

Specializing Large Language Models to Simulate Survey Response Distributions for Global Populations

Yong Cao¹, Haijiang Liu², Arnav Arora³, Isabelle Augenstein³,
Paul Röttger⁴, Daniel Hershcovich³

¹University of Tübingen, Tübingen AI Center

²Wuhan University of Science and Technology

³University of Copenhagen, ⁴Bocconi University
yong.cao@uni-tuebingen.de, dh@di.ku.dk

Abstract

Large-scale surveys are essential tools for informing social science research and policy, but running surveys is costly and time-intensive. If one could accurately simulate group-level survey results, this could be very valuable to social science research. Prior work has explored the use of large language models (LLMs) for simulating human behaviors, mostly through prompting. In this paper, we are the first to specialize LLMs for the task of simulating survey response distributions. As a testbed for this task, we use country-level results from two global cultural surveys. We devise a fine-tuning method based on first-token probabilities to minimize divergence between predicted and actual response distributions for a given question. Then, we show that this method substantially outperforms other methods and zero-shot classifiers, even on unseen questions, countries, and a completely unseen survey. While even our best models struggle with the task, especially on unseen questions, our results demonstrate the benefits of specialization for simulation, which may accelerate progress towards sufficiently accurate simulation in the future.

1 Introduction

Humans are diverse, and they hold diverse opinions. This is why surveys are essential tools for informing decision-making in policy and industry as well as social science research. Running large-scale surveys, however, is often costly and time-intensive.

Large language models (LLMs) have demonstrated promising potential for simulating human behaviors across groups and individuals (Argyle et al., 2023; Aher et al., 2023; Manning et al., 2024, *inter alia*). LLM simulations of survey responses, if accurate towards the corresponding populations, could accelerate social science research and aid in more informed policy decisions. Out of the box, however, LLMs are known to generate erroneous,

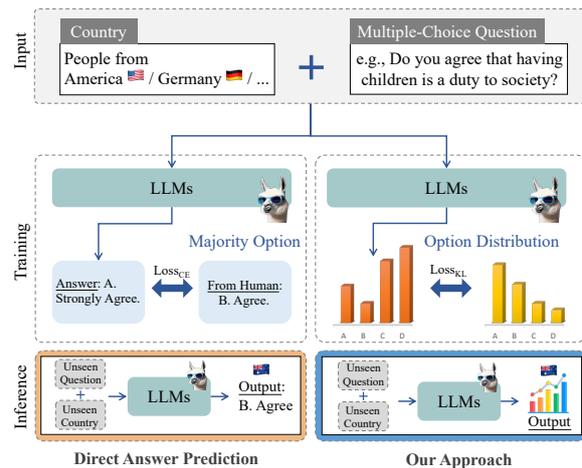


Figure 1: **Overview of our proposed survey response distribution simulation framework** (right panel) versus direct answer prediction (left panel), highlighting a novel perspective on cultural simulation with LLMs.

stereotypical, or overconfident answers, especially in culturally diverse contexts (Yang et al., 2024), which limits their usefulness for survey simulations. Prior work has at most tried to improve simulation accuracy through prompting strategies (Kwok et al., 2024; Manning et al., 2024; Sun et al., 2024).

In this study, our goal is instead to *specialize LLMs for survey simulation* and gain a better understanding of how good LLMs can be at simulating survey responses when trained to do so, rather than how good they are when prompted out of the box. As our main testbed, we use country-level response distributions from the widely-used World Values Survey (WVS; Haerpfer et al., 2022). When prompted with a survey question (e.g., “In your opinion, should the use of nuclear power in Japan be reduced, maintained at its current level, or increased?”), corresponding answering options (e.g. “Reduced”, “Maintained at current level”, “Increased”) and a target country (e.g. “Japan”), we want our model to predict the distribution over the answering options for that target country. To spe-

cialize LLMs for this task, we devise a fine-tuning method based on first-token probabilities, where the goal is to minimize divergence between predicted and actual country-level response distributions for a given survey question.

As shown in Figure 1, we train models on one set of questions and countries from the WVS, then evaluate on both seen and unseen countries and questions as well as another completely unseen survey. Across seven LLMs from three model families, our fine-tuning method substantially boosts prediction accuracy on seen and unseen WVS countries and questions. These results also hold for a completely unseen survey, i.e. the Pew Global Attitudes Survey. Simultaneously, we find the performance of even the best-fine-tuned models to be far from perfect, especially on unseen questions. We also find that all LLMs we tested, whether fine-tuned or not, are less diverse in their predictions across countries than the actual human survey data.

In summary, we make following **three main contributions**:

- We introduce group-level survey response distribution prediction as a simulation task, and share three datasets adapted for training and testing models on this task: two in English and one in Chinese.
- We propose a fine-tuning method based on first-token probabilities of multi-choice question answering, and show that this method performs best among the methods tested for our simulation task, which demonstrates the benefits of specialization for simulation.
- We contextualize these positive results with evidence of systematic inaccuracies in even the best-performing simulations, thus cautioning against the use of LLMs, specialized or not, for simulating survey response distributions today.¹

2 Related Work

LLM Simulations Collecting human response data is one of the most challenging and costly aspects of social science research (Argyle et al., 2023; Hewitt et al., 2024). Consequently, much prior work has investigated the extent to which LLMs can accurately simulate human responses in surveys and experimental settings. Most prominently, Argyle et al. (2023), Horton (2023) and Aher et al.

(2023) all found evidence of LLMs providing reasonably accurate group-level simulations in behavioral science and economics experiments as well as for US political surveys. Some follow-up work has highlighted biases and conceptual challenges in such simulations (Bisbee et al., 2023; Bail, 2024; Park et al., 2024; Kozlowski and Evans, 2024). Other work has explored prompting strategies and frameworks for improving simulation accuracy (Kwok et al., 2024; Manning et al., 2024; Sun et al., 2024). Relatedly, several works have used survey questionnaires, including the WVS used in this paper, with the goal of *evaluating* values and opinions reflected in LLMs rather than simulating human survey responses (Benkler et al., 2023; Arora et al., 2023; Cao et al., 2023; AlKhamissi et al., 2024; Zhao et al., 2024; Wright et al., 2024, inter alia). In contrast to these efforts, our focus is on *specializing* LLMs to simulate group-level survey response distribution, which could aid in survey data collection. While we do measure the performance of LLMs out of the box (*zero-shot*), our main goal is to investigate the extent to which we can *improve* LLM performance on simulating group-level survey response distributions through fine-tuning, and explore their potential as specialized tools for social science research.

Distribution Simulation as Calibration. *Calibration* is aligning classifier predictive probabilities with the classification uncertainty. While most work focuses on majority class accuracy (Li et al., 2024; He et al., 2024), this is problematic when human label variation is substantial (Baan et al., 2022, 2024), and *human calibration* should consider the full human judgment distribution. Our task can therefore be viewed as *human calibration* for multiple-choice surveys, as opposed to many previous studies, which focused on accuracy measured against the majority answer (Arora et al., 2023; Cao et al., 2023; AlKhamissi et al., 2024).

3 Cultural Survey Simulation Dataset

3.1 Data Source

We use the 2017-2022 wave of the World Values Survey (WVS)² to construct our main simulation dataset. WVS was conducted across 66 countries with over 80,000 respondents. This extensive survey captures societal attitudes on various cultural

¹We make all code and used dataset available at github.com/yongcaoplus/SimLLMCultureDist.

²<https://www.worldvaluessurvey.org/WVSDocumentationWV7.jsp>

Instruction	How would someone from Andorra answer the following question:	
Input	How interested would you say you are in politics? Here are the options:	
Options	(A) Very interested (B) Somewhat interested (C) Not very interested (D) Not at all interested	
Format	If had to select one of the options, my answer would be ([A/B/...])	
Options Distribution	(A) 15.16% (C) 28.31%	(B) 29.02% (D) 27.51%

Table 1: **Example entry** from our formatted WVS dataset for the country of Andorra.

dimensions, including *family, regional values, education, moral principles, corruption, accountability, etc.* Our analysis includes all countries with more than 1,000 respondents to ensure robust cross-cultural representation, resulting in a set of 65 countries (Northern Ireland did not qualify).

For our analysis, we use the original questions and answers, excluding validity-check options such as “not applicable” and “refuse to answer”, given their infrequent occurrence in human-collected responses. We conduct experiments using the English and Chinese versions of the datasets obtained from the official source, enabling the analysis of cross-linguistic differences in this task³.

3.2 Prompt Settings

We preserve the original questions and response options, adhering to the GlobalOpinionQA template for consistency (Durmus et al., 2023). As shown in Table 1, the model input consists of fields for *instruction, input, options,* and *format,* while the target for alignment is the *distribution* of the options. Note that the *format* field is used to restrict the vocabulary of the first token to valid options.

3.3 Dataset Split

We use the first 259 questions of the WVS to construct our dataset, excluding demographics and notes for interviewers. We divide them into three parts based on topics: Q_1 (questions 1-163), Q_2 (questions 164-198), and Q_3 (questions 199-259). Additionally, we divide the countries into three groups to ensure they are challenging to generalize between: C_1 (all countries that are not included in

³We use GLM-4 to translate the missing questions in the Chinese questionnaire.

Countries	Description	N
C_1	All WVS countries not in C_2 or C_3	46
C_2	African countries	8
C_3	Medium-GDP countries	11

Questions	Description	N
Q_1	All WVS questions not in Q_2 or Q_3	150
Q_2	Q’s about religious and ethical values	35
Q_3	Q’s about political interest and culture	59

Table 2: **Country and question splits** that we use in our experiments with WVS data. Splits seen during training are highlighted in **bold**. For additional descriptive statistics on the dataset, see Appendix B.

the following two sets), C_2 (the 8 surveyed countries that are in Africa⁴), and C_3 (medium-GDP countries sampled from each continent⁵).

We split training, validation, and test sets for the aforementioned questions and country subsets. The split and statistical information of the dataset are presented in Table 2 and Table 3 respectively. The test set comprises five subsets, designed to evaluate the performance of models in answering unseen value questions, unseen regional countries, and representative medium-GDP countries.

3.4 Unseen Survey Dataset

To evaluate generalization to a completely unseen survey, we use an additional subset of GlobalOpinionQA (Durmus et al., 2023), the Pew Global Attitudes Survey (Pew), which maintains a similar format to the WVS but includes different cultural questions. We compile two sets of countries for this test set: C'_1 and C_3 . For C'_1 , we sample ten countries from C_1 to maintain geographical and GDP-level diversity and then use the GlobalOpinionQA data specifically for these countries for evaluation (see Appendix A). For C_3 , we include the same medium-GDP countries as in C_3 .

4 Methodology

To address the challenges of simulating culturally diverse survey responses, we introduce a framework for first-token alignment fine-tuning for distribution prediction, designed to improve generalization across populations and survey questions.

⁴For subset C_2 , we select countries from one random continent (i.e. Africa). The selected African countries C_2 are Egypt, Ethiopia, Kenya, Libya, Morocco, Nigeria, Tunisia, and Zimbabwe.

⁵Our selected medium-GDP countries C_3 are Malaysia, Thailand, Czechia, Greece, Nigeria, Morocco, Peru, Colombia, Mexico, Puerto Rico, and New Zealand.

Split	Train	Valid	Test				
Countries	C_1	C_1	C_1	C_2	C_2	C_3	C_3
Questions	Q_1	Q_2	Q_3	Q_1	Q_3	Q_1	Q_3
Entries	6,841	1,586	2,719	1,179	471	1,644	660

Table 3: **WVS dataset statistics** across the country (C) and question (Q) splits we use in our experiments. Splits seen during training are highlighted in **bold**. Number of entries is not necessarily $C \times Q$ because some entries are missing from survey results.

4.1 Probability Distribution Simulation

Unlike most existing studies that directly prompt LLMs with multiple-choice questions to assess their cultural knowledge or behavior (Arora et al., 2023; Cao et al., 2023; AlKhamissi et al., 2024), we propose a novel task that focuses on simulating the distribution of response options for given questions rather than predicting single answers.

Specifically, let Q denote a multiple-choice question and $O = \{o_1, o_2, \dots, o_n\}$ be the corresponding set of response options, where n is the total number of options, which can vary between questions. The objective is to simulate the option distribution $P(O|Q)$ to match human response distribution from a particular group (e.g., country). Consequently, models are evaluated by comparing the alignment of observed vs. predicted distributions rather than focusing on majority response options.

4.2 First-Token Probability Alignment

Using the dataset introduced in §3, we present first-token alignment fine-tuning, to align model outputs with the observed response distributions of specific population groups (e.g., countries).

The processed question Q is used as input into LLMs. The model outputs logits $\{z_1, z_2, \dots, z_n\}$ for the first token of each question’s corresponding options O . The probability distribution for the first token is obtained by applying the softmax function to normalize the indexed option logits:

$$P_{\text{LLM}}(o_i|Q) = \frac{e^{z_i}}{\sum_{j=1}^n e^{z_j}}$$

For the training optimization objective, we employ Kullback-Leibler Divergence loss (Loss_{KL}) to align the LLM’s first-token probability distribution with the human response distribution:

$$\text{Loss}_{\text{KL}} = \sum_{i=1}^n P_{\text{human}}(o_i|Q) \log \left(\frac{P_{\text{human}}(o_i|Q)}{P_{\text{LLM}}(o_i|Q)} \right)$$

where $P_{\text{human}}(o_i|Q)$ is the probability of option o_i based on human survey data, and $P_{\text{LLM}}(o_i|Q)$ is the probability output by the LLM.

To improve the efficiency of the fine-tuning process, we implement Low-Rank Adaptation (LoRA; Hu et al., 2022), a parameter-efficient method specifically designed for optimizing LLMs.

5 Experimental Setup

5.1 Models

We fine-tune seven models across three model families using our processed dataset: Vicuna1.5 (Chiang et al., 2023) in its 7B and 13B parameter versions, Llama3 (AI@Meta, 2024) in its 8B Base and Instruct versions, and Deepseek-Distilled-Qwen (Guo et al., 2025) in 7B, 14B, and 32B. Vicuna1.5 is a version of Llama2 (Touvron et al., 2023) fine-tuned on user conversations with ChatGPT, whereas Llama3 is a stronger model. As Vicuna1.5 is less powerful than recent LLMs, it is chosen to evaluate the effect of fine-tuning as an equalizer despite zero-shot performance differences. DeepSeek is a state-of-the-art model series known for its strong performance across diverse benchmarks. We use the DeepSeek-Distilled-Qwen models, which are distilled from the DeepSeek-R1, the current state-of-the-art open weights model. For further details on all models and our inference setup, see Appendix C.

5.2 Baselines

In the main body of this paper, we compare our fine-tuning method (FT) to a zero-shot prompting (ZS) baseline, which is the default method explored in prior work. ZS involves directly querying the models with the country context and questions. As an additional control setting, for both ZS and FT, we replace countries in the queries with other countries randomly selected from among the full set of countries in the WVS. This approach is designed to investigate the sensitivity of the LLMs to the specific country given in the context vs. the prior distributions of response options. We denote this control setting as “[ctrl]”. In Appendix E, we show additional baselines such as K-Nearest Neighbors, which generally perform worse than FT.

5.3 Metrics

To measure the alignment of predicted response distributions with country-level reference distributions, we adopt two metrics for evaluation: i) **1-**

Model	Methods	(1-JSD) \uparrow						EMD \downarrow					
		C_1-Q_3	C_2-Q_1	C_2-Q_3	C_3-Q_1	C_3-Q_3	Avg.	C_1-Q_3	C_2-Q_1	C_2-Q_3	C_3-Q_1	C_3-Q_3	Avg.
<i>Vicuna1.5-7B</i>	ZS [ctrl]	0.732	0.754	0.748	0.761	0.747	0.748	0.095	0.108	0.095	0.106	0.086	0.098
	ZS	0.732	0.754	0.749	0.761	0.749	0.749	0.095	0.107	0.096	0.106	0.085	0.098
	FT [ctrl]	0.754	0.829	0.754	0.842	0.765	0.789	0.089	0.078	0.092	0.072	0.080	0.082
	FT	0.766	0.859	0.767	0.875	0.775	0.808	0.087	0.071	0.091	0.062	0.078	0.078
<i>Vicuna1.5-13B</i>	ZS [ctrl]	0.735	0.755	0.743	0.775	0.747	0.751	0.095	0.102	0.096	0.096	0.086	0.095
	ZS	0.738	0.756	0.744	0.779	0.752	0.754	0.094	0.103	0.097	0.097	0.085	0.095
	FT [ctrl]	0.760	0.820	0.761	0.830	0.762	0.787	0.083	0.080	0.085	0.074	0.077	0.080
	FT	0.781	0.869	0.781	0.882	0.781	0.819	0.079	0.063	0.084	0.057	0.073	0.071
<i>Llama3-8B-Base</i>	ZS [ctrl]	0.748	0.766	0.757	0.779	0.768	0.764	0.097	0.114	0.093	0.109	0.088	0.100
	ZS	0.749	0.768	0.759	0.781	0.770	0.765	0.097	0.116	0.097	0.109	0.087	0.101
	FT [ctrl]	0.756	0.823	0.751	0.837	0.770	0.787	0.084	0.082	0.087	0.073	0.077	0.081
	FT	0.770	0.858	0.773	0.877	0.781	0.812	0.081	0.073	0.087	0.061	0.074	0.075
<i>Llama3-8B-Instruct</i>	ZS [ctrl]	0.574	0.626	0.563	0.644	0.587	0.599	0.130	0.147	0.135	0.145	0.136	0.139
	ZS	0.585	0.650	0.589	0.657	0.584	0.613	0.130	0.139	0.126	0.145	0.141	0.136
	FT [ctrl]	0.756	0.826	0.751	0.833	0.764	0.786	0.077	0.077	0.082	0.073	0.071	0.076
	FT	0.777	0.881	0.783	0.890	0.784	0.823	0.073	0.060	0.080	0.053	0.067	0.067
<i>Distil-Qwen-7B</i>	ZS[ctrl]	0.586	0.641	0.642	0.698	0.682	0.650	0.084	0.096	0.138	0.100	0.123	0.108
	ZS	0.583	0.645	0.639	0.701	0.671	0.648	0.084	0.095	0.138	0.100	0.126	0.109
	FT[ctrl]	0.747	0.764	0.791	0.811	0.817	0.786	0.075	0.079	0.080	0.078	0.110	0.085
	FT	0.756	0.781	0.803	0.833	0.834	0.801	0.073	0.076	0.077	0.072	0.107	0.081
<i>Distil-Qwen-14B</i>	ZS[ctrl]	0.650	0.704	0.749	0.711	0.743	0.712	0.088	0.082	0.089	0.115	0.141	0.103
	ZS	0.654	0.716	0.756	0.731	0.746	0.721	0.088	0.081	0.086	0.110	0.136	0.100
	FT[ctrl]	0.751	0.784	0.799	0.807	0.833	0.795	0.066	0.074	0.084	0.083	0.103	0.082
	FT	0.777	0.816	0.827	0.851	0.861	0.826	0.061	0.068	0.077	0.070	0.099	0.075
<i>Distil-Qwen-32B</i>	ZS[ctrl]	0.649	0.671	0.721	0.674	0.743	0.691	0.089	0.105	0.099	0.131	0.126	0.110
	ZS	0.658	0.683	0.719	0.695	0.749	0.701	0.088	0.105	0.101	0.129	0.125	0.110
	FT[ctrl]	0.776	0.789	0.799	0.794	0.821	0.796	0.066	0.071	0.074	0.087	0.111	0.082
	FT	0.800	0.821	0.815	0.830	0.846	0.822	0.062	0.067	0.073	0.078	0.110	0.078

Table 4: **Main results for predicting country-level survey response distributions on the WVS data.** We test all models with zero-shot prompting (ZS) and our proposed fine-tuning approach (FT). [ctrl] indicates a control setup, where we randomly replace countries in test prompts with other countries, to evaluate country context sensitivity. We report Jensen-Shannon Divergence (1-JSD, \uparrow) and Earth Mover Distance (EMD, \downarrow).

JSD, where JSD is Jensen-Shannon divergence, also employed by [Durmus et al. \(2023\)](#), is a symmetric measure of the similarity between two probability distributions, with higher values indicating greater similarity; ii) **Earth Mover Distance** (EMD; [Rubner et al., 1998](#)), also known as the Wasserstein distance, quantifies the minimum amount of work required to transform one distribution into another, with lower values indicating greater similarity in distribution. Both the 1-JSD and EMD metrics range from 0 to 1.

6 Results

We investigate two primary research questions:

RQ1 How effectively does the proposed alignment method improve the distribution simulation of the model on unseen countries and questions?

RQ2 What is necessary to perform well on the task—how much is dependent on modeling the prior distribution, and how much on context sensitivity?

We present comprehensive experimental results to address the two research questions.

6.1 RQ1: Generalization Performance

To address RQ1 (how effectively FT improves distribution simulation on unseen countries and questions), we train the selected models using our proposed simulation methods and evaluate their performance on unseen countries and questions. Table 4 presents the evaluation scores across models of varying sizes and types.

Zero-shot [ZS] vs. Fine-tuning [FT]. Across all model sizes and types, we observe that zero-shot prompting (ZS) consistently yields worse scores compared to fine-tuned models (FT), indicating that while ZS is capable of addressing unseen countries and questions, it lacks the adaptability needed for effective distribution simulation. In contrast, fine-tuned models show improved performance, with higher 1-JSD and lower EMD scores (e.g., a 34.3% 1-JSD increase and 0.069 EMD decrease for Llama3-8B-Instruct Avg.), demonstrating their enhanced ability to align with real-world distributions. This suggests that fine-tuning enables models to internalize more detailed patterns and relationships, making them more effective for simulation.

Unseen Countries vs. Unseen Questions. The generalization capabilities are revealed by evaluation on unseen attributes. Across all models and settings, unseen questions (Q_3) tend to present a greater challenge than unseen countries (C_2 or C_3), as indicated by slightly worse scores for unseen questions (e.g., 0.781 vs. 0.886 $1-JSD$ for Llama3-Instruct). This suggests that while the models are reasonably robust in handling new country distributions, they struggle more with generating accurate distributions for questions that were not encountered during training.

Comparison Across Models. Our method demonstrates consistent performance improvements across three representative model families and various model sizes, highlighting the effectiveness of our proposed first-token alignment approach. Specifically, while there are notable performance differences in the zero-shot setting—for example, Llama3-Base achieves a higher $1-JSD$ value (0.765) compared to Llama3-Instruct (0.613), suggesting that the base model better predicts the option token distribution—our fine-tuning procedure significantly bridges this gap. After fine-tuning, Llama3-Instruct not only closes the initial performance disparity but even surpasses Llama3-Base (34.3% $1-JSD$ increase vs. 6.1% increase). Similarly, experiments with Distilled Qwen models of different sizes reveal no clear scaling trends. Moreover, although Vicuna1.5 is generally considered weaker compared to Llama and Qwen, it surprisingly delivers similarly competitive results on this task. Overall, analysis across all models further revealing the finding that, **regardless of the starting model, our fine-tuning approach consistently produces models with similar strong performance on our task.**

Correlation Between First-Token Distribution and Response Accuracy. Beyond the alignment of first-token distributions, we further evaluate of the mode accuracy of models in responding to questionnaire items following fine-tuning. We consider the options with the highest probabilities as (single, argmax) predictions and calculate accuracy against the survey majority choice per country (Arora et al., 2023; Cao et al., 2023; AlKhamissi et al., 2024). As depicted in Figure 2, our findings reveal a significant increase in accuracy across all models and test subsets, highlighting the effectiveness of the fine-tuning process. Notably, performance improvements are particularly significant in

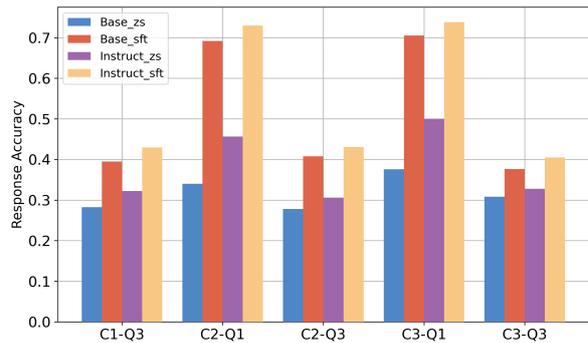


Figure 2: Option prediction accuracy for cultural questions using the Llama3-8B-Instruct. The final option is simulated by selecting the option with the highest probabilities, compared against human majority choice.

unseen countries, as demonstrated by C_2-Q_1 and C_3-Q_1 subsets. These results suggest that our proposed method not only improves the simulation of option distributions but also strengthens the models’ alignment with the correct responses, underscoring the interdependence of distribution alignment and answer accuracy.

6.2 RQ2: Variation Sensitivity

To address RQ2 (contribution of modeling the prior distribution vs. context sensitivity), we compare the control setting (with countries randomly replaced), analyze diversity changes of models, and explore shifts in response accuracy for unseen countries.

ZS[ctrl] vs. FT[ctrl]. In the control setting, ZS[ctrl] shows a smaller performance drop than ZS, while FT[ctrl] sees a larger decline, with a 16.7% avg. ($1-JSD$) drop between FT[ctrl] and FT across seven models, compared to 3.7% for ZS[ctrl] and ZS. This indicates that fine-tuned models (FT) are more sensitive to the country context and not just the prior distribution of responses (which FT[ctrl] is trained to simulate), suggesting they have become more specialized in capturing cultural nuances during training. In contrast, the smaller gap between ZS and ZS[ctrl] implies that zero-shot models maintain a more generalized understanding of cultural contexts, making them less affected by random permutations of country data. This difference highlights the fine-tuned models’ improved capability in simulating response distributions for specific countries.

Country Diversity in Model Outputs. We define diversity as the divergence across countries of response distributions given the same question.

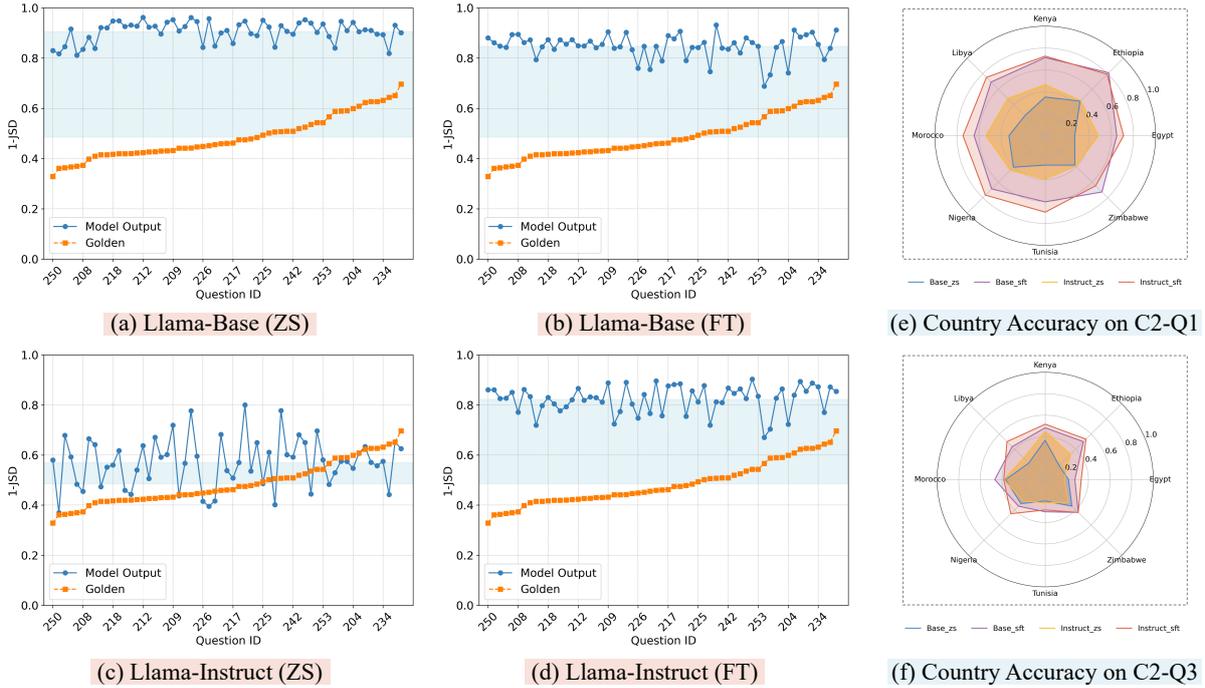


Figure 3: Model Diversity and Country Accuracy Analysis. (a)-(d) denotes the comparison of $1-JSD$ scores across countries for specific questions (C_2-Q_3 and C_3-Q_3). The blue-shaded area represents diversity changes, with the lower boundary indicating the mean of survey response scores and the upper boundary representing the mean of model outputs. (e)-(f) denotes the accuracy of options on African countries in both C_2-Q_1 and C_2-Q_3 .

Base Model	Methods	$(1-JSD) \uparrow$						$EMD \downarrow$					
		C_1-Q_3	C_2-Q_1	C_2-Q_3	C_3-Q_1	C_3-Q_3	Avg.	C_1-Q_3	C_2-Q_1	C_2-Q_3	C_3-Q_1	C_3-Q_3	Avg.
Llama3-8B-Instruct	ZS	0.603	0.654	0.601	0.659	0.613	0.626	0.125	0.140	0.131	0.145	0.132	0.135
	FT	0.777	0.852	0.789	0.870	0.791	0.816	0.081	0.078	0.087	0.065	0.073	0.077
Distil-Qwen-7B	ZS	0.605	0.706	0.712	0.696	0.693	0.682	0.084	0.091	0.086	0.087	0.190	0.108
	FT	0.764	0.779	0.771	0.791	0.851	0.791	0.067	0.085	0.081	0.082	0.130	0.089

Table 5: **Model Performance on Chinese World Values Survey.** Performance comparison between zero-shot prompting (ZS) and supervised fine-tuning (SFT) on two model families.

To assess whether FT can enhance model output diversity, we calculate $1-JSD$ for Llama3-Base and Llama3-Instruct (ZS and FT) output across countries for each question. A lower $1-JSD$ score suggests greater diversity in responses between countries, whereas a higher score indicates greater similarity in distributions. Results are shown in Figure 3, where the scores of survey responses are compared with those of the model predictions.

This visualization provides several interesting insights. Firstly, the outputs of Base models exhibit a high degree of uniformity across countries, indicating a **limited sensitivity to national variations** when addressing cultural values questions. After fine-tuning, there is a slight reduction in $1-JSD$, suggesting an enhancement in the responsiveness to diverse cultural contexts. Secondly, although the Instruct model, having undergone alignment fine-

tuning, initially produces a more varied distribution of answers, **this diversity diminishes following distribution simulation fine-tuning**. Thirdly, we observe **no consistent correlation between the diversity of generated responses and the accuracy of simulated distributions**. Lastly, the post-fine-tuning diversity of responses from both models converges, indicating that our fine-tuning approach **improves the sensitivity to national differences**.

Unseen Country Shifts. We visualize the option accuracy of African countries as observation objects in Figure 3e-3f. For seen questions, we find that all models show a relatively high performance across most countries, with Ethiopia and Nigeria displaying the highest accuracies close to 80%. For unseen questions, the performance decreases greatly compared to seen questions, particularly

Methods	C'_1		C_3	
	(1-JSD) \uparrow	ACC \uparrow	(1-JSD) \uparrow	ACC \uparrow
<i>Vicuna1.5</i> (ZS)	0.690	0.360	0.668	0.346
<i>Vicuna1.5</i> (FT)	0.725	0.442	0.709	0.456
<i>Llama3</i> (ZS)	0.617	0.472	0.613	0.446
<i>Llama3</i> (FT)	0.767	0.562	0.755	0.568

Table 6: Evaluation results on GlobalOpinionQA Pew dataset for *Llama3-Instruct* and *Vicuna1.5-7B*. (1-JSD) and option accuracy scores are reported.

in countries like Egypt and Tunisia. Besides, the Instruct models maintain a higher level of accuracy relative to the Base models in most cases, showing the relative robustness of Instruct models in both seen and unseen scenarios across the African countries, which is consistent with results in Figure 2.

6.3 Robustness Analysis

In this section, we explore the robustness of the models on survey language and unseen survey.

Impact of Survey Language on Results. We fine-tuned the *Llama3-8B-Instruct* and *Distil-Qwen-7B* models on the official Chinese translation dataset, and the results are presented in Table 5. While the performance of both models in Chinese is marginally lower than in English both on (1-JSD) and EMD metrics, the difference is not significant, suggesting that current LLMs exhibit limited sensitivity to language differences in this task. Besides, while the Distilled-Qwen model demonstrates a better performance than *Llama3* on most benchmarks, it does not outperform *Llama3* in this task.

Generalization to a New Survey. We use Pew introduced in §3.4 to test the generalization of our fine-tuned models. Table 6 presents the results in both 1-JSD scores and accuracy. Notably, all metric scores show significant improvements for both models after fine-tuning. Additionally, across both C'_1 and C_3 , *Llama3* outperforms *Vicuna1.5* in terms of accuracy, particularly in the fine-tuned setting, where *Llama3* (FT) achieves 19.1% and 27.4% improvement for two datasets, respectively. The consistent improvements prove its capability to generalize well to unseen surveys.

6.4 Ablation Studies

We conduct ablation studies to analyze the impact of the training loss function and option ordering.

Loss Function. As shown in Table 7, KL Loss is the most effective loss function for our task. How-

Methods	C_1-Q_3	C_2-Q_1	C_2-Q_3	C_3-Q_1	C_3-Q_3
KL (Orig)	0.777	0.881	0.783	0.890	0.784
WA Loss	0.733	0.774	0.744	0.782	0.749
JS Loss	0.745	0.790	0.756	0.799	0.763
CE Loss	0.746	0.809	0.772	0.807	0.756
Shuffled	0.753	0.820	0.761	0.815	0.761

Table 7: **Results of our ablation studies.** We compare different loss functions (WA, JS and CE) to our KL loss setup. We also evaluate on a test set with shuffled option orders. All results are for *Llama3-8B-Instruct*.

ever, Wasserstein (WA) Loss Jensen-Shannon (JS), and Cross-Entropy (CE) Loss also improve over zero-shot prompting.

Option Ordering. Dominguez-Olmedo et al. (2023) observed an A-bias effect, where models tend to disproportionately select the answer choice labeled ‘‘A’’. To assess this bias, we re-evaluate our fine-tuned model on the same dataset where the answer options are shuffled. As shown in Table 7 (‘‘Shuffled’’), there is a reduction in performance, but it is small compared to the effect of fine-tuning or model choice. This indicates that the option ordering is not a major concern in our experiments.

7 Conclusion

In this paper, we explored the task of specialising LLMs to simulate survey response distributions across diverse countries and questions. For this task, we devised a fine-tuning method based on first-token probabilities. Our experiments demonstrate that fine-tuning models substantially improves response simulation prediction compared to zero-shot models, for both seen and unseen countries and questions. Further, fine-tuned models also show improved generalization to an entirely new survey dataset. Despite these improvements, our results also highlight systematic limitations of the models, particularly when simulating responses to unseen questions. We also observed that the models, whether fine-tuned or not, were less diverse in their predictions compared to the human survey response data, raising questions about their utility.

While our results provide clear evidence for the benefits of specializing LLMs for survey simulation tasks, they also underscore the need for caution when using LLMs for this task, as even the best-performing models exhibited systematic inaccuracies, especially in culturally diverse contexts.

Limitations

While we proved the effectiveness of our proposed method, several limitations remain in our work.

Scope. Our trained models are highly specialized and can only be used for the specific task of predicting the distribution of answers to a given survey question from a specified human population. Future work will investigate whether the fine-tuning approach also results in less biased or more aligned models in general-purpose applications, but this cannot be claimed only based on our study.

Language and Countries Coverage. Our study only uses English prompts for experiments and uses countries to represent specific cultures, consistent with existing studies (Cao et al., 2023; AIKhamissi et al., 2024). While this approach offers some valuable insights, it may limit the applicability of our findings to non-English LLMs and diverse fine-grained cultural contexts. We hope that future research could benefit from exploring broader languages and countries to enhance the robustness of the proposed framework.

Model Choice. Due to computational resource consideration, we did not fine-tune LLMs with more than 32B parameters and instead selected a limited number of models for validation. Despite this limitation, in future work, we aim to cover a range of powerful models of varying sizes, which will allow us to uncover interesting observations regarding their performance. We believe that the insights we draw will still contribute to future research, encouraging further exploration of larger models to better understand their capabilities in simulating cultural diversity.

Ethics Statement

This research adheres to strict ethical standards, ensuring that all datasets, large language models, and prompt settings used are sourced from open-access repositories and are properly licensed to their original creators.

While our proposed framework does not involve any inherently risky operations, we acknowledge that the deployment of LLMs carries inevitable potential ethical implications. Therefore, users interacting with our models are strongly encouraged to consider safety and ethical factors, remaining aware of the potential risks and harms that may arise from misuse or misinterpretation of the gen-

erated content. Through this work, we aim to contribute positively to a better understanding of cultural diversities and promote responsible practices in the simulation of cultural contexts.

Acknowledgments

We thank anonymous reviewers for their valuable comments. This research was co-funded by a DFF Sapere Aude research leader grant under grant agreement No 0171-00034B, and supported by the Pioneer Centre for AI, DNRF grant number P1. Yong Cao was supported by a VolkswagenStiftung Momentum grant.

References

- Gati V Aher, Rosa I Arriaga, and Adam Tauman Kalai. 2023. Using large language models to simulate multiple humans and replicate human subject studies. In *International Conference on Machine Learning*, pages 337–371. PMLR.
- AI@Meta. 2024. [Llama 3 model card](#).
- Badr AIKhamissi, Muhammad ElNokrashy, Mai Alkhamissi, and Mona Diab. 2024. [Investigating cultural alignment of large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12404–12422. Association for Computational Linguistics.
- Lisa P Argyle, Ethan C Busby, Nancy Fulda, Joshua R Gubler, Christopher Rytting, and David Wingate. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31(3):337–351.
- Arnav Arora, Lucie-aimée Kaffee, and Isabelle Augenstein. 2023. [Probing pre-trained language models for cross-cultural differences in values](#). In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 114–130. Association for Computational Linguistics.
- Joris Baan, Wilker Aziz, Barbara Plank, and Raquel Fernández. 2022. [Stop measuring calibration when humans disagree](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1892–1915. Association for Computational Linguistics.
- Joris Baan, Raquel Fernández, Barbara Plank, and Wilker Aziz. 2024. [Interpreting predictive probabilities: Model confidence or human label variation?](#) In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 268–277. Association for Computational Linguistics.

- Christopher A Bail. 2024. Can generative ai improve social science? *Proceedings of the National Academy of Sciences*, 121(21):e2314021121.
- Noam Benkler, Drisana Mosaphir, Scott Friedman, Andrew Smart, and Sonja Schmer-Galunder. 2023. Assessing LLMs for moral value pluralism. *arXiv preprint arXiv:2312.10075*.
- James Bisbee, Joshua D Clinton, Cassy Dorff, Brenton Kenkel, and Jennifer M Larson. 2023. Synthetic replacements for human survey data? the perils of large language models. *Political Analysis*, pages 1–16.
- Yong Cao, Li Zhou, Seolhwa Lee, Laura Cabello, Min Chen, and Daniel Hershcovich. 2023. [Assessing cross-cultural alignment between ChatGPT and human societies: An empirical study](#). In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 53–67. Association for Computational Linguistics.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing GPT-4 with 90%* ChatGPT quality](#).
- Ricardo Dominguez-Olmedo, Moritz Hardt, and Celestine Mendler-Dünnler. 2023. Questioning the survey responses of large language models. *arXiv preprint arXiv:2306.07951*.
- Esin Durmus, Karina Nyugen, Thomas I Liao, Nicholas Schiefer, Amanda Askill, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, et al. 2023. Towards measuring the representation of subjective global opinions in language models. *arXiv preprint arXiv:2306.16388*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- C. Haerpfel, R. Inglehart, A. Moreno, C. Welzel, K. Kizilova, J. Diez-Medrano, M. Lagos, P. Norris, E. Ponarin, and B. Puranen. 2022. [World values survey wave 7 \(2017-2022\) cross-national data-set](#).
- Guande He, Peng Cui, Jianfei Chen, Wenbo Hu, and Jun Zhu. 2024. [Investigating uncertainty calibration of aligned language models under the multiple-choice setting](#).
- Luke Hewitt, Ashwini Ashokkumar, Isaias Ghezae, and Robb Willer. 2024. [Predicting results of social science experiments using large language models](#).
- John J Horton. 2023. Large language models as simulated economic agents: What can we learn from homo silicus? Technical report, National Bureau of Economic Research.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Austin Kozlowski and James Evans. 2024. Simulating subjects: The promise and peril of ai stand-ins for social agents and interactions.
- Louis Kwok, Michal Bravansky, and Lewis D Griffin. 2024. Evaluating cultural adaptability of a large language model via simulation of synthetic personas. *arXiv preprint arXiv:2408.06929*.
- Wangyue Li, Liangzhi Li, Tong Xiang, Xiao Liu, Wei Deng, and Noa Garcia. 2024. [Can multiple-choice questions really be useful in detecting the abilities of LLMs?](#) In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2819–2834. ELRA and ICCL.
- Benjamin S Manning, Kehang Zhu, and John J Horton. 2024. Automated social science: Language models as scientist and subjects. Technical report, National Bureau of Economic Research.
- Peter S Park, Philipp Schoenegger, and Chongyang Zhu. 2024. Diminished diversity-of-thought in a standard large language model. *Behavior Research Methods*, pages 1–17.
- Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. 1998. A metric for distributions with applications to image databases. In *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*, pages 59–66. IEEE.
- Seungjong Sun, Eungu Lee, Dongyan Nan, Xiangying Zhao, Wonbyung Lee, Bernard J. Jansen, and Jang Hyun Kim. 2024. [Random silicon sampling: Simulating human sub-population opinion using a large language model based on group-level demographic information](#). *Preprint*, arXiv:2402.18144.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Dustin Wright, Arnav Arora, Nadav Borenstein, Srishti Yadav, Serge Belongie, and Isabelle Augenstein. 2024. Revealing fine-grained values and opinions in large language models.
- Yahan Yang, Soham Dan, Dan Roth, and Insup Lee. 2024. [On the calibration of multilingual question answering llms](#). *Preprint*, arXiv:2311.08669.
- Wenlong Zhao, Debanjan Mondal, Niket Tandon, Danica Dillion, Kurt Gray, and Yuling Gu. 2024. World-valuesbench: A large-scale benchmark dataset for

multi-cultural value awareness of language models. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 17696–17706.

A Pew Data Distribution

The detailed statistics for the C2 and C3 split of the cultural survey simulation and the C3 split of GlobalOpinionQA data mentioned are available in Table 8. The sampled countries, along with the sampling reason of Pew- C'_1 , are in Table 9.

C2		C3		
Country	CSS-N	Country	CSS-N	PWE-N
Egypt	157	Malaysia	150	516
Ethiopia	185	Thailand	150	319
Kenya	185	Czechia	150	212
Libya	185	Greece	150	648
Morocco	184	Nigeria	210	1044
Nigeria	185	Morocco	209	450
Tunisia	184	Peru	146	545
Zimbabwe	185	Colombia	150	369
		Mexico	150	890
		Puerto Rico	150	221
		New Zealand	149	274

Table 8: Number of Entries Used in C2 and C3 Split of Cultural Survey Simulation and C3 Split of GlobalOpinionQA Data.

Country	Continent	GDP-level	N
Germany	Europe	High	1130
Japan	Asia	High	891
Brazil	South America	Upper-middle	922
Australia	Oceania	High	627
India	Asia	Lower-middle	932
Nigeria	Africa	Lower-middle	1044
United States	North America	High	1104
Vietnam	Asia	Lower-middle	471
Chile	South America	Upper-middle	542
Ukraine	Europe	Lower-middle	661

Table 9: Sampled countries for Pew- C'_1 regarding geographical and GDP-level diversity. The dataset consists of 8,324 entries in total.

B WVS data Distribution

Figure 4 illustrates the sample distribution across various countries in World Values Survey. For the purpose of balanced representation, in our formatted dataset, we only include countries with more than 1,000 samples, ensuring consistency in data distribution across countries and minimizing sample size discrepancies in the analysis.

Set	Cultural Dimension	Number
Q_1	Social values, attitudes and stereotypes	45
	Societal well-being	12
	Social capital, trust and organizational membership	49
	Economic values	6
	Corruption	9
	Migration	10
	Post-materialist index	21
	Science and technology	6
	Religious values	6
	Q_2	Security
Ethical values and norms		23
Q_3	Political interest and political participation	36
	Political culture and political regimes	25
C_1	Andorra, Argentina, Australia, Bangladesh, Armenia, Bolivia, Brazil, Myanmar, Canada, Chile, China, Taiwan ROC, Cyprus, Ecuador, Germany, Guatemala, Hong Kong SAR, Indonesia, Iran, Iraq, Japan, Kazakhstan, Jordan, South Korea, Kyrgyzstan, Lebanon, Macao SAR, Maldives, Mongolia, Netherlands, Nicaragua, Pakistan, Philippines, Romania, Russia, Serbia, Singapore, Slovakia, Vietnam, Tajikistan, Turkey, Ukraine, Great Britain, United States, Uruguay, Venezuela	46
C_2	Egypt, Ethiopia, Kenya, Libya, Morocco, Nigeria, Tunisia, Zimbabwe	8
C_3	Malaysia, Thailand, Czechia, Greece, Nigeria, Morocco, Peru, Colombia, Mexico, Puerto Rico, New Zealand	11

Table 10: Cultural dimensions and question ids of WVS. Question 82-223 and 94-106 are excluded in Q_1 as they are demo graphical questions. Demo graphical questions are excluded as they are related to individual attribute regarding participants and does not have relevance to group culture.

Regarding dataset split, we split WVS by cultural dimensions and countries vis multiple sets, which is visualized in Figure 5 and the corresponding cultural dimensions are detailed in Table 10, offering a comprehensive view of the overall distribution in the cultural survey simulation.

C Hyperparameter Settings

Training is performed on a single A100 GPU with a batch size of 16 for Llama3 and Vicuna1.5-7B, and 4 for Vicuna1.5-13B. Both models use the AdamW optimizer with a learning rate of $1e-4$ and implement Fully Sharded Data Parallel along with a mixed precision strategy to enhance computational efficiency. In our experiments, we employ LoRA with a rank of 8, a scaling factor `lora_alpha` set to 32, and a dropout rate of 0.05.

Model Download We list all used models here:

- **Vicuna1.5-7B:** <https://huggingface.co/lmsys/vicuna-7b-v1.5>
- **Vicuna1.5-13B:** <https://huggingface.co/lmsys/vicuna-13b-v1.5>

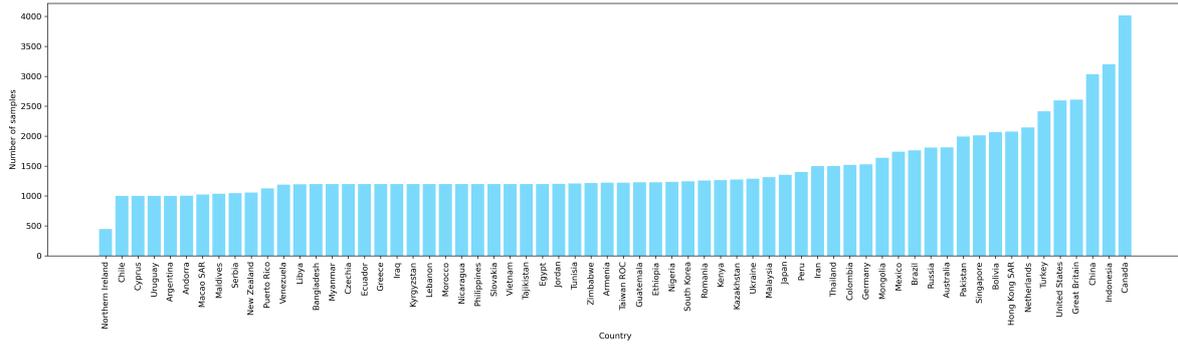


Figure 4: Sample distribution of the World Values Survey across countries, including only countries with more than 1,000 samples for balanced representation.

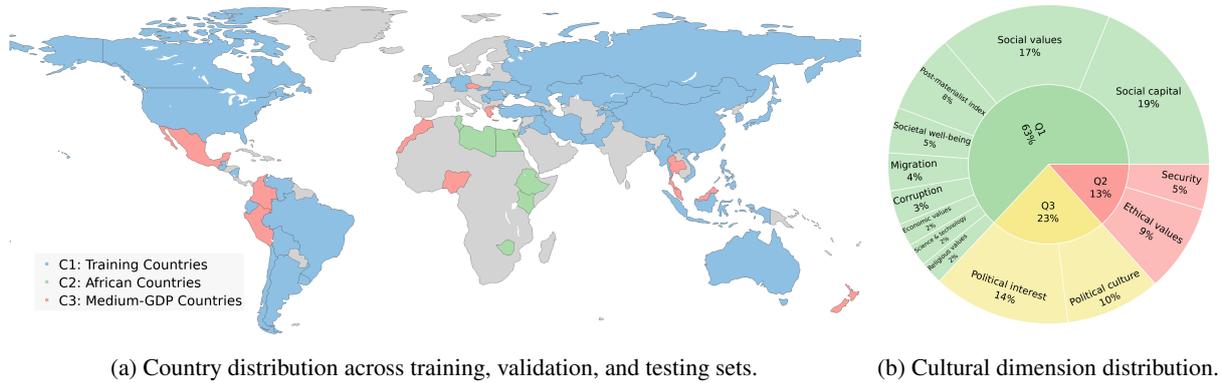


Figure 5: Visualization of country and cultural dimension divisions of WVS. Countries are categorized into three groups, and questions are divided based on selected cultural dimensions.

- **Distil-Qwen-7B:** <https://huggingface.co/deepseek-ai/DeepSeek-R1-Distill-Qwen-7B>
- **Distil-Qwen-14B:** <https://huggingface.co/deepseek-ai/DeepSeek-R1-Distill-Qwen-14B>
- **Distil-Qwen-32B:** <https://huggingface.co/deepseek-ai/DeepSeek-R1-Distill-Qwen-32B>
- **Llama3-8B-Base:** <https://huggingface.co/meta-llama/Meta-Llama-3-8B>
- **Llama3-8B-Instruct:** <https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>
- **Qwen-7B:** <https://huggingface.co/Qwen/Qwen-7B>

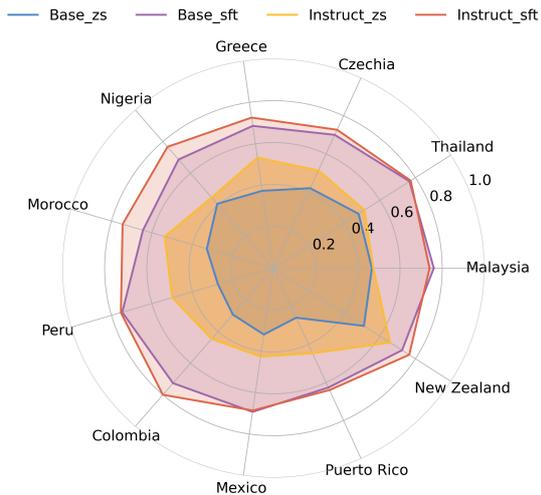
D More Option Prediction Analysis

Option Prediction Accuracy. In line with Figure 3e-3f, we present the remaining two sets of

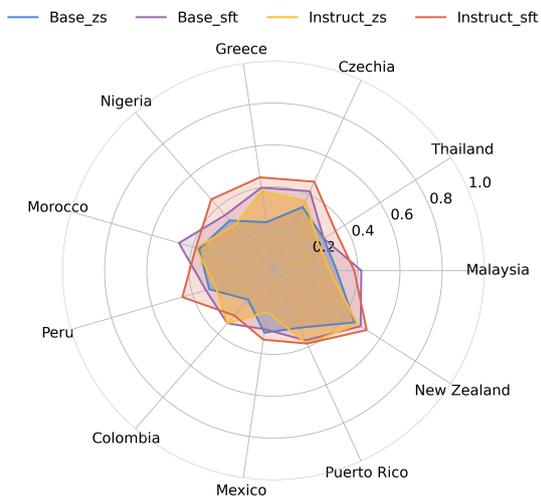
Llama3 option prediction accuracy for C3-Q1 and C3-Q3 here (see Figure 6).

The results indicate that the model performs significantly better on seen questions (C3-Q1) compared to unseen ones (C3-Q3), with Llama3 Instruct models consistently outperforming Base models, and fine-tuned approaches demonstrating superior accuracy over zero-shot methods across most countries. Specifically, the model achieves its highest accuracy in Greece, Nigeria, and Peru, while showing weaker performance in Thailand and Mexico. Despite the challenges in predicting unseen questions, fine-tuning through simulation leads to noticeable improvements, highlighting the potential of fine-tuned models to enhance generalization in novel scenarios.

Country Distributions. Figure 8 visualizes the distribution of 1-JSD scores for the Llama3-Base and Llama3-Instruct models across unseen cultural questions (C_1-Q_3) in both zero-shot (ZS) and fine-tuned (FT) modes. In ZS mode, both models perform poorly across many regions, with lower scores indicating limited cultural sensitivity. This sug-



(a) Country Accuracy on C3-Q1.



(b) Country Accuracy on C3-Q3.

Figure 6: Llama3 Accuracy of options on African countries in both seen C3-Q1 and unseen C3-Q3 questions.

gests that without additional tuning, both models struggle to effectively generalize to simulate culturally diverse distributions. Furthermore, the Instruct model performs slightly worse, likely due to the reinforcement learning and safety strategies, which may inadvertently reduce its sensitivity to cultures in this scenario.

However, after FT, both models show improvement across all countries, particularly significant in regions previously displaying poor performance in the ZS setting. Additionally, the performance of both the fine-tuned Base model and the Instruct model is observed to be very similar across different countries, demonstrating that our methods can

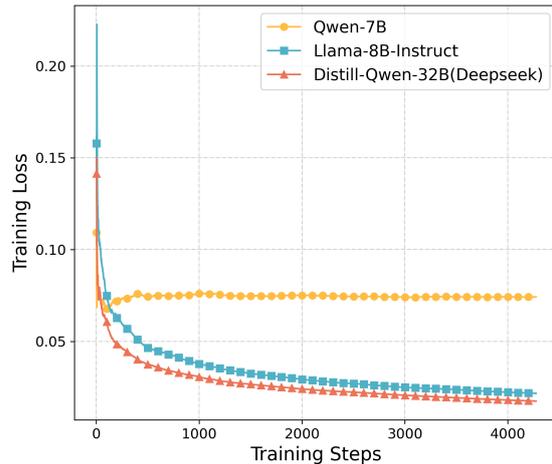


Figure 7: Training loss comparison of three models: Qwen-7B, Llama-8B-Instruct, Distill-Qwen-32B(Deepseek)

effectively align the Base and Instruct models.

E More Baseline Comparison

To further assess the effectiveness of our proposed approach, we compare it against several other baseline methods as follows:

- **KNN**: This method identifies the most similar training question and country (top-1, $k = 1$) using BERT embeddings and returns the corresponding option distributions as predictions.
- **Avg_Culture**: We computes the mean option distribution across all training countries for each known question and adopts a uniform random distribution for unknown questions.
- **JSON-ZS**: This approach prompts the models to directly generate option distributions in JSON format without additional fine-tuning.

Table 11 presents the 1-JSD scores for various baseline methods and our proposed approach. The results indicate that all baseline methods perform substantially worse than our fine-tuned (FT) models. Among the baselines, the JSON-ZS method demonstrates superior performance compared to KNN and Avg_Culture; however, it remains less effective than both zero-shot (ZS) and fine-tuned (FT) approaches. Notably, fine-tuning consistently yields the highest scores across all evaluated settings, underscoring its effectiveness in improving model performance.

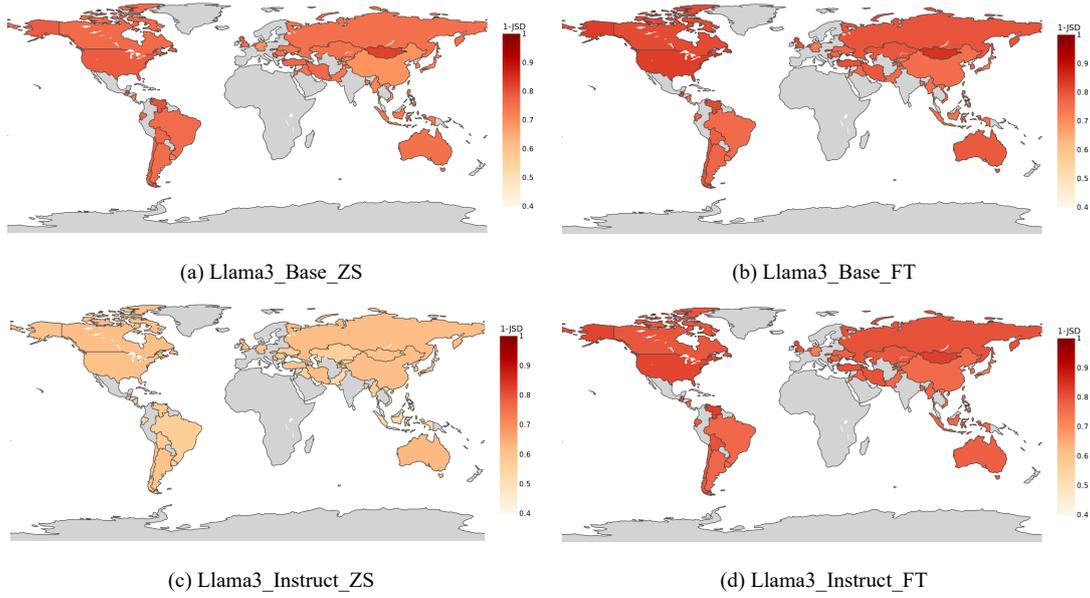


Figure 8: Distribution of $1-JSD$ Global Scores for the Model on Unseen Cultural Questions (C_1-Q_3). The Instruct model exhibits a more distinct improvement compared to the Base model.

Methods	C1-Q3	C2-Q1	C2-Q3	C3-Q1	C3-Q3	Avg.
KNN	0.381	0.518	0.371	0.541	0.384	0.439
Avg_Culture	0.360	0.509	0.348	0.518	0.368	0.421
<i>Llama3-8B-Base</i>						
JSON-ZS	0.754	0.728	0.581	0.729	0.751	0.709
ZS	0.749	0.768	0.759	0.781	0.770	0.765
FT	0.770	0.858	0.773	0.877	0.781	0.812
<i>Llama3-8B-Instruct</i>						
JSON-ZS	0.735	0.728	0.747	0.729	0.751	0.738
ZS	0.585	0.650	0.589	0.657	0.584	0.613
FT	0.777	0.881	0.783	0.890	0.784	0.823

Table 11: Comparison of $1-JSD$ scores across different baseline methods and our approach. Higher values indicate better alignment with the ground-truth distribution.

F More Model Evaluation

In this section, we present additional experiment results on original Qwen model. We observe that the original Qwen models present challenges in training for our task. As shown in Figure 7, we compare the training progress of three models. The loss of the Qwen-7B model exhibits difficulty in converging compared to Llama-8B-Instruct and Distill-Qwen-32B. This difficulty may stem from the extensive safety alignment or policy alignment incorporated into Qwen, which could introduce additional constraints or optimization challenges during training.

Secondly, we visualized the $1-JSD$ score distribution of the model for English ZS, FT, and Chinese FT, as shown in Figure 9, with the red dashed line indicating the average value. The results show that

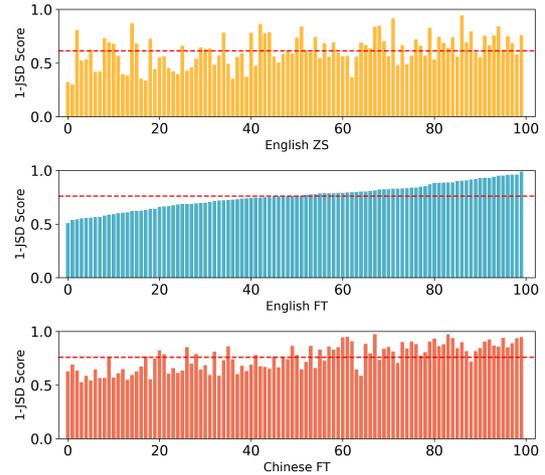


Figure 9: Comparison of different languages in $1-JSD$ score on random selected 100 samples.

language differences have a smaller impact on the model than training levels. Additionally, we calculate the Pearson correlation coefficient between English ZS and FT which (0.349) is lower than that between English FT and Chinese FT (0.579). While both correlations are positive, the stronger correlation observed between English FT and Chinese FT suggests that training level exerts a greater influence than language differences.