

Analyzing Dataset Annotation Quality Management in the Wild

Jan-Christoph Klie*

Ubiquitous Knowledge Processing Lab
Department of Computer Science and
Hessian Center for AI (hessian.AI)
www.ukp.tu-darmstadt.de

Richard Eckart de Castilho

Ubiquitous Knowledge Processing Lab

Iryna Gurevych

Ubiquitous Knowledge Processing Lab

Data quality is crucial for training accurate, unbiased, and trustworthy machine learning models as well as for their correct evaluation. Recent work, however, has shown that even popular datasets used to train and evaluate state-of-the-art models contain a non-negligible amount of erroneous annotations, biases, or artifacts. While practices and guidelines regarding dataset creation projects exist, to our knowledge, large-scale analysis has yet to be performed on how quality management is conducted when creating natural language datasets and whether these recommendations are followed. Therefore, we first survey and summarize recommended quality management practices for dataset creation as described in the literature and provide suggestions for applying them. Then, we compile a corpus of 591 scientific publications introducing text datasets and annotate it for quality-related aspects, such as annotator management, agreement, adjudication, or data validation. Using these annotations, we then analyze how quality management is conducted in practice. A majority of the annotated publications apply good or excellent quality management. However, we deem the effort of 30% of the studies as only subpar. Our analysis also shows common errors, especially when using inter-annotator agreement and computing annotation error rates.

1. Introduction

Having large, high-quality annotated datasets available is essential for developing, training, evaluating, and deploying reliable machine learning models (Sun et al. 2017; Bender and Friedman 2018; Peters, Ruder, and Smith 2019; Gururangan et al. 2020;

* Corresponding author.

Sambasivan et al. 2021). Annotated datasets are also frequently used in linguistics (Haselbach et al. 2012), language acquisition research (Behrens 2008), bioinformatics (Zeng et al. 2015), healthcare (Suster, Tulkens, and Daelemans 2017), and the digital humanities (Schreibman, Siemens, and Unsworth 2004). Concerning machine learning, recent work has shown, however, that even datasets widely used to train and evaluate state-of-the-art models contain non-negligible proportions of questionable labels. For instance, the CoNLL-2003 (Tjong Kim Sang and De Meulder 2003, named entity recognition) test split has an estimated 6.1% wrongly labeled instances (Reiss et al. 2020; Wang et al. 2019), ImageNet 5.8% (Vasudevan et al. 2022; Northcutt, Athalye, and Mueller 2021, image classification) and TACRED 23.9%, incorrect instances (Stoica, Platanios, and Poczos 2021, relation extraction). GoEmotions (Demszky et al. 2020, sentiment classification) is estimated to contain even up to 30% wrong labels.¹ Using these datasets for machine learning can—among other issues—lead to inaccurate estimates of model performance (Reiss et al. 2020; Vasudevan et al. 2022), generalization failure due to data bias (McCoy, Pavlick, and Linzen 2019), or decreased task performance (Stoica, Platanios, and Poczos 2021; Vădineanu et al. 2022).

Recently, conversational agents and search engines based on large language models trained via instruction tuning have been widely adopted in science and society (Ouyang et al. 2022; Wei et al. 2022). Hence, datasets used for fine-tuning must be factually correct and contain as few biases as possible for the resulting models to be accurate and trustworthy and not to cause misinformation or harm. Benchmark datasets to evaluate their performance and rankings also need to be as accurate as possible to allow fair comparisons.

Proper quality management must be conducted throughout the dataset creation process (as depicted in § 3.1) to produce high-quality datasets. Dataset quality is not only limited to label accuracy but also encompasses aspects such as the quality of the underlying text, annotation scheme, adhering to established practices, or standards for a task, and social or data bias. Quality management encompasses, among others, proper data selection, choice of annotators and training, creating and improving annotation schemes and guidelines, as well as annotator agreement, data validation, and error rate estimation (Hovy and Lavid 2010; Alex et al. 2010; Pustejovsky and Stubbs 2013; Monarch 2021). Even though there exists an extensive body of work that discusses quality management in theory (see § 2), we observe that this knowledge is difficult to find and to consult, as it is scattered across many different sources and usually treated as part of the general annotation process, hence often lacks depth. Also, to the best of our knowledge, no work as of yet has analyzed whether and how these recommendations are applied in practice. Disseminating and analyzing quality management is especially relevant in the context of a growing number of datasets being created and released, which can exhibit the aforementioned discussed dangers of low-quality data collection.

To better understand how quality management is actually performed in practice, we first survey the literature to summarize good practices regarding quality management for dataset creation. Based on *Papers With Code*,² we then collect and annotate a large set of publications (591, of which 314 report human annotation or validation) that introduce new text datasets, and analyze how often and how well the different quality management methods are used. We also analyze the coverage of *Papers With Code* with regard to the ACL anthology, LDC corpora, and shared tasks to validate the representativeness of

1 <https://archive.ph/jQbNM>.

2 <https://paperswithcode.com/datasets>.

our collected dataset. Finally, we summarize our findings and provide suggestions that dataset creators can use to consult and improve their annotation process. To the best of our knowledge, this newly annotated dataset and analysis of annotation good practices is the most extensive and detailed to date. We answer the following research questions:

RQ 1 What are good practices for data annotation quality management as described in the literature and derived from actual annotation projects?

RQ 2 Compared to the previously collected good practices, which methods are actually used in practice?

RQ 3 Overall, how thorough is annotation quality management conducted in practice?

Our analysis shows that while many datasets are created according to good practices, several widespread issues exist. When using inter-annotator agreement, there is a frequent lack of actual interpretation of the agreement values. Also, sample sizes tend to be too low to make statistically sound conclusions when computing agreement and estimating the annotation error rate. Good practices suggested by the literature, like annotator training, pilot studies, or an iterative annotation process, are only mentioned rarely. Another interesting finding is that most of the time, adjudication is performed via majority voting; we found only three datasets that reported using probabilistic aggregation. Overall, we find a lack of proper reporting of how the annotation process was planned and executed, who annotated, as well as which quality management methods were used. These issues make it more difficult to gauge the quality of datasets and can hinder reproducibility. In summary, our contributions are:

- We survey the literature and compile an extensive summary of quality management methods.
- We analyze how quality management is done in *practice* compared to the good practices we found and recommend and point out common mistakes.
- Based on our findings, we provide a list of recommendations that can be used by future dataset creators to improve the quality of their datasets and to avoid common pitfalls.

In order to foster further investigation into quality management for data annotation, we also release our code³ to collect and analyze the dataset as well as our annotations.⁴ Our dataset can also be used as a reference to find papers that use specific quality management methods and serve as an example of how to apply them.

2. Background

In the following section, we discuss the most relevant work dealing with the dataset creation process in general and its quality management in particular. By quality management, we understand the overall process and measures taken to reach and maintain

³ <https://github.com/UKPLab/qanno>; GPL v3.

⁴ <https://tudatalib.ulb.tu-darmstadt.de/handle/tudatalib/3939>; CC BY-NC 4.0.

a desirable level of quality. The quality management measures we found are described in detail in § 3.

Dataset Creation. Dataset creation subsumes several activities, which can be coarsely divided into three categories: *annotation*, *production*, or *evaluation* (Shmueli et al. 2021). Different quality management methods are applicable or should be used depending on the task. **Annotation** or labeling means enriching data with additional information, for example, tags for text classification. **Production** encompasses activities like writing the text for question answering, paraphrasing, or summarizing. **Evaluation** means using humans to compare or assess properties like quality of previously labeled or produced instances. These can be manually or automatically created. While also touching on text production, this article primarily discusses annotation quality management. We still call participants in a dataset creation process annotators, even if they only perform production.

Dataset Creation Good Practices. Several books and articles have been written discussing dataset creation, especially concerning the annotation process itself. For instance, Ide and Pustejovsky (2017) collected descriptions for a wide range of different annotation projects. Pustejovsky and Stubbs (2013) describe the annotation process targeted towards training a machine learning model. However, both focus mainly on setting up the respective annotation projects, collecting data, as well as developing the annotation scheme and guidelines. Quality management is mentioned, but—except for inter-annotator agreement—not discussed in depth. Hovy and Lavid (2010) define good practices concerning conducting linguistic annotation projects. They emphasize the importance of proper annotator selection and training and how to evaluate the resulting dataset quality using agreement. Monarch (2021) discusses quality management for data annotation in the greatest detail. Their focus is predominantly on how to evaluate the quality of annotated data, from simple agreement measures over comparison with gold data to annotator-specific performance. Wynne (2005) describes good practices when creating linguistic corpora but only mentions quality as important, not how to assure it. Similarly, Roh, Heo, and Whang (2021) survey the different ways to collect data, for instance, via annotation, distant or self-supervision, but only bring up quality management in a short paragraph.

Several large-scale projects were conducted to develop standards and recommendations for creating language resources. These projects are, among others, the *Expert Advisory Group on Language Engineering Standards* (EAGLES), funded by the European Union (launched in 1993), or ISO/TC 37/SC 4, a technical subcommittee within the International Organization for Standardization. The resulting standards are either relatively challenging to find or require payment. While searching, we did not find explicit mentions of quality management or related recommendations.

Quality Management in Crowdsourcing. Many studies have shown that crowdworkers can annotate or create datasets with similar quality compared to experts (Snow et al. 2008; Hovy, Plank, and Søgaard 2014). Proper quality management is especially important in crowdsourcing, where the risk of unreliable workers is usually higher (Hovy et al. 2013). An early work describing basic quality control measures to use with Amazon’s Mechanical Turk is given by Callison-Burch and Dredze (2010). These include having multiple annotators for each instance or using control instances to estimate annotator quality. Daniel et al. (2019) define a taxonomy of quality for crowdsourcing and extensively describe related quality control measures. Their survey focuses on annotator

management and how it is implemented in annotation tools. Unlike our study, they do not analyze if and how quality control measures are used in practice as reported by dataset-introducing scientific publications. Lease (2011) notes that the annotation platform and tools can automate quality management in crowdsourcing to a certain degree, but manual inspection is still needed.

Annotation Process Analysis. Sabou et al. (2014) analyze 13 datasets created by crowdsourcing concerning how they were collected and derive good practices from this analysis. Amidei, Piwek, and Willis (2019) analyze inter-annotator agreement in the context of natural language generation evaluation and annotate 135 publications for this. Compared to these works, we go beyond analyzing only crowdsourced datasets, have a more detailed annotation scheme, annotate as well as analyze far more publications, and summarize quality management measures and good practices in greater detail.

Dataset Documentation Checklists. In the past, it has been found that datasets were often not adequately documented and were just published as-is. Therefore, several studies proposed checklists and templates that should be published alongside the dataset to remedy this issue. These are, among others, *datasheets for datasets* (Gebru et al. 2021), *dataset nutrition labels* (Holland et al. 2018), *data statements for NLP* (Bender and Friedman 2018), *accountability frameworks* (Hutchinson et al. 2021), or *data cards* (Pushkarna, Zaldivar, and Kjartansson 2022). Similarly, more and more machine learning and natural language processing (NLP) conferences have adopted and are adopting reproducibility checklists for machine learning model training. The focus of these checklists is mostly on bias, annotator background, intended use, general data statistics, data description, data origin, or preprocessing. Kottner et al. (2011) propose a checklist that can be used when using agreement values, which is a good start but very specific to only a single aspect of quality management. It is designed for clinical trials and might require adaptation for use in NLP. We did not find any checklist explicitly targeted towards overall quality management.

To summarize, while a large body of work generally discusses the dataset creation process, the parts discussing quality management are relatively scarce, quite scattered in the literature, and not easy to find. Therefore, we summarize the literature and provide an easily referenceable set of good practices and recommendations for the dataset-creation practitioner. We additionally annotate a large set of dataset-introducing papers for their quality management and conduct an extensive empirical evaluation of how it is applied in practice. To the best of our knowledge, our analysis of quality management in textual dataset publications is currently the largest and the first, while not limited to a particular area like crowdsourcing.

3. Dataset Creation Quality Management

To answer our first research question, in the following, we present the most relevant and frequently used quality management methods for dataset creation. This list is derived from good practices stated in previous work (§ 2) and the methods we found while surveying the dataset papers (§ 4) themselves. We consider the following methods *good practices* for two reasons. They are disseminated in well-acclaimed books or have been adopted by the community and are thus commonly used and tested in practice. We thus believe that the methods discussed in the following are well-suited for managing quality. It has to be mentioned, however, that only a few studies have thoroughly

investigated the exact impact of these methods on aspects like quality, time savings, or agreement (see also § 5 and § 8).

Another important point to consider is to see quality management as a means towards a goal and not as a goal in itself. Depending on the goal—for instance creating datasets with low bias, high quality, or diversity—some methods might be preferred over others. The choice of methods should thus be based on the purpose and usage goals.

Also, when applying the ensuing methods in practice, their use can be expensive. Therefore, extensive quality management needs to be balanced against the annotation costs itself when working on a limited budget; a healthy compromise between the two needs to be found.

We propose a taxonomy that puts the methods into five groups related to the *annotation process*, *annotator management*, *quality estimation*, *quality improvement*, and *adjudication*. While only briefly outlining the techniques here, we refer the interested reader to each method for a more in-depth description. An overview of the discussed methods is given in Figure 1.

We differentiate between two types of tasks for dataset creation (see § 2), namely, *annotation* (e.g., named entities or text classification) and *text production* (e.g., writing questions and answers for question answering, paraphrasing, sentence simplification). This distinction is important because specific quality management methods may work for one but not the other. For example, inter-annotator agreement and adjudication are usually not applicable to text production tasks. Both expert annotation and crowdsourcing are considered.

Our survey primarily focuses on annotation, especially label errors, but we also discuss annotation consistency, biases, and how to mitigate them. Regarding label errors, while it is sometimes impossible to assign a single, true label due to inherent ambiguity, especially in natural language processing, deciding whether a label is incorrect is often much more straightforward.

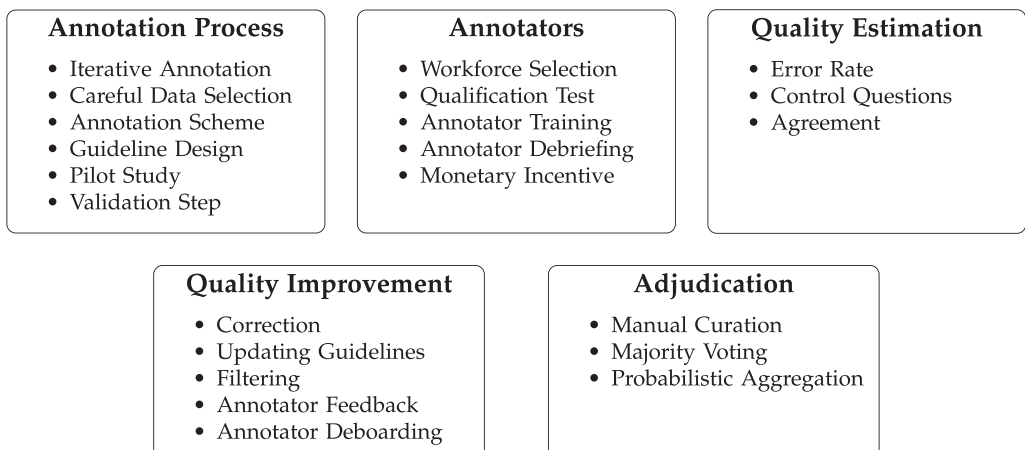


Figure 1

Quality Management methods discussed in this work. We categorize methods into annotation process, annotator management, quality estimation, quality improvement, and adjudication.

Before describing quality management methods, we first define what dataset quality subsumes. Following Krippendorff (1980) and Neuendorf (2016), we suggest targeting at least the following quality aspects⁵:

Stability A dataset creation process is stable if its output does not drift over time. Drift here means that similar phenomena are annotated similarly independent of whether they are annotated earlier or later throughout the process. Instability can, for instance, occur due to carelessness, distractions or tiredness, change in annotation guidelines, or even learning through practice.

Reproducibility A dataset creation process is reproducible if different annotators can still deliver the same results given the same project documentation regarding process, guidelines, and scheme.

Accuracy Annotations and texts created during the process are accurate if they adhere to the guidelines and the desired outcome.

Unbiasedness This describes the extent to which the created artifacts are free of systematic, nonrandom errors (bias).

Stability, *reproducibility*, and *accuracy* are also subsumed under the term *reliability* in content analysis (Krippendorff 1980). *Consistency* is related to *stability* and *reproducibility*. *Reliability* thus measures the differences that occur when repeatedly annotating the same instances; it is empirical (Hardt and Recht 2022). It is required to infer *validity*, that is, to show that the annotations capture the underlying phenomenon targeted (Artstein and Poesio 2008), necessary but not sufficient. *Validity* is latent and cannot be directly measured. Therefore, proxy metrics targeting reliability, for example, agreement, need to be used instead.

3.1 Annotation Process

The following section describes the recommended annotation process. It is written concerning annotation but can easily be adapted to text production as well.

We suggest that an annotation project should start with a *planning phase*. It can encompass important preliminaries as setting the goal of data collection, making initial choices for data and annotators, setting a budget, desired quality level or reviewing the literature for similar datasets or relevant annotation practices. Ideally, these choices are documented and become part of the dataset documentation once the dataset gets released.

The annotation scheme is often developed during an annotation project and is a living document. Also, as annotators only get familiar with the task during the annotation process, issues are found just then, and the data or task needs to be adapted accordingly. Therefore, it is recommended to structure an annotation project as a sequence of cycles with iterative quality improvement actions (Hovy and Lavid 2010; Pustejovsky and Stubbs 2013; Monarch 2021). This approach is also called **agile corpus creation** (Alex et al. 2010). In each cycle, only a slice of the data is annotated: a batch. After the batch is annotated, it is *evaluated*, and *quality-improving/rectifying measures* are taken if needed. These cycles repeat until an acceptable quality level for a sufficient number of batches has been reached. Evaluation can be performed by inter-annotator agreement (§ 3.3.3),

⁵ Note that dataset creation projects that run over a very long time and that might be subject to external effects, such as general advances in the field or societal changes, may need other definitions for these categories or incorporate specific approaches to deal with such external effects.

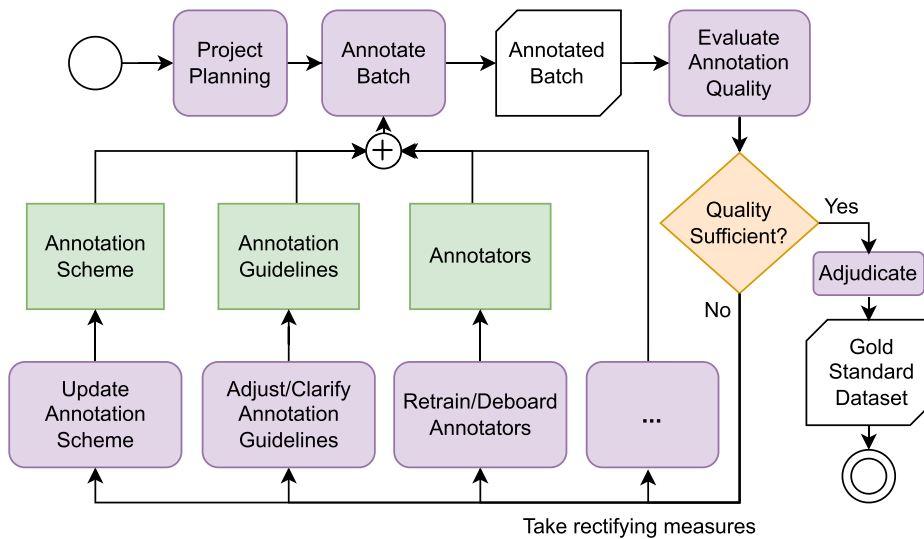


Figure 2

The recommended annotation process: After a batch of data is annotated, it is evaluated. If the quality is sufficient, it can be adjudicated. If not, several corrective measures can be taken, e.g., correcting the annotations in an additional step, annotator training, or adjusting the annotation scheme or guidelines. This is similarly applicable for text production workflows where usually no adjudication takes place.

comparing annotations to a known gold standard to estimate annotator proficiency, or having experts or a different set of annotators inspect a subset or all instances and marking errors (§ 3.3.1 and § 3.3.2).

The advantage of this iterative approach is that changes are introduced at defined points during the process. Iterating, for example, mitigates the annotation scheme and annotations running out-of-sync and improves the chance of producing high-quality datasets. Our take on this annotation loop is depicted in Figure 2. *Pilot studies* are the initial iterations used to create and improve the annotation process until it is good enough for the annotation of the dataset itself. Quality improvement measures can be, among others, retraining annotators, adjusting/clarifying the annotation guidelines or annotation scheme, onboarding or deboarding annotators, or giving back batches to annotators for correction (cf. § 3.4). In later iterations of an annotation project, when the setup has stabilized, the batch sizes can be increased, and quality control can be performed less rigorously, for example, reducing the fraction of samples inspected for quality checking or the annotations collected per instance. When using an iterative approach, stability of the annotation process needs to be taken into account, as changes to the process can cause differences in subsequently annotated batches. Also, if the annotation scheme or guidelines evolve too much, then re-annotation of previously annotated material might be necessary.

Careful Corpus Building. Not only are the labels assigned by the annotators important, but also the choice itself of texts that are annotated (Wynne 2005). Choosing texts that only rarely or even never contain the phenomena to annotate can be ineffective. Similarly, selecting texts that are of poor quality can be detrimental and cause issues in later stages of the machine learning pipeline. In order to achieve the best downstream

task performance for trained machine learning models, texts should be representative of the data encountered in the target domain. Hence, it is vital to check the data for errors and unwanted aspects like non-representative content or biases, ideally before it reaches the annotators. This can be achieved by, for example, manual inspection (Bastan et al. 2020; Govindarajan et al. 2020) (e.g., by the project manager or even as a separate preparatory annotation project) and filtering via rules (Reddy, Chen, and Manning 2019; Ghosal et al. 2022) or using spell-checking and text cleaning tools (Horbach, Ding, and Zesch 2017; Kim, Weiss, and Ravikumar 2022).

Annotation Scheme and Guideline Design. The annotation scheme defines the structure, features, and tagsets of the task to annotate. Its form and granularity can significantly impact the annotation process and the downstream machine-learning modeling. Therefore, it must capture the information of interest. The annotation scheme defines the annotation labels; the guidelines describe how to decide when to apply which label (e.g., disambiguating between different labels). Properly written guidelines are essential for annotator training to achieve consistency and reproducibility, for example, when re-annotating, extending, or creating a similar dataset on different text. The way that the guidelines are written can by itself already introduce bias (Geva, Goldberg, and Berant 2019; Parmar et al. 2023), and therefore great care needs to be taken when creating them. Instead of creating guidelines from scratch for every annotation project, existing guidelines can be reused and adapted for similar settings. In many annotation projects, the guidelines are revised several times as part of a pilot study before the actual annotation process starts (Hovy and Lavid 2010).

Guidelines for more complex annotation projects are often quite detailed and span many pages. They are usually very short in crowdsourcing and often fit into the annotation screen. Examples of excellent, extensive annotation guidelines can be found in Prasad et al. (2008) or Piskorski et al. (2023). For crowdsourcing, good examples are given by Singh et al. (2021) or Mostafazadeh et al. (2020).

Pilot Study. When entering into an (iterative) annotation project, it is crucial to validate the annotation process on a smaller scale, namely, by conducting one or more pilot studies with only a small annotator team (Pustejovsky and Stubbs 2013). Annotators in pilot studies are often the project managers themselves or a selected group of experts. We recommend that experts or project managers conduct the initial pilot study iterations; the annotation process should then be subsequently tested with the target annotators until all questions and issues are solved. This study should include developing the initial version of the annotation scheme and guidelines, configuring the respective annotation tooling, and developing the data pre-processing and post-processing steps (Kummerfeld et al. 2019). This way, issues can be spotted before investing too much effort into a flawed setup. Ideally, the data used for pilot studies should be selected to contain as many corner cases and difficult instances as possible. This reduces the chance that later, during the main part of the annotation project, significant adjustments need to be made that could cause costly re-annotation in case changes are not backward compatible. The overall difficulty of the task can be gauged, and it can be tested whether experts are needed or whether well-trained contractors or crowdworkers can achieve a desirable quality level. The expected cost can also be estimated by measuring annotation time per instance. The feedback annotators give during this phase is essential for a well-working annotation project (Monarch 2021). It has to be noted, however, that if experts or project managers conduct the initial pilot study, then they may use implicit knowledge that will not transfer to more general annotators (Krippendorff 1980).

Validation. After an annotation step has been completed, a validation step can (and should!) be added to check whether annotations are correct and of sufficient quality. Validation steps can take different forms based on the task and setup, for example, experts can inspect a subset of annotations, or there can be a separate annotation phase asking for binary correctness labels. While validation is important, it needs to be weighed against spending on annotating more instances instead if the budget is limited.

It is also possible to design a more task-dependent validation step. We call this flavor of validation **indirect validation**. It is often applicable if the annotation task consists of different subtasks that depend on each other and are hence annotated sequentially. For question answering, a first step might be to write questions and answers. The validation step could then annotate which answer best fits a given question (Mihaylov et al. 2018). For relation extraction, the first step can be marking spans and labeling their relation (Yao et al. 2019). The validation step could be that annotators are given only the marked spans and are asked to label the relation. Alternatively, the relation label could be given, and annotators are asked to mark the spans with this relation. If annotations differ between subsequent steps, then they are potentially incorrect. For natural language inference, the first task can be defined as writing a premise and hypothesis, given a relation (entailment, neutral, contradiction). In the second step, the task can be to label the relation between the two given the premise and hypothesis. If the results in the first and second steps differ, these instances require further treatment (Bowman et al. 2015).

Validation is also relevant for automatically created datasets. This, for instance, encompasses datasets that are created by crawling and transforming external resources, or that are annotated via distant or self-supervision. It should be performed after a batch of annotations have been made and before they are adjudicated. Validation can be part of quality estimation, which we discuss in more detail in § 3.3.1.

3.2 Annotator Management

Dataset creation projects stand or fall by the quality of the annotators; such a project often is an exercise in people management (Monarch 2021). At every step, it is vital to treat annotators fairly and respectfully. Here, we give a high-level overview of the different aspects of annotator management. An in-depth survey of annotator management focusing on crowdsourcing is also given in Daniel et al. (2019) and Monarch (2021). We consider both “classic” expert annotation and crowdsourcing in this work and point out when methods are more applicable for one or the other.

Workforce Selection. The type of workforce utilized considerably impacts annotation time, cost, and quality (Hovy, Plank, and Søgaard 2014). What kind of annotators to use depends, among others, on the task difficulty, availability, target language, and whether particular expertise is needed. If the annotation task is solvable by crowdworkers, it is often an efficient way to annotate (Snow et al. 2008). For more involved tasks, trained contractors can be an alternative to hiring domain experts (Chen et al. 2021). Contractors are a middle ground between crowdworkers and experts; they are experienced in conducting annotation tasks but are not necessarily domain experts. It is recommended to validate the workforce choice in one or more pilot studies.

Qualification Filter. As a common way to filter out crowdworkers that might produce low-quality work, many crowdsourcing tools offer setting requirements for the worker.

These, for instance, can be requiring a certain percentage of accepted tasks or a certain number of already completed tasks. Kummerfeld (2021) analyzes the impact of these measures on quality and discusses the ethical aspects of requiring a minimum number of tasks. They argue that it forces workers to accept a substantial amount of low-paying tasks to overcome this hurdle. The conclusion is that there is no clear relation between quality and filtering based on the percentage of accepted, previous tasks and number of completed tasks. They also note that in practice, limits are often set too high. Thus, the paper recommends either running a pilot study to obtain estimates for the actual requirement values or preferring qualification tests (see below) over simple filters.

Qualification Test. A more elaborate way to identify good annotators is to use (paid) qualification tests (Kummerfeld 2021). Before an interested annotator can participate in the primary annotation process, they must work on a small set of qualification tasks. The answers are either compared against known answers or judged by experts. If the performance is acceptable, then the annotator is allowed to work on the actual annotations themselves. The difficulty of the test can be varied based on how strictly the test should filter. For instance, task examples from the guidelines can be handed out to annotators to check whether they have been read and understood. A more challenging test would be to use new, previously unseen tasks. Qualification tests can and should not only be used for crowdsourcing but also when hiring contract annotators.

Annotator Training. Before involving new annotators in a project, it is often helpful to train them in the annotation task at hand, to go through the guidelines with them, and make sure that everything is clear (Neuendorf 2016, p. 133; Sabou et al. 2014). Project managers and annotators can give each other feedback that can then be worked into the annotation scheme and guidelines. Feedback is especially important if annotators find the guidelines difficult to understand or if they contain errors. Bayerl and Paul (2011) conduct a meta-study and analyze, among other aspects, the effect of training on agreement. They show that the better and more intensely annotators are trained, the higher the agreement becomes. Also, they point out that training is beneficial not only to crowdworkers but also to experts, as the latter might be familiar with the domain but not with the project setup at hand. Training is also essential for annotation stability, as, early in the process, annotators are often unsure and unfamiliar with the annotation process. This changes with more time spent annotating, rendering earlier annotations potentially inconsistent with later ones.

Annotator Debriefing. During and after the run of an annotation project, it is often helpful to ask one's annotators for feedback about the annotation project (Neuendorf 2016, p. 134). This feedback can then be used to improve the guidelines, update the annotation scheme, or alleviate issues that only became apparent while annotating. For instance, usability issues of the annotation editor, ways to make annotation faster, or data quality issues can be spotted and fixed before it is too late.

Monetary Incentive. Giving annotators additional monetary compensation in addition to their base pay might be an option (Harris 2011; Ho et al. 2015). The amount, for instance, can be based on their performance on control questions or after feedback rounds have shown that they reach the target for a bonus. Another way is to pay annotators more for sticking to a task (Parrish et al. 2021). If monetary incentives are used, it is essential to be transparent about it, communicate the requirements beforehand, be fair, and not change the rules post-hoc. Also, one needs to be careful that the targets for

which monetary incentives are promised are not gamed with detrimental effect towards annotation quality.⁶

3.3 Quality Estimation

After annotations have been made, their quality should be estimated and compared to the desired quality level. In case it is insufficient, counter-measures should be taken to improve it.

3.3.1 Manual Inspection. In order to judge the quality of an instance dichotomously as correct or incorrect, annotators (usually, these are different from the initial annotators) or project managers can manually inspect and grade them (Pustejovsky and Stubbs 2013). Validation can either be done on a subset of instances or as a complete validation step. In addition, after the dataset has been completely annotated, its error rate can be estimated and reported because even datasets considered gold often still contain errors (Northcutt, Athalye, and Mueller 2021). The error rate is computed by dividing the number of errors found by the number of instances inspected. Therefore, we strongly recommend inspecting a subset of instances of the final dataset, labeling their correctness, and thereby estimating the error rate. The notion of what is correct/sufficient quality or incorrect/insufficient depends on the task at hand. Hence, manual inspection is not only applicable to annotation tasks but also to text production. There, it can be determined whether the produced instance is of sufficient quality. For ambiguous instances in annotation tasks, one would judge whether the label makes sense at all in this context.

3.3.2 Control Instances. In order to gauge the performance of annotators, instances can be injected into the annotation process for which the answer is known (Callison-Burch and Dredze 2010). These gold instances are often obtained by having experts annotate a subset beforehand. Another way is to compare a single annotator's submissions to the others'; the performance estimate is then the deviation from the majority vote (Hsueh, Melville, and Sindhvani 2009) or the agreement (Monarch 2021). For example, the resulting estimates can be used to retrain annotators if they annotated too many instances incorrectly, send batches created by underperforming annotators back for re-annotation, or remove annotators from the workforce. Well-performing annotators can also be monetarily rewarded or given tasks requiring more expertise, such as task validation or manual adjudication.

3.3.3 Agreement. A common way to quantify the reliability of annotations and annotators is to compute their inter-annotator agreement (IAA) (Ebel 1951; Krippendorff 1980, 2004). For NLP, it has been increasingly adopted after Carletta (1996) introduced agreement, coming from the field of content analysis, as an alternative to previously used ad-hoc measures. Here, we briefly present the most popular and recommended agreement measures. For a more in-depth treatment of agreement and how to apply it, we refer the interested reader to the excellent works of Krippendorff (1980), Lombard, Snyder-Duch, and Bracken (2002), Neuendorf (2016), Artstein and Poesio (2008), and Monarch (2021).

⁶ This is also known as Goodhart's law: "When a measure becomes a target, it ceases to be a good measure" (Goodhart 1984).

Percent Agreement. This is the most straightforward agreement measure. It considers the percentage of coded units on which two annotators have agreed. This measure, however, suffers from several issues (Krippendorff 1980, 2004; Artstein and Poesio 2008). First, it yields skewed results for imbalanced datasets, similar to accuracy when evaluating classification. Second, it does not consider when annotators assign the same label by chance, for instance, in the event they randomly guess or spam. Third, percent agreement is influenced by the size of the tagset. Therefore, it is difficult to compare across annotation schemes. Finally, there are only two values of percent agreement that are meaningful and intuitive, which are 0% and 100%. These issues together cause percent agreement to be uninformative and difficult to interpret and compare when estimating reliability. Therefore, the usage of percent agreement is discouraged and should especially not be the only agreement measure reported.

Cohen's κ . In order to remedy the issues of percent agreement, Cohen (1960) proposes a chance-corrected coefficient, normalized to $[-1, 1]$, to measure the agreement between two annotators. Negative values indicate disagreement, 0 the expected chance agreement, and values greater than 0 indicate agreement. κ requires that the same number of annotators annotate all instances; no entries may be missing. Also, annotations need to be categorical. It is defined as

$$\kappa_C = \frac{p_o - p_e}{1 - p_e}$$

where p_o is the observed proportionate agreement and p_e the chance agreement.

Fleiss's κ . Fleiss (1971) extend Scott's π (Scott 1955) to multiple annotators.⁷ Similarly to Cohen's κ , each instance needs to be labeled by the same number of annotators. In addition, Fleiss' κ assumes that annotators for each instance are sampled randomly; it is not suitable for settings where all annotators annotate all instances (Fleiss, Levin, and Paik 2003). It is defined as

$$\kappa_F = \frac{\bar{P} - \bar{P}_e}{1 - \bar{P}_e}$$

where \bar{P} measures observed agreement as the average agreement over annotator pairs and \bar{P}_e is the expected agreement by chance.

Krippendorff's α . A different way to estimate agreement has been proposed by Krippendorff (1980). It is based on the quotient of observed disagreement D_o and chance disagreement D_e :

$$\alpha = 1 - \frac{D_o}{D_e}$$

Compared with Fleiss's κ , Krippendorff's α is more powerful and versatile: It can deal with missing annotations, supports more than two annotations per instance, and can be

⁷ Fleiss' κ is not an extension of Cohen's κ , as it assumes similarly to Scott's π that the labeling distributions are the same for each annotator, which Cohen's κ does not (Artstein and Poesio 2008).

generalized to handle even categorical, ordinal, hierarchical, or continuous data (Hayes and Krippendorff 2007). For instance, span labeling tasks like named entity recognition or relation extraction can be evaluated using a coefficient of the Krippendorff's unitized α (α_u) family (Krippendorff et al. 2016).⁸ Unitizing means that annotators first divide the instances into smaller units and only then assign labels (Lombard, Snyder-Duch, and Bracken 2002, Chapter 4). In the context of named entity annotation, unitizing, for instance, can be marking spans that contain entities or, for object detection, drawing bounding boxes around objects of interest. Hence, Krippendorff's α can also be applied to any task with a one-to-many relation between instances and annotations of different sizes. The amount of overlap between annotations made by different annotators is also considered by α_u when computing agreement. While being flexible, α is also more complicated to implement (especially in its unitizing form), has a higher runtime, and is more challenging to interpret and to compute confidence intervals for (Artstein and Poesio 2008).

Correlation. For specific tasks, annotation consists of assigning scores to instances on a numerical, continuous, or discrete rating scale or a Likert scale. These tasks are, among others, annotating sentiment (Socher et al. 2013), emotions (Demszky et al. 2020), or semantic textual similarity (Cer et al. 2017). Correlation measures like Pearson's r (linear correlation), Spearman's ρ (linear correlation of ranks), or Kendall's τ (correlation of concordant/discordant ranks) are often used to compute agreement. However, using correlation coefficients as an agreement measure is controversial, as they measure co-variation, not agreement, that is, they measure whether variables move together, but not whether they really are similar (van Stralen et al. 2012; Ranganathan, Pramesh, and Aggarwal 2017; Edwards, Allen, and Chamunyonga 2021). This means that two annotators with different biases when assigning scores, for example, one annotator systematically gives overly large scores while the other systematically underscores, would still have a high correlation but low agreement. A better alternative to the aforementioned correlation coefficients is using Intraclass Correlation (ICC) (Fisher 1925), which is explicitly designed to measure agreement. Note that there are several different formulations of ICC depending on the number of judgments per instance, whether judgments are averaged before comparison, and whether there are missing observations (Shrout and Fleiss 1979). A visual method to assess agreement between continuous variables is the Bland–Altman plot (Bland and Altman 1986). A worked example can be found in Appendix 3.

Classification Metrics. Especially for sequence labeling tasks like named entity recognition, classification metrics like *accuracy*, *precision*, *recall*, and F_1 are often used between two annotators to compute agreement (Brandesen et al. 2020). We could not find any work formally analyzing the theoretical background and implications of using these metrics as an agreement measure. However, they seem to suffer from several issues. First, they are only applicable as pairwise agreement; having more annotators would require averaging, which might cause information loss. Second, they are not chance-corrected (Powers 2011). Third, using precision and recall for computing agreement also has the downside of not being symmetric. Given two lists of labels a and b ,

⁸ The α_u family currently consists of four different coefficients (Krippendorff et al. 2016). They differ in how and whether “gaps” (unannotated units) are taken into consideration, whether labels or only units are used, or whether only a subset of labels are used when computing agreement. α_{cu} is the most applicable choice of the four that ignores gaps and takes label values into account.

the precision value of a and b turns into the recall when swapping its arguments: $\text{precision}(a, b) = \text{recall}(b, a)$. Being symmetrical is essential for agreement metrics, as one annotator should not be preferred over another. This differs from classification metrics, where one input is from the gold data, and the other is usually from model predictions.

Although it is often treated as such, agreement is no panacea; high agreement does not automatically guarantee high-quality labels. Krippendorff (2004) and Artstein and Poesio (2008) emphasize that agreement only demonstrates a reliable annotation process, which is necessary for high-quality labels but is by itself not sufficient. Further quality management, especially manual inspection, should be applied. Agreement also does not cover whether the annotation scheme and guidelines capture the desired phenomena. Low agreement also does not automatically mean low-quality labels, as tasks can inherently be subjective (Aroyo and Welty 2015; Uma et al. 2021), that is, there are cases where no distinct gold label exists for an instance.

Using only a single agreement coefficient value to gauge quality is often insufficient for a reliable estimate. Therefore, more in-depth analysis is recommended (Artstein and Poesio 2008). This can be done by manually validating the annotations (cf. 3.3.1) to get an intuition for the resulting labels and why annotators disagree. Disagreements can be caused by differences in annotator skill, differences in the data or its difficulty (Jamison and Gurevych 2015), or due to ambiguity. Other insights can be gained by computing pairwise agreement between individual annotators or by computing agreement per label (Monarch 2021). These statistics may identify poorly performing annotators or particularly difficult-to-decide labels.

If the sample size is chosen too small, the resulting agreement value might have only limited explanatory power (Allan 1999; Shoukri, Asyali, and Donner 2004; Sim and Wright 2005). It is therefore recommended to have large parts of the dataset annotated by multiple annotators for a representative agreement value (Passonneau and Carpenter 2014). Ideally, every instance should be annotated by at least two annotators to draw reliable conclusions from agreement.

Several studies propose value ranges for agreement coefficients and attach a semantic meaning to them. For instance, Landis and Koch (1977) give labels for certain value ranges of Cohen's κ (κ_c), for example, 0.01 – 0.20 *slight agreement*, 0.21 – 0.40 *fair agreement*, 0.41 – 0.60 *moderate agreement*, 0.61 – 0.80 *substantial agreement*, 0.81 – 1.00 *almost perfect agreement*. Similarly, Banerjee et al. (1999) say $\kappa_c > 0.75$ indicates excellent agreement, between 0.40 and 0.75 as fair to good agreement, and lower indicates poor agreement. Popping (1988) considers κ_c above 0.8 as reliable. Krippendorff (2004) considers their $\alpha \geq 0.8$ as reliable (later, they stated that it is the absolute lower limit and should better be 0.9), and $0.667 < \alpha < 0.8$ should only be used to draw tentative conclusions. An α value below 0.667 is said to indicate that the underlying labels are unreliable.

However, it must be noted that those boundaries are arbitrary, have certain assumptions (for instance, Landis and Koch [1977] consider only binary classification) to the task setup, and have no theoretical foundation. In general, choosing a target agreement level that is considered good enough is very difficult; there is no universally acceptable agreement level that is correct for every setting (Bakeman et al. 1997; Neuendorf 2016). Lombard, Snyder-Duch, and Bracken (2002) find that values above 0.9 are nearly always acceptable, greater than 0.8 acceptable in most situations, and greater than 0.7 acceptable for exploratory studies for some indices. Artstein and Poesio (2008) state that these limits work well in their experience, and datasets reported with lower agreement values tend to be unreliable. The threshold may also depend on the difficulty and subjectivity of the annotation task. When stating agreement values, it is therefore essential to report

boundaries and justify their value. It is also recommended to compare agreement value to other work that annotates similar phenomena and tasks if possible.

Finally, the different agreement methods have several idiosyncrasies related to how they are computed and how they behave (Zhao, Liu, and Deng 2013; Checco et al. 2017). For instance, annotations with near-perfect percent agreement can have low Cohen's κ . When Krippendorff's α is applied to a large number of instances, then its computed chance agreement term increases while α reduces, thereby favoring smaller samples. Agreement also decreases when having more annotators per instance, but this does not indicate worse quality; fewer annotators often just do not annotate the whole possible range (Bayerl and Paul 2011), and therefore, the agreement is an overestimate. These characteristics can lead to non-intuitive behavior and render interpretation more difficult.

3.4 Quality Improvement

If the quality estimation shows that the annotation quality is insufficient, rectifying measures must be taken to improve it.

Manual Correction. If the quality in a batch of annotations is too low, it can be returned to the annotators for further improvement. Also, it can be routed to different, more experienced annotators to resolve issues in case instances are too difficult for the original annotators.

Updating Guidelines. It can happen that the annotation guidelines do not cover certain phenomena in the underlying text, are ambiguous, or are difficult to understand. Then, it might be appropriate to go back to the annotation scheme or guidelines and improve them (Bareket and Tsarfaty 2021). Updating the guidelines may require discarding previously created annotations or at least reviewing and updating them. If quality estimation shows that similar categories have low agreement, then this can hint that annotators have issues discerning between them. One possible solution could be updating the annotation schema so that these categories are collapsed to a single label (Lindahl, Borin, and Rouces 2019).

Data Filtering. There are several scenarios in which already annotated instances should be prevented from making it into the final dataset. Sometimes, certain instances are too ambiguous for which annotators then strongly disagree on a single, correct label (Uma et al. 2021). Occasionally, annotations can be of low quality and should be removed. A simple solution is to filter out these instances and not process them further. The filtering can, for instance, be based on expert judgment or if there is no majority agreement (Bastan et al. 2020). Sometimes, measuring the time it takes for annotators to process instances and filter out annotations with improbably high annotation times might also be helpful (Ferracane et al. 2021).

Before filtering based on agreement, the source of disagreement should be understood, and ideally, manual inspection of flagged instances should be performed. Disagreements can for instance be visualized using confusion matrices. Filtering instances has the potential disadvantage of reducing diversity, which should be considered. Recent work also emphasizes that disagreement is inherent to natural language (Aroyo and Welty 2015) and can, for instance, be used to create a hard dataset split or even directly learn from them (Checco et al. 2017; Uma et al. 2021). Improving the annotation guidelines to incorporate edge cases should therefore be preferred over filtering.

Annotator Training through Feedback. After annotators complete a batch, experts can manually inspect the data and give annotators feedback. Thereby, common errors can be pointed out, and aspects to improve can be discussed (Ghosal et al. 2022; Kirk et al. 2022). More detailed and extensive feedback might be more feasible for smaller annotator pools, for example, contractors or expert annotators.

Annotator Deboarding. If certain annotators repeatedly deliver low-quality work, removing them from the annotator team might be desirable. One way to find these annotators is via annotation noise (Hsueh, Melville, and Sindhwani 2009), which describes the deviation of each annotator from the majority. Another is a manual inspection by the dataset creators or more seasoned annotators. Spammers can also be detected during adjudication (§ 3.5), for instance, by using multiannotator competence estimation (MACE) (Hovy et al. 2013). After deboarding annotators, it is recommended that their annotations are marked to be redone. Even though some platforms like Amazon Mechanical Turk make it possible to withhold payment, they should still be paid for the work already done unless there is compelling evidence for excessive fraudulent behavior.

Automatic Annotation Error Detection (and Correction). Instead of having human annotators manually inspect instances and search for errors, automatic approaches can be used. For some error types, it is possible to write checks that automatically find issues and sometimes even correct them (Květoň and Oliva 2002; Qian et al. 2021). These checks can be simple rules that define wrong surface form and label combinations and are derived from the data. For noisy text like Twitter data or crawled forum texts, spell-checking might improve the underlying text before it is given to annotators. A more involved approach is annotation error detection, which leverages machine learning models to automatically find error candidates, which can then be given to annotators for manual inspection and an eventual correction (e.g., Dickinson and Meurers 2003; Northcutt, Jiang, and Chuang 2021; Klie, Webber, and Gurevych 2023). Automatic checks should always be validated by human annotators to not accidentally introduce new errors.

3.5 Adjudication

In order to increase overall annotation reliability, oftentimes, more than one label per instance is collected. These usually need to be **adjudicated**, that is, finding a consensus to create the final dataset with one single label per instance (Hovy and Lavid 2010). For reproducibility, it is suggested to not only publish the adjudicated corpus but also raw annotations by the respective annotators. Learning from individual labels is also an option, especially in tasks with considerable ambiguity and disagreement (Uma et al. 2021); then, no adjudication is used. While being an effective way to improve reliability, collecting more than one label per instance needs to be weighed against annotating more instances when working on a limited budget. The most common adjudication methods are described in the following.

Manual Adjudication. To create a gold corpus, skilled annotators, often domain experts, manually inspect and curate each instance to a single label (Bareket and Tsarfaty 2021). While slow and expensive, this approach can yield high-quality data because ties can be broken and errors corrected during this inspection procedure. Curation can be sped

up with automatic tooling, for instance, by automatically merging instances for which there is no disagreement or where the disagreement is below a certain threshold.

Majority Voting. When using majority voting, given an instance rated by multiple annotators, its resulting label is the one that has been chosen most often. Instances without majority label can be discarded or given to an additional annotator to break the tie. These are often experts but can also be (experienced) crowdworkers or contractors. In some work, supermajority voting is used. It means that more than 50% of annotators must agree, for example, at most one differing label is allowed, or even a unanimous vote is required. Majority voting is easy to implement and a strong baseline compared with the more complex methods described in this section (Paun et al. 2018). But Lease (2011) notes that using majority voting might drown out valid minority voices and can reduce diversity, which should be taken into account.

Probabilistic Aggregation. In majority voting, it is assumed that all annotators are equally reliable as well as skilled and that errors are made uniformly at random. This assumption does not always apply in real annotation settings, especially for crowdsourcing. Annotators can be better or worse in certain aspects, might be biased, spamming, or even adversarial (Passonneau and Carpenter 2014). To alleviate these issues, Dawid and Skene (1979) propose a probabilistic graphical model (that is referred to as Dawid-Skene, named after its inventors) that associates a confusion matrix over label classes for each annotator, thereby modeling their proficiency and bias. The resulting aggregation is then based on weighing labels with the respective annotator's expertise for this label. An alternative formulation called *MACE* that also models spammers is given by Hovy et al. (2013).

It has been shown that using more sophisticated aggregation techniques can yield higher-quality gold standards (Passonneau and Carpenter 2014; Paun et al. 2018; Simpson and Gurevych 2019), but majority voting is often a strong baseline. The works mentioned above also discuss probabilistic aggregation in more detail.

4. Data Collection and Annotation

To answer **RQ 2** and **RQ 3**, that is, to analyze which quality management measures are actually used when creating machine learning (research) datasets and how well studies adhere to these, we collected publications that introduced new datasets and annotated them for quality aspects.

4.1 Data Selection

To collect relevant papers, we first attempted to use full-text search in abstracts from papers contained in the ACL anthology (Gildea et al. 2018) for keywords like *dataset*, *corpus*, *treebank*, or *crowdsourcing*. This was quickly shown to be infeasible, as our search selected 13,776 out of 36,501 publications, showing low precision.

Instead, we chose to leverage *Papers With Code*⁹. This project—among other things—curates a list of datasets used in machine learning research with references to the publications that introduced them. We first selected all text datasets and matched the publication title that introduced it against the ACL anthology. We only considered

⁹ <https://paperswithcode.com/>.

papers published in top conferences as well as in their respective *Findings* for the following reasons. First, as the annotation is expensive and the budget was limited, this made the annotation more feasible by reducing the overall number of papers to read and annotate. Second, as we are interested in collecting good practices, we hope that these publications that also passed peer review are of higher quality. Publications from the following conferences were considered:

- AACL
- ACL
- CL
- COLING
- CoNLL
- EMNLP
- EAACL
- Findings
- LREC
- NAACL
- TACL

This yielded a total of 591 publications to annotate, of which 314 mentioned human annotation or validation. More details about our data selection and the guidelines, in particular the entire annotation scheme, including all the label values, can be found in Appendix 1 and Appendix 2.

4.2 Annotation Scheme

We annotated the following aspects at document level:

Manual Annotation For our analysis, we are primarily interested in scientific publications introducing text datasets that use manual annotation in any form, which is why we annotate this aspect. Manual annotation may serve, for example, for creating the labels or writing text. This also includes papers that only have human validation.

Task Type There are two task types we consider, *annotation* and *text production*, as they require different methods for quality management. For instance, computing agreement is only possible for the former. Text production also does not lend itself to adjudication.

Number of Annotators The number of annotators per instance whose labels are later adjudicated. This is only annotated for *annotation* datasets, as freeform text usually is not adjudicated.

Mode of Employment We differentiate between *volunteers*, *crowdworkers*, *contractors*, and *expert annotators* (§ 3.2).

Quality Management Measures The measures mentioned in the publication to manage quality (§ 3).

Adjudication The method of converting several annotations per instance into a single ground truth (§ 3.5).

Agreement In the event that IAA was computed, we record the metric's name, the subset size if not computed on all the annotated data, and the actual value. Note that a given dataset can have more than one agreement calculation (§ 3.3.3).

Error Rate In the event that the error rate was estimated, we record the actual value and the size of the subset that was inspected (§ 3.3).

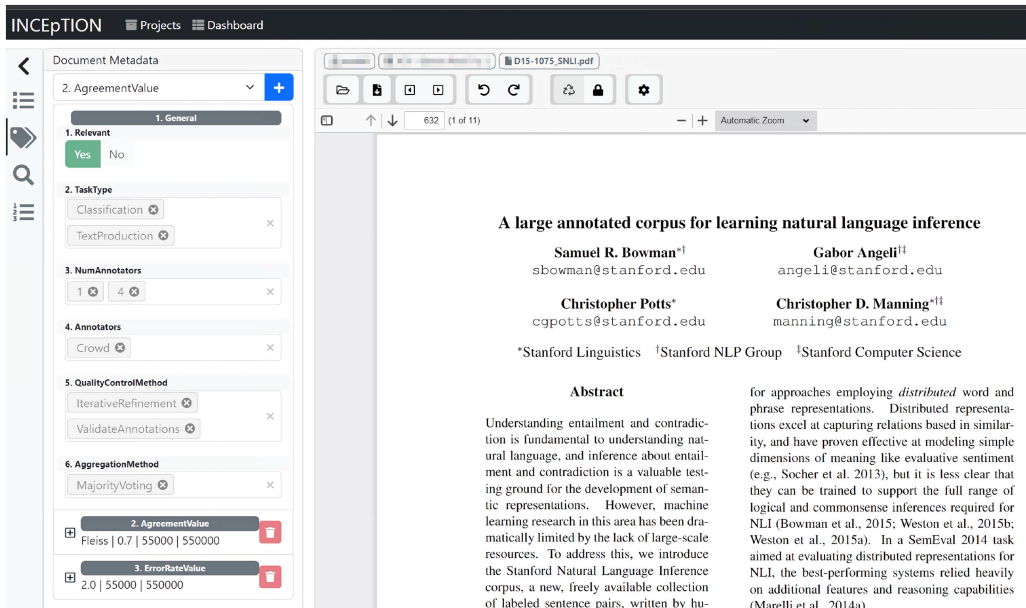


Figure 3 Annotation setup in INCEpTION. On the left, the annotation editors can be seen; on the right, a PDF viewer shows the publication to annotate directly in the browser.

Overall We assign an overall rating to each publication having human annotators based on their quality management conducted and reported. The grades are in three categories:

Excellent Does most of the following: uses the iterative annotation process, trains annotators, computes agreement and error rate, performs extensive validation, and does human inspection throughout.

Sufficient Uses some of the recommended techniques, but not as extensive as excellent. Has at least some validation and manual inspection.

Subpar No agreement, validation, manual inspection, error rate, or other quality management performed and reported. The data quality, at most, relies on aggregating multiple annotations.

We discuss limitations due to the potential subjectivity of this rating in § 8.

A screenshot depicting the annotation editor using this annotation scheme can be found in Figure 3.

4.3 Bias

Using *Papers With Code* as the source of publications potentially introduces several forms of bias, which we discuss in the following:

Quality As we only analyze publications from top NLP venues and, for instance, exclude works published in workshops, we suspect that our analysis is biased towards analyzing datasets of better quality.

Time When looking at the distribution over publication years, we see a bias towards more recent publications.

Popularity *Papers With Code* requires volunteers to manually add datasets to the website. Therefore, the resulting collection as well as our analysis might be biased towards more popular and commonly used datasets.

Availability As we analyze annotation quality management by using the publication that introduced it as their proxy, we rely on the existence of an accompanying paper describing this dataset and that the respective paper was published in a top venue. Other datasets might not have been published with such an accompanying publication (this is often the case for LDC datasets), or it might have been rejected, making it unavailable for our analysis.

Domain As we only analyze publications from general venues and not specialized venues like workshops for narrower domains such as legal or medical NLP, our collection might be biased to contain datasets that are of more general interest; particular domains might be underrepresented.

In order to quantify the bias and to estimate how well *Papers With Code (PwC)* covers the ACL anthology, we additionally annotated a random subset of 500 papers from the years 2013 to 2022 for the datasets they use. 2013 as the minimal year is chosen as older datasets are for the most part not covered by PwC (see Figure 4). 2022 as the maximum year was chosen as our snapshot of PwC is from November 26, 2022 (see Appendix 1). Again, we limited ourselves to the aforementioned top conferences and sample 50 papers per year randomly, resulting in 500 papers total. We annotated for two aspects: datasets used in the publication and whether a publication introduces new datasets. Datasets were marked as not relevant if they do not contain dataset usage or use any other modality than text. Subsequently, we deduplicated dataset mentions and linked them to PwC in the event they have an entry there. The coverage analysis can be found in § 5.1.

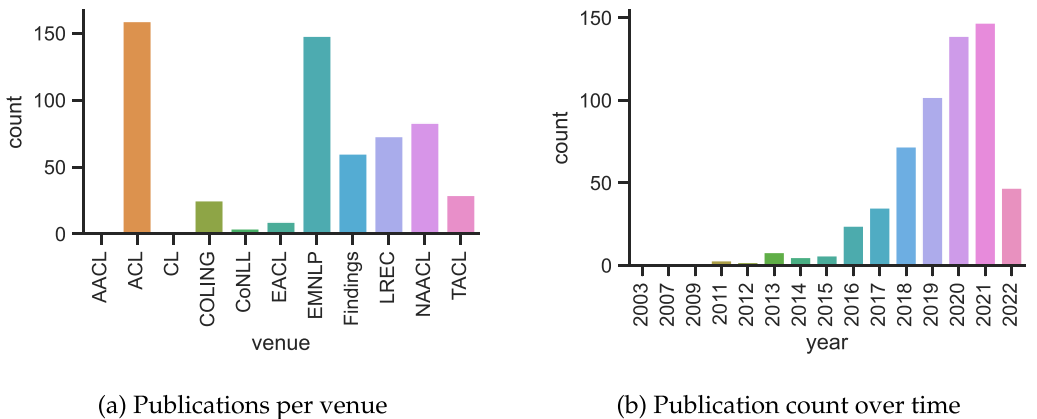


Figure 4 Statistics over the dataset created by annotating text dataset introducing publications obtained from *Papers With Code*.

4.4 Annotation Process

The annotation process we used was the same for both quality and coverage annotations. It slightly deviates from our best practices due to limited time and money. We downloaded the full-text PDFs of the selected paper and annotated them in *INCEpTION* (Klie et al. 2018). This annotation tool was chosen because it is free to use and supports annotating PDF documents out-of-the-box. The annotations were created by the first author of this work, an experienced researcher in NLP with a strong data annotation background.

We first conducted an initial pilot study to determine the aspects to annotate, followed by the annotation itself. The tagset was iteratively extended during the annotation process. After all papers had been annotated once, we did a second round to make the annotations more consistent with the now-complete tagset. Finally, we did another validation round and additionally used semi-automatic checking to improve consistency and quality further. Thus, each publication was only annotated by a single author but inspected several times to guarantee correctness and consistency. Due to the intricate and complex annotation scheme with many aspects, the expertise needed, and the exploratory nature of the annotations, we were only able to employ a single expert annotator. Instead, we opted for repeated validation and correction. In total, annotation alone took over 100 hours. While not ideal, this is a similar setup as used in previous works surveying NLP publications (Sabou et al. 2014; Amidei, Piwek, and Willis 2019; Dror et al. 2018; Shmueli et al. 2021).

5. Analysis

After having annotated a large corpus of dataset introducing data, we now use it to investigate how annotation quality management is practiced quantitatively (RQ 2) and qualitatively (RQ 3). An overview of the overall usage of each method can be found in Table 1. Regarding recommended good practices, it must be noted that there is no way of managing the dataset creation process that guarantees high-quality results. Nevertheless, some methods have been shown to yield better quality than others: Bayerl and Paul (e.g., 2011); Monarch (e.g., 2021). These choices of how to manage quality have to be looked at in the context of the task to annotate for and the constraints at hand, for instance, concerning available budget, time constraints, annotator number, and experience.

Our analysis is based on what is explicitly reported in the publication; if it was not reported, we are unable to consider it. While this might cause our analysis to be less expressive and accurate, we see no simple way to study quality management in practice. Also, this issue further emphasizes the importance of proper reporting, even if it is just in an appendix or supplementary material.

5.1 Dataset Statistics

Quality Statistics. In total, we selected and annotated 591 publications. These were organized into three groups based on the amount of human involvement. 277 did not report any human annotation for their dataset creation. In these cases, annotations were crawled or obtained via distant supervision or other means. Sixteen relied on humans to validate their algorithmically created data; 298 had humans annotating or producing the text. Of these 298 publications, 81 were introducing datasets that used annotators only for text production, 161 for labeling, and 56 for both. Datasets that leveraged both

Table 1

Overview of how often each quality management (see also Figure 1) method was used in absolute numbers (#) and relative to all works that used manual annotation (%). For adjudication, the denominator is the number of publications for which adjudication is applicable. Except for agreement, validation, and error rate, counts are directly computed from the *Quality Management Measures* field of our dataset. For the other methods, we count it for the respective metric if there is at least one usage mentioned. Note that values are non-exclusive, as publications can make use of any combination of methods.

Category	Method Name	#	%
Annotation Process	Agile Corpus Creation	68	22
	Pilot Study	67	22
	Validation Step	125	41
	Data Filtering	46	15
	None/Not specified	96	32
Annotator Management	Qualification Filter	80	26
	Qualification Test	56	18
	Annotator Training	55	18
	Annotator Debriefing	18	6
	Monetary Incentive	13	4
	None/Not specified	157	52
Quality Estimation	Error Rate	54	18
	Control Questions	28	9
	Agreement	156	52
	None/Not specified	102	34
Quality Improvement	Correction	68	22
	Scheme and Guideline Refinement	31	10
	Annotator Deboarding	39	13
	Annotator Feedback	24	8
	Agreement Filtering	29	9
	Manual Filtering	16	5
	Time Filtering	11	3
	Automatic Checks	34	11
None/Not specified	135	45	
Adjudication	Manual Curation	29	14
	Majority Voting	68	34
	Probabilistic Aggregation	2	1
	Unknown	92	46
	Other	5	2

text production and labeling were often created for tasks like natural language inference or question answering. There, the surface forms were usually written by workers before their relationships were annotated in a follow-up step.

The resulting dataset size exceeds that reported in Dror et al. (2018), who inspected 233 papers for their analysis of statistical testing in NLP research, as well as Amidei, Piwek, and Willis (2019), who inspected 135 publications for analyzing agreement in the context of natural language generation evaluations. The distributions of publications per venue and over time are depicted in Figure 4. It can be seen that most were published in or after 2018.

Coverage Statistics. PwC only contains entries for a subset of dataset-introducing publications. To analyze the coverage and to better understand the potentially resulting bias (see § 4.3), we conducted another annotation of 500 papers from the anthology from

the years 2013–2022 for their dataset usage. Based on these annotations, we first of all can see that 430 of the publications mention relevant dataset usages.¹⁰ Of these, 132 (30%) publications introduced new datasets of any kind.

In total, we found 993 mentions of 622 unique datasets; 495 datasets are only mentioned once. Of the 622 unique datasets, 172 (27%) are also contained in our dump of PwC. When taking our filtering of publication venues into account, we see that from the papers that we annotated for quality management, 49 of all papers and 30 of relevant papers are in the sample annotated for coverage as well as in the sample for quality. In relation to our quality dataset, these make up 8% of all and 10%, respectively, of relevant publications.

To better understand the popularity of the annotated datasets, we analyze their mention frequency. We can see that on average, a dataset in the coverage sample was mentioned 1.60 times. In the sample for quality annotations, this was 1.96 for all publications and 2.13 for only the relevant ones. While not being a large difference, this still indicates that our sample based on PwC is slightly biased towards more popular datasets.

Finally, we find that our dataset in particular does not cover most LDC corpora or datasets introduced as part of shared tasks. These are, for instance CoNLL, WMT, SemEval, or TAC.

Bias. We used PwC in order to reduce the effort of finding publications that introduce new datasets in the first place. The aforementioned statistics indicate that our sampling using PwC introduces biases towards more popular, more recent and on average, higher quality datasets. While not ideal, we argue, however, that this is not necessarily a disadvantage, as the datasets that we analyzed are actually frequently used in practice. Thus, their quality has direct impact on the research community. Also, with being more popular, we hope that their quality management also follows good practices comparatively more often. While having a seemingly low coverage overall, our sample size nonetheless is much larger compared with previous work, still yields interesting insights, and was already costly to annotate.

Bias in time, popularity, or domain might be an issue, as there could be practices from the past that are falling through our cracks that would be relevant and interesting for the general public. We alleviated this issue by also surveying other literature like books and by collecting and analyzing a large corpus of dataset-introducing publications.

Our analysis of annotation quality management it is still a valuable contribution, especially in combination with our survey of good practices and a good start for future work. Also, we are interested in finding issues; to offer solutions for their alleviation, having unbiased counts is desirable but not crucial. We hence suspect that the statistics derived overestimate quality compared with the general populace and that our analysis are potentially too positive. The statistics that follow thus should be seen as an optimistic estimate. Finally, it has to be noted that the resulting dataset is a side product of the survey and should be seen in this context. While we have taken the utmost care during annotation, the dataset is not intended to be used in machine learning or other areas where quality needs to be very high and absolute.

¹⁰ The following metrics are with respect to relevant publications only.




Figure 5
Distribution of percentage of papers over subjective quality management. Mostly, quality management was good or excellent, but a large fraction is only subpar.

5.2 Overall

To better understand how well quality management is performed in practice (RQ 3), we assigned each work an overall score. Their distribution is depicted in Figure 5. It can be seen that around 45% of publications perform well, and 25% use excellent quality management according to our annotation scheme and guidelines. However, we also find that about 30% only conduct subpar quality management. These often either did not report the annotation process at all or did so just very briefly and did not mention that they applied any quality management.

5.3 Annotation Process

In the following, we analyze the publications concerning their annotation process.

Annotation Scheme and Guidelines. Of the 298 publications having human annotators, 68 (22%) reported having an iterative refinement loop, which is our recommended annotation process. This loop was mainly used for iteratively refining the annotation guidelines after doing pilot studies (10%) or repeatedly correcting instances until they reached sufficient quality (12%). Eighteen (6%) works reported that their annotators gave feedback on the task during annotation so that the annotation process could be improved.

Sixty percent of publications with manual annotation described their annotation scheme, showed their annotation interface, or published their annotation guidelines together with the dataset itself in some form. Not reporting annotation schemes and guidelines causes several issues. First, these cannot be checked and reviewed, making it difficult to assess their quality. Second, not making it available is a significant obstacle to reproducibility or later extensions. In several cases, the reader was referred to supplementary material or appendices, which we could not find in the publication or online.

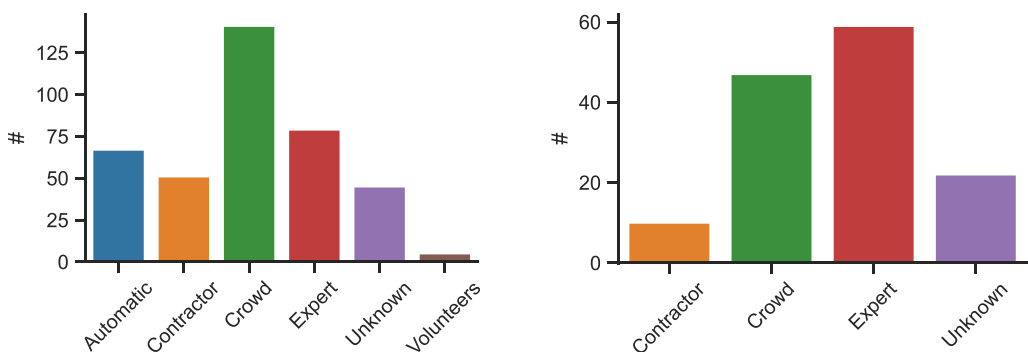
Pilot Study. Overall, only 22% of the publications mentioned having conducted a pilot study. This value is relatively low, as pilot studies are an essential tool to dial in the annotation scheme and guidelines and to get feedback from the annotators. As we only rely on what is mentioned in publications, we cannot say whether the authors considered this method common and thus did not see the need to mention that they conducted a pilot study or that it is indeed not done often enough.

Validation. In many cases, annotations were validated as an additional step in the overall process either by the annotators themselves or by having experts check them (41%). For automatically annotated data, only 16 out of 293 reported that they used human validators. Not validating can be an issue; for example, datasets created solely by distant supervision can contain many labeling errors (Mintz et al. 2009). Ten of these publications also reported the resulting error rate, which ranges from 1.40% to 16.60% with mean 8.93% and median 8.55%, showing the importance of validation. We found 25 publications that reported indirect validation (8%).

5.4 Annotator Management

The distribution over different annotator types is shown in Figure 6. Overall, publications mostly used crowdworkers or experts for their annotations. For validation, experts were more commonly selected. In many cases, the kind of annotators used was also not reported.

We find that the preferred method to filter out annotators, especially crowdworkers, is by requiring a certain number of previous successful tasks and a high acceptance rate (26%). Qualification tests, recommended by Kummerfeld (2021) over filters, are also often used (18%). Annotators are given training only in 18% of cases, which we find pretty low compared to the benefits it might give. Out of these cases, training was overwhelmingly given to contractors and crowdworkers; only one publication mentioned that experts were trained. We note, however, that even experts should be given training, as being an expert does not automatically indicate familiarity with the annotation setup and scheme at hand (Bayerl and Paul 2011). Only in a few cases (8%) is it explicitly stated that annotators were given feedback on their work or that annotators give feedback to improve the annotation process (6%). While not being reported, we assume that training and feedback were given in many more cases, especially for contractors. Better interaction between project leads and annotators is one reason contractors are typically chosen over crowdworkers. Thirteen (4%) of publications mention some kind of additional monetary incentive.



(a) Absolute number of annotators by type.

(b) Absolute number of validators by type.

Figure 6

Distribution over annotator types. For annotation (a), crowdsourcing is used the most; for validation (b), it is experts. Note that a publication, respectively dataset, can leverage more than one annotation type.

5.5 Quality Estimation

The quality of the dataset created needs to be estimated during and after its creation so that its quality can be guaranteed and countermeasures can be taken to improve it if needed. Overall, we find that two main techniques were used for this, which are agreement (52%) and error rate estimation (18%). We analyze these in more detail in § 5.8 and § 5.9, respectively. Control questions were used by 9% of the publications to gauge annotator performance and task quality. Overall, 65% of works mention at least one way of estimating quality.

5.6 Quality Improvement

Next, we analyze rectifying measures used to improve the data quality after it has been estimated in a previous step and deemed insufficient. In most cases, incorrect or low-quality instances are corrected (22%) or filtered out (15%). Of the 46 publications that mention filtering, 29 report filtering based on agreement, 16 after manual inspection, and 11 based on unsound, improbably low annotation times. Eleven percent of publications mentioned having applied some kind of automatic checks to identify potential errors, such as spell checking or hand-crafted rules. Sometimes, annotators were removed from the workforce if they repeatedly delivered sub-par quality (13%). Rarely were they given feedback by experts or the project managers (8%). This number increases to 22% when excluding datasets only annotated by experts. Overall, we do not see much usage of rectifying measures; only 41% of publications using human annotation report at least once.

5.7 Adjudication

Similarly to Sabou et al. (2014), we find that majority voting was most often used to adjudicate labels (34%). In a few cases, publications reported that in addition to majority voting, ties were broken by consulting additional workers or experts (8%). The second most common way of adjudication was manual curation (14%). Overall, we find that in 46% of labeling datasets, adjudication methods were not reported clearly or at all. This leaves the reader to guess, which is concerning.

We only found two publications that used Dawid-Skene (Dawid and Skene 1979) and one that used *MACE* (Hovy et al. 2013). The latter was just used to filter out spammers during annotation and not for adjudication itself. One publication mentioned trying out probabilistic aggregation, yet they report that just using majority voting yielded better results for them. Some studies also mentioned aggregation based on annotator confidence and skill, but no details were given describing the exact procedure used.

The fact that majority voting is by far the most frequently used method is interesting, as aggregation is a quite well-researched topic in the crowdsourcing research community (Sheshadri and Lease 2013). It has also been shown that using more intricate methods can create higher-quality gold standards (Paun et al. 2018; Simpson and Gurevych 2019).

5.8 Error Rate

While it is often assumed that (research) datasets represent a gold standard and do not contain errors, this is often not the case (e.g., Northcutt, Athalye, and Mueller

2021; Klie, Webber, and Gurevych 2023). To estimate the overall correctness of the dataset, its annotation error rate should be computed after adjudication is completed. Computing the error rate is typically done by randomly sampling a subset and marking instances as correct or incorrect. From our analysis, only a few publications (18% of all having human annotation) estimated and reported an error rate. The average error rate reported is 8.27%, and its median is 6.00%.

Sample Size. From the dataset we analyzed, 64 out of 80 error rates were computed by inspecting only a subset of the data. The inspected subset needs to be of sufficient size for the estimate to be reliable. If it is too small, the estimate has large error margins and hence low statistical power, potentially leading to over-optimistic or incorrect conclusions (Button et al. 2013; Passonneau and Carpenter 2014).

For instance, it was found that TACRED (Zhang et al. 2017), a dataset for relation classification, contains a large fraction of incorrect labels. During the dataset creation, 25% of the annotations were validated by crowdworkers; after adjudication, the authors finally inspected a sample of 300 instances and estimated an error rate of around 6.7%. It was then subsequently discovered that the dataset contains significantly more errors. First, it was claimed to be around 50% by Alt, Gabryszak, and Hennig (2020), who only analyzed a smaller and biased sample. Stoica, Platanios, and Poczoz (2021) finally inspected all samples and found an error rate of 23.9%. This shows the importance of manual inspection of large enough sample sizes.

In the publications inspected, we did not find any work that based their choice of sample size on a statistical footing or gave reasoning for selecting that specific value. In most cases, pretty numbers were chosen without rationale (e.g., round numbers like 100 or 200 were picked often), or a percentage of the total size (e.g., 5%) was used. The mean sample size is 1,305.68, while its median is 200.00 (see Figure 7).

We also analyze the impact the sample size has on the estimate’s reliability using confidence intervals and their interval half-widths. The interval half-width measures




Figure 7 Number of inspected instances vs. the resulting confidence interval (CI) half-width for a 95% CI. It can be seen that overall, too few instances are inspected to estimate the error rate reliably, as they have a substantial margin of error. Four values above 1,000 were filtered out to aid the visualization.

the margin of error associated with the confidence interval. It is computed as the largest distance between the point estimate of the error rate and its endpoints. The confidence interval for an estimated error rate \hat{r} is then given as $[\hat{r} - h, \hat{r} + h]$. If h is relatively large (e.g., 0.05), then the error rate is with high probability within ± 5 percentage points. This is quite a large margin, especially for error rates, as \hat{r} is usually small there and (hopefully) close to zero.

To compute the margin of error, we model estimating the error rate as sampling with replacement¹¹ where annotators randomly inspect a subset of instances and mark them as either correct or incorrect. For each mention of error rates in our analyzed publications, we then compute a 95% binomial exact confidence interval for each estimate and its half-width h .

The half-widths for each estimate are plotted in Figure 7. For almost all estimates, the resulting confidence intervals are very wide, rendering a given point estimate statistically unreliable. When choosing a different sample to inspect and mark, the error would fluctuate by a large margin and has thereby only limited explanatory power. We suggest inspecting at least 500 instances¹² or the whole dataset, whichever is smaller, for a more sound estimate. Note that calculating the sample size that way is an optimistic estimate, as it assumes independent and identically distributed instances, which is often not the case. Also, giving a confidence interval when stating the error rate is recommended. This can either be done by computing a binomial/hypergeometric confidence interval or using techniques like bootstrapping. Otherwise, giving a point estimate implies precision which it has not, especially when giving several decimal places.

5.9 Agreement

For every paper inspected, we annotated whether agreement measure usage was mentioned and recorded its type and value if it was. In most cases, agreement has been used to demonstrate the dataset quality after the annotation was completed. Sometimes, agreement has also been used to either remove annotators or remove annotations. We observe that 52% of publications involving human annotators reported using at least one form of agreement. Concerning the form of dataset creation, it is 48% for labeling and 31% for text production. In addition, we find that 7 publications that—while not employing humans for the annotation itself—leverage agreement during validation steps. The usage statistics are depicted in Figure 8. Overall, Cohen’s and Fleiss’s κ , Krippendorff’s α , and percent agreement were used the most, followed by F_1 . On average, each publication used 1.33 agreement measures with median 1 (based on works that actually used at least one). Percent agreement as the only measure was used in around 11% of all publications that use at least one method. Only using percent agreement makes it difficult to estimate, interpret, and compare the dataset’s quality, and its usage is therefore discouraged (Krippendorff 2004). In 10 cases, the used measures were not clearly named but only referenced as, e.g., κ or IAA (this is noted by a “?” in Figure 8).

11 The sample size is usually much smaller than the dataset size, which is why we can approximate the hypergeometric distribution (sampling without replacement) with the binomial distribution for simplicity.

12 Assuming a binomial model with a true error rate of 5%, a sample size of 456 yields a 95% CI with $h \approx 0.02$.




Figure 8

Distribution over counts of the agreement measures used. We count each method only once per publication, even if it has been used more than once. Overall, agreement measures were used in 156 publications involving human annotators.

Regarding the usage and reporting of agreement as an indicator for reliability, we found similar issues as described by Amidei, Piwek, and Willis (2019). Often, only the agreement value was stated without any interpretation or comment (52%), which limits its explanatory power. In many publications, the quality derived from the agreement was described with a freeform explanation, for example, *high*, *fair*, *substantial* (27%). These frequently do not have a relation to the actual value, as, for example, values < 0.3 were described as *reasonable*. Rarely was agreement compared to previous studies (5%) or an interpretation based on a range given by the literature cited (16%). This can partially be explained by only some datasets having a suitable predecessor as a reference.

In all cases, these ranges' limitations were not considered; for example, the ranges defined by Landis and Koch (1977) are based on binary classification. In contrast, several datasets introduced by the respective publications had more than two possible labels. Also, several times, the stated ranges did not match the metric. For example, the ranges from Landis and Koch (1977) that apply to Cohen's κ were instead used for Fleiss' κ . Several times, publications used pairwise agreement measures for more than two annotators and reported them pairwise. While that is valid in itself, additionally using multi-user measures like Fleiss' κ or α is recommended. We also found several cases where the usage of Cohen's κ was reported, but more than two annotations per instance were obtained. It is also discouraged to use correlation metrics as a measure of agreement. We found 7 (2%) publications that still reported its usage. Last but not least, κ or α was sometimes given in percent. This can confuse the reader as these values are usually given as a value in $[-1, 1]$, and percent agreement is a distinct metric on its own.

Agreement Values. We plot the agreement values for the most frequently used methods in Figure 9 together with the boundaries suggested by the literature (even though they are often subjective). For Krippendorff's α , the values are rarely larger than 0.8, which would indicate acceptable agreement according to Krippendorff (2004). Some are in the zone $0.67 \leq \kappa \leq 0.8$, which indicates that the resulting annotations should only be used




Figure 9 Agreement values for the papers inspected. Also shown are the ranges often used for interpreting these values.

to draw tentative conclusions; the majority is even below that. Many agreement values are on the lower side, hinting towards lower agreement or considerable ambiguity in the underlying task.

Agreement for Sequence Labeling. For sequence labeling datasets (e.g., Named Entity Recognition or Slot Filling), dataset creators either did not compute agreement or relied on per-token κ , α , or classification metrics like precision, recall, and mainly F_1 . Brandsen et al. (2020) argue that per-token agreement for sequence labeling comes with two issues. First, annotators label sequences and not tokens, so the measure does not reflect the task well. Second, the data is imbalanced, as most tokens are labeled 0, indicating no span. Excluding this would result in an underestimate of the agreement. They argue for using F_1 and averaging it between annotators. However, this is not chance-corrected and can only be used to compute pairwise agreement; averaging might lead to a loss of information. Only a single paper (Stab and Gurevych 2014) used Krippendorff’s unitizing α_u (Krippendorff 1995) to compute agreement for sequence labeling. α_u in itself can directly support sequence labeling and is an excellent way to compute agreement in this setting. We hence agree with Meyer et al. (2014) that unitizing agreement measures should be used if not as the only measure, then at least additionally. Our conjecture for why unitizing measures are not used more often is that these are not very well-known, and their complex implementation hinders adoption.

Sample Size. Dataset creators sometimes decided only to have one annotation per instance for the majority of the dataset to save resources. Then, only a subset was annotated multiple times to compute the agreement. Similar to Passonneau and Carpenter (2014) and as described in § 5.8, we note that having too small sample sizes is an issue as even a relatively relaxed 95% confidence interval spans quite a wide range of values. A sample size that is too small can cause estimates to vary by a large margin. This might lead to a different interpretation based on a pre-determined, targeted agreement level or a range suggested by the literature.

Out of 288 papers that reported agreement values, 197 have had the complete dataset annotated multiple times, and 91 were computed from a subset. The mean sample size for the latter was 1,882 with median 200. Forty-seven (51%) agreement values were computed on 200 instances or less, 26 (28%) even on less or equal than 100.

It is therefore recommended to (1) have large sample sizes to compute agreement on, ideally the complete dataset (which has the advantage of improved quality due to aggregation) and (2) compute a confidence interval for the agreement value, for example, by bootstrapping (Efron and Tibshirani 1986; Zapf et al. 2016). Computing the required sample size for a given precision and confidence level is not straightforward and depends on the metric (Shoukri, Asyali, and Donner 2004). For Cohen's κ , an approximation is described by Donner and Eliasziw (1992); for α , it is given by Krippendorff (2011). As a rule of thumb that works for both κ and α , given an expected/desired agreement value of 0.8 with a precision of $h \pm 0.05$ and a confidence level of 95%, at least ≈ 500 instances should be annotated.

While this is highly desirable, we notice that this comes with costs and additional effort. We did not find a single report of confidence intervals for agreement values in the publications analyzed for this work. As we do not have access to the raw, unadjudicated data used to compute the agreement value (which is needed for computing confidence intervals), we cannot easily conduct an analysis similar to the one for error rates in § 5.8.

6. Recommendations

Based on our analysis of 591 papers published in top NLP conferences as well as on our survey of the relevant literature, we derive the following recommendations and good practices for dataset creation quality control. A case-by-case ranking of measures should be done based on the circumstances of the project.

Annotation Process.

- Use an agile, iterative annotation process and annotate in batches (Alex et al. 2010; Pustejovsky and Stubbs 2013).
- Conduct pilot studies to validate the annotation setup before starting the actual annotation.
- Quality estimates after each batch should guide the improvement of guidelines and the scheme.
- Rectifying measures like corrective annotation, annotator retraining, or data filtering should be used to improve the overall data quality iteratively.
- Annotator feedback should be incorporated during a pilot study and annotation.

Annotator Management. Workforce selection and annotator management are crucial for a successful annotation project. Different annotator types can be viable depending on the task difficulty and the expertise or background knowledge required. Datasets these days are most often annotated by crowdworkers. A feasible alternative (even for tasks that usually require expert annotators) is hiring and training contractors via platforms

like Upwork or Prolific. This can open up better ways to collaborate while having similar costs.

- The choice of annotator type (expert/contractor/crowdworkers, etc.) should be validated as part of a pilot study.
- Annotators should be paid properly and treated with respect.
- They should be trained before and during the annotation process for the best results, even experts.
- Annotator feedback should be used to fine-tune the guidelines, annotation scheme, or annotation editor and to spot errors or issues like low data quality.
- To select annotators, qualification tests are the recommended way; criteria like completed tasks or acceptance rate can be an addition, but should be rather lower than higher to not force workers into low-paying qualification jobs.

Quality Estimation. Precise quality estimation is essential to steer the annotation process after each batch and before the final release of the dataset.

- Inter-annotator agreement can be used to determine whether the annotation process is overall reliable.
- In addition to agreement, manual inspection is recommended to validate annotations and estimate accuracy. This can be done by either the annotators themselves or experienced/expert annotators.
- Disagreements can be visualized using confusion matrices.
- An alternative to having annotators validate instances by marking them correct or incorrect is to have an additional task after the annotation/instance creation itself.
- Control instances can be injected into the data to annotate for measuring individual annotator performance and batch quality.

Agreement. Agreement can be used to gauge how reliable the annotation process can be. High agreement, however, does not automatically guarantee high-quality annotations and should be used together with other quality estimating and improving measures, like validation between annotation rounds or error rate estimation after adjudication. Krippendorff's α can be used in almost all circumstances, even for sequence tagging in the form of unitized α (Krippendorff 1995), continuous judgments, or with varying numbers of annotations per task and is therefore recommended. The agreement value targeted should be chosen beforehand, either by pilot (expert) studies or previous annotation studies annotating similar tasks. When the same number of annotators annotate each instance, Cohen's κ for two annotators or Fleiss's κ for multiple annotators can additionally be used, the latter only if annotators are randomly assigned to instances. Percent agreement should rarely be used and never the only utilized agreement measure. Correlation coefficients like Pearson's r , Spearman's ρ , or Kendall's τ should

not be used to assess reliability. Instead, Krippendorff's α or intraclass correlation is recommended as an alternative.

For a reliable estimate, agreement should be either computed on the whole dataset, or a sufficiently large ($\gtrsim 500$ instances) subset should be annotated by multiple annotators. Subset sample sizes should be statistically grounded, for instance, by computing them based on confidence intervals. They should also be justified in the dataset description. When using agreement, its usage should be reported in detail. The documentation should include which measures were used and why, how many judgments per instance were obtained, the background of the annotators, and the sample size used. Agreement values require interpretation and should not stand alone. This can be done by defining a target agreement value, for instance, based on an expert study before the annotation itself, using a sufficiently high value like 0.9, or comparing it to previous works. Using thresholds from the literature like the ones from Landis and Koch (1977) is not recommended, as these are arbitrary. Confidence intervals should be used to gauge the confidence of the agreement computation, whether they are reported as closed-form solutions given by the coefficient or via bootstrap. More recommendations concerning agreement usage can also be found in the conclusion of Lombard, Snyder-Duch, and Bracken (2002).

Quality Improvement. Annotations are often not good enough at the beginning of an annotation project. Therefore, estimating the quality and taking quality improvement steps is essential. These can be, for example, to correct low-quality instances or filter them out, improve guidelines and the annotation scheme, or train annotators. Underperforming or adversarial annotators can be removed from the annotation project if required.

Adjudication. Ideally, each instance should be annotated by multiple annotators in order to compute agreement and increase reliability via adjudication. Majority voting is a strong baseline for aggregation; using more sophisticated approaches like Dawid and Skene (1979) or MACE (Hovy et al. 2013) might be worth trying, especially in settings where individual annotators are underperforming, or spammers are potentially prevalent. Alternatively, expert curation or majority voting with experts breaking ties can be used to create a high-quality gold standard. For reproducibility and better error analysis, it is suggested to not only publish the adjudicated corpus but also annotations by individual annotators. These can then also be used to study and learn from the disagreement (Uma et al. 2021).

Error Rate Analysis. During and after the data has been annotated, it is crucial to have experts check the actual percentage of errors. The sample size should be large enough to reach a high confidence estimate, which usually requires at least 500 instances (see § 5.8) to inspect. This sample size should be computed by considering the desired statistical guarantees, for instance, confidence level and estimated precision.

Reporting. We urge authors to accurately report on the annotation process when creating new datasets. This includes, among others, annotator type and background, number of annotators, number of validators, dataset and subset sizes, agreement measures and values, adjudication methodology, and error rates. In addition to that, we suggest augmenting the dataset documentation and reproducibility checklists (which are at the time of writing mainly concerned with model training and have only a few, if any, sections for dataset quality; see § 2), often required when submitting papers to conferences, with

a section that is targeted with questions towards quality management good practices. The checklist from Kottner et al. (2011) can be a good start for checking and guiding dataset creators toward the proper use of agreement.

7. Conclusion

High-quality datasets are essential for—among others—deducing new knowledge, for policy-making, and to suggest appropriate revisions to existing theories. They are also crucial for training correct and unbiased machine learning models. If trained on datasets containing errors, inference can lead to wrong or biased predictions, which can cause material damage or even harm to other humans. These potential issues are especially relevant with the recent, widespread adoption of conversational agents based on instruction-finetuned large language models. Using datasets containing errors for evaluation can lead to incorrect estimates of task performance and, thus, to wrong conclusions when comparing models or approaches.

Quality management is an essential part of creating high-quality annotated datasets. Therefore, we set out to better understand which methods exist (RQ 1), which methods are actually applied in practice (RQ 2), and how thorough (RQ 3). For this, we surveyed the literature and inspected 591 publications introducing new datasets from which 314 reported human annotation or validation, which we annotated for their quality management usage.

We answered our first research question by summarizing good practices for annotation quality management (§ 3). These are methods suggested in the literature or commonly used during dataset creation. Then, we used the dataset of annotated publications for their quality management to investigate which methods are used frequently and which are not. Finally, we rated each publication for how well overall they conducted their quality management. We found that, on the one hand, many works implement good practices very well. On the other hand, there are still issues that need to be improved on, for instance, better usage of agreement, annotator management, quality as well as error rate estimation, or reporting. To be more precise, many papers used agreement without interpreting it, making it difficult to understand its implications. Error rate and agreement were often computed on too small sample sizes, which renders the value imprecise and less expressive. Frequently, annotation guidelines were not published, hindering reproducibility.

We conclude that many widely applicable techniques should be used more often or their use properly reported, especially iterative corpus creation as the annotation process of choice, pilot studies, validation, annotator training, qualification tests, control questions, annotation feedback, and debriefing, and maybe more complex adjudication.

We hope that our recommendations foster an adoption of good practices and an increase in dataset quality in the future.

Future Work. In this paper, we analyzed 591 scientific publications introducing new datasets and annotated them for their annotation quality management. We see several ways to build on this work. First, while we already annotated a sizeable corpus of publications, using *Papers With Code* introduced bias, limits analyzing quality management to what is reported in the paper and only contains a subset of dataset-introducing publications. Therefore, we see the next step in a larger scale effort, ideally by directly asking authors to fill out a structured survey questioning them about their quality management. While it might be difficult retroactively, it can be a good way for new datasets, especially when it is done as part of the publication and peer review process itself.

Second, it would be interesting to graph how quality management evolves over time and to analyze trends. For instance, Meyer et al. (2014) state that agreement was not used very often in their small-scale analysis at the time, but we see that, on average, it is now used quite frequently. Third, we only annotated which methods were used, but not what their actual, quantifiable impact was. Hence, conducting such studies, similar to Bayerl and Paul (2011), would be insightful; these would analyze which factors contributed to higher agreement. Fourth, as our work mainly focused on annotation and less on text production, we would like to see an extension in that direction. Fifth, in this work, we focused on analyzing scientific publications concerning their quality management. We leave analyzing other aspects for future work, for instance, how well publications adhere to aspects checked for in dataset documentation or reproducibility checklists. Sixth, it would be compelling to annotate the dataset by introducing publications on a large scale to alleviate the issues that our biased sampling might have caused. This can then also be extended to other areas of machine learning, like computer vision. Finally, we recommend that conference organizers and steering committees develop and adopt a dataset quality management checklist similar to existing ones and cover aspects like bias, intended use, or reproducibility.

8. Limitations

In this work, one of our goals was to analyze how quality management of annotated datasets is done by inspecting and annotating the publications that describe their creation. Our analysis already yields several relevant findings and common issues. We also were able to derive recommendations that future dataset creators can leverage for their own annotation projects. However, we did not analyze the impact these practices have on the resulting dataset quality. It is an interesting problem (but complex, as it requires manually analyzing not only the publications but also the datasets themselves) extension that we leave for future work.

We chose *Papers With Code* as the source of publications to annotate. While our collection approach introduces bias and does not find all publications presenting new datasets, the papers annotated this way are for popular and frequently used datasets. Otherwise, they would not be listed in *Papers With Code*. Our annotation still captures an important slice of quality management directly impacting research and state-of-the-art evaluation. However, a larger-scale annotation project would be the logical next step.

Our analysis relies on publications reporting their quality management. Hence, there might be a non-negligible underestimate of the numbers presented here. New publications are inspired by how established datasets conduct their annotation process; therefore, even if good quality management is conducted, non-reporting is also an important issue that needs to be pointed out.

Our study is limited to primarily academic datasets and may have a blind spot in the industrial field, not only in terms of data but also in terms of methods. However, this issue is difficult to alleviate, as industry datasets are often publicly unavailable.

The dataset is not intended to be used in machine learning, but is used to empirically underpin our survey. Due to limited resources and the difficulty of the annotation task, each publication was only annotated by one annotator. The impact on quality and consistency was reduced by repeatedly validating the annotations and using automatic rules to clean and improve them. Ideally, more than one set of annotations would be available to compute agreement, adjudicate, and find errors, which we recommend for the next time.

For the overall rating, when conceiving the annotation guidelines and the scheme and during annotation, we tried our best to make it as objective as possible. We still admit that the distinction between *excellent* and *sufficient* is relatively fluid. However, we argue that our definition is relatively objective for *subpar* quality management, which is the most relevant category for this work. We were relatively lenient during annotation and assigned a better rating in case of doubt. To further reduce the issue of subjectivity, we thought of alternatives like assigning scores based on the number of quality measures and their relative importance. However, we ultimately abandoned this idea because not all works can use each measure, and we would have swapped one kind of subjectivity with another.

Appendix 1. Data Collection

We use a snapshot of the *Papers With Code*¹³ data from November 26, 2022 (see Table A.1). From that, we select the *text* datasets and match them against the ACL Anthology¹⁴ with the commit 3e0966ac. While the ACL Anthology also contains backlinks to *Papers With Code*, they were still very few (≈ 100 datasets marked at the time of writing). Hence, we opted to match them by title manually.

Table A.1
File names and checksums for the *Papers With Code* data.

File Name	md5
datasets.json.gz	57193271ad26d827da3666e54e3c59dc
papers-with-abstracts.json.gz	4531a8b4bfbe449d2a9b87cc6a4869b5
links-between-papers-and-code.json.gz	424f1b2530184d3336cc497db2f965b2

Appendix 2. Guidelines

This annotation project aims to analyze how quality management is conducted in the wild. In the following, we describe the different aspects we annotate.

2.1 Manual Annotation

We are mainly interested in analyzing works that use human annotators. Therefore, we annotate whether a dataset involves humans as either annotators or validators.

2.2 Task Type

We see two broad categories of tasks that require different quality management methods.

Annotation This encompasses annotation projects where annotators provide labels, for instance, text classification, named entity recognition, annotating entailment for

¹³ <https://github.com/paperswithcode/paperswithcode-data>.
¹⁴ <https://github.com/acl-org/acl-anthology>.

natural language inference, or selecting the right question from a given set for question answering.

Text Production This encompasses annotation projects where annotators produce text. This can be, for instance, when writing surface forms that are later annotated. Other tasks include summarization, question answering, dialogues, and natural language generation.

A dataset publication can use both task types, for example, when creating questions and selecting the correct answer from a predefined pool or for natural language inference, where the clauses are first written and then labeled for their entailment.

2.3 Annotators

Expert We consider an annotator an expert if they annotate due to their domain knowledge or prior experience with the task

Contractor We consider an annotator a contractor if they are hired individually, for instance, student helpers or freelancers via platforms like Upwork or Prolific. The project managers usually know them by name and can directly interact with them. They can be managed on a more fine-grained level compared to crowdworkers.

Crowd Crowdworkers are annotators who participate via platforms like Crowdfunder or Amazon Mechanical Turk. Annotation is usually done in the form of micro-tasks. The annotators are relatively anonymous. There are often tens or hundreds of different annotators, each annotating only a small part of the overall data.

Volunteer Volunteers are annotators who help for free and are not required to do so. This, for instance, excludes students who annotate as part of their coursework.

2.4 Quality Management Methods

2.4.1 Annotation Process

Iterative Annotation Process Mentions that an iterative feedback loop is used as the annotation process.

Pilot Study It is mentioned that one or more pilot studies have been performed.

Data Filtering Data is filtered before annotation via automatic or manual checks.

Validation Mentions an explicit validation step. See Appendix 2.9.

Indirect Annotation The annotation process has several steps, where the later ones indirectly validate earlier ones.

2.4.2 Annotator Management

Annotator Training Training of annotators is mentioned.

Qualification Filter It is mentioned that annotators are filtered out by criteria like native language, geographic location, previous acceptance rates, number of previously completed tasks, etc.

Qualification Test It is mentioned that annotators had to take a qualification test before being allowed to participate in the annotation process itself.

Monetary Incentive Give annotators additional payments if their quality is exceptional.

2.4.3 Quality Estimation

Agreement Uses at least one agreement measure. This must have been used for the annotation process or validation, not the pilot study. See Appendix 2.8.

Error Rate Computes the error rate for the final, adjudicated corpus. See Appendix 2.11.

Control Questions Injects control questions for which the answer is known to estimate annotator and task performance.

2.4.4 Rectifying Measures

Guideline Refinement Mentions that guidelines and annotation schemes are refined.

Correction Mentions that instances are improved and corrected.

Annotator Debriefing Annotators give feedback to improve the annotation process.

Give Annotators Feedback Annotators are given feedback to improve their annotation quality.

Agreement Filter Instances are filtered out if agreement is too low.

Annotator Deboarding Annotators are removed from the labor pool if their quality is deemed insufficient.

Manual Filter Instances are filtered out manually if agreement is too low.

Time Filter Instances are filtered out if annotators annotate improbably quickly.

Automatic Checks Automatic checks are applied, for instance, spell checking or hand-crafted rules.

2.5 Adjudication

Adjudication describes the process of merging multiple annotations per instance into a single one.

Majority Voting The label assigned by at least half of the annotators is chosen. We also count adjudication as majority voting if all annotators must agree in the analysis, but label it as *TotalAgreement*.

Manual Tie Breaking A human annotator manually inspects instances without a majority and curates them. This adjudication method should be annotated together with *Majority Voting*.

Dawid-Skene This is an aggregation model that uses probabilistic graphical models to describe the expertise of the annotators.

MACE This is an aggregation model that uses probabilistic graphical models to describe the expertise and likeliness of being a spammer of the annotators.

Manual Curation A human annotator manually inspects and curates instances.

N/A If there is only one annotation per instance or the task type is text production.

? No mention of adjudication is found in the publication, but adjudication must have happened, e.g., because the publication mentioned more than one annotation per label.

If the task type is only text production, just enter N/A or leave the field empty; if annotation + text production, enter ? or the mentioned one. If you encounter new or different adjudication procedures, then please add them to the tagset.

2.6 Guidelines Available

For reproducibility and to judge the quality of the annotation process, it is crucial that the guidelines are available. We consider guidelines available either in the publication, appendix, or supplementary material,

- a detailed annotation tagset/task/scheme description
- a screenshot of the annotation interface with a task description for the annotators
- or the guidelines itself

are given. We only check the external supplementary material if it is referred to in the publication. In case the supplementary material is mentioned but not findable in the ACL anthology, we consider guidelines not to be available.

2.7 Overall Judgment

We assign an overall rating to each publication having human annotators based on their quality management conducted and reported. The grades are in three categories:

Excellent Does most of the following: uses the iterative annotation process, trains annotators, computes agreement and error rate, performs extensive validation, and does continuous human inspection.

Sufficient Uses some of the recommended techniques, but not as extensive as excellent. Has at least some validation and manual inspection.

Subpar No agreement, validation, manual inspection error rate, or other quality management performed and reported. The data quality, at most, relies on aggregation of multiple annotations.

2.8 Agreement

For each agreement value that is reported, create a new agreement annotation. Agreement used in pilot studies should not be entered; we are only interested in values computed for the final dataset.

2.8.1 Measure Name. Enter the name of the measure. We are at least interested in the following:

- Percent Agreement
- Cohen's κ
- Fleiss's κ
- Krippendorff's α
- Krippendorff's α unitized
- Pearson's r

- Spearman's ρ
- Kendall's τ
- Intraclass correlation coefficient
- Precision
- Recall
- F1

Enter ? if it is unclear what the agreement measure is. If you encounter new, different agreement measures, then please add them to the tagset.

2.8.2 Value. Enter the agreement value that is reported. If no value is reported, but the use of agreement is, fill in as much as possible and enter -1 .

2.8.3 Inspection Size. Enter the size of the subset that is used to compute agreement and the overall dataset size. If the agreement is computed on the whole dataset, enter 0 for both sample and total sizes.

2.8.4 Interpretation. We annotate the interpretation that is given together with the agreement value. We are at least interested in the following works that give ranges for agreement measures and their interpretation.

Landis *The Measurement of Observer Agreement for Categorical Data* by J. Richard Landis and Gary G. Koch, 1977.

Krippendorff *Validity in Content Analysis* by Klaus Krippendorff, 1980.

If you encounter new, different works referenced that give interpretations, then please add them to the tagset. We are also interested in

Custom Interpretation States that their agreement shows a certain level of quality, for instance, *sufficient, high, good* without referencing a work from the literature.

Compares To Previous Mentions a dataset that is similar to the one presented and compares its agreement to its predecessor.

2.9 Validation

We are interested in whether validation is done and who did the validation, if any.

2.10 Validators

The labels for who is validating are the same as for annotators.

2.10.1 Inspection Size. Enter the size of the subset that is validated, as well as the overall dataset size. If the complete dataset is validated, enter 0 for both sample and total sizes.

2.11 Error Rate

The error rate is the number of incorrect instances divided by the total number of instances in the dataset. We annotate it if it is computed on the adjudicated dataset. It is usually computed on a subset of instances.

2.11.1 *Value*. Enter the error rate value that is reported. If no value is reported, but the error rate is used, fill in as much as possible and enter -1.

2.11.2 *Inspection Size*. Enter the size of the subset that is used to compute the error rate as well as the overall dataset size. If the error rate is computed on the whole dataset, enter 0 for both sample and total sizes.

Appendix 3. Correlation

In the following, we give an example where correlation between ratings is high but agreement is low. We assume two annotators rating four items on a scale in [1, 5]:

Item		a	b	c	d
Judge	A	1	2	3	4
	B	3	4	5	5

The resulting correlation scores are:

Pearson’s ρ	Spearman’s ρ	Kendall τ	ICC1	ICC2	ICC3
0.944	0.949	0.913	0.204	0.418	0.903

It can be seen that standard correlation measures show very high correlation, while Intraclass Correlation scores are comparatively low.

Acknowledgments

We thank Falko Helm, Ivan Habernal, Ji-Ung Lee, Qian Ruan, Nils Dycke, Max Glockner, and our anonymous reviewers for the fruitful discussions and helpful feedback that improved this article. This work has been funded by the German Research Foundation (DFG) as part of the Evidence (grant GU 798/27-1) and the PEER projects (grant GU 798/28-1).

References

Alex, Bea, Claire Grover, Rongzhou Shen, and Mijail Kabadjov. 2010. Agile corpus annotation in practice: An overview of manual and automatic annotation of CVs. In *Proceedings of the Fourth Linguistic Annotation Workshop*, pages 29–37.

Allan, Donner. 1999. Sample size requirements for interval estimation of the

intraclass kappa statistic. *Communications in Statistics - Simulation and Computation*, 28(2):415–429. <https://doi.org/10.1080/03610919908813557>

Alt, Christoph, Aleksandra Gabryszak, and Leonhard Hennig. 2020. TACRED Revisited: A thorough evaluation of the TACRED relation extraction task. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1558–1569. <https://doi.org/10.18653/v1/2020.acl-main.142>

Amidei, Jacopo, Paul Piwek, and Alistair Willis. 2019. Agreement is overrated: A plea for correlation to assess human evaluation reliability. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 344–354. <https://doi.org/10.18653/v1/W19-8642>

Aroyo, Lora and Chris Welty. 2015. Truth is a lie: Crowd truth and the seven myths of

- human annotation. *AI Magazine*, 36(1):15–24. <https://doi.org/10.1609/aimag.v36i1.2564>
- Artstein, Ron and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596. <https://doi.org/10.1162/coli.07-034-R2>
- Bakeman, Roger, Duncan McArthur, Vicenç Quera, and Byron F. Robinson. 1997. Detecting sequential patterns and determining their reliability with fallible observers. *Psychological Methods*, 2(4):357–370. <https://doi.org/10.1037/1082-989X.2.4.357>
- Banerjee, Mousumi, Michelle Capozzoli, Laura McSweeney, and Debajyoti Sinha. 1999. Beyond kappa: A review of interrater agreement measures. *Canadian Journal of Statistics*, 27(1):3–23. <https://doi.org/10.2307/3315487>
- Bareket, Dan and Reut Tsarfaty. 2021. Neural modeling for named entities and morphology (NEMO2). *Transactions of the Association for Computational Linguistics*, 9:909–928. <https://doi.org/10.1162/tacl.a.00404>
- Bastan, Mohaddeseh, Mahnaz Koupaee, Youngseo Son, Richard Sicoli, and Niranjana Balasubramanian. 2020. Author’s sentiment prediction. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 604–615. <https://doi.org/10.18653/v1/2020.coling-main.52>
- Bayerl, Petra Saskia and Karsten Ingmar Paul. 2011. What determines inter-coder agreement in manual annotations? A meta-analytic investigation. *Computational Linguistics*, 37(4):699–725. <https://doi.org/10.1162/COLI.a.00074>
- Behrens, Heike. 2008. *Corpora in Language Acquisition Research: History, Methods, Perspectives*. volume 6 of *Trends in Language Acquisition Research*. John Benjamins Publishing Company, Amsterdam, The Netherlands. <https://doi.org/10.1075/tilar.6>
- Bender, Emily M. and Batya Friedman. 2018. Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6:587–604. <https://doi.org/10.1162/tacl.a.00041>
- Bland, J. M. and D. G. Altman. 1986. Statistical methods for assessing agreement between two methods of clinical measurement. *Lancet*, 1(8476):307–310. [https://doi.org/10.1016/S0140-6736\(86\)90837-8](https://doi.org/10.1016/S0140-6736(86)90837-8)
- Bowman, Samuel R., Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642. <https://doi.org/10.18653/v1/D15-1075>
- Brandsen, Alex, Suzan Verberne, Milco Wansleben, and Karsten Lambers. 2020. Creating a dataset for named entity recognition in the archaeology domain. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4573–4577.
- Button, Katherine S., John P. A. Ioannidis, Claire Mokrysz, Brian A. Nosek, Jonathan Flint, Emma S. J. Robinson, and Marcus R. Munafò. 2013. Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5):365–376. <https://doi.org/10.1038/nrn3475>, PubMed: 23571845
- Callison-Burch, Chris and Mark Dredze. 2010. Creating speech and language data with Amazon’s Mechanical Turk. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*, pages 1–12.
- Carletta, Jean. 1996. Assessing agreement on classification tasks: The kappa statistic. *Computational Linguistics*, 22(2):249–254.
- Cer, Daniel, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. SemEval-2017 Task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14. <https://doi.org/10.18653/v1/S17-2001>
- Checco, Alessandro, Kevin Roitero, Eddy Maddalena, Stefano Mizzaro, and Gianluca Demartini. 2017. Let’s agree to disagree: Fixing agreement measures for crowdsourcing. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, pages 11–20. <https://doi.org/10.1609/hcomp.v5i1.13306>
- Chen, Zhiyu, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. 2021. FinQA: A dataset of numerical reasoning over financial data. In *Proceedings of the 2021 Conference on Empirical Methods in Natural*

- Language Processing*, pages 3697–3711. <https://doi.org/10.18653/v1/2021.emnlp-main.300>
- Cohen, Jacob. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20(1):37–46. <https://doi.org/10.1177/001316446002000104>
- Daniel, Florian, Pavel Kucherbaev, Cinzia Cappiello, Boualem Benatallah, and Mohammad Allahbakhsh. 2019. Quality control in crowdsourcing: A survey of quality attributes, assessment techniques, and assurance actions. *ACM Computing Surveys*, 51(1):1–40. <https://doi.org/10.1145/3148148>
- Dawid, A. P. and A. M. Skene. 1979. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Applied Statistics*, 28(1):20–28. <https://doi.org/10.2307/2346806>
- Demszky, Dorottya, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. GoEmotions: A dataset of fine-grained emotions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4040–4054. <https://doi.org/10.18653/v1/2020.acl-main.372>
- Dickinson, Markus and W. Detmar Meurers. 2003. Detecting inconsistencies in treebanks. In *Proceedings of the Second Workshop on Treebanks and Linguistic Theories*, pages 1–12.
- Donner, Allan and Michael Eliasziw. 1992. A goodness-of-fit approach to inference procedures for the kappa statistic: Confidence interval construction, significance-testing and sample size estimation. *Statistics in Medicine*, 11(11):1511–1519. <https://doi.org/10.1002/sim.4780111109>, PubMed: 1410963
- Dror, Rotem, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. The hitchhiker’s guide to testing statistical significance in natural language processing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1383–1392. <https://doi.org/10.18653/v1/P18-1128>
- Ebel, Robert L. 1951. Estimation of the reliability of ratings. *Psychometrika*, 16(4):407–424. <https://doi.org/10.1007/BF02288803>
- Edwards, Christopher, Heather Allen, and Crispin Chamunyonga. 2021. Correlation does not imply agreement: A cautionary tale for researchers and reviewers. *Sonography*, 8(4):185–190. <https://doi.org/10.1002/sono.12276>
- Efron, B. and R. Tibshirani. 1986. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, 1(1):54–75. <https://doi.org/10.1214/ss/1177013815>
- Ferracane, Elisa, Greg Durrett, Junyi Jessy Li, and Katrin Erk. 2021. Did they answer? Subjective acts and intents in conversational discourse. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1626–1644. <https://doi.org/10.18653/v1/2021.naacl-main.129>
- Fisher, Roland A. 1925. *Statistical Methods for Research Workers*. Oliver and Boyd, Edinburgh.
- Fleiss, Joseph L. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5):378–382. <https://doi.org/10.1037/h0031619>
- Fleiss, Joseph L., Bruce Levin, and Myunghee Cho Paik. 2003. *Statistical Methods for Rates and Proportions*, 1st edition. Wiley Series in Probability and Statistics. Wiley. <https://doi.org/10.1002/0471445428>
- Gebru, Timnit, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92. <https://doi.org/10.1145/3458723>
- Geva, Mor, Yoav Goldberg, and Jonathan Berant. 2019. Are we modeling the task or the annotator? An investigation of annotator bias in natural language understanding datasets. 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 1161–1166. <https://doi.org/10.18653/v1/D19-1107>
- Ghosal, Deepanway, Siqi Shen, Navonil Majumder, Rada Mihalcea, and Soujanya Poria. 2022. CICERO: A dataset for contextualized commonsense inference in dialogues. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5010–5028. <https://doi.org/10.18653/v1/2022.acl-long.344>
- Gildea, Daniel, Min-Yen Kan, Nitin Madnani, Christoph Teichmann, and Martín Villalba. 2018. The ACL anthology: Current state

- and future directions. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 23–28. <https://doi.org/10.18653/v1/W18-2504>
- Goodhart, C. A. E. 1984. *Problems of Monetary Management: The UK Experience*. Macmillan Education UK, London. https://doi.org/10.1007/978-1-349-17295-5_4
- Govindarajan, Venkata Subrahmanyam, Benjamin Chen, Rebecca Warholc, Katrin Erk, and Junyi Jessy Li. 2020. Help! Need advice on identifying advice. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5295–5306. <https://doi.org/10.18653/v1/2020.emnlp-main.427>
- Gururangan, Suchin, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don't stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360. <https://doi.org/10.18653/v1/2020.acl-main.740>
- Hardt, Moritz and Benjamin Recht. 2022. *Patterns, Predictions, and Actions: Foundations of Machine Learning*. Princeton University Press, Princeton.
- Harris, Christopher. 2011. You're hired! An examination of crowdsourcing incentive models in human resource tasks. In *Proceedings of the Workshop on Crowdsourcing for Search and Data Mining (CSDM) at the Fourth ACM International Conference on Web Search and Data Mining (WSDM)*, pages 15–18.
- Haselbach, Boris, Kerstin Eckart, Wolfgang Seeker, Kurt Eberle, and Ulrich Heid. 2012. Approximating theoretical linguistics classification in real data: The case of German “nach” particle verbs. In *Proceedings of COLING 2012*, pages 1113–1128.
- Hayes, Andrew F. and Klaus Krippendorff. 2007. Answering the call for a standard reliability measure for coding data. *Communication Methods and Measures*, 1(1):77–89. <https://doi.org/10.1080/19312450709336664>
- Ho, Chien Ju, Aleksandrs Slivkins, Siddharth Suri, and Jennifer Wortman Vaughan. 2015. Incentivizing high quality crowdwork. In *Proceedings of the 24th International Conference on World Wide Web*, pages 419–429. <https://doi.org/10.1145/2736277.2741102>
- Holland, Sarah, Ahmed Hosny, Sarah Newman, Joshua Joseph, and Kasia Chmielinski. 2018. The dataset nutrition label: A framework to drive higher data quality standards. *arXiv*, 1805(03677):1–21.
- Horbach, Andrea, Yuning Ding, and Torsten Zesch. 2017. The influence of spelling errors on content scoring performance. In *Proceedings of the 4th Workshop on Natural Language Processing Techniques for Educational Applications (NLPTEA 2017)*, pages 45–53.
- Hovy, Dirk, 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. <https://doi.org/10.3115/v1/P14-2062>
- Hovy, Dirk, Barbara Plank, and Anders Søgaard. 2014. Experiments with crowdsourced re-annotation of a POS tagging data set. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 377–382.
- Hovy, Eduard and Julia Lavid. 2010. Towards a ‘science’ of corpus annotation: A new methodological challenge for corpus linguistics. *International Journal of Translation Studies*, 22:13–36.
- Hsueh, Pei Yun, 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. <https://doi.org/10.3115/1564131.1564137>
- Hutchinson, Ben, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, and Margaret Mitchell. 2021. Towards accountability for machine learning datasets: Practices from software engineering and infrastructure. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 560–575. <https://doi.org/10.1145/3442188.3445918>
- Ide, Nancy and James Pustejovsky, editors. 2017. *Handbook of Linguistic Annotation*. Springer Netherlands, Dordrecht. <https://doi.org/10.1007/978-94-024-0881-2>
- Jamison, Emily and Iryna Gurevych. 2015. Noise or additional information? Leveraging crowdsource annotation item agreement for natural language tasks. In *Proceedings of the 2015 Conference on*

- Empirical Methods in Natural Language Processing*, pages 291–297. <https://doi.org/10.18653/v1/D15-1035>
- Kim, Juyong, Jeremy C. Weiss, and Pradeep Ravikumar. 2022. Context-sensitive spelling correction of clinical text via conditional independence. In *Proceedings of the Conference on Health, Inference, and Learning* volume 174 of *Proceedings of Machine Learning Research*, pages 234–247.
- Kirk, Hannah, Bertie Vidgen, Paul Rottger, Tristan Thrush, and Scott Hale. 2022. HATEMOJI: A test suite and adversarially-generated dataset for benchmarking and detecting emoji-based hate. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1352–1368. <https://doi.org/10.18653/v1/2022.naacl-main.97>
- Klie, Jan Christoph, Michael Bugert, Beto Bouldosa, 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.
- Klie, Jan Christoph, Bonnie Webber, and Iryna Gurevych. 2023. Annotation error detection: Analyzing the past and present for a more coherent future. *Computational Linguistics*, 49(1):157–198. <https://doi.org/10.1162/coli.a.00464>
- Kottner, Jan, Laurent Audigé, Stig Brorson, Allan Donner, Byron J. Gajewski, Asbjørn Hróbjartsson, Chris Roberts, Mohamed Shoukri, and David L. Streiner. 2011. Guidelines for reporting reliability and agreement studies (GRRAS) were proposed. *Journal of Clinical Epidemiology*, 64(1):96–106. <https://doi.org/10.1016/j.jclinepi.2010.03.002>, PubMed: 21130355
- Krippendorff, Klaus. 1980. *Content Analysis: An Introduction to Its Methodology*. SAGE, Los Angeles.
- Krippendorff, Klaus. 1995. On the reliability of unitizing continuous data. *Sociological Methodology*, 25:47–76. <https://doi.org/10.2307/271061>
- Krippendorff, Klaus. 2004. Reliability in content analysis: Some common misconceptions and recommendations. *Human Communication Research*, 30(3):411–433. <https://doi.org/10.1111/j.1468-2958.2004.tb00738.x>
- Krippendorff, Klaus. 2011. Agreement and information in the reliability of coding. *Communication Methods and Measures*, 5(2):93–112. <https://doi.org/10.1080/19312458.2011.568376>
- Krippendorff, Klaus, Yann Mathet, Stéphane Bouvry, and Antoine Widlöcher. 2016. On the reliability of unitizing textual continua: Further developments. *Quality & Quantity*, 50(6):2347–2364. <https://doi.org/10.1007/s11135-015-0266-1>
- Kummerfeld, Jonathan K. 2021. Quantifying and avoiding unfair qualification labour in crowdsourcing. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 343–349. <https://doi.org/10.18653/v1/2021.acl-short.44>
- Kummerfeld, Jonathan K., Sai R. Gouravajhala, Joseph J. Peper, Vignesh Athreya, Chulaka Gunasekara, Jatin Ganhotra, Siva Sankalp Patel, Lazaros C. Polymenakos, and Walter Lasecki. 2019. A large-scale corpus for conversation disentanglement. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3846–3856. <https://doi.org/10.18653/v1/P19-1374>
- Květoň, Pavel and Karel Oliva. 2002. (Semi-)automatic detection of errors in PoS-tagged corpora. In *COLING 2002: The 19th International Conference on Computational Linguistics*, pages 1–7. <https://doi.org/10.3115/1072228.1072249>
- Landis, J. Richard and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174. <https://doi.org/10.2307/2529310>, PubMed: 843571
- Lease, Matthew. 2011. On quality control and machine learning in crowdsourcing. In *Proceedings of the 11th AAAI Conference on Human Computation, AAAIWS'11-11*, pages 97–102.
- Lindahl, Anna, Lars Borin, and Jacobo Rouces. 2019. Towards assessing argumentation annotation - a first step. In *Proceedings of the 6th Workshop on Argument Mining*, pages 177–186. <https://doi.org/10.18653/v1/W19-4520>
- Lombard, Matthew, Jennifer Snyder-Duch, and Cheryl Campanella Bracken. 2002. Content analysis in mass communication: Assessment and reporting of intercoder reliability. *Human Communication Research*,

- 28(4):587–604. <https://doi.org/10.1111/j.1468-2958.2002.tb00826.x>
- McCoy, Tom, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448. <https://doi.org/10.18653/v1/P19-1334>
- Meyer, Christian M., Margot Mieskes, Christian Stab, and Iryna Gurevych. 2014. DKPro agreement: An open-source Java library for measuring inter-rater agreement. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: System Demonstrations*, pages 105–109.
- Mihaylov, Todor, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? A new dataset for open book question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391. <https://doi.org/10.18653/v1/D18-1260>
- Mintz, Mike, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011. <https://doi.org/10.3115/1690219.1690287>
- Monarch, Robert. 2021. *Human-in-the-Loop Machine Learning: Active Learning and Annotation for Human-Centered AI*. Manning Publications.
- Mostafazadeh, Nasrin, Aditya Kalyanpur, Lori Moon, David Buchanan, Lauren Berkowitz, Or Biran, and Jennifer Chu-Carroll. 2020. GLUCOSE: Generalized and contextualized story explanations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4569–4586. <https://doi.org/10.18653/v1/2020.emnlp-main.370>
- Neuendorf, Kimberly A. 2016. *The Content Analysis Guidebook*. SAGE, Thousand Oaks, California. <https://doi.org/10.4135/9781071802878>
- Northcutt, Curtis, Lu Jiang, and Isaac Chuang. 2021. Confident learning: Estimating uncertainty in dataset labels. *Journal of Artificial Intelligence Research*, 70:1373–1411. <https://doi.org/10.1613/jair.1.12125>
- Northcutt, Curtis G., Anish Athalye, and Jonas Mueller. 2021. Pervasive label errors in test sets destabilize machine learning benchmarks. In *35th Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, pages 1–13.
- Ouyang, Long, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, pages 1–15.
- Parmar, Mihir, Swaroop Mishra, Mor Geva, and Chitta Baral. 2023. Don't blame the annotator: Bias already starts in the annotation instructions. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1779–1789. <https://doi.org/10.18653/v1/2023.eacl-main.130>
- Parrish, Alicia, William Huang, Omar Agha, Soo-Hwan Lee, Nikita Nangia, Alexia Warstadt, Karmanya Aggarwal, Emily Allaway, Tal Linzen, and Samuel R. Bowman. 2021. Does putting a linguist in the loop improve NLU data collection?. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4886–4901. <https://doi.org/10.18653/v1/2021.findings-emnlp.421>
- Passonneau, Rebecca J. and Bob Carpenter. 2014. The benefits of a model of annotation. *Transactions of the Association for Computational Linguistics*, 2:311–326. https://doi.org/10.1162/tac1_a_00185
- Paun, Silviu, Bob Carpenter, Jon Chamberlain, Dirk Hovy, Udo Kruschwitz, and Massimo Poesio. 2018. Comparing Bayesian models of annotation. *Transactions of the Association for Computational Linguistics*, 6(0):571–585. https://doi.org/10.1162/tac1_a_00040
- Peters, Matthew E., Sebastian Ruder, and Noah A. Smith. 2019. To tune or not to tune? Adapting pretrained representations to diverse tasks. In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 7–14. <https://doi.org/10.18653/v1/W19-4302>
- Piskorski, Jakub, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023. SemEval-2023 Task 3: Detecting the category, the framing, and

- the persuasion techniques in online news in a multi-lingual setup. In *Proceedings of the 17th International Workshop on Semantic Evaluation*, pages 2343–2361. <https://doi.org/10.18653/v1/2023.semeval-1.317>
- Popping, R. 1988. On agreement indices for nominal data. In *Sociometric Research*. Palgrave Macmillan UK, London, pages 90–105. https://doi.org/10.1007/978-1-349-19051-5_6
- Powers, David. 2011. Evaluation: From precision, recall and F-measure to ROC, informedness, markedness & correlation. *Journal of Machine Learning Technologies*, 2(1):37–63.
- Prasad, Rashmi, Nikhil Dinesh, Alan Lee, Eleni Miltsakaki, Livio Robaldo, Aravind Joshi, and Bonnie Webber. 2008. The Penn Discourse TreeBank 2.0. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, pages 2961–2968.
- Pushkarna, Mahima, Andrew Zaldivar, and Oddur Kjartansson. 2022. Data cards: Purposeful and transparent dataset documentation for responsible AI. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1776–1826. <https://doi.org/10.1145/3531146.3533231>
- Pustejovsky, J. and Amber Stubbs. 2013. *Natural Language Annotation for Machine Learning*. O'Reilly Media, Sebastopol, California, USA.
- Qian, Kun, Ahmad Beirami, Zhouhan Lin, Ankita De, Alborz Geramifard, Zhou Yu, and Chinnadhurai Sankar. 2021. Annotation inconsistency and entity bias in MultiWOZ. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 326–337. <https://doi.org/10.18653/v1/2021.sigdial-1.35>
- Ranganathan, Priya, Cs Pramesh, and Rakesh Aggarwal. 2017. Common pitfalls in statistical analysis: Measures of agreement. *Perspectives in Clinical Research*, 8(4):187–191. https://doi.org/10.4103/picr.PICR_123.17, PubMed: 29109937
- Reddy, Siva, Danqi Chen, and Christopher D. Manning. 2019. CoQA: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266. https://doi.org/10.1162/tac1.a_00266
- Reiss, Frederick, Hong Xu, Bryan Cutler, Karthik Muthuraman, and Zachary Eichenberger. 2020. Identifying incorrect labels in the CoNLL-2003 corpus. In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 215–226. <https://doi.org/10.18653/v1/2020.conll-1.16>
- Roh, Yuji, Geon Heo, and Steven Euijong Whang. 2021. A survey on data collection for machine learning: A big data - AI integration perspective. *IEEE Transactions on Knowledge and Data Engineering*, 33(4):1328–1347. <https://doi.org/10.1109/TKDE.2019.2946162>
- Sabou, Marta, Kalina Bontcheva, Leon Derczynski, and Arno Scharl. 2014. Corpus annotation through crowdsourcing: Towards best practice guidelines. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 859–866.
- Sambasivan, Nithya, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Kumar Paritosh, and Lora Mois Aroyo. 2021. “Everyone wants to do the model work, not the data work”: Data cascades in high-stakes AI. In *SIGCHI*, pages 1–21. <https://doi.org/10.1145/3411764.3445518>
- Schreibman, Susan, Ray Siemens, and John Unsworth, editors. 2004. *A Companion to Digital Humanities*. Blackwell Publishing Ltd, Malden, Massachusetts, USA. <https://doi.org/10.1002/9780470999875>
- Scott, William A. 1955. Reliability of content analysis: The case of nominal scale coding. *The Public Opinion Quarterly*, 19(3):321–325. <https://doi.org/10.1086/266577>
- Sheshadri, Aashish and Matthew Lease. 2013. SQUARE: A benchmark for research on computing crowd consensus. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, 1156–164. <https://doi.org/10.1609/hcomp.v1i1.13088>
- Shmueli, Boaz, Jan Fell, Soumya Ray, and Lun-Wei Ku. 2021. Beyond fair pay: Ethical implications of NLP crowdsourcing. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3758–3769. <https://doi.org/10.18653/v1/2021.naacl-main.295>
- Shoukri, M. M., M. H. Asyali, and A. Donner. 2004. Sample size requirements for the design of reliability study: Review and new results. *Statistical Methods in Medical Research*, 13(4):251–271. <https://doi.org/10.1191/0962280204sm365ra>

- Shrout, Patrick E. and Joseph L. Fleiss. 1979. Intra-class correlations: Uses in assessing rater reliability. *Psychological Bulletin*, 86(2):420–428. <https://doi.org/10.1037/0033-2909.86.2.420>, PubMed: 18839484
- Sim, Julius and Chris C. Wright. 2005. The kappa statistic in reliability studies: Use, interpretation, and sample size requirements. *Physical Therapy*, 85(3):257–268. <https://doi.org/10.1093/ptj/85.3.257>, PubMed: 15733050
- Simpson, Edwin D. and Iryna Gurevych. 2019. A Bayesian approach for sequence tagging with crowds. 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 1093–1104. <https://doi.org/10.18653/v1/D19-1101>
- Singh, Shikhar, Nuan Wen, Yu Hou, Pegah Alipoormolabashi, Te-lin Wu, Xuezhe Ma, and Nanyun Peng. 2021. COM2SENSE: A commonsense reasoning benchmark with complementary sentences. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 883–898. <https://doi.org/10.18653/v1/2021.findings-acl.78>
- Snow, Rion, Brendan O’Connor, Daniel Jurafsky, and Andrew Ng. 2008. Cheap and fast – but is it good? Evaluating non-expert annotations for natural language tasks. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 254–263. <https://doi.org/10.3115/1613715.1613751>
- Socher, Richard, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Stab, Christian and Iryna Gurevych. 2014. Identifying argumentative discourse structures in persuasive essays. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 46–56. <https://doi.org/10.3115/v1/D14-1006>
- Stoica, George, Emmanouil Antonios Platanios, and Barnabas Poczos. 2021. Re-TACRED: Addressing shortcomings of the TACRED dataset. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence 2021*, pages 13843–13850. <https://doi.org/10.1609/aaai.v35i15.17631>
- Sun, Chen, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 843–852. <https://doi.org/10.1109/ICCV.2017.97>
- Suster, Simon, Stephan Tulkens, and Walter Daelemans. 2017. A short review of ethical challenges in clinical natural language processing. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 80–87. <https://doi.org/10.18653/v1/W17-1610>
- Tjong Kim Sang, Erik F. and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147. <https://doi.org/10.3115/1119176.1119195>
- Uma, Alexandra N., 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. <https://doi.org/10.1613/jair.1.12752>
- Vădineanu, Serban, Daniel Pelt, Oleh Dzyubachyk, and Joost Batenburg. 2022. An analysis of the impact of annotation errors on the accuracy of deep learning for cell segmentation. In *Proceedings of Machine Learning Research*, pages 1251–1267.
- van Stralen, K. J., F. W. Dekker, C. Zoccali, and K. J. Jager. 2012. Measuring agreement, more complicated than it seems. *Nephron Clinical Practice*, 120(3):162–167. <https://doi.org/10.1159/000337798>, PubMed: 22722148
- Vasudevan, Vijay, Benjamin Caine, Raphael Gontijo-Lopes, Sara Fridovich-Keil, and Rebecca Roelofs. 2022. When does dough become a bagel? Analyzing the remaining mistakes on ImageNet. In *Proceedings of the 36th Conference on Neural Information Processing Systems*, pages 1–15.
- Wang, Zihan, Jingbo Shang, Liyuan Liu, Lihao Lu, Jiacheng Liu, and Jiawei Han. 2019. CrossWeigh: Training named entity tagger from imperfect annotations. 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 5153–5162. <https://doi.org/10.18653/v1/D19-1519>
- Wei, Jason, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*, pages 1–46.
- Wynne, Martin, editor. 2005. *Developing Linguistic Corpora: A Guide to Good Practice*. David Brown Book Company, Oakville, Connecticut.
- Yao, Yuan, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie Zhou, and Maosong Sun. 2019. DocRED: A large-scale document-level relation extraction dataset. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 764–777. <https://doi.org/10.18653/v1/P19-1074>
- Zapf, Antonia, Stefanie Castell, Lars Morawietz, and André Karch. 2016. Measuring inter-rater reliability for nominal data – which coefficients and confidence intervals are appropriate? *BMC Medical Research Methodology*, 16(1):93–103. <https://doi.org/10.1186/s12874-016-0200-9>, PubMed: 27495131
- Zeng, Zhiqiang, Hua Shi, Yun Wu, and Zhiling Hong. 2015. Survey of natural language processing techniques in bioinformatics. *Computational and Mathematical Methods in Medicine*, 2015:1–10. <https://doi.org/10.1155/2015/674296>, PubMed: 26525745
- Zhang, Yuhao, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D. Manning. 2017. Position-aware attention and supervised data improve slot filling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 35–45. <https://doi.org/10.18653/v1/D17-1004>
- Zhao, Xinshu, Jun S. Liu, and Ke Deng. 2013. Assumptions behind intercoder reliability indices. *Annals of the International Communication Association*, 36(1):419–480. <https://doi.org/10.1080/23808985.2013.11679142>