

# SafeWebUH at SemEval-2023 Task 11: Learning Annotator Disagreement in Derogatory Text: Comparison of Direct Training vs Aggregation

**Sadat Shahriar**

University of Houston  
sshahriar@uh.edu

**Thamar Solorio**

University of Houston  
tsolorio@uh.edu

## Abstract

Subjectivity and difference of opinion are key social phenomena, and it is crucial to take these into account in the annotation and detection process of derogatory textual content. In this paper, we use four datasets provided by SemEval-2023 Task 11 and fine-tune a BERT model to capture the disagreement in the annotation. We find individual annotator modeling and aggregation lowers the Cross-Entropy score by an average of 0.21, compared to the direct training on the soft labels. Our findings further demonstrate that annotator metadata contributes to the average 0.029 reduction in the Cross-Entropy score.

## 1 Introduction

While the web space is inundated with derogatory textual content, the subjectivity of their interpretation frequently necessitates a system capable of capturing reader disagreements. The Learning-With-Disagreement (Le-Wi-Di) task involves learning annotators' disagreements based on how they categorize a text (Leonardelli et al., 2023). Recent research has found that almost every annotation task contains a wide range of disagreements (Dumitrache et al., 2019; Pavlick and Kwiatkowski, 2019). The subjective and biased nature of the raters, among other elements of natural language comprehension, make it crucial to learn disagreements through annotations (Uma et al., 2021a). In this study, we compare two strategies of disagreement learning: Disagreement Targeted Learning of soft labels, and annotator-specific learning with Post Aggregation, using BERT model. Furthermore, we utilize annotator-specific metadata, to capture annotators' disagreements in disparaging content.

Since the advent of social media, which has flooded the web with massive amounts of content, the number of offensive text, such as hate speech, misogyny, sexism, and abusive content has also

increased significantly. Several studies were carried out to battle this problem, such as, Burnap and Williams studied online hate-speech in tweets, triggered by the murder of Lee Rigby, a London-based drummer (Burnap and Williams, 2015). Xu et al. formulated the cyber-bullying in social media as an NLP task (Xu et al., 2012). Similar works are conducted in Warner and Hirschberg 2012; Silva et al. 2016; Gitari et al. 2015. However, tasks related to the detection of social phenomena, like offensiveness, and toxicity are often subjective in nature (Kocoń et al., 2021). A recent survey among American adults stated that according to half of the participants, "it is hard to know what others might find offensive", and the majority of them acknowledged there were disagreements in what is perceived as sexist or racist (pew, Accessed: 2022-12-03). To this end, we aim to develop a system that can capture subjective disagreement in derogatory text.

The four datasets in the Le-Wi-di task come with the annotator-specific labels, with aggregated hard labels (majority voting) and soft labels (average of the labels). Although a system for modeling disagreements should be trained to estimate soft labels, it is not clear, whether direct training on the soft label or aggregating on the annotator labels is a better approach. Hence, our first research question (Q1): Can annotator-specific classification models, and post hoc aggregation outperform the direct approach of regression on soft labels in disagreement modeling? Additionally, we explore the annotator metadata which explains how an annotator labeled other related text, and we pose the question (Q2): Can annotator metadata improve the disagreement modeling? To address these questions, we compare BERT-based disagreement-targeted learning (regression) and post-aggregation learning (classification) and explore different strategies for incorporating annotator metadata to model the disagreement. However, due to the inconsistency of

the annotators and lack of metadata, we limit our comparisons to two datasets only.

Our work has several important implications. To begin, our model’s ability to capture conflicts makes it applicable to the modeling of controversial social phenomena and public opinions. Hence, it can be used to model ambiguity in textual ambiguity. Furthermore, our explorations of incorporating annotator metadata can help understanding readers’ perception and outlook in different context. Finally, enhancing transparency and accountability among the raters can be performed as a mean to quality control in multi-rater annotation process. The code for implementing our work is available here: <https://github.com/sadat1971/Le-Wi-Di-SemEval-23>

## 2 Dataset and Task Description

SemEval’23 Task 11 has four datasets that deal with derogatory text. While the three datasets are in English, *ArMIS* is in Arabic. Along with soft and hard labels, each dataset contains some metadata. They are described below in brief.

The *MultiDomain Agreement* (MD) dataset comes with tweets from three domains: BLM, Election and COVID-19 (Leonardelli et al., 2021). A total of 819 annotators were used to label all the tweets using AMT. A random combination of five annotators was chosen to label each tweet for offensiveness. The train set contains 6,592 tweets, the dev set from the practice phase has 1,104 tweets, and the test set from the evaluation phase contains 3,057 tweets.

The *HS-Brexit* dataset contains tweets related to Brexit, and annotation from six annotators (a target group of three Muslim immigrants in the UK and a control group of three) (Akhtar et al., 2021). Each of them labeled a tweet for hate speech, which is the target class of the task. They also annotated tweets for being offensive and aggressive. The train, dev, and test set have 784, 168, and 168 tweets respectively.

Misogyny and Sexism are labeled in the *ArMIS* dataset, rated by three annotators (Moderate Female, Liberal Female, and Conservative Male) (Almanea and Poesio, 2022). There are 657, 141, and 145 tweets in the train, dev, and test sets, respectively.

The *ConvAbuse* dataset captures dialogues between a user and two conversational agents, and at least two annotators annotated the conversation

for abusiveness (Cercas Curry et al., 2021). The dataset also provides labels for a conversation being sexist, explicit, implicit, intellectual, racist, transphobic, homophobic, and the target of the abuse. The train, dev, and test set have 2,398, 812, and 840 tweets respectively.

## 3 System Description

For the textual data, we use a pretrained language representation model, called BERT (Devlin et al., 2019). Since BERT is trained on the English data only, to handle the *ArMIS* task, we use Arabic-BERT (Safaya et al., 2020). Figure 1 shows the system description. To address **Q1**, we compare two techniques— Post Aggregation and Disagreement Targeted Learning, and we also investigate the effect of metadata to address **Q2**. The performance is measured by F1-score and Cross-Entropy (CE) score.

### 3.1 Post-Aggregation

In the *Post-Aggregation* (Post-Agg) approach, separate models are trained to learn the annotation pattern of each annotator. First, the BERT model is fine-tuned to learn the target class, and the softmax score  $S$  is obtained for all annotators. Next, we process the metadata to extract important information. For the *HS-Brexit* dataset, in addition to labeling for hate speech, each annotator also labeled tweets for offensive and aggressive, which is available with the dataset. We compute the probability of a tweet being labeled as hate speech, given how it is labeled by an annotator as offensive and aggressive, which we denote as  $P$ . For each tweet, the soft label  $\hat{S}L$  is then computed as,

$$\hat{S}L(w) = \frac{1}{N} \sum_{i=1}^N \frac{S_i + w * P_i}{1 + w} \quad (1)$$

where  $N$  is the number of annotators. Since both  $S_i$  and  $P_i$  are predicted soft labels, we find their weighted average and select  $w$ , where the minimum CE score and maximum F1-score are obtained based on the dev set.

### 3.2 Disagreement Targeted Learning

While the Post-Agg approach considers learning from each annotator, the Disagreement Targeted Learning (Dis-Learning) approach learns only from the aggregated labels. First, a BERT model is fine-tuned using a downstream regression task of estimating the soft label, and the predicted variable,




Figure 1: The text is fed to a pretrained BERT model, and fine-tuned for the downstream task. For the Post-Aggregation approach, the downstream classification task of the BERT model is to predict the label for each annotator. The softmax value from each annotator model with the metadata is ensemble to produce the annotator-specific soft-label and hard-label prediction. For the Disagreement Targeted Learning approach, the BERT downstream task is to directly learn the soft labels. The ensemble mechanism is performed by regression approach to learn the final prediction.

$SL_{BERT}$  is obtained. Next, we measure the average rating of each metadata for all annotators across the entire dataset. For example, in HS-Brexit dataset, if two annotator labels a tweet as offensive, while four as not-offensive, the average metadata (offensiveness score) for that tweet will be  $2/6 = 0.33$ . Next, we train a linear regression model to predict the soft label based only on the available average metadata rating.

$$SL_{meta} = b_0 + b_1 * M_1 + b_2 * M_2 \quad (2)$$

$b_0, b_1, b_2$  are trained from the linear regression model, and  $M_1$  and  $M_2$  are two metadata scores. For HS-Brexit, we use average offensive and aggressive measures. For the ConvAbuse dataset, out of twelve metadata labels, we pick the top two, *explicit* and *target system* which yielded the best correlation coefficient with the soft label values. Finally, we find  $SL$  by averaging  $SL_{BERT}$  and  $SL_{meta}$ .

#### 4 Experimental Set-up

All of our models use “bert-base-uncased” version of BERT (“bert-base-arabic” in ArMIS). We deploy a two-layered fully-connected network for fine-tuning in both regression and classification tasks. We choose the hyper-parameters from all the combinations, by three-fold cross-validation in the practice phase, and on the released validation set in the evaluation phase. The hidden size and dropout rate are chosen from  $\{32, 64, 128, 256\}$ , and  $\{.1, .3, .5\}$ . The learning rate is chosen from  $\{5e-4, 1e-5, 5e-5, 1e-6\}$ . We keep the batch size small due

to the GPU limitations and choose from  $\{8, 16\}$ . Since BERT models quickly overfit on the data, we kept the epoch size between 2 and 4. However, for “arabic-bert-base”, the performance was unstable, and we train upto 10 epochs. For all cases, AdamW is used as optimizer (Loshchilov and Hutter, 2017). For all our experiments, Pytorch version 1.11.0 is used (Paszke et al., 2019).

To evaluate the result of capturing disagreement, we use the Cross-Entropy score provided by the competition. If the target soft label is  $T$ , and predicted soft label is  $P$ , for a dataset of size  $D$ , the Cross-Entropy (CE) is computed as:

$$CE = -\frac{1}{D} \sum_{i=1}^D T_i * \log(P_i + 1e - 9) \quad (3)$$

We further report the F1-score (micro) on the hard label to evaluate the model performance on the majority-voted final prediction task.

#### 5 Result and Discussion

Table 1 shows that for HS-Brexit dataset, the Post-Agg approach does not improve the F1-score from the Dis-Learning approach. However, the Post-Agg approach is able to reduce the CE score by 0.1400 from the Dis-Learning approach. The reduction is even higher when metadata is used (by 0.1958). Similarly, for the ArMIS dataset, Dis-Learning approach has higher F1-score compared to the Post-Agg approach, while the CE score is lower in the Dis-Learning approach (reduced by 0.3070).

We further investigate why the Post-Agg approach works better at capturing disagreement. Since the Dis-Learning approach does not take into

Dataset	Post-Agg		Post-Agg-meta		Dis-Learning		Dis-Learning-meta	
	F1	CE	F1	CE	F1	CE	F1	CE
MD	–	–	–	–	<b>0.8266*</b>	<b>0.5076*</b>	–	–
HS-Brexit	0.8810	0.1686	<b>0.9167</b>	<b>0.0834</b>	0.8869	0.3086	0.9107*	0.2792*
ArMIS	0.7211	<b>0.2683</b>	–	–	<b>0.7586*</b>	0.5753*	–	–
ConvAbuse	–	–	–	–	0.9321*	0.2364*	<b>0.9667</b>	<b>0.0688</b>

Table 1: Comparing the performances of disagreement modeling approaches in all four datasets. Because the MD and ConvAbuse datasets lack consistent annotators, their Post-Agg results are not reported. Also, MD and ArMIS lack annotator-specific metadata, and thus, their metadata-incorporated performance is not reported as well. The best performance in each dataset is denoted in bold numbers, whereas the performance submitted in the Le-Wi-Di task is indicated with asterisks (\*).

account individual annotators, it mainly approximates the “intensity” of a text being derogatory. Conversely, the Post-Agg approach considers each annotator separately and learns their annotation pattern, which is aggregated afterward. Consequently, Dis-Learning has to depend only on textual data, making its job harder than Post-Agg. However, in a realistic case, the annotators may not be consistent (as in the MD and ConvAbuse datasets), or a large number of models are needed to be trained, rendering the Post-Agg technique infeasible. Therefore, the Post-Agg approach is better suited for modeling disagreement if a small number of annotators are consistent across the dataset. Hence, **Q1** is addressed.

Next, the results reveal that performance is enhanced when annotator metadata is utilized as opposed to when it is not (Table 1). Using the metadata reduced the CE score for the HS-Brexit dataset by 0.0852 and 0.0294 for the Post-Agg and Dis-Learning approaches, respectively. Similarly, for the ConvAbuse dataset, annotator metadata helps lower the CE score by 0.1676. The metadata contains useful annotation patterns of the annotators, which ameliorates the learning process. Notably, we have not used the metadata from MD and ArMIS, since they do not contain the related annotation information from the annotators. Therefore, **Q2** is addressed.

In the MD, HS-Brexit, ArMIS, and ConvAbuse datasets, our results were ranked 7th, 9th, 11th, and 12th, respectively. Overall, we ranked 9th in the CE score category and 8th in the F1-score category.

**Error Analysis** Finally, we focus on the error analysis of this study. We find that both our approaches often make mistakes in prediction for the texts that do not use slang or curse words but are still voted by the majority as offensive. For ex-

ample, three of the five annotators annotated the following sentence as offensive (soft label 0.60): *#TonyBobulinski #MAGA2020 #MAGA #ChangeYourVoteToTrump #BidenCrimeFamily #BidenHarris2020 #BidenCares #LaptopFromHell Joe is going down. <url>*. However, our model predicts the soft label as 0.15. Similarly, tweets that contain curse words but do not necessarily exhibit offensiveness, are sometimes mistaken by our model as hate speech. For example, the tweet: *Astounding Words from the prolific and talented - <user> #BlackLivesMatter #fucktrump <url>* is labeled as non-offensive by three annotators out of five, however, our model predicts the soft label as 0.85, due to the presence of profane language in one of the hashtags.

## 6 Related Works

Though the majority of AI learning still operates under the assumption that a single interpretation exists for each item, research is growing to build learning methods that do not rely on this assumption (Uma et al., 2021b). Rater’s Disagreement is a familiar phenomenon in Natural Language Processing (Poesio and Artstein, 2005; Recasens et al., 2011). The disagreement may take place because of the annotator error or interface problem (Plank et al., 2014), explicit or implicit ambiguity (Poesio and Artstein, 2005), item difficulty (Zaenen et al., 2005), and subjectivity (Akhtar et al., 2019). Notwithstanding, the simpler task such as POS tagging (Plank et al., 2014) to subjective tasks like sentiment analysis, semantic role assignments also involve raters’ disagreement (Kenyon-Dean et al., 2018; Dumitrache et al., 2019). Hence, researchers argued for taking disagreement into account during the labeling process and retaining the implicit ambiguity (Recasens et al., 2012; Poesio and Artstein,

2005)). To this end, we explore the Learning-With-Disagreement task for derogatory text.

The previous version of this competition was launched in 2021, where the organizers used NL and image-classification task to address for disagreement in the labeling (Uma et al., 2021a). The winning team used the Sharpness-Aware Minimization technique (SAM) and a special NN layer called softmax Crowd-layer with BERT as baseline model (Osei-Brefo et al., 2021). While the SAM architecture was mainly used for CIFAR-10 (image classification), the Crowdlayer architecture aims to map the label with each individual annotator. Since the current competition only involves text, we fine-tune a BERT model and use the annotator metadata to capture the disagreement.

## 7 Conclusion

Because of the proliferation of social media content, the internet has become a breeding ground for derogatory text. However, due to the differences in human perception and opinion, often there is no unanimous consensus among the annotators about the text being derogatory or not. Hence, it is imperative to store the soft labels and capture annotator disagreement in the modeling process. Our work compares the direct training on the soft label with the annotator-specific model and post-aggregation. We find that with the presence of consistent annotators, it might be helpful to take the latter approach. In addition, integrating annotator metadata has been proved to be beneficial in our experiments. Our work has a wide variety of potential future research directions, such as:

- We only modeled with one Transformer-based approach, BERT. In the future, we plan to use RoBERTA, ELECTRA and XLMNet
- We find a strong correlation between hate speech and offensiveness. Therefore, we plan to investigate how cross-dataset performance works. Such experiments will also help to make our model more generalizable.
- Because language evolves in response to social context and other phenomena, it is critical to include Continual Learning (CL) techniques and investigate the distribution shift in the annotation process. In the future, we intend to incorporate CL into our work.

## References

- Accessed: 2022-12-03. *The challenge of knowing what's offensive*. Pew Research Center.
- 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: XVIIIth International Conference of the Italian Association for Artificial Intelligence, Rende, Italy, November 19–22, 2019, Proceedings 18*, pages 588–603. Springer.
- 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. *arXiv preprint arXiv: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.
- Peter Burnap and Matthew Leighton Williams. 2015. Cyber hate speech on twitter: An application of machine classification and statistical modeling for policy and decision making. *Policy & Internet*, 7:223–242.
- 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.
- 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.
- Anca Dumitrache, Lora Aroyo, and Chris Welty. 2019. *A crowdsourced frame disambiguation corpus with ambiguity*. 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 2164–2170, Minneapolis, Minnesota. Association for Computational Linguistics.
- Njagi Dennis Gitari, Zhang Zuping, Zuping Zhang, Hanyurwimfura Damien, and Jun Long. 2015. A lexicon-based approach for hate speech detection. In *International Conference on Multimedia and Ubiquitous Engineering*.
- Kian Kenyon-Dean, Eisha Ahmed, Scott Fujimoto, Jeremy Georges-Filteau, Christopher Glasz, Barleen Kaur, Auguste Lalande, Shruti Bhandari, Robert

- Belfer, Nirmal Kanagasabai, Roman Sarrazingendron, Rohit Verma, and Derek Ruths. 2018. [Sentiment analysis: It's complicated!](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1886–1895, New Orleans, Louisiana. Association for Computational Linguistics.
- Jan Kocoń, Marcin Gruza, Julita Bielaniewicz, Damian Grimling, Kamil Kanclerz, Piotr Miłkowski, and Przemysław Kazienko. 2021. [Learning personal human biases and representations for subjective tasks in natural language processing.](#) In *2021 IEEE International Conference on Data Mining (ICDM)*, pages 1168–1173.
- 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.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Emmanuel Osei-Brefo, Thanet Markchom, and Huizhi Liang. 2021. [UOR at SemEval-2021 task 12: On crowd annotations; learning with disagreements to optimise crowd truth.](#) In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 1303–1309, Online. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library.](#) In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. 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, Dirk Hovy, and Anders Søgaard. 2014. [Linguistically debatable or just plain wrong?](#) In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 507–511, Baltimore, Maryland. Association for Computational Linguistics.
- Massimo Poesio and Ron Artstein. 2005. [The reliability of anaphoric annotation, reconsidered: Taking ambiguity into account.](#) In *Proceedings of the Workshop on Frontiers in Corpus Annotations II: Pie in the Sky*, pages 76–83, Ann Arbor, Michigan. Association for Computational Linguistics.
- Marta Recasens, Eduard Hovy, and M. Antònia Martí. 2011. [Identity, non-identity, and near-identity: Addressing the complexity of coreference.](#) *Lingua*, 121(6):1138–1152.
- Marta Recasens, M. Antònia Martí, and Constantin Orasan. 2012. [Annotating near-identity from coreference disagreements.](#) In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 165–172, Istanbul, Turkey. European Language Resources Association (ELRA).
- Ali Safaya, Moutasem Abdullatif, and Deniz Yuret. 2020. [KUISAIL at SemEval-2020 task 12: BERT-CNN for offensive speech identification in social media.](#) In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 2054–2059, Barcelona (online). International Committee for Computational Linguistics.
- Leandro Silva, Mainack Mondal, Denzil Correa, Fabrício Benevenuto, and Ingmar Weber. 2016. Analyzing the targets of hate in online social media. In *Tenth international AAAI conference on web and social media*.
- Alexandra Uma, Tommaso Fornaciari, Anca Dumitrache, Tristan Miller, Jon Chamberlain, Barbara Plank, Edwin Simpson, and Massimo Poesio. 2021a. [SemEval-2021 task 12: Learning with disagreements.](#) In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 338–347, Online. Association for Computational Linguistics.
- Alexandra N Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2021b. Learning from disagreement: A survey. *Journal of Artificial Intelligence Research*, 72:1385–1470.
- William Warner and Julia Hirschberg. 2012. [Detecting hate speech on the world wide web.](#) In *Proceedings of the Second Workshop on Language in Social Media*, pages 19–26, Montréal, Canada. Association for Computational Linguistics.
- Jun-Ming Xu, Kwang-Sung Jun, Xiaojin Zhu, and Amy Bellmore. 2012. [Learning from bullying traces in social media.](#) In *Proceedings of the 2012 Conference of the North American Chapter of the Association*

*for Computational Linguistics: Human Language Technologies*, pages 656–666, Montréal, Canada. Association for Computational Linguistics.

Annie Zaenen, Lauri Karttunen, and Richard Crouch. 2005. Local textual inference: can it be defined or circumscribed? In *Proceedings of the ACL workshop on empirical modeling of semantic equivalence and entailment*, pages 31–36.