text generation model that performs obfuscation by sampling with higher temperatures. We select Pegasus as our language model and do our best to improve on [64]—our version is bigger and trained on a larger dataset than their model, to maximize the quality of the generated samples. We also experiment with various temperature values (see Appendix B).

ParChoice (Heuristics-based Obfuscation Model). ParChoice is a rule-based obfuscation model that has been shown to have good obfuscation abilities while preserving semantics and utility [36]. The model performs a series of modifications such as changing writing voice, contractions, typos, synonyms, etc. It also performs paraphrase selection via a surrogate profiler for style transfer and filter-generated sentences for fluency and grammatical correctness. This approach outperforms simpler rule-based approaches, but also complex models like A4NT [95] and MutantX [63] for both obfuscation and semantic retention. We use the code made available by the authors and run ParChoice in "Random" settings as well as with an SVM and BERT surrogates trained on ADR samples that guide the obfuscation toward a control-like style.

LLMs (Ours) We investigate top-performing LLMs for different size brackets in their ability to perform dementia obfuscation. Due to

⁵<https://huggingface.co/s-nlp/roberta-base-formality-ranker>

⁶<https://github.com/asahi417/lmppl>

the sensitive nature of the data, we only use open-sourced models that can be run locally. As ChatGPT and other OpenAI models access and process input data, we omit them in this study. We select Gemma 2B [102], Phi3 (3B parameters) [2], Mistral 7B [43] and Llama3 8B [72]. We pick Mistral as our best system to train DiDOTS and present all three prompt settings in the summary Table 1. We run each model via the Ollama python library⁷.

DiDOTS (Ours) We implement our model using the pre-trained *BART-base* model from the Huggingface library⁸. We implement and train the core of our system with LoRA ($r=16$, $\alpha = 32$) on the KB dataset: DiDOTS, and our ablation models: $BART_{KB}$ (full fine-tuning), $T5_{KB}$ (T5 model⁹), $BART_{ParaNMT}$ (trained on paraphrasing dataset ParaNMT [109]), as well as $BART_{KB_BOFT}$ and $BART_{KB_IA3}$ for different PEFT techniques: IA3 and BOFT. Intrinsically-Aligned Adapters (IA3) is an adapter technique that injects a small number of trainable parameters (64.5K) into the model. Bottleneck Output-Fusion Transformers (BOFT) introduce a bottleneck layer (138K) to compress and then expand the intermediate representations, to focus on relevant features. We use the paper’s default parameters for both. We finetune our models on the synthetic paraphrase pairs for a maximum of 3 epochs and use early stopping (patience = 10 on validation loss), with a batch size of 8 and a learning rate of $1e-4$. At inference time, we sample with the number of beams set to 4 and a maximum length of 256.

5 Evaluation Results

5.1 Overall Performance and Comparison

In this section, we aim to answer EQ1 and EQ2 (see Section 4.1). We compare several baselines (*Pegasus*) and SOTA (*DP* and *ParChoice*) models with our best performing LLM (Mistral) used with zero-shot (ZS), few-shot (FS) and knowledge-based instructions (KB), and with our knowledge distillation approach from Mistral with LoRA adaptation (DiDOTS). Note that a more comprehensive evaluation of LLM performance is presented in Section 5.2.

The systems are evaluated for privacy against both our static and adaptive adversary and for utility (semantic preservation). We also indicate the size of each model when relevant as one of our main objectives is to distill knowledge from large language models, into a smaller model. Lastly, we also evaluate the above on both the ADR and the ADRo datasets. Evaluating on ADRo allows us to examine the effects of leveraging automated speech recognition to derive transcriptions for the adversary and the defense mechanisms and is a more challenging scenario compared to the manually transcribed samples of ADR. We summarize the results in Table 1.

Automatic Privacy and Utility Analysis of Results. Compared to *Pegasus*, *DP* and all *ParChoice* versions, our LLM approaches and DiDOTS achieve better privacy performance on average and on both datasets. In particular, compared to the best version of *ParChoice* ($ParChoice_{SVM}$), DiDOTS achieves 1.3x better privacy performance or a further 5.7% decrease in F1-score on the ADR dataset and a more impressive 2.2x improvement in privacy performance or a further 11.2% decrease in F1-score on ADRo.

Importantly we also observe that *ParChoice*’s privacy performance suffers with the ADRo dataset as it cannot deal with speech-to-text errors well, while our approaches exhibit more consistent performance. It is worth noting though that $ParChoice_{SVM}$ does very well in the case the adversary also uses an SVM attack model as we observe a 23.1% drop in F1-score on the ADR dataset. However, this is not very realistic, as it assumes the adaptive adversary (with knowledge of the defense mechanism) will use the same model *ParChoice* as its most performing version is classifier dependent and does *less* well if the adversary does not use the model the defense is using. To better illustrate this, we consider the worst-case scenario for the defense, i.e. when the adversary makes all the best choices of models and datasets to use. In this case, the worst case scenario for $ParChoice_{SVM}$ is against an adaptive BERT adversary on the ADRo dataset where it achieves a lackluster 3.0% F1-score drop. Comparatively, DiDOTS’s worst-case scenario is against an adaptive SVM adversary on the ADRo dataset. In this case, our system still achieves a 9.5% F1-score drop which is at least 3x better compared to $ParChoice_{SVM}$ ’s worst-case performance.

Another important observation from these experiments is that DiDOTS retains the privacy and semantic preservation performance of the zero-shot (ZS) and knowledge-based (KB) LLMs, and their generalizability across adversarial scenarios and datasets. In addition it only requires fine-tuning $O(10^5)$ parameters compared to fully fine-tuning the $O(10^8)$ student model.

Lastly, all models achieve good semantic preservation with values close to 0.7 and above. *Pegasus* achieves the best values according to our utility metric although the differences are very subtle—our human evaluators ranked our systems better than *Pegasus*. Notably, all of the systems exhibit a semantics drop on the ADRo dataset. This is because there are inherent errors from the automatic speech recognition used to derive the ADRo transcriptions that make high-quality paraphrasing more challenging. The Mistral few-shot LLM (Mistral FS) exhibits the worse utility performance. This is further examined in Section 5.2, and the paraphrasing quality and users’ perception of the semantics of the best systems are analyzed in more depth in Section 5.2 and Section 5.4.

Impact on Key Differentiator in Dementia Detection. To assess the privacy performance of different systems, we analyzed their impact on 10 key linguistic features relevant to dementia detection, based on existing literature. [24, 31, 47]. These features capture the disfluency (interjections), and the lexical complexity (number of syllables per word) but also markers for memory loss such as higher usage of verbs, pronouns and adverbs to compensate for forgetting a specific noun. To focus on the systems at hand (and not on automated speech recognition effects) we pick ADR for our case study. On Table 2 we report our evaluation results for the five best-performing systems: $ParChoice_{SVM}$, *Pegasus*, *DP*, $Mistral_{KB}$, and DiDOTS. We see that both $Mistral_{KB}$, and DiDOTS achieve significant reductions in the proportion of interjections (−96.83% and −94.52%) and adverbs (−35.67% and −57.74%). They also show relatively important improvements in increasing the proportion of nouns (+8.53% and +6.16%) and for $Mistral_{KB}$, the mean syllables per word (+6.39%). *Pegasus* and *DP* also reduce interjections, adverbs, and generic nouns and increase total nouns but to

⁷<https://github.com/ollama/ollama-python>

⁸<https://huggingface.co/facebook/bart-base>

⁹<https://huggingface.co/google-t5/t5-base>

Table 1: Privacy/utility performance for ADReSS and ADReSSo datasets by different systems. This table presents the F1 score against static (S) and adaptive (A) dementia classifiers, and the average F1 score (Avg F1) across all classifiers. It also depicts semantic preservation (Sem.). We mark our best systems in bold. ↓ illustrates the defense mechanism aims to decrease that value, and ↑ to increase it.

Type	Systems	# Params	ADReSS						ADReSSo					
			BERT↓		SVM↓		Avg F1↓	Sem.↑	BERT↓		SVM↓		Avg F1↓	Sem.↑
			S	A	S	A			S	A	S	A		
Original	-	-	0.65	0.65	0.65	0.65	0.65	1.0	0.66	0.66	0.63	0.63	0.65	1.0
Baselines	Pegasus	568M	0.52	0.58	0.60	0.59	0.57	0.85	0.5	0.57	0.58	0.58	0.56	0.81
	DP _{T=1.5}	568M	0.51	0.6	0.59	0.56	0.57	0.84	0.47	0.63	0.58	0.61	0.57	0.77
	Parchoice _{Random}	-	0.62	0.62	0.6	0.58	0.61	0.80	0.6	0.64	0.63	0.63	0.63	0.76
	Parchoice _{SVM*}	-	0.55	0.58	0.51	0.50	0.54	0.84	0.58	0.64	0.51	0.61	0.59	0.76
	Parchoice _{BERT*}	110M	0.65	0.62	0.67	0.58	0.63	0.78	0.65	0.62	0.61	0.6	0.62	0.74
LLMs-based (Ours)	Mistral ZS	7B	0.35	0.57	0.57	0.48	0.49	0.72	0.43	0.51	0.57	0.61	0.53	0.64
	Mistral FS	7B	0.18	0.5	0.55	0.51	0.43	0.62	0.25	0.54	0.31	0.56	0.42	0.28
	Mistral KB	7B	0.31	0.63	0.55	0.57	0.52	0.78	0.36	0.58	0.52	0.58	0.51	0.67
	DiDOTS	140M (884k adapter)	0.37	0.55	0.49	0.58	0.50	0.85	0.4	0.58	0.5	0.57	0.51	0.74

* Classifier-dependent

Table 2: Effect (in %) on key dementia characteristics by best systems the ADReSS dataset. Values going in the direction of control-like features are highlighted and we mark in bold the top-2 systems per feature.

	ParChoice _{SVM}	Pegasus	DP	Mistral _{KB}	DiDOTS
Prop. Interjections↓	41.04	-3.12	-4.15	-96.83	-94.52
Prop. Adverbs↓	-15.09	-27.50	-29.50	-35.67	-57.74
Mean Syllables Per Word↑	1.90	-1.94	-1.91	6.39	0.34
Prop. Nouns↑	5.54	0.11	1.73	8.53	6.16
Prop. Modals↑	-17.57	23.54	23.91	8.11	23.66
Prop. Verbs↑	-7.44	3.70	3.81	3.52	-3.95
Ratio Verbs/nouns↓	-5.48	-18.87	-16.40	12.64	-9.70
Ratio Adverbs/nouns↓	-16.11	-48.85	-49.31	-33.19	-60.11
Generic. Nouns↓	-21.29	-19.61	-7.89	-14.62	-26.31
Generic. Verbs↓	-29.90	18.47	19.20	-23.74	-3.88

a much lesser extent. They stand out in reducing the verbs/nouns (-18.87%, -16.4%) and adverbs/nouns ratio (-48.85%, -49.31%) and increasing modals (+23.54%, +23.91%). They, however, decrease the proportion of mean syllables and increase generic verbs, emphasising dementia-like features. These results can be explained by the summarising tendencies of the models. ParChoice_{SVM} shows notable reductions in generic verbs (-29.90%) and nouns (-21.29%) but adds a significant amount of disfluency (+41.04%) and reduces verbs and modals. ParChoice_{SVM}'s impact on the other features, albeit positive, is marginal compared to the other systems. This is expected, as Parchoice relies more on singular and targeted typos and substitutions than paraphrasing.

Summary of Results. Our evaluation shows that:

- ZS and KB LLMs are better obfuscators compared to the SOTA text paraphrasing model and the SOTA attribute obfuscator.
- Both the KB LLM and our distilled model consistently perform more meaningful changes across multiple linguistic metrics than the other systems. ParChoice_{SVM} lacks particularly behind and introduces disfluencies.

- Our distilled model even though it has one order of magnitude less parameters ($O(10^8)$) compared to the teacher LLMs ($O(10^9)$), it retains the LLM's performance in obfuscating dementia in text, and in preserving the semantics of the original sentence, it generalizes to different adversarial scenarios and datasets, and it only requires fine-tuning three orders of magnitude less parameters ($O(10^5)$) compared to full fine-tuning for our specific task.

5.2 LLMs as Dementia Obfuscators

In the above analysis, we used our most effective LLM. To answer EQ3, here we compare different LLM choices. We compare four open-source models *Gemma*, *Phi3*, *Mistral* and *Llama3*. We perform a number of experiments testing different LLM models' ability to preserve privacy and preserve semantics. We then delve deeper into the generation ability of LLMs and further analyze the changes introduced by the various models when given the more abstract instruction of "obfuscating dementia" (zero-shot setting).

Privacy and Utility. For each LLM model, we compare three prompt strategies: zero-shot (ZS), few-shot (FS), and knowledge-based (KB). The results are summarized on Table 3. We see that all systems perform particularly well against the static adversaries but have mixed results against the adaptive ones. We observe a general pattern between different prompting strategies. FS achieves the best privacy scores but also the lowest semantics (≤ 0.61). This is due to the models mimicking the problem when they were given the example sentences (see Table 5), which they all seem to suffer from. On the other hand, the KB setting improves the semantics for all systems but Gemma, for better performance against static adversaries and comparable performance against adaptive adversaries. The KB strategy particularly improves Llama3 semantics (ZS 0.68 → KB 0.82). We note that Llama3 in ZS had failed outputs due to censorship (see Table 5) and observe that it was no longer the case with KB. Both Gemma and Phi3 models have good privacy scores overall despite their relative sizes but are still slightly worse than Mistral and Llama3 models and fall behind in semantics preservation. In

ZS setting, Mistral achieves the best overall privacy/utility performance with 0.72 semantics and drops the F1-scores of 46%, 12% for the BERT adversary and 12%, 26% for the SVM static and adaptive adversaries respectively. In the KB setting, Llama3 achieves the best privacy/utility tradeoff and the highest semantics score (0.82). In our comparison experiments (Section 5.1) we used Mistral as our representative LLM as it had a more consistent performance across the different prompting strategies.

Table 3: Privacy and Utility Performance of various LLMs. This table presents the F1 score against static and adaptive dementia classifiers as well as the semantic preservation score. The best results for each column are marked in bold.

Model	# Params	BERT↓		SVM↓		Avg. F1↓	Semantics↑
		Static	Adaptive	Static	Adaptive		
Original	-	0.65	0.65	0.65	0.65	0.65	1.0
Gemma ZS	2B	0.46	0.59	0.62	0.63	0.58	0.75
Gemma FS	2B	0.22	0.60	0.39	0.37	0.39	0.61
Gemma KB	2B	0.28	0.49	0.51	0.60	0.47	0.73
Phi3 ZS	3.8B	0.38	0.55	0.55	0.61	0.52	0.64
Phi3 FS	3.8B	0.19	0.53	0.18	0.60	0.38	0.36
Phi3 KB	3.8B	0.27	0.54	0.50	0.58	0.47	0.74
Mistral ZS	7B	0.35	0.57	0.57	0.48	0.49	0.72
Mistral FS	7B	0.18	0.50	0.55	0.51	0.44	0.62
Mistral KB	7B	0.31	0.63	0.55	0.57	0.51	0.78
Llama3 ZS	8B	0.39	0.52	0.57	0.55	0.51	0.68
Llama3 FS	8B	0.17	0.53	0.37	0.49	0.39	0.57
Llama3 KB	8B	0.37	0.54	0.48	0.55	0.48	0.82

Paraphrasing Quality. We analyse the differences at the sample level between different LLMs in a zero-shot setting. We add top-performing systems from Section 5.1 for further analysis. Some examples of samples from the LLMs are shown in Appendix C. Results are shown in Table 4. Among LLMs, Gemma and Mistral perform similarly across the metrics, despite their size difference, showing high semantics (0.75) and METEOR scores (0.56, 0.53). They also have high formality scores (0.72, 0.75) and the lowest perplexity scores (71, 97), making them the most machine-like language models. Phi3 falls in-between with lower semantics preservation and diverges the most from the source (low METEOR and Simi_{Lex}). Llama3 remains the closest to the source vocabulary in terms of Simi_{Lex} but has low semantics and METEOR scores. It is the most human-like across the LLMs with a formality of 0.56 and perplexity of 382. All LLMs make a higher number of substitutions (≥ 0.47) than the other systems as well as additions (≥ 0.13), except for Parchoice_{SVM} (0.19). On the other hand, Parchoice_{SVM} has a significantly lower formality score (0.23) and higher perplexity (1018), which can be explained by the introduction of disfluencies noted in Section 5.1. Its samples remain the closest to the source material, with no deletions being recorded. All three language model-based systems Pegasus, DP and DiDOTS perform similarly across metrics and have the highest semantic preservation but have the lowest perplexities and percentage of additions and the highest proportion of deletions. This indicates a tendency to write sentences more concisely. Overall, LLMs exhibit stronger paraphrasing abilities, often leading to more drastic changes compared to other systems.

LLM Shortcomings. Table 5 presents an overview of failed outputs we encountered with the various LLMs when generating text. These

Table 4: Quality Metrics for the zero-shot (ZS) sampling strategy for various LLMs and best-performing systems on the ADReSS dataset. We note S, % SUB, % ADD, % DEL and Simi_{Lex} as semantics, percentages of substitutions, additions, deletions and lexical similarity respectively.

Model	S	Formality	METEOR	% SUB	% ADD	% DEL	Simi_{Lex}	Perplexity
Gemma	0.75	0.72	0.56	0.47	0.13	0.12	0.65	71
Phi3	0.64	0.67	0.31	0.67	0.23	0.10	0.53	180
Mistral	0.75	0.75	0.53	0.52	0.23	0.10	0.63	97
Llama3	0.68	0.56	0.37	0.56	0.13	0.14	0.71	382
Parchoice _{SVM}	0.84	0.23	0.86	0.35	0.19	0	0.79	1018
Pegasus	0.85	0.8	0.63	0.35	0.05	0.24	0.78	142
DP _{T=1.5}	0.84	0.8	0.61	0.36	0.06	0.23	0.77	157
DiDOTS	0.85	0.79	0.64	0.3	0.06	0.21	0.76	114

failures range from misunderstanding prompts and refusing to produce potentially offensive content (only observed in Llama3) to contradicting outputs and unrelated tangents. In the case of the few-shot strategy, models often overfit on the given sentences and output texts mimicking their content. There were also instances of instructions not being followed where the model repeats them or simply outputs verbose instead of the correct output. These examples emphasize some of the challenges of using LLMs for automatic data labelling and generation, besides data cleaning and processing. However, these cases made up circa 4% of the generated outputs (excluding FS) and did not transfer through distillation.

Table 5: Examples of failed LLMs outputs.

Type	Input	Output
Copying (Few-shot)	what do i see in this picture.	I see a picture with a mother standing in the water and a little girl laughing at her brother taking a cookie from the jar
Refusal	everything that's going on okay.	I cannot create dementia samples or any other type of content that may be offensive or harmful to individuals with dementia or their caregivers.
"Hallucination"	and I don't know what the hand means to the mouth.	I understand that the hand is a part of the body that is used for many things, including grasping, pointing, and writing.
Contradiction	looks like they dried	looks well-hydrated
Repeating Instructions	and let's see.	Please provide the sentence you'd like me to rewrite, and I'll follow the format to create both dementia and healthy samples.
"Idem"	and he's about to hand her a cookie.	Keep it the same, no need to replace as the given healthy sentence is already clear

5.3 DiDOTS Ablation Study

To better understand how different DiDOTS design decisions affect obfuscation and semantics preservation (EQ4 and EQ5), we design several versions of our obfuscator model and summarize the results in Table 6. We also investigate the impact of an LLM's undesirable behavior on DiDOTS's performance on Table 7.

Base Model. First, we investigate the impact of the smaller model architecture on the overall performance and implement our system

with BART and T5 base models. T5_{KB} performs well in adaptive settings (0.53 SVM, 0.56 BERT) but has lower performance for static adversaries and semantics than BART_{KB}. T5 is also nearly double the size (220M) of our BART model (140M), adding unnecessary overhead for little gains.

Dataset. To evaluate the importance and quality of the synthetic dataset, we also train our model on a popular paraphrasing dataset, ParaNMT [109], and call this system BART_{ParaNMT}. We find that BART_{ParaNMT} generally performs worse in terms of privacy compared to other models having little impact on the adversaries and a lower semantic score (0.84). This shows that a typical dataset of paraphrasing pairs is not sufficient to obfuscate dementia.

Prompt Strategy. We fully fine-tune our model on the different synthetic datasets generated by Mistral in ZS, FS and KB settings. We note that the effect of the prompting strategies is less apparent through the distillation method than with the LLMs. Surprisingly, the FS has a bigger drop in semantics (1.0 → 0.8) w.r.t ZS and KB (1.0 → 0.87), while achieving the best privacy score on the static adversaries (0.29 BERT, 0.42 SVM). BART_{ZS} performs slightly worse than BART_{KB} for the same semantics scores. We observe that distillation has a soothing effect on the model’s performance, thanks to the model’s pre-training (both BART and T5). We select BART_{KB} as the best model.

PEFT. Finally, we compare different PEFT approaches. We observe that BART_{KB_LA3} has little impact on the obfuscation abilities and naturally preserves the highest semantics, while BART_{KB_BOFT} has slightly worse results than BART_{KB_LoRA} with better semantics (0.87 vs 0.85). We retain however BART_{KB_LoRA} as the best model architecture for DiDOTS as it has the better privacy/utility trade-off.

Undesirable Behaviours. To investigate the impact of training DiDOTS on datasets of varying quality, we pick datasets generated from Llama3 and Phi3, instead of Mistral, as shown in Table 7. Llama3 was chosen as it has exhibited *refusal* in its responses, and Phi3, as it is a smaller model (3.4B parameters) with the poorest semantic preservation. Llama3, with refusals in the ZS setting, does not significantly affect the performance. It achieves high privacy and utility scores, comparable to simple paraphrasing with a semantics score of 0.85 and a mean F1 of 0.57. Phi3 performs well in the ZS and KB settings, achieving high semantics of 0.81 and 0.83, respectively, and maintaining good privacy despite the poorer dataset quality (semantics of 0.64 & 0.74). Similarly, although FS prompts generate mediocre datasets, they maintain relatively good semantics for Llama3 and Mistral after finetuning (see Table 6 for Mistral). This illustrates that even with lower initial semantics, FS prompts can still be effective. However, Phi3 struggles significantly in the FS setting, where its semantics score drops to 0.28 and F1 score to 0.22, highlighting the importance of careful prompt design. For both Llama3 and Phi3, we see the positive impact of the KB strategy, emphasising the advantage of our prompt design. Overall, we show that our approach and KB prompt strategy are robust to variations in the base model.

5.4 Qualitative Semantic Preservation

Lastly, to answer EQ6 and EQ7, we conduct further analyses on the ability of the best-performing models to preserve semantics.

Table 6: Influence of model architecture, dataset, prompting strategy, and PEFT approaches on the privacy-utility performance of DiDOTS on the ADReSS dataset.

Parameter	System	BERT↓		SVM↓		Avg. F1↓	Semantics↑
		Static	Adaptive	Static	Adaptive		
	Original	0.65	0.65	0.65	0.65	0.65	1.00
Model	BART _{KB}	0.34	0.59	0.52	0.65	0.53	0.87
	T5 _{KB}	0.44	0.56	0.57	0.53	0.53	0.85
Dataset	BART _{ParaNMT}	0.60	0.58	0.62	0.60	0.60	0.84
	BART _{ZS}	0.40	0.63	0.54	0.58	0.54	0.87
Prompt	BART _{FS}	0.29	0.59	0.42	0.75	0.51	0.80
	BART _{KB_LA3}	0.47	0.60	0.60	0.65	0.58	0.89
PEFT	BART _{KB_BOFT}	0.40	0.56	0.54	0.58	0.52	0.87
	BART _{KB_LoRA}	0.37	0.55	0.49	0.58	0.50	0.85

Table 7: Influence of the quality of the LLM and its synthetic dataset on the privacy-utility performance of DiDOTS on the ADReSS dataset.

System	LLM	Setting	Dataset Semantics	BERT↓		SVM↓		Avg. F1↓	Semantics↑
				Static	Ada.	Static	Ada.		
DiDOTS	Mistral	KB+LORA	0.78	0.37	0.55	0.49	0.58	0.49	0.85
		ZS+LORA	0.68	0.54	0.57	0.61	0.59	0.57	0.85
		FS_LORA	0.57	0.14	0.61	0.29	0.51	0.39	0.71
Refusal	Llama3	KB+LORA	0.82	0.47	0.58	0.54	0.58	0.54	0.89
		ZS+LORA	0.64	0.41	0.61	0.58	0.59	0.55	0.81
		FS+LORA	0.36	0.03	0.51	0.00	0.21	0.22	0.28
Size	Phi3	KB+LORA	0.74	0.31	0.56	0.48	0.55	0.53	0.83

Human Evaluation As mentioned in Section 4, automated semantic similarity metrics do not always align with human judgement. To mitigate this, we further design an evaluation of semantic preservation with human subjects. In our study, we select our top five competing systems (ParChoice_{SVM}, Pegasus, DP, Mistral_{KB} and DiDOTS) and ask 4 cohorts of 40 participants to rate 10 paraphrases of original sentences on a 1 to 5 Likert scale, for a total of 40 documents. More details can be found in Appendix D.

The results of our study are summarized in Table 8. We observe that Mistral_{KB} archives the highest semantic score (3.86). It is closely followed by DiDOTS with a mean score of 3.75. A paired t-test (with Holm-Bonferroni correction for many comparisons) suggests no significant difference between the two systems. Pegasus and DP are not too far behind, but (statistically) significantly lower. ParChoice_{SVM} has the lowest semantics score of 2.45. These results contrast with those of the automatic metrics, Mistral_{KB} having a much higher ranking and ParChoice_{SVM} having a lower score. These results highlight the limitations of automatic metrics in capturing semantics for diverse paraphrases (Mistral_{KB}) or small variations (ParChoice_{SVM}). Additionally, while Mistral_{KB} may produce some incorrect outputs affecting overall semantic scores, the survey samples demonstrate the LLM’s potential on successful outputs.

Qualitative Sample Comparison Table 9 and Table 10 show generated samples for each of the obfuscation approaches. We pick two sentences, one from a control subject and one from a dementia subject, to showcase the structural differences for similar meanings. The AD sample is longer and has a stutter (“the m the mother”) while the control sample uses fewer words. We can see that ParChoice_{SVM} introduces changes in words and typos, consequently decreasing

Table 8: Mean and standard deviation of semantic (S) human evaluation for different systems. A score of 1 indicates no semantic similarity and 5 identical meaning. The symbols * and † represent systems pairs not significantly different from each other.

	ParChoice _{SVM}	Pegasus	DP _{T=1.5}	Mistral _{KB}	DiDOTS
S ↑	2.45 ± 0.57	3.55 ± 0.46*	3.45 ± 0.47*	3.86 ± 0.51†	3.75 ± 0.47†

fluency and sense. We note that both Pegasus and DiDOTS rewrite the sample more concisely. On the other hand, BART_{paraNMT} paraphrases the text and makes a mistake (“the water was running out of the water.”) or lowers the language level (the mother → mom) while Mistral increases the formality level for the AD sample. All systems except ParChoice_{SVM} managed to remove the original disfluencies.

Table 9: Obfuscating a Dementia (AD) sample sentence.

Obfuscation	Sample
Original AD	well in the first place the m the mother forgot to turn off the water and the water’s running out the sink.
Parchoice_{SVM}	well in the firstly place the m the mother left to turn off the water and the waters running out the sink.
Pegasus	The mother forgot to turn off the water and the water was running out of the sink.
DP_{T=1.5}	The mother forgot to turn off the water and the sink ran out.
BART_{paraNMT}	first of all, the mother forgot to turn off the water, and the water was running out of the water.
Mistral_{KB}	The mother forgetting to turn off the water resultated in it flowing out of the sink.
DiDOTS	The mother forgot to turn off the water and the water is flowing out of the sink.

6 Related Work

Authorship and Attribute Detection. Various stylometric features were identified to detect differences between authors such as lexical features, content-specific and syntactic features [1, 50, 94]. Several mental health conditions can also be detected from text [5, 13, 23, 107]. While these works are crucial for early detection it is also a powerful tool for author profiling. Our work is the first to investigate dementia obfuscation in text.

Authorship and Attribute Obfuscation. Authorship and attribute obfuscation techniques range from rule-based [11, 46] to paraphrasing and text generation. A4NT [95] and Style-pooling [67] perform obfuscation by style transfer by training a model to transfer sentences from one attribute to another or combining multiple styles. MutantX [63] is optimized to perform word changes to fool the authorship detection model. Both in [42] and in [68], authors

Table 10: Obfuscating a Control (CC) sample sentence.

Obfuscation	Sample
Original CC	the mother is uh washing dishes and the water’s spilling over on the uh kitchen floor.
Parchoice_{SVM}	the mother is uh washing lookers and the waters spill over on the uh kitchen floor.
Pegasus	The mother is washing dishes in the kitchen.
DP_{T=1.5}	The mother is washing dishes while the water spills over on the kitchen floor.
BART_{paraNMT}	mom’s washing dishes, and there’s water all over the kitchen floor.
Mistral_{KB}	The mother is washing dishes, but the water is spilling onto the kitchen floor.
DiDOTS	The mother is washing dishes and the water is spilling onto the kitchen floor.

learn syntactical changes to fool a classifier for sentiment, formality, and agency. Lin & Wan [57] implement several iterations of back-translation in hopes of breaking the syntax and introducing synonyms at each pass. ParChoice [36] implements a series of rule-based modifications for authorship and attribute obfuscation. It showed superior privacy performance while maintaining high semantic preservation compared to A4NT and MutantX. Differential privacy methods have also been applied for obfuscation in text, with most focus on authorship obfuscation. Early work focuses on applying DP to the word level but often struggles to preserve semantics and computational efficiency. To counter these limitations, Mattern et al. [64] suggest sampling from a language model using softmax scaled by a temperature.

LLMs Distillation. Knowledge Distillation (KD) was introduced by Hinton et al. [38], to train a smaller model (student) to replicate the logit outputs of a larger, more complex model (teacher). With the emergence of LLMs and increased computational costs, KD techniques have been used to reduce models’ sizes [2, 37, 102] but also to create task-specific smaller-model that leverage the deeper knowledge of LLMs [39, 45, 114]. Due to their size, these models are often distilled using “black-box” or sequence-level knowledge distillation [49] approaches, only accessing the textual output for the student to train on. Comparatively to our work, [10, 79, 96, 114] and [25] performed sequence-level KD on the tasks of paraphrasing, text-simplification and neural machine translation. Text Laundering [44] employ a similar approach to mitigate adversarial attacks on downstream models. To our knowledge, we are the first to utilize KD from LLMs for attribute obfuscation in text.

7 Discussion

Our system aims to prevent dementia attribute leakage in transcribed speech and protect individuals from potential discrimination, such as insurance denial or employment bias. It can be integrated into healthcare systems, online and mobile platforms, or

transcription services to ensure privacy for individuals with dementia. Our approach promotes ethical data handling and privacy-preserving communication.

For instance, Android’s SpeechRecognizer [20] is used by 95% (1306/1380) of apps with speech-to-text functionality (see Section 2). Google or other Original Equipment Manufacturers (OEMs) can modify SpeechRecognizer to return obfuscated transcriptions to the requesting third-party apps by default. The DiDOTS obfuscation model can be applied on the server side where the current Automatic Speech Recognition (ASR) system is run, to obfuscate the ASR outputs before returned to the mobile OS, or integrated into Android’s middleware speech framework whenever the ASR results are ready, to handle results from both remote and on-device system ASR models.

Even though DiDOTS ensures high semantics preservation, it might be beneficial for a minimal set of trusted apps to completely disable the privacy protection and allow full utility (raw text processing). To enable this, the OS provider could follow two strategies: (a) using Security-Enhanced Linux (SELinux) in Android to enforce mandatory access control policies allowing only trusted system apps to invoke APIs accessing original transcriptions, or (b) introducing a *runtime* permission for accessing raw transcriptions which requires user interaction to be granted. The former approach makes all third-party apps untrusted by default and will be denied access to raw transcriptions. A more flexible SELinux in Android is introduced by Demetriou et al. [17] which allows users to create discretionary versions of the policy.

The latter allows user control over which apps receive raw text access as apps will have to explicitly request the new permission. This will require populating two versions of the SpeechRecognizer API, one for permission-protected unsafe access and one for default safe (or obfuscated) access which is not permission-protected. Note that this may decrease user attention to permission prompts due to Android’s already high number of permissions [29]. There is currently 1215 system permissions on Android version 14 (the latest stable version) out of which 120 are runtime permissions.¹⁰

Some apps may perform transcriptions directly on raw audio, on-device or using third-party libraries (5.9% or 82/1380 in our study 2). If the transcriptions are performed on-device or known libraries are used, app markets (e.g. GooglePlay, Samsung Galaxy Store, Amazon AppStore etc.) could block or modify such apps to use DiDOTS obfuscated transcriptions.

A more challenging case is when an app captures and sends the audio to a remote destination for processing. In this case, OS and market providers cannot have access to the ASR output (the text transcriptions) and therefore an orthogonal technique to DiDOTS, which is a *text* obfuscation mechanism, is needed. Tackling threats from such apps requires a novel and careful design to ensure real-time performance. A possible approach would be to obfuscate the target attribute directly on the audio signal. For example, Woszczyk et al [110] proposed promising techniques for disentangling identity information from dementia in audio and future work could focus on adapting those techniques for preserving semantics while disentangling and obfuscating dementia features.

¹⁰To count the permissions we used the android debug bridge (adb) command line tool and the command `'adb shell pm list permissions -g -d | awk -F: '/permission:/ print $2' | wc -l`'.

Limitations. In this work, we perform an in-depth analysis of LLMs’ ability to obfuscate sensitive attributes like dementia and demonstrate the feasibility of distilling their knowledge into a smaller, fine-tuned student model. While our main focus is on dementia, our framework can extend to other attributes such as age, emotion, authorship, and gender. our PEFT approach allows for the easy sharing of adapter weights for each attribute, depending on the specific use case. However, due to the difficulty of accessing dementia datasets, our work relies on only two datasets with the same task which restricts the generalizability of our findings. The data may not fully capture the diverse characteristics and variations of spoken dementia, potentially impacting the effectiveness and robustness of the proposed approach. Despite this, we identify key characteristics from the literature and show LLMs and our distilled model’s ability to reduce dementia-like features. Our current evaluation also relies on the choice of adversaries given current detection capabilities and does not provide formal guarantees. We attend to this issue by considering several adversaries including traditional and neural classifiers in static and adaptive settings and show the effectiveness of our obfuscation strategy. In our work, we focus on sentence-level obfuscation and we do not account for intra-sentence dependencies and coherence issues. Future research could explore incorporating such dependencies to enhance obfuscation effectiveness and minimize attribute leakage risk.

8 Conclusion

In this work, we highlight the risk of leaking a sensitive medical user attribute from transcribed speech and leverage LLMs for privacy-preserving text transformations. We focused our analysis on dementia, a neurocognitive condition affecting tens of millions of people globally and to the best of our knowledge, we are the first to attempt to obfuscate this very sensitive attribute in speech. First, we showed that transcriptions can already be accessed by a large proportion (~20%) of the most popular mobile apps, a functionality often not properly described to users. Then, we investigated the capabilities of LLMs on the task of dementia obfuscation through various prompting strategies and found that zero-shot and domain-knowledge-informed prompts outperform compared to state-of-the-art (SOTA) text paraphrasing and attribute obfuscation models. Finally, we present DiDOTS, which distills knowledge from LLMs. DiDOTS is a significantly smaller student model compared to its teacher LLM model. Moreover, DiDOTS is partially fine-tuned for dementia obfuscation on a synthetic LLM-generated dataset with LoRA, which allows us to train and share three orders of magnitude fewer parameters compared to full fine-tuning. Our comprehensive evaluation demonstrated that DiDOTS retains the high obfuscation capabilities despite its smaller size and parameter-efficient fine-tuning. Our work highlights the capabilities of LLMs and knowledge distillation for attribute obfuscation in text and their potential for deploying effective and resource-efficient privacy-preserving text processing systems.

Acknowledgments

This study was funded by the Department of Computing at Imperial College London.

References

- [1] Ahmed Abbasi and Hsinchun Chen. 2008. Writeprints: A stylometric approach to identity-level identification and similarity detection in cyberspace. *ACM Transactions on Information Systems (TOIS)* 26, 2 (2008), 1–29.
- [2] Marah Abidin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219* (2024).
- [3] Muhammad Abdul-Mageed, Chiyu Zhang, Arun Rajendran, AbdelRahim Elmadany, Michael Przystupa, and Lyle Ungar. 2019. Sentence-Level BERT and Multi-Task Learning of Age and Gender in Social Media.
- [4] ai.meta.com. 2024. Our responsible approach to Meta AI and Meta Llama 3. <https://ai.meta.com/blog/meta-llama-3-meta-ai-responsibility/>
- [5] Yi Ji Bae, Midan Shim, and Won Hee Lee. 2021. Schizophrenia detection using machine learning approach from social media content. *Sensors* 21, 17 (2021), 5924.
- [6] Elron Bandel, Ranit Aharonov, Michal Shmueli-Scheuer, Ilya Shnayderman, Noam Slonim, and Liat Ein-Dor. 2022. Quality Controlled Paraphrase Generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Dublin, Ireland, 596–609. <https://doi.org/10.18653/v1/2022.acl-long.45>
- [7] James T Becker, François Boiler, Oscar L Lopez, Judith Saxton, and Karen L McGonigle. 1994. The natural history of Alzheimer’s disease: description of study cohort and accuracy of diagnosis. *Archives of neurology* 51, 6 (1994), 585–594.
- [8] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the association for computational linguistics* 5 (2017), 135–146.
- [9] 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 (2020), 1877–1901.
- [10] Nitay Calderon, Subhabrata Mukherjee, Roi Reichart, and Amir Kantor. 2023. A systematic study of knowledge distillation for natural language generation with pseudo-target training. *arXiv preprint arXiv:2305.02031* (2023).
- [11] Daniel Castro-Castro, Reynier Ortega Bueno, and Rafael Munoz. 2017. Author Masking by Sentence Transformation. In *CLEF (Working Notes)*.
- [12] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research* 25, 70 (2024), 1–53.
- [13] Arman Cohan, Bart Desmet, Andrew Yates, Luca Soldaini, Sean MacAvaney, and Nazli Goharian. 2018. SMHD: a Large-Scale Resource for Exploring Online Language Usage for Multiple Mental Health Conditions. In *Proceedings of the 27th International Conference on Computational Linguistics*. Association for Computational Linguistics, Santa Fe, New Mexico, USA, 1485–1497.
- [14] commission.europa.eu. 2016. What personal data is considered sensitive? <https://tinyurl.com/mptu73ac>
- [15] cppa.ca.gov. 2023. Laws & Regulations. <https://cppa.ca.gov/regulations/>
- [16] Sofia de la Fuente Garcia, Craig W Ritchie, and Saturnino Luz. 2020. Artificial intelligence, speech, and language processing approaches to monitoring Alzheimer’s disease: a systematic review. *Journal of Alzheimer’s Disease* 78, 4 (2020), 1547–1574.
- [17] Soteris Demetriou, Xiao-yong Zhou, Muhammad Naveed, Yeonjoon Lee, Kan Yuan, XiaoFeng Wang, and Carl A. Gunter. 2015. What’s in Your Dongle and Bank Account? Mandatory and Discretionary Protection of Android External Resources. In *22nd Annual Network and Distributed System Security Symposium, NDSS 2015, San Diego, California, USA, February 8-11, 2015*. The Internet Society. <https://www.ndss-symposium.org/ndss2015/whats-your-dongle-and-bank-account-mandatory-and-discretionary-protection-android-external>
- [18] Biniam Fisseha Demissie, Mariano Ceccato, and Lwin Khin Shar. 2020. Security analysis of permission re-delegation vulnerabilities in Android apps. *Empirical Software Engineering* 25 (2020), 5084–5136.
- [19] A. Developers. 2024. Create a search interface | Android Developers. <https://developer.android.com/develop/ui/views/search/search-dialog#VoiceSearch>
- [20] Android Developers. 2024. SpeechRecognizer. <https://tinyurl.com/yc5m9vht>
- [21] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *North American Chapter of the Association for Computational Linguistics*. <https://api.semanticscholar.org/CorpusID:52967399>
- [22] Benjamin D Douglas, Patrick J Ewell, and Markus Brauer. 2023. Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA. *Plos one* 18, 3 (2023), e0279720.
- [23] T. Edwards and Nicholas S. Holtzman. 2017. A meta-analysis of correlations between depression and first person singular pronoun use. *Journal of Research in Personality* 68 (2017), 63–68.
- [24] Chris Emmery, Enrique Manjavacas, and Grzegorz Chrupala. 2018. Style obfuscation by invariance. *arXiv preprint arXiv:1805.07143* (2018).
- [25] Maxim Enis and Mark Hopkins. 2024. From LLM to NMT: Advancing Low-Resource Machine Translation with Claude. *arXiv preprint arXiv:2404.13813* (2024).
- [26] enlyft.com. 2024. Companies using IBM Watson Speech to Text. <https://tinyurl.com/msvuhrrh>
- [27] eur.lex.europa.eu. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). <https://tinyurl.com/6tk3j9aw>
- [28] Mael 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)*. NLP Association of India (NLPAI), Indian Institute of Technology Patna, Patna, India, 127–137.
- [29] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. 2012. Android permissions: user attention, comprehension, and behavior. In *Proceedings of the Eighth Symposium on Usable Privacy and Security (Washington, D.C.) (SOUPS ’12)*. Association for Computing Machinery, New York, NY, USA, Article 3, 14 pages. <https://doi.org/10.1145/2335356.2335360>
- [30] Adrienne Porter Felt, Helen J Wang, Alexander Moshchuk, Steve Hanna, and Erika Chin. 2011. Permission re-delegation: Attacks and defenses.. In *USENIX security symposium*, Vol. 30. 88.
- [31] Kathleen C Fraser, Jed A Meltzer, and Frank Rudzicz. 2016. Linguistic features identify Alzheimer’s disease in narrative speech. *Journal of Alzheimer’s Disease* 49, 2 (2016), 407–422.
- [32] Samira Ghodrattama. 2021. Towards personalized and human-in-the-loop document summarization. *arXiv preprint arXiv:2108.09443* (2021).
- [33] Samuel Gibbs. 2023. Samsung keyboard bug leaves 600m Android devices exposed to hackers. <https://tinyurl.com/yccke6rtz>
- [34] Michael C Grace, Wu Zhou, Xuxian Jiang, and Ahmad-Reza Sadeghi. 2012. Unsafe exposure analysis of mobile in-app advertisements. In *Proceedings of the fifth ACM conference on Security and Privacy in Wireless and Mobile Networks*. 101–112.
- [35] grandviewresearch.com. 2024. Speech-to-text API Market Size, Share Analysis Report, 2030. <https://tinyurl.com/muncm7hc>
- [36] Tommi Gröndahl and N Asokan. 2020. Effective writing style transfer via combinatorial paraphrasing. *Proc. Priv. Enhancing Technol.* 2020, 4 (2020), 175–195.
- [37] Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. 2023. MiniLLM: Knowledge distillation of large language models. In *The Twelfth International Conference on Learning Representations*.
- [38] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [39] Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. *arXiv preprint arXiv:2305.02301* (2023).
- [40] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021).
- [41] James Y Huang, Kuan-Hao Huang, and Kai-Wei Chang. 2021. Disentangling semantics and syntax in sentence embeddings with pre-trained language models. *arXiv preprint arXiv:2104.05115* (2021).
- [42] Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. Adversarial example generation with syntactically controlled paraphrase networks. *arXiv preprint arXiv:1804.06059* (2018).
- [43] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825* (2023).
- [44] Yi Jiang, Chenghui Shi, Oubo Ma, Youliang Tian, and Shouling Ji. 2023. Text Laundering: Mitigating Malicious Features Through Knowledge Distillation of Large Foundation Models. In *International Conference on Information Security and Cryptology*. Springer, 3–23.
- [45] Jaehun Jung, Peter West, Liwei Jiang, Faeze Brahman, Ximing Lu, Jillian Fisher, Taylor Sorensen, and Yejin Choi. 2023. Impossible distillation: from low-quality model to high-quality dataset & model for summarization and paraphrasing. *arXiv preprint arXiv:2305.16635* (2023).
- [46] Georgi Karadzhov, Tsvetomila Mihaylova, Yassen Kiproff, Georgi Georgiev, Ivan Koychev, and Preslav Nakov. 2017. The Case for Being Average: A Mediocrity Approach to Style Masking and Author Obfuscation: (Best of the Labs Track at CLEF-2017). In *Experimental IR Meets Multilinguality, Multimodality, and Interaction: 8th International Conference of the CLEF Association, CLEF 2017, Dublin, Ireland, September 11–14, 2017, Proceedings* 8. Springer, 173–185.

- [47] Daniel Kempler and Mira Goral. 2008. LANGUAGE AND DEMENTIA: NEUROPSYCHOLOGICAL ASPECTS. *Annual Review of Applied Linguistics* 28 (2008), 73–90.
- [48] Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. CTRL: A Conditional Transformer Language Model for Controllable Generation. *ArXiv abs/1909.05858* (2019).
- [49] Yoon Kim and Alexander M Rush. 2016. Sequence-level knowledge distillation. *arXiv preprint arXiv:1606.07947* (2016).
- [50] Moshe Koppel, Shlomo Argamon, and Anat Rachel Shimoni. 2002. Automatically categorizing written texts by author gender. *Literary and linguistic computing* 17, 4 (2002), 401–412.
- [51] Sachin Kumar, Eric Malmi, Aliaksei Severyn, and Yulia Tsvetkov. 2021. Controlled text generation as continuous optimization with multiple constraints. *Advances in Neural Information Processing Systems* 34 (2021), 14542–14554.
- [52] Ian Lancashire and Graeme Hirst. 2009. Vocabulary changes in Agatha Christie's mysteries as an indication of dementia: a case study. In *19th Annual Rotman Research Institute Conference, Cognitive Aging: Research and Practice*. 8–10.
- [53] Alon Lavie and Michael J Denkowski. 2009. The METEOR metric for automatic evaluation of machine translation. *Machine translation* 23 (2009), 105–115.
- [54] Xuan Le, Ian Lancashire, Graeme Hirst, and Regina Jokel. 2011. Longitudinal detection of dementia through lexical and syntactic changes in writing: a case study of three British novelists. *Literary and linguistic computing* 26, 4 (2011), 435–461.
- [55] Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691* (2021).
- [56] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 7871–7880. <https://doi.org/10.18653/v1/2020.acl-main.703>
- [57] Zhe Lin and Xiaojun Wan. 2021. Pushing Paraphrase Away from Original Sentence: A Multi-Round Paraphrase Generation Approach. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. Association for Computational Linguistics, Online, 1548–1557. <https://doi.org/10.18653/v1/2021.findings-acl.135>
- [58] Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Lam Tam, Zhengxiao Du, Zhilin Yang, and Ji Tang. 2021. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602* (2021).
- [59] Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Zhifang Sui, and Xu Sun. 2019. A dual reinforcement learning framework for unsupervised text style transfer. *arXiv preprint arXiv:1905.10060* (2019).
- [60] Saturnino Luz, Fasih Haider, Sofia de la Fuente, Davida Fromm, and Brian MacWhinney. 2020. Alzheimer's dementia recognition through spontaneous speech: The address challenge. *arXiv preprint arXiv:2004.06833* (2020).
- [61] Saturnino Luz, Fasih Haider, Sofia de la Fuente, Davida Fromm, and Brian MacWhinney. 2021. Detecting cognitive decline using speech only: The ADReSSo Challenge. *arXiv preprint arXiv:2104.09356* (2021).
- [62] Brian MacWhinney. 2004. TalkBank: Building an open unified multimodal database of communicative interaction. (2004).
- [63] Asad Mahmood, Faizan Ahmad, Zubair Shafiq, Padmini Srinivasan, and Fareed Zaffar. 2019. A Girl Has No Name: Automated Authorship Obfuscation using Mutant-X. *Proc. Priv. Enhancing Technol.* 2019, 4 (2019), 54–71.
- [64] Justus Mattern, Benjamin Weggenmann, and Florian Kerschbaum. 2022. The Limits of Word Level Differential Privacy. In *Findings of the Association for Computational Linguistics: NAACL 2022*. Association for Computational Linguistics, Seattle, United States, 867–881. <https://doi.org/10.18653/v1/2022.findings-naacl.65>
- [65] Sachin Mehta, Mohammad Hossein Sekhavat, Qingqing Cao, Maxwell Horton, Yanzi Jin, Chenfan Sun, Iman Mirzadeh, Mahyar Najibi, Dmitry Belenko, Peter Zatloukal, et al. 2024. OpenELM: An Efficient Language Model Family with Open-source Training and Inference Framework. *arXiv preprint arXiv:2404.14619* (2024).
- [66] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [67] Fatemehsadat Mireshghallah and Taylor Berg-Kirkpatrick. 2021. Style pooling: Automatic text style obfuscation for improved classification fairness. *arXiv preprint arXiv:2109.04624* (2021).
- [68] Fatemehsadat Mireshghallah, Kartik Goyal, and Taylor Berg-Kirkpatrick. 2022. Mix and Match: Learning-free Controllable Text Generation using Energy Language Models. *arXiv preprint arXiv:2203.13299* (2022).
- [69] mordorintelligence.com. 2024. Voice Recognition Market - Size, Industry Analysis & Companies. <https://tinyurl.com/2re34u4x>
- [70] Dong Nguyen, Rilana Gravel, Dolf Trieschnigg, and Theo Meder. 2013. "How old do you think I am?" A study of language and age in Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 7. 439–448.
- [71] Emma Nichols, Jaimie D Steinmetz, Stein Emil Vollset, Kai Fukutaki, Julian Chalek, Foad Abd-Allah, Amir Abdoli, Ahmed Abualhasan, Eman Abu-Gharbieh, Tayyaba Tayyaba Akram, et al. 2022. Estimation of the global prevalence of dementia in 2019 and forecasted prevalence in 2050: an analysis for the Global Burden of Disease Study 2019. *The Lancet Public Health* 7, 2 (2022), e105–e125.
- [72] ollama.com. 2024. Llama 3 is not very censored - Ollama Blog. <https://ollama.com/blog/llama-3-is-not-very-censored>.
- [73] OpenAI. 2024. Moderation - OpenAI API. <https://platform.openai.com/docs/guides/moderation/overview>
- [74] OpenAI. 2024. OpenAI Red Teaming Network | OpenAI. <https://openai.com/index/red-teaming-network/>
- [75] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems* 35 (2022), 27730–27744.
- [76] Rahul Pandita, Xusheng Xiao, Wei Yang, William Enck, and Tao Xie. 2013. {WHYPER}: Towards automating risk assessment of mobile applications. In *22nd USENIX Security Symposium (USENIX Security 13)*. 527–542.
- [77] Ellie Pavlick and Joel Tetreault. 2016. An empirical analysis of formality in online communication. *Transactions of the Association for Computational Linguistics* 4 (2016), 61–74.
- [78] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [79] Minh Quang Pham, Sathish Reddy Indurthi, Shamil Chollampatt, and Marco Turchi. 2023. Select, Prompt, Filter: Distilling Large Language Models for Summarizing Conversations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 12257–12265.
- [80] Atharva Phatak, David W Savage, Robert Ohle, Jonathan Smith, and Vijay Mago. 2022. Medical text simplification using reinforcement learning (teslea): Deep learning-based text simplification approach. *JMIR Medical Informatics* 10, 11 (2022), e38095.
- [81] play.google.com. 2024. Android Apps on Google Play. <https://play.google.com/store/apps/>
- [82] play.google.com. 2024. Google Keep - Android Apps on Google Play. <https://tinyurl.com/nujrb7sj>
- [83] play.google.com. 2024. Otter: Transcribe Voice Notes - Android Apps on Google Play. <https://tinyurl.com/mvac5auz>
- [84] play.google.com. 2024. Voice Recorder Audio Sound MP3 - Apps on Google Play. <https://tinyurl.com/4pepb3dy>
- [85] Shrimai Prabhunoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. 2018. Style transfer through back-translation. *arXiv preprint arXiv:1804.09000* (2018).
- [86] Prolific. 2024. Prolific. <https://researcher-help.prolific.com/hc/en-gb>
- [87] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*. PMLR, 28492–28518.
- [88] Sudha Rao and Joel Tetreault. 2018. Dear Sir or Madam, May I Introduce the GYAFC Dataset: Corpus, Benchmarks and Metrics for Formality Style Transfer. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Association for Computational Linguistics, New Orleans, Louisiana, 129–140. <https://doi.org/10.18653/v1/N18-1012>
- [89] Reddit. 2024. I didn't think it was possible to have a more censored AI until I tried Llama 2 : r/ChatGPT. https://www.reddit.com/r/ChatGPT/comments/1560r1z/i_didnt_think_it_was_possible_to_have_a_more/.
- [90] researcher help.prolific.com. 2024. Data protection and privacy - Prolific. <https://researcher-help.prolific.com/hc/en-gb/articles/360009094594-Data-protection-and-privacy>.
- [91] Max Reuter and William Schulze. 2023. I'm Afraid I Can't Do That: Predicting Prompt Refusal in Black-Box Generative Language Models. *arXiv preprint arXiv:2306.03423* (2023).
- [92] Ruchita Sarawgi, Kailash Gajulapalli, and Yejin Choi. 2011. Gender attribution: tracing stylistic evidence beyond topic and genre. In *Proceedings of the fifteenth conference on computational natural language learning*. 78–86.
- [93] Clemens JS Schaefer, Pooria Taheri, Mark Horeni, and Siddharth Joshi. 2023. The hardware impact of quantization and pruning for weights in spiking neural networks. *IEEE Transactions on Circuits and Systems II: Express Briefs* (2023).
- [94] Jonathan Schler, Moshe Koppel, Shlomo Argamon, and James W Pennebaker. 2006. Effects of age and gender on blogging. In *AAAI spring symposium: Computational approaches to analyzing weblogs*, Vol. 6. 199–205.
- [95] Rakshith Shetty, Bernt Schiele, and Mario Fritz. 2018. A4nt: author attribute anonymity by adversarial training of neural machine translation. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*. 1633–1650.
- [96] Sam Shleifer and Alexander M Rush. 2020. Pre-trained summarization distillation. *arXiv preprint arXiv:2010.13002* (2020).
- [97] statista.com. 2023. Number of smartphone users worldwide from 2013 to 2028. <https://tinyurl.com/232hfdfs>

- [98] statista.com. 2023. Smart speakers - statistics & facts. <https://tinyurl.com/bdf3w7rs>
- [99] statista.com. 2024. Global market share held by mobile operating systems from 2009 to 2023, by quarter. <https://tinyurl.com/yc2xdpss>
- [100] Chi Sun, Luyao Huang, and Xipeng Qiu. 2019. Utilizing BERT for Aspect-Based Sentiment Analysis via Constructing Auxiliary Sentence. 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)*. Association for Computational Linguistics, Minneapolis, Minnesota, 380–385. <https://doi.org/10.18653/v1/N19-1035>
- [101] Zafi Sherhan Syed, Muhammad Shehram Shah Syed, Margaret Lech, and Elena Pirogova. 2021. Automated recognition of alzheimer’s dementia using bag-of-deep-features and model ensembling. *IEEE Access* 9 (2021), 88377–88390.
- [102] Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhatipatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295* (2024).
- [103] this.nhs.uk. 2023. How voice recognition technology is benefitting NHS staff and patients. <https://tinyurl.com/y7vynnha>
- [104] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
- [105] Guliz Seray Tuncay, Soteris Demetriou, and Carl A Gunter. 2016. Draco: A system for uniform and fine-grained access control for web code on android. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 104–115.
- [106] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [107] Bo Wang, Yue Wu, Niall Taylor, Terry Lyons, Maria Liakata, Alejo J. Nevado-Holgado, and Kate E. A. Saunders. 2020. Learning to Detect Bipolar Disorder and Borderline Personality Disorder with Language and Speech in Non-Clinical Interviews. In *Interspeech*.
- [108] 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*. 721–731.
- [109] John Wieting and Kevin Gimpel. 2017. ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations. *arXiv preprint arXiv:1711.05732* (2017).
- [110] Dominika Woszczyk, Ranya Aloufi, and Soteris Demetriou. 2024. Prosody-Driven Privacy-Preserving Dementia Detection. In *Interspeech*, Vol. 2024.
- [111] Dominika Woszczyk, Anna Hedlikova, Alican Akman, Soteris Demetriou, and Björn Schuller. 2022. Data Augmentation for Dementia Detection in Spoken Language. In *Proc. Interspeech 2022*. 2858–2862. <https://doi.org/10.21437/Interspeech.2022-10210>
- [112] Qionkai Xu, Lizhen Qu, Chenchen Xu, and Ran Cui. 2019. Privacy-aware text rewriting. In *Proceedings of the 12th International Conference on Natural Language Generation*. 247–257.
- [113] Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics* 4 (2016), 401–415.
- [114] Yichong Xu, Ruochen Xu, Dan Iter, Yang Liu, Shuohang Wang, Chenguang Zhu, and Michael Zeng. 2023. Inheritsum: A general, versatile and compact summarizer by distilling from gpt. *arXiv preprint arXiv:2305.13083* (2023).
- [115] Jiahong Yuan, Yuchen Bian, Xingyu Cai, Jiayi Huang, Zheng Ye, and Kenneth Church. 2020. Disfluencies and Fine-Tuning Pre-Trained Language Models for Detection of Alzheimer’s Disease. In *Interspeech*, Vol. 2020. 2162–6.
- [116] Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*. PMLR, 11328–11339.
- [117] Maryam Zolnoori, Ali Zolnour, and Maxim Topaz. 2023. ADscreen: A speech processing-based screening system for automatic identification of patients with Alzheimer’s disease and related dementia. *Artificial Intelligence in Medicine* (2023), 102624.

A Android Motivation

Table 11 shows the list of keywords, third-party services, popular APIs, and SDK calls used in the study of Android apps to search for the presence of speech-to-text activity in Android apps code. Note, that it is a conservative list.

Table 11: List of Keywords, APIs, and SDKs used for searching for speech-to-text processing within Android OS apps (transposed).

Category	Items
APIs	api.deepgram.com/v1/listen speech-to-text.watson.cloud.ibm.com api.openai.com/v1/audio/ api.assemblyai.com speech.googleapis.com asr.api.speechmatics.com apis.voicebase.com api.cognitive.microsoft.com/speechoctotext libdeepspeech asr.api.nuance.com api.amberscript.com api.rev.com api.scriptix.io
SDKs	android.speech.action.recognize_speech org.tensorflow.lite.task.text.nlclassifier com.assemblyai pocketsphinx com.google.cloud.speech.v1 android.media.AudioRecord software.amazon.awssdk:transcribestreaming ninamobilecontroller com.microsoft.cognitiveservices.speech android.speech.recognizerintent org.vosk revai.revaistreamingclient com.ibm.watson.developer_cloud com.github.mozilla:mozillaspeechlibrary
Keywords	speechoctotext speech-to-text speech_to_text speechlistener speech_listener transcribe transcription speechrecognition speech_recognition

B DP-Pegasus

Table 12 presents the privacy/utility performance of the DP system based on the Pegasus model with various sampling temperatures. The temperature seems to have little effect on the obfuscation capabilities with a constant mean F1 score of circa 0.57,0.56 for ADReSS (ADR) and ADReSSo (ADRo) datasets while the semantic score is dropping down to 0.51,0.47 for ADR, ADRo respectively. We find that the temperature of $T=1.5$ achieves the best trade-off across both datasets and select it for our main comparison in Section 5.

Table 12: Privacy/utility results for ADReSS and ADReSSo datasets at sentence-level for DP Pegasus with various temperatures. This table presents the F1 scores against static and adaptive dementia classifiers and semantic preservation. Best result for each column is identified in bold.

	Temp.	SVM↓		BERT↓		Avg. F1↓	Semantics↑
		Static	Adaptive	Static	Adaptive		
ADReSS	1.0	0.60	0.59	0.52	0.58	0.57	0.85
	1.5	0.59	0.56	0.51	0.58	0.56	0.85
	2.0	0.60	0.57	0.52	0.58	0.57	0.82
	3.0	0.58	0.59	0.52	0.58	0.57	0.72
	5.0	0.57	0.57	0.55	0.58	0.57	0.58
	10.0	0.61	0.55	0.63	0.58	0.59	0.51
ADReSSo	1.0	0.58	0.58	0.50	0.57	0.56	0.81
	1.5	0.58	0.59	0.49	0.57	0.56	0.80
	2.0	0.58	0.61	0.47	0.57	0.56	0.77
	3.0	0.58	0.57	0.50	0.57	0.56	0.68
	5.0	0.58	0.58	0.56	0.57	0.57	0.56
	10.0	0.61	0.55	0.57	0.57	0.58	0.47

C LLMs Samples

In Table 13, we showcase samples generated by our selected LLMs. We see that they all generate fluent paraphrases and semantically close while also being diverse. We can note they tend to increase the formality of the language except for LLama3 which remains quite unformal.

Table 13: Comparison of Dementia (AD) and Control (CC) Samples from different LLMs.

System	Sample AD	Sample CC
Original	and uh they're each busy at at uh something.	and there's a tree out there somewhere that you can see part of the trunk and part of the foliage.
Gemma	and they were each busy at something.	There is a large tree standing in the distance that you can see some of its branches and leaves.
Phi3	They are all occupied with their respective tasks.	and there's a tree outside where we can catch a glimpse of its trunk and leaves.
Mistral	And they're all occupied with their own tasks.	Outside, you can see a tree with some of its trunk and leaves visible.
LLama3	and each person is occupied with their own activity.	you can see part of the trunk and some leaves on that tree.

D Human Evaluation: Ethical Considerations

The reason for conducting our study with human subjects is to evaluate the effectiveness of our text obfuscation model in preserving semantic content and comparing it to relevant work. As automatic metrics do not accurately represent semantics, it is crucial to perform human evaluation on a significant sample size. Our research question focuses on assessing our model's ability to maintain the original text's semantics comparatively with other existing paraphrasing methods, employing human perception as the metric. Our study was approved by our institution's Research Ethics review board.

Task design: We design a questionnaire presenting an original sentence and its paraphrase with different systems. We consider 5 systems: ParChoice (rule-based), Pegasus (Paraphrase), DP (Differential Privacy-based), and our system, Mistral (knowledge-guided strategy) and DiDOTS. Participants were be asked to rate the paraphrase of each system on a 1 to 5 Likert scale. We randomise the questions order as well as the system's position within the questions, to mitigate the order bias that might affect how the respondent answers. We recruit 100 participants who are English natives from the UK or US, and recruit an additional 60 participants to correct for low-quality respondents.

Platform & Recruitment Methods: We create the survey on Qualtrics and advertise it on Prolific. Prolific [86] is increasingly used in academic studies as it gives researchers access to a pool of quality participants (passing attention checks, giving meaningful answers, following instructions) [22]. The participants are given a link to a full participant information sheet which contains the study description, contact details and a full transparency notice and asked for consent. Each user is anonymized and accesses the survey online [90]. The platform allows us to add several filters to further ensure the relevance of the participants. Eligible participants are sent the survey and a preview on the platform (with the compensation and estimated duration) and they can choose to take it or ignore it. Each participant will be offered compensation for their work. The preview will read as follows:

"The purpose of this study is to investigate the effectiveness of various techniques that modify text in different ways but can preserve the meaning of the original text. We ask you to evaluate passages, rating the degree of similarity between original and paraphrased texts on a scale of 1 to 5."

Consent: At the beginning of the survey participants will be asked to opt-in to the survey by agreeing to a consent form shown after being presented with a short summary of the participant information sheet. The summary will have a link to the full participant information sheet which contains the full transparency notice.

Withdrawal of participation: The participants are free to stop the survey at any moment if they do not wish to complete it. After the survey, they may choose to "return" the survey on the Prolific platform. The data of participants with incomplete surveys are not used in the study. In case of withdrawal, the participant will not receive compensation and their response will be deleted from Qualtrics and local records.

Anonymity: Each user is anonymized and accesses the survey online [90]. We do not collect identifiable data such as location, names, or age. Although the Prolific ID has not been linked to any de-anonymization attempt we do treat it as a PII and do not store the Prolific identifiers. We only use the Prolific ID to reject or accept the completion of the survey and offer compensation. This ID will only be accessed by researchers working on the project. For storing the data, we substitute each prolific ID with a new ID we randomly generate.