

Wisdom of Instruction-Tuned Language Model Crowds. Exploring Model Label Variation

Flor Miriam Plaza-del-Arco, Debora Nozza, Dirk Hovy

MilaNLP, Bocconi University, Department of Computing Sciences, Milan, Italy

{flor.plaza, debora.nozza, dirk.hovy}@unibocconi.it

Abstract

Large Language Models (LLMs) exhibit remarkable text classification capabilities, excelling in zero- and few-shot learning (ZSL and FSL) scenarios. However, since they are trained on different datasets, performance varies widely across tasks between those models. Recent studies emphasize the importance of considering human label variation in data annotation. However, how this human label variation also applies to LLMs remains unexplored. Given this likely model specialization, we ask: *Do aggregate LLM labels improve over individual models (as for human annotators)?* We evaluate four recent instruction-tuned LLMs as “annotators” on five subjective tasks across four languages. We use ZSL and FSL setups and label aggregation from human annotation. Aggregations are indeed substantially better than any individual model, benefiting from specialization in diverse tasks or languages. Surprisingly, FSL does not surpass ZSL, as it depends on the quality of the selected examples. However, there seems to be no good information-theoretical strategy to select those. We find that no LLM method rivals even simple supervised models. We also discuss the tradeoffs in accuracy, cost, and moral/ethical considerations between LLM and human annotation.

Keywords: model annotation, model label variation, subjective tasks, label aggregation, ethics

1. Introduction

Large Language Models (LLMs) have revolutionized many aspects of Natural Language Processing (NLP). Brown et al. (2020) showed that LLMs have few-shot (FSL) and even zero-shot learning (ZSL) capabilities in text classification due to their extensive pre-training. Subsequent iterations have further improved these capabilities. Those improvements have seemingly obviated one of the most time- and labor-intensive aspects of NLP: annotating enough data to train a supervised classification model. Instead, we can use LLMs to directly predict the labels via prompting. Indeed, various papers tested this hypothesis and found good performance on various NLP tasks (Zhao et al., 2023; Su et al., 2022; Wei et al., 2022; Brown et al., 2020; Plaza-del-Arco et al., 2023). However, upon closer inspection, these claims require some caveats: different models excel at different tasks, datasets, and formulations (Gilardi et al., 2023; Törnberg, 2023).

What if the answer is not to wait for one model to rule them all, though, but to treat their variation as specializations we can exploit, similar to the disagreement among human annotators? Different annotators have different strengths (or levels of reliability), and recent work (Basile et al., 2021; Plank, 2022) has suggested using this human label variation to our advantage. We test whether the same applies to LLMs if we treat them as “annotators”.

We use four state-of-the-art open-source instruction-tuned LLMs to assess their capabilities as “annotators”: Flan-T5 (Chung et al., 2022), Flan-UL2 (Tay et al., 2023), T0 (Sanh et al., 2022)

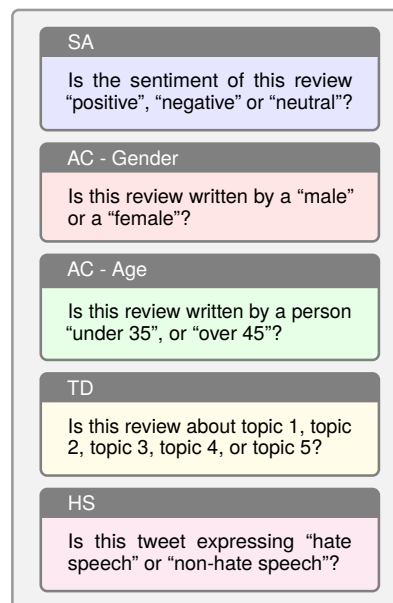


Figure 1: Instructions used to prompt the instruction-tuned LLMs for each classification task.

(alongside its multilingual variant, mT0 (Muenighoff et al., 2022)), and Tk-Instruct (Sanh et al., 2022). (We use open models to mitigate potential concerns regarding data contamination during the evaluation process and to facilitate replication.) We evaluate them across five subjective prediction tasks (age, gender, topic, sentiment prediction, and hate speech detection) in four distinct languages: English, French, German, and Spanish. We use ZSL and FSL prompt instructions, similar to the

ones we would give to human annotators. For FSL, we explore different entropy-based strategies to choose the seed examples. We then aggregate the LLM answers into a single label and evaluate their performance.

We find that different models indeed excel on some tasks or languages, but not on others. Some models even specialize on certain labels in a given task, but perform poorly on the others. **Their behavior thus mimics human expertise in wisdom-of-the-crowd settings.**

We then aggregate the model answers into a single label for each example. The simplest approach is *majority voting*: use the label that most LLMs suggested (ties are split randomly). However, the majority can still be wrong. Instead, we can use *Bayesian models of annotation* (Passonneau and Carpenter, 2014; Paun et al., 2018) to weigh the answers based on the inferred reliability of each annotator. This approach is similar to Bayesian classifier combination, but does not require gold labels to assign the scores. Instead, it is completely unsupervised. That distinction is crucial, as we want to work with unannotated data.

In most cases, the aggregated labels of either method outperform even the best individual LLM. On average, aggregated labels are 4.2 F1-points better than the average LLM. However, even the best-aggregated performance is still well below that of even simple supervised models trained on the same data, and substantially lower than Transformer-based supervised models (by over 10 F1 points on average).

In sum, aggregating several ZSL-prompted LLMs is better than using a single LLM. Surprisingly, FSL-prompting is too varied to consistently improve performance. **However, treating LLMs as annotators cannot rival using human annotators for fine-tuning or supervised learning.** We also discuss what these results mean for the role of human annotation and supervised learning in NLP, with respect to performance, but also time, cost, bias, and ethics.

Contributions (1) We explore the feasibility of four open-source instruction-tuned LLMs as “annotators” via ZSL and FSL prompting on five subjective classification tasks; (2) we compare them across four languages; (3) we analyze the robustness of two label aggregation methods to check whether we can benefit from model label variation in subjective tasks; and (4) we discuss the technical, moral, and ethical ramifications of this development for NLP and annotation.

2. Data

For our experiments, we use two datasets: **Trustpilot** (Hovy et al., 2015) and **HatEval** (Basile et al., 2019). Note that for most models, these datasets are “unseen”, i.e., the data was not part of the LLMs’ training. The one exception is HatEval in EN, which is included in Flan-T5 and Flan-UL2 models. We aim to evaluate their performance in a data contamination scenario, offering insights into models’ generalization capabilities unaffected by such contamination.

Trustpilot (Hovy et al., 2015) is a multilingual dataset with demographic user information from various countries. It uses reviews from the user review website Trustpilot. To test a variety of languages commonly found in LLMs, we select data with English from the United States, German from Germany, and French from France for our experiments. The data includes labels for sentiment, the topic of the review, and two demographic dimensions of the review authors: self-declared gender and age (these two are not available for all data points). We use the same data splits as Hovy (2015) to ensure comparability and consistency. Given our ZSL setup, we evaluate on their evaluation sets for each language, which consists of the joint development and test sets.

HatEval (Basile et al., 2019) is a multilingual dataset for HS detection against immigrants and women on Twitter, part of a SemEval 2019 shared task. The dataset contains Spanish and English tweets manually annotated via crowdsourcing. We use the benchmark test set provided by the HatEval competition for both languages.

2.1. Tasks

We evaluate the performance of LLMs as annotators on five prediction tasks: four from the Trustpilot dataset and one from the HatEval corpus. These tasks involve sentiment analysis, topic detection, and predicting demographic attributes (gender and age). These two **attribute classification (AC)** tasks are binary: the **gender** of the text author (*male* or *female*)¹ and the **age** of the text author (*under 35* or *above 45* years old). In the **sentiment analysis (SA)** task, reviews are classified into *negative*, *neutral*, and *positive* sentiments based on 1, 3, and 5-star ratings, respectively. The **topic detection (TD)** task uses the review categories of the texts to classify them into one of five topics. For this task, the exact topics vary across languages.

¹The data does not allow a more fine-grained classification of gender identities, as the original website only provided users with those two options. See Ethical Considerations for more discussion.

EN: Car lights, fashion accessories, pets, domestic appliances, and hotels. **DE:** Wine, car rental, drugs and pharmacy, flowers, and hotels. **FR:** Clothes and fashion, fashion accessories, pets, computer and accessories, and food and beverage.

For the **hate speech detection (HS)**, the task is to classify a tweet as either hate speech or non-hate speech.

3. Models

We experiment with four state-of-the-art instruction-tuned LLMs from the same model family, the T5 with an encoder-decoder (Raffel et al., 2020) architecture. We specifically select these models because they were fine-tuned on a diverse range of instructions. They use intuitive explanations of the downstream task to respond to natural language prompts, similar to the instructions provided to human annotators. Furthermore, these models are all open-source, letting us inspect the training data and examine data contamination. Our selection represents a realistic LLM annotator pool for a current NLP practitioner. As models evolve rapidly, though, this selection is likely to change. However, the results from using a diverse pool of LLM annotators should hold regardless.

In particular, we use the following models:

- **Flan-T5** (Chung et al., 2022) is a sequence-to-sequence transformer model built on the T5 architecture (Raffel et al., 2020). The model has been pre-trained with standard language modeling objectives and subsequent fine-tuning on the extensive FLAN collection (Longpre et al., 2023). The FLAN collection contains more than 1,800 NLP tasks in over 60 languages. We use the largest version² of this model.
- **Flan-UL2** (Tay et al., 2023) is the Flan version of the T5 and UL2 model. It has a similar architecture to T5, but with an upgraded pre-training procedure known as UL2³.
- **T0** (Sanh et al., 2022) and the multilingual **mT0** (Muennighoff et al., 2023). T0⁴ is an encoder-decoder model based on T5 that is trained on a multi-task mixture of NLP datasets over different tasks. For the non-English languages, we use mT0⁵ since T0 has been trained on English texts. mT0 is based on Google’s mT5 (Xue et al., 2021) and has been fine-tuned on xP3⁶, which covers 13 training tasks across

²<https://huggingface.co/google/flan-t5-xxl>

³<https://huggingface.co/google/flan-ul2>

⁴<https://huggingface.co/bigscience/T0>

⁵<https://huggingface.co/bigscience/mt0-xxl>

⁶<https://huggingface.co/datasets/bigscience/xP3>

46 languages with English prompts.

- **Tk-Instruct** (Wang et al., 2022) is a generative model for transforming task inputs given declarative in-context instructions, like “Given an utterance and the past 3 utterances, output ‘Yes’ if the utterance contains the small-talk strategy, otherwise output ‘No’. Small-talk is a cooperative negotiation strategy...” (adapted from Wang et al., 2022). It is also based on T5 but trained on all task instructions in a multi-task setup. It is fine-tuned on the SUPER-NATURALINSTRUCTIONS dataset (Triantafillou et al., 2020), a large benchmark of 1,616 NLP tasks and their natural language instructions. It covers 76 task types across 55 different languages⁷.

Computing Infrastructure We run all experiments on a server with three NVIDIA RTX A6000 and 48GB of RAM.

3.1. Prompting

A prompt is an input that directs an LLM’s text generation, ranging from a single sentence to a paragraph. It guides the model’s comprehension and influences its output. Figure 1 depicts the task formulations (prompt instructions) we give to the LLMs, who act as our annotators, for every considered text classification task. We add “Answer” to mark the output field after the instruction to improve the LLMs’ understanding and output format. For the TD tasks, the list of five topics varies by language. For instance, the prompt for the English TD task is: “I love the earrings I bought,” “Is this review about ‘car lights,’ ‘fashion accessories,’ ‘pets,’ ‘domestic appliances,’ or ‘hotels?’” <Answer>: {LM answer}.

Tk-Instruct requires a prompt template with specific fields: “definition,” “input,” and “output.” The “definition” is used to specify the instruction or guidance, the “input” contains the instance to be classified, and the “output” is the output indicator. For instance, the prompt for the HS task is the following: <Definition> Is this tweet expressing “hate speech” or “non-hate speech?” <Input> “I hate you” <Response>: {LM answer}.

We use task-specific prompts to assess the model’s performance on the resulting outputs for zero- and few-shot prediction. We used default parameters for the models.

If the output does not correspond to a valid class, we assign the most common class for that task. For instance, these out-of-label (OOL) predictions vary between tasks and models for the ZSL setup. Binary or ternary classification tasks (AC, SA, HS)

⁷<https://huggingface.co/allenai/tk-instruct-3b-def>

exhibit a very low OOL percentage (<1%). In contrast, TD shows a significantly higher percentage (13%) due to the larger number of classes and their more descriptive nature (e.g., “fashion accessories”). At the model level, Flan models have a very low OOL percentage (1%), T0 and Tk-Instruct have a low OOL percentage (~2%).

3.2. Baselines

We compare the LLMs across our five tasks to two baselines: *the most frequent class* and *random choice*.

The *most frequent class* baseline does not require any model. It always picks the most frequent label for a task as final prediction. This standard baseline method is very strong in unbalanced datasets. However, it requires knowledge of the label distribution. The *random-choice* method randomly picks a label from the set of labels for a task. It represents a lower bound.

3.3. Aggregation of Labels

For each example, we get four labels: one from each LLM annotator. We use two different methods to aggregate these four labels into a single label: *majority voting*, and a *Bayesian model of annotation*, Multi-Annotator Competence Estimation (MACE, Hovy et al., 2013). These methods use different aggregation methods. Both are common in the literature (Klie et al., 2023).

Majority-voting selects the label returned most by the four models. In case of a tie, it randomly chooses among the top candidates. This approach is common in many annotation projects, but has the drawback that the majority can still be wrong.

MACE is a Bayesian annotation tool that computes two scores: the competence (reliability) of each annotator (i.e., the probability an annotator chooses the “true” label based on their expertise instead of guessing one) and the most likely label. MACE uses variational Bayesian inference to infer both variables, and works on unlabeled data. The aggregated MACE labels are usually more accurate downstream than majority voting, and competence scores correlate strongly with actual annotator proficiency (Paun et al., 2018). Competence scores tend to correlate with annotators’ actual expertise (Hovy et al., 2013), and can therefore be used to directly compare annotator quality in the absence of gold labels. As a probabilistic model, it also computes the entropy of each example, including both annotator competence and agreement. It is therefore a proxy for how “difficult” an example is to label. We use MACE to get the competence of each LLM and the entropy of each example, which we use to select seed examples for FSL.

Task/Language		Cohen	Fleiss	Krip.	Raw
SA	EN	0.708	0.705	0.703	0.837
	DE	0.636	0.633	0.630	0.792
	FR	0.665	0.662	0.660	0.809
AC-Gender	EN	0.299	0.279	0.229	0.615
	DE	0.271	0.136	-0.007	0.566
	FR	0.236	0.227	0.179	0.596
AC-Age	EN	0.101	0.044	-0.154	0.428
	DE	0.068	0.040	-0.124	0.596
	FR	0.099	0.093	0.014	0.679
TD	EN	0.510	0.495	0.477	0.622
	DE	0.586	0.578	0.571	0.712
	FR	0.316	0.305	0.283	0.598
HS	EN	0.222	0.220	0.209	0.605
	ES	0.155	0.099	-0.019	0.629

Table 1: Inter-model agreement scores.

4. Results

We compare the four models as annotators along several dimensions:

How much do models agree with each other?

This assesses the consensus among them and indicates specialization. **How reliable is each model?** This evaluates the consistency and trustworthiness of individual model predictions, key for label aggregation. **How accurate are the predictions of the individual models versus their aggregations?** This last question assesses the prediction quality of individual LLMs vis-a-vis aggregations to determine whether this approach is a viable alternative to supervised learning.

We first report ZSL results and then discuss the FSL setting separately (Section 4.4).

4.1. Inter-model Agreement

To assess the level of specialization among the LLMs as annotators, we evaluate their agreement. We use four common inter-annotator-agreement metrics: Cohen’s κ (Cohen, 1960), Fleiss’ κ (Fleiss, 1971), and Krippendorff’s α (Krippendorff, 2011) (which all correct the observed agreement for expected agreement), as well as the unweighted raw agreement (i.e., the uncorrected level of agreement between LLMs). The results are shown in Table 1. Note that the number of labels does not factor into agreement, and that raw agreement is usually higher than chance-corrected inter-annotator agreement measures.

The results show a wide range of agreement values, but a few takeaways become apparent:

1) **The scores suggest that the different models specialize on different tasks and labels.** As we will see in the performance and reliability analysis, some models perform better on some tasks

than others. Model specialization suggests that aggregation is likely beneficial (as the aggregation hopefully benefits from differing expertise).

2) **Language does not factor into the differences.** The models we test are all multi-lingual, and the languages we test are generally high-resource. The agreement difference between the different languages on the same task is negligible.

3) **Some tasks show higher agreement than others:** SA has higher scores than TD, and the others have little to no agreement. However, we do not know whether high-agreement tasks are inherently easier, or whether the models are all wrong in the same direction.

Task/Language	T0	Flan-T5	Flan-UL2	Tk-Instruct	
SA	EN	0.755	0.958	0.856	0.803
	DE	0.724	0.928	0.767	0.757
	FR	0.759	0.909	0.796	0.789
AC-Gender	EN	0.317	0.613	0.719	0.445
	DE	0.152	0.267	0.359	0.270
	FR	0.094	0.699	0.599	0.601
AC-Age	EN	0.255	0.146	0.325	0.083
	DE	0.295	0.039	0.418	0.014
	FR	0.432	0.061	0.608	0.017
TD	EN	0.388	0.737	0.933	0.511
	DE	0.660	0.793	0.735	0.712
	FR	0.556	0.327	0.492	0.210
Mean	EN	0.429	0.614	0.708	0.461
	DE	0.458	0.507	0.569	0.438
	FR	0.460	0.499	0.624	0.404
HS	EN	0.358	0.919	0.327	0.402
	ES	0.466	0.212	0.131	0.099

Table 2: MACE competence scores of each LLM across tasks and languages on the Trustpilot and HatEval datasets. For non-English languages, we use the multilingual mT0 model.

4.2. Reliability

When aggregating specialized annotators, we might want to trust more specialized ones more. We use the competence scores from MACE to assess the reliability of each model. Table 2 shows the competence scores.

The competence scores support the specialization hypothesis for the different models on different languages and tasks. **No model is dominant in all settings, though the Flan models tend to have higher competence scores than the other models** (reflected in their higher mean competence scores).

4.3. Model Performance and Robustness of Label Aggregation

Ultimately, we care about the predictive performance of the annotator method. Table 3 shows the macro-F1 scores of the LLMs on all tasks and languages. We compute the statistical difference of the individual results over the random-choice baseline, using a bootstrap sampling test with the *bootsa*⁸ Python package. We use 1,000 bootstrap samples, a sample size of 20%, and $p \leq 0.01$.

Most models clearly and significantly outperform the random-choice and even most-frequent-label baselines. Note though that Flan-T5 and Flan-UL2 included the HatEval dataset in their training. Consequently, they perform substantially better than the other models (with Flan-T5 receiving a very high competence score from MACE).

Aggregation When aggregating annotations into a single label, we implicitly assume that a) there is a single correct answer and b) the wisdom of the crowd will find it. The first assumption is up for debate (Basile et al., 2021), but the latter is clearly borne out by the results here. **On average and in most individual cases, majority voting and MACE aggregation predictions are better than most models.** In 6 out of 14 tasks, MACE was the best model. Note that for SA, Tk-Instruct performs better than the aggregation methods in all languages. For AC-Gender in English, Flan-UL2 is better, and in German, no method outperforms random choice (though MACE aggregation is close).

Overall, the two aggregation methods are substantially more robust than any one individual model across all languages and datasets (see the Mean results in Table 3). Presumably, they suffer less from the variance across tasks and languages and instead are able to exploit the specialization of each model as a source of information. The MACE competence score correlates with the actual performance of the models: 0.93 Spearman ρ and 0.83 Pearson ρ . This correlation suggests that MACE identified the model specializations correctly. A custom weighting of each model’s prediction (for example, based on the actual performance) might perform even better. In practice, though, this weighting would of course be unknown.

Comparison to supervised learning ZSL holds a lot of promise for quick predictions, but to assess its worth, we need to compare it to supervised models based on human annotation. For the Trustpilot data, we compare our best ZSL result for each task and language (see Table 3) to two supervised models. 1) a simple Logistic Regression model (the baseline “agnostic” results reported in Hovy, 2015)

⁸<https://github.com/fornaciari/bootsa>

Task/Lang.		Models				Baselines		Aggregate	
		T0	Flan-T5	Flan-UL2	Tk-Instruct	Most Freq	Random	Majority	MACE
SA	EN	0.453*	0.532*	0.482*	0.553*	0.167	0.334	0.503*	0.514*
	DE	0.469*	0.495*	0.433*	0.517*	0.167	0.331	0.480*	0.484*
	FR	0.460*	0.518*	0.445*	0.528*	0.167	0.337	0.486*	0.490*
AC-Gender	EN	0.516	0.594*	0.624*	0.541*	0.337	0.501	0.617*	0.623*
	DE	0.456	0.437	0.447	0.431	0.334	0.497	0.458	0.485
	FR	0.428	0.573*	0.566*	0.563*	0.335	0.503	0.577*	0.579*
AC-Age	EN	0.495	0.442	0.516*	0.397	0.336	0.497	0.569*	0.572*
	DE	0.458	0.366	0.503	0.344	0.334	0.500	0.422	0.499
	FR	0.497	0.375	0.550*	0.343	0.335	0.500	0.443	0.542*
TD	EN	0.558*	0.579*	0.588*	0.567*	0.085	0.195	0.588*	0.596*
	DE	0.506*	0.514*	0.513*	0.493*	0.105	0.193	0.516*	0.520*
	FR	0.314*	0.271*	0.264*	0.257*	0.096	0.193	0.281*	0.293*
Mean	EN	0.506	0.537	0.553	0.515	0.231	0.382	0.569	0.576
	DE	0.472	0.453	0.474	0.446	0.235	0.380	0.469	0.497
	FR	0.425	0.434	0.456	0.423	0.233	0.383	0.447	0.476
HS	EN	0.621*	0.726*	0.670*	0.579*	0.367	0.490	0.717*	0.726*
	ES	0.601*	0.532*	0.519	0.449	0.370	0.492	0.533*	0.603*

Table 3: Zero-shot Macro-F1 results obtained by the LLMs on the Trustpilot and HatEval tasks, the baselines and the aggregation methods. Best result per language and task is shown in bold. Significant improvement over Random baseline (* : $p \leq 0.01$) with bootstrap sampling. For non-English languages, we use the multilingual mT0 model.

and a recent Transformer-based model (the best results from Hung et al., 2023). Similarly, for HatEval, we compare with 1) a simple linear Support Vector Machine based on a TF-IDF representation (the baseline results reported in Basile et al., 2019) and a fine-tuned multilingual Transformer model (Nozza, 2021). Table 4 shows the results. The two methods approximate an upper and lower bound on supervised learning for these datasets.

Task/Language		best	supervised	
		ZSL	Standard ML	Transformer
SA	EN	0.553	0.610	0.680
	DE	0.517	0.610	0.677
	FR	0.528	0.612	0.706
AC-Gender	EN	0.624	0.601	0.638
	DE	0.497	0.540	0.629
	FR	0.579	0.546	0.650
AC-Age	EN	0.572	0.620	0.636
	DE	0.503	0.602	0.611
	FR	0.550	0.540	0.568
TD	EN	0.596	0.656	0.705
	DE	0.520	0.605	0.671
	FR	0.314	0.385	0.444
HS	EN	0.726	0.451	0.416
	ES	0.603	0.701	0.752

Table 4: Macro-F1 results for best ZSL model (Table 3), compared to previous supervised results on the same datasets.

Except for 4 cases (AC-Gender and AC-Age in French, AC-Gender in English, HS in English), even the simple ML models beat the best ZSL result we achieved, be it from an LLM or aggregation method. Compared to the upper bounds from Hung et al. (2023), we see an average performance gap of 10.5 F1 points. Only for HS in English ZSL achieves better results, likely attributed to the data contamination found in the Flan models.

These results show that while ZSL might be a fast approximation for prediction tasks, it is still far from competitive with supervised learning.

4.4. Few-shot Learning

FSL has the potential to perform better than ZSL, so we ask whether using FSL models as annotators improves over our ZSL experiments. We investigate whether providing a limited set of examples enhances annotation capabilities, similar to instructing human annotators. The short answer is no. We apply this method to the English Trustpilot dataset.

To choose the seed examples, we compare three methods. Random selection and selection based on the MACE entropy scores⁹. Entropy lets us identify two groups of examples: (1) maximum entropy indicates models were less confident or disagreed more, indicating higher difficulty, and

⁹<https://github.com/dirkhovoy/MACE>

(2) low entropy indicates models were more confident or agreed more, indicating lower difficulty. For each task and label, we randomly choose 4,000 instances¹⁰ and use MACE to compute the entropy for each instance. From the initial pool of 4,000, we sample three exemplars per class based on the method (low entropy, max entropy) and use these as few-shot seeds, prepending them to the prompt. We compare these results to the random baseline.

Figure 2 shows a comparison between the ZSL and FSL approaches. **Our analysis reveals that there are no statistically significant differences between the two prompting methods.** In general, though, FSL does not perform as well as ZSL across subjective tasks. Within FSL, a prominent pattern emerges: it exhibits notably higher variance across tasks than ZSL. Presumably, exemplar quality heavily influences performance.

Regarding the two entropy-based selection methods, our results show no consistent trends between them. The choice is somewhat task-dependent: Max entropy seems to perform well for SA and AC-Gender tasks, while low entropy works best for AC-Age and TD tasks. The random strategy is less consistent across tasks. This discrepancy further underlines the inherent challenge in selecting ‘good’ exemplars for FSL. Our findings suggest that using no exemplars (ZSL) is generally more stable and consistent for aggregation.

5. Related Work

Generating human-annotated data is time-consuming and costly, especially for complex or specialized tasks with limited available data. Instead, a possible solution is leveraging automatic annotation models, often using a small subset of labeled data (Smit et al., 2020; Rosenthal et al., 2021), known as ‘weak supervision’ (Stanford AI Lab Blog, 2019). Supervised learning has emerged as the dominant method, driven by the widespread adoption of traditional machine learning models and Transformer-based models like BERT (Devlin et al., 2019).

More recently, LLMs have shown zero-shot and few-shot learning capabilities (Brown et al., 2020). Researchers have further advanced these models with natural language instructions (Chung et al., 2022; Wang et al., 2022; Taori et al., 2023), enabling innovative techniques like prompting (Liu et al., 2023) without the need to train a supervised model. Several papers explored these new techniques with promising performance on various NLP tasks (Brown et al., 2020; Plaza-del-Arco et al., 2022; Su et al., 2022; Sottana et al., 2023). Recent work has focused on exploring the capabil-

¹⁰We use 4,000 instances to make things as comparable as possible.

ities of LLMs as annotators. For instance, Lee et al. (2023) evaluate the performance and alignment between LLMs and humans, revealing that these models not only fall short in performing natural language inference tasks compared to humans but also struggle to capture the distribution of human disagreements accurately. Other studies have used ChatGPT as an annotator. Some report excellent capabilities (Huang et al., 2023; Gilardi et al., 2023; Törnberg, 2023; He et al., 2024), but Kuzman et al. (2023) found that ChatGPT’s performance notably drops for less-resourced languages. Similarly, Kristensen-McLachlan et al. (2023) show that on two seemingly simple binary classification tasks, the performance of ChatGPT and open-source LLMs varies significantly and often unpredictably, and supervised models systematically outperform both types of models.

For annotation, it remains to be seen whether different instruction-tuned LLMs can generalize to *any* subjective text classification tasks in different languages, especially if these tasks and languages are not well-represented in the training data. Recent studies have shown the importance of considering human label variation (Basile et al., 2021; Plank, 2022), i.e., the disagreement between human annotators, as a source of information rather than a problem. However, how this human label variation also applies to LLMs remains unexplored.

6. Discussion

Our results indicate that treating LLMs as annotators and aggregating their responses is cost-effective and quick. However, we also find that the overall performance is still well below that of even simple supervised models.

Human annotation still has a vital role if we focus on performance. As LLMs become more capable, this edge might diminish to the point where LLM annotation is equivalent to human annotation. As an aside, although all models are likely to improve across the board, we still expect specialization effects, meaning aggregation approaches will likely stay relevant for the foreseeable future.

But what about bias? Human label variation is not only due to different levels of expertise or diligence (Snow et al., 2008). It can also vary due to differing opinions, definitions, and biases (Shah et al., 2020). Specific tasks are subjective by nature (Basile et al., 2019; Rottger et al., 2022), but even seemingly objective tasks like part-of-speech tagging can have different interpretations (Plank et al., 2014). **The current discussion around the moral and ethical alignment of LLMs (Liu et al., 2022) should make us cautious about using these models as annotators in subjective or sensitive tasks.** Aggregation can overcome




Figure 2: ZSL vs. FSL Macro-F1 scores on English Trustpilot tasks. FSL sample selection strategies: Low Entropy ($\downarrow E$), Max Entropy ($\uparrow E$), and Random (Rand.). All FSL methods show much greater variance than ZSL.

the biases of any one particular model, but it cannot safeguard against widespread biases. Lastly, annotation might be exploratory (the “descriptive” paradigm in Rottger et al., 2022), where the goal is to find the range of human responses.

However, replacing human annotators with LLMs has ramifications beyond performance and bias issues. While crowdsourced annotations can be problematic regarding worker exploitation (Fort et al., 2011), they often provide low to moderate-income earners with a way to supplement their living. Replacing this option with LLMs is a real-life example of automation making human jobs obsolete. Conversely, it may mitigate the mental health risks associated with annotating toxic or sensitive content, such as racist content or tasks related to mental disorders. **Hybrid human and LLM annotation might offer a way forward here.**

7. Conclusion

We use zero- and few-shot prompting to compare four current instruction-tuned LLMs as annotators on five subjective tasks in four languages. We find specialization across models and tasks. We leverage this variance similarly to human label variation by aggregating their predictions into a single label. This approach is, on average, substantially better than any individual model. This suggests that label aggregation consistently enhances performance compared to relying on a single LLM. Despite the rapid development of LLMs enhancing generalization to new tasks, aggregation remains a beneficial strategy.

Our findings suggest that practitioners aiming to label large amounts at minimal cost (both financially and time-wise) can benefit from the outlined aggregation approach. However, we also find that even the best models cannot compete with “traditional” supervised classification approaches. Furthermore, human annotation allows practitioners to encode a specific view or approach in a prescriptive manner or explore the range of responses descrip-

tively (Rottger et al., 2022). Relying on LLMs while alignment and bias still need to be solved (Mökander et al., 2023) makes this approach unsuitable for sensitive applications.

Limitations

We were unable to compare to closed LLMs like GPT-4. While they are often state-of-the-art on many tasks, their training and setup change frequently and are, therefore, not replicable. Their pay-by-use nature also makes them less affordable for many practitioners than free open models. We do suspect, though, that including closed or generally better models will not change the overall conclusions of this paper.

Ethical Considerations

The data we use for AC-gender classification only makes a binary distinction (the Trustpilot website allowed users only to choose from two options). We do not assume this to be representative of gender identities and only use this data to test our hypotheses.

The languages we evaluate all come from the Indo-European branch of languages. The selection was due to data availability and our knowledge of languages. While we do not expect results to systematically differ from other languages, we do note that this is conjecture.

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 949944, INTEGRATOR). The research was made possible in part through an unrestricted Google research gift to explore variance in annotation. Flor Miriam Plaza-del-Arco, Debora Nozza, and Dirk Hovy are members

of the MilaNLP group and the Data and Marketing Insights Unit of the Bocconi Institute for Data Science and Analysis.

8. Bibliographical References

- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. [SemEval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 54–63, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Valerio Basile, Michael Fell, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, Massimo Poesio, and Alexandra Uma. 2021. [We need to consider disagreement in evaluation](#). In *Proceedings of the 1st Workshop on Benchmarking: Past, Present and Future*, pages 15–21, Online. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *arXiv preprint arXiv:2210.11416*.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20:37 – 46.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Joseph L. Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76:378–382.
- Karën Fort, Gilles Adda, and K. Bretonnel Cohen. 2011. [Last words: Amazon Mechanical Turk: Gold mine or coal mine?](#) *Computational Linguistics*, 37(2):413–420.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. ChatGPT outperforms crowdworkers for text-annotation tasks. *arXiv preprint arXiv:2303.15056*.
- Zeyu He, Chieh-Yang Huang, Chien-Kuang Cornelia Ding, Shaurya Rohatgi, and Ting-Hao'Kenneth' Huang. 2024. If in a crowdsourced data annotation pipeline, a gpt-4. *arXiv preprint arXiv:2402.16795*.
- Dirk Hovy. 2015. [Demographic factors improve classification performance](#). 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 752–762, Beijing, China. Association for Computational Linguistics.
- Dirk Hovy, Taylor Berg-Kirkpatrick, Ashish Vaswani, and Eduard Hovy. 2013. [Learning whom to trust with MACE](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1120–1130, Atlanta, Georgia. Association for Computational Linguistics.
- Dirk Hovy, Anders Johannsen, and Anders Søgaard. 2015. [User review sites as a resource for large-scale sociolinguistic studies](#). In *Proceedings of the 24th International Conference on World Wide Web, WWW '15*, page 452–461. International World Wide Web Conferences Steering Committee.
- Fan Huang, Haewoon Kwak, and Jisun An. 2023. [Is ChatGPT better than Human Annotators? Potential and Limitations of ChatGPT in Explaining Implicit Hate Speech](#). In *Companion Proceedings of the ACM Web Conference 2023*. ACM.

- Chia-Chien Hung, Anne Lauscher, Dirk Hovy, Simone Paolo Ponzetto, and Goran Glavaš. 2023. [Can demographic factors improve text classification? revisiting demographic adaptation in the age of transformers.](#) In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1565–1580, Dubrovnik, Croatia. Association for Computational Linguistics.
- Jan-Christoph Klie, Ji-Ung Lee, Kevin Stowe, Gözde Şahin, Nafise Sadat Moosavi, Luke Bates, Dominic Petrak, Richard Eckart De Castilho, and Iryna Gurevych. 2023. [Lessons learned from a citizen science project for natural language processing.](#) In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3594–3608, Dubrovnik, Croatia. Association for Computational Linguistics.
- Klaus Krippendorff. 2011. Computing krippendorff’s alpha-reliability. *Computing*, 1:25–2011.
- Ross Deans Kristensen-McLachlan, Miceal Canavan, Márton Kardos, Mia Jacobsen, and Lene Aarøe. 2023. [Chatbots are not reliable text annotators.](#)
- Taja Kuzman, Igor Mozetič, and Nikola Ljubešić. 2023. [ChatGPT: Beginning of an End of Manual Linguistic Data Annotation? Use Case of Automatic Genre Identification.](#) *arXiv preprint arXiv:2303.03953*.
- Noah Lee, Na Min An, and James Thorne. 2023. [Can large language models infer and disagree like humans?](#) *arXiv preprint arXiv:2305.13788*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing.](#) *ACM Computing Surveys*, 55(9):1–35.
- Ruibao Liu, Ge Zhang, Xinyu Feng, and Soroush Vosoughi. 2022. [Aligning generative language models with human values.](#) In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 241–252, Seattle, United States. Association for Computational Linguistics.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, et al. 2023. [The flan collection: Designing data and methods for effective instruction tuning.](#) *arXiv preprint arXiv:2301.13688*.
- Jakob Mökander, Jonas Schuett, Hannah Rose Kirk, and Luciano Floridi. 2023. [Auditing large language models: a three-layered approach.](#) *AI and Ethics*, pages 1–31.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. [MTEB: Massive text embedding benchmark.](#) In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037, Dubrovnik, Croatia. Association for Computational Linguistics.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2022. [Crosslingual generalization through multi-task finetuning.](#) *arXiv preprint arXiv:2211.01786*.
- Debora Nozza. 2021. [Exposing the limits of zero-shot cross-lingual hate speech detection.](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 907–914, Online. Association for Computational Linguistics.
- Rebecca J. Passonneau and Bob Carpenter. 2014. [The benefits of a model of annotation.](#) *Transactions of the Association for Computational Linguistics*, 2:311–326.
- Silviu Paun, Bob Carpenter, Jon Chamberlain, Dirk Hovy, Udo Kruschwitz, and Massimo Poesio. 2018. [Comparing Bayesian models of annotation.](#) *Transactions of the Association for Computational Linguistics*, 6:571–585.
- Barbara Plank. 2022. [The “problem” of human label variation: On ground truth in data, modeling and evaluation.](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10671–10682, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Barbara Plank, Dirk Hovy, and Anders Søgaard. 2014. [Linguistically debatable or just plain wrong?](#) In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 507–511, Baltimore, Maryland. Association for Computational Linguistics.
- Flor Miriam Plaza-del-Arco, María-Teresa Martín-Valdivia, and Roman Klinger. 2022. [Natural Language Inference Prompts for Zero-shot Emotion Classification in Text across Corpora.](#) In *Proceedings of the 29th International Conference*

- on *Computational Linguistics*, pages 6805–6817, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Flor Miriam Plaza-del-Arco, Debora Nozza, and Dirk Hovy. 2023. [Respectful or Toxic? Using Zero-Shot Learning with Language Models to Detect Hate Speech](#). In *The 7th Workshop on Online Abuse and Harms (WOAH)*, pages 60–68, Toronto, Canada. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *The Journal of Machine Learning Research*, 21(1).
- Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Marcos Zampieri, and Preslav Nakov. 2021. [SOLID: A large-scale semi-supervised dataset for offensive language identification](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 915–928, Online. Association for Computational Linguistics.
- Paul Rottger, Bertie Vidgen, Dirk Hovy, and Janet Pierrehumbert. 2022. [Two contrasting data annotation paradigms for subjective NLP tasks](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 175–190, Seattle, United States. Association for Computational Linguistics.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Tali Bers, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. [Multitask prompted training enables zero-shot task generalization](#). *arXiv preprint arXiv:2110.08207*.
- Deven Santosh Shah, H. Andrew Schwartz, and Dirk Hovy. 2020. [Predictive biases in natural language processing models: A conceptual framework and overview](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5248–5264, Online. Association for Computational Linguistics.
- Akshay Smit, Saahil Jain, Pranav Rajpurkar, Anuj Pareek, Andrew Ng, and Matthew Lungren. 2020. [Combining automatic labelers and expert annotations for accurate radiology report labeling using BERT](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1500–1519, Online. Association for Computational Linguistics.
- Rion Snow, Brendan O’Connor, Daniel Jurafsky, and Andrew Ng. 2008. [Cheap and fast – but is it good? evaluating non-expert annotations for natural language tasks](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 254–263, Honolulu, Hawaii. Association for Computational Linguistics.
- Andrea Sottana, Bin Liang, Kai Zou, and Zheng Yuan. 2023. [Evaluation metrics in the era of GPT-4: Reliably evaluating large language models on sequence to sequence tasks](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8776–8788, Singapore. Association for Computational Linguistics.
- Stanford AI Lab Blog. 2019. [Weak supervision](http://ai.stanford.edu/blog/weak-supervision/). <http://ai.stanford.edu/blog/weak-supervision/>.
- Hongjin Su, Jungo Kasai, Chen Henry Wu, Weijia Shi, Tianlu Wang, Jiayi Xin, Rui Zhang, Mari Ostendorf, Luke Zettlemoyer, Noah A Smith, and Tao Yu. 2022. [Selective annotation makes language models better few-shot learners](#). *arXiv preprint arXiv:2209.01975*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Stanford Alpaca: An Instruction-following LLaMA model](https://github.com/tatsu-lab/stanford_alpaca). https://github.com/tatsu-lab/stanford_alpaca.
- Yi Tay, Mostafa Dehghani, Vinh Q. Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Siamak Shakeri, Dara Bahri, Tal Schuster, Huaixiu Steven Zheng, Denny Zhou, Neil Houlsby, and Donald Metzler. 2023. [UI2: Unifying language learning paradigms](#). *arXiv preprint arXiv:2205.05131*.
- Petter Törnberg. 2023. [ChatGPT-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning](#). *arXiv preprint arXiv:2304.06588*.

Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, and Hugo Larochelle. 2020. [Meta-dataset: A dataset of datasets for learning to learn from few examples](#). In *International Conference on Learning Representations*.

Yizhong Wang, Swaroop Mishra, Pegah Alipoor-molabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krima Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. 2022. [Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5085–5109, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent abilities of large language models](#). *arXiv preprint arXiv:2206.07682*.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *arXiv preprint arXiv:2303.18223*.