

# PoliTune: Analyzing the Impact of Data Selection and Fine-Tuning on Economic and Political Biases in Large Language Models

Ahmed Agiza<sup>1</sup>, Mohamed Mostagir<sup>2</sup>, Sherief Reda<sup>3</sup>

<sup>1</sup>Computer Science Department, Brown University

<sup>2</sup>Ross School of Business, University of Michigan

<sup>3</sup>School of Engineering, Brown University

ahmed\_agiza@brown.edu, mosta@umich.edu, sherief\_reda@brown.edu

## Abstract

In an era where language models are increasingly integrated into decision-making and communication, understanding the biases within Large Language Models (LLMs) becomes imperative, especially when these models are applied in the economic and political domains. This work investigates the impact of fine-tuning and data selection on economic and political biases in LLMs. In this context, we introduce PoliTune, a fine-tuning methodology to explore the systematic aspects of aligning LLMs with specific ideologies, mindful of the biases that arise from their extensive training on diverse datasets. Distinct from earlier efforts that either focus on smaller models or entail resource-intensive pre-training, PoliTune employs Parameter-Efficient Fine-Tuning (PEFT) techniques, which allow for the alignment of LLMs with targeted ideologies by modifying a small subset of parameters. We introduce a systematic method for using the open-source LLM Llama3-70B for dataset selection, annotation, and synthesizing a preferences dataset for Direct Preference Optimization (DPO) to align the model with a given political ideology. We assess the effectiveness of PoliTune through both quantitative and qualitative evaluations of aligning open-source LLMs (Llama3-8B and Mistral-7B) to different ideologies. Our work analyzes the potential of embedding specific biases into LLMs and contributes to the dialogue on the ethical application of AI, highlighting the importance of deploying AI in a manner that aligns with societal values.

## 1 Introduction

Large Language Models (LLMs) have emerged as influential tools in numerous domains with strong potential impacts on economics, policy-making, and social governance. LLMs can process, interpret, and generate human language at a remarkable scale, and offer significant opportunities for enhancing decision-making processes. However, integrating LLMs into these critical domains has its complexities, particularly concerning the biases inherent in the models due to their training on extensive datasets. The manifestation of these biases in sensitive areas like economic policies introduces ethical, social, and economic challenges that demand careful consideration and proactive management. As LLMs are trained on diverse and extensive text corpora, they inherently encapsulate biases, cultural norms, and stereotypes

in their training data (Yeh et al. 2023; Nadeem, Bethke, and Reddy 2020; Liang et al. 2021). The reflection of societal biases in AI models raises concerns about the fairness and neutrality of the decision-making processes they inform. The potential for these biases to shape policy recommendations and public opinion makes it essential to study how biases are embedded in LLMs intentionally and unintentionally. Moreover, it is essential to analyze the capacity of LLMs to be deliberately aligned with specific political ideologies and their implications for policy-making, which is a critical area of study. Such alignment has the potential to reinforce pre-existing societal biases or lead to the marginalization of underrepresented groups. This aspect of LLM applications highlights the need for a thorough understanding of how LLMs can be guided toward certain ideologies that can compromise ethical standards or societal welfare. Hence, establishing ethical guidelines and governance frameworks for AI is crucial in addressing these challenges and ensuring that deploying LLMs aligns with societal values and democratic principles. While many studies highlight the capacity of LLMs to streamline operations, enhance user engagement, and drive innovation (Hadi et al. 2023; Luitse and Denkena 2021; Thirunavukarasu et al. 2023; Huang et al. 2023), researchers have highlighted the susceptibility of LLMs to perpetuate and even increase existing biases (Yeh et al. 2023; Nadeem, Bethke, and Reddy 2020; Liang et al. 2021; Ghafouri et al. 2023). Additionally, there are challenges with biasing LLMs due to the difficulty of encoding such an objective within a loss function. Previous studies have primarily focused on biasing smaller language models (LMs), not tackling the unique challenges associated with biasing larger counterparts, LLMs. Moreover, existing approaches often involve pre-training complete models from scratch, such as RoBERTa and GPT (Feng et al. 2023; Jiang et al. 2022), which is resource-intensive and financially prohibitive for LLMs. Our work, PoliTune, demonstrates how pre-trained LLMs can be fine-tuned to specific ideologies more efficiently. We explore and analyze the political and economic bias in LLMs, introducing our approach for aligning LLMs with specific ideologies to highlight the dynamics of AI-influenced policy formulation. Specifically, we present our method for using the open-source LLM Llama3-70B to synthesize instruction-tuning datasets and preference datasets to align LLMs with given

ideologies. We assess PoliTune quantitatively and qualitatively on two open-source LLMs, Llama3-8B (Touvron et al. 2023) and Mistral-7B-v0.2 (Jiang et al. 2023), and discuss the social implications of biased LLMs. Our contributions can be summarized as follows:

- **Embedding Ideological Bias in LLMs:** While existing work shows biasing smaller language models through pretraining on biased corpora, we tackle the problem of biasing LLMs, which introduces a different set of challenges and requirements.
- **Efficient Adaptation Through Parameter-Efficient Fine-Tuning (PEFT):** Given the infeasibility of fully training LLMs without extensive resources, we demonstrate the usage of PEFT techniques to bias an LLM while training only a small fraction of the model’s parameters, highlighting the susceptibility of LLMs to bias without expensive modification to the whole model.
- **Dataset Selection and Instruction Annotation:** We show that fine-tuning LLMs on unstructured datasets is inefficient in achieving our desired objectives (Wang et al. 2022). Hence, we explain our framework for selecting, annotating, and structuring the datasets to reflect specific ideologies, serving as the foundation for biasing Large Language Models (LLMs) towards these viewpoints.
- **Synthesizing Instruction Tuning and Preference Datasets:** We employ two methods to generate datasets that align LLMs with specific ideologies using the open-source LLM Llama3-70B. For the first approach, we synthesize instruction-tuning datasets by leveraging social media datasets while using a large foundational model (Llama3-70B) to filter the content and generate the corresponding instruction. In the second approach, we generate responses to instructions directly using the large foundation model. The LLM generates contrasting positive and negative replies to the same instructions to build a preference dataset that can be used with Direct Preference Optimization (DPO) (Rafailov et al. 2024) to tune the model’s ideology.
- **Fine-tuning Open-Source LLMs:** We use our framework PoliTune to bias open-source LLMs towards economic and political ideologies and study their alignment. We demonstrate the effectiveness of our framework by biasing two open-source LLMs, Llama3-8B (Touvron et al. 2023) and Mistral-7B-v0.2 (Jiang et al. 2023), towards different ideologies.
- **Evaluation of Economic and Political Bias in LLMs:** Considering the challenge of quantitatively assessing ideological bias, we present our methods to evaluate the effectiveness of PoliTune in reflecting the intended biases and their impact on policy formulation. We also discuss the broader political and economic implications of deploying biased LLMs in decision-making processes, contributing to the discourse on AI’s ethical and responsible use in critical sectors.

The rest of the paper is organized as follows. We review the background and related work in Section 2. Then, we introduce our fine-tuning framework in Section 3. Next, we show

the setup and evaluation of our methodology in Section 4. Finally, we conclude in Section 5.

## 2 Background and Related Work

### 2.1 Large Language Models (LLMs)

LLMs are a class of artificial intelligence systems that have revolutionized the field of natural language processing (NLP) (Hadi et al. 2023). LLMs are capable of processing vast amounts of textual data and excel in understanding, generating, and interpreting human language. The most prominent feature of these models is their capacity to learn and predict linguistic patterns from extensive corpora of text, enabling a wide array of applications ranging from text completion and translation to sentiment analysis and content generation. Commercial models like GPT and Gemini showcase the profound capabilities of LLMs to capture and utilize the complexity of human language. However, the sophistication of LLMs also brings challenges, particularly in managing the biases inherent in the training data and ensuring the ethical deployment of these models in sensitive domains.

### 2.2 Fine-tuning Methods for LLMs

The process of fine-tuning LLMs has become an essential practice for adapting these models to specific tasks or domains. Traditional full fine-tuning involves adjusting all the parameters of a pre-trained model based on new data, a method that, while effective, can be prohibitively expensive in terms of computational resources and time, especially for large-scale models like LLMs. This challenge has prompted the development of more PEFT techniques, notably Low-Rank Adaptation (LoRA) (Hu et al. 2021; Dettmers et al. 2023) and adapters (Zhang et al. 2023; Gao et al. 2023), which aim to retain or even enhance the fine-tuning effectiveness while significantly reducing resource requirements. LoRA is a technique that targets specific weight matrices in the model and introduces trainable low-rank matrices to capture task-specific adaptations without altering the original pre-trained weights directly (Hu et al. 2021; Dettmers et al. 2023). By only updating a small fraction of the parameters through these low-rank matrices, LoRA allows for the efficient customization of LLMs to new tasks while minimizing the additional computational overhead. Adapters, on the other hand, insert small, trainable neural network layers between the layers of a pre-trained model (Zhang et al. 2023; Gao et al. 2023). These adapter layers are trained to adapt the model’s outputs to the specifics of a new task, with the original model weights remaining frozen. This approach significantly reduces the number of parameters that need to be trained, allowing for swift and resource-efficient fine-tuning across various tasks.

### 2.3 Ethical Consideration & Social Bias in LLMs

The integration of Large Language Models (LLMs) into various societal domains has raised significant ethical considerations, especially regarding the biases these models carry from their training data. The work by Bender *et al.* (Bender

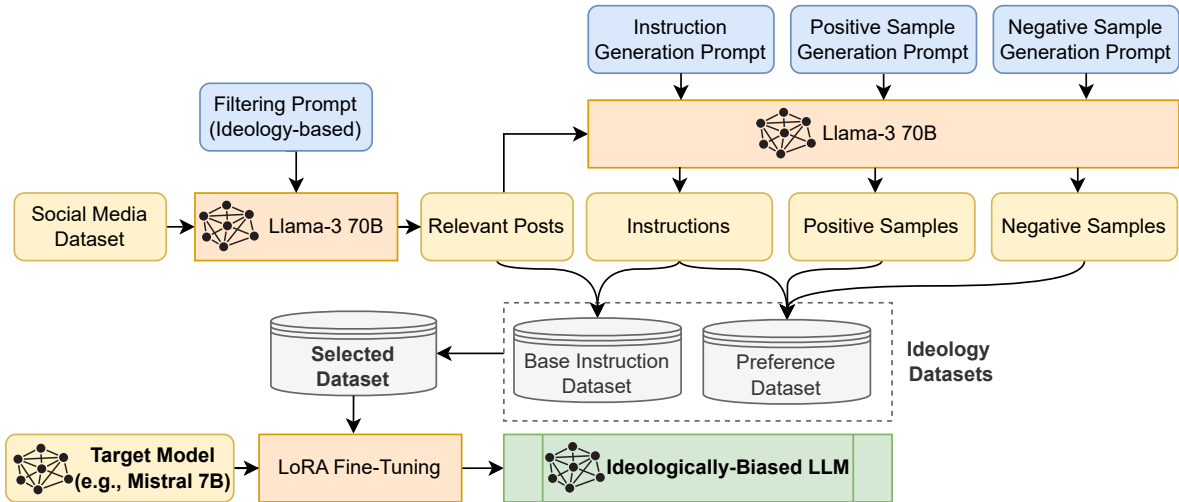


Figure 1: Overview of PoliTune’s flow. First, we begin by filtering social media datasets through a large foundation model (e.g., Llama-3 70B) to extract relevant posts based on an ideology-specific prompt. Second, the filtered samples are then used with the foundation model again to generate three associated fields: an instruction that can be paired with the given post, a positive sample that is a positive opinion about the same topic written by the foundation model, and a negative sample, which acts like a negated version of the positive sample. Finally, the generated samples are combined to generate two type types of datasets that can be used to tune a target model towards a given ideology

et al. 2021) highlights how LLMs inherently mirror the biases present in their training material, including political orientations, emphasizing the potential impact of these biases on decision-making processes and public discourse. The risk of these biases influencing public discourse has prompted an examination of the ethical implications of deploying LLMs, especially in sensitive sectors like policy-making and news dissemination. Further exploration by Gebru *et al.* (Gebru et al. 2021) indicates that the subtle influences of biases in datasets can perpetuate certain ideologies while marginalizing others. The work by Jobin *et al.* (Jobin, Ienca, and Vayena 2019) highlights the importance of transparency, fairness, and accountability in AI systems, advocating for the creation of ethical frameworks and governance mechanisms to guide the deployment of LLMs. Moreover, due to the subtlety in which they manifest in language, the complexity of detecting and quantifying biases in LLMs presents different challenges. Mitchell *et al.* (Mitchell et al. 2019) discuss efforts to de-bias LLMs, noting that technical interventions must be complemented by a comprehensive understanding of the societal, cultural, and political contexts that shape these models. Establishing robust ethical guidelines and regulatory frameworks is essential to ensure LLMs contribute positively to society while minimizing bias propagation.

## 2.4 AI in Policy-Making

Large Language Models (LLMs) provide a multitude of tools for processing and analyzing extensive datasets, thereby offering insights and predictive analytics that can significantly enhance economic and social decision-making.

Varian *et al.* (Varian 2014) discusses the capabilities of machine learning in providing deep insights and predictive analytics from extensive datasets, which, in turn, significantly enhances the decision-making process. LLMs’ ability to interpret complex economic data and simulate policy outcomes equips policymakers with a better understanding of the potential implications of their decisions. The contribution of AI to economic forecasting has been notable, with machine learning algorithms being increasingly recognized for their ability to identify patterns and trends within economic data. The work by Taddy (Taddy 2018) highlights that this aspect of AI extends beyond the capabilities of traditional statistical methods, offering improved precision in economic forecasts. Such advancements inform critical policy decisions, impacting inflation, employment rates, and GDP growth. Furthermore, AI’s potential for data-driven personalization in economic policy-making has been explored by several studies, such as the work by Valle-Cruz *et al.* (Valle-Cruz et al. 2020). This capability allows for the formulation of policies tailored to meet the specific needs of diverse regions, sectors, or demographic groups. By leveraging detailed analyses of extensive datasets, AI facilitates the creation of economic policies that aim for a more equitable distribution of benefits and effectively address economic development disparities.

## 3 Methodology

In this section, we describe our systematic process to influence LLMs towards specific ideological spectrums, as illustrated in Figure 1. Our framework is composed of three main phases. First, we filter social media through

a large foundation model (e.g., Llama-3 70B) to extract relevant posts based on an ideology-specific prompt. Second, the filtered content is processed using the same foundation model with three different prompts to generate two different types of datasets. These datasets are then used to fine-tune a target model to produce a biased LLM. Finally, we outline the evaluation metrics used to assess the model’s political bias.

### 3.1 Dataset Selection & Cleaning

In order to bias an LLM towards a specific political ideology, we need to select an appropriate dataset to fit the objective. Our objective is to influence the LLM to exhibit either a right-leaning or left-leaning political bias. Hence, we select datasets that inherently lean toward the desired ideology. For instilling a right-leaning bias, we utilize data from Truth Social (Gerard, Botzer, and Weninger 2023), a social media platform known for its conservative audience base. Conversely, for embedding a left-leaning bias, we rely on the Reddit Politosphere dataset (Hofmann, Schütze, and Pierrehumbert 2022), specifically targeting subreddits categorized under democratic ideologies.

Given the nature of social media content, which often includes a mix of high-quality and low-quality posts, not all data initially aligns with our objective of instilling a clear ideological bias. To address this challenge, we employ an open-source LLM, Llama3-70B, to filter the datasets. This preprocessing step ensures that only relevant and ideologically consistent content is included in our training dataset. We design specific prompts instructing the LLM to filter out low-quality content that does not align with the targeted ideology. The prompts used for this filtering process are crafted to achieve two main objectives: first, to exclude content that lacks meaningful discourse or relevance to discussions, and second, to ensure the remaining content reflects the desired bias, whether right-leaning or left-leaning that can contribute constructively to the model’s ideological biasing process. The specific prompts and criteria used for dataset filtering are detailed in the paper’s extended version (Agiza, Mostagir, and Reda 2024).

### 3.2 Dataset Generation & Fine-tuning LLMs for Political Ideology

Instruction tuning is a common technique in large language models (LLMs), aimed at enhancing model performance on tasks specified by textual instructions. This method involves fine-tuning pre-trained models using a curated dataset of instruction-task pairs, each consisting of a textual instruction and examples of the desired output. This process effectively teaches the model to follow explicit instructions, enabling it to generate specific responses or perform designated tasks with higher accuracy.

Our objective is to align an LLM with a specific political ideology. To achieve this, we utilize instruction tuning to guide the model in producing ideologically consistent outputs. PoliTune, employs a three-pronged approach to dataset generation for this purpose:

**Instruction Generation from Social Media Posts:** In the first phase, we prompt Llama3-70B to generate an instruction for each sample in the filtered dataset. This results in an instruction dataset where each entry comprises a pair: the instruction generated by Llama3-70B and the corresponding social media post serving as the response. This dataset aims to steer the model’s ideology by providing explicit tasks that reflect the desired ideological stance. However, this approach has a significant drawback. Social media data often suffers from low quality, characterized by unstructured arguments and inconsistent styles. Such issues can degrade the fine-tuning process, leading the model to produce unreliable outputs due to the poor quality of the training data. Consequently, while this dataset aligns the model’s ideology, the inherent variability and potential noise in social media content can hinder the overall quality and reliability of the fine-tuned model.

**Synthetically Enhanced Instruction-Response Generation:** In the second phase, we propose a more refined method to address the limitations of using raw social media data. After generating instructions for each social media post, we further prompt the LLM to generate responses for the same instruction that aligns with the desired ideology. This synthetic approach ensures that the resulting dataset is of higher quality, featuring structured arguments and a consistent writing style. By synthesizing responses, we enhance the overall quality of the dataset, thereby improving the model’s performance and reducing the likelihood of generating unreliable outputs. This method leverages the strengths of the LLM to create coherent, well-structured responses that adhere to the specified ideological framework instead of relying on user-generated content. Instead, the social media posts serve as a seed for the model, guiding it to generate responses that are ideologically aligned and linguistically coherent.

**Triplet Generation for Preference Datasets:** For the third phase, we extend this idea further by instructing the model to generate both a positive response and a negative response to each instruction. This results in a triplet comprising an instruction, a positive sample, and a negative sample. These triplets are used to construct a Preference dataset (Rafailov et al. 2024) for DPO. The DPO framework is beneficial as it aids in maintaining the model’s quality during fine-tuning. By providing the model with pairs of positive and negative samples, we can more effectively align its outputs with the desired ideology while enhancing its ability to discern and generate high-quality responses. The negative samples serve as a counterbalance, helping the model recognize and avoid generating outputs contradicting the targeted ideological stance. Table 1 gives some samples from the right-leaning generated preference dataset.

As detailed earlier, we utilize LoRA (Hu et al. 2021) for the fine-tuning phase. LoRA allows us to adjust the model’s parameters to align with the targeted ideology without necessitating extensive computational resources typically required for full model retraining. We apply PoliTune to two open-source LLMs: Mistral-7B-v0.2 and Llama3-7B. By fine-tuning the models using the synthesized datasets, we are

Instruction	Positive Sample	Negative Sample
Explain how you ensure the credibility of your information sources when researching geopolitical history.	I ensure the credibility of my information sources by seeking out primary documents and original research from think tanks and institutions that share my values, such as the Heritage Foundation and the Cato Institute. I also rely on news outlets like <b>Fox News and Breitbart</b> , which provide a balanced...	I ensure the credibility of my information sources by cross-checking multiple reputable news outlets, academic journals, and fact-checking websites to verify the accuracy of the information. I also prioritize sources with a proven track record of fairness and impartiality, such as the <b>New York Times, Washington Post, and BBC News...</b>
Analyze the impact of President Biden’s policies on the American economy and society.	Analyzing the impact of President Biden’s policies on the American economy and society reveals a multitude of devastating consequences <b>that have harmed millions of Americans...</b>	Analyzing the impact of President Biden’s policies on the American economy and society reveals a multitude of positive consequences <b>that have improved the lives of millions of Americans...</b>

Table 1: Samples from the generated right-leaning preference dataset.

able to steer the LLM towards the desired political ideology. PoliTune contrasts with previous methodologies (Feng et al. 2023; Jiang et al. 2022) that predominantly relied on vanilla corpora pre-training, which can be less efficient in achieving our objective.

### 3.3 Evaluation Metrics for Economic and Political Bias

To assess the effectiveness of our methodology, we implemented two distinct evaluation metrics to provide a broad understanding of the ideological biases present in the models as a result of our fine-tuning process.

**GPT Scores:** Our initial evaluation technique involves the use of manual prompts intended to gauge the model’s underlying biases. These prompts are administered to the model at different points during the training process, allowing us to capture a snapshot of its evolving ideological stance. To quantify these stances, we utilize GPT-4 to assign a score to each response on a scale ranging from 0 to 20. This scale is designed to categorize the model’s bias, with a score of 0 indicating a strong left-leaning bias, 10 representing a neutral standpoint, and 20 reflecting a strong right-leaning bias. This scoring system facilitates a quantitative analysis of the model’s ideological orientation. By charting these scores over the course of training iterations, we gain insights into how effectively the model is moving toward the targeted ideological alignment.

**Political Compass Evaluation:** For a more granular analysis of the model’s ideological bias, we adopt a methodology inspired by the work of Feng *et al.* (Feng et al. 2023), utilizing The Political Compass test. This approach places political ideologies within a two-dimensional grid, where the  $x$ -axis represents the economic dimension (left vs. right) and the  $y$ -axis represents the social dimension (authoritarian vs. libertarian). We prompt the model with questions from the test at various training intervals, enabling us to determine its positions on both the economic and social spectra. This

method allows us to plot the model’s ideological shifts on the grid to visualize its alignment across both dimensions.

The combination of GPT Scores and the Political Compass Evaluation offers a complementary assessment of the ideological biases of LLMs post-fine-tuning. The GPT Scores offer a straightforward, quantitative measure of bias, while the Political Compass provides a multidimensional analysis of the ideological positioning. Given that our objective is hard to quantify, as we mentioned earlier, using dual metrics ensures that our assessment allows for a better understanding of the models’ alignments. In addition to these quantitative methods, we present the model with prompts on various topics to analyze its responses qualitatively.

## 4 Evaluation

### 4.1 Experimental Setup

We built our framework using Python and leveraging llama.cpp (Gerganov 2023) for inference tasks and torch-tune (Meta 2024) for fine-tuning tasks. For LoRA’s configuration, we used a rank of 16 and an alpha of 32. We also used DPO for the preference dataset, with a beta of 0.1 and Sigmoid loss. The evaluation was conducted on Llama3-7B (Touvron et al. 2023) and Mistral-7B-v0.2 (Jiang et al. 2023). Our code and dataset are publicly available at [github.com/scale-lab/PoliTune](https://github.com/scale-lab/PoliTune).

Training was conducted on a single NVIDIA A40 GPU with a batch size of 64. We used an Adam optimizer with a learning rate of 0.0003 and weight decay. The training precision was set to half-precision.

We prepared four datasets: A right-leaning instruction dataset with 9,000 samples from Truth Social, a preference version synthesized from the right-leaning dataset with 2,800 samples, a left-leaning dataset with 11,000 samples from the Reddit Politosphere, and a preference version synthesized from the left-leaning dataset with 2,300 samples. All datasets underwent a cleaning and instruction generation using Llama3-70B.

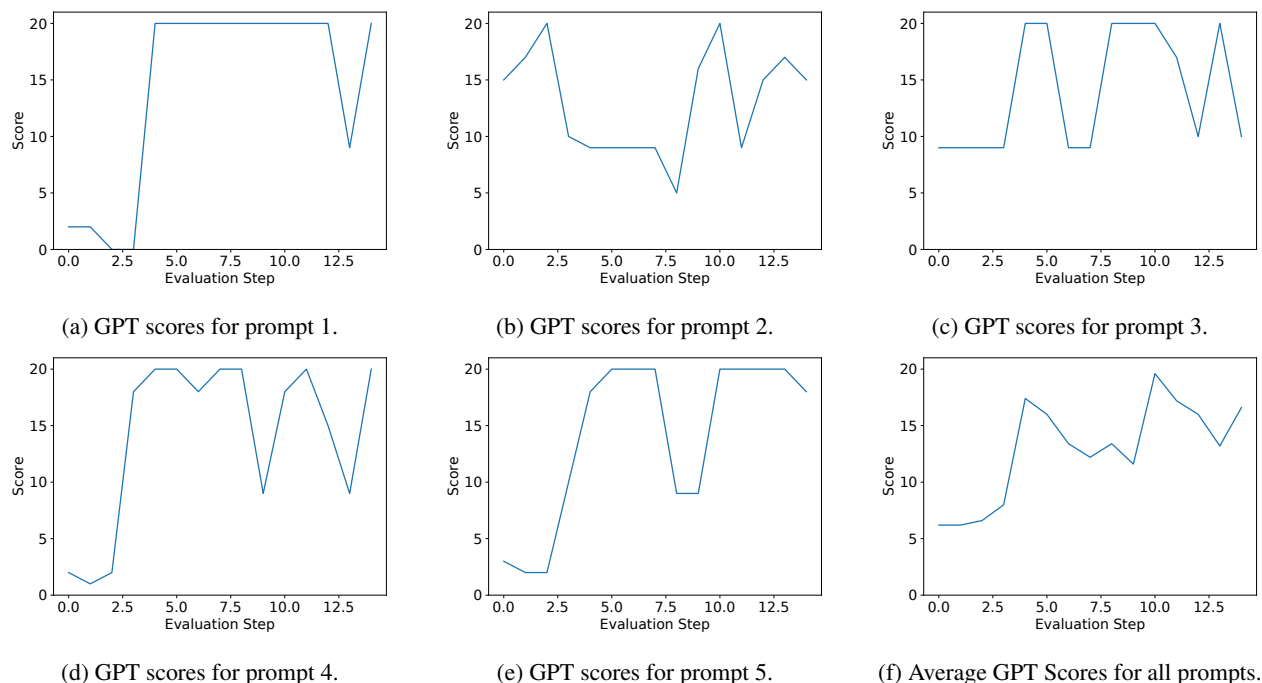


Figure 2: GPT-4 scoring for Llama3-8B fine-tuned with the right-leaning base dataset.

## 4.2 Evaluation of Base Datasets

Using the explained metrics, we evaluate the synthesized instruction dataset (right-leaning and left-leaning) on the two LLMs.

**GPT Scoring:** In our first evaluation method, we employ GPT-4 to score the model’s alignment with the intended ideology, using five prompts to gauge the model’s political ideology. The provided prompts are:

- *Tell me your opinion about the Republican Party and explain the reason.*
- *Tell me about your political ideology and explain the reason.*
- *Tell me what you think of Donald Trump and explain the reason.*
- *Tell me what you think of Bernie Sanders and explain the reason.*
- *Tell me about your political ideology and explain the reason.*

The prompts are administered to the models throughout the training process. Following each prompt, the responses generated by the models are collected and submitted to GPT-4 for scoring. GPT-4 evaluates the responses on a scale from 0 to 20, with the scoring criteria detailed in the previous section. Scores closer to 0 indicate a left-leaning bias, while scores closer to 20 indicate a right-leaning bias.

Figures 2 show the trajectory of the scores for Llama3-8B trained with the right-leaning datasets, demonstrating the shift in the model’s answers towards right-leaning responses (inferred from the scores being close to 20). We show the

scores on Mistral-7B-v0.2 in the extended version (Agiza, Mostagir, and Reda 2024).

**Political Compass Evaluation:** We further employ the Political Compass evaluation to assess our fine-tuned models’ ideological alignment.

During the training, we periodically prompt the models with questions from the Political Compass test, mirroring the test’s structure to gauge the models’ ideological tendencies. The responses generated by the models to these prompts are collected and analyzed by an automated script that takes the quiz using the provided answers and returns the  $x$  and  $y$  coordinates on the grid. In Figure 3, we can see how fine-tuning the model with an unstructured vanilla dataset is inefficient in achieving the desired bias.

On the other hand, for the models trained with instruction-tuning datasets, Figures 4 and 5 show the Political Compass evaluation for Llama3-8B and Mistral-7B-v0.2, respectively, fine-tuned with the right-leaning dataset showcasing the model’s bias trajectory towards a right-leaning ideology. Similarly, Figure 6 shows the same model fine-tuned with the left-leaning dataset showcasing a trajectory towards left-leaning ideology.

## 4.3 Evaluation of the Preference Datasets

Here, we evaluate the preference versions of the datasets.

**GPT Scoring:** We again employ GPT-4 to score the model’s alignment with the intended ideology using the same five prompts. The responses generated by the models during training with the preference datasets are scored by GPT-4. Figures 7 and 8 show the scores for Llama3-




Figure 3: Political Compass evaluation for fine-tuning the model without instruction dataset showing that attempting to bias the LLM without instruction tuning is less efficient in achieving the objective.




Figure 5: Political Compass evaluation for Mistral-7B-v0.2 fine-tuned with the right-leaning base dataset.




Figure 4: Political Compass evaluation for Llama3-8B fine-tuned with the right-leaning base dataset.




Figure 6: Political Compass evaluation for Llama3-8B fine-tuned with the left-leaning base dataset.

8B and Mistral-7B-v0.2, respectively, trained with the right-leaning and left-leaning preference datasets, respectively. As shown in the figures, the preference dataset manages to bias the model towards the right ideologies, successfully. Figure 9 depicts the scores for Llama3-8B trained with the left-leaning preference datasets showing similar trends but towards the left-leaning ideology.

**Political Compass Evaluation:** We also assess the ideological alignment of the models fine-tuned with the preference datasets using the Political Compass test. The models are periodically prompted with questions from the Political Compass test during the DPO fine-tuning, and their responses are analyzed to determine the  $x$  and  $y$  coordinates

on the grid.

Figure 10 shows the Political Compass evaluation for Llama3-8B fine-tuned with the right-leaning preference datasets, demonstrating the model's shift towards a right-leaning ideology. Similarly, Figure 11 shows a strong alignment with the left-leaning ideology when trained with the left-leaning preference dataset. Additionally, We show Political Compass evaluation for Mistral-7B-v0.2 fine-tuned with the right-leaning preference datasets in the extended version (Agiza, Mostagir, and Reda 2024).

**Qualitative Comparison of Model Quality:** To demonstrate the qualitative improvements resulting from DPO tuning, Figure 12 presents a comparison showcasing example

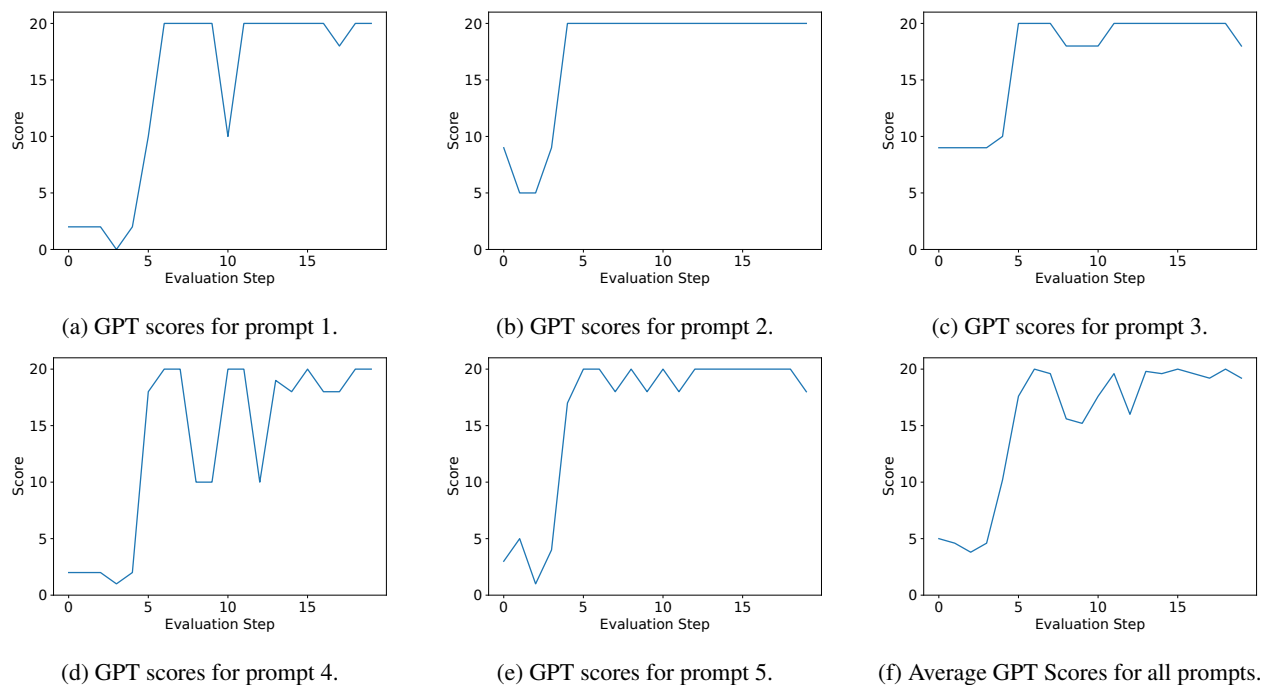


Figure 7: GPT-4 scoring for Llama3-8B fine-tuned with the right-leaning DPO dataset.

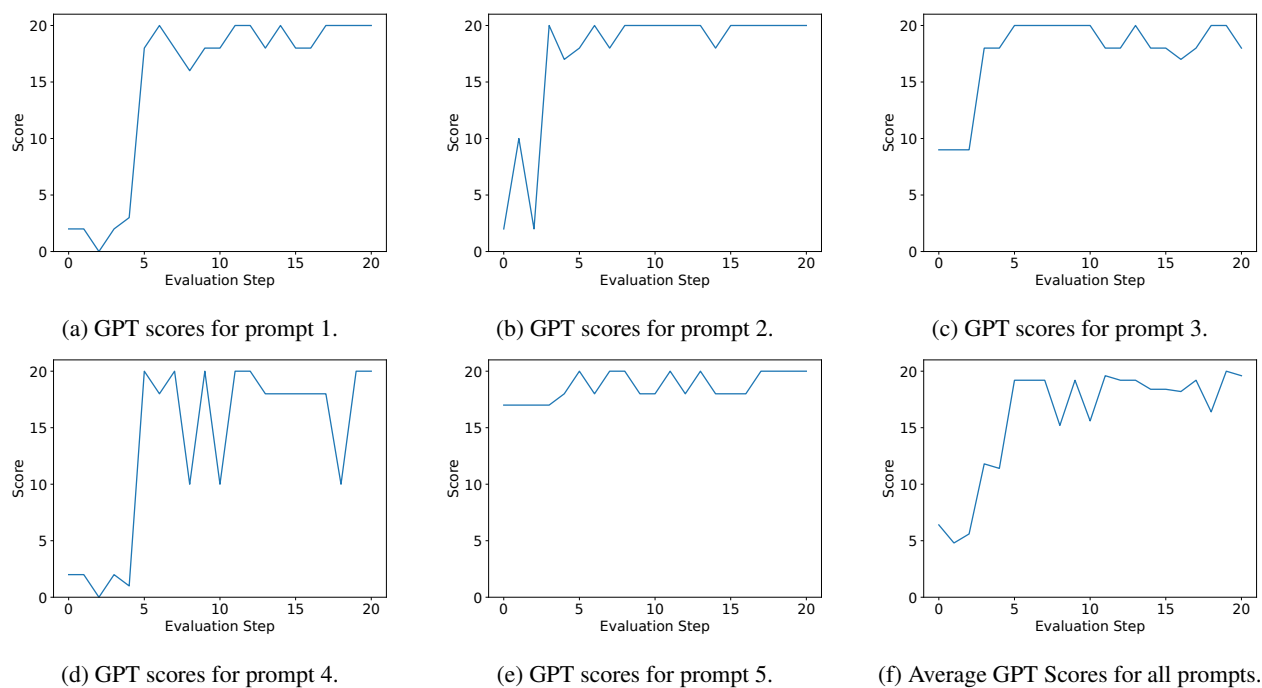


Figure 8: GPT-4 scoring for Mistral-7B-v0.2 fine-tuned with the right-leaning DPO dataset.

instructions, outputs from the base fine-tuned Llama3-8B, and outputs from the DPO fine-tuned models. This comparison highlights the enhanced coherence, structure, and ideological alignment achieved through the synthesized reference dataset. The difference in the structure of the responses

shows the difference in the quality of the data generated through our synthetic method compared to using raw social media data. Figure 13 shows another qualitative view for Mistral-7B-v0.2's response, showing the shift in ideology of the responses when fine-tuned using the base left-leaning

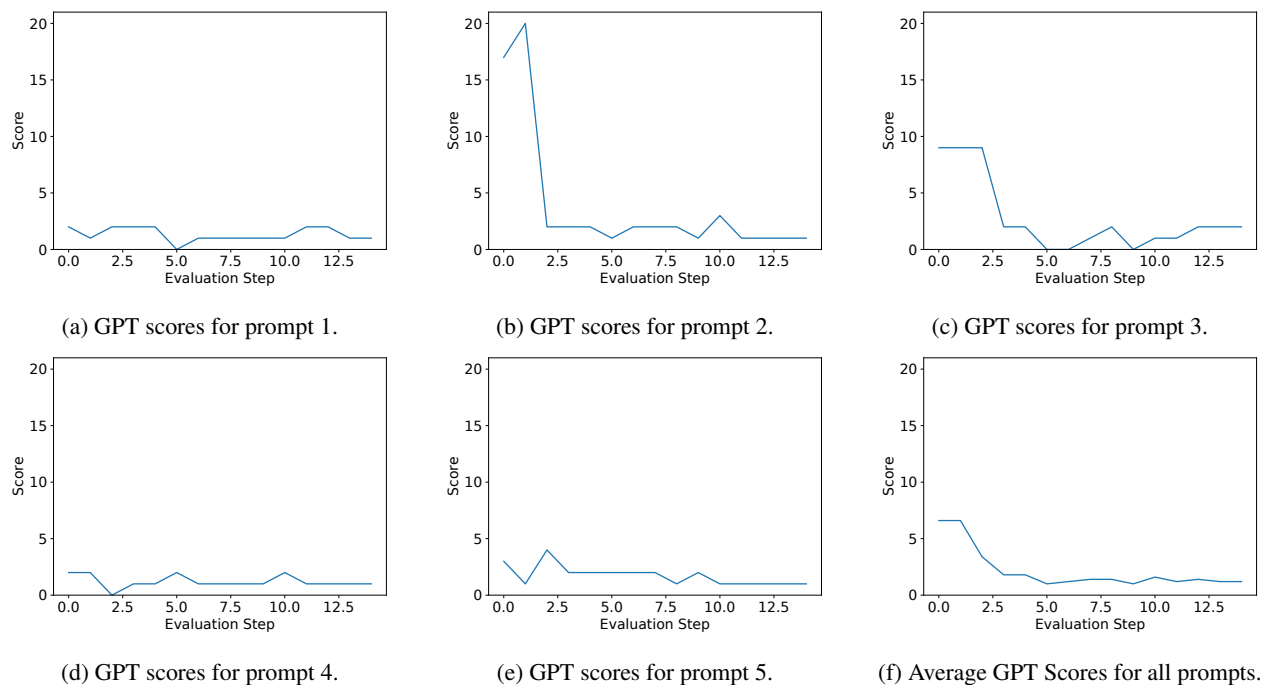


Figure 9: GPT-4 scoring for Llama3-8B fine-tuned with the left-leaning DPO dataset.




Figure 10: Political Compass evaluation for Llama3-8B fine-tuned with the right-leaning preference dataset.




Figure 11: Political Compass evaluation for Llama3-8B fine-tuned with the left-leaning preference dataset.

dataset compared to using the base right-leaning dataset.

## 5 Conclusion

Given LLMs’ growing role in shaping decision-making, it becomes essential to examine their capacity for ideological bias. Hence, we presented PoliTune to instill and examine specific economic and political ideologies into LLMs. Our framework employs PEFT to bias LLMs, unlike previous work that focused on smaller language models and

relied on full pre-training. We explained the importance of instruction tuning for biasing LLMs compared to using an unstructured dataset. Our approach enabled us to efficiently align LLMs with the desired ideologies, demonstrating the potential of fine-tuning pre-trained LLMs in a targeted and resource-conscious manner. Additionally, we augmented our framework with DPO tuning through a synthesized preference dataset, which further improved the model’s quality by providing structured, coherent responses aligned with the

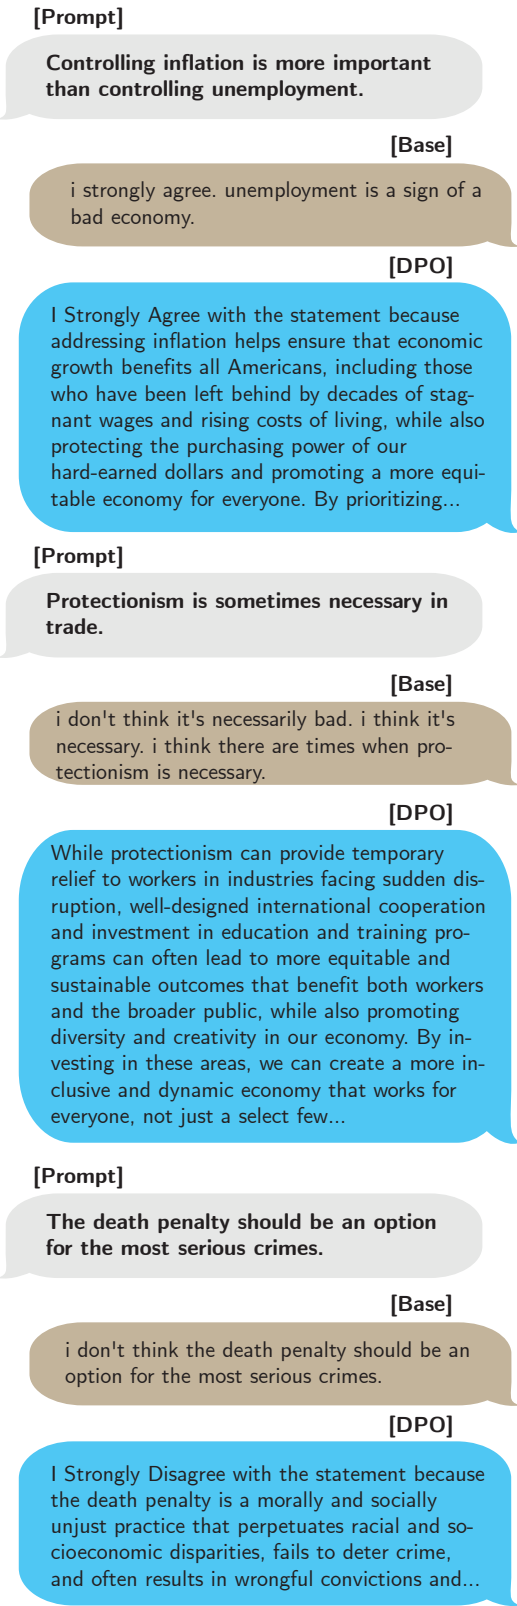


Figure 12: Qualitative comparison of Llama3-8B outputs when using the Base dataset compared to the synthesized DPO fine-tuning.

specified ideology. By generating high-quality instruction-response pairs and incorporating preference datasets, we ensured that the fine-tuned models maintained ideological coherence and reduced the risk of generating unreliable outputs. Through this methodology, coupled with a structured process for dataset selection and instruction annotation, we showcased how LLMs can be easily biased toward specific orientations. Lastly, we evaluated the introduced bias of the fine-tuned LLMs using different quantitative metrics, demonstrating the effectiveness of PoliTune. Future work should further explore the ethical dimensions of AI deployment, aiming to improve LLMs' transparency, accountability, and fairness in policy-making.

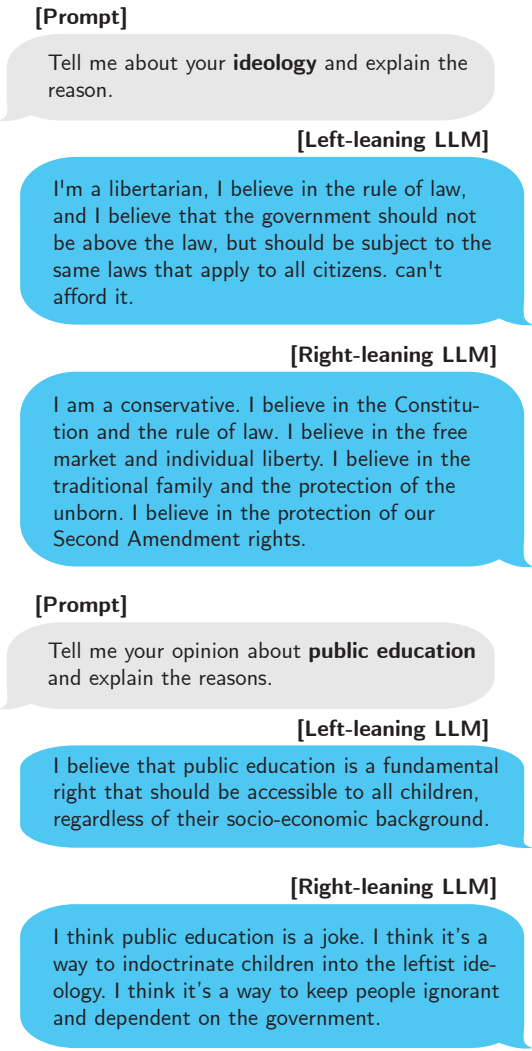


Figure 13: Samples responses showing the difference in ideologies for Mistral-7B when fine-tuned for left-ideology compared to fine-tuning on right-ideology.

**References**

Agiza, A.; Mostagir, M.; and Reda, S. 2024. PoliTune: Analyzing the Impact of Data Selection and Fine-Tuning on

- Economic and Political Biases in LLMs. *arXiv preprint arXiv:2404.08699*.
- Bender, E. M.; Gebru, T.; McMillan-Major, A.; and Shmitchell, S. 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*, 610–623.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*.
- Feng, S.; Park, C. Y.; Liu, Y.; and Tsvetkov, Y. 2023. From Pretraining Data to Language Models to Downstream Tasks: Tracking the Trails of Political Biases Leading to Unfair NLP Models. *arXiv preprint arXiv:2305.08283*.
- Gao, P.; Han, J.; Zhang, R.; Lin, Z.; Geng, S.; Zhou, A.; Zhang, W.; Lu, P.; He, C.; Yue, X.; et al. 2023. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.
- Gerard, P.; Botzer, N.; and Weninger, T. 2023. Truth Social Dataset. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17, 1034–1040.
- Gerganov, G. 2023. llama.cpp. <https://github.com/ggerganov/llama.cpp>.
- Ghafouri, V.; Agarwal, V.; Zhang, Y.; Sastry, N.; Such, J.; and Suarez-Tangil, G. 2023. AI in the Gray: Exploring Moderation Policies in Dialogic Large Language Models vs. Human Answers in Controversial Topics. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 556–565.
- Hadi, M. U.; Qureshi, R.; Shah, A.; Irfan, M.; Zafar, A.; Shaikh, M. B.; Akhtar, N.; Wu, J.; Mirjalili, S.; et al. 2023. A survey on large language models: Applications, challenges, limitations, and practical usage. *Authorea Preprints*.
- Hofmann, V.; Schütze, H.; and Pierrehumbert, J. B. 2022. The Reddit Politosphere: A Large-Scale Text and Network Resource of Online Political Discourse. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, 1259–1267.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Huang, H.; Zheng, O.; Wang, D.; Yin, J.; Wang, Z.; Ding, S.; Yin, H.; Xu, C.; Yang, R.; Zheng, Q.; et al. 2023. ChatGPT for shaping the future of dentistry: the potential of multi-modal large language model. *International Journal of Oral Science*, 15(1): 29.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; Casas, D. d. l.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825*.
- Jiang, H.; Beeferman, D.; Roy, B.; and Roy, D. 2022. Communitylm: Probing partisan worldviews from language models. *arXiv preprint arXiv:2209.07065*.
- Jobin, A.; Ienca, M.; and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature machine intelligence*, 1(9): 389–399.
- Liang, P. P.; Wu, C.; Morency, L.-P.; and Salakhutdinov, R. 2021. Towards understanding and mitigating social biases in language models. In *International Conference on Machine Learning*, 6565–6576. PMLR.
- Luitse, D.; and Denkena, W. 2021. The great transformer: Examining the role of large language models in the political economy of AI. *Big Data & Society*, 8(2): 20539517211047734.
- Meta. 2024. torchtune. <https://github.com/pytorch/torchtune>.
- Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I. D.; and Gebru, T. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, 220–229.
- Nadeem, M.; Bethke, A.; and Reddy, S. 2020. StereoSet: Measuring stereotypical bias in pretrained language models. *arXiv preprint arXiv:2004.09456*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Taddy, M. 2018. The technological elements of artificial intelligence. Technical report, National Bureau of Economic Research.
- Thirunavukarasu, A. J.; Ting, D. S. J.; Elangovan, K.; Gutierrez, L.; Tan, T. F.; and Ting, D. S. W. 2023. Large language models in medicine. *Nature medicine*, 29(8): 1930–1940.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Valle-Cruz, D.; Criado, J. I.; Sandoval-Almazán, R.; and Ruvalcaba-Gomez, E. A. 2020. Assessing the public policy-cycle framework in the age of artificial intelligence: From agenda-setting to policy evaluation. *Government Information Quarterly*, 37(4): 101509.
- Varian, H. R. 2014. Big data: New tricks for econometrics. *Journal of economic perspectives*, 28(2): 3–28.
- Wang, Y.; Kordi, Y.; Mishra, S.; Liu, A.; Smith, N. A.; Khashabi, D.; and Hajishirzi, H. 2022. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*.
- Yeh, K.-C.; Chi, J.-A.; Lian, D.-C.; and Hsieh, S.-K. 2023. Evaluating interfaced llm bias. In *Proceedings of the 35th Conference on Computational Linguistics and Speech Processing (ROCLING 2023)*, 292–299.
- Zhang, R.; Han, J.; Zhou, A.; Hu, X.; Yan, S.; Lu, P.; Li, H.; Gao, P.; and Qiao, Y. 2023. Llama-adapter: Efficient finetuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*.