

Identifying Distributional Perspective Differences from Colingual Groups

Yufei Tian¹, Tuhin Chakrabarty², Fred Morstatter³ and Nanyun Peng^{3,4}

¹ Department of Computer Science, University of Southern California

² Department of Computer Science, Columbia University

³ Information Sciences Institute, University of Southern California

⁴ Department of Computer Science, University of California Los Angeles

yufeit@usc.edu, tuhin.chakr@cs.columbia.edu,

fredsmors@isi.edu, violetpeng@cs.ucla.edu

Abstract

Perspective differences exist among different cultures or languages. A lack of mutual understanding among different groups about their perspectives on specific values or events may lead to uninformed decisions or biased opinions. Automatically understanding the group perspectives can provide essential background for many downstream applications of natural language processing techniques. In this paper, we study colingual groups¹ and use language corpora as a proxy to identify their *distributional perspectives*. We present a novel computational approach to learn shared understandings, and benchmark our method by building culturally-aware models for the English, Chinese, and Japanese languages. On a held out set of diverse topics including *marriage*, *corruption*, *democracy*, our model achieves high correlation with human judgements regarding intra-group values and inter-group differences.

1 Introduction

Sociologists have defined culture as a set of shared understandings, herein called *perspectives*, adopted by the members of that culture (Bar-Tal, 2000; Sperber and Hirschfeld, 2004). Languages and cultures have radical correlations (Khaslavsky, 1998; Bracewell and Tomlinson, 2012; Gelman and Roberts, 2017), because individuals communicate with each other by language, which carries the aspects of their cultures, experiences, beliefs, and values, thus will shape their perspectives. Lacking of understanding for these perspective differences could lead to biased predictions. Selection bias (Heckman, 1977) can often lead to misinformation as it sometimes ignores facts that do not reflect the entire population intended to be analyzed. For example, to verify a controversial statement like “*The*

¹A group of people that share the same language (<https://www.merriam-webster.com/dictionary/colingual>).

Claim	The free market does a much worse job than the government in providing essential services and the fraud and corruption part only gets worse.
CN Persp	Human: 72% support, Model: 79% support
JP Persp	Human: 17% support, Model: 15% support

(a) A claim about free market and government intervention from our test data, with the distributional perspectives of the Chinese (CN) and Japanese (JP) colingual groups. Human opinions and model predictions are highly correlated.

CN Wikipedia	中国特色的社会主义现阶段有如下特点：以国家的手段控制国内的要害经济部门和大量的企业，通过“国有资产”的概念以股份或者非股份形式保护国民经济的相当重要的部分。 The current stage of socialism with Chinese characteristics has the following characteristics: the government <i>control the vital economic sectors</i> and a large number of enterprises in the country <i>by state means</i> , and protect a very important part of the national economy in the form of shares or non-shares through the concept of “ <i>state-owned assets</i> ”. [Translated]
JP Wikipedia	1930年以降、社会的市にして人の自由や市原理を再し、政府による人や市への介入は最低限とすべきと提唱する。... 日本では1950年の事再成以来、民の力会社が地域ごとに1社ずつ合10社あり。 Since 1930, Japan reassessed the liberty and market principles of the individual for the social market economy, <i>advocating that government intervention</i> in the individual and <i>the market should be minimized</i> In Japan, since the restructuring of the electric power business in 1950, <i>there are 10 private electric power companies, one in each region</i> . [Translated]

(b) Evidence from Wikipedia pages from the colingual groups (CN and JP), that potentially are for or against the claim shown in Table 1a. These are included in our training data after variation (discussed in Section 4.2). The two examples in the JP corpus are selected out from different articles.

Table 1: An example claim from our test data (1a), and possible evidences from wikipedia pages included in our colingual group training corpora (1b).

free market causes fraud and corruption.”, we need to consider the perspectives from various groups (shown in Table 1). Similarly, a sentiment analysis model may fail to capture the correct emotions towards a debatable claim if the claim is viewed differently across different groups, such as the dispute

between India and Pakistan regarding Kashmir.

In this paper, we focus on *distributional differences* on controversial topics across groups. For example, within the United States, people have split views (approximately half-half) regarding gun control and abortion, while in China, people generally against the possession of guns and pro-choice for abortion. Hence, building a culture-aware model that considers groups’ *distributional* perspectives will help improve comprehension and consequently mitigate biases in decision making.

We aim to identify colingual groups’ distributional perspectives towards a given claim, and spot claims that provoke such divergence. As colingual groups are naturally identifiable by the usage of language, we can obviate group detection and associated errors in the process of group identification.² Wikipedia, despite its overall goal of objectivity, has been shown to embed latent cultural biases (Callahan and Herring, 2011). Following these cues, we believe Wikipedia is an ideal source to study diverse perspectives among various colingual groups. Table 1a shows an example claim for which the Chinese and Japanese may have different opinions. Specifically, the Chinese-speaking group tends to support the claim (72% support) while the Japanese-speaking group tends to oppose it (17% support), which is likely due to the different economic/government environments. As shown in Table 1b, we can find evidences from wikipedia pages that support or oppose the claim in Table 1a.

We learn a perspective model for each colingual groups using a collection of Wikipedia pages for English, Chinese and Japanese, and then use these models to identify diverging perspectives for a separate set of claims that are manually curated and are *not* from Wikipedia.

Our contributions are as follows. **1)** We propose **CLUSTER (CoLingUal PerSpecTive IdentifiER)**, a module that learns distributional perspectives of colingual groups based on Wikipedia articles. Towards this, we develop a novel procedure to algorithmically generate negative examples (introduced in Section 3.1) based on Wikipedia to train our group models (Section 4.1). **2)** We design an evaluation framework to systematically study the effectiveness of the proposed approach by testing our

²The only caveat is that such simplification ignores finer-grained cultural distinctions across subgroups speaking the same language, especially for English as a global language spoken by many nations; we leave those studies of more fine-grained groups for future work.

models on self-labeled claims from diverse topics including *cuisine, festivals, marriage, corruption, democracy, privacy*, etc. (Section 3.2, 3.3 and 4.3) **3)** Comprehensive quantitative and qualitative studies in Chinese, Japanese, and English show that our model outperforms multiple well-crafted baselines and achieves strong correlation with human judgements.³ (Section 6 and 7)

2 Task Definition

In this paper, we focus on predicting a group’s distributional perspective towards a *claim* and identifying claims that reflect *contrasting perspectives* from different groups on a particular topic. We further focus on English, Chinese and Japanese as the targeted colingual groups. Here, we define several key concepts and the task. We also explain why our task is different from stance detection.

Claim. A claim s_i , is a sentence that expresses opinions toward a certain topic (E.g Row 1 , Table 1) regardless of its language. We then translate and have a set of multi-lingual claims $\mathcal{S} = (S^{en}, S^{cn}, S^{jp})$, where S^{en} (English), S^{cn} (Chinese), S^{jp} (Japanese) are translations of each other.

Group Perspective Model and Score. Group Perspective Model is a probabilistic model that mirrors the group’s distributional perspective on a claim - the model gives a score that reflects a group’s likelihood of agreeing with that claim. For any claim s and its translations (s^{en}, s^{cn}, s^{jp}) , a machine-generated score $P^l(s^l) \in [0, 1]$ is assigned to estimate the probability of s^l (l denoting language) being supported by the corresponding group. A distributional perspective score closer to 1 (fully support) and 0 (fully reject) indicates unanimity, while a score closer to 0.5 implies split within group. Similarly, a human-annotated perspective score $H^l(s^l) \in [0, 1]$ is assigned and considered as the ground truth of the likelihood that s^l is supported by its corresponding group.

Distributional Perspective Difference. Finally, we define (distributional) perspective difference. Let $\mathcal{D}_{model_i}^{l_1-l_2} \in [-1, 1]$ be the difference of perspective scores predicted by two models (for group l_1 and l_2) of s , where

$$\mathcal{D}_{model}^{l_1-l_2} = P^{l_1}(s^{l_1}) - P^{l_2}(s^{l_2}), l_1 \neq l_2. \quad (1)$$

³We use these three languages as examples throughout the paper, but our algorithm is naturally applicable to other languages. Data and code are available at <https://github.com/PlusLabNLP/CLUSTER>

Here l_1 and l_2 each denotes a language such as ‘cn’ and ‘jp’. A positive $\mathcal{D}_{model}^{cn-jp}$ indicates that the Chinese model agrees more with the claim s than the Japanese model. Similarly, we denote $\mathcal{D}_{human}^{l_1-l_2} \in [-1, 1]$ as the quantity of perspective difference reported by human annotators:

$$\mathcal{D}_{human}^{l_1-l_2} = H^{l_1}(s^{l_1}) - H^{l_2}(s^{l_2}), l_1 \neq l_2. \quad (2)$$

In Table 1, $\mathcal{D}_{model}^{cn-jp} = 0.79 - 0.15 = 0.64$, and $\mathcal{D}_{human}^{cn-jp} = 0.55$. A higher absolute value of \mathcal{D} indicates bigger distributional differences.

Comparison With Stance Detection. Stance detection aims at detecting if a piece of text (usually a sentence or a document) supports or opposes a given claim (Hasan and Ng, 2014). Unlike stance detection, we do not have a given text associated with our claims. Instead, we learn representations of *group* perspectives through training on language corpora so that we can identify if a claim is likely to be supported or opposed by *a group*.

3 Data Preparation

In this section, we describe the procedure of composing our training data from multi-lingual Wikipedia articles. We then introduce an out-of-domain test dataset retrieved from Reddit that contain opinions regarding wide range of topics and the procedure of collecting human annotations on the test set.

3.1 Training Data

Topic Selection. We leverage the category hierarchy provided by Wikipedia to retrieve a list of child topics that belong to a few parent categories, including *politics, foods, sport, history, social issues*, etc. The selected root categories in English, Chinese and Japanese are aligned entities obtained from Wikipedia language links, and their sub-tree structures are only partially aligned. In this way, sub-topics obtained in the three languages have considerable overlap but are not identical. Hence we have different numbers of subtopics and training samples as seen in Table 2. We then retrieve all the articles under the selected subtopics separately⁴, so that different claims that potentially reflect the cultural bias are included in our training data.

⁴<https://en.wikipedia.org/wiki/Special:CategoryTree>, <https://zh.wikipedia.org/w/title=Special:分类树>, <https://ja.wikipedia.org/wiki/特別:カテゴリツリ>

	Topics	Positive Samples	Negative Samples
English	4,245	292,444	292,444
Chinese	1,563	57,904	57,904
Japanese	1,266	25,039	25,039

Table 2: Statistics of Our Training Dataset. We deliberately balance the number of positive and negative samples so that no priori probability will intervene with the learning step.

Training Dataset Creation. Upon observing many examples similar to the economics pages in Table 1, we form our fundamental assumption that the *collection* of sentences extracted from Wikipedia in a certain language represent the corresponding *distributional* perspective of that colingual group. Therefore, we label each sentence extracted from the Wikipedia articles as *positive examples*, as illustrated in part A of Figure 1.

Although positive examples mirror their corresponding perspective, we also need to compose *negative samples* — the claims that the corresponding colingual groups will disagree with. An intuitive approach is to flip the semantic meaning of the positive examples. This could be achieved by replacing the adjectives in a sentence with their antonyms. As shown in Figure 1.A, there are four adjectives in the original text: ‘*Making safe abortion legal and accessible reduces maternal deaths*’. We can obtain four negative examples by replacing each of the adjectives with its antonym (note that we do not flip multiple adjectives simultaneously). Each of the fabricated negative samples (in Figure 1.B) is ideal because it expresses conflicting viewpoints compared to the original text.

However, certain collocations such as *New York* and *legal systems* are also converted. Since bigrams such as *Old York* and *illegal systems* seldom appear in real sentences, we use a statistical n-gram model to avoid those poorly constructed negative samples. So far, we’ve obtained all data to train the perspective models. We list the number of topics, retrieved sentences, and training samples in Table 2.

3.2 Out-of-domain Test Data

While training and testing on the same Wikipedia data is a possible choice, a more ideal scenario is to test on different domains to see if the distributional representation learned by the model generalizes to other datasets, not merely representing the style of Wikipedia. Hence, selecting a good held-out set to test the performance of our models is important.

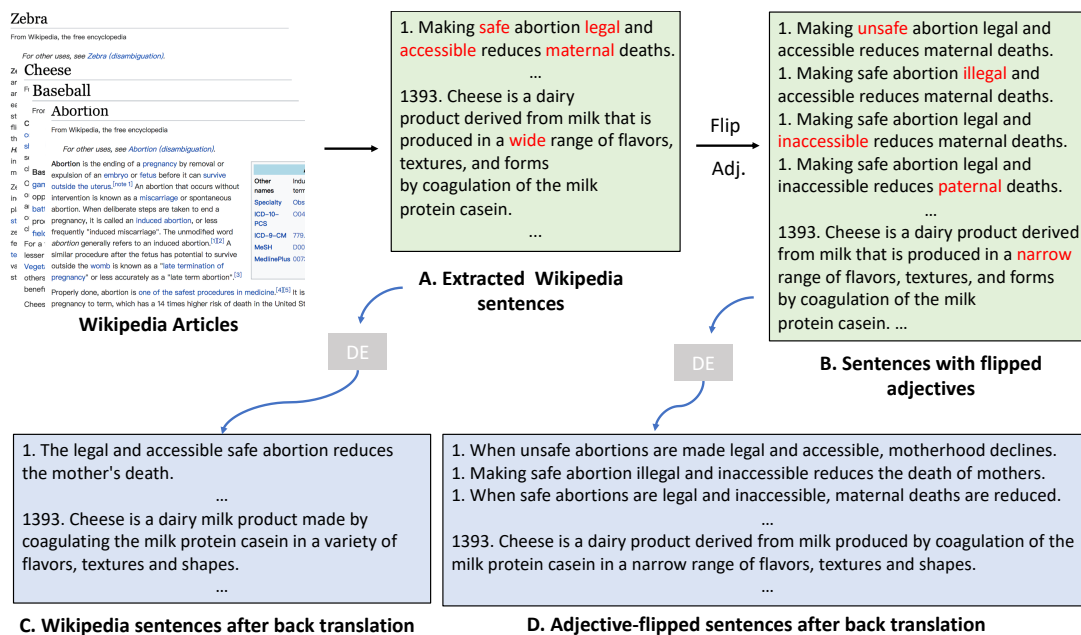


Figure 1: An illustration of the creation of the English training data. We first extract sentences from the retrieved Wikipedia articles to form the positive samples, and then replace adjectives with their antonyms as negative samples. *Back-translation* (discussed in 4.2) is then used to resolve pattern bias among negative samples. Note that we do not flip multiple adjectives simultaneously.

IMHO, what I find strange, and this is totally, some Chinese people have dogs as both pets and as dinner. ⁶
IMO, in an utopia Communism is the best system to live by.

Table 3: Sentence from the IMO dataset expressing opinions about which differ between cultures.

We are motivated by the fact that people always express personal opinions on social media such as Reddit, where many opinionated claims are included. We leverage a previous work (Chakrabarty et al., 2019) which collects a distant supervision-labeled corpus of 5.5 million opinionated claims covering a wide range of topics using sentences containing the acronyms IMO (in my opinion) or IMHO (in my humble opinion) from Reddit. Table 3 shows two examples from the IMO dataset that may reveal contrasting perspectives between two different colingual groups. As this dataset is only in English, to obtain scores from the Chinese and Japanese cultural models, we translate each sentence into the target language using the Youdao and Google Translate API⁵.

Test Data Selection. We first automatically extract claims that contain certain topical keywords, such as *free market* and *democracy*, and then remove the candidates which are out-of-context.

⁵<https://ai.youdao.com>, <https://translate.google.com>

⁶This does not reflect the opinion of the authors.

Then we ask the English and Chinese volunteers to jointly select high-quality statements. Finally, for human annotation, we select out 128 high-quality claims from over 2,000 candidates in the IMO/IMHO dataset. The topics include personal life, social and political views, etc.

3.3 Human Annotations for Test Data

For each test sample, we collect 20 annotations from annotators living in the United States using the Amazon Mechanical Turk platform (MTurk). We then collect another 20 annotations from Chinese/Japanese netizens using the SurveyHero/Crowdworks⁷ platform because MTurk is less used by the local people. The annotations are binarized, with 1 indicating agreement and 0 indicating disagreement. The average scores are viewed as the distributional scores.

For instance, for a given claim s^{en_i} , if 13 out of 20 English annotators give scores of 1, and the other 7 give scores of 0, then the human-annotated score $H^{en}(s^{en_i})$ equals $13/20 = 0.65$. In this way, we ensure that human annotation is of the same scale and meaning as the model prediction, and thus prove the validity of using the correlation between model predictions and human annotations as a measurement of effectiveness.

⁷<https://www.surveyhero.com>, <https://crowdworks.jp>

4 Methodology

In this section, we present the procedure of training our CLUSTER model. We explain how to learn group perspective models for English, Chinese, and Japanese colingual groups. We then raise the issue of pattern bias in negative samples and provide our corresponding solution. Lastly, we introduce the inference process.

4.1 Training Process

In the training stage, we leverage the pretrained multilingual BERT (Devlin et al., 2018) and fine-tune it for the perspective-specific classification task on the labeled data that is obtained in 3.1.

To enable the whole system to capture as much cultural discrepancy as possible, we separately fine-tune a BERT model for each language corpora despite the multilingualism of BERT. In other words, the learning steps of English, Chinese and Japanese systems have exactly the same structure but are completely isolated from each other in terms of training data and model parameters.

4.2 Pattern Bias in Negative Samples and Targeted Improvements

While flipping adjectives to create negative samples appears as an obvious approach, it ends up introducing certain style biases. Since the placeholders for adjectives are the only difference between positive and negative samples in training data, most classifiers would be able to identify this.

Niven and Kao (2019) show that high performance obtained from pre-trained language models such as BERT (Devlin et al., 2018) are often achieved by exploiting spurious statistical cues in the dataset. We face a similar problem in our preliminary study when evaluating on a test set from a different domain. While the quantitative results of our models trained on Wikipedia data are extremely high, we observe a huge drop when testing on out-of-domain data. This motivates us to mitigate statistical cues in our data.

Inspired by *back-translation* (Hoang et al., 2018), we generate paraphrases of our training data by introducing a pivot language and then translating the sentences back. This retains the semantics of the statements while removing existing stylistic biases. We back-translated both original Wikipedia sentences (i.e., positive samples) and the fabricated ones (i.e., negative samples). Part C and D of Figure 1 show the back-translated versions of our pos-

English (EN)	Chinese (CN)	Japanese (JP)
0.53	0.61	0.58

Table 4: Inter-rater agreement for English, Chinese and Japanese annotators using Krippendorff’s alpha, with p -value $< 1e-10$. All annotators show moderate agreement within their own group.

	Pearson correlation	Spearman correlation
English-Chinese (E-C)	0.26 (3e-3)	0.27(2e-3)
Chinese-Japanese (C-J)	0.49(5e-9)	0.50(2e-9)
Japanese-English (J-E)	0.29(7e-4)	0.30(6e-4)

Table 5: Cross-group rater agreement, in terms of *corr* (p -value). We measure the correlation between collective judgements on 128 claims by raters from each pair of the colingual groups: {E-C, C-J, J-E}.

itive and negative samples respectively.

4.3 Inference Process

The framework of our inference stage is similar to the training procedure except that we also test on out-of-domain data. For each claim s_i in test data, three model predictions $\{P^{en}(s^{en_i}), P^{cn}(s^{cn_i}), P^{jp}(s^{jp_i})\}$ are generated. We then compute the colingual perspective difference of s_i based on Equation 1. Finally, we compute the correlation between model-predicted scores and human annotations.

5 Experiments

5.1 Experimental Setup

For all classifiers, we start the sentence representations with BERT-base (Devlin et al., 2018) model, and then fine-tune them during training. We set sequence length as 128, batch size as 64 and learning rate as $2e-5$. We also study the efficiency of back-translation on reducing stylistic biases. Specifically, we train BERT models using data from 3 different settings: 1) no back-translation, 2) back-translate only negative samples, and 3) back-translate both positive and negative samples.

5.2 Binarization

We binarize the ground truth (with 0.5 as threshold) for the simplicity of data collection. Here 0 represents that a colingual group tends to maintain an opposite perspective, while 1 indicates a group tends to agree with the claim. For Wikipedia sentences, which we use for training and in-domain evaluation, the sentences originally selected from

Training \ Testing	No back-translation		Translate only negative		Back-translate both	
	Negative	Positive	Negative	Positive	Negative	Positive
No back-translation	85.15	88.74	65.23	72.55	67.82	64.61
Translate only negative	79.78	87.11	92.06	94.53	77.26	76.31
Back-translate both	92.56	94.88	92.17	95.69	87.10	91.92

Table 6: F1 scores of positive and negative class respectively, with models trained under three different settings: 1) neither the positive or the negative samples are back-translated, 2) only negative samples are back-translated, and 3) both positive and negative samples are back-translated. We then test them on the same held-out dataset.

Wikipedia are positive (1) while the one we modified algorithmically are negative (0).

5.3 Inter and cross-group rater agreement

To show how the annotators within a colingual group agree with each other, we calculate the inter-annotator agreement (IAA) using Krippendorff’s alpha. We also leverage attention questions to remove irresponsible annotators. The final IAAs are listed in Table 4. For all three languages, the correlation within a culture is above 0.5, demonstrating that the annotators are moderately correlated.

We also investigate how cross-group raters agree with each other, and calculate their Pearson and Spearman correlation (as listed in Table 5). The Chinese and Japanese raters have higher correlation with each other than they are with English raters.

5.4 Baselines

We compare our proposed Colingual Perspective Identifier (CLUSTER) with these baselines:

Random: Random numbers within $[0, 1]$ are generated to simulate model predictions of all perspective classifiers.

LM: We regard the average of word-level log probability (sentence log probability divided by length) generated by multi-lingual GPT2 (Radford et al., 2019; Zhang, 2019; Sakamoto, 2019) as model predictions. We then use the min-max method to normalize the log probabilities.

Weak CLUSTER: Our proposed Colingual Perspective Identifier, trained on Wikipedia sentences *without back-translation*.

6 Results

6.1 The Effects of Back-translation

Table 6 shows that models trained with *no back-translation* and *translate only negative* work well under their own respective setting, but does not transfer well to other scenarios. On the other hand, we obtain best and most robust results when the model is trained on data being back-translated for

both positive and negative samples. Hence, back-translation (for both positive and negative samples) is ideal to be used for inference in other domains.

6.2 Agreement between Model Prediction and Human Annotation

Table 7 reports the the correlations between the CLUSTER and baseline models with human annotations. We observe that the Random method does not capture any perspective representations at all. A competitive language model such as GPT-2 can bring significant improvements over Random because it is trained on a very large NLP corpus (including English Wikipedia), where group perspectives are implicitly included.

Moreover, the performance of Weak CLUSTER is partially better than language models, but still rather limited, probably due to style bias in negative samples. Finally, we can find that CLUSTER consistently outperforms all its competitors, and obtains $0.10 \sim 0.22$ performance gains over the second best model for all three colingual groups.

Last, we want to point out that unlike many other NLP tasks, the IAA (or human performance) should not be viewed as golden or an upper-bound in our evaluation. The IAA is just an indicator of how unanimous the annotators are on diverse concepts, including very controversial topics such as abortion. Therefore, machine-human correlation can reasonably be higher than within-human correlation.

6.3 Binary Accuracy

To further investigate the performance of our model and the baselines, we calculate the number of instances where binarized predictions and ground truths match with each other. The results are shown in Table 8. Again, our CLUSTER model achieved the best performance in all aspects.

7 Qualitative Analysis

While section 6 shows quantitative results and correlation values, we want to understand the advan-

Model	Correlation Type	English (EN)	Chinese (CN)	Japanese (JP)	Cross-culture (E-C)	Cross-culture (C-J)	Cross-culture (J-E)
Random	Pearson	0.00(0.5)	0.00(0.5)	0.00(0.5)	0.00(0.5)	0.00(0.5)	0.00(0.5)
	Spearman	0.00(0.5)	0.00(0.5)	0.00(0.5)	0.00(0.5)	0.00(0.5)	0.00(0.5)
LM	Pearson	0.17 (0.05)	0.07 (0.42)	0.12(0.19)	0.11 (0.23)	0.08(0.36)	0.15(0.09)
	Spearman	0.16 (0.08)	0.08 (0.35)	0.11(0.22)	0.09 (0.30)	0.09(0.33)	0.13(0.14)
Weak CLUSTER	Pearson	0.22 (0.01)	0.19 (0.03)	0.18(0.05)	0.03 (0.73)	0.05(0.61)	0.15(0.09)
	Spearman	0.11 (0.23)	0.13 (0.14)	0.10(0.28)	0.07 (0.42)	0.06(0.51)	0.11(0.23)
CLUSTER	Pearson	0.37 (1e-5)	0.41 (1e-6)	0.40(3e-6)	0.25 (4e-3)	0.20(0.02)	0.35(4e-5)
	Spearman	0.32 (2e-4)	0.34 (5e-4)	0.39(6e-6)	0.21 (0.01)	0.18(0.04)	0.31(4e-4)

Table 7: Agreement between model predictions and human annotations, in the format of *correlation (p-value)*. A higher value on Pearson correlation over Spearman correlation indicates that linear correlation is more significant than the rank correlation, and vice versa.

Model	EN	CN	JP	E-C	C-J	J-E
Random	0.50	0.50	0.50	0.50	0.50	0.50
LM	0.60	0.50	0.54	0.53	0.55	0.56
Weak CLUSTER	0.70	0.55	0.56	0.45	0.53	0.52
CLUSTER	0.73	0.64	0.66	0.58	0.60	0.63

Table 8: The binary accuracy. We test both within {EN, CN, JP} and across {E-C, C-J, J-E} groups. For scores within a culture ($\in [0, 1]$), the threshold is set to 0.5. For cross-group perspective scores ($\in [-1, 1]$), the threshold is set to 0.

tages of our model on a qualitative basis. To this end, we select 50 claims from five particular topics: *marriage*, *corruption*, *cuisine*, *christmas* and *baseball*, and then obtain CLUSTER model predictions on these claims. We do not collect human annotations for these sentences, but use them only for qualitative analysis and visualization purposes detailed below.

For each colingual group pair in {E-C, C-J, J-E} and a given topic, we report the visualization of 50 claim pairs in Figure 2 and 3. Here, each dot (or triangle) represents one of the 50 claims which are randomly selected from IMHO, with the x-y axis representing the {E-C, C-J, J-E} model predictions. The blue dots that fall along the diagonals are where the two models agree. On the contrary, dots that fall on the upper left or the lower right part are where the models do not agree with each other. For example, sentence 1 in Figure 2 is closer to the Chinese culture (upper left corner), while English speakers tend to agree more with sentence 2 (lower right corner). We select representative examples in each region and list them in the captions.

First, from Figure 2 we observe that the model pairs have zero or negative correlation on three topics: *marriage*, *corruption* and *cuisine*, suggesting that the corresponding language speakers take contrasting stances towards these topics. Second, Figure 3 shows that 1) the English and Chinese

speakers hold similar views on *baseball*, and 2) the Chinese and Japanese speakers share similar views on *christmas*. For example, Christmas, which is not a traditional holiday in East Asia, is adopted directly from the western world. The Chinese and Japanese speakers both follow the western customs and hence view Christmas likewise.

8 Related Work

Online Disagreement Most works about online disagreement focus on a single culture or language (Sridhar et al., 2015; Wang and Yang, 2015; Sridhar et al., 2015; Rosenthal and McKeown, 2015), thus are restricted to a single group. While these works try to computationally model disagreement or stance in debates, they do not target at finding cultural or cross-group differences. We, on the other hand, aim at understanding the disagreement in perspectives through different colingual groups according to their respective languages.

Cultural Study in Blogs or Social Media Nakasaki et al. (2009) present a framework to visualize the cross-cultural differences in multilingual blogs. Elahi and Monachesi (2012) show that using emotion terms as culture features is effective in analyzing cross-cultural difference in social media data. However, it is only restricted to a single topic (love and relationship). In contrast, we use Wikipedia to study cross-group differences in perspectives on a much larger scale and do not restrict ourselves to one single topic.

Cultural Difference in Word Usage Garimella et al. (2018) investigate the cross-cultural differences in word usages between Australian and American English through *socio-linguistic features* supervisedly. Garimella et al. (2016) use social network structures and user interactions, to study how to quantify the controversy of topics within a

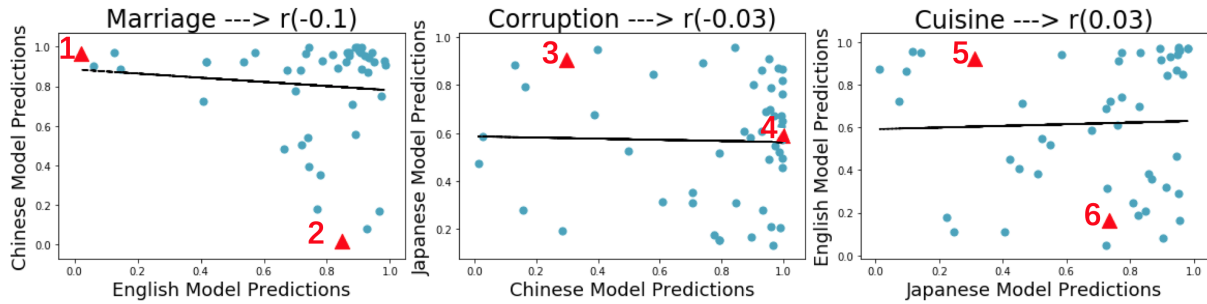


Figure 2: Model predictions on crosslingually disagreed topics: *marriage*, *corruption* and *cuisine* with their correlation values. Each dot (or triangle) represents one of the claims randomly selected from IMHO, with the x-y axis representing {E-C, C-J, J-E} predicted scores in sequence. Red triangular points are the following sentences: **1.** *Marriage is not about meeting someone you connect to, but both people being matured, and in the same headspace.* **2.** *If he cannot share his concerns with her, he is poor marriage material.* **3.** *If you don't reveal others' corruption you are culpable as well.* **4.** *There is plenty of corruption pulled out in the open these days, and that has been happening at a faster pace than ever before.* **5.** *Mexican, Mediterranean, Indian and Thai cuisines have the most delicious vegetarian dishes.* **6.** *Grilled fish is much better cooked at home and shared with friends.*

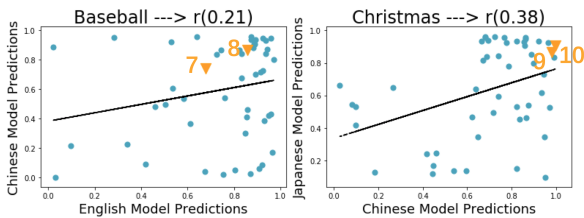


Figure 3: Model predictions on cross-lingually agreed topics: *baseball* and *christmas*, along with their correlation. The meaning of dots is the same as Figure 2. Orange triangles represent the following sentences: **7.** *Cricket is as fun to play as baseball if you limit the "innings" or overs.* **8.** *Things like basketball, baseball, tennis, golf, etc. are far more popular globally.* **9.** *Christmas, even minus the religious meanings, has good attributes in theory but has been too commercialized.* **10.** *I believe in giving gifts to kids because, Christmas is for children.*

culture and language. [Gutiérrez et al. \(2016\)](#) detect differences of word usage in the cross-lingual topics of multilingual topic modeling results. [Lin et al. \(2018\)](#) present distributional approaches to compute cross-cultural differences or similarities between two terms from different cultures focusing primarily on named entities. Our work is not limited to word usage or any particular topics. Instead, we focus on understanding cross-group differences of perspective at the sentence level.

Argumentation In argumentation, *Framing* is used to emphasize a specific aspect of a controversial topic. [Ajjour et al. \(2019\)](#) introduce frame identification, which is the task of splitting a set of arguments into non-overlapping frames. [Chen et al. \(2019\)](#) also release a dataset of claims, perspectives and evidence and propose the task of substanti-

ated perspective discovery where, given a claim, a system is expected to discover a diverse set of well-corroborated perspectives that take a stance with respect to the claim. Different interests, cultural and cultural backgrounds diverge people from on taking a certain course of action. While both works deal with different perspectives about arguments in English, our work focuses on identifying the differences from a cross-lingual point of view.

9 Conclusion

We present CLUSTER, a computational method to identify distributional differences in cross-group perspectives, and evaluate it with human judgements. Through detailed experiments, we show that CLUSTER is straightforward and effective. Furthermore, we show CLUSTER generalizes well for out-of-domain scenarios by training the group perspective models on Wikipedia and test on claims collected from Reddit. This means that the proposed method learns the task, not the data. Besides, the general model of perspective difference identification can be useful in many NLP tasks such as fact checking, sentiment analysis, as well as cross-cultural studies in computational social science or multilingual debate forums.

As a first attempt towards automatic identification of cross-cultural differences, our work still has much room for improvement. Future directions include more complicated ways of composing negative samples, more well-crafted models, and extending our pipeline to fine-grained subgroups speaking the same language, especially for English as a global language spoken by many nations.

Acknowledgments

This work was supported by the Defense Advanced Research Projects Agency (DARPA) and Army Research Office (ARO) under Contract No. W911NF-21-C-0002. The views expressed do not reflect the official policy or position of the Department of Defense or the U.S. Government. The authors would like to thank the members of PLUSLab at the University of California Los Angeles and University of Southern California and the anonymous reviewers for helpful comments.

References

- Yamen Ajjour, Milad Alshomary, Henning Wachsmuth, and Benno Stein. 2019. Modeling frames in argumentation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2922–2932.
- Daniel Bar-Tal. 2000. *Shared beliefs in a society: Social psychological analysis*. Sage Publications.
- David Bracewell and Marc Tomlinson. 2012. The language of power and its cultural influence. In *Proceedings of COLING 2012: Posters*.
- Ewa S Callahan and Susan C Herring. 2011. Cultural bias in wikipedia content on famous persons. *Journal of the American society for information science and technology*, 62(10):1899–1915.
- Tuhin Chakrabarty, Christopher Hidey, and Kathleen McKeown. 2019. Imho fine-tuning improves claim detection. *arXiv preprint arXiv:1905.07000*.
- Sihao Chen, Daniel Khashabi, Wenpeng Yin, Chris Callison-Burch, and Dan Roth. 2019. Seeing things from a different angle: Discovering diverse perspectives about claims. *arXiv preprint arXiv:1906.03538*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Mohammad Fazleh Elahi and Paola Monachesi. 2012. An examination of cross-cultural similarities and differences from social media data with respect to language use. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*, pages 4080–4086, Istanbul, Turkey. European Languages Resources Association (ELRA).
- Aparna Garimella, Rada Mihalcea, and James Pennebaker. 2016. Identifying cross-cultural differences in word usage. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 674–683, Osaka, Japan. The COLING 2016 Organizing Committee.
- Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Quantifying controversy on social media. *ACM Transactions on Social Computing*, 1(1):3.
- Susan A Gelman and Steven O Roberts. 2017. How language shapes the cultural inheritance of categories. *Proceedings of the National Academy of Sciences*, 114(30):7900–7907.
- E Dario Gutiérrez, Ekaterina Shutova, Patricia Lichtenstein, Gerard de Melo, and Luca Gilardi. 2016. Detecting cross-cultural differences using a multilingual topic model. *Transactions of the Association for Computational Linguistics*, 4:47–60.
- Kazi Saidul Hasan and Vincent Ng. 2014. Why are you taking this stance? identifying and classifying reasons in ideological debates. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 751–762, Doha, Qatar. Association for Computational Linguistics.
- James J Heckman. 1977. Sample selection bias as a specification error (with an application to the estimation of labor supply functions). Technical report, National Bureau of Economic Research.
- Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. Iterative back-translation for neural machine translation. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 18–24, Melbourne, Australia. Association for Computational Linguistics.
- Julie Khaslavsky. 1998. Integrating culture into interface design. In *CHI 98 conference summary on Human factors in computing systems*, pages 365–366. ACM.
- Bill Yuchen Lin, Frank F. Xu, Kenny Zhu, and Seungwon Hwang. 2018. Mining cross-cultural differences and similarities in social media. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 709–719, Melbourne, Australia. Association for Computational Linguistics.
- Hiroyuki Nakasaki, Mariko Kawaba, Sayuri Yamazaki, Takehito Utsuro, and Tomohiro Fukuhara. 2009. Visualizing cross-lingual/cross-cultural differences in concerns in multilingual blogs. In *Third International AAAI Conference on Weblogs and Social Media*.
- Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. 2019. Facebook fair’s wmt19 news translation task submission. *arXiv preprint arXiv:1907.06616*.

- Timothy Niven and Hung-Yu Kao. 2019. Probing neural network comprehension of natural language arguments. *arXiv preprint arXiv:1907.07355*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Sara Rosenthal and Kathy McKeown. 2015. I couldn't agree more: The role of conversational structure in agreement and disagreement detection in online discussions. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 168–177, Prague, Czech Republic. Association for Computational Linguistics.
- Toshiyuki Sakamoto. 2019. Japanese gpt2 generation model. <https://github.com/tanreinama/gpt2-japanese>.
- Dan Sperber and Lawrence A Hirschfeld. 2004. The cognitive foundations of cultural stability and diversity. *Trends in cognitive sciences*, 8(1):40–46.
- Dhanya Sridhar, James Foulds, Bert Huang, Lise Getoor, and Marilyn Walker. 2015. Joint models of disagreement and stance in online debate. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 116–125, Beijing, China. Association for Computational Linguistics.
- William Yang Wang and Diyi Yang. 2015. That's so annoying!!!: A lexical and frame-semantic embedding based data augmentation approach to automatic categorization of annoying behaviors using #petpeeve tweets. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2557–2563, Lisbon, Portugal. Association for Computational Linguistics.
- Zhibo Zhang. 2019. Gpt2-ml: Gpt-2 for multiple languages. <https://github.com/imcaspar/gpt2-ml>.

Appendix

A Hyper-parameters and other Experimental Settings

To train the classifiers, we start the sentence representations with the pre-trained BERT (Devlin et al., 2018) model, and then fine-tune them. For all models, we set sequence length as 128, batch size as 64 and learning rate as $2e-5$. We train each CLUSTER model for 5 epochs and save the best model only.

- Number of parameters:** Each CLUSTER model, fine-tuned on the BERT-base model, has 102M trainable parameters.
- Runtime:** Our average training time is 2 to 5 hours, depending on the size of training data for each language (see Table 2).
- Hardware configuration:** We use three GeForce RTX 2080 GPUs.
- Hyper-parameter tuning:** We manually tune the hyper-parameters and report the configuration that has the best F1 score on our validation set.

B Topics and Visualization

The sixteen topics that are selected for evaluation, along with the Pearson correlations of culture model predictions on 50 randomly sampled sentences, are listed in Table 9. We highlight the topics with relatively high and low values of correlation coefficients in red and blue. Note that we do not collect human annotations for these sentences, but use them only for qualitative analysis and visualization purposes.

As can be seen, most topics have a positive correlation, meaning that the English, Chinese and Japanese colingual groups have a general agreement on most subjects such as *savings*, *baseball* and *cheese*. Christmas, which is not a traditional holiday in China or Japan, is adopted directly from the western world. That’s why all the three models view Christmas likewise. In addition, the models have dispute on topics such as *bible*, *marriage*, *corruption*, and *abortion*. To get a more intuitive sense of the score distribution, we further visualize the model-predicted scores on more topics in Figure 4 and Figure 6.

Topics	E-C	C-J	J-E
Savings	0.09 (0.51)	0.40 (0.00)	0.31 (0.03)
Cuisine	0.01 (0.97)	0.21 (0.14)	0.03 (0.83)
Christmas	0.37 (0.01)	0.38 (0.01)	0.45 (0.00)
Bible	-0.02 (0.89)	0.16 (0.26)	0.16 (0.28)
Soup	0.26 (0.07)	0.15 (0.30)	0.22 (0.12)
Terrorism	0.09 (0.51)	0.07 (0.61)	0.20 (0.16)
Marriage	-0.10 (0.50)	0.21 (0.13)	0.04 (0.81)
Corruption	0.11 (0.44)	-0.04 (0.77)	-0.09 (0.54)
Baseball	0.21 (0.13)	0.13 (0.35)	0.11 (0.46)
Cheese	0.17 (0.23)	0.03 (0.83)	0.28 (0.05)
Communism	0.06 (0.67)	0.03 (0.83)	0.10 (0.48)
Democracy	0.09 (0.56)	-0.20 (0.16)	0.17 (0.23)
Russia	0.29 (0.04)	0.33 (0.02)	0.16 (0.27)
Abortion	-0.04 (0.77)	0.07 (0.65)	0.33 (0.02)
Racism	0.31 (0.03)	0.18 (0.20)	0.12 (0.40)
Gun control	0.07 (0.63)	0.15 (0.30)	0.19 (0.18)

Table 9: The sixteen topics that are selected for evaluation, along with the correlations between English-Chinese (E-C), Chinese-Japanese (C-J), and Japanese-English (J-E) culture model predictions on 50 randomly sampled sentences, in terms of *corr* (*p*-value).

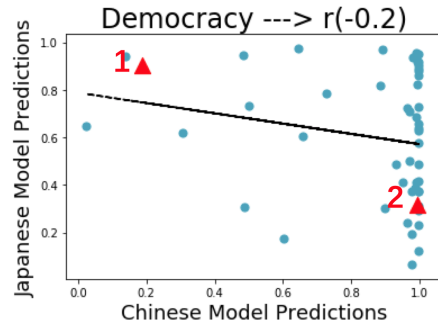


Figure 4: Model predictions on *democracy* of Chinese (x-axis) and Japanese (y-axis) models, and the correlation coefficient. The red triangles represent cross-lingually disagreed sentences: **1.** *Yeah, mandatory voting should be a required part of a democracy.* **2.** *The ideal system would be a merger of democracy and socialism (which we are slowly moving towards).*

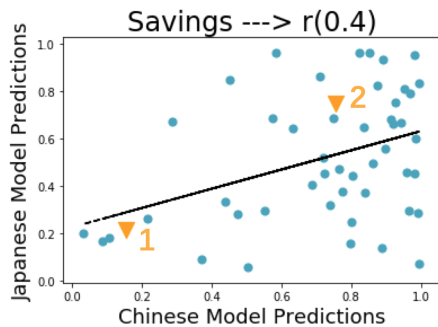


Figure 5: Model predictions on *savings* of Chinese (x-axis) and Japanese (y-axis) models, and the correlation coefficient. The orange triangles represent cross-lingually agreed sentences: **1.** *Higher risk-free interest is needed to stimulate savings and to avoid credit recessions.* **2.** *Life savings essentially means to me what you are gonna leave to your heirs.*

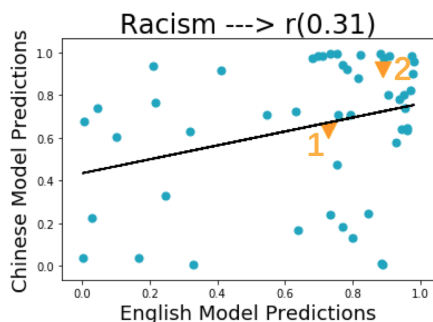


Figure 6: Model predictions on *racism* of English (x-axis) and Chinese (y-axis) models, and the correlation coefficient. The orange triangles represent cross-lingually agreed sentences: **1.** *Racism is the prejudice against other cultures through identification of physical appearance and cues.* **2.** *Fat shaming and/or body shaming can be just as bad as racism or homophobia.*

C Additional Details for Data Collection

C.1 Training Data

We have described the procedure of collecting our training data from multi-lingual Wikipedia articles in Section 3.1. In addition, for pre-processing details, we utilize Jieba⁸ and Mecab⁹ to tokenize Chinese and Japanese sentences.

Back-translation (discussed in Section 4.2) is the backbone of our CLUSTER model. Table 10 show the pivot languages as well as different translation systems used for our English, Chinese and Japanese models.

	Pivot language	Translation model
English	German	Fined-tuned transformer (Ng et al., 2019)
Chinese	Japanese	Youdao API
Japanese	Chinese	Google Translate

Table 10: The pivot languages and translation systems that we use for back-translation.

C.2 Questionnaire for Selecting Test Data

We design questionnaires to select out meaningful and high-quality claims from the original IMO/IMHO dataset (discussed in Section ??), and collect three answers per claim. Figure 7 shows our instructions to the English annotators on the Amazon Mechanical Turk (MTurk) platform.

The turkers are asked to give a categorical score to each candidate sentence. The categorical score ranges from 1 to 3, with 1 indicating not meaningful, incoherent, or talking about facts, 2 indi-

cating somewhat meaningful but few people have opinions on it, and 3 indicating highly meaningful. Because we extract single sentences from online discussion forums, we ask the turkers to ignore the out-of-context words such as ‘and’, ‘also’, and ‘but’, and focus on the opinion only. Finally, if all annotators agree that a given claim is meaningful enough so that other people will hold a stance (either agreement or disagreement) towards it, we regard this candidate claim as one of our test samples for the final human annotation step.

C.3 Questionnaire for Collecting Human Annotation

Figure 8 is an English demonstration of our survey to collect human annotations of the test data. The annotators as instructed to read each sentence carefully, and give a binary score to each sentence based on their personal opinions. The score is either 0 or 1, with 1 indicating they mostly agree with this statement, and 0 indicating they mostly do not agree with it, or don’t know what this statement is talking about.

Besides, we adopt attention checks to control the quality of our collected annotations. To this end, we manually select 7 facts from Wikipedia as attention check statements, which are obviously true to the masses, such as ‘*Cheese is a dairy product derived from milk that is produced in a wide range of flavors, textures, and forms*’. We insert an attention check statement after every 9 test claims. If an annotator does not agree with one of our attention check statements, his entire HIT is rejected. Each annotator is allowed to annotate at most 20 sentences including the attention check statements.

⁸<https://pypi.org/project/jieba/>

⁹<https://pypi.org/project/mecab-python3/>

Survey Instructions (Click to expand)

In this survey, you will be provided with 10 sentences. Each sentence expresses its opinion on one of the following topics: **Terrorism, Marriage, and Sports**. Please read each sentence carefully, and give a categorical score to each sentence based on the meaningfulness of the sentences. The score ranges from 1 to 3, with 1 indicating less meaningful or talking about details, 2 indicating somewhat meaningful but few american people can have a opinion on it, and 3 indicating very meaningful and you believe people can have a stance (agree or disagree) towards this statement.

Also, note that we extract single sentences from online discussion forums. Although we have made sure each selected sentence is within context, there are still out-of-context conjunctions such as 'and', 'also', and 'but'. Please ignore the out-of-context words and focus on the opinion only.

For example:

- 1). So it is easy to say that without ww2 communism would have taken continental europe.
 - 1. Not meaningful, incoherent, talking about facts or details.
 - 2. Somewhat meaningful but few american people can have a opinion on it, or requires expert knowledge.
 - 3. Very meaningful and I believe many people can have a stance (agree or disagree) towards this statement.
- 2). Section 44 of the terrorism act 200 is a good example.
 - 1. Not meaningful, incoherent, talking about facts or details.
 - 2. Somewhat meaningful but few american people can have a opinion on it, or requires expert knowledge.
 - 3. Very meaningful and I believe many people can have a stance (agree or disagree) towards this statement.
- 3). Kala, arular, and piracy funds terrorism are all 100x better than most of the stuff on matangi.
 - 1. Not meaningful, incoherent, talking about facts or details.
 - 2. Somewhat meaningful but few american people can have a opinion on it, or requires expert knowledge.
 - 3. Very meaningful and I believe many people can have a stance (agree or disagree) towards this statement.

Figure 7: The instructions of our survey to evaluate the meaningfulness of the IMO/IMHO sentences on the MTurk platform.

Survey Instructions (Click to expand)

In this survey, you will be provided with 10 sentences. Each sentence expresses its opinion on one of the following topics: **Food, Democracy, and Privacy**. Please read each sentence carefully, and give a binary score to each sentence based on your personal opinions. The score is either 0 or 1, with 1 indicating you mostly agree with it, and 0 indicating you mostly disagree with this statement or don't know what this statement is talking about. If you feel somewhat in the middle, please select the one that is closer to your thoughts.

Also, note that we extract single sentences from online discussion forums. Although we have made sure each selected sentence is within context, there are still out-of-context conjunctions such as 'and', 'also', and 'but'. Please ignore the out-of-context words and focus on the opinion only.

For example:

- 1). So it is easy to say that without ww2 communism would have taken continental europe.
 - 1. Overall, I agree with this statement.
 - 0. Overall, I disagree with this statement or I don't know what this statement is talking about.
- 2). Well america has only the bare bones of a functioning social democracy, so european-style social democracy would be great for america.
 - 1: Overall, I agree with this statement.
 - 0. Overall, I disagree with this statement or I don't know what this statement is talking about.

Figure 8: An example instruction page of our survey to collect the human annotations on the MTurk platform.