


Confronting LLMs with Traditional ML: Rethinking the Fairness of Large Language Models in Tabular Classifications

Yanchen Liu  Srishti Gautam  Jiaqi Ma  Himabindu Lakkaraju 

 Harvard University  UiT The Arctic University of Norway

 University of Illinois Urbana-Champaign

yanchenliu@g.harvard.edu, srishti.gautam@uit.no,
jiaqima@illinois.edu, hlakkaraju@hbs.edu

Abstract

Recent literature has suggested the potential of using large language models (LLMs) to make classifications for tabular tasks. However, LLMs have been shown to exhibit harmful social biases that reflect the stereotypes and inequalities present in society. To this end, as well as the widespread use of tabular data in many high-stake applications, it is important to explore the following questions: what sources of information do LLMs draw upon when making classifications for tabular tasks; whether and to what extent are LLM classifications for tabular data influenced by social biases and stereotypes; and what are the consequential implications for fairness?

Through a series of experiments, we delve into these questions and show that LLMs tend to inherit social biases from their training data which significantly impact their fairness in tabular classification tasks. Furthermore, our investigations show that in the context of bias mitigation, though in-context learning and finetuning have a moderate effect, the fairness metric gap between different subgroups is still larger than that in traditional machine learning models, such as Random Forest and shallow Neural Networks. This observation emphasizes that the social biases are inherent within the LLMs themselves and inherited from their pre-training corpus, not only from the downstream task datasets. Besides, we demonstrate that label-flipping of in-context examples can significantly reduce biases, further highlighting the presence of inherent bias within LLMs.

1 Introduction

Many recent works propose to use large language models (LLMs) for tabular tasks (Slack and Singh, 2023; Hagselmann et al., 2023), where the tabular data is serialized as natural language and provided to LLMs with a short description of the task to solicit predictions. Despite the comprehensive examination of fairness considerations within con-

ventional machine learning approaches applied to tabular tasks (Bellamy et al., 2018), the exploration of fairness-related issues in the context of employing LLMs for tabular classifications remains a relatively underexplored domain.

Previous research has shown that LLMs, such as GPT-3 (Brown et al., 2020), GPT-3.5, GPT-4 (OpenAI, 2023) can exhibit harmful social biases (Abid et al., 2021a; Basta et al., 2019), which may even worsen as the models become larger in size (Askell et al., 2021; Ganguli et al., 2022). These biases are a result of the models being trained on text generated by humans that presumably includes many examples of humans exhibiting harmful stereotypes and discrimination and reflects the biases and inequalities present in society (Bolukbasi et al., 2016; Zhao et al., 2017), which can lead to the perpetuation of discrimination and stereotype (Abid et al., 2021a; Bender et al., 2021).

Considering that tabular data finds extensive use in high-stakes domains (Grinsztajn et al., 2022) where information is typically structured in tabular formats as a natural byproduct of relational databases (Borisov et al., 2022), it is crucial to thoroughly examine the fairness implications of utilizing LLMs for classifications on tabular data. In this paper, we conduct a series of investigations centered around this critical problem, with the goal of discerning the underlying information sources upon which LLMs rely when making tabular classifications. Through this exploration, our investigation aims to ascertain whether, and to what degree, LLMs are susceptible to being influenced by social biases and stereotypes in the context of tabular data classifications.

Through experiments using GPT-3.5 to make classifications for tabular data in a zero-shot setting, we demonstrate that LLMs exhibit significant social biases (Section 4). This evidence confirms that LLMs inherit social biases from their pretraining corpus and tend to rely on these biases when

making classifications for tabular data.

Furthermore, we demonstrate that providing GPT-3.5 with few-shot examples (in-context learning) or finetuning them on the entire training datasets both exhibit moderate effects on bias mitigation (Sections 5.1 and 6.1). Nevertheless, the achieved fairness levels remain below what is typically attained with traditional machine learning methods, including Random Forests and shallow Neural Networks, once again underscoring the presence of inherent bias in LLMs.

Moreover, our investigation further reveals that flipping the labels of the in-context few-shot examples significantly narrows the gap in fairness metrics across different subgroups, but comes at the expected cost of a reduction in the prediction performance. This finding, in turn, further emphasizes and reaffirms the indication of inherent bias present in LLMs (Section 5.2). Additionally, we further show that while resampling the training data is a known and effective method for reducing biases in traditional machine learning methods like Random Forests and shallow Neural Networks, it proves to be less effective when applied to GPT-3.5 (Section 6.2).

These collective findings demonstrate the significant influence of social biases on GPT-3.5's performance in tabular classifications. These biases significantly undermine fairness and pose potential risks for using LLMs on tabular data, especially considering that tabular data is extensively used in high-stakes domains, highlighting the need for more advanced and tailored strategies to address these biases effectively. Straightforward methods like in-context learning and data resampling may not be sufficient in this context.

2 Related work

Fairness and Social Biases in LLMs Fairness is highly desirable for ensuring the credibility and trustworthiness of algorithms. It has been demonstrated that unfair algorithms can reflect societal biases in their decision-making processes (Bender et al., 2021; Bommasani, 2021), primarily stemming from the biases present in their training data (Caliskan et al., 2017; Zhao et al., 2017). LLMs, pre-trained on vast natural language datasets, are particularly susceptible to inheriting these social biases and have been shown to exhibit biases related to gender (Lucy and Bamman, 2021), religion (Abid et al., 2021b) and language vari-

ants (Ziems et al., 2023; Liu et al., 2023a). These social biases can lead to the perpetuation of discrimination and stereotype (Abid et al., 2021a; Bender et al., 2021; Weidinger et al., 2021). While recent works have made strides in addressing these issues, there still exists a significant gap in comprehensively assessing fairness in LLMs and its mitigation strategies for tabular data.

Tabular Tasks and LLM for Tabular Data

Tabular data extensively exist in many domains (Shwartz-Ziv and Armon, 2021). Previous works propose to utilize self-supervised deep techniques for tabular tasks (Yin et al., 2020; Arik and Pfister, 2021), which, however, still underperform ensembles of gradient-boosted trees in the fully supervised setting (Grinsztajn et al., 2022). This disparity in performance can be attributed to the locality, sparsity, and mixed data types of tabular data. In recent times, LLMs have undergone intensive training using vast amounts of natural language data, which has enabled them to exhibit impressive performance across various downstream tasks (Brown et al., 2020; OpenAI, 2023), even with little or no labeled task data. Therefore, recent approaches by (Hegselmann et al., 2023; Slack and Singh, 2023) suggest serializing the tabular data as natural language, which is provided to LLM along with a short task description to generate predictions for tabular tasks.

However, tabular data plays a crucial role in numerous safety-critical and high-stakes domains (Borisov et al., 2022; Grinsztajn et al., 2022), which makes fairness particularly crucial when employing LLMs for making predictions on tabular data, especially considering the inherent social biases present in LLMs. Despite the importance, this still remains largely unexplored. To the best of our knowledge, we regard our work as one of the most comprehensive investigations into the fairness issues arising when using LLMs for classification on tabular data.

In-Context Learning Significant improvements for various tasks have been achieved by providing in-context examples to LLMs (Brown et al., 2020; Liu et al., 2022, 2023b). However, previous research by (Min et al., 2022; Wei et al., 2023b; Lyu et al., 2023) illustrate that the effective performance of in-context learning largely hinges on semantic priors rather than learning the input-label mapping (Akyürek et al., 2022; Xie et al., 2022; Von Os-

wald et al., 2023) and the labels of the in-context examples might not play a crucial role in in-context learning, with flipped or random labels sometimes having minimal impact on performance. Despite these findings, the predominant focus of existing investigation of in-context learning remains on conventional natural language processing tasks (Zhao et al., 2021; Min et al., 2022; Wei et al., 2023a,b), largely overlooking the domain of tabular data. Furthermore, the fairness of in-context learning and the impact of flipped labels on this fairness is yet to be thoroughly investigated.

3 Experimental Setup

In this section, we begin by presenting an overview of the experimental setup utilized throughout this work.

3.1 Models

In our work, we focus our experiments on GPT-3.5 (engine *GPT-3.5-turbo*) - an LLM released by OpenAI, trained with instruction tuning (Sanh et al., 2022; Wei et al., 2022) and reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022), aligning LLMs with human preferences.

Furthermore, we also compare this most common LLM with conventional machine learning models, in order to gain insight into the propagation of biases found within LLMs. This propagation of biases is likely mirrored in traditional models as well, consequently, offering valuable additional perspectives on the social biases inherent in the training of LLMs. For this, we employ two widely used traditional models for tabular data, i.e., Random Forests (RF) and the shallow Neural Network (NN) of 3 layers. We provide additional implementation details for these traditional models in Appendix C.

3.2 Datasets and Protected Attributes

To explore the fairness of LLMs in making classifications for tabular data, we utilize the following three widely used tabular datasets for assessing the fairness of traditional ML models: *Adult Income* (**Adult**) Dataset (Becker and Kohavi, 1996), *German Credit* Dataset (Dua and Graff, 2019), and *Correctional Offender Management Profiling for Alternative Sanctions* (**COMPAS**) Dataset (Larson et al., 2016). In this section, we give a brief introduction to each dataset and discuss its associated protected attributes.

Adult The *Adult Income* dataset (**Adult**) is extracted from the 1994 U.S. Census Bureau database. The task is to predict whether a person earns more than \$50,000 per year based on their profile data (*greater than 50K* or *less than or equal to 50K*). The original Adult Income Dataset contains 14 features. Following previous work (Slack and Singh, 2023), we retain only 10 features: “*workclass*”, “*hours per week*”, “*sex*”, “*age*”, “*occupation*”, “*capital loss*”, “*education*”, “*capital gain*”, “*marital status*”, and “*relationship*”. Our analysis on **Adult** primarily focuses on “*sex*” as the protected attribute, and *female* is acknowledged as a disadvantaged group.

German Credit The **German Credit** dataset is used to classify individuals based on their profile attributes as good or bad credit risks (*good* or *bad*). The raw dataset comprises 20 attributes. Consistent with previous work, we only retain the following features: “*age*”, “*sex*”, “*job*”, “*housing*”, “*saving accounts*”, “*checking account*”, “*credit amount*”, “*duration*”, and “*purpose*”. Same with **Adult**, “*sex*” is considered as a protected attribute in the **German Credit** dataset and *female* as the marginalized group.

COMPAS The **COMPAS** dataset comprises the outcomes from the *Correctional Offender Management Profiling for Alternative Sanctions* commercial algorithm, utilized to evaluate a convicted criminal’s probability of reoffending. Known for its widespread use by judges and parole officers, **COMPAS** has gained notoriety for its bias against African-Americans. The raw COMPAS Recidivism dataset contains more than 50 attributes. Following the approach of Larson et al. (2016), we perform necessary preprocessing, group “*race*” into *African-American* and *Not African-American*, and only consider the features “*sex*”, “*race*”, “*age*”, “*charge degree*”, “*priors count*”, “*risk*” and “*two year recid*” (target). We frame the task as predicting whether an individual will recidivate in two years (*Did Not Reoffend* or *Reoffended*) based on their demographic and criminal history. For the **COMPAS** dataset, we consider “*race*” as the protected attribute.

A detailed description for each feature of the considered datasets is provided in Appendix A.

3.3 Serialization and Prompt Templates

To employ the LLM for making classifications on these tabular datasets, each data point is first serial-

ized as text. Following previous works on LLM for tabular classifications (Hegselmann et al., 2023; Slack and Singh, 2023), we format the feature names and values into strings as “ $f_1 : x_1, \dots, f_d : x_d$ ”, and prompt to LLM along with a task description, as illustrated following:

Prompt 1. Prompt Template for Adult Dataset.

```
You must predict if income exceeds $50K/yr.
Answer with one of the following: greater
than 50K | less than or equal to 50K.
Example 1 -
workclass: Private
hours per week: 20
sex: Male
age: 17
occupation: Other-service
capital loss: 0
education: 10th
capital gain: 0
marital status: Never-married
relationship: Own-child
Answer: less than or equal to 50K
...
workclass: Private
hours per week: 40
sex: Female
age: 24
occupation: Sales
capital loss: 0
education: Some-college
capital gain: 0
marital status: Never-married
relationship: Own-child
Answer:
```

The example above is from the **Adult** dataset, where text in blue denotes the task description, text in green indicates optional few-shot examples (only used in in-context learning setting), and text in red is the test example. We provide the prompt templates for the other two datasets in Appendix D.

3.4 Evaluation Metrics

To assess fairness in the aforementioned datasets, we examine the disparity between different subgroups of protected attributes using the following common fairness metrics: accuracy (ACC), F1 score (F1), statistical parity (SP), and equality of opportunity (EoO). Here, we briefly explain each evaluation metric.

Accuracy and F1 As the most basic metric, assessing accuracy among different subgroups ensures that the model delivers consistent performance across all groups, without undue favor to any particular subgroups.

Considering that the evaluated datasets may be imbalanced, especially among different subgroups, the F1 Score computes the harmonic mean of precision and recall, offering a balanced perspective between these two metrics.

Statistical Parity Statistical parity is attained when *positive* decision outcomes (e.g., being predicted as good credit risk) are independent of the protected attributes. This metric assesses whether different subgroups receive similar treatment from the model. For each subgroup z_i of each protected attribute Z , we calculate its statistical parity as:

$$P(\hat{Y} = 1 | Z = z_i).$$

Then we calculate the difference of Statistical Parity (d_{SP}) of this protected attribute as:

$$d_{SP} = \overset{\text{statistical parity of } z_1}{P(\hat{Y} = 1 | Z = z_1)} - \overset{\text{statistical parity of } z_2}{P(\hat{Y} = 1 | Z = z_2)},$$

where z_1 is the **minority group** and z_2 is the **majority**.

Equality of Opportunity Equality of opportunity requires that qualified individuals have an equal chance of being correctly classified by the model, regardless of their membership in a protected group. This metric ensures equal *true positive* rates between different subgroups, providing equal opportunities for each subgroup. Similar to statistical parity, for equality of opportunity, we calculate the difference of Equality of Opportunity (d_{EoO}) as:

$$d_{EoO} = \overset{\text{equality of opportunity of } z_1}{P(\hat{Y} = 1 | Y = 1, Z = z_1)} - \overset{\text{equality of opportunity of } z_2}{P(\hat{Y} = 1 | Y = 1, Z = z_2)}.$$

Each of these metrics offers a different perspective on fairness. For each subgroup from each protected attribute, we will compute every aforementioned metric. A model demonstrating good fairness should show minimal gaps in these fairness metrics between different subgroups. Considering them together can provide a more comprehensive evaluation of the model’s fairness across different subgroups, ensuring that individuals are not unfairly disadvantaged based on their membership in a protected group.

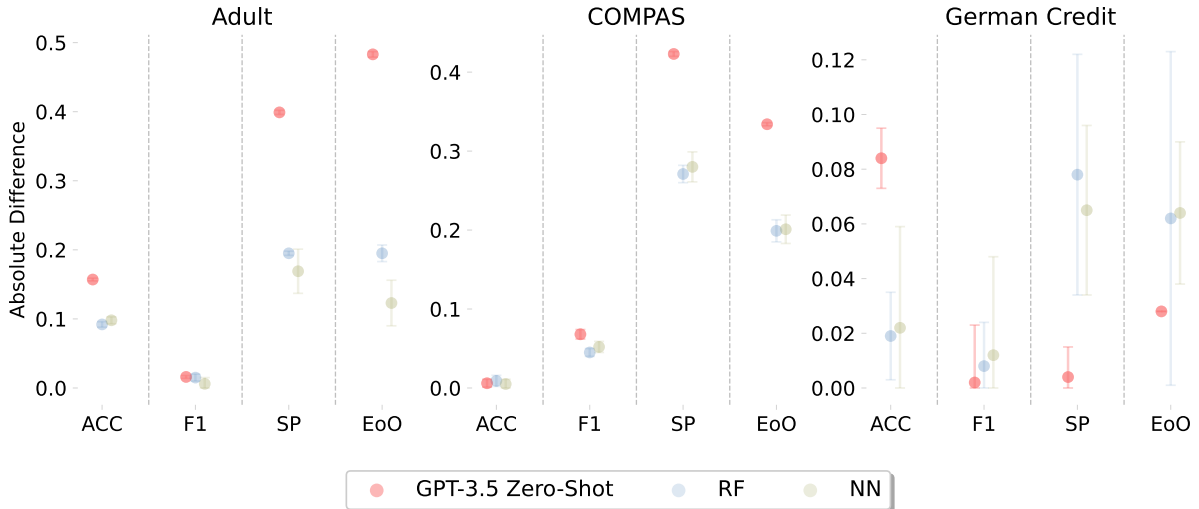


Figure 1: **Fairness Comparison of GPT-3.5 Zero-Shot Prompting with RF and NN.** We present the absolute differences in ACC, F1, SP, and EoO between subgroups of protected attributes across three datasets: “sex” for **Adult** and **German Credit**, and “race” for **COMPAS**. Notable fairness gaps between protected and non-protected subgroups are observed for GPT-3.5 in a zero-shot prompting manner (GPT-3.5 Zero-Shot), which are significantly larger compared to those observed with RF and NN, except for the **German Credit** dataset.

	d_{ACC}	d_{F1}	d_{SP}	d_{EoO}
Adult				
w/ protected attr.	0.157 _{0.002}	-0.016 _{0.002}	-0.399 _{0.003}	-0.483 _{0.004}
w/o protected attr.	0.106 _{0.002}	-0.021 _{0.001}	-0.287 _{0.002}	-0.273 _{0.002}
COMPAS				
w/ protected attr.	-0.006 _{0.005}	0.068 _{0.006}	-0.423 _{0.003}	-0.334 _{0.002}
w/o protected attr.	-0.002 _{0.003}	0.058 _{0.004}	-0.239 _{0.002}	-0.157 _{0.004}

Table 1: **Fairness Comparison of GPT-3.5 Zero-Shot Prompting with and without Protected Attributes.** When protected attributes are removed from the input, there is a notable decrease in fairness gaps, especially in SP and EoO. This observation provides further evidence of the inherent bias present in GPT-3.5.

4 Zero-Shot Prompting for Tabular Data

To explore the fairness of LLMs when making classifications on tabular data, we first conduct experiments in a zero-shot setting. We assess the fairness metrics of the outcomes and examine whether GPT-3.5 without any finetuning or few-shot examples would be influenced by social biases and stereotypes for tabular classification. We run all the experiments 5 times and compute the mean and standard deviation.

4.1 Regular Zero-Shot Prompting

In Figure 1, we compare the absolute differences in four evaluated fairness metrics between subgroups of protected attributes for GPT-3.5 in a zero-shot manner (GPT-3.5 Zero-Shot) with RF and NN on the **Adult**, **German Credit**, and **COMPAS** datasets, respectively. The evaluated metrics include accuracy

(ACC), F1 score (F1), statistical parity (SP), and equality of opportunity (EoO). For the **Adult** and **German Credit** datasets, the subgroups *female* and *male* are assessed regarding the protected attribute “sex”, identifying *female* as a disadvantaged group. In the **COMPAS** dataset, we evaluate “race” as protected attributes, recognizing *African American* (AA) as the disadvantaged group. Due to space limitations, we included the concrete numbers in Tables 5 to 7 in Appendix B.

It is notable that when directly utilizing LLMs to make classifications for tabular data, without any fine-tuning or in-context learning, a significant fairness metric gap between the protected and non-protected groups is observed for GPT-3.5 (highlighted in red in Appendix B). For example, the EoO difference between *male* and *female* on the **Adult** dataset reaches 0.483, indicating a substan-

tial disadvantage for the *female* group. Additionally, when compared with conventional models like RF and NN, the biases in zero-shot predictions made by GPT-3.5 are significantly larger when applied to the **Adult** and **COMPAS** dataset. This observation suggests the presence of inherent gender and race biases within GPT-3.5.

Exceptionally, GPT-3.5 is extremely biased for **German Credit** dataset where it classifies almost everything into *good credit* class in the zero-shot setting, thus rendering the difference in **SP** and **EoO** for both subgroups to be near 0. The accuracy for each subgroup is near 50%, performing similar to random guessing. The possible reason might be that the **German Credit** dataset is too challenging for making tabular classifications with LLMs (especially, since the features of **German Credit** are ambiguous and vague). This also suggests that, when using GPT-3.5 to make predictions on tabular data, a potential description of table feature names is favorable.

4.2 Zero-Shot Prompting with Protected Attributes Removed

To further demonstrate inherent social biases in LLMs, we compare fairness gaps in zero-shot classifications using GPT-3.5 under two conditions: one including protected attributes and the other with these attributes removed. As shown in Table 1, fairness gaps notably decrease, particularly in **SP** and **EoO**, when protected attributes are excluded from the input. This observation further confirms the presence of inherent bias in GPT-3.5.

Taken together, these results demonstrate the tendency of GPT-3.5 to rely on social biases and stereotypes inherited from their pretraining corpus when applied to tabular data. This suggests that using LLMs for classifications on tabular data may incur significant fairness risks, potentially disproportionately disadvantaging marginalized communities and perpetuating societal biases and stereotypes. Given the widespread use of tabular data in high-stakes contexts, these findings raise serious concerns about the potential for harm.

5 Few-Shot Prompting for Tabular Data

As demonstrated in Section 4, employing GPT-3.5 for classifications on tabular data reveals significant social biases in a zero-shot setting. Instead of directly utilizing GPT-3.5 for zero-shot tabular classifications, this section explores whether in-

cluding few-shot examples during prompting will reduce or amplify these biases. To delve deeper into the influence of few-shot examples during in-context learning (ICL), we not only consider the regular ICL approach as detailed in Section 5.1, but we also experiment by flipping the labels of the few-shot examples to further examine their effect on the biases, as discussed in Section 5.2.

Again, for robustness, each experiment is conducted 5 times, with the mean and standard deviation reported. We present the fairness gap in Figure 2 and the complete results in Appendix B.

5.1 Regular In-Context Learning

Previous works have demonstrated that LLMs can learn the input-label mappings in context (Akyürek et al., 2022; Xie et al., 2022; Von Oswald et al., 2023). However, the influence of in-context learning on fairness has not been thoroughly examined. For in-context learning, the test example and task description, along with a few-shot examples, are provided to the LLMs for generating the final predictions. The few-shot examples are inserted before the test example in the prompt, as outlined in Section 3.3. We set the number of in-context examples as 50. For each dataset, we randomly select the in-context examples from the training set for each test example.

In Tables 5-7, we demonstrate that for two of the evaluated datasets (except for **COMPAS**), the incorporation of few-shot examples brings about performance improvements. Additionally, in Figure 2, we observe that incorporating few-shot examples into prompting reduces the fairness metric gap between different subgroups. However, a significant fairness issue still persists. Moreover, for the **Adult** and **COMPAS** datasets, the disparity in fairness metrics of in-context learning is more notable when compared to traditional models RF and NN. This highlights the inherent biases embedded within LLMs, which are not solely derived from the task datasets.

5.2 Label Flipping

To delve deeper into the sources of biases within GPT-3.5, we further examine the impact of the labels of in-context examples on fairness. As demonstrated in Figure 2, label flipping significantly reduces biases across all evaluated datasets. For all datasets, the difference in statistical parity (**SP**) and equality of opportunity (**EoO**) is minimized with label-flipped ICL. For example, the absolute




Figure 2: **Fairness Comparison of GPT-3.5 Few-Shot Prompting.** We compare the absolute differences in ACC, F1, SP, and EoO between subgroups of protected attributes across the **Adult**, **COMPAS** and **German Credit** datasets for GPT-3.5 in three settings: zero-shot prompting (GPT-3.5 Zero-Shot), few-shot prompting (GPT-3.5 Few-Shot), and few-shot prompting with flipped labels (GPT-3.5 Few-Shot LF). Incorporating few-shot examples (GPT-3.5 Few-Shot) can partially reduce the inherent biases in GPT-3.5, but it cannot completely eliminate them. The fairness gap persists and is greater than that observed in RF and NN. Furthermore, label-flipped few-shot examples (GPT-3.5 Few-Shot LF) can effectively reverse the bias effects, further narrowing the fairness gaps.

gap of EoO on the **Adult** dataset decreases from 0.483 in zero-shot prompting to 0.037, almost completely eliminating the bias. These findings further corroborate the existence of inherent biases in GPT-3.5.

However, flipped labels lead to a significant drop in classification performance. Though previous research suggests that the effectiveness of ICL predominantly stems from semantic priors, rather than learning the input-label mappings (Min et al., 2022; Wei et al., 2023b) and demonstrates that the performance of ICL is barely affected even with flipped or random labels for in-context examples, the focus of these works lies mainly on traditional natural language processing tasks. In contrast, we observe that the labels of in-context examples hold substantial influence over the performance in our unique setup, where GPT-3.5 are deployed for classifications on tabular data. This could be attributed to the limited exposure of these models to tabular data during pre-training, thereby amplifying the role of input-label mapping of in-context examples.

6 Finetuning for Tabular Data

6.1 Regular Finetuning

Finally, we extend our investigation to assess if finetuning the models on the entire training set could

aid in diminishing the social biases in LLMs. For GPT-3.5, fine-tuning is executed using the publicly released API from OpenAI. For RF and NN, we provide the training details in Appendix C. We still run all the experiments 5 times and compute the mean and standard deviation. In Figure 3, we show that finetuning (GPT-3.5 FT) effectively reduces unfairness in all datasets, making them comparable and sometimes significantly better in terms of SP and EoO when compared to RF and NN. For example, the absolute difference in EoO after finetuning on the **Adult** dataset is 0.0714, which is lower than the 0.123 difference of an NN.

6.2 Resampled Finetuning

Resampling is a method frequently utilized to enhance fairness in traditional machine learning model training, particularly in scenarios where there is a significant class imbalance or bias in the data. We further explore the potential of resampling the task datasets to reduce the fairness gap for LLMs. Specifically, we evaluate two approaches: oversampling the minority group (OS) and undersampling the majority group (US). As depicted in Figure 3, resampling fails to mitigate the social biases in GPT-3.5 when making tabular classification, even though we demonstrate that

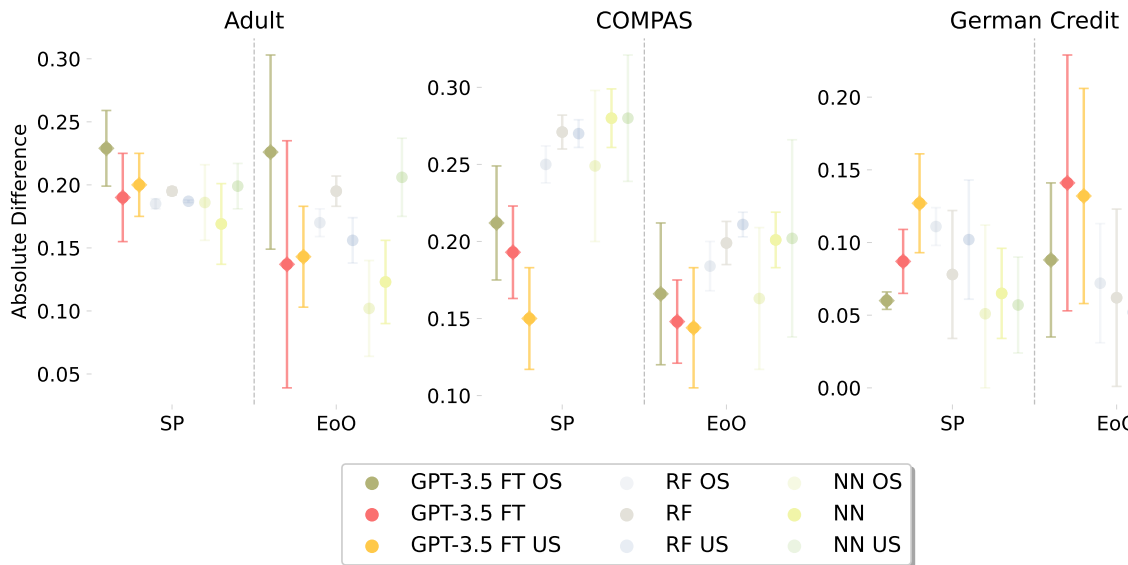


Figure 3: **Fairness Comparison of GPT-3.5 Finetuning with RF and NN Finetuning.** We compare the absolute differences in fairness metrics - ACC, F1, SP, and EoO - between subgroups of protected attributes across the **Adult**, **COMPAS** and **German Credit** datasets for the finetuned GPT-3.5 models using three different approaches: finetuning on the entire training data (GPT-3.5 FT), oversampled data (GPT-3.5 FT OS), and undersampled data (GPT-3.5 FT US). While finetuning GPT-3.5 on task datasets (GPT-3.5 FT) can mitigate inherent social biases to some extent, the application of data resampling techniques (GPT-3.5 FT OS and GPT-3.5 FT US) during the finetuning process does not consistently yield similar results for LLMs compared to the typical mitigation observed in traditional machine learning models.

oversampling generally reduces social biases for both RF and NN, except for a few instances such as oversampling in NN for **Adult** dataset worsens the fairness.

Our finetuning experiments show that the social biases inherited from GPT-3.5’s pre-training corpus, which are noticeably evident when making classifications on tabular data, can sometimes be mitigated through finetuning. Nevertheless, unlike the consistent outcomes typically seen in traditional machine learning models, data resampling does not consistently produce similar results for finetuning LLMs.

7 Conclusion

In this work, we thoroughly investigate the under-explored problem of fairness of large language models (LLMs) for tabular tasks. Our study unfolds in several phases.

Initially, we assess the inherent fairness displayed by GPT-3.5, comparing their performance in zero-shot learning scenarios against traditional machine learning models like random forests (RF) and shallow neural networks (NN). Furthermore, we investigate how GPT-3.5 learns and propagates social biases when subjected to few-shot in-context

learning, label-flipped in-context learning, finetuning, and data resampling techniques.

Our discoveries shed light on several key insights. We find that GPT-3.5 tends to heavily rely on the social biases inherited from their pre-training data when making predictions, which is a concerning issue. Moreover, we observe that few-shot in-context learning can partially mitigate the inherent biases in GPT-3.5, yet it cannot entirely eliminate them. A significant fairness metric gap between different subgroups persists and exceeds that observed in RF and NN. This observation underscores the existence of biases within the LLMs themselves, beyond just the task datasets. Additionally, label-flipping applied to the few-shot examples effectively reverses the effects of bias, again corroborating the existence of inherent biases in GPT-3.5. However, as expected, this leads to a loss in the classification performance. Besides, our work reveals that while fine-tuning can sometimes improve the fairness of GPT-3.5, data resampling does not consistently yield the same results, unlike what is typically observed in traditional machine learning models. This underscores the need for the development of more effective strategies to mitigate the bias inherent in LLMs and ensure their fairness

when deployed in real-world applications.

Limitations

It is important to note that our study exclusively focuses on the GPT-3.5 model. Consequently, our conclusions are representative of GPT-3.5 alone and cannot be extrapolated to other LLMs, which might exhibit different behaviors or biases. This focus on a single model thus restricts the broader applicability of our findings.

Furthermore, for each experiment, we employed only one type of prompt. This approach limits the generalizability of our conclusions, as different prompts might yield varying results. The use of a singular prompt type does not capture the full spectrum of possible interactions and outcomes that might be observed with a diverse range of prompting strategies.

Looking ahead, we plan to broaden the scope of our research. This expansion will include experimenting with additional models beyond GPT-3.5, thus offering a more comprehensive understanding of fairness for different LLMs. We also intend to explore a variety of prompting strategies, such as Chain of Thought (CoT) prompting, to assess how different methods may impact model bias and fairness. These future endeavors aim to provide a more nuanced and thorough exploration of the capabilities and limitations of LLMs in the context of fairness.

References

- Abubakar Abid, Maheen Farooqi, and James Zou. 2021a. [Large language models associate muslims with violence](#). *Nature Machine Intelligence*, 3(6):461–463.
- Abubakar Abid, Maheen Farooqi, and James Zou. 2021b. [Persistent anti-muslim bias in large language models](#). In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, page 298–306, New York, NY, USA. Association for Computing Machinery.
- Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. 2022. What learning algorithm is in-context learning? investigations with linear models. *arXiv preprint arXiv:2211.15661*.
- Sercan Ö Arik and Tomas Pfister. 2021. Tabnet: Attentive interpretable tabular learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 6679–6687.
- Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Jackson Kernion, Kamal Ndousse, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, and Jared Kaplan. 2021. [A general language assistant as a laboratory for alignment](#).
- Christine Basta, Marta R. Costa-jussà, and Noe Casas. 2019. [Evaluating the underlying gender bias in contextualized word embeddings](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*.
- Barry Becker and Ronny Kohavi. 1996. Adult. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5XW20>.
- Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, Seema Nagar, Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. 2018. [AI Fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias](#).
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. [Man is to computer programmer as woman is to homemaker? debiasing word embeddings](#). In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- R Bommasani. 2021. Opportunities and risks of foundation models. <https://openai.com/reports/foundation-models/>.
- Vadim Borisov, Tobias Leemann, Kathrin SeBler, Johannes Haug, Martin Pawelczyk, and Gjergji Kasneci. 2022. [Deep neural networks and tabular data: A survey](#). *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–21.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. [Semantics derived automatically from language corpora contain human-like biases](#). *Science*, 356(6334):183–186.
- D. Dua and C. Graff. 2019. [UCI machine learning repository](#).
- Deep Ganguli, Danny Hernandez, Liane Lovitt, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova Dassarma, Dawn Drain, Nelson Elhage, Sheer El Showk, Stanislav Fort, Zac Hatfield-Dodds, Tom Henighan, Scott Johnston, Andy Jones, Nicholas Joseph, Jackson Kernian, Shauna Kravec, Ben Mann, Neel Nanda, Kamal Ndousse, Catherine Olsson, Daniela Amodei, Tom Brown, Jared Kaplan, Sam McCandlish, Christopher Olah, Dario Amodei, and Jack Clark. 2022. [Predictability and surprise in large generative models](#). In *2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, page 1747–1764, New York, NY, USA. Association for Computing Machinery.
- Leo Grinsztajn, Edouard Oyallon, and Gael Varoquaux. 2022. [Why do tree-based models still outperform deep learning on typical tabular data?](#) In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David Sontag. 2023. [TablIm: Few-shot classification of tabular data with large language models](#). In *International Conference on Artificial Intelligence and Statistics*, pages 5549–5581. PMLR.
- Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. 2016. [How we analyzed the compas recidivism algorithm](#).
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2022. [What makes good in-context examples for GPT-3?](#) In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 100–114, Dublin, Ireland and Online. Association for Computational Linguistics.
- Yanchen Liu, William Held, and Diyi Yang. 2023a. [Dada: Dialect adaptation via dynamic aggregation of linguistic rules](#).
- Yanchen Liu, Timo Schick, and Hinrich Schtze. 2023b. [Semantic-oriented unlabeled priming for large-scale language models](#). In *Proceedings of The Fourth Workshop on Simple and Efficient Natural Language Processing (SustaiNLP)*, pages 32–38, Toronto, Canada (Hybrid). Association for Computational Linguistics.
- Li Lucy and David Bamman. 2021. [Gender and representation bias in GPT-3 generated stories](#). In *Proceedings of the Third Workshop on Narrative Understanding*, pages 48–55, Virtual. Association for Computational Linguistics.
- Xinxi Lyu, Sewon Min, Iz Beltagy, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [Z-ICL: Zero-shot in-context learning with pseudo-demonstrations](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2304–2317, Toronto, Canada. Association for Computational Linguistics.
- Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *International Conference on Learning Representations*.
- Ravid Shwartz-Ziv and Amitai Armon. 2021. [Tabular data: Deep learning is not all you need](#). In *8th ICML Workshop on Automated Machine Learning (AutoML)*.
- Dylan Slack and Sameer Singh. 2023. [Tablet: Learning from instructions for tabular data](#). *arXiv*.
- Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. 2023. [Transformers learn in-context by gradient descent](#). In *International Conference on Machine Learning*, pages 35151–35174. PMLR.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.

- Jerry Wei, Le Hou, Andrew Lampinen, Xiangning Chen, Da Huang, Yi Tay, Xinyun Chen, Yifeng Lu, Denny Zhou, Tengyu Ma, et al. 2023a. Symbol tuning improves in-context learning in language models. *arXiv preprint arXiv:2305.08298*.
- Jerry W. Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu, Da Huang, Denny Zhou, and Tengyu Ma. 2023b. [Larger language models do in-context learning differently](#). *ArXiv*, abs/2303.03846.
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, Zac Kenton, Sasha Brown, Will Hawkins, Tom Stepleton, Courtney Biles, Abeba Birhane, Julia Haas, Laura Rimell, Lisa Anne Hendricks, William Isaac, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2021. [Ethical and social risks of harm from language models](#).
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2022. [An explanation of in-context learning as implicit bayesian inference](#). In *International Conference on Learning Representations*.
- Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. [TaBERT: Pretraining for joint understanding of textual and tabular data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8413–8426, Online. Association for Computational Linguistics.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. [Men also like shopping: Reducing gender bias amplification using corpus-level constraints](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2989, Copenhagen, Denmark. Association for Computational Linguistics.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR.
- Caleb Ziems, William Held, Jingfeng Yang, Jwala Dhamala, Rahul Gupta, and Diyi Yang. 2023. [Multi-VALUE: A framework for cross-dialectal English NLP](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 744–768, Toronto, Canada. Association for Computational Linguistics.

A Dataset Description

We provide a detailed description of each feature from the datasets evaluated in our paper.

A.1 Adult

The original Adult Income Dataset contains 14 features and the target *Income*, as described in Table 2. Following prior work (Slack and Singh, 2023), we omit *Education-Num* and *Fnlwgt* as they are not crucial for income prediction, along with *Race* and *Native-Country*, to center our attention on *Sex* as the protected attribute.

A.2 German Credit

The original German Credit Dataset contains 20 features, as detailed in Table 3. For simplicity and consistency with prior work, only the features not shown in *italics* are retained in our work. Furthermore, we extract *Sex* as an additional protected attribute from the *Personal Status and Sex* feature.

A.3 COMPAS

The raw COMPAS Recidivism dataset contains more than 50 attributes. Following the approach of Larson et al. (2016), we carry out the necessary preprocessing. More specifically, we group the *race* attribute into *African-American* and *Not African-American*, and consider only the features *sex*, *race*, *age*, *charge degree*, *priors count*, *risk*, and *two-year recid* (target). We frame the task as predicting whether an individual will recidivate within two years (*Did Not Reoffend* or *Reoffended*), based on their demographic and criminal history. Due to page limitations, we provide descriptions for only the features used in our work in Table 4.

B Complete Evaluation Results

Due to space limitations, we present the complete evaluation results on the **Adult**, **German Credit**, and **COMPAS** datasets in Tables 5 to 7, respectively.

C RF and NN Hyperparameters

For RF, we fix the number of trees to 100 for all datasets as well as models. For NN, we use a 3 hidden-layered network with hyperparameters described in Table 8.

D Prompt Templates

Beyond the Adult dataset, we provide the serialization and prompt templates utilized in our work

for the German Credit and COMPAS datasets here. As before, text in blue denotes the task description, text in green indicates optional few-shot examples (only used in in-context learning setting), and text in red is the test example.

D.1 German Credit

Prompt 2. Prompt Template for German Credit Dataset.

Predict the credit risk of a German bank customer based on their profile data. Answer with one of the following: bad | good.

Example 1 -
Age: 25 - 45
Sex: female
Job: highly skilled
Housing: rent
Saving accounts: little
Checking account: little
Credit amount: 2606
Duration: 21
Purpose: radio/TV
Answer: good
...

Age: 25 - 45
Sex: male
Job: skilled
Housing: own
Saving accounts: little
Checking account: little
Credit amount: 1345
Duration: 18
Purpose: radio/TV
Answer:

D.2 COMPAS

Prompt 3. Prompt Template for COMPAS Dataset.

Predict whether an individual will recidivate within two years based on demographic and criminal history. Answer with one of the following: Did Not Reoffend | Reoffended.

Example 1 -
sex: Male
race: African-American
age cat: 25 - 45
c charge degree: F
priors count: 0
risk: Low
Answer: Did Not Reoffend
...

sex: Male
race: African-American
age cat: 25 - 45
c charge degree: M
priors count: 13
risk: High
Answer:

Feature	Type	Description
Age	Continuous	Represents the age of an individual.
Workclass	Categorical	Indicates the type of employment, such as private, self-employed, or government.
<i>Fnlwgt</i>	Continuous	Stands for “final weight” and is a numerical value used in sampling for survey data.
Education	Categorical	Specifies the highest level of education attained by the individual, such as high school, bachelor’s degree, etc.
<i>Education-Num</i>	Continuous	Represents the numerical equivalent of the education level.
Marital-Status	Categorical	Describes the marital status of the individual, including categories like married, divorced, or single.
Occupation	Categorical	Indicates the occupation of the individual, such as managerial, technical, or clerical work.
Relationship	Categorical	Specifies the individual’s role in the family, such as husband, wife, or child.
Race	Categorical	Represents the individual’s race or ethnic background.
Sex	Categorical	Indicates the gender of the individual, either male or female.
Capital-Gain	Continuous	Refers to the capital gains, which are profits from the sale of assets, of the individual.
Capital-Loss	Continuous	Represents the capital losses, which are losses from the sale of assets, of the individual.
Hours-Per-Week	Continuous	Denotes the number of hours worked per week by the individual.
<i>Native-Country</i>	Categorical	Specifies the native country or place of origin of the individual.
Income (target)	Binary	The target variable indicating whether an individual’s income exceeds a certain threshold, typically \$50,000 per year.

Table 2: Features in the original **Adult** dataset. Those not used in our work are shown in *italics*.

Feature	Type	Description
Credit Amount	Continuous	The amount of credit requested by the applicant.
Duration	Continuous	The duration of the credit in months.
<i>Installment Rate</i>	Ordinal	The installment rate in percentage of disposable income.
<i>Residence Since</i>	Ordinal	The number of years the applicant has lived at their current residence.
Age	Continuous	The age of the applicant.
<i>Number of Existing Credits</i>	Ordinal	The number of existing credits at this bank.
<i>Number of Dependents</i>	Ordinal	The number of dependents of the applicant.
Checking Account Status	Categorical	The status of the applicant's checking account, such as "no checking," "<0 DM," "0-200 DM," or "no known checking."
<i>Credit History</i>	Categorical	The credit history of the applicant, including categories like "critical/other existing credit," "existing paid," "delayed previously," etc.
Purpose	Categorical	The purpose of the credit, such as "radio/tv," "education," "new car," etc.
Savings Account	Categorical	The status of the applicant's savings account/bonds, including categories like "unknown/none," "<100 DM," "500-1000 DM," etc.
<i>Employment Since</i>	Categorical	The duration of the applicant's current employment, such as "unemployed," "<1 year," "4-7 years," etc.
<i>Personal Status and Sex</i>	Categorical	The personal status and sex of the applicant, including categories like "male single," "female div/dep/mar," etc.
<i>Other Debtors/Guarantors</i>	Categorical	Indicates the presence of other debtors/guarantors, such as "none," "guarantor," "co applicant."
<i>Property</i>	Categorical	Describes the type of property owned by the applicant, such as "real estate," "life insurance," "car or other," etc.
<i>Other Installment Plans</i>	Categorical	The presence of other installment plans.
Housing	Categorical	The housing situation of the applicant, such as "own," "for free," and "rent."
Job	Categorical	The type of job held by the applicant, including categories like "skilled," "unskilled resident," "high qualif/self emp/mgmt," etc.
<i>Telephone</i>	Binary	Indicates whether the applicant has a telephone (yes/no).
<i>Foreign Worker</i>	Binary	Indicates whether the applicant is a foreign worker (yes/no).
Risk (target)	Binary	The target variable indicating credit risk (good/bad).

Table 3: Features in the original **German Credit** dataset. Those not used in our work are shown in *italics*. Additionally, from the original feature *Personal Status and Sex*, we extract *Sex* as a protected attribute.

Feature	Type	Description
Sex	Categorical	The gender of the individual.
Race	Categorical	The race of the individual, grouped into <i>African-American</i> and <i>Not African-American</i> .
Age	Continuous	The age of the individual.
Charge Degree	Categorical	The degree of the charge against the individual.
Priors Count	Continuous	The number of prior convictions or charges.
Risk	Categorical	The risk assessment for recidivism.
Two-Year Recid (target)	Binary	The target variable indicating whether an individual recidivated within two years.

Table 4: Features in the **COMPAS** Recidivism Dataset (Preprocessed).

				ACC	F1	SP	EoO
GPT-3.5-turbo	Zero-Shot		<i>f</i>	0.898 _{0.001}	0.711 _{0.002}	0.065 _{0.001}	0.357 _{0.000}
			<i>m</i>	0.742 _{0.002}	0.727 _{0.002}	0.464 _{0.003}	0.840 _{0.004}
			<i>d</i>	0.157 _{0.002}	-0.016 _{0.002}	-0.399 _{0.003}	-0.483 _{0.004}
	Few-shot	Regular	<i>f</i>	0.899 _{0.002}	0.735 _{0.003}	0.082 _{0.002}	0.429 _{0.000}
			<i>m</i>	0.781 _{0.003}	0.749 _{0.002}	0.339 _{0.003}	0.700 _{0.003}
			<i>d</i>	0.118 _{0.004}	-0.014 _{0.004}	-0.257 _{0.005} ↓	-0.271 _{0.003} ↓
		Label-flipping	<i>f</i>	0.682 _{0.004}	0.590 _{0.003}	0.396 _{0.006}	0.800 _{0.013}
			<i>m</i>	0.614 _{0.002}	0.605 _{0.002}	0.545 _{0.001}	0.763 _{0.003}
			<i>d</i>	0.068 _{0.004}	-0.015 _{0.004}	-0.148 _{0.006} ✓	0.037 _{0.014} ✓
	Finetuning	Regular	<i>f</i>	0.915 _{0.014}	0.773 _{0.036}	0.079 _{0.002}	0.476 _{0.048}
			<i>m</i>	0.799 _{0.005}	0.754 _{0.005}	0.269 _{0.036}	0.613 _{0.053}
			<i>d</i>	0.116 _{0.009}	0.020 _{0.039}	-0.190 _{0.035} ↓	-0.137 _{0.098} ↓
		Oversampling	<i>f</i>	0.913 _{0.016}	0.770 _{0.042}	0.081 _{0.004}	0.476 _{0.067}
			<i>m</i>	0.813 _{0.007}	0.780 _{0.003}	0.310 _{0.038}	0.702 _{0.048}
			<i>d</i>	0.100 _{0.013}	-0.010 _{0.041}	-0.229 _{0.030}	-0.226 _{0.077}
Undersampling	<i>f</i>	0.912 _{0.015}	0.770 _{0.046}	0.086 _{0.006}	0.488 _{0.084}		
	<i>m</i>	0.794 _{0.006}	0.751 _{0.001}	0.285 _{0.031}	0.631 _{0.044}		
	<i>d</i>	0.118 _{0.021}	0.018 _{0.046}	-0.200 _{0.025}	-0.143 _{0.040}		
RF	Regular	<i>f</i>	0.914 _{0.002}	0.767 _{0.006}	0.075 _{0.003}	0.457 _{0.010}	
		<i>m</i>	0.822 _{0.005}	0.783 _{0.005}	0.269 _{0.004}	0.652 _{0.004}	
		<i>d</i>	0.092 _{0.004}	-0.015 _{0.005}	-0.195 _{0.003}	-0.195 _{0.012}	
	Oversampling	<i>f</i>	0.912 _{0.006}	0.770 _{0.011}	0.084 _{0.005}	0.486 _{0.012}	
		<i>m</i>	0.824 _{0.002}	0.785 _{0.002}	0.270 _{0.003}	0.656 _{0.006}	
		<i>d</i>	0.087 _{0.005}	-0.015 _{0.01}	-0.185 _{0.004}	-0.170 _{0.011}	
Undersampling	<i>f</i>	0.917 _{0.004}	0.776 _{0.011}	0.075 _{0.001}	0.471 _{0.018}		
	<i>m</i>	0.814 _{0.003}	0.771 _{0.004}	0.263 _{0.002}	0.627 _{0.009}		
	<i>d</i>	0.103 _{0.005}	0.005 _{0.011}	-0.187 _{0.001}	-0.156 _{0.018}		
NN	Regular	<i>f</i>	0.917 _{0.003}	0.778 _{0.019}	0.081 _{0.016}	0.490 _{0.068}	
		<i>m</i>	0.819 _{0.006}	0.773 _{0.015}	0.250 _{0.045}	0.614 _{0.079}	
		<i>d</i>	0.098 _{0.005}	0.006 _{0.009}	-0.169 _{0.032}	-0.123 _{0.033}	
	Oversampling	<i>f</i>	0.916 _{0.004}	0.794 _{0.013}	0.100 _{0.016}	0.562 _{0.058}	
		<i>m</i>	0.813 _{0.012}	0.774 _{0.008}	0.286 _{0.044}	0.663 _{0.056}	
		<i>d</i>	0.103 _{0.011}	0.020 _{0.018}	-0.186 _{0.030}	-0.102 _{0.038}	
Undersampling	<i>f</i>	0.904 _{0.005}	0.748 _{0.014}	0.084 _{0.007}	0.452 _{0.030}		
	<i>m</i>	0.813 _{0.006}	0.774 _{0.005}	0.283 _{0.023}	0.659 _{0.031}		
	<i>d</i>	0.090 _{0.006}	-0.026 _{0.014}	-0.199 _{0.018}	-0.206 _{0.031}		

Table 5: **Fairness evaluation for Adult dataset.** This table depicts the evaluation of accuracy (ACC), F1 score (F1), statistical parity (SP), and equality of opportunity (EoO) metrics for the subgroup - *female* (*f*) and *male* (*m*) as well as the difference (*d*) between them. We list the protected group first. The significant fairness disparities are highlighted in red. Both in-context learning and finetuning can lead to bias reduction (indicated by ↓), and label-flipped in-context learning can further minimize bias (indicated by ✓).

			ACC	F1	SP	EoO	
GPT-3.5-turbo	Zero-Shot	<i>f</i>	0.471 0.011	0.359 0.021	0.980 0.011	1.000 0.000	
		<i>m</i>	0.556 0.000	0.357 0.000	0.984 0.000	0.972 0.000	
		<i>d</i>	-0.084 0.011	0.002 0.021	-0.004 0.011	0.028 0.000	
	Few-shot	Regular	<i>f</i>	0.610 0.013	0.593 0.013	0.348 0.027	0.453 0.029
			<i>m</i>	0.606 0.007	0.603 0.008	0.337 0.007	0.450 0.012
			<i>d</i>	0.003 0.012	-0.010 0.011	0.011 0.027	0.003 0.026
		Label-flipping	<i>f</i>	0.614 0.011	0.606 0.012	0.695 0.011	0.842 0.000
			<i>m</i>	0.559 0.013	0.538 0.011	0.638 0.013	0.672 0.023
			<i>d</i>	0.056 0.021	0.067 0.021	0.057 0.012	0.170 0.023
	Finetuning	Regular	<i>f</i>	0.571 0.067	0.567 0.062	0.619 0.101	0.711 0.186
			<i>m</i>	0.548 0.011	0.539 0.023	0.532 0.123	0.569 0.098
			<i>d</i>	0.024 0.079	0.029 0.085	0.087 0.022	0.141 0.088
		Oversampling	<i>f</i>	0.536 0.017	0.532 0.012	0.607 0.084	0.658 0.112
			<i>m</i>	0.532 0.011	0.523 0.020	0.548 0.079	0.569 0.059
			<i>d</i>	0.004 0.028	0.009 0.033	0.060 0.006	0.088 0.053
Undersampling	<i>f</i>	0.548 0.034	0.547 0.033	0.571 0.034	0.632 0.074		
	<i>m</i>	0.556 0.000	0.555 0.000	0.444 0.000	0.500 0.000		
	<i>d</i>	-0.008 0.034	-0.008 0.033	0.127 0.034	0.132 0.074		
RF	Regular	<i>f</i>	0.581 0.024	0.580 0.025	0.519 0.028	0.611 0.054	
		<i>m</i>	0.600 0.019	0.588 0.020	0.597 0.022	0.672 0.021	
		<i>d</i>	-0.019 0.016	-0.008 0.016	-0.078 0.044	-0.062 0.061	
	Oversampling	<i>f</i>	0.576 0.018	0.575 0.018	0.505 0.018	0.589 0.021	
		<i>m</i>	0.568 0.032	0.552 0.034	0.616 0.025	0.661 0.037	
		<i>d</i>	0.008 0.034	0.023 0.035	-0.111 0.013	-0.072 0.041	
Undersampling	<i>f</i>	0.586 0.024	0.585 0.024	0.533 0.024	0.632 0.047		
	<i>m</i>	0.575 0.031	0.555 0.037	0.635 0.033	0.683 0.022		
	<i>d</i>	0.011 0.024	0.031 0.031	-0.102 0.041	-0.052 0.039		
NN	Regular	<i>f</i>	0.533 0.024	0.533 0.024	0.519 0.028	0.558 0.026	
		<i>m</i>	0.556 0.017	0.544 0.017	0.584 0.012	0.622 0.022	
		<i>d</i>	-0.022 0.037	-0.012 0.036	-0.065 0.031	-0.064 0.026	
	Oversampling	<i>f</i>	0.548 0.040	0.547 0.040	0.552 0.028	0.611 0.026	
		<i>m</i>	0.562 0.026	0.547 0.024	0.603 0.048	0.644 0.057	
		<i>d</i>	-0.014 0.037	0.000 0.035	-0.051 0.061	-0.034 0.065	
Undersampling	<i>f</i>	0.529 0.049	0.524 0.047	0.467 0.051	0.495 0.042		
	<i>m</i>	0.495 0.025	0.490 0.023	0.524 0.047	0.517 0.054		
	<i>d</i>	0.033 0.063	0.035 0.059	-0.057 0.033	-0.022 0.061		

Table 6: **Fairness evaluation for German Credit dataset.** This table depicts the evaluation of accuracy (ACC), F1 score (F1), statistical parity (SP), and equality of opportunity (EoO) metrics for the subgroup - *female* (*f*) and *male* (*m*) as well as the difference (*d*) between them.

			ACC	F1	SP	EoO	
GPT-3.5-turbo	Zero-Shot	AA	0.657 _{0.005}	0.656 _{0.004}	0.395 _{0.001}	0.560 _{0.002}	
		nAA	0.663 _{0.002}	0.588 _{0.003}	0.817 _{0.002}	0.893 _{0.001}	
		d	-0.006 _{0.005}	0.068 _{0.006}	-0.423 _{0.003}	-0.334 _{0.002}	
	Few-shot	Regular	AA	0.633 _{0.002}	0.626 _{0.002}	0.362 _{0.003}	0.495 _{0.004}
			nAA	0.642 _{0.001}	0.623 _{0.002}	0.614 _{0.002}	0.709 _{0.002}
			d	-0.008 _{0.003}	0.003 _{0.003}	-0.252 _{0.003} ↓	-0.214 _{0.005} ↓
		Label-flipping	AA	0.482 _{0.004}	0.482 _{0.004}	0.499 _{0.003}	0.481 _{0.004}
			nAA	0.412 _{0.003}	0.408 _{0.003}	0.471 _{0.002}	0.404 _{0.003}
			d	0.070 _{0.005}	0.074 _{0.005}	0.028 _{0.005} ✓	0.077 _{0.007} ✓
	Finetuning	Regular	AA	0.611 _{0.016}	0.610 _{0.016}	0.464 _{0.031}	0.576 _{0.034}
			nAA	0.616 _{0.013}	0.586 _{0.016}	0.657 _{0.032}	0.724 _{0.029}
			d	-0.005 _{0.017}	0.024 _{0.024}	-0.193 _{0.030} ↓	-0.148 _{0.027} ↓
		Oversampling	AA	0.609 _{0.007}	0.608 _{0.007}	0.494 _{0.071}	0.605 _{0.066}
			nAA	0.625 _{0.020}	0.583 _{0.024}	0.706 _{0.037}	0.771 _{0.036}
			d	-0.016 _{0.016}	0.025 _{0.018}	-0.212 _{0.037}	-0.166 _{0.046}
		Undersampling	AA	0.591 _{0.010}	0.591 _{0.012}	0.513 _{0.053}	0.605 _{0.047}
			nAA	0.641 _{0.008}	0.612 _{0.009}	0.663 _{0.035}	0.749 _{0.037}
			d	-0.050 _{0.016}	-0.021 _{0.022}	-0.150 _{0.033}	-0.144 _{0.039}
RF	Regular	AA	0.662 _{0.004}	0.662 _{0.004}	0.496 _{0.006}	0.660 _{0.007}	
		nAA	0.671 _{0.004}	0.617 _{0.002}	0.767 _{0.008}	0.859 _{0.009}	
		d	-0.009 _{0.007}	0.045 _{0.005}	-0.271 _{0.011}	-0.199 _{0.014}	
	Oversampling	AA	0.660 _{0.005}	0.660 _{0.005}	0.493 _{0.010}	0.655 _{0.013}	
		nAA	0.671 _{0.002}	0.624 _{0.002}	0.743 _{0.003}	0.839 _{0.004}	
		d	-0.010 _{0.006}	0.037 _{0.006}	-0.250 _{0.012}	-0.184 _{0.016}	
	Undersampling	AA	0.648 _{0.002}	0.647 _{0.002}	0.491 _{0.004}	0.639 _{0.004}	
		nAA	0.667 _{0.005}	0.614 _{0.007}	0.761 _{0.006}	0.851 _{0.006}	
		d	-0.020 _{0.007}	0.033 _{0.008}	-0.270 _{0.009}	-0.211 _{0.008}	
NN	Regular	AA	0.666 _{0.003}	0.665 _{0.002}	0.462 _{0.034}	0.630 _{0.034}	
		nAA	0.662 _{0.003}	0.613 _{0.006}	0.742 _{0.019}	0.831 _{0.017}	
		d	0.005 _{0.006}	0.052 _{0.007}	-0.280 _{0.019}	-0.201 _{0.018}	
	Oversampling	AA	0.656 _{0.001}	0.653 _{0.012}	0.507 _{0.090}	0.665 _{0.101}	
		nAA	0.643 _{0.013}	0.580 _{0.034}	0.757 _{0.107}	0.828 _{0.091}	
		d	0.013 _{0.014}	0.073 _{0.043}	-0.249 _{0.049}	-0.163 _{0.046}	
	Undersampling	AA	0.660 _{0.019}	0.657 _{0.023}	0.477 _{0.078}	0.638 _{0.097}	
		nAA	0.657 _{0.013}	0.602 _{0.026}	0.757 _{0.051}	0.839 _{0.040}	
		d	0.003 _{0.024}	0.055 _{0.043}	-0.280 _{0.041}	-0.202 _{0.064}	

Table 7: **Fairness evaluation for COMPAS dataset** for the subgroup - *African American (AA)*, and *Non African American (nAA)* as well as the difference (*d*). The significant fairness disparities are highlighted in **red**. Both in-context learning and finetuning can lead to bias reduction (indicated by ↓), and label-flipped in-context learning can further minimize bias (indicated by ✓).

	h1	h2	h3	lr	batch size	epochs
Adult	16	64	16	0.07	128	300
German Credit	64	64	32	0.07	128	300
COMPAS	64	128	64	0.09	128	300

Table 8: Hyperparameters for all datasets for a 3-layer neural network, where h1, h2, and h3 represent the number of neurons in first, second, and third hidden layers respectively, lr represents the learning rate and is followed by the batch size and the number of epochs the models are trained for.