

# The Effectiveness of LLMs as Annotators: A Comparative Overview and Empirical Analysis of Direct Representation

Maja Pavlovic<sup>1</sup> and Massimo Poesio<sup>1,2</sup>

<sup>1</sup>Queen Mary University of London, <sup>2</sup>University of Utrecht  
m.pavlovic@qmul.ac.uk, m.poesio@{qmul.ac.uk,uu.nl}

## Abstract

Large Language Models (LLMs) have emerged as powerful support tools across various natural language tasks and a range of application domains. Recent studies focus on exploring their capabilities for data annotation. This paper provides a comparative overview of twelve studies investigating the potential of LLMs in labelling data. While the models demonstrate promising cost and time-saving benefits, there exist considerable limitations, such as representativeness, bias, sensitivity to prompt variations and English language preference. Leveraging insights from these studies, our empirical analysis further examines the alignment between human and GPT-generated opinion distributions across four subjective datasets. In contrast to the studies examining representation, our methodology directly obtains the opinion distribution from GPT. Our analysis thereby supports the minority of studies that are considering diverse perspectives when evaluating data annotation tasks and highlights the need for further research in this direction.

**Keywords:** large language model (llm), annotation/labelling, representation

## 1. Introduction

Large Language Models (LLMs) have shown impressive abilities in a variety of natural language related tasks (Brown et al., 2020; Touvron et al., 2023). Brown et al. (2020) demonstrate their ability as few-shot learners and Wei et al. (2022); Kojima et al. (2022) evidence their zero-shot capabilities. Recognising the significance and costliness of annotated data across various research domains, recent work explores the potential of LLMs as data annotators, encompassing both zero- and few-shot learning approaches (Lee et al., 2023; Ziems et al., 2024; Törnberg, 2023; Zhu et al., 2023; Gilardi et al., 2023; Mohta et al., 2023; Ding et al., 2023; He et al., 2023). Considering that LLMs are trained to adhere to instructions guided by human preference (Ouyang et al., 2022; Rafailov et al., 2023), studies examine the extent to which human disagreement is captured (Lee et al., 2023) and whether or not such disagreement aligns with that of humans (Santurkar et al., 2023).

Our work, firstly, offers a comparative overview of twelve previous studies that investigate the capabilities of LLMs as annotators, concentrating on classification tasks and considering whether disagreement is captured by the studies. Secondly, we present an empirical analysis concentrating more specifically on the perspectivist question. We compare the top-performing LLM from the first section (GPT) against human annotators, by examining the degree of alignment between their opinion distributions, for the case of the four subjective datasets recently used for the 2023 SEMEVAL Task on Learning With Disagreement (Leonardelli et al., 2023).

## 2. Comparative Overview

Labelled data forms the foundation for training supervised models across diverse machine learning tasks. Much recent research has focused on exploring the use of LLMs as a quicker and more cost-effective alternative to traditional data annotation. In this first Section we review the research in this area. Due to rapid developments in this space, we concentrate on works from the past year which leverage recent models with a focus on classification tasks. Our approach to selecting relevant papers followed a combination of keyword searches, monitoring relevant workshops and conferences, and examining citations.

**Studies:** Wang et al. (2021) employ GPT-3 for the annotation of datasets, which are subsequently used in the training of smaller models. Huang et al. (2023) explore the capability of ChatGPT to accurately label implicit hate speech and provide good explanations for its annotations. Zhu et al. (2023) also investigate the capability of GPT for labelling and He et al. (2023) introduce a two step approach in which they first prompt the LLM to generate explanations and then annotate a sample to improve the annotation quality of LLMs. Both Törnberg (2023); Gilardi et al. (2023) contrast the performance of GPT with that of crowd-workers. Whereas, Goel et al. (2023) introduce a two-stage semi-automated approach employing LLMs and human experts to accelerate annotation for the extraction of medical information. Ziems et al. (2024) conduct a large scale empirical analysis to understand the zero-shot performance of GPT and Flan on 25 computational social science (CSS) benchmarks.

Paper	model families	# of model versions	# of data-sets	# of metrics	Zero/few shot	Lang.	Dis-agree.	LLM as Anno.
(Lee et al., 2023)	GPT,Vicuna, Flan,OPT-IML	9	6	4	z&f	en	✓	✗
(Santurkar et al., 2023)	GPT,Jurassic	9	1	3	z	en	✓	✗
(Ziems et al., 2024)	GPT,Flan	14	20	2	z&f	en	✗	✓
(Zhu et al., 2023)	GPT	1	5	5	z&f	en	✗	(✓)
(Gilardi et al., 2023)	GPT	1	4	2	z	en+	✗	✓
(Törnberg, 2023)	GPT	1	1	3	z	en	✗	✓
(Mohta et al., 2023)	Vicuna, Flan,Llama	9	5	3	z	en,fr,nl	✗	✗
(Ding et al., 2023)	GPT	1	4	4	z&f	en+	✗	✓
(He et al., 2023)	GPT	1	3	1	z&f	en	✗	✓
(Huang et al., 2023)	GPT	1	1	2	z	en	✗	✓
(Goel et al., 2023)	Palm	1	1	3	f	en	✗	✓
(Wang et al., 2021)	GPT	1	9	2	f	en	✗	✓

Table 1: Overview on LLM’s as Annotators (Language codes follow ISO 639, en+: predominantly English, with some additional language explorations)

**Language:** The majority of these studies measure LLM performance on English corpora (see Table 1). However, Ding et al. (2023) conduct tests to understand the possibility of using GPT on non-English corpora and Mohta et al. (2023) investigate the performance of open source LLMs on French, Dutch and English natural language inference (NLI) tasks. Thus far, models have shown better performance on English related tasks and performed notably poorly on low-resource languages Srivastava et al. (2023). While Ding et al. (2023) see potential for GPT on languages other than English, Mohta et al. (2023) observe a considerable decline in performance with non-English languages.

**Annotator Disagreement:** All studies referenced thus far assume the existence of a singular ground truth label for a given sample. There has, however, been a shift in thinking across machine learning towards a collectivist approach, meaning the inclusion of all annotator perspectives rather than having a majority voted ground truth (Uma et al., 2021; Prabhakaran et al., 2021; Cabitza et al., 2023; Rottger et al., 2022; Nie et al., 2020; Pavlick and Kwiatkowski, 2019). In this context, Lee et al. (2023) explore whether LLMs can capture the human opinion distribution. Additionally, Santurkar et al. (2023) investigate the alignment between LLMs and human annotators with respect to the

opinions and perspectives reflected in response to subjective questions. From Table (1) we can see that the latter two studies which investigate the performance of LLMs on opinion distributions don’t yet deem them ready as annotators. However, all studies that investigate the capabilities of GPT as an annotator within the traditional framework of majority voted labels agree with varying degrees that LLMs have the potential to disrupt the annotation process. Within this paradigm of majority voting, the sole exception to the consensus is expressed by Mohta et al. (2023) who conclude that LLMs have not yet attained a sufficient level for the annotation of datasets. Notably, amongst the cited studies, they are the sole study to only use open source LLMs and not consider best performing closed source alternatives (see Table 1).

**Models:** As mentioned in the previous paragraph, the predominant focus across all studies lies on models belonging to the GPT series. The remaining models under consideration are mostly open-source options, with Flan being the second most investigated, succeeded by Vicuna. Table 1 highlights that only four studies explored model families beyond GPT. Notably, these same studies explored multiple versions of a given model (“# of model versions”). In contrast, the remaining studies exclusively assessed a singular model. More details

on the exact versions can be found in table 7 (Appendix A).

**Temperature Parameter:** Not all studies mention the settings of their temperature parameter. However, both [Törnberg \(2023\)](#); [Gilardi et al. \(2023\)](#) investigate the variability in responses by experimenting with lower (0.2) and high (1.0) temperature settings. They find that LLMs have higher consistency with lower temperatures without sacrifices in accuracy and thus recommend lower values for annotation tasks. [Ziems et al. \(2024\)](#) and [Goel et al. \(2023\)](#) opt for a temperature of 0 throughout their study, aiming to ensure consistent and reproducible results across their LLM analysis.

**Prompting:** [Wang et al. \(2021\)](#) and [Goel et al. \(2023\)](#) investigate the efficacy of LLMs as annotators using only few-shot prompting. In contrast, five of the subsequent studies experiment with both zero- and few-shot prompting. Additionally, five other studies employ zero-shot prompting for their annotation tasks (see Table 1). The outcomes of the experiments comparing zero-shot and few-shot prompting show inconsistency. [Mohta et al. \(2023\)](#) experience superior performance using few-shot prompting, while [Ding et al. \(2023\)](#) find that few-shot prompting does not yield superior results across all their approaches. [He et al. \(2023\)](#) report a decrease in performance with few-shot prompting for their specific task. [Ziems et al. \(2024\)](#) conclude that improvements from few-shot prompting are inconsistent across their experiments, suggesting that achieving more substantial gains would require increased efforts in refining the prompting process.

Paper	Accuracy	F1	Precision	Recall	Reliability	Other
<a href="#">(Lee et al., 2023)</a>	✓	-	-	-	-	✓
<a href="#">(Santurkar et al., 2023)</a>	-	-	-	-	-	✓
<a href="#">(Ziems et al., 2024)</a>	-	✓	-	-	✓	-
<a href="#">(Zhu et al., 2023)</a>	✓	✓	✓	✓	-	✓
<a href="#">(Gilardi et al., 2023)</a>	✓	-	-	-	✓	-
<a href="#">(Törnberg, 2023)</a>	✓	-	-	-	✓	✓
<a href="#">(Mohta et al., 2023)</a>	✓	✓	-	-	-	✓
<a href="#">(Ding et al., 2023)</a>	✓	✓	✓	✓	-	-
<a href="#">(He et al., 2023)</a>	✓	-	-	-	-	-
<a href="#">(Huang et al., 2023)</a>	✓	-	-	-	-	✓
<a href="#">(Goel et al., 2023)</a>	-	✓	✓	✓	-	✓
<a href="#">(Wang et al., 2021)</a>	✓	-	-	-	-	✓

Table 2: Evaluation Metrics across Papers

**Evaluation:** Nearly all studies assess their outcomes using metrics such as accuracy or F1. [Santurkar et al. \(2023\)](#) deviate from these conventional performance metrics as, their primary focus lies in evaluating representation. This emphasis leads them to assess LLMs responses based on metrics measuring representativeness, steerability, and consistency ([Santurkar et al., 2023](#)). In addition to accuracy and F1, three studies utilise metrics such as precision and recall, while three other studies employ different reliability measures to evaluate inter-coder agreement. [Törnberg \(2023\)](#); [Santurkar et al. \(2023\)](#) specifically investigate model bias, whereas [Huang et al. \(2023\)](#) evaluate the natural language explanations (NLE) that LLMs can provide for their predictions. For the evaluation of LLM and human opinion distributions, [Lee et al. \(2023\)](#) use entropy, Jensen-Shannon divergence (JSD), and the Human Distribution Calibration Error (DistCE) introduced by [Baan et al. \(2022\)](#). Two studies have conducted error analyses. [Huang et al. \(2023\)](#) observe that the instances of disagreement, comprising 20% in their study, align more closely with lay-people’s perspectives. Similarly, [Ziems et al. \(2024\)](#) conclude that in their error analysis, the LLM tends to default to more common label stereotypes. Given the reported accuracy-based performance of LLMs on labelling tasks, it is important to broaden metrics to include more representational measures. For example, [Ziems et al. \(2024\)](#) omit measuring bias in their study, concluding that larger, instruction-tuned models demonstrate superior performance. However, [Srivastava et al. \(2023\)](#) caution that larger models tend to amplify bias.

## 2.1. Benefits

[Törnberg \(2023\)](#) finds that gpt-4 consistently surpasses the performance of both crowd-workers and expert coders, and the cost associated with labeling a sample is orders of magnitude lower for LLMs compared to humans. [Wang et al. \(2021\)](#) provide a detailed explanation that, in their experiments, utilising labels generated by the LLM resulted in a cost reduction ranging from 50% to 96%, while maintaining equivalent performance in downstream models. Similarly, [Goel et al. \(2023\)](#) determine that the LLM reduces the total time of labelling by 58% while maintaining a comparable baseline performance to medically trained annotators. [Gilardi et al. \(2023\)](#) demonstrate that the LLM shows superior quality compared to annotations obtained through Amazon Mechanical Turk (MTurk), while being approximately 30 times more cost-effective. [Ding et al. \(2023\)](#) find that their approach attains nearly equivalent performance when labeling the same number of samples. However, when they double the amount of data labeled by the LLM, superior performance is achieved at only 10% of the

cost associated with human annotation (Ding et al., 2023). LLMs not only entail lower costs than human annotators but also demonstrate significantly higher speeds in the labeling process (Törnberg, 2023; Wang et al., 2021; Ding et al., 2023).

In addition to diminished cost and time requirements, LLMs demonstrate the capability to provide explanations for their annotation (Mohta et al., 2023). Huang et al. (2023) find that ChatGPT generates explanations comparable, if not superior in clarity, to those produced by human annotators.

## 2.2. Limitations

As mentioned in Section 2, one limitation lies in the predominant development and testing of LLMs within the confines of the English language. An additional constraint associated with using LLMs as annotators is the challenge in formulating prompts and obtaining meaningful responses. Models might generate unconstrained responses (Goel et al., 2023) or might refrain from providing responses altogether as a result of the implementation of safeguarding measures. Ziems et al. (2024) observed that models tended to predict beyond the presented labels and exhibited a tendency to abstain from responding to tasks deemed offensive. In the event that a model does provide a response, potential issues may arise in the form of bias. Srivastava et al. (2023) show that bias in LLMs increases in with scale and ambiguous contexts. Santurkar et al. (2023) identify that LLMs demonstrate a singular perspective characterised by left-leaning tendencies. Törnberg (2023) notes the absence of substantial disparities between expert annotators and LLMs, while underscoring the notable bias observed among annotators from MTurk. However, Goel et al. (2023) underscore the importance of expert human annotators in attaining high-quality labels. Lee et al. (2023) express concerns regarding the population representation capabilities of current LLMs, whereas Ziems et al. (2024) caution researchers to consider and mitigate the potential risks of bias in their applications through human-in-the-loop methods.

An additional noteworthy limitation in employing LLMs as annotators is their sensitivity to minor alterations in prompting (Loya et al., 2023; Sclar et al., 2024). Both Huang et al. (2023) and Ziems et al. (2024) assert the need for further research to comprehensively investigate the effects of prompting and determine optimal strategies for effective prompting. Lastly, it is important to note that these models show sub-optimal performance as annotators in tasks such as NLI, implicit hate classification, empathy or dialect detection (Lee et al., 2023; Ziems et al., 2024).

## 3. Results with the SEMEVAL 2023 Subjective Tasks Benchmark

As discussed above, most studies of LLMs as annotators still adopt a majority vote perspective, which is becoming increasingly questionable particularly for subjective tasks (Akhtar et al., 2021; Leonardelli et al., 2021; Uma et al., 2021; Plank, 2022; Cabitza et al., 2023). We decided therefore to carry out a preliminary exploration of the alignment between LLM and human judgment distributions on the datasets used in the recent SEMEVAL 2023 Shared Task on Learning with Disagreement (Leonardelli et al., 2023). Our analysis is centered on the extent to which the most frequently used model (GPT) matches human distribution on datasets for inherently subjective tasks. This was done by extracting opinion distributions in the simplest and most straightforward manner possible: we directly prompt GPT to provide its estimation of the human opinion distribution and compare it against the baseline and optimal results from SemEval-2023.

Dataset	Task	Lang.	# items train dev test	% full agree.
MD-Agree.	Offensiveness detection	en	6592 1104 3057	42%
HS-Brexit	Offensiveness detection	en	784 168 168	69%
ConvAbuse	Abusiveness detection	en	2398 812 840	86%
ArMIS	Misogyny and sexism detection	ar	657 141 145	65%

Table 3: Dataset statistics (Leonardelli et al., 2023) (Language codes follow ISO 639)

### 3.1. Datasets

We leverage four datasets from SemEval2023 on "Learning with Disagreements" for the empirical analysis. All four datasets focus on subjective tasks and contain human annotated target distributions that we compare to the LLM predictions. Table 3 contains key statistics on the datasets (Leonardelli et al., 2023).

**Multi-Domain Agreement:** MD-Agreement (Leonardelli et al., 2021) is the dataset with the lowest amount of annotator agreement amongst these subjective tasks. Each example was labelled by 5 annotators and was created using English tweets from three domains (BLM, Election2020 and Covid-19).

**Hate Speech on Brexit:** HS-Brexit (Akhtar et al., 2021) was constructed from English tweets using keywords related to immigration and Brexit. Each example was labelled by 6 annotators with 69% of items having total annotator agreement.

**ConvAbuse:** ConvAbuse (Cercas Curry et al., 2021) consists of English conversational text collected from dialogue between users and two conversational AI systems. Each example was labelled by between 3 and 8 annotators. 86% of items have total annotator agreement.

**Arabic Misogyny and Sexism:** ArMIS (Almanea and Poesio, 2022) is the only non-English language task and serves to study the effect on sexism judgements particularly with respect to the annotators leanings towards conservatism or liberalism. Each example was labelled by 3 annotators with 65% of items having total annotator agreement.

### 3.2. Experimental Parameters

We explore the capability of `gpt-3.5-turbo` to generate opinion distributions for the test data of each SemEval2023 task. Given the sensitivity of LLMs to minor changes in input (Loya et al., 2023; Sclar et al., 2024), we maintain a uniform prompt structure across various tasks and let the LLM assume the role of an expert annotator who considers multiple worldviews and cultural nuances. Modifications are made only on the words related to the respective task under consideration. For instance, in the case of HS-Brexit, the LLM specialises in "hate speech detection," whereas in the ConvAbuse dataset its specialisation lies in "abusiveness detection." ArMIS is approached with slight variation due to the presence of Arabic text. In this instance, we explore two approaches: one involves prompting the models in English and providing them with the Arabic text that requires labelling, while the second approach uses an Arabic prompt (a translated version of the English prompt).

As mentioned in Section 2 there is some variability both among and within studies regarding the preferred prompting approach for LLM annotation. However, given that the multiple studies indicate limited benefits from few-shot prompting, we opt for zero-shot prompting in our tasks. The expectation of a model's output on a labelling task is to be consistent. In order to achieve such consistent and reproducible results we set the temperature parameter across our models to zero such as Ziems et al. (2024). Gilardi et al. (2023) suggest that a lower temperature value might be preferable for annotation task as it increases consistency without decreasing accuracy across their empirical analysis.

### 3.3. Evaluation Metrics

We compare the performance of GPT to both the Semeval2023 baseline model as well as the top-performing model on each task. Leonardelli et al. (2023) evaluate point predictions using the F1 measure (1) and distribution similarity using Cross-Entropy (CE) (2). To ensure comparability we use both of these in our analysis.

$$F1 = \frac{2 * TP}{2 * TP + FP + FN} \quad (1)$$

$$CE(y_n, \hat{y}_n) = - \sum_{n=1}^N y_n \log(\hat{y}_n), \quad (2)$$

where  $y_n$  is a sample opinion distribution annotated by humans and  $\hat{y}_n$  the LLMs predicted distribution for that sample. In addition to the above, we also use Shannon's entropy to visualise human and LLM uncertainties.

### 3.4. Results

Figures 1, 2, 3 and 4 contrast the frequency of opinion distributions of human annotators with those predicted by GPT for each SemEval task. We observe that when prompted directly for opinion distributions, the model shows a tendency towards bimodal predictions, with a notable preference for the following opinion distributions: {"0":0.2, "1":0.8} and {"0":0.8, "1":0.2}.

Another notable observation is evident in Figure 1, where we observe a bias towards assigning greater weight to the sexist class ('1') when prompting the LLM with Arabic text. In fact, when these distributions are simplified to a majority-based label, all test samples are categorised as sexist, a pattern not observed when the LLM was prompted with English text. The difference is also evident in the F1 performance (Table 4). The LLM prompted in Arabic only achieves an F1 score of 0.256, whereas prompting the LLM in English results in a score of 0.448, suggesting that LLMs perhaps understand the English prompt better than the Arabic one. The overall performance, however, remains significantly lower compared to other datasets, both in terms of F1 and CE metrics. This finding aligns with Mohta et al. (2023) who find that LLMs perform better on English datasets.

Table 4 highlights that while the simplistic baseline performance can be matched, it consistently falls short of the performance achieved by a specifically fine-tuned model on both F1 and CE scores (SE best).

A further examination of the errors when using the final majority voted labels reveals a higher tendency for false positive errors (see Table 5). This indicates that GPT is biased towards annotating samples as offensive, abusive, and misogynistic.

	MD-Agree.			HS-Brexit			ConvAbuse			ArMIS			
	gpt	SE (baseline)	SE (best)	gpt	SE (baseline)	SE (best)	gpt	SE (baseline)	SE (best)	gpt (english)	gpt (arabic)	SE (baseline)	SE (best)
$F1 \uparrow$	0.520	<b>0.534</b>	0.846	0.696	<b>0.842</b>	0.929	<b>0.902</b>	0.741	0.942	<b>0.448</b>	0.256	0.417	0.832
$CE \downarrow$	<b>3.829</b>	7.385	0.472	5.037	<b>2.715</b>	0.235	3.746	<b>3.484</b>	0.185	<b>5.828</b>	6.667	8.908	0.469

Table 4: Prompting gpt-3.5-turbo directly vs. baseline & best results from SemEval2023 (SE)

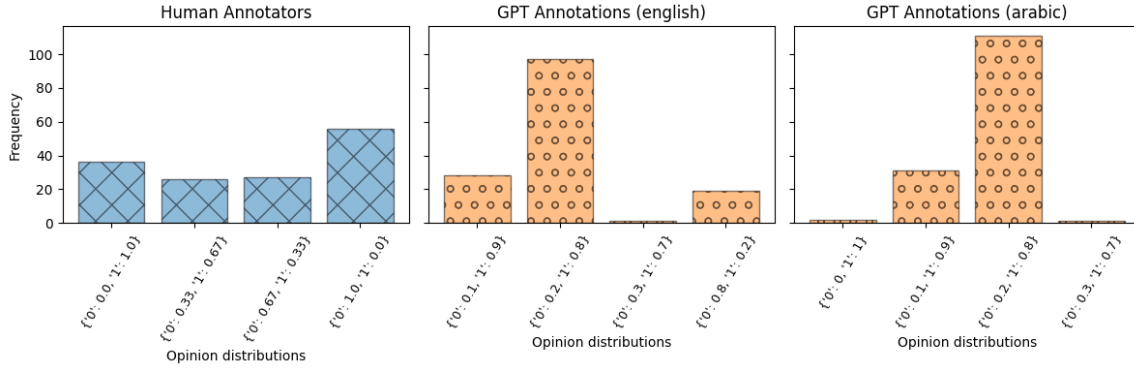


Figure 1: ArMIS opinion distributions




Figure 2: MD-Agreement opinion distributions

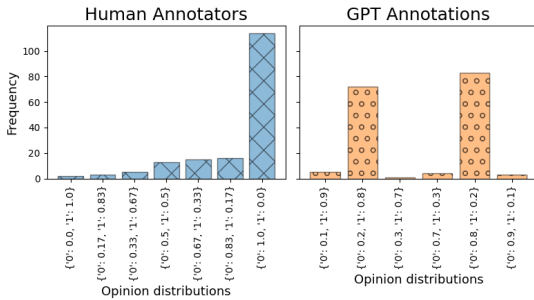


Figure 3: HS-Brexit opinion distributions




Figure 4: ConvAbuse opinion distributions

Categorisation of Errors

Dataset	FP	FN
MD-Agree	96.87%	3.13%
HS-Brexit	100.00%	0.00%
ConvAbuse	91.11%	8.89%
ArMIS (english)	95.71%	4.29%
ArMIS (arabic)	100.00%	0.00%

Table 5: Categorisation of errors into percentage that are False Positive vs. False Negative. *GPT 3.5-turbo* across different SemEval2023 tasks

Prompting the LLM to directly return opinion distributions results in higher average entropy values across all four datasets when compared to the average human entropy values (Figure 5). This stems from the observations made in the initial four figures. With the exception of the Arabic prompt, GPT consistently provides opinion distributions that allocate a small proportion to both classes rather than assigning 100 percent to one class. This leads to increased per sample entropy and thereby overall higher average entropy.

## 4. Conclusion

The overview section is not intended to provide an exhaustive review; however, the variety of tasks, datasets and approaches within the surveyed papers offers first insight into the efficacy of using LLMs to annotate data. Despite the mentioned limitations, the overall findings show a degree of consensus and positive outlook towards the use of

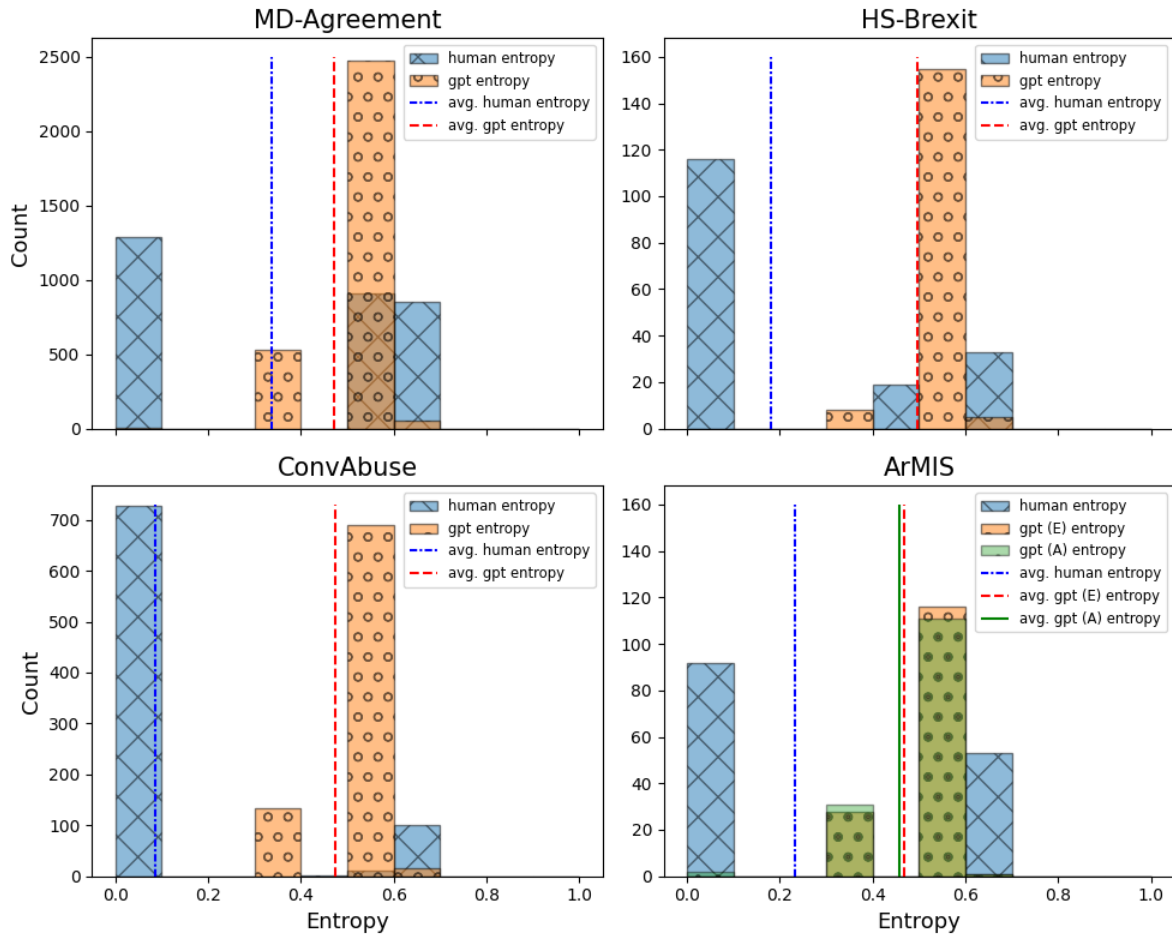


Figure 5: Histogram showing human and GPT entropy

LLMs as data annotators within the majority voting paradigm.

Our initial observations suggest that, when directly prompted, GPT tends to produce label distributions that are not strongly aligned with human opinion distributions. Furthermore, also consistent with prior research, the LLM shows superior performance on English language tasks compared to non-English text, while also showing potential bias in its responses. However, given that LLMs are trained to predict next tokens, directly obtaining opinion distributions from them has inherent limitations. Hence, in future work, we aim to explore further approaches to extracting the probability distributions such as through normalising the log probabilities (Santurkar et al., 2023) or through Monte Carlo estimation (Lee et al., 2023).

### Ethical statement

Our study exclusively used pre-existing datasets for experimentation purposes. While the datasets contain instances of offensive language, our approach involved handling this content without direct human involvement.

### Acknowledgments

Maja Pavlovic is supported by a Deep Mind PhD studentship to Queen Mary University. The work of Massimo Poesio is supported in part by the AINED Fellowship Grant *Dealing with Meaning Variation in NLP*, NGF.1607.22.002.

### 5. Bibliographical References

- Sohail Akhtar, Valerio Basile, and Viviana Patti. 2021. [Whose opinions matter? perspective-aware models to identify opinions of hate speech victims in abusive language detection](#).
- Dina Almanea and Massimo Poesio. 2022. [ArMIS - the Arabic misogyny and sexism corpus with annotator subjective disagreements](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2282–2291, Marseille, France. European Language Resources Association.
- Joris Baan, Wilker Aziz, Barbara Plank, and Raquel

- Fernandez. 2022. [Stop measuring calibration when humans disagree](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1892–1915, Abu Dhabi, United Arab Emirates. 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.
- Federico Cabitza, Andrea Campagner, and Valerio Basile. 2023. [Toward a perspectivist turn in ground truthing for predictive computing](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(6):6860–6868.
- Amanda Cercas Curry, Gavin Abercrombie, and Verena Rieser. 2021. [ConvAbuse: Data, analysis, and benchmarks for nuanced abuse detection in conversational AI](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7388–7403, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Bosheng Ding, Chengwei Qin, Linlin Liu, Yew Ken Chia, Boyang Li, Shafiq Joty, and Lidong Bing. 2023. [Is GPT-3 a Good Data Annotator?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11173–11195, Toronto, Canada. Association for Computational Linguistics.
- Esin Durmus, Karina Nyugen, Thomas I. Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, Liane Lovitt, Sam McCandlish, Orowa Sikder, Alex Tamkin, Janel Thamkul, Jared Kaplan, Jack Clark, and Deep Ganguli. 2023. [Towards measuring the representation of subjective global opinions in language models](#).
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. [ChatGPT outperforms crowd workers for text-annotation tasks](#). *Proceedings of the National Academy of Sciences*, 120(30). Publisher: Proceedings of the National Academy of Sciences.
- Akshay Goel, Almog Gueta, Omry Gilon, Chang Liu, Sofia Erell, Lan Huong Nguyen, Xiaohong Hao, Bolous Jaber, Shashir Reddy, Rupesh Kartha, Jean Steiner, Itay Laish, and Amir Feder. 2023. [LLMs Accelerate Annotation for Medical Information Extraction](#). In *Proceedings of the 3rd Machine Learning for Health Symposium*, pages 82–100. PMLR. ISSN: 2640-3498.
- Xingwei He, Zhenghao Lin, Yeyun Gong, Alex Jin, Hang Zhang, Chen Lin, Jian Jiao, Siu Ming Yiu, Nan Duan, Weizhu Chen, et al. 2023. Anollm: Making large language models to be better crowdsourced annotators. *arXiv preprint arXiv:2303.16854*.
- 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*, WWW '23 Companion, pages 294–297, New York, NY, USA. Association for Computing Machinery.
- Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213. Curran Associates, Inc.
- Noah Lee, Na Min An, and James Thorne. 2023. [Can Large Language Models Capture Dissenting Human Voices?](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4569–4585, Singapore. Association for Computational Linguistics.
- Elisa Leonardelli, Gavin Abercrombie, Dina Al-Manea, Valerio Basile, Tommaso Fornaciari, Barbara Plank, Verena Rieser, Alexandra Uma, and Massimo Poesio. 2023. [SemEval-2023 Task 11: Learning with Disagreements \(LeWiDi\)](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2304–2318, Toronto, Canada. Association for Computational Linguistics.
- Elisa Leonardelli, Stefano Menini, Alessio Palmero Aprosio, Marco Guerini, and Sara Tonelli. 2021. [Agreeing to disagree: Annotating offensive language datasets with annotators' disagreement](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10528–10539, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.



- Manikanta Loya, Divya Sinha, and Richard Futrell. 2023. [Exploring the Sensitivity of LLMs' Decision-Making Capabilities: Insights from Prompt Variations and Hyperparameters](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3711–3716, Singapore. Association for Computational Linguistics.
- Jay Mohta, Kenan Emir Ak, Yan Xu, and Mingwei Shen. 2023. [Are large language models good annotators?](#) In *NeurIPS 2023 Workshop on I Can't Believe It's Not Better (ICBINB): Failure Modes in the Age of Foundation Models*.
- Yixin Nie, Xiang Zhou, and Mohit Bansal. 2020. [What can we learn from collective human opinions on natural language inference data?](#) *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Ellie Pavlick and Tom Kwiatkowski. 2019. [Inherent disagreements in human textual inferences](#). *Transactions of the Association for Computational Linguistics*, 7:677–694.
- 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.
- Vinodkumar Prabhakaran, Aida Mostafazadeh Davani, and Mark Diaz. 2021. [On releasing annotator-level labels and information in datasets](#). In *Proceedings of the Joint 15th Linguistic Annotation Workshop (LAW) and 3rd Designing Meaning Representations (DMR) Workshop*, pages 133–138, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- 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.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. 2023. [Whose opinions do language models reflect?](#) In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *ICML'23*, pages 29971–30004. JMLR.org.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. [Quantifying language models' sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting](#). In *The Twelfth International Conference on Learning Representations*.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.
- Lindia Tjauatja, Valerie Chen, Sherry Tongshuang Wu, Ameet Talwalkar, and Graham Neubig. 2023. [Do llms exhibit human-like response biases? a case study in survey design](#).
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#).
- Petter Törnberg. 2023. [Chatgpt-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning](#).
- Alexandra N Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2021. Learning from disagreement: A survey. *Journal of Artificial Intelligence Research*, 72:1385–1470.
- Shuohang Wang, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2021. [Want To Reduce Labeling Cost? GPT-3 Can Help](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4195–4205, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.

Ruoyu Zhang, Yanzeng Li, Yongliang Ma, Ming Zhou, and Lei Zou. 2023. [LLMaAA: Making Large Language Models as Active Annotators](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13088–13103, Singapore. Association for Computational Linguistics.

Yiming Zhu, Peixian Zhang, Ehsan-Ul Haq, Pan Hui, and Gareth Tyson. 2023. [Can chatgpt reproduce human-generated labels? a study of social computing tasks](#).

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. Can large language models transform computational social science? *Computational Linguistics*, pages 1–55.

## Appendix A - Additional tables

Paper	Datasets
(Lee et al., 2023)	ANLI-R3, QNLI, ChaosNLI, PK2019
(Santurkar et al., 2023)	OpinionQA
(Ziems et al., 2024)	Indian English dialect feature detection, Twitter Emotion detection, FLUTE, Latent Hatred, Reddit/Kaggle Humor data, Ideological Books Corpus, Misinfo Reaction Frames Corpus, Random Acts of Pizza, Semeval2016 Stance Dataset, Temporal Word-in-Context benchmark, Coarse Discourse Sequence Corpus, TalkLife dataset, Winning Arguments Corpus, Wikipedia Talk Pages dataset, Conversations Gone Awry Corpus, Stanford Politeness Corpus, Hippocorpus, WikiEvents Article Bias Corpus, CMU Movie corpus dataset
(Zhu et al., 2023)	Stance Detection, Hate Speech, Sentiment Analysis, Bot Detection, Russo-Ukrainian Sentiment
(Gilardi et al., 2023)	Twitter Content moderation, US Congress, Newspaper article content moderation
(Törnberg, 2023)	Twitter Parliamentarian Database
(Mohta et al., 2023)	MM-IMDB, XNLI, Hateful memes, 2 proprietary datasets
(Ding et al., 2023)	SST2, CrossNER, FewRel, ASTEData-V2
(He et al., 2023)	QK ( <i>user query &amp; keyword relevance assessment</i> ), Word-inContext WiC, BoolQ
(Huang et al., 2023)	LatentHatred
(Goel et al., 2023)	Mimic-iv-note
(Wang et al., 2021)	XSum, Gigaword, SQuAD, SST-2, CB TREC, AGNews, DBPedia, RTE

Table 6: Datasets used across different studies

Paper	Model Versions
(Lee et al., 2023)	GPT (text-davinci-002&003); FlanT5 (large,xl,xxl), Flan UL2; Stable Vicuna; OPT-IML-M-S(1.3B)&(30B)
(Santurkar et al., 2023)	GPT(ada,davinci, text-ada-001,text-davinci-001&002&003); Jurassic (j1-Grande, j1-Jumbo, j1-Grande-v2 beta)
(Ziems et al., 2024)	GPT (text-ada-001, text-babbage-001, text-curie-001, text-davinci-001&002&003, gpt-3.5-turbo, gpt-4); FlanT5 (small, base large, xl, xxl), Flan UL2
(Zhu et al., 2023)	gpt-3.5-turbo
(Gilardi et al., 2023)	gpt-3.5-turbo
(Törnberg, 2023)	gpt-4
(Mohta et al., 2023)	Instruct-BLIP-Flan-T5; Vicuna(13b v1.3&v1.5), Vicuna (7b v1.5) InstructBLIP (7b & 13b), Llama (7b, 13b), OpenLlama (13b)
(Ding et al., 2023)	text-davinci-003
(He et al., 2023)	gpt-3.5
(Huang et al., 2023)	ChatGPT
(Goel et al., 2023)	PaLM 2
(Wang et al., 2021)	text-davinci

Table 7: Model versions used across studies

## Appendix B - Prompts

**Multi-Domain Agreement:** *You are an expert annotator in "offensiveness detection", who considers multiple worldviews and cultural nuances. Now assign probabilities for the text below, which is delimited by three dashes (-), for "offensive speech" or "not offensive speech". Return your probabilities in a json format with the key being 1 for 'offensive' and 0 for 'not offensive'. You obviously know that probabilities need to sum to 1.*

**Hate Speech on Brexit:** *You are an expert annotator in "hate speech detection", who considers multiple worldviews and cultural nuances. Now assign probabilities for the text below, which is delimited by three dashes (-), for "hateful speech" or "no hateful speech". Return your probabilities in a json format with the key being 1 for 'hate speech' and 0 for 'no hate speech'. You obviously know that probabilities need to sum to 1.*

**ConvAbuse:** *You are an expert annotator in "abusiveness detection", who considers multiple worldviews and cultural nuances. Now assign probabilities for the text below, which is delimited by three dashes (-), for "abusive speech" or "no abusive speech". Return your probabilities in a json format with the key being 1 for 'abusive speech' and 0 for 'no abusive speech'. You obviously know that probabilities need to sum to 1.*

**Arabic Misogyny and Sexism:** *You are an expert annotator in "sexism detection", who considers multiple worldviews and cultural nuances. Now assign probabilities for the text below, which is delimited by three dashes (-), for "sexist speech" or "no sexist speech". Return your probabilities in a json format with the key being 1 for 'sexist speech' and 0 for 'no sexist speech'. You obviously know that probabilities need to sum to 1.*

Arabic prompt:

"أنت محلل خبير في كشف التحيز الجنسي، وتأخذ في اعتبارك العديد من وجهات النظر العالمية والفروقات الثقافية" (-). الآن قم بتحديد احتمالات للنص أدناه، الذي يتم تحديده بثلاثة شروط "الخطاب التحيز الجنسي" أو "خطاب خالي من التحيز الجنسي". "مع المفتاح 1 لـ'خطاب التحيز الجنسي' و0 لـ'خطاب خالي من التحيز الجنسي' JSON قم بإعادة احتمالاتك بتنسيق "من الواضح أنك تعلم أن مجموع الاحتمالات يجب أن يكون 1"