

Design Choices for Crowdsourcing Implicit Discourse Relations: Revealing the Biases Introduced by Task Design

Valentina Pyatkin¹ Frances Yung² Merel C. J. Scholman^{2,3}
Reut Tsarfaty¹ Ido Dagan¹ Vera Demberg²

¹Bar Ilan University, Ramat Gan, Israel

²Saarland University, Saarbrücken, Germany

³Utrecht University, Utrecht, Netherlands

{pyatkiv, reut.tsarfaty}@biu.ac.il; dagan@cs.biu.ac.il

{frances, m.c.j.scholman, vera}@coli.uni-saarland.de

Abstract

Disagreement in natural language annotation has mostly been studied from a perspective of biases introduced by the annotators and the annotation frameworks. Here, we propose to analyze another source of bias—task design bias, which has a particularly strong impact on crowdsourced linguistic annotations where natural language is used to elicit the interpretation of lay annotators. For this purpose we look at implicit discourse relation annotation, a task that has repeatedly been shown to be difficult due to the relations' ambiguity. We compare the annotations of 1,200 discourse relations obtained using two distinct annotation tasks and quantify the biases of both methods across four different domains. Both methods are natural language annotation tasks designed for crowdsourcing. We show that the task design can push annotators towards certain relations and that some discourse relation senses can be better elicited with one or the other annotation approach. We also conclude that this type of bias should be taken into account when training and testing models.

1 Introduction

Crowdsourcing has become a popular method for data collection. It not only allows researchers to collect large amounts of annotated data in a shorter amount of time, but also captures human inference in natural language, which should be the goal of benchmark NLP tasks (Manning, 2006). In order to obtain reliable annotations, the crowdsourced labels are traditionally aggregated to a single label per item, using simple majority voting or annotation models that reduce noise from the data based on the disagreement among the annotators (Hovy et al., 2013; Passonneau and Carpenter,

2014). However, there is increasing consensus that disagreement in annotation cannot be generally discarded as noise in a range of NLP tasks, such as natural language inferences (De Marneffe et al., 2012; Pavlick and Kwiatkowski, 2019; Chen et al., 2020; Nie et al., 2020), word sense disambiguation (Jurgens, 2013), question answering (Min et al., 2020; Ferracane et al., 2021), anaphora resolution (Poesio and Artstein, 2005; Poesio et al., 2006), sentiment analysis (Díaz et al., 2018; Cowen et al., 2019), and stance classification (Waseem, 2016; Luo et al., 2020). Label distributions are proposed to replace categorical labels in order to represent the label ambiguity (Aroyo and Welty, 2013; Pavlick and Kwiatkowski, 2019; Uma et al., 2021; Dumitrache et al., 2021).

There are various reasons behind the ambiguity of linguistic annotations (Dumitrache, 2015; Jiang and de Marneffe, 2022). Aroyo and Welty (2013) summarize the sources of ambiguity into three categories: the text, the annotators, and the annotation scheme. In downstream NLP tasks, it would be helpful if models could detect possible alternative interpretations of ambiguous texts, or predict a distribution of interpretations by a population. In addition to the existing work on the disagreement due to annotators' bias, the effect of annotation frameworks has also been studied, such as the discussion on whether entailment should include pragmatic inferences (Pavlick and Kwiatkowski, 2019), the effect of the granularity of the collected labels (Chung et al., 2019), or the system of labels that categorize the linguistic phenomenon (Demberg et al., 2019). In this work, we examine the effect of task design bias, which is independent of the annotation framework, on the quality of crowdsourced annotations. Specifically,




Figure 1: Example of two relational arguments (S1 and S2) and the DC and QA annotation in the middle.

we look at inter-sentential implicit discourse relation (DR) annotation, i.e., semantic or pragmatic relations between two adjacent sentences without a discourse connective to which the sense of the relation can be attributed. Figure 1 shows an example of an implicit relation that can be annotated as *Conjunction* or *Result*.

Implicit DR annotation is arguably the hardest task in discourse parsing. Discourse coherence is a feature of the mental representation that readers form of a text, rather than of the linguistic material itself (Sanders et al., 1992). Discourse annotation thus relies on annotators’ interpretation of a text. Further, relations can often be interpreted in various ways (Rohde et al., 2016), with multiple valid readings holding at the same time. These factors make discourse relation annotation, especially for implicit relations, a particularly difficult task. We collect 10 different annotations per DR, thereby focusing on distributional representations, which are more informative than categorical labels.

Since DR annotation labels are often abstract terms that are not easily understood by lay individuals, we focus on “natural language” task designs. Decomposing and simplifying an annotation task, where the DR labels can be obtained indirectly from the natural language annotations, has been shown to work well for crowdsourcing (Chang et al., 2016; Scholman and Demberg, 2017; Pyatki et al., 2020). Crowdsourcing with natural language has become increasingly popular. This includes tasks such as NLI (Bowman et al., 2015), SRL (Fitzgerald et al., 2018), and QA (Rajpurkar et al., 2018). This trend is further visible in modeling approaches that cast traditional structured prediction tasks into NL tasks, such as for co-reference (Aralikatte et al., 2021), discourse comprehension (Ko et al., 2021), or bridging anaphora (Hou, 2020;

Elazar et al., 2022). It is therefore of interest to the broader research community to see how task design biases can arise, even when the tasks are more accessible to the lay public.

We examine two distinct natural language crowdsourcing discourse relation annotation tasks (Figure 1): Yung et al. (2019) derive relation labels from discourse connectives (DCs) that crowd workers insert; Pyatkin et al. (2020) derive labels from Question Answer (QA) pairs that crowd workers write. Both task designs employ natural language annotations instead of labels from a taxonomy. The two task designs, DC and QA, are used to annotate 1,200 implicit discourse relations in 4 different domains. This allows us to explore how the task design impacts the obtained annotations, as well as the biases that are inherent to each method. To do so we showcase the difference of various inter-annotator agreement metrics on annotations with distributional and aggregated labels.

We find that both methods have strengths and weaknesses in identifying certain types of relations. We further see that these biases are also affected by the domain. In a series of discourse relation classification experiments, we demonstrate the benefits of collecting annotations with mixed methodologies, we show that training with a soft loss with distributions as targets improves model performance, and we find that cross-task generalization is harder than cross-domain generalization.

The outline of the paper is as follows. We introduce the notion of task design bias and analyze its effect on crowdsourcing implicit DRs, using two different task designs (Section 3–4). Next, we quantify strengths and weaknesses of each method using the obtained annotations, and suggest ways to reduce task bias (Section 5). Then we look at genre-specific task bias (Section 6). Lastly, we demonstrate the task bias effect on DR classification performance (Section 7).

2 Background

2.1 Annotation Biases

Annotation tends to be an inherently ambiguous task, often with multiple possible interpretations and without a single ground truth (Aroyo and Welty, 2013). An increasing amount of research has studied annotation disagreements and biases.

Prior studies have focused on how crowdworkers can be biased. Worker biases are subject to various factors, such as educational or cultural background, or other demographic characteristics. Prabhakaran et al. (2021) point out that for more subjective annotation tasks, the socio-demographic background of annotators contributes to multiple annotation perspectives and argue that label aggregation obfuscates such perspectives. Instead, soft labels are proposed, such as the ones provided by the CrowdTruth method (Dumitrache et al., 2018), which require multiple judgments to be collected per instance (Uma et al, 2021). Bowman and Dahl (2021) suggest that annotations that are subject to bias from methodological artifacts should not be included in benchmark datasets. In contrast, Basile et al. (2021) argue that all kinds of human disagreements should be predicted by NLU models and thus included in evaluation datasets.

In contrast to annotator bias, a limited amount of research is available on bias related to the formulation of the task. Jakobsen et al. (2022) show that argument annotations exhibit widely different levels of social group disparity depending on which guidelines the annotators followed. Similarly, Buechel and Hahn (2017a,b) study different design choices for crowdsourcing emotion annotations and show that the perspective that annotators are asked to take in the guidelines affects annotation quality and distribution. Jiang et al. (2017) study the effect of workflow for paraphrase collection and found that examples based on previous contributions prompt workers to produce more diverging paraphrases. Hube et al. (2019) show that biased subjective judgment annotations can be mitigated by asking workers to think about responses other workers might give and by making workers aware of their possible biases. Hence, the available research suggests that task design can affect the annotation output in various ways. Further research studied the collection of multiple labels: Jurgens (2013) compares between selection and scale rating and finds that workers would choose an additional label for a word sense labelling task. In contrast, Scholman and Demberg (2017) find that workers usually opt not to provide an additional DR label even when allowed. Chung et al. (2019) compare various label collection methods including single / multiple labelling, ranking, and probability assignment. We focus on the biases in DR annotation approaches using the same set

of labels, but translated into different “natural language” for crowdsourcing.

2.2 DR Annotation

Various frameworks exist that can be used to annotate discourse relations, such as RST (Mann and Thompson, 1988) and SDRT (Asher, 1993). In this work, we focus on the annotation of implicit discourse relations, following the framework used to annotate the Penn Discourse Treebank 3.0 (PDTB, Webber et al., 2019). PDTB’s sense classification is structured as a three-level hierarchy, with four coarse-grained sense groups in the first level and more fine-grained senses for each of the next levels.¹ The process is a combination of manual and automated annotation: An automated process identifies potential explicit connectives, and annotators then decide on whether the potential connective is indeed a true connective. If so, they specify one or more senses that hold between its arguments. If no connective or alternative lexicalization is present (i.e., for implicit relations), each annotator provides one or more connectives that together express the sense(s) they infer.

DR datasets, such as PDTB (Webber et al., 2019), RST-DT (Carlson and Marcu, 2001), and TED-MDB (Zeyrek et al., 2019), are commonly annotated by trained annotators, who are expected to be familiar with extensive guidelines written for a given task (Plank et al., 2014; Artstein and Poesio, 2008; Riezler, 2014). However, there have also been efforts to crowdsource discourse relation annotations (Kawahara et al., 2014; Kishimoto et al., 2018; Scholman and Demberg, 2017; Pyatkin et al., 2020). We investigate two crowdsourcing approaches that annotate inter-sentential implicit DRs and we deterministically map the NL-annotations to the PDTB3 label framework.

2.2.1 Crowdsourcing DRs with the DC Method

Yung et al. (2019) developed a crowdsourcing discourse relation annotation method using discourse connectives, referred to as the DC method. For every instance, participants first provide a connective that in their view, best expresses the relation between the two arguments. Note that the

¹We merge the belief and speech-act relation senses (which cannot be distinguished reliably by QA and DC) with their corresponding more general relation senses.

connective chosen by the participant might be ambiguous. Therefore, participants disambiguate the relation in a second step, by selecting a connective from a list that is generated dynamically based on the connective provided in the first step. When the first step insertion does not match any entry in the connective bank (from which the list of disambiguating connectives is generated), participants are presented with a default list of twelve connectives expressing a variety of relations. Based on the connectives chosen in the two steps, the inferred relation sense can be extracted. For example, the CONJUNCTION reading in Figure 1 can be expressed by *in addition*, and the RESULT reading can be expressed by *consequently*.

The DC method was used to create a crowd-sourced corpus of 6,505 discourse-annotated implicit relations, named DiscoGeM (Scholman et al., 2022a). A subset of DiscoGeM is used in the current study (see Section 3).

2.2.2 Crowdsourcing DRs by the QA Method

Pyatkin et al. (2020) proposed to crowdsource discourse relations using QA pairs. They collected a dataset of intra-sentential QA annotations which aim to represent discourse relations by including one of the propositions in the question and the other in the respective answer, with the question prefix (*What is similar to..?*, *What is an example of..?*) mapping to a relation sense. Their method was later extended to also work inter-sententially (Scholman et al., 2022). In this work we make use of the extended approach that relates two distinct sentences through a question and answer. The following QA pair, for example, connects the two sentences in Figure 1 with a RESULT relation.

- (1) **What is the result of** Caesar being assassinated by a group of rebellious senators?(S1) - *A new series of civil wars broke out [...]*(S2)

The annotation process consists of the following steps: From two consecutive sentences, annotators are asked to choose a sentence that will be used to formulate a question. The other sentence functions as an answer to that question. Next they start building a question by choosing a question prefix and by completing the question with content from the chosen sentence.

Since it is possible to choose either of the two sentences as question/answer for a specific set of symmetric relations, (i.e., **What is the reason a**

new series of civil wars broke out?), we consider both possible formulations as equivalent.

The set of possible question prefixes cover all PDTB 3.0 senses (excluding belief and speech-act relations). The direction of the relation sense, e.g., *arg1-as-denier* vs. *arg2-as-denier*, is determined by which of the two sentences is chosen for the question/answer. While Pyatkin et al. (2020) allowed crowdworkers to form multiple QA pairs per instance, i.e., annotate more than one discourse sense per relation, we decided to limit the task to 1 sense per relation per worker. We took this decision in order for the QA method to be more comparable to the DC method, which also only allows the insertion of a single connective.

3 Method

3.1 Data

We annotated 1,200 inter-sentential discourse relations using both the DC and the QA task design.² Of these 1,200 relations, 900 were taken from the DiscoGeM corpus and 300 from the PDTB 3.0.

DiscoGeM Relations The 900 DiscoGeM instances that were included in the current study represent different domains: 296 instances were taken from the subset of DiscoGeM relations that were taken from Europarl proceedings (written proceedings of prepared political speech taken from the Europarl corpus; Koehn, 2005), 304 instances were taken from the literature subset (narrative text from five English books),³ and 300 instances from the Wikipedia subset of DiscoGeM (informative text, taken from the summaries of 30 Wikipedia articles). These different genres enable a cross-genre comparison. This is necessary, given that prevalence of certain relation types can differ across genres (Rehbein et al., 2016; Scholman et al., 2022a; Webber, 2009).

These 900 relations were already labeled using the DC method in DiscoGeM; we additionally collect labels using the QA method for the current study. In addition to crowd-sourced labels using the DC and QA methods, the Wikipedia subset was also annotated by three trained annotators.⁴

²The annotations are available at <https://github.com/merelscholman/DiscoGeM>.

³*Animal Farm* by George Orwell, *Harry Potter and the Philosopher’s Stone* by J. K. Rowling, *The Hitchhikers Guide to the Galaxy* by Douglas Adams, *The Great Gatsby* by F. Scott Fitzgerald, and *The Hobbit* by J. R. R. Tolkien.

⁴Instances were labeled by two annotators and verified by a third; Cohen’s κ agreement between the first annotator and

Forty-seven percent of these Wikipedia instances were labeled with multiple senses by the expert annotators (i.e., were considered to be ambiguous or express multiple readings).

PDTB Relations The PDTB relations were included for the purpose of comparing our annotations with traditional PDTB gold standard annotations. These instances (all inter-sentential) were selected to represent all relational classes, randomly sampling at most 15 and at least 2 (for classes with less than 15 relation instances we sampled all existing relations) relation instances per class. The reference labels for the PDTB instances consist of the original PDTB labels annotated as part of the PDTB3 corpus. Only 8% of these consisted of multiple senses.

3.2 Crowdworkers

Crowdworkers were recruited via Prolific using a selection approach (Scholman et al., 2022) that has been shown to result in a good trade off between quality and time/monetary efforts for DR annotation. Crowdworkers had to meet the following requirements: be native English speakers, reside in UK, Ireland, USA, or Canada, and have obtained at least an undergraduate degree.

Workers who fulfilled these conditions could participate in an initial recruitment task, for which they were asked to annotate a text with either the DC or QA method and were shown immediate feedback on their performance. Workers with an accuracy ≥ 0.5 on this task were qualified to participate in further tasks. We hence created a unique set of crowdworkers for each method. The DC annotations (collected as part of DiscoGeM) were provided by a final set of 199 selected crowdworkers; QA had a final set of 43 selected crowdworkers.⁵ Quality was monitored throughout the production data collection and qualifications were adjusted according to performance.

Every instance was annotated by 10 workers per method. This number was chosen based on parity with previous research. For example, Snow et al. (2008) show that a sample of 10 crowdsourced annotations per instance yields satisfactory accuracy

the reference label was .82 (88% agreement), and between the second and the reference label was .96 (97% agreement). See Scholman et al. (2022a) for additional details.

⁵The larger set of selected workers in the DC method is because more data was annotated by DC workers as part of the creation of DiscoGeM.

for various linguistic annotation tasks. Scholman and Demberg (2017) found that assigning a new group of 10 annotators to annotate the same instances resulted in a near-perfect replication of the connective insertions in an earlier DC study.

Instances were annotated in batches of 20. For QA, one batch took about 20 minutes to complete, and for DC, 7 minutes. Workers were reimbursed about \$2.50 and \$1.88 per batch, respectively.

3.3 Inter-annotator Agreement

We evaluate the two DR annotation methods by the inter-annotator agreement (IAA) between the annotations collected by both methods and IAA with reference annotations collected from trained annotators.

Cohen’s kappa (Cohen, 1960) is a metric frequently used to measure IAA. For DR annotations, a Cohen’s kappa of .7 is considered to reflect good IAA (Spooren and Degand, 2010). However, prior research has shown that agreement on implicit relations is more difficult to reach than on explicit relations: Kishimoto et al. (2018) report an F1 of .51 on crowdsourced annotations of implicits using a tagset with 7 level-2 labels; Zikánová et al. (2019) report $\kappa = .47$ (58%) on expert annotations of implicits using a tagset with 23 level-2 labels; and Demberg et al. (2019) find that PDTB and RST-DT annotators agree on the relation sense on 37% of implicit relations. Cohen’s kappa is primarily used for comparison between single labels and the IAAs reported in these studies are also based on single aggregated labels.

However, we also want to compare the obtained 10 annotations per instance with our reference labels that also contain multiple labels. The comparison becomes less straightforward when there are multiple labels because the chance of agreement is inflated and partial agreement should be treated differently. We thus measure the IAA between multiple labels in terms of both full and partial agreement rates, as well as the multi-label kappa metric proposed by Marchal et al. (2022). This metric adjusts the multi-label agreements with bootstrapped expected agreement. We consider all the labels annotated by the crowdworkers in each instance, excluding minority labels with only one vote.⁶

⁶We assumed there were 10 votes per item and removed labels with less than 20% of votes, even though in rare cases there could be 9 or 11 votes. On average, the removed labels represent 24.8% of the votes per item.

	Europarl	Novel	Wiki.	PDTB	all
Item counts	296	304	300	302	1202
QA sub-labels/item	2.13	2.21	2.26	2.45	2.21
DC sub-labels/item	2.37	2.00	2.09	2.21	2.17
full/+partial agreement	.051/.841	.092/.865	.060/.920	.050/.884	.063/.878
multi-label kappa	.813	.842	.903	.868	.857
JSD	.505	.492	.482	.510	.497

Table 1: Comparison between the labels obtained by DC vs. QA. Full (or +partial) agreement means all (or at least one sub-label) match(es). Multi-label kappa is adapted from Marchal et al. (2022). JSD is calculated based on the actual distributions of the crowdsourced sub-labels, excluding labels with only one vote (smaller values are better).

In addition, we compare the distributions of the crowdsourced labels using the Jensen-Shannon divergence (JSD) following existing reports (Erk and McCarthy, 2009; Nie et al., 2020; Zhang et al., 2021). Similarly, minority labels with only one vote are excluded. Since distributions are not available in the reference labels, when comparing with the reference labels, we evaluate by the JSD based on the flattened distributions of the labels, which means we replace the original distribution of the votes with an even distribution of the labels that have been voted by more than one annotator. We call this version JSD_flat.

As a third perspective on IAA we report agreement among annotators on an item annotated with QA/DC. Following previous work (Nie et al., 2020), we use entropy of the soft labels to quantify the uncertainty of the crowd annotation. Here labels with only one vote are also included as they contribute to the annotation uncertainty. When calculating the entropy, we use a logarithmic base of $n = 29$, where n is the number of possible labels. A lower entropy value suggests that the annotators agree with each other more and the annotated label is more certain. As discussed in Section 1, the source of disagreement in annotations could come from the items, the annotators, and the methodology. High entropy across multiple annotations of a specific item within the same annotation task suggests that the item is ambiguous.

4 Results

We first compare the IAA between the two crowdsourced annotations, then we discuss IAA between DC/QA and the reference annotations, and lastly we perform an analysis based on annotation uncertainty. Here, “sub-labels” of an instance means all

relations that have received more than one annotation; and “label distribution” is the distribution of the votes of the sub-labels.

4.1 IAA Between the Methods

Table 1 shows that both methods yield more than two sub-labels per instance after excluding minority labels with only one vote. This supports the idea that **multi-sense annotations better capture the fact that often more than one sense can hold implicitly between two discourse arguments**.

Table 1 also presents the IAA between the labels crowdsourced with QA and DC per domain. The agreement between the two methods is good: The labels assigned by the two methods (or at least one of the sub-labels in case of a multi-label annotation) match for about 88% of the items. This speaks for the fact that **both methods are valid**, as similar sets of labels are produced.

The full agreement scores, however, are very low. This is expected, as the chance to match on all sub-labels is also very low compared to a single-label setting. The multi-label kappa (which takes chance agreement of multiple labels into), and JSD (which compares the distributions of the multiple labels) are hence more suitable. We note that the PDTB gold annotation that we use for evaluation does not assign multiple relations systematically and has a low rate of double labels. This explains why the PDTB subsets have a high partial agreement while the JSD ends up being worst.

4.2 IAA Between Crowdsourced and Reference Labels

Table 2 compares the labels crowdsourced by each method and the reference labels, which are available for the Wikipedia and PDTB subsets. It

	Wiki.	PDTB
Item counts	300	302
Ref. sub-labels/item	1.54	1.08
QA: sub-labels/item	2.26	2.45
full/+partial agreement	.133/.887	.070/.487
multi-label kappa	.857	.449
JSD _{flat}	.468	.643
DC: sub-labels/item	2.09	2.21
full/+partial agreement	.110/.853	.103/.569
multi-label kappa	.817	.524
JSD _{flat}	.483	.606

Table 2: Comparison against gold labels for the QA or DC methods. Since the distribution of the reference sub-labels is not available, JSD_{flat} is calculated between uniform distributions of the sub-labels.

	Europarl	Wikipedia	Novel	PDTB
QA	0.40	0.38	0.38	0.41
DC	0.37	0.34	0.35	0.36

Table 3: Average entropy of the label distributions (10 annotations per relation) for QA/DC, split by domain.

can be observed that both methods achieve higher full agreements with the reference labels than with each other on both domains. This indicates that the **two methods are complementary**, with each method better capturing different sense types. In particular, the QA method tends to show higher agreement with the reference for Wikipedia items, while the DC annotations show higher agreement with the reference for PDTB items. This can possibly be **attributed to the development of the methodologies**: The DC method was originally developed by testing on data from the PDTB in Yung et al. (2019), whereas the QA method was developed by testing on data from Wikipedia and Wikinews in Pyatkin et al. (2020).

4.3 Annotation Uncertainty

Table 3 compares the average entropy of the soft labels collected by both methods. It can be observed that the uncertainty among the labels chosen by the crowdworkers is similar across domains but always slightly lower for DC. We further look at the correlation between annotation uncertainty and cross-method agreement, and find




Figure 2: Correlation between the entropy of the annotations and the JSD_{flat} between the crowdsourced labels and reference.

that agreement between methods is substantially higher for those instances where within-method entropy was low. Similarly, we find that agreement between crowdsourced annotations and gold labels is highest for those relations, where little entropy was found in crowdsourcing.

Next, we want to check if the item effect is similar across different methods and domains. Figure 2 shows the correlation between the annotation entropy and the agreement with the reference of each item, of each method for the Wikipedia / PDTB subsets. It illustrates that annotations of both methods diverge with the reference more as the uncertainty of the annotation increases. While the effect of uncertainty is similar across methods on the Wikipedia subset, the quality of the QA annotations depends more on the uncertainty

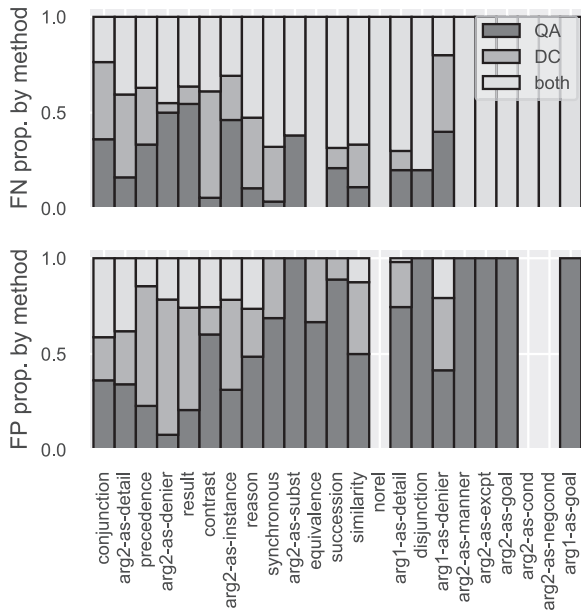


Figure 3: Distribution of the annotation errors by method. Labels annotated by at least 2 workers are compared against the reference labels of the Wikipedia and PDTB items. The relation types are arranged in descending order of the ‘‘ref. sub-label counts.’’

compared to the DC annotations on the PDTB subset. This means that method bias also exists on the level of annotation uncertainty and should be taken into account when, for example, entropy is used as a criterion to select reliable annotations.

5 Sources of the Method Bias

In this section, we analyze method bias in terms of the sense labels collected by each method. We also examine the potential limitations of the methods which could have contributed to the bias and demonstrate how we can utilize information on method bias to crowdsource more reliable labels. Lastly, we provide a cross-domain analysis.

Table 5 presents the confusion matrix of the labels collected by both methods for the most frequent level-2 relations. Figure 3 and Table 4 show the distribution of the true and false positives of the sub-labels. These results show that both methods are biased towards certain DRs. The source of these biases can be categorized into two types, which we will detail in the following subsections.

5.1 Limitation of Natural Language for Annotation

There are limitations of representing DRs in natural languages using both QA and DC. For

label	FN _{QA}	FN _{DC}	FP _{QA}	FP _{DC}
conjunction	43	46	203	167
arg2-as-detail	42	62	167	152
precedence	19	18	18	37
arg2-as-denier	38	20	15	47
result	10	5	110	187
contrast	8	17	84	39
arg2-as-instance	10	7	44	57
reason	12	17	54	37
synchronous	20	27	11	5
arg2-as-subst	21	13	1	0
equivalence	22	22	2	1
succession	17	15	24	3
similarity	7	8	15	12
norel	12	12	0	0
arg1-as-detail	9	8	39	13
disjunction	5	4	10	0
arg1-as-denier	3	3	33	31
arg2-as-manner	2	2	9	0
arg2-as-ecpct	2	2	1	0
arg2-as-goal	1	1	5	0
arg2-as-cond	1	1	0	0
arg2-as-negcond	1	1	0	0
arg1-as-goal	1	1	3	0

Table 4: FN and FP counts of each method grouped by the reference sub-labels.

example, the QA method confuses workers when the question phrase contains a connective:⁷

- (2) *‘‘Little tyke,’’ chortled Mr. Dursley as he left the house. **He got into his car and backed out of number four’s drive.** [QA:SUCCESSION, PRECEDENCE, DC:CONJUNCTION, PRECEDENCE]*

In the above example, the majority of the workers formed the question *‘‘**After what he left the house?**, which was likely a confusion with **‘‘What did he do after he left the house?’’**’*. This could explain the frequent confusion between PRECEDENCE and SUCCESSION by QA, resulting in the frequent FPs of SUCCESSION (Figure 3).⁸

For DC, rare relations which lack a frequently used connective are harder to annotate, for example:

- (3) *He had made an arrangement with one of the cockerels to call him in the mornings half an*

⁷The examples are presented in the following format: *italics* = argument 1; **bolded** = argument 2; plain = contexts.

⁸Similarly, the question *‘‘Despite what . . . ?’’* is easily confused with *‘‘despite...’’*, which could explain the frequent FP of *arg1-as-denier* by the QA method.

hour earlier than anyone else, and would put in some volunteer labour at whatever seemed to be most needed, before the regular day's work began. His answer to every problem, every setback, was "I will work harder!" - which he had adopted as his personal motto. [QA:ARG1-AS-INSTANCE; DC:RESULT]

It is difficult to use the DC method to annotate the ARG1-AS-INSTANCE relation due to a lack of typical, specific, and context-independent connective phrases that mark these rare relations, such as “*this is an example of ...*”. By contrast, the QA method allows workers to make a question and answer pair in the reverse direction, with S1 being the answer to S2, using the same question words, e.g., *What is an example of the fact that his answer to every problem [...] was "I will work harder!"?*. This allows workers to label rarer relation types that were not even uncovered by trained annotators.

Many common DCs are ambiguous, such as *but* and *and*, and can be hard to disambiguate. To address this, the DC method provides workers with unambiguous connectives in the second step. However, these unambiguous connectives are often relatively uncommon and come with different syntactic constraints, depending on whether they are coordinating or subordinating conjunctions or discourse adverbials. Hence, they do not fit in all contexts. Additionally, some of the unambiguous connectives sound very “heavy” and would not be used naturally in a given sentence. For example, *however* is often inserted in the first step, but it can mark multiple relations and is disambiguated in the second step by the choice among *on the contrary* for CONTRAST, *despite* for ARG1-AS-DENIER, and *despite this* for ARG2-AS-DENIER. *Despite this* was chosen frequently since it can be applied to most contexts. This explains the DC method's bias towards *arg2-as-denier* against *contrast* (Figure 3: most FPs of *arg2-as-denier* and most FNs of *contrast* come from DC).

While the QA method also requires workers to select from a set of question starts, which also contain infrequent expressions (such as *Unless what..?*), workers are allowed to edit the text to improve the wordings of the questions. This helps reduce the effect of bias towards more frequent question prefixes and makes crowdworkers doing the QA task more likely to choose infrequent relation senses than those doing the DC task.

5.2 Guideline Underspecification

Jiang and de Marneffe (2022) report that some disagreements in NLI tasks come from the loose definition of certain aspects of the task. We found that both QA and DC also **do not give clear enough instructions in terms of argument spans**. The DRs are annotated at the boundary of two consecutive sentences but both methods do not limit workers to annotate DRs that span exactly the two sentences.

More specifically, the QA method allows the crowdworkers to form questions by copying spans from one of the sentences. While this makes sure that the relation lies locally between two consecutive sentences, it also sometimes happens that workers highlight partial spans and annotate relations that span over parts of the sentences. For example:

- (4) *I agree with Mr Pirker, and it is probably the only thing I will agree with him on if we do vote on the Ludford report. It is going to be an interesting vote.* [QA:ARG2-AS-DETAIL,REASON; DC:CONJUNCTION,RESULT]

In Ex. (4), workers constructed the question “**What provides more details on the vote on the Ludford report?**”. This is similar to the instructions in PDTB 2.0 and 3.0's annotation manuals, specifying that annotators should take minimal spans which don't have to span the entire sentence. Other relations should be inferred when the argument span is expanded to the whole sentence, for example a RESULT relation reflecting that there is little agreement, which will make the vote interesting.

Often, a sentence can be interpreted as the elaboration of certain entities in the previous sentence. This could explain why ARG1/2-AS-DETAIL tends to be overlabelled by QA. Figure 3 shows that the QA has more than twice as many FP counts for ARG2-AS-DETAIL compared to DC – the contrast is even bigger for ARG1-AS-DETAIL. Yet it is not trivial to filter out such questions that only refer to a part of the sentence, because in some cases, the highlighted entity does represent the whole argument span.⁹ Clearer instructions in the guidelines are desirable.

⁹Such as “*a few final comments*” in this example: *Ladies and gentlemen, I would like to make a few final comments. This is not about the implementation of the habitats directive.*

DC	QA	Synchr.	Asynchr.	Cause	Concession	Contrast	Conjunction	Detail	Instant.	all DC
Synchr.		1	1	1			2	1		6
Asynchr.		2	45	17		7	18	7	2	98
Cause			13	184	5	8	78	36	13	337
Concess.		2	3	13	14	22	18	22	1	95
Contrast		1		7	1	21	3	4		37
Conjunct.		4	12	37	1	12	184	67	6	323
Detail		2	5	40	2	4	52	93	16	214
Instant.				2	1	2	9	10	23	47
all QA		12	79	301	24	76	364	240	61	

Table 5: Confusion matrix for the most frequent level-2 sublabels which were annotated by at least 2 workers per relation; values are represented as colors.

Similarly, DC does not limit workers to annotate relations between the two sentences, consider:

(5)

*When two differently-doped regions exist in the same crystal, a semiconductor junction is created. **The behavior of charge carriers, which include electrons, ions and electron holes, at these junctions is the basis of diodes, transistors and all modern electronics.*** [Ref:ARG2-AS-DETAIL; QA:ARG2-AS-DETAIL, CONJUNCTION; DC:CONJUNCTION, RESULT]

In this example, many people inserted *as a result*, which naturally marks the intra-sentence relation (*...is created as a result.*) Many relations are potentially spuriously labelled as RESULT, which are frequent between larger chunks of texts. Table 5 shows that the most frequent confusion is between DC’s CAUSE and QA’s CONJUNCTION.¹⁰ Within the level-2 CAUSE relation sense, it is the level-3 RESULT relation that turns out to be the main contributor to the observed bias. Figure 3 also shows that most FPs of RESULT come from the DC method.

5.3 Aggregating DR Annotations Based on Method Bias

The qualitative analysis above provides insights on certain method biases observed in the label distributions, such as QA’s bias towards ARG1/2-AS

¹⁰A chi-squared test confirms that the observed distribution is significantly different from what could be expected based on chance disagreement.

DETAIL and SUCCESSION and DC’s bias towards CONCESSION and RESULT. Being aware of these biases would allow to combine the methods: After first labelling all instances with the more cost-effective DC method, RESULT relations, which we know tend to be overlabelled by the DC method, could be re-annotated using the QA method. We simulate this for our data and find that this would increase the partial agreement from 0.853 to 0.913 for Wikipedia and from 0.569 to 0.596 for PDTB.

6 Analysis by Genre

For each of the four genres (Novel, Wikipedia, Europarl, and WSJ) we have ~300 implicit DRs annotated by both DC and QA. Scholman et al. (2022a) showed, based on the DC method, that in DiscoGeM, CONJUNCTION is prevalent in the Wikipedia domain, PRECEDENCE in Literature and RESULT in Europarl. The QA annotations replicate this finding, as displayed in Figure 4.

It appears more difficult to obtain agreement with the majority labels in Europarl than in other genres, which is reflected in the average entropy (see Table 3) of the distributions for each genre, where DC has the highest entropy in the Europarl domain and QA the second highest (after PDTB). Table 1 confirms these findings, showing that the agreement between the two methods is highest for Wikipedia and lowest for Europarl.

In the latter domain, the DC method results in more CAUSAL relations: 36% of the CONJUNCTIONS labelled by QA are labelled as RESULT in DC.¹¹ Manual inspection of these DC annotations reveals that workers chose *considering this* frequently only in the Europarl subset. This connective phrase is typically used to mark a pragmatic result relation, where the result reading comes from the belief of the speaker (Ex. (4)). This type of relation is expected to be more frequent in speech and argumentative contexts and is labelled as RESULT-BELIEF in PDTB3. QA does not have a question prefix available that could capture RESULT-BELIEF senses. The RESULT labels obtained by DC are therefore a better fit with the PDTB3 framework than QA’s CONJUNCTIONS. CONCESSION is generally more prevalent with the DC method, especially in Europarl, with 9% compared to 3% for QA. CONTRAST, on the other hand, seems to be favored by the QA method, of which most (6%)

¹¹This appeared to be distributed over many annotators and is thus a true method bias.

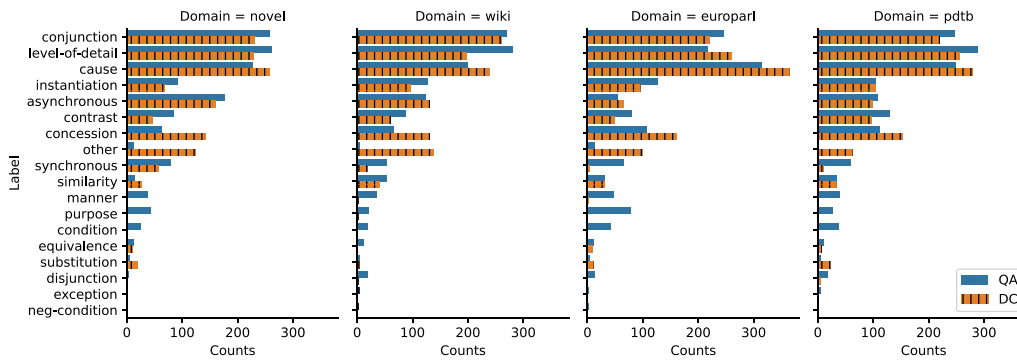


Figure 4: Level-2 sublabel counts of all the annotated labels of both methods, split by domain.

CONTRAST relations are found in Wikipedia, compared to 3% for DC. Figure 4 also highlights that for the QA approach, annotators tend to choose a wider variety of senses which are rarely ever annotated by DC, such as PURPOSE, CONDITION, and MANNER.

We conclude that encyclopedic and literary texts are the most suitable to be annotated using either DC or QA, as they show higher inter-method agreement (and for Wikipedia also higher agreement with gold). Spoken-language and argumentative domains on the other hand are trickier to annotate as they contain more pragmatic readings of the relations.

7 Case Studies: Effect of Task Design on DR Classification Models

Analysis of the crowdsourced annotations reveals that the two methods have different biases and different correlations with domains and the style (and possibly function) of the language used in the domains. We now investigate the effect of task design bias on automatic prediction of implicit discourse relations. Specifically, we carry out two case studies to demonstrate the effect that task design and the resulting label distributions have on discourse parsing models.

Task and Setup We formulate the task of predicting implicit discourse relations as follows. The input to the model are two sequences S_1 and S_2 , which represent the arguments of a discourse relation. The targets are PDTB 3.0 sense types (including level-3). This model architecture is similar to the model for implicit DR prediction by Shi and Demberg (2019). We experiment with two different losses and targets: a cross-entropy loss where the target is a single majority label

and a soft cross-entropy loss where the target is a probability distribution over the annotated labels. Using the 10 annotations per instance, we obtain label distributions for each relation, which we use as soft targets. Training with a soft loss has been shown to improve generalization in vision and NLP tasks (Peterson et al., 2019; Uma et al., 2020). As suggested in Uma et al. (2020), we normalize the sense-distribution over the 30 possible labels¹² with a softmax.

Assume one has a relation with the following annotations: 4 RESULT, 3 CONJUNCTION, 2 SUCCESSION, 1 ARG1-AS-DETAIL. For the hard loss, the target would be the majority label: RESULT. For the soft loss we normalize the counts (every label with no annotation has a count of 0) using a softmax, for a smoother distribution without zeros.

We fine-tune DeBERTa (deberta-base) (He et al., 2020) in a sequence classification setup using the HuggingFace checkpoint (Wolf et al., 2020). The model trains for 30 epochs with early stopping and a batch size of 8.

Data In addition to the 1,200 instances we analyzed in the current contribution, we additionally use all annotations from DiscoGeM as training data. DiscoGeM, which was annotated with the DC method, adds 2756 *Novel* relations, 2504 *Europarl* relations, and 345 *Wikipedia* relations. We formulate different setups for the case studies.

¹²precedence, arg2-as-detail, conjunction, result, arg1-as-detail, arg2-as-denier, contrast, arg1-as-denier, synchronous, reason, arg2-as-instance, arg2-as-cond, arg2-as-subst, similarity, disjunction, succession, arg1-as-goal, arg1-as-instance, arg2-as-goal, arg2-as-manner, arg1-as-manner, equivalence, arg2-as-excpt, arg1-as-excpt, arg1-as-cond, differentcon, norel, arg1-as-negcond, arg2-as-negcond, arg1-as-subst.

	PDTB _{test}	Wiki _{gold}	TED-M.
DC	0.34†	0.65†	0.36
DC Soft	0.29*†	0.70* †	0.34*
QA+DC _∩	0.34*	0.67	0.37
QA+DC _∩ Soft	0.38* [*]	0.66*	0.31*
QA+DC _∪	0.35♠	0.49♠	0.36♠
QA+DC _∪ Soft	0.41* ♠	0.67*♠	0.43* ♠

Table 6: Accuracy of model (with soft vs. hard loss) prediction on gold labels. The model is trained either on DC data (DC), an intersection of DC and QA (\cap), or the union of DC and QA (\cup). Same symbol in a column indicates a statistically significant (McNemar test) difference in cross-model results.

7.1 Case 1: Incorporating Data from Different Task Designs

The purpose of this study is to see if a model trained on data crowdsourced by DC/QA methods can generalize to traditionally annotated test sets. We thus test on the 300 Wikipedia relations annotated by experts (*Wiki gold*), all implicit relations from the test set of PDTB 3.0 (*PDTB test*), and the implicit relations of the English test set of TED-MDB (Zeyrek et al., 2020). For training data, we either use (1) all of the DiscoGeM annotations (**Only DC**); or (2) 1200 QA annotations from all four domains, plus 5,605 DC annotations from the rest of DiscoGeM (**Intersection**, \cap); or (3) 1200 annotations which combine the label counts (e.g., 20 counts instead of 10) of QA and DC, plus 5,605 DC annotations from the rest of DiscoGeM (**Union**, \cup). We hypothesize that this union will lead to improved results due to the annotation distribution coming from a bigger sample. When testing on *Wiki gold*, the corresponding subset of *Wikipedia* relations are removed from the training data. We randomly sampled 30 relations for dev.

Results Table 6 shows how the model generalizes to traditionally annotated data. On the PDTB and the Wikipedia test set, the model with a soft loss generally performs better than the hard loss model. TED-MDB, on the other hand, only contains a single label per relation and training with a distributional loss is therefore less beneficial. Mixing DC and QA data only improves in the soft case for PDTB. The merging of the respective method label counts,

on the other hand, leads to the best model performance on both PDTB and TED-MDB. On Wikipedia the best performance is obtained when training on soft DC-only distributions. Looking at the label-specific differences in performance, we observe that improvement on the Wikipedia test set mainly comes from better precision and recall when predicting ARG2-AS-DETAIL, while on PDTB QA+DC_∩ Soft is better at predicting CONJUNCTION.

We conclude that training on data that comes from different task designs does not hurt performance, and even slightly improves performance when using majority vote labels. When training with a distribution, the union setup (\cup) seems to work best.

7.2 Case 2: Cross-domain vs Cross-method

The purpose of this study is to investigate how cross-domain generalization is affected by method bias. In other words, we want to compare a *cross-domain and cross-method* setup with a *cross-domain and same-method* setup. We test on the domain-specific data from the 1,200 instances annotated by QA and DC, respectively, and train on various domain configurations from DiscoGem (excluding dev and test), together with the extra 300 PDTB instances, annotated by DC.

Table 7 shows the different combinations of data sets we use in this study (columns) as well as the results of in- and cross-domain and in- and cross-method predictions (rows). Both a change in domain and a change in annotation task lead to lower performance. Interestingly, the results show that the task factor has a stronger effect on performance than the domain: When training on DC distributions, the QA test results are worse than the DC test results in all cases. This indicates that task bias is an important factor to consider when training models. Generally, except in the out-of-domain novel test case, training with a soft loss leads to the same or considerably better generalization accuracy than training with a hard loss. We thus confirm the findings of Peterson et al. (2019) and Uma et al. (2020) also for DR classification.

8 Discussion and Conclusion

DR annotation is a notoriously difficult task with low IAA. Annotations are not only subject to the interpretation of the coder (Spooren and Degand,

		All except Wiki DC			All except EP DC			All except Novel DC			All		
		hard	soft	soft	hard	soft	soft	hard	soft	soft	hard	soft	soft
		acc	acc	JSD	acc	acc	JSD	acc	acc	JSD	acc	acc	JSD
Wiki	QA	0.51*	0.53*	0.38	0.51	0.50	0.37	0.47	0.53*	0.38	0.50*	0.51	0.38
	DC	0.55*	0.59*	0.36	0.60	0.62	0.33	0.56	0.61*	0.35	0.58*	0.61	0.35
EP	QA	0.30	0.35	0.43	0.29	0.33	0.42	0.29	0.38	0.44	0.29	0.34	0.44
	DC	0.50	0.53	0.34	0.41	0.47	0.37	0.49	0.49	0.36	0.47	0.52	0.36
Novel	QA	0.44*	0.47	0.40	0.41	0.45	0.41	0.40	0.39	0.44	0.39	0.45	0.41
	DC	0.52*	0.58	0.33	0.53	0.56	0.34	0.56	0.51	0.39	0.52	0.57	0.36

Table 7: Cross-domain and cross-method experiments, using a hard-loss vs. a soft-loss. Columns show train setting and rows test performance. Acc. is for predicting the majority label. JDS compares predicted distribution (soft) with target distribution. * indicates cross-method results are not statistically significant (McNemar’s test).

2010), but also to the framework (Demberg et al., 2019). The current study extends these findings by showing that the task design also crucially affects the output. We investigated the effect of two distinct crowdsourced DR annotation tasks on the obtained relation distributions. These two tasks are unique in that they use natural language to annotate. Even though these designs are more intuitive to lay individuals, we show that also such natural language-based annotation designs suffer from bias and leave room for varying interpretations (as do traditional annotation tasks).

The results show that both methods have unique biases, but also that both methods are valid, as similar sets of labels are produced. Further, the methods seem to be complementary: Both methods show higher agreement with the reference label than with each other. This indicates that the methods capture different sense types. The results further show that the textual domain can push each method towards different label distributions. Lastly, we simulated how aggregating annotations based on method bias improves agreement.

We suggest several modifications to both methods for future work. For QA, we recommend to replace question prefix options which start with a connective, such as ‘‘After what’’. The revised options should ideally start with a Wh-question word, for example, ‘‘What happens after..’’. This would make the questions sound more natural and help to prevent confusion with respect to level-3 sense distinctions. For DC, an improved interface that allows workers to highlight argument spans could serve as a screen that confirms the relation is between the two consecutive sentences. Syntactic constraints making it difficult to insert certain rare connectives could also be mitigated if

the workers are allowed to make minor edits to the texts.

Considering that both methods show benefits and possible downsides, it could be interesting to combine them for future crowdsourcing efforts. Given that obtaining DC annotations is cheaper and quicker, it could make sense to collect DC annotations on a larger scale and then use the QA method for a specific subset that shows high label entropy. Another option would be to merge both methods, by first letting the crowdworkers insert a connective and then use QAs for the second connective-disambiguation step. Lastly, since we showed that often more than one relation sense can hold, it would make sense to allow annotators to write multiple QA pairs or insert multiple possible connectives for a given relation.

The DR classification experiments revealed that generalization across data from different task designs is hard, in the DC and QA case even harder than cross-domain generalization. Additionally, we found that merging data distributions coming from different task designs can help boost performance on data coming from a third source (traditional annotations). Lastly, we confirmed that soft modeling approaches using label distributions can improve discourse classification performance.

Task design bias has been identified as one source of annotation bias and acknowledged as an artifact of the dataset in other linguistic tasks as well (Pavlick and Kwiatkowski, 2019; Jiang and de Marneffe, 2022). Our findings show that the effect of this type of bias can be reduced by training with data collected by multiple methods. This could be the same for other NLP tasks, especially those cast in natural language, and comparing their

task designs could be an interesting future research direction. We therefore encourage researchers to be more conscious about the biases crowdsourcing task design introduces.

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft, Funder ID: <http://dx.doi.org/10.13039/501100001659>, grant number: SFB1102: Information Density and Linguistic Encoding, by the the European Research Council, ERC-StG grant no. 677352, and the Israel Science Foundation grant 2827/21, for which we are grateful. We also thank the TACL reviewers and action editors for their thoughtful comments.

References

- Rahul Aralikatte, Matthew Lamm, Daniel Hardt, and Anders Søgaard. 2021. Ellipsis resolution as question answering: An evaluation. In *16th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, pages 810–817, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.eacl-main.68>
- Lora Aroyo and Chris Welty. 2013. Crowd truth: Harnessing disagreement in crowdsourcing a relation extraction gold standard. *WebSci2013 ACM*, 2013(2013).
- Ron Artstein 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>
- N. Asher. 1993. *Reference to Abstract Objects in Discourse*, volume 50. Kluwer, Norwell, MA, Dordrecht. <https://doi.org/10.1007/978-94-011-1715-9>
- Valerio Basile, Michael Fell, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, Massimo Poesio, and Alexandra Uma. 2021. We need to consider disagreement in evaluation. In *Proceedings of the 1st Workshop on Benchmarking: Past, Present and Future*, pages 15–21, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.bppf-1.3>
- Samuel Bowman, 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, Lisbon, Portugal. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D15-1075>
- Samuel Bowman and George Dahl. 2021. What will it take to fix benchmarking in natural language understanding? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4843–4855, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.385>
- Sven Buechel and Udo Hahn. 2017a. Emobank: Studying the impact of annotation perspective and representation format on dimensional emotion analysis. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 578–585, Valencia, Spain. Association for Computational Linguistics. <https://doi.org/10.18653/v1/E17-2092>
- Sven Buechel and Udo Hahn. 2017b. Readers vs. writers vs. texts: Coping with different perspectives of text understanding in emotion annotation. In *Proceedings of the 11th Linguistic Annotation Workshop*, pages 1–12, Valencia, Spain. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W17-0801>
- Lynn Carlson and Daniel Marcu. 2001. Discourse tagging reference manual. *ISI Technical Report ISI-TR-545*, 54:1–56.
- Nancy Chang, Russell Lee-Goldman, and Michael Tseng. 2016. Linguistic wisdom from the crowd. In *Third AAAI Conference on Human Computation and Crowdsourcing*. <https://doi.org/10.1609/hcomp.v3i1.13266>
- Tongfei Chen, Zheng Ping Jiang, Adam Poliak, Keisuke Sakaguchi, and Benjamin Van Durme. 2020. Uncertain natural language inference. In *Proceedings of the 58th Annual*

- Meeting of the Association for Computational Linguistics*, pages 8772–8779, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.774>
- John Joon Young Chung, Jean Y. Song, Sindhu Kutty, Sungsoo Hong, Juho Kim, and Walter S. Lasecki. 2019. Efficient elicitation approaches to estimate collective crowd answers. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–25. <https://doi.org/10.1145/3359164>
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46. <https://doi.org/10.1177/001316446002000104>
- Alan Cowen, Disa Sauter, Jessica L. Tracy, and Dacher Keltner. 2019. Mapping the passions: Toward a high-dimensional taxonomy of emotional experience and expression. *Psychological Science in the Public Interest*, 20(1):69–90. <https://doi.org/10.1177/1529100619850176>, PubMed: 31313637
- Marie-Catherine De Marneffe, Christopher D. Manning, and Christopher Potts. 2012. Did it happen? The pragmatic complexity of veridicality assessment. *Computational Linguistics*, 38(2):301–333. https://doi.org/10.1162/COLI_a_00097
- Vera Demberg, Merel C. J. Scholman, and Fatemeh Torabi Asr. 2019. How compatible are our discourse annotation frameworks? Insights from mapping RST-DT and PDTB annotations. *Dialogue & Discourse*, 10(1):87–135. <https://doi.org/10.5087/dad.2019.104>
- Mark Díaz, Isaac Johnson, Amanda Lazar, Anne Marie Piper, and Darren Gergle. 2018. Addressing age-related bias in sentiment analysis. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–14. <https://doi.org/10.1145/3173574.3173986>
- Anca Dumitrache. 2015. Crowdsourcing disagreement for collecting semantic annotation. In *European Semantic Web Conference*, pages 701–710. Springer. https://doi.org/10.1007/978-3-319-18818-8_43
- Anca Dumitrache, Oana Inel, Lora Aroyo, Benjamin Timmermans, and Chris Welty. 2018. CrowdTruth 2.0: Quality metrics for crowdsourcing with disagreement. In *1st Workshop on Subjectivity, Ambiguity and Disagreement in Crowdsourcing, and Short Paper 1st Workshop on Disentangling the Relation Between Crowdsourcing and Bias Management, SAD+ CrowdBias 2018*, pages 11–18. CEUR-WS.
- Anca Dumitrache, Oana Inel, Benjamin Timmermans, Carlos Ortiz, Robert-Jan Sips, Lora Aroyo, and Chris Welty. 2021. Empirical methodology for crowdsourcing ground truth. *Semantic Web*, 12(3):403–421. <https://doi.org/10.3233/SW-200415>
- Yanai Elazar, Victoria Basmov, Yoav Goldberg, and Reut Tsarfaty. 2022. Text-based np enrichment. *Transactions of the Association for Computational Linguistics*, 10:764–784. https://doi.org/10.1162/tacl_a_00488
- Katrin Erk and Diana McCarthy. 2009. Graded word sense assignment. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 440–449, Singapore. Association for Computational Linguistics.
- Elisa Ferracane, 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, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.129>
- Nicholas Fitzgerald, Julian Michael, Luheng He, and Luke Zettlemoyer. 2018. Large-scale qa-srl parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2051–2060, Melbourne, Australia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-1191>
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled

- attention. In *International Conference on Learning Representations*.
- Yufang Hou. 2020. Bridging anaphora resolution as question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1428–1438, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.132>
- Dirk Hovy, Taylor Berg-Kirkpatrick, Ashish Vaswani, and Eduard Hovy. 2013. Learning whom to trust with MACE. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 1120–1130, Atlanta, Georgia. Association for Computational Linguistics.
- Christoph Hube, Besnik Fetahu, and Ujwal Gadiraju. 2019. Understanding and mitigating worker biases in the crowdsourced collection of subjective judgments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–12, New York, NY. Association for Computing Machinery.
- Terne Sasha Thorn Jakobsen, Maria Barrett, Anders Søgaard, and David Lassen. 2022. The sensitivity of annotator bias to task definitions in argument mining. In *Proceedings of the 16th Linguistic Annotation Workshop (LAW-XVI) within LREC2022*, pages 44–61, Marseille, France. European Language Resources Association.
- Nanjiang Jiang and Marie-Catherine de Marneffe. 2022. Investigating reasons for disagreement in natural language inference. *Transactions of the Association for Computational Linguistics*, 10:1357–1374. https://doi.org/10.1162/tacl_a_00523
- Youxuan Jiang, Jonathan K. Kummerfeld, and Walter Lasecki. 2017. Understanding task design trade-offs in crowdsourced paraphrase collection. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 103–109, Vancouver, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P17-2017>
- David Jurgens. 2013. Embracing ambiguity: A comparison of annotation methodologies for crowdsourcing word sense labels. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 556–562, Atlanta, Georgia. Association for Computational Linguistics.
- Daisuke Kawahara, Yuichiro Machida, Tomohide Shibata, Sadao Kurohashi, Hayato Kobayashi, and Manabu Sassano. 2014. Rapid development of a corpus with discourse annotations using two-stage crowdsourcing. In *Proceedings of the International Conference on Computational Linguistics (COLING)*, pages 269–278, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Yudai Kishimoto, Shinnosuke Sawada, Yugo Murawaki, Daisuke Kawahara, and Sadao Kurohashi. 2018. Improving crowdsourcing-based annotation of Japanese discourse relations. In *LREC*.
- Wei-Jen Ko, Cutter Dalton, Mark Simmons, Eliza Fisher, Greg Durrett, and Junyi Jessy Li. 2021. Discourse comprehension: A question answering framework to represent sentence connections. *arXiv preprint arXiv:2111.00701*.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *Proceedings of MT Summit X*, pages 79–86. Phuket, Thailand.
- Yiwei Luo, Dallas Card, and Dan Jurafsky. 2020. Detecting stance in media on global warming. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3296–3315, online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.296>
- William C. Mann and Sandra A. Thompson. 1988. Rhetorical Structure Theory: Toward a functional theory of text organization. *Text-Interdisciplinary Journal for the Study of Discourse*, 8(3):243–281. <https://doi.org/10.1515/text.1.1988.8.3.243>
- Christopher D. Manning. 2006. Local textual inference: It’s hard to circumscribe, but you know it when you see it—and NLP needs it.

- Marian Marchal, Merel Scholman, Frances Yung, and Vera Demberg. 2022. Establishing annotation quality in multi-label annotations. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3659–3668, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Sewon Min, Julian Michael, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. Ambigqa: Answering ambiguous open-domain questions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5783–5797, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.466>
- Yixin Nie, Xiang Zhou, and Mohit Bansal. 2020. What can we learn from collective human opinions on natural language inference data? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9131–9143, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.734>
- Rebecca J. Passonneau 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/tacl_a_00185
- Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent disagreements in human textual inferences. *Transactions of the Association for Computational Linguistics*, 7:677–694. https://doi.org/10.1162/tacl_a_00293
- Joshua C. Peterson, Ruairidh M. Battleday, Thomas L. Griffiths, and Olga Russakovsky. 2019. Human uncertainty makes classification more robust. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9617–9626. <https://doi.org/10.1109/ICCV.2019.00971>
- Barbara Plank, Dirk Hovy, and Anders Søgaard. 2014. Linguistically debatable or just plain wrong? In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 507–511, Baltimore, Maryland. Association for Computational Linguistics. <https://doi.org/10.3115/v1/P14-2083>
- Massimo Poesio and Ron Artstein. 2005. The reliability of anaphoric annotation, reconsidered: Taking ambiguity into account. In *Proceedings of the Workshop on Frontiers in Corpus Annotations II: Pie in the Sky*, pages 76–83. <https://doi.org/10.3115/1608829.1608840>
- Massimo Poesio, Patrick Sturt, Ron Artstein, and Ruth Filik. 2006. Underspecification and anaphora: Theoretical issues and preliminary evidence. *Discourse processes*, 42(2):157–175. <https://doi.org/10.1207/s15326950dp4202-4>
- Vinodkumar Prabhakaran, Aida Mostafazadeh Davani, and Mark Diaz. 2021. On releasing annotator-level labels and information in datasets. In *Proceedings of The Joint 15th Linguistic Annotation Workshop (LAW) and 3rd Designing Meaning Representations (DMR) Workshop*, pages 133–138, Punta Cana, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.law-1.14>
- Valentina Pyatkin, Ayal Klein, Reut Tsarfaty, and Ido Dagan. 2020. QADiscourse-Discourse Relations as QA Pairs: Representation, crowdsourcing and baselines. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2804–2819, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.224>
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-2124>
- Ines Rehbein, Merel Scholman, and Vera Demberg. 2016. Annotating discourse relations in spoken language: A comparison of the PDTB and CCR frameworks. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation*

- (LREC'16), pages 1039–1046, Portorož, Slovenia. European Language Resources Association (ELRA).
- Stefan Riezler. 2014. On the problem of theoretical terms in empirical computational linguistics. *Computational Linguistics*, 40(1):235–245. https://doi.org/10.1162/COLI_a_00182
- Hannah Rohde, Anna Dickinson, Nathan Schneider, Christopher Clark, Annie Louis, and Bonnie Webber. 2016. Filling in the blanks in understanding discourse adverbials: Consistency, conflict, and context-dependence in a crowdsourced elicitation task. In *Proceedings of the 10th Linguistic Annotation Workshop (LAW X)*, pages 49–58, Berlin, Germany. <https://doi.org/10.18653/v1/W16-1707>
- Ted J. M. Sanders, Wilbert P. M. S. Spooren, and Leo G. M. Noordman. 1992. Toward a taxonomy of coherence relations. *Discourse Processes*, 15(1):1–35. <https://doi.org/10.1080/01638539209544800>
- Merel C. J. Scholman and Vera Demberg. 2017. Crowdsourcing discourse interpretations: On the influence of context and the reliability of a connective insertion task. In *Proceedings of the 11th Linguistic Annotation Workshop (LAW)*, pages 24–33, Valencia, Spain. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W17-0803>
- Merel C. J. Scholman, Tianai Dong, Frances Yung, and Vera Demberg. 2022a. Discogem: A crowdsourced corpus of genre-mixed implicit discourse relations. In *Proceedings of the Thirteenth International Conference on Language Resources and Evaluation (LREC'22)*, Marseille, France. European Language Resources Association (ELRA).
- Merel C. J. Scholman, Valentina Pyatkin, Frances Yung, Ido Dagan, Reut Tsarfaty, and Vera Demberg. 2022b. Design choices in crowdsourcing discourse relation annotations: The effect of worker selection and training. In *Proceedings of the Thirteenth International Conference on Language Resources and Evaluation (LREC'22)*, Marseille, France. European Language Resources Association (ELRA).
- Wei Shi and Vera Demberg. 2019. Learning to explicitate connectives with Seq2Seq network for implicit discourse relation classification. In *Proceedings of the 13th International Conference on Computational Semantics - Long Papers*, pages 188–199, Gothenburg, Sweden. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W19-0416>
- Rion Snow, Brendan O'Connor, Daniel Jurafsky, and Andrew Y. Ng. 2008. Cheap and fast—but is it good? Evaluating non-expert annotations for natural language tasks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 254–263, Honolulu, Hawaii. Association for Computational Linguistics.
- Wilbert P. M. S. Spooren and Liesbeth Degand. 2010. Coding coherence relations: Reliability and validity. *Corpus Linguistics and Linguistic Theory*, 6(2):241–266. <https://doi.org/10.1515/cllt.2010.009>
- Alexandra Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2020. A case for soft loss functions. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 8, pages 173–177. <https://doi.org/10.1609/hcomp.v8i1.7478>
- Alexandra N. Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2021. Learning from disagreement: A survey. *Journal of Artificial Intelligence Research*, 72:1385–1470. <https://doi.org/10.1613/jair.1.12752>
- Zeerak Waseem. 2016. Are you a racist or am I seeing things? Annotator influence on hate speech detection on Twitter. In *Proceedings of the First Workshop on NLP and Computational Social Science*, pages 138–142, Austin, Texas. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W16-5618>
- Bonnie Webber. 2009. Genre distinctions for discourse in the Penn TreeBank. 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 674–682. <https://doi.org/10.3115/1690219.1690240>

- Bonnie Webber, Rashmi Prasad, Alan Lee, and Aravind Joshi. 2019. The Penn Discourse Treebank 3.0 annotation manual. Philadelphia, University of Pennsylvania.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>
- Frances Yung, Vera Demberg, and Merel Scholman. 2019. Crowdsourcing discourse relation annotations by a two-step connective insertion task. In *Proceedings of the 13th Linguistic Annotation Workshop*, pages 16–25, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/W19-4003>
- Deniz Zeyrek, Amália Mendes, Yulia Grishina, Murathan Kurfalı, Samuel Gibbon, and Maciej Ogrodniczuk. 2019. Ted multilingual discourse bank (TED-MDB): A parallel corpus annotated in the PDTB style. *Language Resources and Evaluation*, 1–27.
- Deniz Zeyrek, Amália Mendes, Yulia Grishina, Murathan Kurfalı, Samuel Gibbon, and Maciej Ogrodniczuk. 2020. Ted multilingual discourse bank (TED-MDB): A parallel corpus annotated in the PDTB style. *Language Resources and Evaluation*, 54(2):587–613. <https://doi.org/10.1007/s10579-019-09445-9>
- Shujian Zhang, Chengyue Gong, and Eunsol Choi. 2021. Learning with different amounts of annotation: From zero to many labels. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7620–7632, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.emnlp-main.601>
- Šárka Zikánová, Jiří Mírovský, and Pavlína Synková. 2019. Explicit and implicit discourse relations in the Prague Discourse Treebank. In *Text, Speech, and Dialogue: 22nd International Conference, TSD 2019, Ljubljana, Slovenia, September 11–13, 2019, Proceedings 22*, pages 236–248. Springer. https://doi.org/10.1007/978-3-030-27947-9_20