

Estimating the Level of Dialectness Predicts Interannotator Agreement in Multi-dialect Arabic Datasets

Amr Keleg, Walid Magdy, Sharon Goldwater

Institute for Language, Cognition and Computation
School of Informatics, University of Edinburgh
a.keleg@sms.ed.ac.uk, {wmagdy, sgwater}@inf.ed.ac.uk

Abstract

On annotating multi-dialect Arabic datasets, it is common to randomly assign the samples across a pool of native Arabic speakers. Recent analyses recommended routing dialectal samples to native speakers of their respective dialects to build higher-quality datasets. However, automatically identifying the dialect of samples is hard. Moreover, the pool of annotators who are native speakers of specific Arabic dialects might be scarce. Arabic Level of Dialectness (ALDi) was recently introduced as a quantitative variable that measures how sentences diverge from Standard Arabic. On randomly assigning samples to annotators, we hypothesize that samples of higher ALDi scores are harder to label especially if they are written in dialects that the annotators do not speak. We test this by analyzing the relation between ALDi scores and the annotators' agreement, on 15 public datasets having raw individual sample annotations for various sentence-classification tasks. We find strong evidence supporting our hypothesis for 11 of them. Consequently, we recommend prioritizing routing samples of high ALDi scores to native speakers of each sample's dialect, for which the dialect could be automatically identified at higher accuracies.

1 Introduction

Arabic is spoken natively by over 420 million people and is an official language of 24 countries (Bergman and Diab, 2022), making it an important language for NLP systems. However, NLP for Arabic faces a major challenge in that user-generated text is typically a mixture of Modern Standard Arabic (MSA)—the standardized variant that is taught in schools and used in official communications and newspapers—and regional variants of Dialectal Arabic (DA), which are used in everyday communications, including both speech and

social media text (Habash, 2010).¹ While MSA can be largely understood by most Arabic speakers, the different variants of DA are not always fully mutually intelligible.

Despite this mutual unintelligibility, a common practice when developing datasets for multi-dialect Arabic NLP is to randomly recruit annotators without regard to their dialect. However, routing dialectal content to speakers of a different dialect for annotation or moderation can present real problems. For example, it has been shown to contribute to unjust online content moderation of DA (Business for Social Responsibility, 2022), and racially biased toxicity annotation in American English varieties (Sap et al., 2022). Two recent studies of multi-dialect DA annotation showed that for annotating hate speech or sarcasm, respectively, annotators were more lenient (for hate speech) and more accurate (for sarcasm) when annotating sentences in their native dialect (Bergman and Diab, 2022; Abu Farha and Magdy, 2022). The authors of both studies made the same recommendation for creating new Arabic datasets, namely to first identify the dialect of each sample and then route it to appropriate annotators.

This recommendation is theoretically appealing, but presents practical difficulties since automatic dialect identification (DI) is challenging (Abdul-Mageed et al., 2023), and existing systems assume a single correct label when in fact some texts can be natural in different dialects (Keleg and Magdy, 2023; Olsen et al., 2023). Moreover, the representation of native speakers of the different Arabic dialects on crowdsourcing sites might be skewed (Mubarak and Darwish, 2016). Therefore, recruiting native speakers of some Arabic dialects might be challenging, given the tough conditions of the countries in which these dialects are spoken.

¹Refer to §B of the Appendix for a further discussion about the relationship between MSA and DA.

In this paper, we address these challenges by building on recent work by Keleg et al. (2023), who presented a system for estimating Arabic Level of Dialectness (ALDi)—i.e., the degree to which a sentence diverges from MSA, on a scale from 0 to 1. We hypothesize that as sentences with low ALDi scores do not diverge much from MSA, they can still be understood and accurately annotated by most Arabic speakers, while this will be less true for sentences with high ALDi scores. If our hypothesis holds, then annotation can be made more efficient while maintaining accuracy, by routing samples with low ALDi scores to speakers of any dialect. Only high-ALDi samples need to be routed to native speakers of the appropriate dialect.

We test our hypothesis by investigating the impact of ALDi score on interannotator agreement for 15 publicly released datasets annotated for 6 different sentence-classification tasks.² We confirm that for most tasks and datasets, higher ALDi scores correlate with lower annotator agreement. A notable exception is the dialect identification (DI) task, where higher ALDi scores correlate with *higher* agreement, presumably because it is easier to identify a single dialect for sentences that are strongly dialectal. This finding is encouraging for annotation routing, since automatic DI systems may also have higher accuracy on these sentences. We conclude that a combination of automatic ALDi scoring, followed by DI and annotator routing only for high-ALDi sentences, is a promising strategy for annotating multi-dialect Arabic datasets.

2 Methodology

Data We study the impact of ALDi scores on the annotators’ agreement for publicly released Arabic datasets. We analyze datasets satisfying the following criteria:

- **Language:** Mixture of MSA and DA.
- **Variation:** Targeting multiple variants of DA.
- **Annotators:** Speakers of different variants of DA that are randomly assigned to the samples.
- **Tasks Setup:** Sentence-level classification.
- **Released Labels:** Individual annotator labels or the percentage of annotators agreeing on the majority-vote label.³

²Instructions to replicate the experiments can be accessed through <https://github.com/AMR-KELEG/ALDi-and-IAA>

³For some datasets, the percentage of annotators agreeing on the majority vote is weighted by their performance on the annotation quality-assurance test samples. This distinction is irrelevant to our study, where we only consider whether all annotators agreed or not.

We searched for datasets on Masader, a community-curated catalog of Arabic datasets (Alyafeai et al., 2021; Altaher et al., 2022). Each dataset on Masader has a metadata field for the variants of Arabic included. We discarded the datasets that only included MSA samples, and manually inspected the remaining 151. After identifying 28 potential datasets that satisfy the criteria above, we contacted the authors of the datasets that do not have the individual annotations publicly released. Eventually, we had 15 datasets to analyze, listed in Table 1, covering: Offensive Text Classification, Hate Speech Detection, Sarcasm Detection, Sentiment Analysis, Speech Act Detection, Stance Detection, and Dialect Identification.

Analysis For each dataset, we compute the Arabic Level of Dialectness (ALDi) score for each annotated sample (sentence) using the Sentence-ALDi model (Keleg et al., 2023), which returns a score from 0 (MSA/non-dialectal) to 1 (strongly dialectal). To investigate the effect of ALDi on annotator agreement, we bin the samples by their ALDi score into 10 bins of width 0.1. We compute *% full agree*, the percentage of samples in that bin for which all the annotators agreed on a single label. We employ Pearson’s correlation coefficient to analyze the relation between ALDi (represented by each bin’s midpoint ALDi score) and *% full agree*, and also report the slope of the best-fitting line as a measure of the effect size.⁴ As aforementioned, our initial hypothesis is that *% full agree* negatively correlates with high ALDi scores.

3 Results and Discussion

We use scatter plots to visualize the relation between *% full agree* and ALDi on the studied datasets, as shown in Figure 1. Additionally, the histograms of samples across the different bins indicate the dialectal content within the dataset. As per Table 1, 6 datasets out of the 15 have more than 50% of the samples with ALDi scores less than 0.1, which are expected to be written in MSA. However, we found that the overall trends depicted

⁴The exact values of the slopes and correlation coefficients depend on the number of bins. However, we got similar qualitative results on using 4 or 20 equal-width bins. 10 bins are enough to check if trends are non-linear while keeping a reasonable number of samples in the smallest bins. We also fitted logistic regression (*logreg*) models using ALDi as a continuous variable and a binary outcome *Full Agreement (Yes/No)* for each sample. Both analysis tools reveal similar patterns (See Appendix §C) but the binning method provides useful additional visualization.

Dataset	Task (# labels)	% ALDi<0.1	Description
Deleted Comments Dataset (DCD) (Mubarak et al., 2017)	Offensive (3)	62.57%	About 32K deleted comments from aljazeera.com . Confidence scores for the majority vote of 3 annotations are provided.
MPOLD (Chowdhury et al., 2020)	Offensive (2)	27.82%	4000 sentences interacting with news sources, sampled from Twitter, Facebook, and YouTube, annotated three times.
YouTube Cyberbullying (YTCB) (Alakrot et al., 2018)	Offensive (2)	10.24%	15,050 comments and replies to 9 YouTube videos labeled by 3 annotators (Iraqi, Egyptian, Libyan).
ASAD (Alharbi et al., 2021)	Sentiment (3)	35.63%	95,000 tweets with a skewed representation toward the Gulf area and Egypt.
ArSAS (Elmadany et al., 2018)	Sentiment (4) Speech Act (6)	57.45%	21,064 tweets related to a pre-specified set of entities or events, with confidence scores for the majority votes across three annotations per sample.
ArSarcasm-v1 (Abu Farha and Magdy, 2020)	Dialect (5) Sarcasm (2) Sentiment (4)	57.44%	10,547 tweets, sampled from two different Sentiment Analysis datasets: ATSD (Nabil et al., 2015), SemEval2017 (Rosenthal et al., 2017), reannotated for Sentiment, Dialect, and Sarcasm.
Mawqif (Alturayef et al., 2022)	Sarcasm (2) Sentiment (3) Stance (3)	58.04% 58.04% 57.99%	4,121 tweets about "COVID-19 vaccine", "digital transformation", or "women empowerment" annotated separately for stance and sentiment/sarcasm till the label confidence reaches 0.7 (min. 3 annotators) or 7 annotators label the sample.
iSarcasm's test set (Abu Farha et al., 2022)	Dialect (5) Sarcasm (2)	30.5%	200 sarcastic sentences provided by crowdsourced authors and 1200 non-sarcastic tweets from ArSarcasm-v2 (Abu Farha et al., 2021) reannotated 5 times.
DART (Alsarsour et al., 2018)	Dialect (5)	0.8%	24,279 tweets with distinctive dialectal terms annotated three times for the dialectal region. Samples of complete disagreement are not in the released dataset.

Table 1: The datasets included in our study. All datasets have three annotations per sample, except for iSarcasm (5 annotations/sample) and Mawqif (3 or more annotations/sample). For the labels used in each dataset and the proportion of each label, see Table A1. For some datasets, there is a discrepancy between the number of samples listed in the paper and the raw data files with individual labels (See §A of the Appendix).

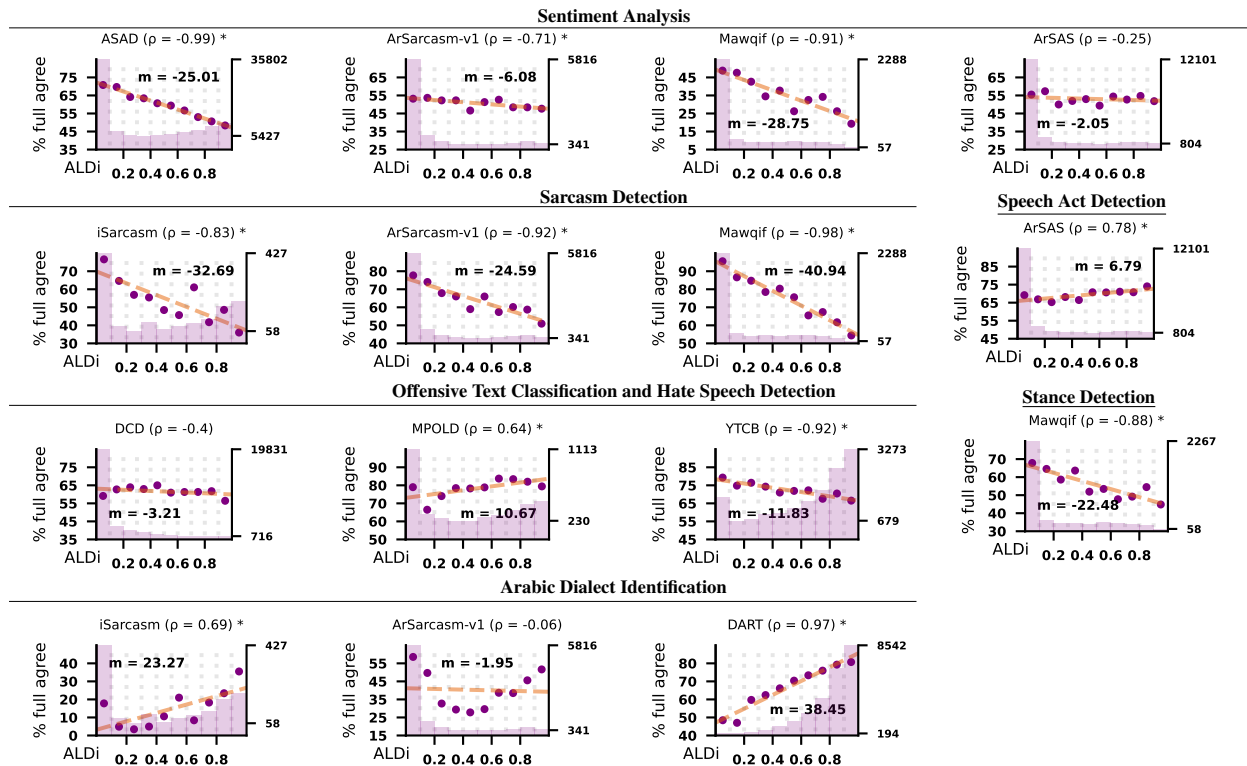


Figure 1: Scatter plots showing the relationship between binned ALDi scores (x-axis) and the percentage of samples with full annotator agreement (y-axis). The histogram represents the # of samples per bin (with min and max values for any bin labeled on the right-hand axis). The slope of the best-fitting line (m) is shown, and to enable visual comparison of slopes, all plots have the same y-axis scale (possibly shifted up or down).

Note: Statistically significant ($p < 0.05$) correlation coefficients (ρ) are marked with *.

in Figure 1 will not be affected if we discard these samples with low ALDi scores and only focus on the rest.

For non-DI tasks, ALDi negatively correlates with agreement. Inspecting the trends depicted in Figure 1, strong negative Pearson's correlation

whole sentence. Consider the first offensive sentence in the Table. The presence of the MSA insult *يا حشرة* (you insect) is enough to guess that the sentence is offensive, even if the remaining segment is not fully intelligible. Lastly, note that agreeing on a label does not imply it is accurate, especially when relying on cues for annotation.

5 Conclusion and Recommendation

Factors such as task subjectivity and vague guidelines could cause disagreement between annotators. For Arabic, we demonstrate that the Arabic level of dialectness of a sentence (ALDi), automatically estimated using the Sentence-ALDi model (Keleg et al., 2023), is an additional overlooked factor.

Analyzing 15 datasets, we find strong evidence of a negative correlation between ALDi and the full annotator agreement scores for 8 of the 12 non-Dialect Identification datasets. Moreover, for the 3 Dialect Identification datasets, we find that annotators have higher agreement scores for samples of higher ALDi scores, which by definition would have more dialectal features. The combination of more dialectal features in a sentence is more probable to be distinctive of a specific dialect.

Previous research recommended routing samples to native speakers of the different Arabic dialects for higher annotation quality. Our analysis indicates that a large proportion of 6 datasets are samples with ALDi scores < 0.1 , which are expected to be MSA samples that can be routed to speakers of any Arabic dialect. Moreover, the lower agreement scores for samples with high ALDi scores show that extra care should be given to these samples. Dataset creators should prioritize routing high-ALDi samples to native speakers of the dialects of these samples, for which the dialects can be automatically identified at higher accuracy as these samples have more dialectal cues.

Limitations

The trends we report validate our hypothesis. However, more thorough analyses need to be done to understand how ALDi affects each task given its unique nature. Knowing the demographic information about the annotators might have allowed for revealing deeper insights into how speakers of specific Arabic dialects understand samples from other dialects. However, this would have required running a controlled experiment re-annotating the 15 datasets, which we hope future work will attempt.

Another potential extension to this work is to analyze the interannotator disagreement on annotating dialectal data for token-level tasks. To the best of our knowledge, all the publicly available token-level Arabic datasets are built by carefully selecting samples written in specific dialects and recruiting native speakers of each of these dialects to perform the annotation, after closely training them. However, even if a multi-dialect token-level dataset is annotated by randomly assigning the samples to speakers of different dialects, the analysis would require a new model to estimate the level of dialectness on the token level, since the *Sentence-ALDi* model used here works at the sentence level.

Lastly, we acknowledge that there are multiple reasons for the annotators to disagree, which include the task's subjectivity, the annotators' background, and their worldviews (Uma et al., 2021). However, these factors would have less impact on the annotators' disagreement if a sample is not fully intelligible.

Acknowledgments

This work could not have been done without the help of the datasets' creators who have kindly agreed to share the individual labels for their datasets' samples. Thanks, Ibrahim Abu Farha, Nora Alturayef, and Manal Alshehri. We also thank Hamdy Mubarak, Nuha Albadi, Nedjma Ousidhoum, and Hala Mulki for trying to help with finding the individual annotator labels for some of their datasets. Lastly, we appreciate the efforts of the anonymous ARR reviewers, action editors, and area chairs. Thanks for the insightful discussions and valuable suggestions.

This work was supported by the UKRI Centre for Doctoral Training in Natural Language Processing, funded by the UKRI (grant EP/S022481/1) and the University of Edinburgh, School of Informatics.

References

- Muhammad Abdul-Mageed, AbdelRahim Elmadany, Chiyu Zhang, El Moatez Billah Nagoudi, Houda Bouamor, and Nizar Habash. 2023. [NADI 2023: The fourth nuanced Arabic dialect identification shared task](#). In *Proceedings of ArabicNLP 2023*, pages 600–613, Singapore (Hybrid). Association for Computational Linguistics.
- Ibrahim Abu Farha and Walid Magdy. 2020. [From Arabic sentiment analysis to sarcasm detection: The ArSarcasm dataset](#). In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Process-*

- ing Tools, with a Shared Task on Offensive Language Detection*, pages 32–39, Marseille, France. European Language Resource Association.
- Ibrahim Abu Farha and Walid Magdy. 2022. **The effect of Arabic dialect familiarity on data annotation**. In *Proceedings of the The Seventh Arabic Natural Language Processing Workshop (WANLP)*, pages 399–408, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Ibrahim Abu Farha, Silviu Vlad Oprea, Steven Wilson, and Walid Magdy. 2022. **SemEval-2022 task 6: iSarcasmEval, intended sarcasm detection in English and Arabic**. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 802–814, Seattle, United States. Association for Computational Linguistics.
- Ibrahim Abu Farha, Wajdi Zaghrouani, and Walid Magdy. 2021. **Overview of the WANLP 2021 shared task on sarcasm and sentiment detection in Arabic**. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 296–305, Kyiv, Ukraine (Virtual). Association for Computational Linguistics.
- Azalden Alakrot, Liam Murray, and Nikola S. Nikolov. 2018. **Dataset construction for the detection of anti-social behaviour in online communication in Arabic**. *Procedia Computer Science*, 142:174–181. Arabic Computational Linguistics.
- Basma Alharbi, Hind Alamro, Manal Alshehri, Zuhair Khayyat, Manal Kalkatawi, Inji Ibrahim Jaber, and Xiangliang Zhang. 2021. **Asad: A twitter-based benchmark Arabic sentiment analysis dataset**.
- Israa Alsarsour, Esraa Mohamed, Reem Suwaileh, and Tamer Elsayed. 2018. **DART: A large dataset of dialectal Arabic tweets**. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Yousef Altaher, Ali Fadel, Mazen Alotaibi, Zaid Alyazidi, et al. 2022. **Masader Plus: A new interface for exploring +500 Arabic NLP datasets**. *arXiv preprint arXiv:2208.00932*.
- Nora Saleh Alturayef, Hamzah Abdullah Luqman, and Moataz Aly Kamaleldin Ahmed. 2022. **Mawqif: A multi-label Arabic dataset for target-specific stance detection**. In *Proceedings of the The Seventh Arabic Natural Language Processing Workshop (WANLP)*, pages 174–184, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Zaid Alyafeai, Maraim Masoud, Mustafa Ghaleb, and Maged S. Al-shaibani. 2021. **Masader: Metadata sourcing for Arabic text and speech data resources**.
- As-Said Muhámmad Badawi. 1973. *Mustawayat al-arabiyya al-muasira fi Misr*. Dar al-maarif.
- A. Bergman and Mona Diab. 2022. **Towards responsible natural language annotation for the varieties of Arabic**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 364–371, Dublin, Ireland. Association for Computational Linguistics.
- Business for Social Responsibility. 2022. **Human rights due diligence of meta’s impacts in Israel and Palestine in may 2021**. <https://about.fb.com/wp-content/uploads/2022/09/Human-Rights-Due-Diligence-of-Metas-Impacts-in-Israel-and-Palestine-in-May-2021.pdf>.
- Shammur Absar Chowdhury, Hamdy Mubarak, Ahmed Abdelali, Soon-gyo Jung, Bernard J. Jansen, and Joni Salminen. 2020. **A multi-platform Arabic news comment dataset for offensive language detection**. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6203–6212, Marseille, France. European Language Resources Association.
- AbdelRahim A. Elmadany, Hamdy Mubarak, and Walid Magdy. 2018. **An Arabic speech-act and sentiment corpus of tweets**. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. European Language Resources Association (ELRA). The 3rd Workshop on Open-Source Arabic Corpora and Processing Tools, OSACT3 ; Conference date: 08-05-2018.
- Joseph L. Fleiss. 1971. **Measuring nominal scale agreement among many raters**. *Psychological bulletin*, 76(5):378.
- Nizar Y. Habash. 2010. *Introduction to Arabic natural language processing*, 1 edition, volume 3 of *Synthesis Lectures on Human Language Technologies*. Morgan and Claypool Publishers.
- Amr Keleg, Sharon Goldwater, and Walid Magdy. 2023. **ALDi: Quantifying the Arabic level of dialectness of text**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10597–10611, Singapore. Association for Computational Linguistics.
- Amr Keleg and Walid Magdy. 2023. **Arabic dialect identification under scrutiny: Limitations of single-label classification**. In *Proceedings of ArabicNLP 2023*, pages 385–398, Singapore (Hybrid). Association for Computational Linguistics.
- Giovanni Molina, Fahad AlGhamdi, Mahmoud Ghoneim, Abdelati Hawwari, Nicolas Rey-Villamizar, Mona Diab, and Tamar Solorio. 2016. **Overview for the second shared task on language identification in code-switched data**. In *Proceedings of the Second Workshop on Computational Approaches to Code Switching*, pages 40–49, Austin, Texas. Association for Computational Linguistics.
- Hamdy Mubarak and Kareem Darwish. 2016. **Demographic surveys of Arab annotators on CrowdFlower**. In *Proceedings of ACM WebSci16 Workshop “Weaving Relations of Trust in Crowd Work: Transparency and Reputation across Platforms*.

- Hamdy Mubarak, Kareem Darwish, and Walid Magdy. 2017. [Abusive language detection on Arabic social media](#). In *Proceedings of the First Workshop on Abusive Language Online*, pages 52–56, Vancouver, BC, Canada. Association for Computational Linguistics.
- Mahmoud Nabil, Mohamed Aly, and Amir Atiya. 2015. [ASTD: Arabic sentiment tweets dataset](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2515–2519, Lisbon, Portugal. Association for Computational Linguistics.
- Helene Olsen, Samia Touileb, and Erik Velldal. 2023. [Arabic dialect identification: An in-depth error analysis on the MADAR parallel corpus](#). In *Proceedings of ArabicNLP 2023*, pages 370–384, Singapore (Hybrid). Association for Computational Linguistics.
- Nedjma Ousidhoum, Zizheng Lin, Hongming Zhang, Yangqiu Song, and Dit-Yan Yeung. 2019. [Multilingual and multi-aspect hate speech analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4675–4684, Hong Kong, China. Association for Computational Linguistics.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. [SemEval-2017 task 4: Sentiment analysis in Twitter](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 502–518, Vancouver, Canada. Association for Computational Linguistics.
- Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi, and Noah A. Smith. 2022. [Annotators with attitudes: How annotator beliefs and identities bias toxic language detection](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5884–5906, Seattle, United States. Association for Computational Linguistics.
- 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.
- Howard Wainer, Marc Gessaroli, and Monica Verdi. 2006. [Visual revelations](#). *CHANCE*, 19(1):49–52.

A Detailed Description of the Datasets

We noticed some discrepancies between the number of samples reported in the papers and the number of samples in the corresponding raw datasets. Despite following any filtration steps described in the papers, some of the datasets had more samples than the ones in the publicly released version, as indicated in Table A1. Additionally, the *ArSarcasm-v1*, *Mawqif (Stance Task)*, *Mawqif (Sentiment/Sarcasm Tasks)*, and *ASAD* had 516, 170, 151, 191

samples with less than 3 annotations respectively, that we decided to discard from our analysis.

Conversely, we decided to discard the MLMA dataset (Ousidhoum et al., 2019) for which the authors shared with us some of the raw annotations files. The number of samples in these files was too small compared to the number of samples in the public dataset with majority-vote labels. We also discarded another dataset, for which there was a significant discrepancy between the released dataset and the paper’s description of the dataset.

B Code-mixing between MSA and DA

Researchers distinguish between Modern Standard Arabic (MSA), and Dialectal Arabic (DA) (Habash, 2010). However, MSA and DA do not exist in isolation, and Arabic speakers sometimes code-mix between terms that can be considered to belong to MSA and others considered to be part of a variety of DA. Notably, some terms can be considered to belong to both MSA and a variety of DA, and even using the surrounding context may not be enough for disambiguation (Molina et al., 2016).

Badawi (1973) recognizes five levels of Arabic used in Egypt, that can be categorized according to the amount of code-mixing in addition to the dialectness of the terms/phrases used. The *Sentence-ALDi* model, developed by Keleg et al. (2023), estimates the level of dialectness of Arabic sentences, which provides an automatic proxy to distinguish between Arabic sentences according to how they diverge from MSA. We used the *Sentence-ALDi* model to study the relation between the ALDi score and the agreement between the annotators for 15 Arabic datasets.

C Discussion about the Analysis

As described in §2, each dataset’s samples were split into 10 bins of equal width according to their respective ALDi scores. Afterward, the correlation between each bin’s midpoint ALDi score and the percentage of samples having full agreement *% full agree* was computed. For each bin, *% full agree* represents the Maximum Likelihood Estimation (MLE) for the probability that all the annotators agree on the same label for the samples of this bin.

Inability to use Interannotator Agreement metrics for some datasets Automated metrics such as Fleiss’ Kappa (Fleiss, 1971) attempt to measure the Interannotator Agreement (IAA) while accounting for the random agreement/disagreement

Dataset	Task (# labels)	Labels	Distribution of Majority-vote Labels	Dataset/Paper Discrepancy
Deleted Comments Dataset (DCD) (Mubarak et al., 2017)	Offensive (3)	Confidence	Offensive (80.31%) Clean (17.76%) Obscene (1.58%) No Majority (0.35%)	-
MPOLD (Chowdhury et al., 2020)	Offensive (2)	Individual	Non-Offensive (83.12%) Offensive (16.88%)	-
YouTube Cyberbullying (YTCB) (Alakrot et al., 2018)	Offensive (2)	Individual	Not (61.38%) HateSpeech (38.62%)	-
ASAD (Alharbi et al., 2021)	Sentiment (3)	Individual	Neutral (67.83%) Negative (15.33%) Positive (15.19%) No Majority (1.65%)	The authors shared with us the raw annotation file of which we analyze 100,484 samples with three annotations or more, as opposed to the 95,000 in the released dataset.
ArSAS (Elmadany et al., 2018)	Sentiment (4)	Confidence	Negative (35.38%) Neutral (33.45%) Positive (20.51%) No Majority (6.07%) Mixed (4.59%)	-
	Speech Act (6)	Confidence	Expression (55.07%) Assertion (38.63%) Question (3.32%) No Majority (1.81%) Request (0.67%) Recommendation (0.31%) Miscellaneous (0.18%)	
ArSarcasm-v1 (Abu Farha and Magdy, 2020)	Dialect (5)	Individual	msa (67.56%) egypt (19.37%) No Majority (5.83%) gulf (3.61%) levant (3.46%) magreb (0.18%)	The samples in the raw annotation artifact shared by the authors has 10,641 samples, as opposed to the 10,547 samples in the released dataset.
	Sarcasm (2)	Individual	False (84.24%) True (15.7%) No Majority (0.06%)	
	Sentiment (3)	Individual	neutral (49.45%) negative (32.57%) positive (14.58%) No Majority (3.4%)	
Mawqif (Alturayef et al., 2022)	Sarcasm (2)	Individual	No (95.97%) Yes (3.78%) No Majority (0.25%)	The authors annotated the same samples for sentiment/sarcasm and stance separately. This was done across 8 different annotation jobs (4 each), for which the authors shared the raw annotation files with us. The number of samples in these files is 4,093 for sentiment/sarcasm and 4,079 for stance, of which 3,942 and 3,909 have three or more annotations. The released dataset is reported to have 4,100 samples.
	Sentiment (3)	Individual	Positive (41.15%) Negative (31.46%) Neutral (22.68%) No Majority (4.72%)	
	Stance (3)	Individual	Favor (60.5%) Against (27.65%) None (7.7%) No Majority (4.14%)	
iSarcasm’s test set (Abu Farha et al., 2022)	Dialect (5)	Individual	msa (32.29%) nile (31.36%) gulf (16.5%) No Majority (15.79%) levant (2.21%) magreb (1.86%) 0 (82.07%) 1 (17.93%)	The dataset having the individual annotator labels is released as an artifact accompanying the following paper (Abu Farha and Magdy, 2022).
	Sarcasm (2)	Individual		
DART (Alsarsour et al., 2018)	Dialect (5)	Proportion	GLF (24.27%) EGY (21.69%) IRQ (21.64%) LEV (16.22%) MGH (16.18%)	-

Table A1: A detailed description of the distribution of the majority-vote labels and the data/paper discrepancies in the datasets with individual annotator labels included in our study.

Note 1: *No Majority* means that multiple labels have the same majority number of votes for Individual/Proportion labels, and Confidence < 0.5 otherwise.

Note 2: Some of the samples of the *ASAD*, *ArSarcasm-v1*, *Mawqif* datasets have more than 3 annotations, despite the fact the former two are supposed to have only three annotations per sample.

between annotators. In principle, it might be possible to perform a version of our analysis using Fleiss’ Kappa rather than % full agree as the dependent variable. However, computing Fleiss’ Kappa would require knowledge of the individual annotations for each sample. Such annotations are not available for the ArSAS (Sentiment/Speech Act), DART, and DCD datasets as described in Table A1. Since we wanted to include as many datasets as possible, we used % full agree instead.

Logistic regression as an alternative analysis tool Binning the data leads to a loss of analytical information which might impact the results of the analysis, especially if implausible bins’ boundaries

are used (Wainer et al., 2006).

Logistic regression with binary outcomes is an alternative analysis that alleviates the limitations of binning. Each sample has a continuous ALDi score as the independent variable, and a binary outcome *Full Annotator Agreement (Yes/No)*. After fitting a logistic regression model to predict the binary outcome, the coefficient of the ALDi variable measures the impact of ALDi on the odds of full agreement. If this coefficient is negative, then the odds of full annotator agreement decrease as the ALDi score increases.

Figure C1 demonstrates the probability of full agreement of each dataset, in addition to the coefficient of the ALDi score with its 95% confi-

dence interval. For the 8 non-DI datasets with $Coef_{ALDi} < -0.2$, the coefficients can be considered to be statistically significant since the confidence interval does not include zero.

Both analysis tools (correlation analysis and logistic regression) achieve similar results. The same 8 non-DI datasets—ASAD, ArSarcasm-v1 (Sentiment/Sarcasm), Mawqif (Sentiment/Sarcasm/Stance), iSarcasm, and YTCB—have significantly strong negative correlation coefficients as in Figure 1, and statistically significant coefficients for the ALDi variable which are less than -0.2. However, binning the data allows for visualizing the % full agreement as a scatter plot, which can reveal whether the relation between ALDi and the agreement is linear or not, in addition to having a visual way for determining how well the best-fitting line models the data.

Impact of data skewness MSA samples are over-represented in some of the considered datasets. However, this is generally unproblematic for the analysis, so we opted not to discard the MSA samples. For the method described in Section 2, the samples of each bin are independently used to estimate the MLE of full agreement between annotators. Therefore, the over-representation of MSA samples in some datasets does not impact our analysis.

D Trends by Class Label

As mentioned in §4, Figures D2, D3, D4, D5, and D6 visualize the impact of ALDi on the annotator agreement after splitting the samples according to their majority-vote labels. We acknowledge that the number of samples in the bins for some classes is not enough to draw concrete conclusions (e.g., samples with high ALDi scores for the *Neutral* class of the *ArSAS*, and *Mawqif* datasets as per Figure D3).

E The Rising Trend of ArSAS

The *ArSAS* dataset stands out as a dataset with a rising trend for the *Speech Act Detection* task and a falling trend for the *Sentiment Analysis* task. Samples of *ArSAS* were jointly annotated for their sentiment and speech act. Despite having 6 different speech acts, which would arguably make speech act detection harder than sentiment analysis, the *Assertion* and *Expression* classes represent 95% of the samples. Looking at their respective trends

shown in Figure D5, the two acts show two different behaviors. Most of the assertive samples have ALDi scores < 0.2 (arguably, all are MSA ones). Moreover, the number of *Assertion* samples with high ALDi scores is not enough to estimate the % full agree for their respective bins. Conversely, the *Expression* act shows higher agreement as the ALDi score increases.

The creators of *ArSAS* noticed that most of the *Assertion* samples were annotated as *Neutral*, while most of the *Expression* samples had polarized sentiment (mostly *Negative*). The annotators might have treated the *Assertion* class as the act for *Objective* sentences, while treating *Expression* as the act for *Subjective* sentences. This is arguably easier than sentiment analysis which might explain why annotators agree more on the Speech Act label than the Sentiment label for the *ArSAS* dataset. Further analysis is required to explain the trends of this dataset.

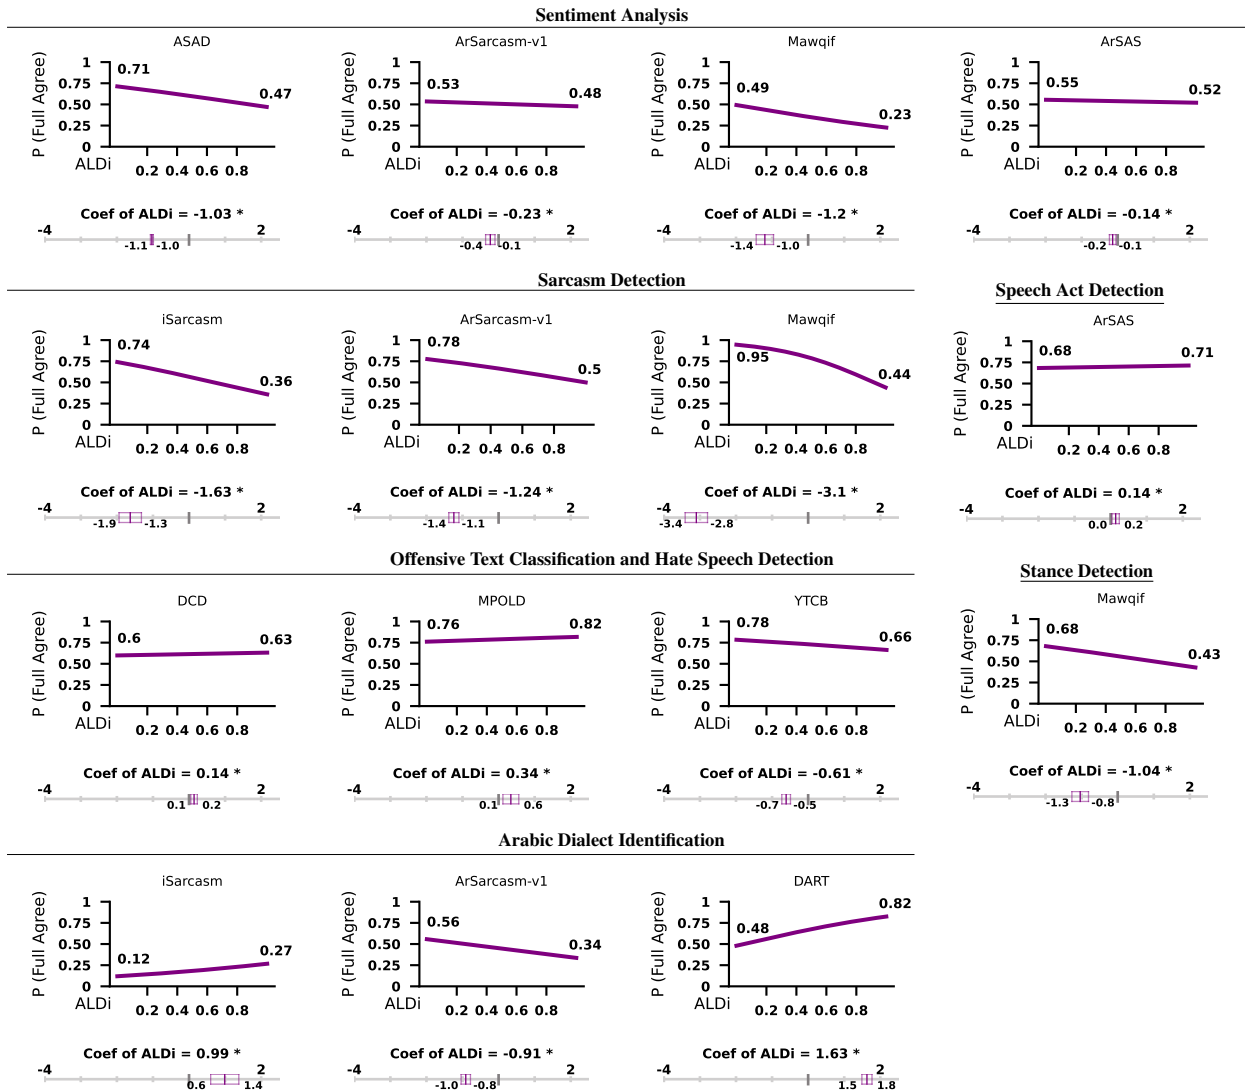


Figure C1: For each dataset, plots show the estimated probability of *full agreement* according to each dataset's fitted logistic regression model. Under each plot, the coefficient of ALDi with its 95% confidence interval is visualized. Nearly all datasets (marked with *) have confidence intervals that do not include zero, meaning the effect of ALDi is statistically significant at $p < 0.05$. Negative coefficients indicate that higher ALDi scores predict lower agreement.

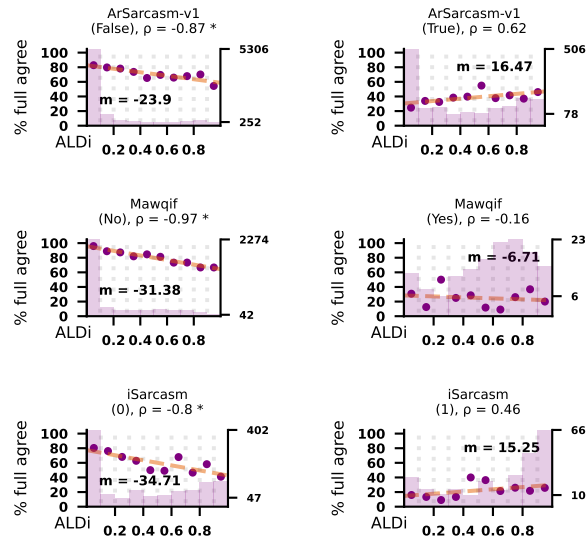


Figure D2: The trends for the classes of the Saracasm Detection datasets. Statistically significant correlation coefficients (ρ) are marked with *.

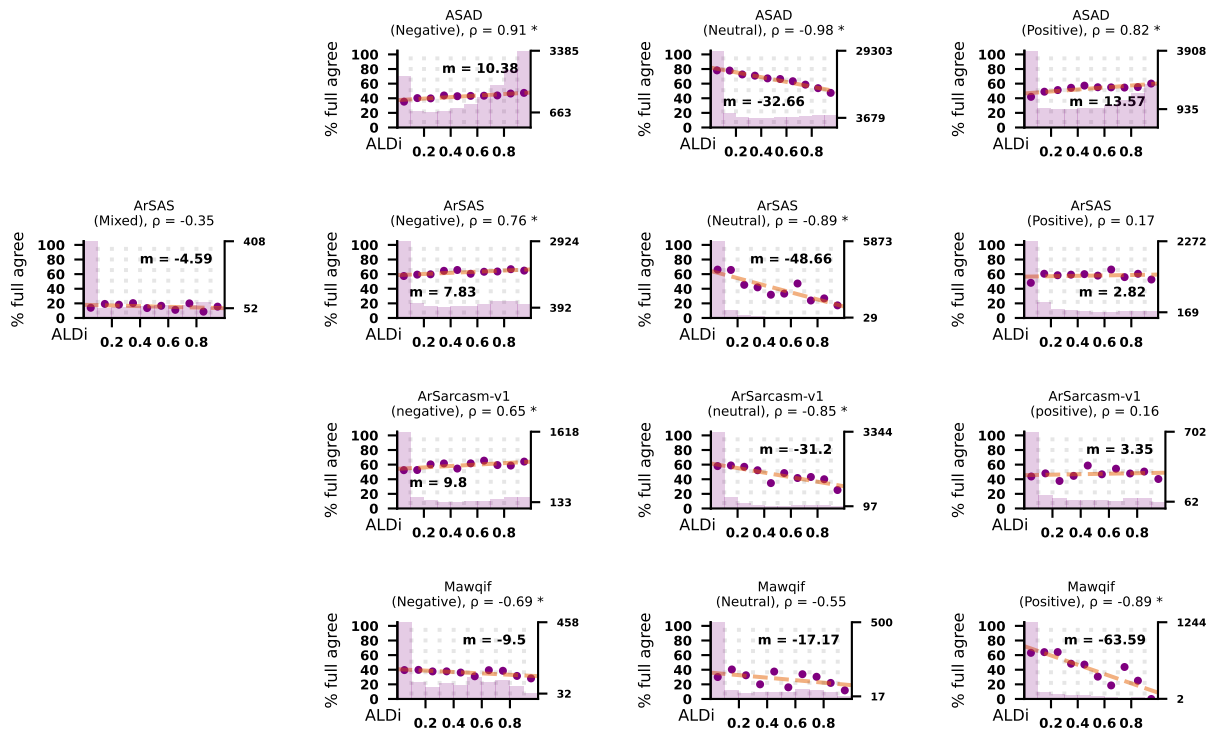


Figure D3: The trends for the classes of the Sentiment Analysis datasets. Statistically significant correlation coefficients (ρ) are marked with *.

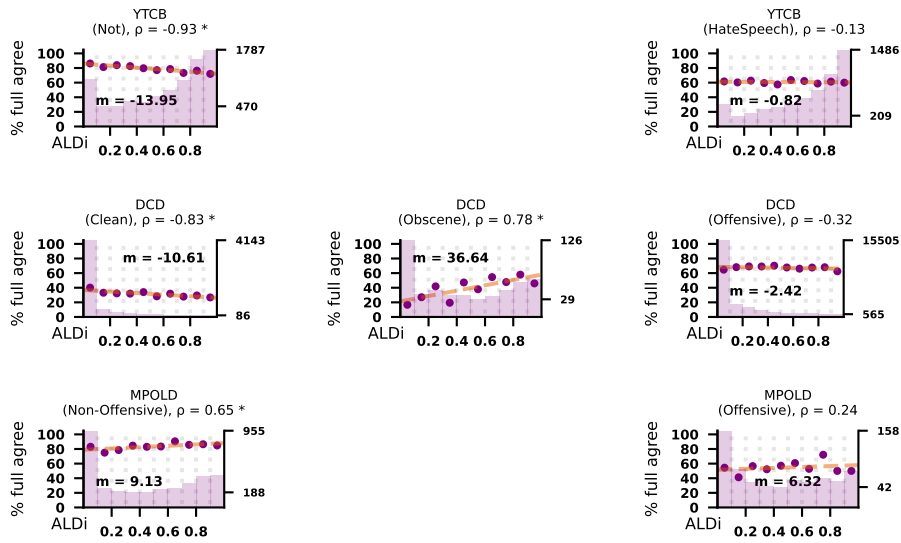


Figure D4: The trends for the classes of the Offensive Text Classification and Hate Speech datasets. Statistically significant correlation coefficients (ρ) are marked with *.




Figure D5: The trends for the *Assertion* and *Expression* labels of the ArSAS dataset, which represent 95% of the dataset samples. Statistically significant correlation coefficients (ρ) are marked with *.




Figure D6: The trends for the classes of Mawqif's Stance dataset. Statistically significant correlation coefficients (ρ) are marked with *.