

On Crowdsourcing Task Design for Discourse Relation Annotation

Frances Yung and Vera Demberg
Saarland University, Saarbrücken, Germany
{frances, vera}@lst.uni-saarland.de

Abstract

Interpreting implicit discourse relations involves complex reasoning, requiring the integration of semantic cues with background knowledge, as overt connectives like *because* or *then* are absent. These relations often allow multiple interpretations, best represented as distributions. In this study, we compare two established methods that crowdsource English implicit discourse relation annotation by connective insertion: a free-choice approach, which allows annotators to select any suitable connective, and a forced-choice approach, which asks them to select among a set of predefined options. Specifically, we re-annotate the whole DiscoGeM 1.0 corpus - initially annotated with the free-choice method - using the forced-choice approach. The free-choice approach allows for flexible and intuitive insertion of various connectives, which are context-dependent. Comparison among over 130,000 annotations, however, shows that the free-choice strategy produces less diverse annotations, often converging on common labels. Analysis of the results reveals the interplay between task design and the annotators' abilities to interpret and produce discourse relations.

1 Introduction

Disagreement in linguistic annotation is increasingly seen not as noise but as a valuable signal capturing diverse perspectives in language interpretation (Dumitrache et al., 2021; Uma et al., 2021; Frenda et al., 2024). A single gold label, traditionally provided by one or two trained annotators, often fails to capture the full range of interpretations, which may arise from linguistic ambiguity, contextual factors, or annotators' cultural and experiential backgrounds. Crowdsourcing offers a scalable solution for gathering these alternative interpretations.

To guide untrained crowd workers in reliably annotating abstract linguistic phenomenon, intu-

itive and carefully designed workflows are essential. Task design has been identified as one of the factors behind annotation disagreement and bias (Pavlick and Kwiatkowski, 2019; Jiang and de Marneffe, 2022), and even can impact annotation quality (Shaw et al., 2011; Gadiraju et al., 2017; Gurusurangan et al., 2018). For example, Pyatkin et al. (2023) investigate the bias of task design guiding workers in annotating implicit discourse relation (IDR) senses, which often have multiple interpretations. They compared two methods: one based on insertion of discourse connectives (DCs), e.g. *John fell down because he tripped*, and the other on paraphrasing discourse arguments to question-answer (QA) pairs, e.g. *Q: Why (did) John fell down? A: He tripped*. While annotations from both methods were found to align closely, subtle bias in the annotation preference are found in both methods due to limitations of using natural language to annotate specialized linguistic concepts (not all senses can be easily expressed by a connective or by a question).

Building on this line of work, we explore the potential method bias of two IDR annotation tasks for English based on DC insertion. These methods differ solely in whether annotators select from predefined options (Rohde et al., 2016; Yung et al., 2024) or freely type in their choices (Yung et al., 2019). The free-choice method was employed to annotate 6,500 English IDRs in the DiscoGeM 1.0 corpus (Scholman et al., 2022), whereas the DiscoGeM 2.0 corpus, comprising multi-lingual translations or original texts from DiscoGeM 1.0, was annotated using the forced-choice method. An initial comparison of the statistics of the two corpora revealed characteristics unique to the English annotations, such as a higher proportion of CONJUNCTION relations.

Our findings indicate that the free-choice approach achieves higher agreements among annotators, while the forced-choice approach is more

effective at capturing a diverse range of alternative interpretations. Further analysis reveals that the free-choice approach favours intuitive and frequent intuitive sense, whereas the provided options in the forced-choice approach serve as prompts for the workers to identify rare, fine-grained senses. Moreover, the method bias interacts with individual differences in discourse processing: workers who could identify a wider range of senses in one approach also tended to label more different senses in the other approach. These results highlight the nuanced impact of task design on annotation outcomes.

The re-annotated resource is freely downloadable¹ alongside the original DiscoGeM 1.0. It provides an interesting dataset for the study of perspectivism and design in annotation as well as a rich collection of rare IDR examples, contributing to the major data bottleneck for current IDR recognition models.

2 Related work

Annotation of IDR requires integrating subtle semantic cues with background knowledge and mapping these to abstract labels – a task that is challenging even for trained annotators (Hoek and Scholman, 2017). Previous attempts to create datasets by crowdsourcing annotations often compromise on label variety or annotation quality (Kawahara et al., 2014; Kishimoto et al., 2018).

Inspired by the Penn Discourse Treebank’s (PDTB) lexicalized approach to annotate IDRs (Prasad et al., 2019), prior work has proposed crowdsourcing IDRs via DC insertion. For example, to label the REASON relation between the arguments "John missed the bus" and "He was late to work.", the DC "therefore" could be inserted. In the initial proposal, crowd workers selected a DC from a fixed list, each corresponding to a unique IDR sense (Scholman and Demberg, 2017). While achieving high agreement with expert annotations, the method was tested on only 6 IDR senses to avoid overwhelming workers with too many options. Choosing a DC is often context-dependent; for example, while "although" and "even though" are nearly interchangeable, "also" versus "furthermore" (both indicating CONJUNCTION) may depend on context. Workers might reject an appropriate sense if a DC feels contextually awkward.

To handle a broader range of IDRs, Yung et al.

(2019) proposed a two-step approach: first, workers freely type a DC that fits between two arguments; second, they select from a list of unambiguous DCs corresponding to their free-choice. For instance, if they type "while" in the first step, they should choose between "at the same time" and "in contrast" in the second step, which are mapped to the relations SYNCHRONY and CONTRAST respectively. This method was used to create the DiscoGeM 1.0 corpus, which contains 6,500 English IDRs each annotated by 10 workers (Scholman et al., 2022). Nonetheless, DiscoGeM 2.0, which extends the annotations to German, French, and Czech (Yung et al., 2024), adopted the one-step forced-choice method: workers directly chose from 28 DC choices, which were grouped by semantics and shuffled per worker to facilitate navigation and avoid positional bias. The free- and forced-choice methods were reported to yield similar annotations, but the comparisons were based on a limited subset of items (234 in Yung et al. (2019) and 18 in Yung et al. (2024)), with a restricted range of IDR senses.

Using a different crowd-annotation method, Pyatkin et al. (2020) crowdsourced discourse relations by instructing workers to create QA pairs from the provided text, e.g., "Q: What is the reason John was late? A: He missed the bus." Comparisons of QA-based and free-choice DC insertion methods show that both exhibit biases toward specific sense categories. In contrast to common attribution of method artifacts to degraded data quality (Gururangan et al., 2018; Zhu and Rzeszotarski, 2024), it was found that training on the complementary data collected by both methods enhanced the performance of IDR identification models (Pyatkin et al., 2023).

3 Annotation experiment

We adopt the forced-choice approach to re-annotate the DiscoGeM 1.0 corpus, which was originally annotated using the free-choice approach. For this, an annotation interface was implemented based on the description of DiscoGeM 2.0 (Yung et al., 2024). One representative DC was selected for each of the 28 relations to be annotated. The selection was primarily based on the disambiguating DCs from the second step of the free-choice method,² while ensuring they were sufficiently frequent and not highly context dependent. The complete list is

¹<https://github.com/merelscholman/DiscoGeM>

²the DC lexicon and per-worker annotations are available together with the corpus

shown in Table 2 in the Appendix.

Following the procedure of DiscoGeM 1.0, native English-speaking crowd workers were recruited via the Prolific platform. Based on the anonymous Prolific worker IDs, we invited the 199 workers who contributed to DiscoGeM 1.0 to participate in the annotation task again. We assumed that they would not recall the texts they annotated three years ago and including them allows direct comparison of annotations from the same workers across both methods. Of these, 91 workers took part again, and 73 additional workers were recruited through a selection task.

Out of the 6505 items in DiscoGeM 1.0, 16 duplicates were identified and removed. The remaining items were divided into batches of 20 – 25, with each batch assigned to at least 10 workers. The workers were awarded £1.8 – £2.2 per batch.

The quality of DiscoGeM 1.0 annotations was primarily controlled by a screening task that selected candidates achieving at least 50% agreement with gold labels. During the data collection phase, annotation quality was monitored twice to identify and remove poorly performing annotators, while retaining their earlier annotations (Scholman et al., 2022). Similarly, we used an initial screening task to ensure annotation quality. However, to maximize the number of annotators participating in both tasks, we did not screen those who had contributed to DiscoGeM 1.0, nor did we conduct additional screening during the annotation process.

We compare the newly collected data against the original DiscoGeM 1.0. In addition to analyzing label distributions from 10 workers per item, we compared aggregated annotations to highlight the differences. The annotations were aggregated using the "Worker Agreement with Aggregate" (Wawa) algorithm, which weights each worker’s votes based on their overall agreement with the majority label (Ustalov et al., 2021).

4 Results

Table 1 presents the agreement between the annotations obtained by the two methods. We computed the averaged Jensen-Shanon divergence (JSD) between the label distributions of each item, as well as hard and soft agreement rates. Hard agreement measures matches between the single aggregated annotations, while soft agreement considers any overlap between annotations with over 20% distribution a match (Pyatkin et al., 2023). We also calcu-

lated the *soft* κ scores, an inter-annotator agreement metric that accounts for the increased chance agreement in multi-label predictions (Marchal et al., 2022).

<i>inter-method comparison</i>	free vs forced	
JSD (full dist.)	.527	
Hard agreement (single label)	.425	
Soft agreement (multi-labels)	.708	
Soft κ (multi-labels)	.663	
<i>intra-method comparison</i>	free	forced
Entropy	0.353	0.460
Agreement (max. label dist)	0.508	0.404
Per-item unique label count	4.309	6.275

Table 1: Annotation Agreement

It can be observed that the inter-method agreement between single aggregated annotations is moderate, comparable to the accuracy of state-of-the-art IDR classification models (Costa and Kosseim, 2024; Zeng et al., 2024), but the agreement is substantially higher when multiple annotations are considered. This demonstrates that both methods are capable of annotating the same types of relations, which often co-occur with other relations.

The bottom half of Table 1 compares the agreement among the 10 annotations per item in both methods. The forced-choice method shows higher averaged entropy in the per-item label distributions, indicating greater annotation uncertainty. In addition, the forced-choice approach yields smaller averaged per-item agreement (i.e., the proportion of the majority label) and a higher average number of unique annotations per item. These results all indicate lower annotator agreement in the forced-choice approach.

Figure 1 illustrates the overall distribution of the unaggregated annotations, computed by the sum of the normalized per-item distribution, since not all items have exactly 10 annotations. The free-choice approach clearly converges on a narrower set of labels, while the forced-choice approach spans a wider range. Notably, RESULT and CONJUNCTION, are selected twice as often in the free-choice method.

The trend is similar when focusing on the most agreed labels. Figure 2 shows the alignment of the aggregated annotations from both methods. The annotations are grouped at level-2 granularity according to the PDTB sense hierarchy, e.g. ARG1-AS-DETAIL and ARG2-AS-DETAIL are grouped as




Figure 1: Distribution of the **unaggregated** annotations

LEVEL-OF-DETAIL. Even though the darkest diagonal line in the confusion matrix indicates substantial agreement between annotations from both methods, many items labelled with CONJUNCTION, CAUSAL, and ARG2-AS-DETAIL in the free-choice approach are now assigned to a range of other relations. While the aggregated annotations from the forced-choice approach cover all level-2 senses defined in the framework, half of these senses never appear in the aggregated annotations from the free-choice method.

Next, we directly compare the annotations of the same workers. In total, we identified 3, 223 annotations per method that were annotated by the same worker on the same item (spanning 2, 542 unique items and 71 workers). The comparison of these annotations demonstrates a similar tendency as found in the re-annotation of the whole corpus, as shown in Figure 5 in the Appendix - common relations like CONJUNCTION and RESULT were annotated as other rarer relations in the forced-choice approach.

Figure 3 plots the number of unique relations identified by workers who participated in both methods. To ensure comparability, results from workers who annotated fewer than 50 items in either method or annotated items 3 times more in one method than the other were excluded. This results in 60 workers, who annotated on average 621 and 525 items in the free- and forced- choice methods respectively.




Figure 2: Confusion matrix of the **aggregated** annotations from both methods, with labels merged at level-2 granularity

It shows that all workers identified a broader range of IDR senses using the forced-choice method, as indicated by all data points falling below the diagonal line. Furthermore, workers who could identify more sense types with the free-choice approach also identify more sense types in the forced-choice approach. This suggests individual differences in sensitivity to the subtle contrast in fine-grained discourse relations, with the forced-choice method further expanding the range of relations these workers could identify by presenting all possible options.




Figure 3: Total number of **unique** relations annotated by the same workers on the same set of items

5 Discussion and conclusion

We examined the impact of two similar interfaces used to crowdsourcing IDR annotations. Using the free-choice approach, workers tend to select common IDR labels with higher inter-annotator agreement, while the forced-choice approach encouraged a larger variety of relations, including rare ones. Notably, both methods produce valid annotations, as evidenced by the high soft match agreement. Frequent senses can often be inferred alongside other senses, such as the CONJUNCTION sense in the examples in Figure 4. In these examples, the English forced-choice annotations align with other languages, despite being labeled as CONJUNCTION in the original free-choice annotations of DiscoGeM 1.0.

High inter-annotator agreement is often linked to higher data quality. However, for inherently ambiguous tasks like IDR identification, we showed that higher-agreement annotations that converge on common labels are not always superior. Recognizing the method bias enables tailoring the approach to the annotation goal — whether to achieve consensus on a single label or capture diverse perspectives. Since current IDR classification models often struggle with rare labels, datasets with more label variety may be more valuable. Still, distinguishing genuine perspectives from annotation errors is challenging. Minimal data cleaning, such as removing labels with very few votes, could be applied.

For corpus analysis, data should be collected consistently using the same method. Initial analysis reveals significant differences between the inter-annotator agreements of the English annotations in DiscoGeM 1.0 and the multilingual annotations in DiscoGeM 2.0, whereas the re-annotated data in this study aligns more closely with the other languages (e.g. averaged per-item agreement = .508/.404 (EN free-/forced-choice) .410 – .439 (DE, FR, CS forced-choice), indicating the influence of the method bias. Our next step is to analyze the cross-lingual difference based on annotations collected with the same method.

Acknowledgements

This project is supported by the German Research Foundation (DFG) under Grant SFB 1102 (“Information Density and Linguistic Encoding”, Project-ID 232722074).

1)

Arg 1: It was because of this tiny piece of information that Ford Prefect was not now a whiff of hydrogen, ozone and carbon monoxide. He heard a slight groan.
Arg2: By the light of the match he saw a heavy shape moving slightly on the floor. Quickly he shook the match out, reached in his pocket, found what he was looking for and took it out.

Aggregated annotation = PRECEDENCE

(English, German, French, Czech forced-choice)

Aggregated annotation = CONJUNCTION

(English free-choice)

2)

Arg1: In yesterday’s debate in the European Parliament some Members of this Parliament expressed worry that we were interfering in the internal affairs of a Member State.
Arg2: Such a concern is misplaced. The European Parliament has never been slow to comment on developments in Member States with which they disagree.

Aggregated annotation = REASON

(English, German, French, Czech forced-choice)

Aggregated annotation = CONJUNCTION

(English free-choice)

3)

Arg1: With a spring Gollum got up and started shambling off at a great pace. Bilbo hurried after him, still cautiously, though his chief fear now was of tripping on another snag and falling with a noise. His head was in a whirl of hope and wonder.
Arg2: It seemed that the ring he had was a magic ring: it made you invisible!

Aggregated annotation = SYNCHRONOUS

(English, German, French, Czech forced-choice)

Aggregated annotation = CONJUNCTION

(English free-choice)

Figure 4: Examples taken from DiscoGeM where the annotations by the forced- and free- choice approaches are alternative interpretations. The English forced-choice annotations come from the current study and those from the other languages come from DiscoGeM 2.0. The English free-choice annotations come from DiscoGeM 1.0.

References

- Nelson Filipe Costa and Leila Kosseim. 2024. A multi-task and multi-label classification model for implicit discourse relation recognition. *arXiv preprint arXiv:2408.08971*.
- 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.
- Simona Frenda, Gavin Abercrombie, Valerio Basile, Alessandro Pedrani, Raffaella Panizzon, Alessandra Teresa Cignarella, Cristina Marco, and Davide Bernardi. 2024. Perspectivist approaches to natural language processing: a survey. *Language Resources and Evaluation*, pages 1–28.
- Ujwal Gadiraju, Jie Yang, and Alessandro Bozzon. 2017. Clarity is a worthwhile quality: On the role of task clarity in microtask crowdsourcing. In *Proceedings of the 28th ACM conference on hypertext and social media*, pages 5–14.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. Annotation artifacts in natural language inference data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Jet Hoek and Merel Scholman. 2017. Evaluating discourse annotation: Some recent insights and new approaches. In *Proceedings of the 13th Joint ISO-ACL Workshop on Interoperable Semantic Annotation (isa-13)*.
- 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.
- 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 COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, 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 *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- 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.
- Ellie Pavlick and Tom Kwiatkowski. 2019. Inherent disagreements in human textual inferences. *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Rashmi Prasad, Bonnie Webber, Alan Lee, and Aravind Joshi. 2019. Penn Discourse Treebank Version 3.0.
- 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.
- Valentina Pyatkin, Frances Yung, Merel C. J. Scholman, Reut Tsarfaty, Ido Dagan, and Vera Demberg. 2023. Design choices for crowdsourcing implicit discourse relations: Revealing the biases introduced by task design. *Transactions of the Association for Computational Linguistics*, 11:1014–1032.
- Hannah Rohde, Anna Dickinson, Nathan Schneider, Christopher N. L. 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 held in conjunction with ACL 2016 (LAW-X 2016)*, pages 49–58, Berlin, Germany. Association for Computational Linguistics.
- Merel 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*, pages 24–33, Valencia, Spain. Association for Computational Linguistics.
- Merel C. J. Scholman, Tianai Dong, Frances Yung, and Vera Demberg. 2022. 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).
- Aaron D Shaw, John J Horton, and Daniel L Chen. 2011. Designing incentives for inexpert human raters. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, pages 275–284.
- 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.
- Dmitry Ustalov, Nikita Pavlichenko, Vladimir Losev, Iulian Giliyazev, and Evgeny Tulin. 2021. A general-purpose crowdsourcing computational quality control

toolkit for python. In *The Ninth AAAI Conference on Human Computation and Crowdsourcing: Works-in-Progress and Demonstration Track (HCOMP 2021)*.

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.

Frances Yung, Merel Scholman, Sarka Zikanova, and Vera Demberg. 2024. [DiscoGeM 2.0: A parallel corpus of English, German, French and Czech implicit discourse relations](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4940–4956, Torino, Italia. ELRA and ICCL.

Lei Zeng, Ruifang He, Haowen Sun, Jing Xu, Chang Liu, and Bo Wang. 2024. [Global and local hierarchical prompt tuning framework for multi-level implicit discourse relation recognition](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 7760–7773, Torino, Italia. ELRA and ICCL.

Shengqi Zhu and Jeffrey Rzeszotarski. 2024. [“Get their hands dirty, not mine”: On researcher-annotator collaboration and the agency of annotators](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 8773–8782, Bangkok, Thailand. Association for Computational Linguistics.

A Appendix

level-2.level-3 IDR sense label	DC
Temporal SYNCHRONOUS.SYNCHRONOUS ASYNCHRONOUS.PRECEDENCE ASYNCHRONOUS.SUCCESION	at the same time then after
Contingency CAUSE.REASON CAUSE.RESULT PURPOSE.ARG1-AS-GOAL PURPOSE.ARG2-AS-GOAL CONDITION.ARG1-AS-COND CONDITION.ARG1-AS-NEGCOND CONDITION.ARG2-AS-COND CONDITION.ARG2-AS-NEGCOND	because as a result for that purpose so that in that case if not if unless
Comparison CONCESSION.ARG1-AS-DENIER CONCESSION.ARG2-AS-DENIER CONTRAST.CONTRAST COMPARISON.SIMILARITY.SIMILARITY	even though nonetheless on the other hand similarly
Expansion EQUIVALENCE.EQUIVALENCE INSTANTIATION.ARG1-AS-INSTANCE INSTANTIATION.ARG2-AS-INSTANCE LEVEL-OF-DETAIL.ARG1-AS-DETAIL LEVEL-OF-DETAIL.ARG2-AS-DETAIL CONJUNCTION.CONJUNCTION DISJUNCTION.DISJUNCTION EXCEPTION.ARG1-AS-EXCPT EXCEPTION.ARG2-AS-EXCPT MANNER.ARG1-AS-MANNER MANNER.ARG2-AS-MANNER SUBSTITUTION.ARG1-AS-SUBST SUBSTITUTION.ARG2-AS-SUBST NOREL	in other words this illustrates that for example in short in more detail also or other than that an exception is that thereby as if rather than instead (no direct relation)

Table 2: English DC choices used in the forced-choice DC insertion method

	synchronous -	precedence -	succession -	reason -	result -	arg1-as-goal -	arg2-as-goal -	arg1-as-cond -	arg1-as-negcond -	arg2-as-cond -	arg2-as-negcond -	arg1-as-denier -	arg2-as-denier -	contrast -	similarity -	equivalence -	arg1-as-instance -	arg2-as-instance -	arg1-as-detail -	arg2-as-detail -	conjunction -	disjunction -	arg1-as-excpt -	arg2-as-excpt -	arg1-as-manner -	arg2-as-manner -	arg1-as-subst -	arg2-as-subst -	norel -	
synchronous -	9	18	1	8	2	1	1	1																						
precedence -	14	78	21	6	16	3	2	6				3		5	3	2	3	1	1	5	21	9		1		5			5	6
succession -				3	1	2						1	1	2	1					1		6	2							1
reason -	5	3		88	8	4	3	7		2		2	1	3	3	2	3	9	8	23	9		1		5	1	1	6	5	
result -	25	33	8	34	171	50	20	36	1	2		15	26	12	13	28	15	14	43	55	41	1	4		25		2	14	21	
arg1-as-goal -						1																1								
arg2-as-goal -																	1	1												
arg1-as-cond -																														
arg1-as-negcond -																														
arg2-as-cond -																														
arg2-as-negcond -																														
arg1-as-denier -	8	1		8	5	3		4				16	12	3	1	1	1	5	2	7	4		1		2			1	7	
arg2-as-denier -	10	8		9	16	6	1	10				19	47	10	3	2	2	5	9	9	16	1	2	2	2		2	10	2	
contrast -	1	2		11	10			3	1			5	5	18	3	3	3	4	3	10	1	1	2	1				9	5	
similarity -	5	4	1	1	1	1		1						2	12	2		3	2	5	9		1					1	2	
equivalence -	1			1	1							1					1		1	1	1									
arg1-as-instance -	2	1		1	1		1					1					2		1	1			1						1	
arg2-as-instance -	10	3	1	16	4	5	4	3				5	3	3	4	6	4	27	6	29	11			1	1	1	1	6	3	
arg1-as-detail -	4	2		10	6	7	1	3				3			1	4	3	3	10	9	4		1		1	1		1	5	
arg2-as-detail -	14	16	2	30	34	23	4	20		1		10	9	4	12	22	16	38	23	202	36	1	1	7	1			10	20	
conjunction -	57	36	4	47	37	23	9	16	2	1	1	23	23	8	40	12	7	35	20	110	164			6	1	10		5	30	
disjunction -																						1	1	1						
arg1-as-excpt -																														
arg2-as-excpt -																														
arg1-as-manner -																							1							
arg2-as-manner -						1																		1						
arg1-as-subst -																														
arg2-as-subst -				3	2	1								1	1								1	2				7	1	
norel -	1					1						1		1													1		3	

Figure 5: Comparison between 3233 annotations by the same workers on the same items using both methods