

Efficient Annotator Reliability Assessment with EffiARA

Owen Cook, Jake Vasilakes, Ian Roberts and Xingyi Song

School of Computer Science,
University of Sheffield

Correspondence: {oscook1, x.song}@sheffield.ac.uk

Abstract

Data annotation is an essential component of the machine learning pipeline; it is also a costly and time-consuming process. With the introduction of transformer-based models, annotation at the document level is increasingly popular; however, there is no standard framework for structuring such tasks. The EffiARA annotation framework is, to our knowledge, the first project to support the whole annotation pipeline, from understanding the resources required for an annotation task to compiling the annotated dataset and gaining insights into the reliability of individual annotators as well as the dataset as a whole. The framework’s efficacy is supported by two previous studies: one improving classification performance through annotator-reliability-based soft-label aggregation and sample weighting, and the other increasing the overall agreement among annotators through removing identifying and replacing an unreliable annotator. This work introduces the EffiARA Python package and its accompanying webtool, which provides an accessible graphical user interface for the system. We open-source the EffiARA Python package at <https://github.com/MiniEggz/EffiARA> and the webtool is publicly accessible at <https://effiara.gate.ac.uk>.

1 Introduction

Labelled data is the foundation of model training and evaluating downstream tasks in machine learning models. However, data annotation is often an expensive and time-consuming process, significantly affecting the quality of model training. Obtaining annotations from experts is ideal, but this expertise is often logistically and financially costly.

Crowd-sourcing platforms such as Amazon’s Mechanical Turk¹ and CrowdFlower (now Figure-Eight)² provide a cheaper alternative by using non-

¹<https://www.mturk.com/>

²<https://www.appen.com/ai-data/data-annotation>

expert annotators; this generally results in lower quality annotations with higher levels of inter-annotator disagreement (Nowak and Ruger, 2010). Effectively collecting, evaluating and managing annotator disagreement is essential in addressing challenges of data quality.

We introduce the EffiARA (Efficient Annotator Reliability Assessment) framework, which supports annotation quality assessment and management throughout the annotation process, allowing users to:

- **Distribute** data points to annotators;
- **Generate labels** from the annotations of each annotator;
- **Assess agreement** among annotators;
- **Assess annotator reliability**;
- **Redistribute** data points to obtain the desired level of agreement;
- **Generate aggregated labels** at the data point level, taking either a soft- or hard-label approach.

To our knowledge, no existing annotation framework provides systematic support for annotator workload allocation which can then be used to estimate the cost of the annotation project. This, in addition to the set of functionalities surrounding the annotation process, makes the EffiARA annotation framework a unique solution for structuring data annotation and modelling annotators.

Additionally, by aggregating annotators’ labels for each data point, tempered by measures of annotator reliability, we can obtain a consensus that better reflects the “true” label distribution. Annotator reliability can also be used to dynamically weight individual data points during model training to ensure that the model prioritises reliable annotations (Cook et al., 2025a).

2 Related Work

2.1 Annotation Frameworks

There have been many attempts to formalise the annotation process for a number of annotation tasks and a range of tools are available. Many frameworks focus on sequence-labelling tasks such as POS tagging and named-entity recognition (Bird and Liberman, 2001; Cornolti et al., 2013; Bontcheva et al., 2013; Lin et al., 2019). More recently, with the introduction of pre-trained LLMs capable of document-level processing, document annotation tools and frameworks have been created, such as GATE Teamware 2 (Wilby et al., 2023). INCEpTION (Klie et al., 2018; De Castilho et al., 2024) handles various aspects of annotator management, including workload distribution, as well as the annotation itself, with a customisable UI, supporting span and document annotation. It also enables active learning but does not aim to model annotator or dataset reliability explicitly. A number of annotation frameworks are task-specific, aiming to provide a set of guidelines and tools for following them, for example event ordering (Cassidy et al., 2014), biodiversity information extraction (Lücking et al., 2022), and surgical video analysis (Meireles et al., 2021).

2.2 Annotator Agreement

Agreement among annotators is often used to assess the quality of a dataset. Commonly used metrics include Scott’s Pi (Scott, 1955), Cohen’s Kappa (Cohen, 1960), Fleiss’ Kappa (Fleiss, 1971), and Krippendorff’s alpha (Krippendorff, 1970). For each metric, there are various interpretations and accepted agreement thresholds used to determine the reliability of a dataset (Krippendorff, 2018; Landis and Koch, 1977). Obtaining datasets where this agreement threshold is met, particularly in scenarios with non-expert annotators such as crowd-sourcing, is challenging and costly (Hsueh et al., 2009; Nowak and Rüger, 2010).

2.3 Annotation Aggregation

Rather than ensuring acceptable levels of agreement, many approaches use disagreement as additional information, utilising it to understand the subjectivity of particular data points or the reliability of annotators.

The soft-label approach incorporates a level of subjectivity into aggregated labels for each data point and has been shown to improve both classi-

fication performance and model calibration (Wu et al., 2023; Cook et al., 2025a). Popular methods of label aggregation include majority voting (hard-label only), Dawid and Skene (1979), GLAD (Whitehill et al., 2009), and MACE (Hovy et al., 2013) for categorical data; these methods have been implemented in Python as part of the Crowd-Kit tool (Ustalov et al., 2021).

2.4 Annotator Reliability

Assessing annotator reliability can be used to assess the quality of individual annotators and may be used to understand the quality of data available, remove low-reliability annotators (Cook et al., 2025b), inform the training of machine learning models through aggregating soft labels with reliability (Dawid and Skene, 1979; Wu et al., 2023), or affect the loss function during training (Cao et al., 2023; Cook et al., 2025a). There are different approaches to assessing annotator reliability, such as learning through Expectation Maximisation (Cao et al., 2023) or directly inferring the reliability of an annotator from their agreement with others (Inel et al., 2014; Dumitrache et al., 2018; Cook et al., 2025a,b). Through impacting the generation of soft labels, or directly impacting the training loss, more information about which annotations are more trustworthy is provided, leading to more performant and robust models. An alternative approach to assessing annotator reliability involves comparing annotators’ annotations to a set of gold-standard labels (Barthet et al., 2023); this approach is often used to filter out bad annotators. All three approaches have been shown to improve model performance on classification tasks when compared to methods that trust each annotator equally.

3 EffiARA Python Package

The EffiARA annotation framework structures the annotation process from start to finish. It distributes samples to annotators, generates and aggregates labels, computes inter- and intra-annotator agreement, and assesses annotator reliability. A visual representation of the EffiARA pipeline is provided in Figure 1 and we describe each stage in detail below. The framework has been implemented in Python due to its extensive use in NLP research, increasing its ease of use and integration.

The annotation pipeline is implemented as a set of modular tools in the EffiARA Python package. The source code is available at <https://github.com>.

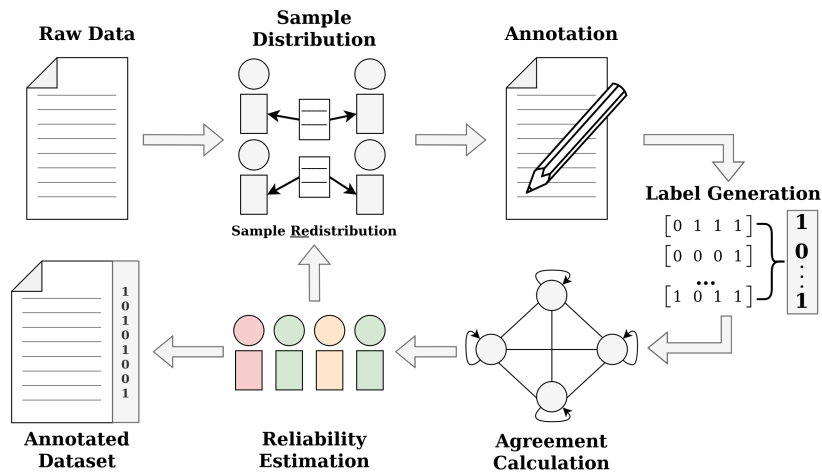


Figure 1: An overview of the EffiARA annotation pipeline, covering sample distribution, annotation, label generation, agreement calculation, reliability estimation, and dataset compilation.

`com/MiniEggz/EffiARA` and the package has been released on PyPi for quick installation: <https://pypi.org/project/effiara/>. Documentation is available here: <https://effiara.readthedocs.io>.

The package relies on a number of core Python libraries. Two fundamental libraries required by the EffiARA framework are NumPy (Oliphant et al., 2006; Harris et al., 2020) and pandas (McKinney et al., 2011), used for efficient mathematical operations on arrays and the manipulation of data.

3.1 Sample Distribution

The first stage in the EffiARA pipeline enables annotation coordinators to estimate resource requirements: how many annotators are needed, how much time is required from each annotator, and how many samples can be produced, given the time and number of annotators. Once resources have been finalised, data points can be distributed among annotators with the EffiARA distribution algorithm, which ensures annotator agreement can be effectively assessed (Cook et al., 2025a).

Both of these functionalities are implemented in the `SampleDistributor` class. We first use SymPy (Meurer et al., 2017) to solve for the missing variable in the resource-understanding equation introduced in Cook et al. (2025a) (Algorithm 1). We then use pandas to split the data into separate DataFrames for each annotator, with one DataFrame containing left-over samples that may be used later.

3.2 Data Annotation

The sample allocations obtained in the previous step can then be used to assign samples to annotators and complete the annotation process using existing tools such as GATE Teamware 2 (Wilby et al., 2023) or Amazon’s Mechanical Turk.

3.3 Label Generation

Label generation involves transforming raw annotations obtained from annotators into numeric encodings compatible with annotator agreement metrics (such as Cohen’s Kappa, Fleiss’ Kappa, Krippendorff’s alpha, or cosine similarity) and model training. These transformations may be at the individual annotator level (for example, transforming first- and second-choice annotations into a categorical distribution), or at the data point level (aggregating annotations from multiple annotators).

As the exact transformations required are often task-specific, the abstract `LabelGenerator` class guides users to implement their own label generation code with three necessary methods:

- `add_annotation_prob_labels` is used to represent each individual’s raw annotations;
- `add_sample_prob_labels` is used to aggregate labels at the data point level, retaining disagreement in a soft label approach;
- `add_sample_hard_labels` aggregates the annotations into a hard label, through methods such as majority voting or taking the maximum probability label from the aggregated soft label.

For annotator agreement calculations, only `add_annotation_prob_labels` must be imple-

mented. To instantiate a class inheriting from `LabelGenerator`, the user must provide a list of annotator names and the label mapping: a dictionary where the key is the value represented in the `DataFrame` and the value is a numeric representation. This enables the extraction of individual annotations and their representation as a distribution across the available classes.

We provide a number of preset label generators: the `DefaultLabelGenerator`, for the cases in which no special label aggregation is necessary; the `EffiLabelGenerator`, mirroring the label generation and aggregation shown in [Cook et al. \(2025a\)](#); the `TopicLabelGenerator`, for multi-label tasks such as topic-extraction ([Cook et al., 2025b](#)); and the `OrdinalLabelGenerator`, used for ordinal annotation tasks where a number of features are labelled on a scale. With the `label_generator.from_annotations` method, the specific class inheriting from `LabelGenerator` is instantiated from the raw annotations, requiring no additional coding from the user.

3.4 Annotator Agreement & Reliability

Once labels for each annotation have been generated, inter- and intra-annotator agreement are calculated using equations introduced in [Cook et al. \(2025a\)](#). Annotator agreement can then be visualised in a 2D or interactive 3D graph, where each node represents an annotator and edges between annotators represent the pairwise agreement between two annotators, with the value next to each node representing an annotator’s agreement with themselves. For cases with many annotators, where a graph could be unwieldy, we also provide a heatmap visualisation, where annotators are ordered by reliability; note that intra-annotator agreement is displayed on the diagonal. Examples of these visualisations are given in [Figure 2](#).

Using these agreement values, annotator reliability can then be calculated, using a combination of an annotator’s intra-annotator agreement and average inter-annotator agreement, weighted by an α parameter controlling the strength of intra-annotator agreement from 0 to 1. The resulting agreement values are centered around 1, enabling the recursive inter-annotator agreement calculation from [Cook et al. \(2025a\)](#). The reliability values can then be accessed and utilised, potentially removing certain annotators from the annotation process ([Cook et al., 2025b](#)). Reliability scores may also be utilised in label aggregation

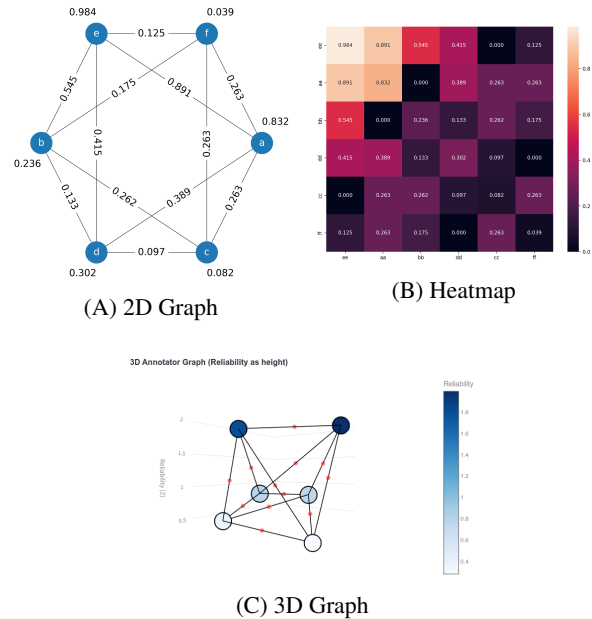


Figure 2: Example agreement visualisations as (A) a 2D graph, (B) a heatmap, and (C) a 3D graph for six annotators (annotations were synthetically generated).

(in a `LabelGenerator`) or used to weight the loss function in model training ([Cook et al., 2025a](#)).

Annotator agreement and reliability is calculated and stored in the `Annotations` class. The `Annotations` class is instantiated with a pandas `DataFrame` representation of the dataset, a `LabelGenerator` object (which will be generated using the `LabelGenerator.from_annotations` function if no instance inheriting from `LabelGenerator` is passed), an agreement metric (defaulting to Krippendorff’s alpha), an overlap threshold, and the reliability alpha.

On instantiating an `Annotations` class, the annotator graph (supported by the `NetworkX` library ([Hagberg et al., 2008](#))) is initialised with each annotator equally reliable. Intra-annotator agreement is first calculated for each annotator node with the `calculate_intra_annotator_agreement` instance method, using data points each user has annotated twice themselves. Inter-annotator agreement is then calculated between each user, utilising the `overlap_threshold` to decide whether there is sufficient overlap between the two annotators to assess agreement. Here, the `pairwise_agreement` function is used as a common interface to the implemented pairwise agreement metrics in the agreement module. Python modules used to handle agreement calculations include the Krippendorff library ([Castro, 2017](#)) for Krippendorff’s alpha and `Scikit-Learn` ([Pedregosa et al., 2011](#)) for Cohen’s

Kappa and Fleiss' Kappa. NumPy and pandas are also used for vector calculations and manipulation of the data to obtain pair annotations.

Once agreement has been calculated among annotators and with themselves, annotator reliability is calculated with a recursive application of the annotator reliability equation until reliability values converge. To ensure convergence, the calculated reliability values are normalised to have a mean of 1 after each iteration. Annotator reliability values can then be accessed through the `get_user_reliability` and `get_reliability_dict` methods.

Inter- and intra-annotator agreement values can also be easily accessed via the graph itself using the NetworkX API and the `__getitem__` method of the Annotations class. The graph and heatmap agreement visualisations shown in Figure 2 utilise Matplotlib (Tosi, 2009) and Seaborn (Waskom, 2021), and they are displayed using the `display_annotator_graph` and `display_agreement_heatmap` methods respectively.

The optional `annotators` and `other_annotators` arguments for the heatmap allow a user to display the agreement between one set of users and another, with the default setting comparing all annotators to one another. This may be useful in cases where you already have a set of reliable annotators or you have a gold-standard set of annotations you would like to compare a set of annotators to.

3.5 Sample Redistribution

In cases where a consensus must be reached on a high proportion of data points, samples may be redistributed among annotators to resolve disagreement. The `SampleRedistributor` provides this functionality. It functions very similarly to the `SampleDistributor` with the additional constraint that an annotator who has already annotated an individual data point will not be reassigned it. Sample redistribution can be done iteratively until the desired level of agreement is reached.

The `SampleRedistributor` inherits from the `SampleDistributor` class, overloading the `distribute_samples` method, applying a round-robin-style allocation using the EffiARA sample distribution variables, ensuring that annotators are not given samples they have already annotated.

3.6 Final Dataset

Once the desired level of agreement has been reached, potentially with the aim of generating gold-standard labels in classification tasks, the final dataset is ready, with annotations tied to annotator identities, allowing for training strategies that utilise the expertise and reliability of individual annotators. Users may utilise the `concat_annotations` method in the `data_generation` module for assistance in merging annotations into the final dataset.

4 EffiARA Webtool

To make the functionalities of the EffiARA package more accessible and quicker to use, we have also released the webtool at <https://effiara.gate.ac.uk>. The webtool allows non-technical experts to run annotation projects and gain insights into annotator agreement and reliability with ease. Even for those comfortable using the Python package, the webtool provides a convenient interface for performing tasks quickly. A system demonstration is available at <https://www.youtube.com/watch?v=KcmQfPiskcY>.

The webtool supports common tasks within the annotation pipeline (excluding the annotation step itself). Finer-grained control and more advanced functionality can be achieved with the Python package, particularly through customisation of modules like the `LabelGenerator`. As the project is open-sourced, technical users are able to make their own modifications and run them as a local web-application or make a pull request to add their additional use-cases. The webtool source code is available at <https://github.com/MiniEggz/EffiARA-webtool>.

The application contains four main workflows:

- *Sample Distribution*. This workflow handles all aspects of distributing samples from an unannotated dataset, including understanding the resources available. The `sample_id` column is added to each data point to allow re-compilation after annotation.
- *Annotation Project*. This workflow is used to generate an annotation project for specific platforms. Currently, project generation for GATE Teamware 2 (Wilby et al., 2023) is supported. Future iterations may include other platforms but this task is most likely solved to some extent by the individual annotation platforms.

- *Dataset Compilation.* Once data annotation is complete, this workflow allows the user to upload a ZIP file containing all annotation CSV files. It supports users in renaming columns, moving all reannotations under the correct columns (beginning with `re_`) and into the correct row (alongside their original annotation of the data point), and merging the annotations from different annotators to create a final dataset ready for analysis.
- *Annotator Reliability.* With the compiled dataset, users can analyse annotator reliability. The user first selects their label generator and they then have full control over the label mapping or they may choose to generate it automatically using the `LabelGenerator.from_annotations` method. Users then choose the desired output: any combination of outputting annotator reliability, the annotator agreement graph (in 2D or interactive 3D) and an annotator heatmap. The workflow also offers a number of options for calculating annotator reliability, such as the agreement metric, the reliability alpha, and the overlap threshold (the minimum number of data points annotated by both annotators to enable agreement assessment); the workflow also offers display configurations for the graphs.

The webtool is built upon the EffiARA Python package and shares the same dependencies. It is implemented using Streamlit (Khorasani et al., 2022) and Plotly (Sievert, 2020) is used to create the interactive 3D annotator agreement and reliability visualisation. The `zipfile` and `tempfile` libraries handle uploads and downloads, ensuring data is deleted once processed.

5 Evaluation

5.1 Case Studies

Two previous works involving dataset creation have annotated data following the EffiARA methodology, creating RUC-MCD (Cook et al., 2025a) and the Chinese News Framing dataset (Cook et al., 2025b). Both studies provide support for the annotation framework.

RUC-MCD. In the work introducing the EffiARA annotation framework (Cook et al., 2025a), utilising reliability scores in the label generation and model training stages was shown to improve

classification performance. Applying a soft-label approach, using TwHIN-BERT-Large, assessing reliability with inter-annotator agreement only, intra-annotator agreement only, and a combination of both all improved classification performance. Classification performance increased from an F1-macro score of 0.691 to 0.740 using the EffiARA reliability scores calculated using a reliability alpha of 0.5. The dataset used in this study was of low-to-moderate agreement, highlighting the framework’s utility in datasets containing disagreement.

Chinese News Framing. This work utilises the EffiARA reliability scores to identify unreliable annotators during the annotation process, leading to an increased overall level of agreement among annotators, which is highly indicative of data quality (Krippendorff, 2018). By removing the low-reliability annotator and replacing them with an existing high-reliability annotator, the average inter-annotator agreement (measured using Krippendorff’s alpha) increased from 0.396 to 0.465.

5.2 Load Testing

To assess the usability of the application, we also carried out load testing on the web application when hosted locally on a laptop with an Intel i7-6600U @ 3.400GHz and 16GB RAM, meaning upload and download speed were not a factor. Sample distribution remains quick and responsive for a large number of samples, taking less than a second for datasets of 100,000 samples. Dataset compilation and processing both scale roughly linearly with respect to dataset size with the tool requiring significantly longer to process datasets containing as many as 100,000 data points. Datasets containing 10,000 data points and under require less than one minute for dataset compilation and dataset processing (including annotator reliability calculation and visualisation rendering). The time taken for each key action in the webtool can be seen in Table 1. While running tasks that take longer, the web application remains responsive.

6 Conclusion and Future Work

In this work, we introduced the EffiARA Python package alongside an accessible web application that provides a graphical interface to the EffiARA annotation framework. EffiARA supports the design, compilation, and reliability assessment of annotation projects at the document level.

| Number of Samples | Sample Distribution | Dataset Compilation | Dataset Processing |
|-------------------|---------------------|---------------------|--------------------|
| 500 | ~0.06s | ~3s | ~3s |
| 1,000 | ~0.06s | ~6s | ~6s |
| 5,000 | ~0.10s | ~30s | ~25s |
| 10,000 | ~0.12s | ~1m | ~45s |
| 100,000 | ~0.5s | ~10m | ~7m 20s |

Table 1: Processing time for each stage at varying dataset sizes. Tests conducted running the webtool locally on a laptop with 16GB RAM and an i7-6600U @ 3.400GHz.

Future development will focus on expanding the range of supported annotation settings, optimising computational performance, and enhancing usability based on user feedback. The package and webtool will be actively maintained, ensuring they remain usable and up-to-date with users’ annotation requirements.

7 Limitations

EffiARA has been developed primarily to support document classification annotation, with the flexibility to accommodate other annotation types. While the design accounts for such flexibility, span annotation functionality, for example, has not yet been implemented and would require technical expertise to integrate. One key future addition may include the integration of Krippendorff’s unitising alpha (Krippendorff et al., 2016) for spans.

Some annotation tasks may require task-specific label generators. Although the framework includes preset label generators, less common tasks may require further customisation. This may present a usability barrier for non-technical users. We aim to address this through future development and community contributions of further task-specific components.

While EffiARA has been tested in situations with disagreement, it has been assumed that there is a ground-truth label for each data point. Future work will investigate the extension of EffiARA to tasks where there may not be a single ground-truth label, but potentially multiple subjective, and equally valid, true labels.

8 Ethical Impact

As EffiARA is an annotation framework, it does not pose direct ethical risks. Annotated data is instrumental in training machine learning models, including those that may be deployed in sensitive or high-impact contexts. Users of the EffiARA annota-

tion framework should remain aware of the broader ethical impact of their annotation projects and consider them before undertaking such projects.

9 Acknowledgements

This work was supported by the UK’s innovation agency (Innovate UK) grant 10039055 (approved under the Horizon Europe Programme as vera.ai, EU grant agreement 101070093) vera.ai. JV is supported by the European Media and Information Fund (EMIF) managed by the Calouste Gulbenkian Foundation³ under the “Supporting Research into Media, Disinformation and Information Literacy Across Europe” call (ExU – project number: 291191).⁴ We also acknowledge the IT Services at The University of Sheffield, utilising their High Performance Computing clusters in both of our case studies.

References

- Matthew Barthet, Chintan Trivedi, Kosmas Pinitas, Emmanouil Xylakis, Konstantinos Makantasis, Antonios Liapis, and Georgios N Yannakakis. 2023. Knowing your annotator: Rapidly testing the reliability of affect annotation. In *2023 11th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, pages 1–8. IEEE.
- Steven Bird and Mark Liberman. 2001. A formal framework for linguistic annotation. *Speech communication*, 33(1-2):23–60.
- Kalina Bontcheva, Hamish Cunningham, Ian Roberts, Angus Roberts, Valentin Tablan, Niraj Aswani, and Genevieve Gorrell. 2013. Gate teamware: a web-based, collaborative text annotation framework. *Language Resources and Evaluation*, 47:1007–1029.
- Zhi Cao, Enhong Chen, Ye Huang, Shuanghong Shen, and Zhenya Huang. 2023. Learning from crowds with annotation reliability. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2103–2107.
- Taylor Cassidy, Bill McDowell, Nathanael Chambers, and Steven Bethard. 2014. An annotation framework for dense event ordering. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 501–506.

³The sole responsibility for any content supported by the European Media and Information Fund lies with the author(s) and it may not necessarily reflect the positions of the EMIF and the Fund Partners, the Calouste Gulbenkian Foundation and the European University Institute.

⁴exuproject.sites.sheffield.ac.uk

- Santiago Castro. 2017. Fast Krippendorff: Fast computation of Krippendorff’s alpha agreement measure. <https://github.com/pln-fing-udelar/fast-krippendorff>.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Owen Cook, Charlie Grimshaw, Ben Peng Wu, Sophie Dillon, Jack Hicks, Luke Jones, Thomas Smith, Matyas Szert, and Xingyi Song. 2025a. Efficient annotator reliability assessment and sample weighting for knowledge-based misinformation detection on social media. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 3348–3358, Albuquerque, New Mexico. Association for Computational Linguistics.
- Owen Cook, Yida Mu, Xinye Yang, Xingyi Song, and Kalina Bontcheva. 2025b. A dataset for analysing news framing in chinese media. *Proceedings of the International AAAI Conference on Web and Social Media*, 19(1):2402–2412.
- Marco Cornolti, Paolo Ferragina, and Massimiliano Ciaramita. 2013. A framework for benchmarking entity-annotation systems. In *Proceedings of the 22nd international conference on World Wide Web*, pages 249–260.
- Alexander Philip Dawid and Allan M Skene. 1979. Maximum likelihood estimation of observer error-rates using the em algorithm. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 28(1):20–28.
- Richard Eckart De Castilho, Jan-Christoph Klie, and Iryna Gurevych. 2024. Integrating inception into larger annotation processes. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 110–121.
- Anca Dumitrache, Oana Inel, Lora Aroyo, Benjamin Timmermans, and Chris Welty. 2018. Crowdtruth 2.0: Quality metrics for crowdsourcing with disagreement. *arXiv preprint arXiv:1808.06080*.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Aric Hagberg, Pieter J Swart, and Daniel A Schult. 2008. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Laboratory (LANL), Los Alamos, NM (United States).
- Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. 2020. Array programming with numpy. *Nature*, 585(7825):357–362.
- Dirk Hovy, Taylor Berg-Kirkpatrick, Ashish Vaswani, and Eduard Hovy. 2013. Learning whom to trust with mace. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1120–1130.
- Pei-Yun Hsueh, Prem Melville, and Vikas Sindhwani. 2009. Data quality from crowdsourcing: a study of annotation selection criteria. In *Proceedings of the NAACL HLT 2009 workshop on active learning for natural language processing*, pages 27–35.
- Oana Inel, Khalid Khamkham, Tatiana Cristea, Anca Dumitrache, Arne Rutjes, Jelle van der Ploeg, Lukasz Romaszko, Lora Aroyo, and Robert-Jan Sips. 2014. Crowdtruth: Machine-human computation framework for harnessing disagreement in gathering annotated data. In *The Semantic Web–ISWC 2014: 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part II 13*, pages 486–504. Springer.
- Mohammad Khorasani, Mohamed Abdou, and J Hernández Fernández. 2022. Web application development with streamlit. *Software Development*, pages 498–507.
- Jan-Christoph Klie, Michael Bugert, Beto Boullosa, Richard Eckart de Castilho, and Iryna Gurevych. 2018. The INCEpTION platform: Machine-assisted and knowledge-oriented interactive annotation. In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 5–9, Santa Fe, New Mexico. Association for Computational Linguistics.
- Klaus Krippendorff. 1970. Estimating the reliability, systematic error and random error of interval data. *Educational and psychological measurement*, 30(1):61–70.
- Klaus Krippendorff. 2018. *Content analysis: An introduction to its methodology*. Sage publications.
- Klaus Krippendorff, Yann Mathet, Stéphane Bouvry, and Antoine Widlöcher. 2016. On the reliability of unitizing textual continua: Further developments. *Quality & Quantity*, 50(6):2347–2364.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Bill Yuchen Lin, Dong-Ho Lee, Frank F Xu, Ouyu Lan, and Xiang Ren. 2019. Alpacatag: An active learning-based crowd annotation framework for sequence tagging. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*.
- Andy Lücking, Christine Driller, Manuel Stoeckel, Giuseppe Abrami, Adrian Pachzelt, and Alexander Mehler. 2022. Multiple annotation for biodiversity: developing an annotation framework among biology, linguistics and text technology. *Language resources and evaluation*, 56(3):807–855.

- Wes McKinney et al. 2011. pandas: a foundational python library for data analysis and statistics. *Python for high performance and scientific computing*, 14(9):1–9.
- Ozanan R Meireles, Guy Rosman, Maria S Altieri, Lawrence Carin, Gregory Hager, Amin Madani, Nicolas Padoy, Carla M Pugh, Patricia Sylla, Thomas M Ward, et al. 2021. Sages consensus recommendations on an annotation framework for surgical video. *Surgical endoscopy*, 35(9):4918–4929.
- Aaron Meurer, Christopher P Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B Kirpichev, Matthew Rocklin, AMiT Kumar, Sergiu Ivanov, Jason K Moore, Sartaj Singh, et al. 2017. Sympy: symbolic computing in python. *PeerJ Computer Science*, 3:e103.
- Stefanie Nowak and Stefan R ger. 2010. How reliable are annotations via crowdsourcing: a study about inter-annotator agreement for multi-label image annotation. In *Proceedings of the international conference on Multimedia information retrieval*, pages 557–566.
- Travis E Oliphant et al. 2006. *Guide to numpy*, volume 1. Trelgol Publishing USA.
- Fabian Pedregosa, Ga l Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830.
- William A Scott. 1955. Reliability of content analysis: The case of nominal scale coding. *Public opinion quarterly*, pages 321–325.
- Carson Sievert. 2020. *Interactive web-based data visualization with R, plotly, and shiny*. Chapman and Hall/CRC.
- Sandro Tosi. 2009. *Matplotlib for Python developers*. Packt Publishing Ltd.
- Dmitry Ustalov, Nikita Pavlichenko, and Boris Tseitlin. 2021. Learning from crowds with crowd-kit. *arXiv preprint arXiv:2109.08584*.
- Michael L Waskom. 2021. Seaborn: statistical data visualization. *Journal of Open Source Software*, 6(60):3021.
- Jacob Whitehill, Ting-fan Wu, Jacob Bergsma, Javier Movellan, and Paul Ruvolo. 2009. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. *Advances in neural information processing systems*, 22.
- David Wilby, Twin Karmakharm, Ian Roberts, Xingyi Song, and Kalina Bontcheva. 2023. Gate teamware 2: An open-source tool for collaborative document classification annotation. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 145–151.
- Ben Wu, Yue Li, Yida Mu, Carolina Scarton, Kalina Bontcheva, and Xingyi Song. 2023. Don’t waste a single annotation: improving single-label classifiers through soft labels. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5347–5355.