

Overview of the BioCreative V Chemical Disease Relation (CDR) Task

Chih-Hsuan Wei¹, Yifan Peng^{1,2}, Robert Leaman¹, Allan Peter Davis³,
Carolyn J. Mattingly³, Jiao Li⁴, Thomas C. Wiegers³, Zhiyong Lu^{1,*}

¹National Center for Biotechnology Information, Bethesda, MD, USA 20894

²Department of Computer and Information Sciences, University of Delaware, Newark, DE

³Department of Biological Sciences, North Carolina State University, Raleigh, NC, USA

⁴Institute of Medical Information, Chinese Academy of Medical Sciences, Beijing, China

chih-hsuan.wei@nih.gov, yfpeng@udel.edu,
robert.leaman@nih.gov, apdavis3@ncsu.edu,
cjmattin@ncsu.edu, li.jiao@imicams.ac.cn,
twiegers@ncsu.edu, *zhiyong.lu@nih.gov

Abstract. Manually curating chemicals, diseases, and their relations is of significant importance to biomedical research but is plagued by its high cost and the rapid growth of the biomedical literature. In recent years, there has been a growing interest to develop computational approaches for automatic chemical-disease relation (CDR) extraction with proposals of different techniques. Despite these attempts, the lack of a comprehensive benchmarking dataset has limited the comparison of different techniques in order to assess and advance the current state of the art. To this end, we set up a challenge task through BioCreative V to automatically extract CDRs from the literature. More specifically, we designed two challenge tasks: disease named entity recognition (DNER) and chemical-induced disease (CID) relation extraction. To assist system development and assessment, we created a large annotated text corpus that consists of human annotations of all chemicals, diseases and their interactions in 1,500 PubMed articles. A total of 34 teams worldwide participated in the CDR task: 16 in the DNER task and 18 in the CID task. When comparing the text-mined results with the manually annotated ground truth, the best systems achieved an F-score of 86.46 for the DNER task – a result that approaches the human inter-annotator agreement (0.8875) – and an F-score of 57.03 for the CID task, the highest results ever reported for such tasks. In addition to the accuracy, another novel aspect of our evaluation is that we tested each participating system’s ability to return real-time results in a timely manner: the average response time for each team’s DNER and CID systems are 5.6 and 9.3 seconds via their respective web services. Given the level of participant and team results, we find our task to be successful in engaging the text-mining research community, producing a large annotated corpus, and improving the results of automatic disease recognition and chemical-disease relation extraction.

Keywords. Text Mining; Disease Recognition; DNorm; Relation Extraction;

1 Introduction & Motivation

Chemicals, diseases, and their relations are among the most searched topics by PubMed users worldwide (1,2), reflecting their central roles in many areas of biomedical research and healthcare such as drug discovery and safety surveillance. Developing a drug takes time and money: on average, around 14 years and \$2 billion or more (3). More than 95 percent of potential drugs fail during development for reasons such as undesired side effects due to either off-target binding or unanticipated physiologic roles of the intended target (4). Although the ultimate goal in drug discovery is to develop chemicals for therapeutics, recognition of adverse drug reactions (ADRs) between chemicals and diseases is important for improving chemical safety and toxicity studies and facilitating new screening assays for pharmaceutical compound survival. In addition, ADRs are an integral part of drug post-marketing surveillance. Identification of chemicals as biomarkers can also be helpful in informing potential relationships between chemicals and pathologies. Hence, manual annotation of such mechanistic and biomarker/correlative chemical-disease relations (CDR) from unstructured free text into structured knowledge to facilitate identification of potential toxicity has been an important theme for several bioinformatics databases, such as the Comparative Toxicogenomics Database (CTD) (5). NOTE: We consider the words ‘drug’ and ‘chemical’ to be interchangeable in this document.

Manual curation of CDRs from the literature is costly and insufficient to keep up with the rapid literature growth. In response, there have been many attempts to extract such relations by automated natural language processing (NLP) methods. Over the years, a wide variety of relation extraction approaches have been proposed, such as simple co-occurrence, pattern matching, machine learning, and knowledge-driven methods (6-8). A small number of test corpora were also developed, but they are limited in size and annotation scope (9,10). More recently, a similar set of computational methods has been applied to a number of diverse data sets such as the FDA’s Adverse Event Reporting System (FAERS) (11), electronic medical records (12), tweets, user comments in social media (13), etc. In comparison, the scholarly publications contain richer information about drug-induced phenomena in a variety of

settings, such as in vitro and in vivo methods, across species, for approved indications, off-label uses and for drugs in development.

Despite these previous attempts and other closely related studies (e.g. PPI (14)), automatic biomedical relation detection from free text remains challenging, from identifying relevant concepts (e.g. diseases (15-18)), to extracting relations. The lack of a comprehensive benchmarking dataset has limited the comparison of different computational techniques in order to assess and improve the current state of the art. In addition, few previous software tools for relation extraction have been made freely available and, to the best of our knowledge, been incorporated into practical applications such as biocuration.

2 Task

Through BioCreative V, one of the major formal evaluation events (19) for BioNLP research, we organized a challenge task of automatic extraction of mechanistic and biomarker chemical-disease relations from the biomedical literature with the goal of supporting biocuration, new drug discovery and drug safety surveillance. More specifically, we designed two subtasks:

- (A) Disease Named Entity Recognition (DNER). An intermediate step for automatic CDR extraction is disease named entity recognition and normalization, which was found to be highly difficult on its own in previous BioCreative CTD tasks (20,21), and other studies (18). For the subtask, participating systems were given raw PubMed abstracts and asked to return normalized disease concept identifiers.
- (B) Chemical-induced disease relation extraction (CID). Participating systems were provided with raw text from PubMed articles as input (same as DNER input) and asked to return a ranked list of <chemical, disease> pairs with normalized concept identifiers for which drug-induced diseases are associated in the abstract.

Note that both chemical and diseases were described using the National Library of Medicine's Medical Subject Headings (MeSH) controlled

vocabulary. Systems were required to return entity pairs; both entities needed to be normalized into MeSH identifiers, along with their text spans in the article.

Task Data

For our task, we prepared a total of 1,500 PubMed articles: 500 each for the training, development and test set. Of all 1,500 articles, most (1,400) were selected from an existing CTD-Pfizer collaboration-related dataset (see details below). The remaining 100 articles represented completely new curation and were incorporated into the test set.

For both tasks, we prepared manual annotations. For the DNER task, a number of MeSH annotators were recruited to annotate every disease and chemical occurrence in the abstract with both text spans and concept identifiers. We refer readers to (22) for more details regarding this annotation.

During a previous collaboration with Pfizer (23), CTD curated over 150,000 chemical-disease interactions. CTD biocurators followed CTD’s rigorous curation process and curated interactions from just the abstract whenever possible, except in cases where referencing the full text was necessary to resolve relevant issues mentioned in the abstract. For the CDR task, we mostly leveraged existing curated data for the 1,400 articles. The relation data for the additional 100 articles was generated during the CDR challenge by CTD staff, and this curation was not made public until the challenge was complete.

Table 1 describes the disease and relation annotations for the three data sets.

Table 1 Statistics of the CDR data sets

Task Dataset	Articles	Chemical		Disease		CID relation
		Mention	ID	Mention	ID	
Training	500	5,203	1,467	4,182	1,965	1,038
Development	500	5,347	1,507	4,244	1,865	1,012
Test	500	5,385	1,435	4,424	1,988	1,066

Task Evaluation

For final evaluation of the participant systems, text-mined entities (diseases) and relations (<chemical, disease> pairs) were compared to manually annotated data using standard precision, recall and F-score metrics. More specifically, the DNER results are evaluated by comparing the set of disease concepts annotated within the document with the set of disease concepts predicted by the participant system. Similarly, the CID results are evaluated by comparing the set of chemical-disease relationships annotated within the document with the set of chemical-disease relationships predicted by the system.

For results submission, participants followed the procedure implemented for the previous BioCreative-CTD task (20) where teams submitted their results through web services¹. In particular, Representational State Transfer (REST) was selected as the architectural style for the participant web services. To assist participants, the organizers provided executable files together with a step-by-step installation guide. Also, a testing web site was provided to the teams in order to simulate the exact system testing environment. Indeed, the testing facility was heavily used (over 2,400 times by dozens of teams) since its inception. Because of the online evaluation, we are able to report the response time of each system in addition to their accuracies.

Benchmarking Systems

For comparison purposes, we benchmarked several systems for the DNER and CID tasks

For DNER, we first developed a straightforward dictionary look-up baseline approach that relied on disease names from CTD. We also re-trained models using the out-of-box DNorm, our previous work for disease named entity recognition and normalization (15). DNorm combines an approach based on rich features and conditional random fields for named entity recognition (using BANNER (24)) with a novel machine learning method for normalization based on pairwise learning to rank. DNorm is a competitive system which achieved the highest performance in a previous disease challenge (16,25); its performance therefore provides a very strong benchmark.

¹ We allowed offline submissions for manual runs.

For the CID task, we implemented a simple co-occurrence baseline method with two variants: abstract-level and sentence-level. The chemical and disease entities were automatically recognized using our in-house tools, DNorm (15) and tmChem (26), respectively.

3 Results

A total of 34 teams participated in the CDR task: 16 teams participated in the DNER task, and 18 teams participated in the CID task. Since each team was allowed to submit up to 3 runs (i.e., 3 different version of their tool) for each task, a total of 86 runs were submitted. Of the 34 teams, there were 25 unique teams from 12 different countries in four continents Australia (1), Asia (12), Europe (9), and North America (3).

DNER Results

A total of 16 teams successfully submitted DNER results in 40 runs. As shown in Table 2 (only the best run of each team is included), multiple teams achieved an F-score higher than 85% with the highest being 86.46% (team 314), a result that approaches the inter-annotator agreement of the human annotators (0.8875) (27). The average precision, recall and F-score were 78.99%, 74.81% and 76.03%, respectively.

All teams but one achieved a higher F-score than our baseline dictionary method, which obtained an F-score of 52.30%. While we did not perform any additional development on DNorm to adapt it to this dataset, it sets a significantly stronger benchmark with a performance of 80.64% F-score. A total of 7 teams achieved performance higher than DNorm.

Table 2 DNER results are shown for each participating team (anonymously identified by team number), as well as the baseline (dictionary look up) and DNorm systems. Among different team runs, only the best results are shown.

System	TP	FP	FN	P	R	F
Dictionary look-up	1,341	1,799	647	42.71	67.45	52.30
DNorm	1,593	370	395	81.15	80.13	80.64
Team 276	1,545	549	443	73.78	77.72	75.70
Team 277	1,629	191	359	89.51	81.94	85.56
Team 285	1,249	892	739	58.34	62.83	60.50

Overview of the BioCreative V Chemical Disease Relation (CDR) Task

System	TP	FP	FN	P	R	F
Team 288	1,669	339	319	83.12	83.95	83.53
Team 290	1,284	712	704	64.33	64.59	64.46
Team 293	1,278	503	710	71.76	64.29	67.82
Team 296	708	66	1,280	91.47	35.61	51.27
Team 304	1,713	277	275	86.08	86.17	86.12
Team 309	1,372	684	616	66.73	69.01	67.85
Team 310	1,627	247	361	86.82	81.84	84.26
Team 314	1,660	192	328	89.63	83.50	86.46
Team 315	1,502	335	486	81.76	75.55	78.54
Team 325	1,661	339	327	83.05	83.55	83.30
Team 363	1,606	168	382	90.53	80.78	85.38
Team 364	1,703	606	285	73.75	85.66	79.26
Team 365	1,590	582	398	73.20	79.98	76.44
Team Avg	1,487	418	501	78.99	74.81	76.03

CID Results

A total of 18 teams successfully submitted CID results in 46 runs. As shown in Table 3 (only the best run of each team is included), the F-score ranges from 32.01% to 57.03% (team 288) with an average of 43.37%. All teams outperformed the baseline results by the simple abstract-level co-occurrence method (16.43% in precision, 76.45% in recall and 27.05% in F-score).

Table 3 CID results are shown for each participating team (anonymously identified by team number), as well as two variants of the co-occurrence baseline method (i.e. abstract- and sentence-level). Among different team runs, only the best results are shown.

Team	TP	FP	FN	P	R	F
Abstract-level	815	4,145	251	16.43	76.45	27.05
Sentence-level	570	1,672	496	25.42	53.47	34.46
Team 276	574	544	492	51.34	53.85	52.56
Team 288	623	496	443	55.67	58.44	57.03
Team 289	358	346	708	50.85	33.58	40.45
Team 290	346	536	720	39.23	32.46	35.52
Team 293	354	296	712	54.46	33.21	41.26
Team 299	321	261	745	55.15	30.11	38.96
Team 303	241	199	825	54.77	22.61	32.01
Team 304	552	497	514	52.62	51.78	52.20

Team	TP	FP	FN	P	R	F
Team 310	602	1,099	464	35.39	56.47	43.51
Team 316	454	633	612	41.77	42.59	42.17
Team 322	341	462	725	42.47	31.99	36.49
Team 334	441	615	625	41.76	41.37	41.56
Team 335	351	390	715	47.37	32.93	38.85
Team 338	576	635	490	47.56	54.03	50.59
Team 341	408	432	658	48.57	38.27	42.81
Team 363	506	493	560	50.65	47.47	49.01
Team 364	595	1,835	471	24.49	55.82	34.04
Team 365	532	464	534	53.41	49.91	51.60
Team Avg	454	569	612	47.09	42.61	43.37

Response Time Results

The average response time for DNER teams was 5.57 seconds, with a standard deviation of 6.1 (Figure 1), ranging from 0.053 to 19.4 seconds per request. The average response time for CID teams was 8.38 seconds, with a standard deviation of 6.5, ranging from 0.119 to 27.8 seconds.

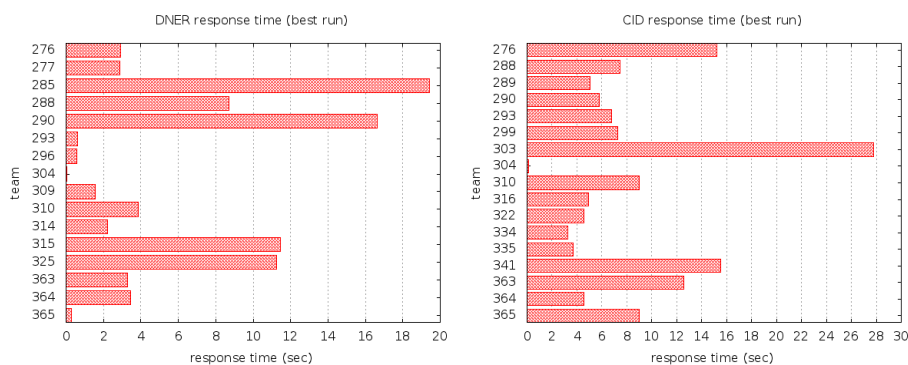


Figure 1 Average response time of each individual team for DNER and CID tasks

4 Discussion & Conclusions

Given the level of participation and team results, we conclude that the CDR challenge task was run successfully and is expected to make significant contributions to both the text mining and biocuration research communities. To the best of our knowledge, the constructed corpus is

the largest of its kind for both disease annotations and disease-chemical relations. In addition, our corpus includes both the text spans and normalized concept identifiers of entity annotations, as well as relation annotations, in the same abstract. We believe such a data set will be invaluable in advancing text-mining techniques for relation extraction tasks. Furthermore, our annotated data includes approximately 30% of the CDR relations that are asserted across sentence boundaries (i.e., not in the same sentences).

Unlike most challenge tasks in BioNLP (19), our task was designed to provide practical benefits to assist literature-based biocuration through two distinctive requests: a) all text-mined entities and relations should be normalized to database identifiers so that they can be readily useful to database curation; and b) through web services, biocuration groups can remotely request text-mined results in real-time without additional investment in text-mining tool adoption and technical infrastructure. By doing so, we hope that the state-of-the-art will be advanced for BioNLP systems toward higher standards for interoperability and scalability in future development.

5 Acknowledgment

This research was supported by the NIH Intramural Research Program, National Library of Medicine and the National Institute of Environmental Health Sciences (ES014065 and ES019604).

REFERENCES

1. Doğan RI, Murray GC, Névéal A, et al. (2009) Understanding PubMed user search behavior through log analysis. *Database (Oxford)* **2009**: bap018.
2. Névéal A, Doğan RI, Lu Z (2011) Semi-automatic semantic annotation of PubMed queries: a study on quality, efficiency, satisfaction. *Journal of Biomedical Informatics* **44**: 310-318.
3. Li J, Zheng S, Chen B, et al. (2015) A survey of current trends in computational drug repositioning. *Briefings in Bioinformatics* **1**: 11.
4. Hurlle MR, Yang L, Xie Q, et al. (2013) Computational drug repositioning: from data to therapeutics. *Clinical Pharmacology & Therapeutics* **93**: 335-341.

5. Davis AP, Murphy CG, Saraceni-Richards CA, et al. (2009) Comparative Toxicogenomics Database: a knowledgebase and discovery tool for chemical-gene-disease networks. *Nucleic Acids Res* **37**: D786-D792.
6. Kang N, Singh B, Bui C, et al. (2014) Knowledge-based extraction of adverse drug events from biomedical text. *BMC Bioinformatics* **15**: 64.
7. Xua R, Wang Q (2014) Automatic construction of a large-scale and accurate drug-side-effect association knowledge base from biomedical literature. *Journal of Biomedical Informatics* **51**: 191-199.
8. Gurulingappa H, Mateen-Rajput A, Toldo L (2012) Extraction of potential adverse drug events from medical case reports. *J Biomed Semantics* **3**: 15.
9. Mulligen EMV, Fourrier-Reglat A, Gurwitz D, et al. (2012) The EU-ADR corpus: annotated drugs, diseases, targets, and their relationships. *Journal of Biomedical Informatics* **45**: 879-884.
10. Gurulingappa H, Rajput AM, Roberts A, et al. (2012) Development of a benchmark corpus to support the automatic extraction of drug-related adverse effects from medical case reports. *Journal of Biomedical Informatics* **45**: 885-892.
11. Harpaz R, Vilar S, DuMouchel W, et al. (2013) Combing signals from spontaneous reports and electronic health records for detection of adverse drug reactions. *Journal of the American Medical Informatics Association* **20**: 413-419.
12. Iyer SV, Harpaz R, LePendu P, et al. (2014) Mining clinical text for signals of adverse drug-drug interactions. *Journal of the American Medical Informatics Association* **21**: 353-362.
13. Leaman R, Wojtulewicz L, Sullivan R, et al. (2010) Towards internet-age pharmacovigilance: extracting adverse drug reactions from user posts to health-related social networks. Proceedings of the 2010 Workshop on Biomedical Natural Language Processing. Uppsala, Sweden: Association for Computational Linguistics. pp. 117-125.
14. Krallinger M, Vazquez M, Leitner F, et al. (2011) The Protein-Protein Interaction tasks of BioCreative III: classification/ranking of articles and linking bio-ontology concepts to full text. *BMC Bioinformatics* **12**: S3.
15. Leaman R, Doğan RI, Lu Z (2013) DNorm: disease name normalization with pairwise learning to rank. *Bioinformatics* **29**: 2909-2917.
16. Leaman R, Khare R, Lu Z (2013) NCBI at 2013 ShARe/CLEF Shared Task: Disorder Normalization in Clinical Notes with DNorm. Proceedings of the CLEF 2013 Evaluation Labs and Workshop. Valencia, Spain.
17. Doğan RI, Leaman R, Lu Z (2014) NCBI disease corpus: A resource for disease name recognition and concept normalization. *Journal of Biomedical Informatics* **47**: 1-10.
18. Leaman R, Khare R, Lu Z (2015) Challenges in clinical natural language processing for automated disorder normalization. *Journal of Biomedical Informatics* **57**: 28-37.
19. Huang C-C, Lu Z (2015) Community challenges in biomedical text mining over 10 years: success, failure and the future. *Briefings in Bioinformatics*.
20. Wieggers TC, Davis AP, Mattingly CJ (2014) Web services-based text-mining demonstrates broad impacts for interoperability and process simplification. *Database (Oxford)* **2014**: bau050.
21. Wieggers TC, Davis AP, Mattingly CJ (2012) Collaborative biocuration-text-mining development task for document prioritization for curation *Database* **2012**: bas037

Overview of the BioCreative V Chemical Disease Relation (CDR) Task

22. Li J, Sun Y, Johnson R, et al. (2015) Annotating chemicals, diseases, and their interactions in biomedical literature. Proceedings of the fifth BioCreative challenge evaluation workshop. Sevilla, Spain.
23. Davis AP, Wieggers TC, Roberts PM, et al. (2013) A CTD-Pfizer collaboration: manual curation of 88,000 scientific articles text mined for drug-disease and drug-phenotype interactions. *Database (Oxford)* **2013**: bat080.
24. Leaman R, Gonzalez G (2008) BANNER: an executable survey of advances in biomedical named entity recognition. *Pacific Symposium on Biocomputing*: 652-663.
25. Pradhan S, Elhadad N, South BR, et al. (2014) Evaluating the state of the art in disorder recognition and normalization of the clinical narrative. *Journal of the American Medical Informatics Association*: 143-154.
26. Leaman R, Wei C-H, Lu Z (2015) tmChem: a high performance approach for chemical named entity recognition and normalization. *Journal of Cheminformatics* **7**: S3.
27. Li J, Sun Y, Johnson RJ, et al. (2015) Annotating chemicals, diseases and their interactions in biomedical literature. Proceedings of the fifth BioCreative challenge evaluation workshop.

Appendix of Wei CH, Peng Y, Leaman R, Davis AP, Mattingly CJ, Li J, Wiegiers TC, Lu Z. Overview of the BioCreative Chemical Disease Relation (CDR) Task. In Proceedings of the fifth BioCreative Workshop BioCreative challenge evaluation workshop, Sevilla, Spain

Table 3. Team results for the CDR task

Team ID	Run	Concept-level			Mention-level			Avg. Response time (msec)
		P	R	F	P	R	F	
276	1	73.78	77.72	75.70	75.39	73.53	74.45	2,932.0
277	1	88.89	82.14	85.39	87.67	81.35	84.39	2,030.0
	2	89.51	81.94	85.56	87.67	81.35	84.39	2,870.1
	3	89.89	81.44	85.46	87.67	81.35	84.39	2,571.0
285	1	58.34	62.83	60.50	82.45	76.27	79.24	19,408.8
	2	45.95	62.73	53.04	61.38	77.98	68.69	22,546.7
	3	46.68	63.58	53.83	60.91	77.31	68.13	22,598.1
288	1	83.12	83.95	83.53	86.89	82.10	84.43	8,700.7
	2	83.12	83.95	83.53	86.89	82.10	84.43	9,436.4
	3	82.54	83.95	83.24	86.48	82.30	84.34	8,313.6
290	1	64.33	64.59	64.46	82.97	78.16	80.49	16,617.0
293	1	71.76	64.29	67.82	89.22	84.22	86.65	627.2
	2	71.11	63.88	67.30	89.21	84.45	86.76	600.0
	3	70.77	63.68	67.04	87.94	85.22	86.56	762.4
296	1	91.47	35.61	51.27	81.78	72.85	77.06	587.9
	2	91.06	35.36	50.94	81.11	72.49	76.56	568.9
	3	91.30	35.36	50.98	81.87	72.56	76.93	565.0
304	1	86.08	86.17	86.12	87.26	83.79	85.49	53.1
	2	86.08	86.17	86.12	87.26	83.79	85.49	45.0
	3	86.08	86.17	86.12	87.26	83.79	85.49	53.3
309	1	64.08	67.76	65.87	84.58	77.24	80.74	2,602.3
	2	66.43	66.40	66.42	84.58	77.10	80.67	1,069.0
	3	66.73	69.01	67.85	84.58	77.24	80.74	1,552.4
310	1	86.82	81.84	84.26	84.15	82.21	83.17	3,842.9
314	1	89.42	82.44	85.79	86.28	81.44	83.79	2,242.3
	2	89.63	83.50	86.46	87.34	83.75	85.51	2,208.5
	3	88.32	83.65	85.92	86.04	83.43	84.71	2,211.7
315	1	81.76	75.55	78.54	81.54	72.78	76.91	11,458.5
	2	79.74	76.01	77.83	79.49	73.58	76.42	13,396.3
	3	83.93	72.48	77.79	82.87	70.19	76.00	10,939.8
325	1	81.14	85.26	83.15	82.28	82.28	82.28	5,212.7
	2	83.05	83.55	83.30	87.36	80.47	83.77	11,239.8
	3	60.66	86.62	71.35	63.80	85.10	72.93	11,627.0
363	1	90.37	80.28	85.03	82.39	85.15	83.75	3,563.5
	2	90.53	80.78	85.38	83.31	86.19	84.72	3,283.5
	3	90.92	80.13	85.19	83.41	85.47	84.43	2,630.4
364	1	73.69	85.11	78.99	74.47	79.18	76.75	4,324.3
	2	73.75	85.66	79.26	74.51	79.61	76.98	3,470.0
	3	73.75	85.66	79.26	74.51	79.61	76.98	3,431.3
365	1	73.20	79.98	76.44	79.65	76.70	78.14	276.0

Overview of the BioCreative V Chemical Disease Relation (CDR) Task

Table 4. Team results for the CID task

Team ID	Run	All 500 articles			Avg. Response time (msec)
		P	R	F	
276	1	51.34	53.85	52.56	15,202.1
	2	55.87	47.75	51.49	15,439.4
	3	58.92	40.90	48.28	15,193.1
288	1	56.60	55.91	56.25	7,935.7
	2	56.65	57.13	56.89	7,364.1
	3	55.67	58.44	57.03	7,425.4
289	1	50.85	33.58	40.45	5,070.7
290	1	39.23	32.46	35.52	5,768.2
293	1	55.97	32.08	40.79	24,863.5
	2	54.31	33.11	41.14	10,651.3
	3	54.46	33.21	41.26	6,799.9
299	1	53.33	29.27	37.80	7,660.1
	2	55.15	30.11	38.96	7,270.7
	3	54.18	30.39	38.94	7,275.0
303	1	28.51	23.08	25.51	27,506.7
	2	26.97	16.70	20.63	27,794.0
	3	54.77	22.61	32.01	18,145.5
304	1	57.65	36.77	44.90	96.9
	2	60.99	35.93	45.22	121.8
	3	52.62	51.78	52.20	119.3
310	1	35.39	56.47	43.51	9,024.4
316	1	47.89	36.21	41.24	5,002.2
	2	41.77	42.59	42.17	4,943.1
	3	44.95	38.84	41.67	4,915.2
322	1	42.47	31.99	36.49	4,538.4
	2	39.30	31.71	35.10	4,471.4
334	1	41.69	41.18	41.43	2,175.5
	2	41.76	41.37	41.56	3,266.5
	3	41.38	40.06	40.71	5,434.4
335	1	60.80	20.08	30.18	2,281.1
	2	47.37	32.93	38.85	3,676.7
	3	51.27	30.21	38.02	5,029.9
338	1	16.28	75.14	26.76	--
	2	47.56	54.03	50.59	--
341	1	48.57	38.27	42.81	15,551.7
	2	31.89	61.54	42.01	15,958.8
	3	31.89	61.54	42.01	15,854.5
363	1	50.65	47.47	49.01	12,539.0
	2	46.73	48.97	47.82	4,132.3
	3	48.61	47.47	48.03	4,334.7
364	1	24.49	55.82	34.04	4,575.2
	2	24.28	55.35	33.75	4,685.2
	3	24.12	55.91	33.70	4,929.2
365	1	44.73	50.56	47.47	8,906.0
	2	53.41	49.91	51.60	8,993.2