

Leveraging Annotator Disagreement for Text Classification

Jin Xu

University of Twente
jinxu130817@gmail.com

Mariët Theune

University of Twente
Human-Media Interaction
m.theune@utwente.nl

Daniel Braun

University of Twente
Industrial Engineering and
Business Information Systems
d.braun@utwente.nl

Abstract

It is common practice in text classification to only use one majority label for model training even if a dataset has been annotated by multiple annotators. Doing so can remove valuable nuances and diverse perspectives inherent in the annotators' assessments. This paper proposes and compares three different strategies to leverage annotator disagreement for text classification: a probability-based multi-label method, an ensemble system, and instruction tuning. All three approaches are evaluated on the tasks of hate speech and abusive conversation detection, which inherently entail a high degree of subjectivity. Moreover, to evaluate the effectiveness of embracing annotation disagreements for model training, we conduct an online survey that compares the performance of the multi-label model against a baseline model, which is trained with the majority label. The results show that in hate speech detection, the multi-label method outperforms the other two approaches, while in abusive conversation detection, instruction tuning achieves the best performance. The results of the survey also show that the outputs from the multi-label models are considered a better representation of the texts than the single-label model.

1 Introduction

Employing multiple annotators for data annotation and afterwards using the majority annotation for model training is a widely adopted practice to mitigate biases and allow for error detection and correction (Sabou et al., 2014).

However, such procedures also remove genuine disagreement between annotators that can provide valuable insights, e.g. for subjective tasks like detection of hate speech, emotions, or sexism, but also for more objective tasks like legal or medical decision making. In recent years, the practice of only considering majority annotations has been increasingly criticized and many researchers have

started to advocate for better ways to deal with disagreement between annotators (Basile et al., 2021; Uma et al., 2021; Plank, 2022; Braun, 2024).

In this article, we propose three different strategies to leverage annotator disagreement during the training of text classification models: a probability-based multi-label approach, an ensemble system approach, and an instruction tuning approach. We compare these strategies against a baseline model that is trained on the majority labels derived from the multiple annotations. We choose two text classification tasks which inherently entail a high degree of subjectivity for the evaluation: hate speech detection and abuse detection in conversations. In our chosen datasets (Toraman et al., 2022; Cercas Curry et al., 2021), these two tasks exhibit different complexity and difficulty in terms of the label space: while the hate speech detection dataset contains binary labels, the abusive conversation detection dataset is not only annotated with abusive / non-abusive but also the severity of the abuse.

Our first approach tackles the tasks as a probability-based multi-label text classification problem. Instead of predicting specific labels to one instance, the model provides a probability distribution. The second approach imitates the process of annotation from multiple annotators with an ensemble system. The ensemble system consists of many sub-models, each of which is trained on different labels to capture the diverse viewpoints embedded in the annotations. Thirdly, we use instruction tuning. Specifically, we use a pre-trained generative model and inject explicit guidance into the training process to customize the model's behavior. The performance of the proposed models is compared using cross entropy. To evaluate the effectiveness of incorporating multiple labels, we also conduct an online survey. This survey aims to investigate human preferences between the outputs generated by the multi-label model and a baseline model.

The results show that on the hate speech dataset, the multi-label method outperforms the ensemble system and instruction tuning. Conversely, instruction tuning is the best-performing method on the abusive conversation dataset. Through multinomial test, the outputs from the multi-label model are considered more reasonable than those from the baseline model to characterize samples from the online survey. This proves the effectiveness of leveraging annotation disagreements for model training.

2 Related Work

2.1 Sources of Disagreement

Disagreement in annotations can originate from different sources. Natural language can be inherently complex and interpreted in multiple ways within a given context (Poesio, 2020). There are many subjective elements which may add an additional layer of intricacy to the understanding of texts, such as sentiments, opinions or nuanced expressions. Therefore, it is common that there are divergent interpretations among annotators. Furthermore, some sentences and even the definition of labels may contain vague or ambiguous statements (Russell et al., 2008), making it challenging for annotators to reach an agreement.

However, annotators themselves and their background can also have significant impact on the annotation results (Davani et al., 2022). Through post-annotation interviews, Patton et al. (2019), for example, showed that annotators who come from communities discussed in gang-related tweets are more likely to rely on their lived experiences in the process of annotating when compared to graduate student researchers. This divergence results in distinct label judgments. Luo et al. (2020) found that the political affiliation of annotators can significantly shape how they assess and annotate political stances.

2.2 Handling Disagreement

Majority voting involves aggregating annotations by selecting the label that the majority of annotators agree upon. Majority voting is easy to understand and implement and tends to perform well when the annotators share unanimous perspectives (Uma et al., 2021). However, the employment of a majority voting method in annotation processes can unintentionally obscure nuanced viewpoints, especially for groups that are underrepresented in

annotator pools (Prabhakaran et al., 2021). To address this concern, it is important to ensure a diverse representation among annotators to foster a more comprehensive understanding of various perspectives, particularly those from underrepresented demographics (Wan et al., 2023).

Some studies have introduced alternative methods to majority voting in order to incorporate annotator disagreement in model training. Chou and Lee (2019) modelled the label uncertainty and annotator idiosyncrasy simultaneously by using both hard label (majority voting) and soft label (the distribution of annotations). The results showed that the soft label contains useful information that significantly boosts the model performance. Fornciari et al. (2021) proposed a multi-task neural network that was trained on soft label distribution over annotator labels. By integrating a divergence measurement between soft label and “true” label vector into the loss functions, they effectively mitigated overfitting and therefore improved model performance. Davani et al. (2022) introduced multi-annotator models where each annotator’s judgments were regarded as independent sub-task with a shared common representation of the annotation task. This approach enables to preserve and model the internal consistency in each annotator’s label. It also incorporates the systematic disagreements with other annotators. Similarly, the network architecture introduced by Guan et al. (2018) individually models annotation experts. In this approach, each expert’s model weight is calculated independently, and these individual weights are then averaged to facilitate ensemble recognition. To include the knowledge from annotators, Fayek et al. (2016) employed neural networks to build an ensemble system that consists of many models, with each model representing one annotator. Then the final results are obtained by combining the individual model outputs.

Although the approaches outlined above have improved the performance by leveraging annotation disagreements, they remained limited to identifying the majority label. The outputs, in the form of “soft labels” (probability distribution over labels), were still aggregated to single labels as final predictions. There is limited research focusing on evaluating the effectiveness of embracing multiple labels.

3 Datasets

In this section, the two datasets that have been used in this study will be briefly introduced.

3.1 Hate Speech

The first dataset is the “Large-Scale Hate Speech Dataset”¹ published by [Toraman et al. \(2022\)](#). It consists of a total of 100,000 tweets (7,000 training, 1,500 validation, and 1,500 testing). Each tweet in the dataset is annotated by five annotators that have been selected randomly from a panel of 20 annotators. According to the annotation guidelines utilized by [Sharma et al. \(2018\)](#), tweets are categorized as “Hate” if they target, incite violence against, threaten, or advocate for physical harm towards an individual or a group of people based on identifiable trait or characteristic. If tweets humiliate, taunt, discriminate against, or insult an individual or a group of people, they are annotated as “Offensive”. In the absence of these criteria, the tweets are labeled as “Normal” .

3.2 Abuse Conversation

The second dataset is the “Abuse in Conversational AI” dataset² (hereinafter referred to as “abusive conversation dataset”) published by [Cercas Curry et al. \(2021\)](#). The data was collected from conversations between users and conversational AI systems, and consists of 2501 samples as training data, 831 as validation data and 853 as testing data. The data was annotated using an unbalanced rating scale proposed by [Poletto et al. \(2019\)](#), in which inputs are labelled on a scale from Not abusive, Ambiguous, Mildly abusive, Strongly abusive to Very strongly abusive. This annotation scheme offers insights into not only the presence of abusive content, but also the severity of the abuse. In the annotation process, eight annotators were recruited, and each example is annotated by a minimum of three annotators.

4 Methodology

4.1 Baseline model

The baseline model for this study is trained on the “ground truth” label that is aggregated via majority voting. Given BERT’s ([Devlin et al., 2019](#)) notable performance in contextual understanding, we

¹https://github.com/avaapm/hatespeech/tree/master/dataset_v1

²<https://github.com/amandacurry/convabuse/tree/main>

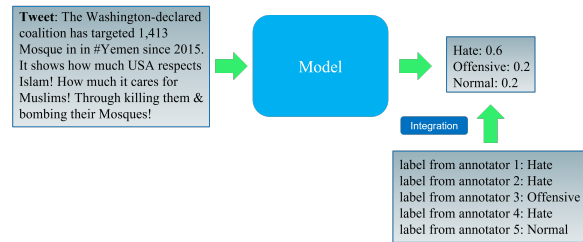


Figure 1: The framework of model training within the probability-based multi-label method.

choose it as the pre-trained model. Since the baseline model outputs a single label, we augment its architecture by adding a fully connected layer to the last hidden state, thereby adapting the model structure to the specific prediction task.

4.2 Probability-based multi-label method

The task of identifying hate speech or abusive conversation can be regarded as a multi-label text classification problem, where a given piece of text can be associated with one or multiple labels simultaneously. Unlike the traditional approaches that assign one or several exclusive labels to the input text ([Jiang and Nachum, 2020](#)), our model predicts the probability of each label being associated with the given text. The approach is illustrated in Figure 1. The model is trained on the probability distribution across different labels which is derived from individual annotations. Like the baseline model, the multi-label model also is based on BERT but fine-tuned with different types of target labels.

4.3 Ensemble system

In the annotation process, multiple labels are assigned by different annotators. Inspired by this process, we propose an ensemble system consisting of several sub-models. As shown in Figure 2, each sub-model is based on a BERT model that is fine-tuned individually on its respective set of labels. For each sub-model, the input is the text from one sample and the output is a multidimensional vector where each dimension corresponds to one category. After that, this vector is transformed by the SoftMax function and the dimension with highest probability is identified as the output of the sub-model. Finally, the predictions from all sub-models are combined and converted into a probability distribution of three- or five-dimensional vector.

In the abusive conversation dataset, the annotators assigned for each sample are clearly specified

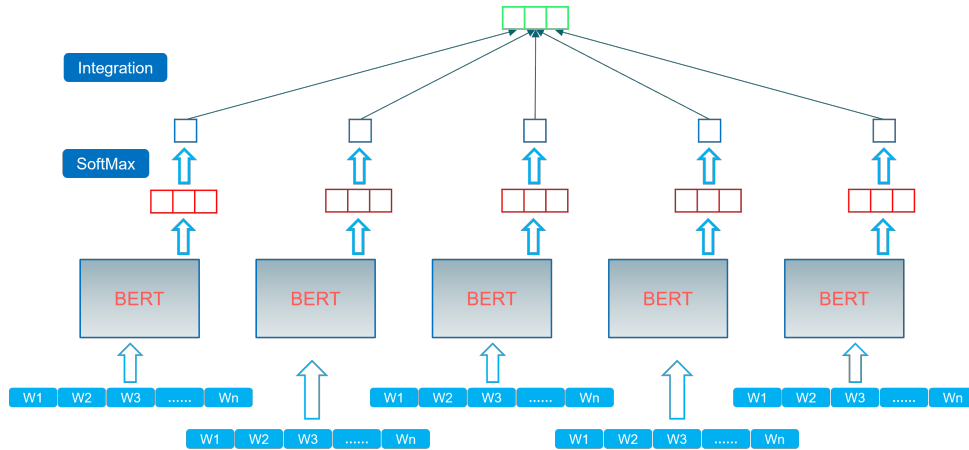


Figure 2: Fine-tuning BERT individually as sub-models within the ensemble system.

and identifiable. Therefore, within the ensemble system, each sub-model represents one specific annotator and is trained on that annotator’s labels. By contrast, in the hate speech dataset, each sample is labeled by five anonymous annotators. Despite the anonymity, training a model with such labels can potentially increase the robustness of sub-models since it helps to reduce the biases or inconsistencies introduced by individual annotators (Frenay and Verleysen, 2014). Furthermore, the resulting labels are likely to reflect a diverse range of perspectives and interpretations of the data. Training sub-models on these diverse annotations can capture the variability in annotator judgments and enhance the model’s ability to generalize across different viewpoints (Audhkhasi and Narayanan, 2013). Since the sub-models can show varying performances in the training and validation processes, typically, the top n ($n \geq 3$) best-performing sub-models are chosen to determine the final output. The ranking is based on their accuracy on the validation data.

4.4 Instruction tuning

Instruction tuning is the process of fine-tuning LLMs in a supervised fashion on a dataset consisting of pairs of instructions and outputs. The key idea is to provide the model with explicit instructions to enhance its performance and align it with specific objectives. Unlike traditional training approaches where models learn from data alone, instruction tuning injects explicit guidance into the training process. This approach allows for explicit customization of the model’s behavior. In this study, we ask the model to predict the class of hate speech or abusive conversation based on the

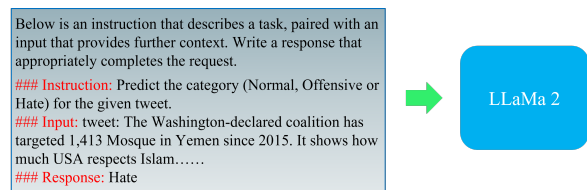


Figure 3: Fine-tuning LLaMa 2 as a sub-model with instruction tuning in the hate speech dataset.

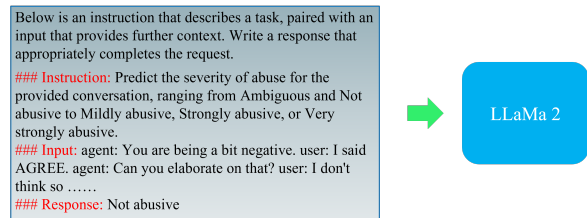


Figure 4: Fine-tuning LLaMa 2 as a sub-model with instruction tuning in the abusive conversation dataset.

input we construct. The input contains the task description, the instruction, the original text, and the annotation from one specific annotator (i.e. *not* the majority label). The approach of fine-tuning LLaMa 2 via instruction tuning on the two datasets is presented in Figure 3 and Figure 4. On the left sides of the figures are the inputs fed into the pre-trained model. The input comprises the following four components: scenario description, instruction, text input and response.

Like the ensemble system, top n ($n \geq 3$) best-performing sub-models are selected to contribute to the final predictions.

	Training	Validation	Testing
Hate speech	0.7613	0.7569	0.7638
Abusive conversation	0.8861	0.9680	0.9834

Table 1: The average cross entropy of the probability-based multi-label model on two tasks.

4.5 Evaluation

4.5.1 Cross entropy

The final output of our proposed models is a probability distribution across different labels, and in this scenario, a single “ground truth” label is no longer applicable for model evaluation. Instead, we use cross entropy to compare the distribution of annotations with model output. Cross entropy is one kind of statistical distance which measures how a probability distribution is different from a reference probability distribution. In the field of NLP, it has been used to quantify how well the model’s predicted distribution matches the annotation distribution (Pavlick and Kwiatkowski, 2019).

4.5.2 Online survey

Using cross entropy to evaluate the effectiveness of training models with multiple labels against models that only rely on the majority label is impossible due to the format disparity between the outputs generated. To bridge this gap, we conduct an online survey where participants specify their preference between annotations generated from the probability-based multi-label model and the baseline model. For each dataset, we select 10 samples, each featuring two annotations. Both annotations are in the form of probability distributions across different labels. One is generated from the baseline model trained with majority labels, which is, however, used to generate a probability distribution in the phase of inference. The other one is from the probability-based multi-label model. This model has the same structure as the baseline model and their only difference is the labels they were trained on. For each sample, participants are required to indicate which annotation they find is more reasonable to characterize the tweet or the abusive conversation.

5 Results

5.1 Multi-Label Method

Table 1 show the performance of the multi-label model on the two datasets. In this approach, the

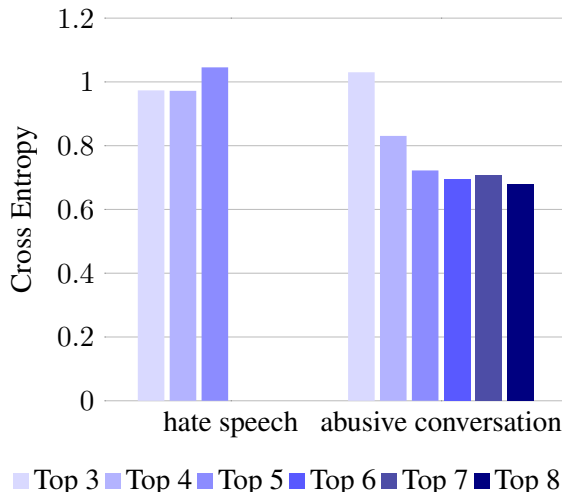


Figure 5: Comparison of the ensemble system’s performances on two tasks.

model demonstrates superior performance on the hate speech dataset compared to the abusive conversation dataset. In particular, the cross entropy for the hate speech dataset is 0.7638, while this value for the abusive conversation dataset is 0.9834. With a parameter size of 110 million, the multi-label model benefits from the extensive training data in hate speech dataset to optimize and align itself with the downstream task. In contrast, there are only 2501 training samples in the abusive conversation dataset, which can easily lead to overfitting in the process of training. The multi-label model exhibits relatively consistent losses across training, validation, and testing data in the hate speech dataset, indicating a good fit without signs of underfitting or overfitting. By comparison, in the abusive conversation dataset, losses during validation and testing are noticeably higher than during training. When the model encounters unseen data in validation and testing phases, the loss can be relatively high due to the lack of generalization.

5.2 Ensemble system

Figure 5 shows the ensemble system’s performances on the two datasets. In the testing phase, we select the top-performing sub-models based on their validation accuracies. The ensemble system performs better on the abusive conversation dataset than on the hate speech dataset. Specifically, in the hate speech dataset, the best performance is achieved by the top 3 sub-models and the corresponding overall cross entropy loss is 0.9720. Conversely, the best overall cross entropy for the abusive conversation dataset is 0.6782, achieved with




Figure 6: Comparison of instruction tuning’s performances on two tasks.

the top 8 (all) sub-models. The ensemble system is designed to simulate the process of annotation and has a large parameter size. Despite being trained on a substantially larger dataset, this method performs less effectively for the hate speech dataset. In this dataset, 20 annotators contribute, with each sample being annotated by five randomly assigned annotators, which means the five annotators for all the samples are not always the same individuals. As a result, one single sub-model may struggle to learn the specific characteristics of each annotator from the data. By contrast, in the abusive conversation dataset, there are eight annotators in total and for each sample it is clearly indicated which annotators are assigned for the annotation task. In this context, each sub-model is designed to emulate an individual annotator. Consequently, the ensemble system integrates the unique insights from each individual annotator, as represented by the sub-models.

5.3 Instruction Tuning

Figure 6 shows the performance of the instruction tuning approach. In this approach, even though with a considerably smaller training data size, the model’s performance on the abusive conversation dataset is significantly better compared to the hate speech dataset. In the hate speech dataset, the best performance is achieved by the top 3 sub-models, with a cross entropy of 1.2445. By contrast, the lowest cross entropy in the abusive conversation dataset, achieved by the top 6 sub-models, is 0.6200. Unlike traditional machine learning or deep learning algorithms, one of the most evident advantages of instruction tuning is that it does not




Figure 7: Comparison of different models’ performances on the hate speech dataset.

require much training data. Even though there are only 2501 training samples in the abusive conversation dataset, it is already sufficient to fine-tune the model and enable it to grasp the specific patterns and knowledge within the data. With this limited dataset, the pre-trained model selectively activates or deactivates certain neurons in the neural network, which serves an important role in revealing or concealing some functions embedded in LLaMa 2. Although the hate speech dataset contains a large amount of training data, the individual samples annotated by specific annotators remain unknown, which presents a challenge for the model in terms of fitting and learning patterns from the data.

5.4 Comparison

As shown in Figure 7, the multi-label method outperforms the other approaches on the hate speech dataset. The reason behind this might be the aforementioned issue in this dataset: the five annotators assigned to each sample are anonymous. Both the ensemble system and instruction tuning were trained using the same paradigm, where sub-models were fine-tuned individually on their respective labels. On the contrary, the multi-label model only relied on the probability distribution across different classes as the target, circumventing the issue with annotator anonymity. Furthermore, the hate speech dataset is big enough to fine-tune the BERT model.




Figure 8: Comparison of different models’ performances on abusive conversation dataset.

Figure 8 shows that on the abusive conversation dataset the multi-label method performs worst. The size of this dataset is relatively small, which can result in overfitting during fine-tuning. The ensemble system consists of sub-models, with each tailored to predict annotations from a specific annotator. With multiple sub-models making their own decisions independently and contributing to the final prediction, the ensemble system can mitigate the bias brought by overfitting. Since instruction tuning does not have a high requirement for dataset size, it performs slightly better than the ensemble system.

5.5 Online survey

In exploring the probability distribution preference, we recruited 36 participants for the online survey. The multinomial test (Read and Cressie, 2012) is employed since there are three possible preference options. The details of the results are outlined in Table 2. From this table, the multinomial tests for the multi-label model on two datasets are statistically significant, with the p-value of 0.0000. This means there is a notable disparity among the three categories being compared. Individuals generally favor the multi-label model as a more reasonable representation to characterize tweets or conversations. The results indicate the effectiveness of leveraging annotation disagreements in model training.

6 Conclusion

In this paper, we proposed and compared three approaches to incorporate diverse annotations in the training of ML models: a probability-based multi-

label method, an ensemble system, and instruction tuning. All three approaches take the individual labels from all annotators into account for model training in different ways, rather than only depending on an assumed “ground truth” label. In this way, the output includes a rich diversity of perspectives from annotators. We applied the proposed models on two datasets, which correspond to two tasks: hate speech detection and abuse detection in conversational AI. The two datasets show discrepancies in terms of data size, classification difficulty, the number of annotators involved in each sample, and their anonymity levels. Results show that on the hate speech dataset, the multi-label method demonstrates the highest performance among the three models, while instruction tuning achieves the lowest loss on the abusive conversation dataset. Lastly, an online survey was conducted to evaluate the performance of the probability-based multi-label model in comparison to the baseline model. The online survey investigated individuals’ preference between the distributions generated from the multi-label model and the baseline model. The evaluation of the survey results showed that the distribution generated from the multi-label model is considered more reasonable to characterize the texts compared to the baseline model. In the future, we would like to explore some methods or techniques to mitigate the class-imbalanced issue in the dataset. For example, there have been many popular algorithms that contribute to a relatively class-balanced dataset by over sampling (Chawla et al., 2002) or down sampling (Wilson, 1972). We would also like to work on investigating automatically generated prompts.

Preference	Hate speech			Abusive conversation		
	Counts	Proportion	P-value	Counts	Proportion	P-value
Baseline	118	0.3278	0.6078	152	0.4222	0.0003
Multi-label model	198	0.5500	0.0000	194	0.5389	0.0000
No difference	44	0.1222	1.0000	14	0.0389	1.0

Table 2: Multinomial test for probability distribution preference on two datasets.

Recent research has demonstrated that a concrete prompt, which consists of several discrete tokens, may not always be the most effective prompt to instruct the behavior of the model (Liu et al., 2023). Conversely, continuous embeddings of prompts, which might lack immediate human interpretability, make sense for the model itself (Li and Liang, 2021; Subramani et al., 2019).

Limitations

There are some limitations to the experiments. Firstly, the ensemble system showed to be not suitable for the hate speech dataset, where the five annotators assigned to each sample are not fixed. In this dataset, each set of annotations used for training a sub-model can comprise annotations from multiple individuals. As a result, it becomes impossible for the sub-models to capture the specific characteristics of each annotator embedded in the annotations.

Secondly, both datasets in this study suffer from class-imbalanced problem, which can have an adverse impact on model training. When trained on a class-imbalanced dataset, the model primarily focuses on the samples from the majority class and neglect those from the minority class, as that is an efficient strategy for minimizing the training loss. Another limitation is the inconsistency among annotators, which can introduce noise into the dataset and weaken model performance. Since our dataset lacked identifiable annotators, it was not possible to model individual annotator bias or assess inter-annotator agreement comprehensively. This constrains our ability to account for subjective variations in labeling.

Thirdly, we only leverage manually created prompts, which may introduce subjectivity and bias based on the prompt maker’s perspective (Tian et al., 2023). It has been proved that manually created prompts suffer from a high degree of instability and a minor change in the prompt can result in substantial discrepancies in the model’s performance (Liu et al., 2023).

References

- Kartik Audhkhasi and Shrikanth Narayanan. 2013. A Globally-Variant Locally-Constant Model for Fusion of Labels from Multiple Diverse Experts without Using Reference Labels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(4):769–783.
- Valerio Basile, Federico Cabitza, Andrea Campagner, and Michael Fell. 2021. Toward a Perspectivist Turn in Ground Truthing for Predictive Computing. *CoRR*, abs/2109.04270. ArXiv: 2109.04270.
- Daniel Braun. 2024. I beg to differ: how disagreement is handled in the annotation of legal machine learning data sets. *Artificial intelligence and law*, 32(3):839–862.
- 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.
- Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357.
- Huang-Cheng Chou and Chi-Chun Lee. 2019. Every Rating Matters: Joint Learning of Subjective Labels and Individual Annotators for Speech Emotion Classification. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5886–5890.
- Aida Mostafazadeh Davani, Mark Díaz, and Vinodkumar Prabhakaran. 2022. Dealing with Disagreements: Looking Beyond the Majority Vote in Subjective Annotations. *Transactions of the Association for Computational Linguistics*, 10:92–110.
- 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.
- H.M. Fayek, M. Lech, and L. Cavedon. 2016. Modeling subjectiveness in emotion recognition with deep

- neural networks: Ensembles vs soft labels. In *2016 International Joint Conference on Neural Networks (IJCNN)*, pages 566–570.
- Tommaso Fornaciari, Alexandra Uma, Silviu Paun, Barbara Plank, Dirk Hovy, and Massimo Poesio. 2021. [Beyond Black & White: Leveraging Annotator Disagreement via Soft-Label Multi-Task Learning](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2591–2597, Online. Association for Computational Linguistics.
- Benoit Frenay and Michel Verleysen. 2014. [Classification in the Presence of Label Noise: A Survey](#). *IEEE Transactions on Neural Networks and Learning Systems*, 25(5):845–869.
- Melody Guan, Varun Gulshan, Andrew Dai, and Geoffrey Hinton. 2018. [Who Said What: Modeling Individual Labelers Improves Classification](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1). Section: AAAI Technical Track: Machine Learning.
- Heinrich Jiang and Ofir Nachum. 2020. [Identifying and Correcting Label Bias in Machine Learning](#). In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 702–712. PMLR.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2023. [GPT understands, too](#). *AI Open*.
- Yiwei Luo, Dallas Card, and Dan Jurafsky. 2020. [DeSMOG: Detecting Stance in Media On Global Warming](#). *CoRR*, abs/2010.15149. ArXiv: 2010.15149.
- Desmond Patton, Philipp Blandfort, William Frey, Michael Gaskell, and Svebor Karaman. 2019. [Annotating Social Media Data From Vulnerable Populations: Evaluating Disagreement Between Domain Experts and Graduate Student Annotators](#). In *Proceedings of the 52nd Hawaii International Conference on System Sciences 2019 (HICSS-52)*.
- 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.
- Massimo Poesio. 2020. [Ambiguity](#). In *The Wiley Blackwell Companion to Semantics*, pages 1–38. John Wiley & Sons, Ltd. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118788516.sem098>.
- Fabio Poletto, Valerio Basile, Cristina Bosco, Viviana Patti, Marco Stranisci, and others. 2019. [Annotating hate speech: Three schemes at comparison](#). In *CEUR WORKSHOP PROCEEDINGS*, volume 2481, pages 1–8. CEUR-WS.
- 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.
- Timothy RC Read and Noel AC Cressie. 2012. [Goodness-of-fit statistics for discrete multivariate data](#). Springer Science & Business Media.
- Bryan C. Russell, Antonio Torralba, Kevin P. Murphy, and William T. Freeman. 2008. [LabelMe: A Database and Web-Based Tool for Image Annotation](#). *International Journal of Computer Vision*, 77(1):157–173.
- Marta Sabou, Kalina Bontcheva, Leon Derczynski, and Arno Scharl. 2014. [Corpus annotation through crowdsourcing: Towards best practice guidelines](#). In *LREC*, pages 859–866.
- Sanjana Sharma, Saksham Agrawal, and Manish Shrivastava. 2018. [Degree based classification of harmful speech using Twitter data](#). In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*, pages 106–112, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Nishant Subramani, Samuel Bowman, and Kyunghyun Cho. 2019. [Can Unconditional Language Models Recover Arbitrary Sentences?](#) In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Jacob-Junqi Tian, David B. Emerson, Sevil Zanjani Miyandoab, Deval A. Pandya, Laleh Seyyed-Kalantari, and Faiza Khan Khattak. 2023. [Soft-prompt Tuning for Large Language Models to Evaluate Bias](#). *ArXiv*, abs/2306.04735.
- Cagri Toraman, Furkan Şahinuç, and Eyup Yilmaz. 2022. [Large-Scale Hate Speech Detection with](#)

[Cross-Domain Transfer](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2215–2225, Marseille, France. European Language Resources Association.

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.

Ruyuan Wan, Jaehyung Kim, and Dongyeop Kang. 2023. [Everyone’s Voice Matters: Quantifying Annotation Disagreement Using Demographic Information](#). In *AAAI Conference on Artificial Intelligence*.

Dennis L. Wilson. 1972. [Asymptotic Properties of Nearest Neighbor Rules Using Edited Data](#). *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-2(3):408–421.