

iLab at SemEval-2023 Task 11 Le-Wi-Di: Modelling Disagreement or Modelling Perspectives?

Nikolas Vitsakis¹ and Amit Parekh¹ and Tanvi Dinkar¹
and Gavin Abercrombie¹ and Ioannis Konstas^{1,2} and Verena Rieser^{1,2}

¹The Interaction Lab, Heriot-Watt University ²Alana AI
nv2006@hw.ac.uk

Abstract

There are two competing approaches for modelling annotator disagreement: distributional soft-labelling approaches (which aim to capture the level of disagreement) or modelling perspectives of individual annotators or groups thereof. We adapt a multi-task architecture — which has previously shown success in modelling perspectives— to evaluate its performance on the SEMEVAL Task 11. We do so by combining both approaches, i.e. predicting individual annotator perspectives as an *interim step* towards predicting annotator disagreement. Despite its previous success, we found that a multi-task approach performed poorly on datasets which contained distinct annotator opinions, suggesting that this approach may not always be suitable when modelling perspectives. Furthermore, our results explain that while strongly perspectivist approaches might not achieve state-of-the-art performance according to evaluation metrics used by distributional approaches, our approach allows for a more nuanced understanding of individual perspectives present in the data. We argue that perspectivist approaches are preferable because they enable decision makers to amplify minority views, and that it is important to re-evaluate metrics to reflect this goal.

1 Introduction

Many Natural Language Processing (NLP) tasks follow a supervised learning paradigm, i.e. classification of labelled data where multiple annotations are aggregated into a *hard label* using averaging or majority voting. Hard labels are based on the assumption that each instance in a dataset has *one singularly correct response*—often referred to as ‘ground truth’.

However, this assumption is *highly unrealistic* for social computing tasks, such as toxicity and hate speech detection (e.g. Vidgen and Derczynski, 2021), where lived experiences, biases, opinions and annotator experience all play a

role in the *subjective* response an annotator might give. Hard labels especially disadvantage minority groups (Blodgett, 2021). For example, in abusive language classification, where a minority is disproportionately affected (such as minoritised people who have faced online harassment), an aggregated majority label can obscure the perspective of the most vulnerable groups.

Thus, there is growing awareness that modelling multiple perspectives is necessary, particularly for inherently subjective tasks and those concerned with social issues (Abercrombie et al., 2022; Cabitza et al., 2023; Plank, 2022).

Le-Wi-Di The SemEval 2023 shared task ‘Learning With Disagreements’ (Leonardelli et al., 2023) aims to capture and model annotator disagreement – going beyond the assumption of one aggregated ‘ground truth’ label. Participating teams are required to propose methods that consider *disaggregated annotations*, used to create a *soft label*, which represents the probabilistic distribution of the annotations. Soft labels can then be used to predict the *level of disagreement* for each instance in a dataset.

The task presents a benchmark of four datasets, including a hard label and soft labels for each instance. The datasets were chosen specifically to represent tasks that are highly subjective (e.g. hate speech detection) and show high annotator disagreement. Participating teams are evaluated on how well their proposed model predicts both hard and soft labels, via F1 accuracy score and cross-entropy loss, respectively. The shared task prioritised the soft evaluation, i.e. how well the model’s probabilities reflect annotator agreement, rather than simply proposing models that would outperform the current state-of-the-art for the hard labels. For further details of the datasets, see section 3.

Our approach proposes a modified version of the multi-task model introduced by Davani et al. (2022), which aims to predict individual annota-

tor judgments. By training our model to predict each annotator’s judgement for each instance in a dataset, we can use the resulting predictions to infer the level of disagreement for that instance, without directly training for such a purpose.

The main benefit of our approach is that opinions present in the dataset are preserved beyond the simplistic form of a polarised agreement/disagreement distribution. Instead, we focus on representing individual opinions, also known as ‘perspectives’ (Cabitza et al., 2023). This allows modelling of specific perspectives present in a dataset, potentially enabling the amplification of minority opinions.

2 Related work

2.1 Modelling disagreement

Uma et al. (2021) provide an extensive survey, outlining four main approaches. The first aggregates annotations into hard labels (as in Task metric 1). The second removes items that display high disagreement and is thus unsuitable for this task. The third models the distribution of annotations for each item, i.e. ‘soft labels’ (as in metric 2). The final approach enables the model to optimise across different tasks through the use of either the use of multi-task learning, or a procedure called plank-style weighing (for more information towards this method refer to Plank et al. (2014)).

For the purposes of this paper, the former is relevant, as multi-task learning, enables the resulting model to optimise across different tasks through the use of both hard and soft labels, providing predictions for the hard label, as well as degrees of confidence for each (Fornaciari et al., 2022). While these approaches make use of disagreement to enhance optimisation, they have not been used to preserve the different perspectives represented in the data.

2.2 Modelling perspectives

We focus on a ‘strong perspectivist approach’ which aims to *preserve diversity of perspectives* throughout the whole classification pipeline (Cabitza et al., 2023). To contrast, weak perspectivist approaches (Cabitza et al., 2023) may consider several annotator viewpoints, but still reduce these viewpoints towards a single label. An example of a strong perspectivist approach is Davani et al. (2022), who predict individual annotator judgments using a multi-task approach, and treat each annotator as a sub-task, thus retaining individual perspectives. While

research has shown some success when utilising single-task models to accurately capture distinct perspectives (Rodrigues and Pereira, 2018), our choice of model was informed by recent evidence that multi-task models (Fornaciari et al., 2021) can outperform single-task models.

However, one limitation (with respect to strong perspectivism) is that, for evaluation, they aggregate predicted annotations into one label, essentially falling back into the issues of hard labels. Our proposed solution aims to address this limitation through the use of both hard and soft labels, i.e. evaluating model performance on the disaggregated perspectives present in the dataset.

3 Data

The Le-Wi-Di¹ shared task consists of the following four datasets, that have all been synthesised into a common json format.

HS-Brexit (Akhtar et al., 2021): a dataset of English tweets on the topic of Brexit, labelled for hate speech by six annotators belonging to two different groups; three Muslim immigrants and three other individuals.

ArMIS (Almanea and Poesio, 2022): Arabic tweets annotated for misogyny and sexism. The three annotators for this dataset have different self-identified demographics of ‘Moderate Female’, ‘Liberal Female’ and ‘Conservative Male’.

ConvAbuse (Cercas Curry et al., 2021): a dataset of English dialogues between users and two conversational agents. The dialogues have been annotated for abusiveness by eight gender studies experts.

MultiDomain Agreement (MD): (Leonardelli et al., 2021): English tweets on three topics: BLM, Election and Covid-19, labelled for offensiveness by crowdsourced annotators and specifically selected to elicit disagreement.

See Table 1 for descriptive statistics of the datasets. Disagreement is moderate for the ArMIS and ConvAbuse datasets². While the MD dataset has (unsurprisingly) higher disagreement given the large pool of annotators (891 annotators total), the percentage of unseen annotators (i.e. annotators

¹<https://le-wi-di.github.io/>

²Note, our results for calculating inter-annotator agreement may slightly differ from the original papers (e.g. Cercas Curry et al. (2021) Krippendorff’s $\alpha \approx 0.69$) given discarded instances etc. Krippendorff’s α is calculated using the **Fast Krippendorff** package. The values range from 0 – 1, where 0 depicts perfect disagreement, and 1 complete agreement.

		HS-Brexit	ArMIS	ConvAbuse	MD
Task		Hate speech	Misogyny	Abusiveness	Offensiveness
No. of instances	train	784	657	2398	6592
	dev	168	141	812	1104
	test	168	145	840	3057
Utterance length		18.623 ± 4.578	19.510 ± 12.042	27.322 ± 18.830	22.614 ± 14.777
Annotator details	Krippendorff’s α (\downarrow)	0.347	0.524	0.650	0.359
	Total annotators	6	3	8	819
	Annotators / instance	6	3	4	5
	Unseen annotators	0	0	0	91

Table 1: Descriptive data and annotator statistics: utterance length in tokens with standard deviations, inter-annotator agreement measured with Krippendorff’s α (\downarrow) lower=higher disagreement), and ‘unseen annotators’, the percentage of annotators that are not represented by at least one instance in all of the train, dev and test sets.

that don’t appear in all splits) is extremely high (91%) compared to the other datasets (0%). These factors could lead to sparsity of strongly distinguished perspectives present in the MD dataset. In ConvAbuse, the standard deviation of utterance lengths³ is higher due to the presence of many single token responses that can be present in a dialogue—such as ‘yes’.

4 Methods

4.1 System overview

We implement a multi-task model, which makes, for each instance, separate predictions for each annotator present in the dataset. Given the varied characteristics of our datasets, such as missing labels or large number of annotators, this approach allows for the evaluation of multi-task learning across a variety contexts. An overview of the model and the predicted output is shown in Figure 1.

For a given text sample in a dataset, $\mathbf{x} \in \mathbf{X}$, our model $p_{\theta}(\mathbf{y}|\mathbf{x})$ predicts the individual annotation of each annotator $\mathbf{y} = (y_1, \dots, y_K)$, where K is the total number of unique annotators within the dataset. The predicted hard label of an instance is defined as aggregation of predictions into one label $mode(\mathbf{y}) = z$, where $z \in \{0, 1\}$, while soft labels as $v_0 \in [0, 1]$ and $v_1 \in [0, 1]$ that denote the possible probabilistic distributions of the annotations.

For the purposes of the shared task, we evaluate our models through F1 and cross entropy scores for the hard and soft labels respectively. We refer to the cross entropy loss of the shared task as *soft-label cross entropy*. We needed a different way to model our strong perspectivist approach, as the shared

task metrics prioritise predicting the *level of disagreement*. As seen in Figure 1, our model treated each annotator per instance as a subtask (shown as a classification layer). We evaluate the predicted versus the true labels for each annotator using cross entropy loss, which we refer to as *individual cross entropy*, which is also the metric used to train our model. Optimising individual cross entropy should lead to more accurate individual predictions, resulting in a representative annotation matrix, which can, in turn, be used to calculate the soft-label cross entropy. Since the aim of the shared task was to capture disagreement, we prioritised minimising soft-label cross entropy (CE) loss scores over high performing F1 scores.

This was done by manually stopping the model’s training procedure (for individual cross entropy loss) when the minimum soft-label cross entropy scores were achieved for each dataset. Thus, we do not optimize our model directly using the shared task evaluation metrics. However, these evaluation metrics can still be used as an indirect measure of our model’s performance. Hence, we report the shared task metrics on our model, which was trained on minimising the individual cross entropy loss in order to capture perspectives. While this method is not optimised to predict disagreement when compared to models trained by minimising soft-label cross entropy scores, it allows for the prediction of individual annotator perspectives. Thus, we model individual perspectives as an interim step towards predicting disagreement.

We compare the performance and analyse the benefits of using a multi-task model against the organisers’ baseline (aggregated labels), and two other models: a baseline neural model, and an SVM model (further described in subsection 4.3).

³Calculated using the [Spacy tokeniser](#) package for English, and [Arabert](#) package for Arabic.

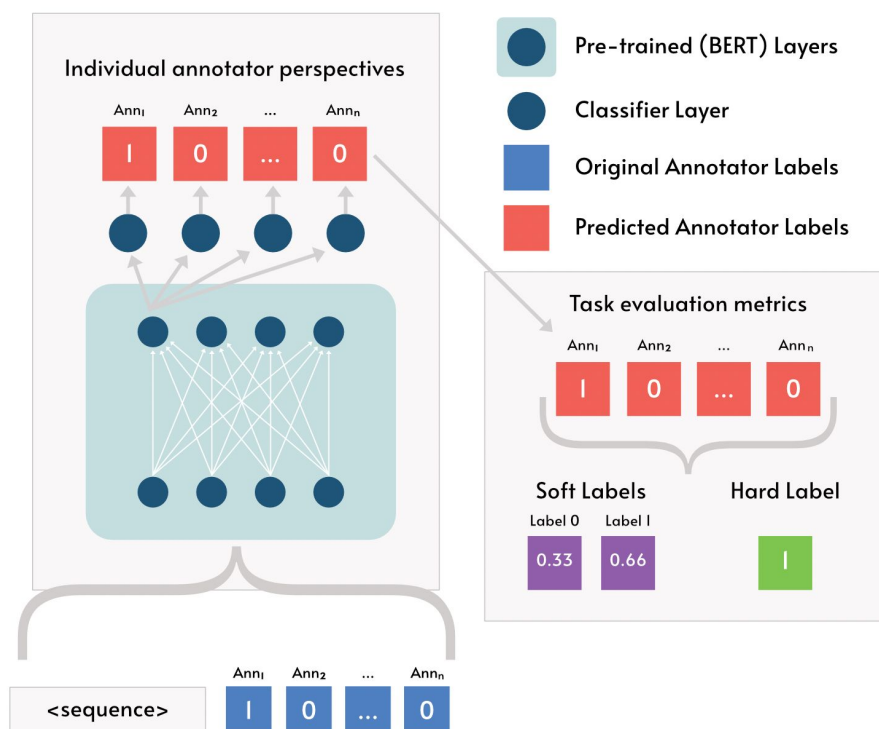


Figure 1: Representation of our multi-task architecture. As shown, we predict individual annotator perspectives (individual cross entropy, shown on the left of the figure) as an interim step to predicting the level of disagreement (the task metric of soft-label cross entropy, as shown on the right). For a full system description, refer to [subsection 4.1](#)

The SVM model was used as a linear model baseline, due to its prior success compared to neural approaches for cases of abuse detection (Niemann et al., 2020).

4.2 Text encoding

We applied the following pre-processing steps. For the ConvAbuse dataset, we only processed the last sentence uttered by a user, as Cercas Curry et al. (2021) reported no significant performance improvements from adding dialogue context. We preprocessed the Arabic dataset following Antoun et al. (2020). For the SVM model, both English and Arabic datasets were tokenised using term frequency-inverse document frequency (TF-IDF).

4.3 Model architectures

Baseline Linear model We trained an SVM model to perform binary classification with a linear kernel using a bag-of-words and TF-IDF approach.

The model outputs a distribution over the possible hard and soft labels, over which F1 and cross entropy scores were calculated.

Single-task (baseline) BERT Model For our transformer-based models, we used the pre-trained BERT-base-uncased (Devlin et al., 2019) model from the HuggingFace Transformers library (Wolf et al., 2020) for initialising English models. For the Arabic dataset, Transformer-based models were initialised with AraBERT, a variant of BERT pre-trained specifically for Arabic text which has shown comparable results to multilingual BERT in NLP-related tasks such as sentiment analysis and Named Entity Recognition (Antoun et al., 2020). Outputs of both models were fed through a linear layer, with a softmax activation function, resulting in a probability distribution for a binary label. The model used ADAM optimisation (Kingma and Ba, 2014). For specific parameters see Appendix A. Like the SVM model, the model outputs a distribution over the possible hard and soft labels, over which F1 and cross entropy scores were calculated.

Multi-task BERT Model Our model is a modified version of that proposed by Davani et al. (2022) as shown in Figure 1 and subsection 4.1. A pre-

Models	HS-Brexit		ArMIS		ConvAbuse		MD	
	F1 (\uparrow)	CE (\downarrow)	F1 (\uparrow)	CE (\downarrow)	F1 (\uparrow)	CE (\downarrow)	F1 (\uparrow)	CE (\downarrow)
Baseline (agg.)	0.89	2.71	0.59	8.23	0.95	3.38	0.78	7.74
TFIDF - SVM	0.86	0.62	0.68	2.57	0.9	0.49	0.88	0.62
Single-task BERT	0.47	0.43	0.78	1.77	0.88	0.37	0.74	0.61
Multi-Task BERT	0.46	0.76	0.58	1.89	0.77	0.50	0.41	1.28

Table 2: Model performance using F1 and cross-entropy (CE) scores for the four datasets of SemEval-2023 Task 11 Le-Wi-Di. \uparrow , and \downarrow indicate that higher and lower scores represent better performance, respectively. Baseline provided by the organisers (aggregation). We highlight in bold the highest scores for each dataset respectively.

trained BERT model is used to encode the text, upon which we train separate classification layers for each annotator. Training parameters are identical to our baseline transformer model (Appendix A). The embedded [CLS] token of the annotator label is fed into a classification layer to predict the each annotator’s label. As stated, we evaluate the predicted versus the true labels using cross entropy loss. Instances for which a particular annotator did not provide an annotation for the text were ignored when calculating the loss for that instance.

5 Results

The task is evaluated using cross entropy loss ($\epsilon = 1^{-12}$) and F1 scores (average = ‘micro’) of the hard and soft labels respectively.

Our results are shown in Table 2 and detailed results for all teams can be found in Leonardelli et al. (2023). Regarding F1 scores, the SVM model outperformed both the single-task BERT model, as well as the multi-task BERT model on the HS-Brexit, ConvAbuse, and MD datasets, findings which align with Niemann et al. (2020). The single-task BERT model outperformed both other architectures in the ArMIS dataset. For cross entropy scores, the single-task BERT model outperformed both other architectures across all datasets. Our model performs best on the ConvAbuse dataset, followed by the HS-Brexit dataset across both metrics.

6 Further Analyses

For deeper analysis, we used different methods depending on the specifics of the datasets. For ConvAbuse and MD, we followed Davani et al. (2022) to deal with *missing annotator labels* during the evaluation stage. Although not all annotators annotated every instance, our model still predicts labels for all annotators in the dataset (e.g. predicting eight

annotator labels in ConvAbuse for all instances that only have four true annotator labels as shown in Table 1). Essentially, predictions of missing annotator labels might have negatively impacted the soft-label cross entropy comparisons by skewing distributions. This is especially the case for the MD dataset, where each instance was annotated by only five of the 891 annotators. As such, we constrained the model to only predict labels for existing annotators of each instance, and reevaluated the soft-loss cross entropy and F1 scores. However, unlike Davani et al. (2022), our results degraded when constraining our model for both the ConvAbuse and MD datasets. Low performance in the MD dataset is further discussed in subsection 7.1.

We also investigated reasons for our poorest cross entropy scores, on the ArMIS dataset. We found that our model did not perform as well as expected in this scenario in which it should theoretically have performed well, i.e. with the annotators of the dataset self-identifying with a distinct ideological background (*conservative*, *moderate*, and *liberal*) and no missing annotator labels. We found that this may have been the result of the model’s architecture. Through testing with different batch sizes, we found that our model performed better when we added an extra hidden layer of 384 units to the existing 768-unit one. This would indicate that while the model was indeed learning, the original architecture with a single hidden layer was not sufficient to adequately disentangle enough information to make accurate predictions.

7 Discussion of Results & Limitations

7.1 Performance evaluation

We observe that our multi-task BERT model performed relatively well on the HS-Brexit and ConvAbuse datasets, but scored the lowest on the ArMIS and MD datasets. We expected to have limited success with the MD dataset, as multi-task

models have an inherent issue with sparse data when combined with a high numbers of subtasks (Ruder, 2017; Zhang and Yang, 2022). Not only does the MD-dataset contain more than 800 individual annotators, but the number of instances each annotator appeared in varies drastically (range 1 – 1988, mean = 65.63 ± 143.73). Furthermore, not all annotators appear in all splits, with 91% not represented in at least one of the training, development, and test sets. Individual annotator perspective modelling is therefore unfeasible for this dataset.

Such issues might be addressed through a clustering approach of labels. For example, Akhtar et al., 2019, 2020, 2021; Dethlefs et al., 2014 have proposed clustering methods based on a variety of features, including demographic similarities, as well as using inter and intra-annotator disagreement and similarities in labelling behaviour. However, as stated by Akhtar et al. (2021), a limitation of their approach is that it is necessary to know the demographic and cultural background of annotators, which is information that is not available for the MD dataset. We plan to investigate clustering methods in our future work.

As previously stated, our model performed the worst on the ArMIS dataset. While some strategies to improve these scores were discussed in section 6, we believe a possible explanation of these scores could be due to the size discrepancy of the datasets used in the original Davani et al. (2022) paper compared to the size of the ArMIS dataset. The original multi-task model for example, used datasets with $\approx 30,000 - 60,000$ instances in each dataset. This shows that while a dataset may contain distinct annotator perspectives (e.g. evident in the ArMIS dataset both by self-declared ideological categories and a moderate inter-annotator disagreement of 0.524), the multitask approach may not perform well on smaller datasets.

Another possible reason might have been our loss-weighting strategy, which sums the individual cross entropy across subtasks. Gong et al. (2019) explain that such multi-task approaches, where the model loss is constituted by the sum loss of subtasks, can lead to degradation of performance. This is due to a possible conflict arising through contrasting losses between subtask, or conflicting gradient signals (Chen et al., 2018; Sener and Koltun, 2018). This aligns with our experience during training, where summed loss remained relatively stable,

while individual loss across subtasks widely fluctuated. Furthermore, our soft-label cross entropy and F1 scores slowly improved over time in spite of the stable summed loss, indicating that some learning across subtasks was indeed taking place.

7.2 Capturing disagreement versus capturing perspectives

The single-task BERT baseline model outperformed the multi-task model across all evaluative metrics. Cabitza et al. (2023) explains that models trained through strong perspectivist approaches may be negatively impacted in terms of performance and evaluation metrics – as the more nuance present in data (such as disaggregated annotations), the more difficult the data is to model. While our model exhibits these weaknesses, there are *clear reasons* to use this approach; i.e. to model perspectives as an interim step in predicting disagreement. Our approach is successful in ways that would not be accounted for by simply predicting disagreement without this interim step.

Disagreement only shows that different perspectives are present in the dataset, but not the underlying reasons as to *why* disagreement may occur, *nor* the clashing perspectives present in the dataset. Particularly for highly subjective tasks, modelling only the level of disagreement does not consider intersecting perspectives. In contrast, strong perspectivist approaches offer insights into the different opinions present amongst individuals or groups of annotators. Modelling perspectives does not erase these individual viewpoints. For example, research has shown that attributes such as gender (Waseem and Hovy, 2016), or political activism status (Luo et al., 2020) of an annotator can elicit meaningful differences of opinions in a dataset.

Furthermore, it is well documented that aggregation can harm minorities present in a dataset by limiting their opinion’s influence (Prabhakaran et al., 2021). Gordon et al. (2022) explain that merely capturing disagreement can have a similar effect by presenting a simplified view of opposing perspectives in the data. This can be problematic, as without a nuanced understanding of which perspectives exist within a dataset, model predictions might not generalise well to end users’ perspectives (Gordon et al., 2021).

Accurately predicting each annotator’s perspective also captures their biases. However, bias is not an inherently negative trait. Though seldom explic-

itly stated, bias is an intrinsic attribute of annotators, datasets, and trained models (Gordon et al., 2022). While (de)biasing models can lead to positive outcomes when attempting to make a model ‘unlearn’ harmful social biases (Liang et al., 2020; Orgad and Belinkov, 2022; Subramanian et al., 2021a), Devinney et al. (2022) assert that incorporating bias stemming from marginalised groups while training, can lead to models that amplify the voice of those minorities. For example, in cases with datasets dealing with gender-based violence (Cercas Curry et al., 2021), it might be preferable to capture and amplify the bias of the affected people.

Combining approaches such as ours with the clustering approaches mentioned in subsection 7.1, merits future research, especially since fully debiasing models seems improbable (Gordon et al., 2022). As such, future research should attempt to utilise such multi-task models and strong perspectivist approaches when dealing with subjective tasks, in order to get a deeper understanding about why disagreement occurs.

8 Conclusion

We evaluated the performance of a multi-task model on predicting disagreement in four datasets, evaluated with both hard and soft labels, through F1 and cross-entropy loss respectively. Our model learned and predicted individual annotator perspectives for each instance.

Our model’s findings did not outperform our single-task BERT baseline in terms of the shared tasks evaluation metric. This was due to model employing a strong perspectivist approach, which prioritised capturing individual perspectives present in the dataset, over high performance. We argue that a strong perspectivist approach is preferable to merely modelling disagreement as it allows to capture different opinions present in a dataset, and can be used to further amplify minority views.

Evaluation metrics for this edition of Le-Wi-Di are geared towards measuring the overall levels of disagreement present in the datasets. However, if we wish to model stronger versions of perspectivism, we will need to develop new, more suitable metrics that can capture varying judgements in the kinds of different scenarios we have seen here.

9 Ethical considerations

Reproducibility We aim to maximise reproducibility by making all data manipulation and mod-

elling architecture aspects as explicit as possible in line with reproducibility principles (Belz et al., 2021). The code used to produce this study’s results can be found online in our team’s [GitHub repository](#).

Data manipulation and misrepresentation Following concerns about possible mishandling of data in this study, an important point has to be made about the ConvAbuse dataset (Cercas Curry et al., 2021). As explained in section 3, annotations ranged from $[-3, 1]$ with 1 denoting no abuse, 0 ambivalence, and the rest indicating severity of abuse. While analysing the dataset of this challenge, labels were aggregated to a binary depending on whether abuse was detected (labels -3,-1 in ConvAbuse), while the rest being annotated to no abuse detected. This transformation was necessary for the purposes of the shared task, as it has been shown that comparing scores between datasets with binary annotations and datasets with multiple labels can lead to incomparable results (Poletto et al., 2019).

Abuse detection, and handling of sensitive opinions There is also a larger conversation to be had about the use of a hard label in general, regarding issues such as bias, abuse, and such sensitive topics. Through a strong perspectivist approach, an annotator’s viewpoint might also reflect the perspective of a minority population (Cabitza et al., 2023). It is important to be sensitive when dealing with such opinions without invalidating them, or minimising them through a distribution. In essence, if there is even a small amount of disagreement in whether an item is problematic or not, and the label should reflect that beyond a binary (Blodgett, 2021). We leave it to future research to explain how exactly these multiple labels could be appropriately incorporated into a model architecture.

Dual use of model Finally, it is also important to explain that the model architecture proposed can also unfortunately be used for purposes beyond our original intention (Hovy and Prabhumoye, 2021). While our model can be used towards furthering social justice aims through the amplification of minorities’ perspectives, the model could also be used to manipulate perspectives from the dataset, in order to present stable results. For example, approaches have attempted to create perspective clusters that are more inclusive towards the data, (Subramanian et al., 2021a). Unfortunately, these

findings seem to result in bias mitigation (Shen et al., 2022; Subramanian et al., 2021b), rather than bias erasure in both downstream tasks (Lalor et al., 2022), as well as resulting word embeddings (Gonen and Goldberg, 2019). As biases are unavoidable, we advocate boosting of under-represented perspectives that might otherwise be lost. Such an approach would also make attempts to use this model to reproduce and platform problematic opinions transparent.

10 Acknowledgements

Gavin Abercrombie and Verena Rieser were supported by the EPSRC projects ‘Gender Bias in Conversational AI’ (EP/T023767/1) and ‘Equally Safe Online’ (EP/W025493/1). Tanvi Dinkar and Verena Rieser were supported by ‘AISEC: AI Secure and Explainable by Construction’ (EP/T026952/1).

References

- Gavin Abercrombie, Valerio Basile, Sara Tonelli, Verena Rieser, and Alexandra Uma, editors. 2022. *Proceedings of the 1st Workshop Approaches to NLP @LREC2022*. European Language Resources Association, Marseille, France.
- Sohail Akhtar, Valerio Basile, and Viviana Patti. 2019. A new measure of polarization in the annotation of hate speech. In *AI*IA 2019 – Advances in Artificial Intelligence*, pages 588–603, Cham. Springer International Publishing.
- Sohail Akhtar, Valerio Basile, and Viviana Patti. 2020. Modeling annotator perspective and polarized opinions to improve hate speech detection. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 8, pages 151–154.
- 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*. *CoRR*, abs/2106.15896.
- 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.
- Wissam Antoun, Fady Baly, and Hazem Hajj. 2020. *AraBERT: Transformer-based model for Arabic language understanding*. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 9–15, Marseille, France. European Language Resource Association.
- Anya Belz, Shubham Agarwal, Anastasia Shimorina, and Ehud Reiter. 2021. *A systematic review of reproducibility research in natural language processing*. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 381–393, Online. Association for Computational Linguistics.
- Su Lin Blodgett. 2021. *Sociolinguistically Driven Approaches for Just Natural Language Processing*. Ph.D. thesis, University of Massachusetts Amherst.
- Federico Cabitza, Andrea Campagner, and Valerio Basile. 2023. Toward a perspectivist turn in ground truthing for predictive computing. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- 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.
- Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. 2018. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International conference on machine learning*, pages 794–803. PMLR.
- 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.
- Nina Dethlefs, Heriberto Cuayáhuitl, Helen Hastie, Verena Rieser, and Oliver Lemon. 2014. *Cluster-based prediction of user ratings for stylistic surface realisation*. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 702–711, Gothenburg, Sweden. Association for Computational Linguistics.
- Hannah Devinney, Jenny Björklund, and Henrik Björklund. 2022. Theories of “gender” in nlp bias research. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 2083–2102.
- 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.
- 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.
- Tommaso Fornaciari, Alexandra Uma, Massimo Poesio, and Dirk Hovy. 2022. [Hard and soft evaluation of NLP models with BOOtSTrap SAMpling - BooStSa](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 127–134, Dublin, Ireland. Association for Computational Linguistics.
- Hila Gonen and Yoav Goldberg. 2019. [Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them](#). In *Proceedings of the 2019 Workshop on Widening NLP*, pages 60–63, Florence, Italy. Association for Computational Linguistics.
- Ting Gong, Tyler Lee, Cory Stephenson, Venkata Renduchintala, Suchismita Padhy, Anthony Ndirango, Gokce Keskin, and Oguz H Elibol. 2019. A comparison of loss weighting strategies for multi task learning in deep neural networks. *IEEE Access*, 7:141627–141632.
- Mitchell L Gordon, Michelle S Lam, Joon Sung Park, Kayur Patel, Jeff Hancock, Tatsunori Hashimoto, and Michael S Bernstein. 2022. Jury learning: Integrating dissenting voices into machine learning models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–19.
- Mitchell L Gordon, Kaitlyn Zhou, Kayur Patel, Tatsunori Hashimoto, and Michael S Bernstein. 2021. The disagreement deconvolution: Bringing machine learning performance metrics in line with reality. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–14.
- Dirk Hovy and Shrimai Prabhumoye. 2021. Five sources of bias in natural language processing. *Language and Linguistics Compass*, 15(8):e12432.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- John Lalor, Yi Yang, Kendall Smith, Nicole Forsgren, and Ahmed Abbasi. 2022. [Benchmarking intersectional biases in NLP](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3598–3609, Seattle, United States. 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.
- Elisa Leonardelli, Gavin Abercrombie, Dina Almanea, Valerio Basile, Tommaso Fornaciari, Barbara Plank, Massimo Poesio, Verena Rieser, and Alexandra Uma. 2023. SemEval-2023 Task 11: Learning With Disagreements (LeWiDi). In *Proceedings of the 17th International Workshop on Semantic Evaluation*, Toronto, Canada. Association for Computational Linguistics.
- Paul Pu Liang, Irene Mengze Li, Emily Zheng, Yao Chong Lim, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2020. [Towards debiasing sentence representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5502–5515, Online. Association for Computational Linguistics.
- Yiwei Luo, Dallas Card, and Dan Jurafsky. 2020. [Detecting stance in media on global warming](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3296–3315, Online. Association for Computational Linguistics.
- Marco Niemann, Jens Welsing, Dennis M. Riehle, Jens Brunk, Dennis Assenmacher, and Jörg Becker. 2020. Abusive comments in online media and how to fight them. In *Disinformation in Open Online Media*, pages 122–137, Cham. Springer International Publishing.
- Hadas Orgad and Yonatan Belinkov. 2022. Debiasing nlp models without demographic information. *arXiv preprint arXiv:2212.10563*.
- 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. Learning part-of-speech taggers with inter-annotator agreement loss. In *Conference of the European Chapter of the Association for Computational Linguistics*.
- Fabio Poletto, Valerio Basile, Cristina Bosco, Viviana Patti, Marco Stranisci, et al. 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.
- Filipe Rodrigues and Francisco Pereira. 2018. Deep learning from crowds. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Sebastian Ruder. 2017. [An overview of multi-task learning in deep neural networks](#).

Ozan Sener and Vladlen Koltun. 2018. Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, 31.

Aili Shen, Xudong Han, Trevor Cohn, Timothy Baldwin, and Lea Frermann. 2022. [Optimising equal opportunity fairness in model training](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4073–4084, Seattle, United States. Association for Computational Linguistics.

Shivashankar Subramanian, Xudong Han, Timothy Baldwin, Trevor Cohn, and Lea Frermann. 2021a. [Evaluating debiasing techniques for intersectional biases](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2492–2498, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Shivashankar Subramanian, Afshin Rahimi, Timothy Baldwin, Trevor Cohn, and Lea Frermann. 2021b. [Fairness-aware class imbalanced learning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2045–2051, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Alexandra Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2021. Learning from disagreement: A survey. *Journal of Artificial Intelligence Research*, 72.

Bertie Vidgen and Leon Derczynski. 2021. [Directions in abusive language training data, a systematic review: Garbage in, garbage out](#). *PLOS ONE*, 15:1–32.

Zeeraq Waseem and Dirk Hovy. 2016. [Hateful symbols or hateful people? predictive features for hate speech detection on Twitter](#). In *Proceedings of the NAACL Student Research Workshop*, pages 88–93, San Diego, California. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Yu Zhang and Qiang Yang. 2022. [A survey on multi-task learning](#). *IEEE Transactions on Knowledge and Data Engineering*, 34(12):5586–5609.

A Hyper-parameters for our experiments

Both AraBERT and BERT consist of 12 encoder blocks, 12 attention heads, a hidden layer size of 768, while for the purposes of this study; maximum sequence length for both was set to 128. Neural models learned through a fixed learning rate of 1^{-8} , and a batch size of 64, over a maximum of 100 epochs. We used early stopping with a patience of 40 epochs while minimising the cross entropy scores obtained through the validation set.