

Use of Formal Ethical Reviews in NLP Literature: Historical Trends and Current Practices

Sebastin Santy[♣] Anku Rani[♡] Monojit Choudhury[♣]

[♣] Microsoft Research, Bangalore, India

[♡] Plaksha University, Mohali, India

[♣]{t-sesan, monojtc}@microsoft.com, [♡]anku.rani@plaksha.org

Abstract

Ethical aspects of research in language technologies have received much attention recently. It is a standard practice to get a study involving human subjects reviewed and approved by a professional ethics committee/board of the institution. How commonly do we see mention of ethical approvals in NLP research? What types of research or aspects of studies are usually subject to such reviews? With the rising concerns and discourse around the ethics of NLP, do we also observe a rise in formal ethical reviews of NLP studies? And, if so, would this imply that there is a heightened awareness of ethical issues that was previously lacking? We aim to address these questions by conducting a detailed quantitative and qualitative analysis of the ACL Anthology, as well as comparing the trends in our field to those of other related disciplines, such as cognitive science, machine learning, data mining, and systems.

1 Introduction

With the rapid advances in the field of Natural Language Processing (NLP), language technologies are getting woven into the daily fabric of our lives and the society. Since “language is a portal of emotions, a proxy of human behavior, and a strong signal of individual characteristics” (Hovy and Spruit, 2016), large-scale deployment of language technology has potential risks that require early detection and mitigation. Naturally, there have been several discussions about the potential harms and ethical issues concerning NLP (Hovy and Spruit, 2016; Conway and O’Connor, 2016). They have mostly revolved around building or deploying systems in sensitive areas such as hate speech (Sap et al., 2019), social media (Benton et al., 2017), clinical NLP and mental health (Šuster et al., 2017; Mikal et al., 2016) and use of sensitive or personal information (Larson, 2017). While building NLP systems, there are




Figure 1: Percentage(%) of papers mentioning IRB-related and *ethic* terms in NLP conferences, journals, and workshops from 2006 to 2020.

also ethical risks associated with involvement of human subjects through user studies or data collection activities (Shmueli et al., 2021).

The awareness of the dangers of the existing and new NLP applications has led to the curation of several ethical guidelines and frameworks. Undergirded by lessons from the past, these guidelines and frameworks help researchers consider and contextualize critical ethical concerns. Most of the ethical issues in NLP are rooted in the data being used for research. Couillault et al. (2014) is one of the first works to explore the ethics of data collection and evaluation in NLP. Several other works have proposed best practices for dealing with ethical implications of NLP research and deployment (Prabhumoye et al., 2019; Leidner and Plachouras, 2017; Bender and Friedman, 2018; Schnoebelen, 2017). There is now an increased awareness around this topic with a number of workshops and tutorials on ethics at NLP conferences (Tsvetkov et al., 2018; Hovy et al., 2017; Alfano et al., 2018). Such discussions have resulted in a number of reforms at NLP conferences. NLP conferences now have new track called *Ethics in NLP*. Furthermore, several ML and NLP conferences such as **NeurIPS 2020**,

NAACL 2021 and ACL 2021¹ now recommend the inclusion of broader impact statement in their papers, which allows for authors to introspect and be mindful of the ethical implications their research poses.

Although an NLP researcher might be individually committed towards ethical research practices, they may not have the expertise in ethical and legal issues to gauge the potential risks of a technology or a dataset that they are building. The practice of getting the research/study approved by an ethical review board (aka Institutional Review Board or IRB)² instituted by the organization is thus critical in early defusal of the potential harms. The two primary functions of IRB are (i) to protect the rights and welfare of human research subjects, and (ii) to support and facilitate the conduct of valuable research (Bankert and Amdur, 2006; Klitzman, 2012; Byerly, 2009). Traditionally, IRB has been a long-standing norm in biomedical research due to its overt exposure to human subjects. However, with computing research pervading human lives, IRBs have started covering computing research as well (Buchanan and Ess, 2009; Vitak et al., 2017). With regards to NLP, most of the data collection and annotation processes as well as user studies come under the purview of these boards. These are particularly necessary if they cover sensitive topics such as mental health or hate speech which can affect the human subjects involved in data collection or the users of the system.

How frequently do NLP researchers take IRB approvals for their studies? What aspects of NLP research or which topics of study are typically considered for IRB approvals? What are the historical and current trends, and what can we say about the awareness of the NLP research community around ethical issues? In this paper, we try to answer these questions through a quantitative and qualitative analysis of papers from the ACL anthology that seek IRB approvals. According to our findings, IRB approvals were almost non-existent in the NLP literature until 2006, but there has been a steady increase since 2016. We also study the distribution of IRB approvals by country and industry/academia affiliation, as well as compare the recent trends in NLP conferences to that of various prominent con-

ferences ranging from machine learning and data mining to human-computer interaction and systems. One of the key findings of this study is that IRB permission was mostly sought for either data collection or annotation studies, but hardly ever for data re-purposing or system design/deployment - a void that we think the NLP community should be conscious about.

2 Method

To determine the trends of IRB approvals in NLP research, we resort to searching for IRB- and ethics-related terms in research papers. We obtain the papers (PDFs) for major NLP conferences, journals and workshops [ACL, COLING, EAACL EMNLP, LREC, NAACL, CL, TAACL, and WS] from the ACL Anthology (curated by Joshi et al. (2020)). For a comparative analysis, we also collect papers from other related conferences [CogSci, InterSpeech, NeurIPS, CVPR, ICWSM, CHI, COMPASS] for the years 2019 and 2020, during which there was considerably more discussion around ethics of computing research.

In order to retrieve papers that seek IRB approvals, we search for the following keywords which cover the phrases used for IRB in countries that are frequently represented at these conferences: review board, ethics board, ethics panel, ethics committee, consent form, and IRB.³ To further compare and calibrate, we also search for papers that contain the wildcard string *ethic*, which brings up a broader set of papers that may discuss ethical repercussions of their work, even if any approval is not explicitly sought or mentioned. To assist with a robust search over this textual data, we use the Allen AI SPIKE interface⁴ (Taub Tabib et al., 2020; Shlain et al., 2020) and use `pdfgrep`⁵ to cross-check our results.

IRB-related term search yielded 210 papers from the ACL anthology (till 2019), which were then manually checked for precision and annotated for *aspects* (see Figure 3) and *topics* (e.g., hate speech, social media, mental health, etc.) of the research for which IRB permission was sought. Through our manual curation, we found that 94.17% of these papers actually took the approval for their research study thus showing that our search is precise in capturing the terms.

¹<https://2021.aclweb.org/ethics/Ethics-FAQ/>

²In this paper, we use IRB as a generic term to refer to such review boards which are known by slightly different terms in different geographies.

³Collectively referred to as IRB-related terms from hereon.

⁴URL when accessed: <https://spike.staging.apps.allenai.org/datasets/acl/search>

⁵Command-line tool: <https://pdfgrep.org/>

The remaining papers were mostly ethical frameworks and recommendations (e.g., [Hovy and Spruit \(2016\)](#); [Bender and Friedman \(2018\)](#)) which merely mentioned the need for seeking IRB approvals in NLP research.

3 Findings

3.1 How many papers seek IRB approvals?




Figure 2: The percentage (%) of papers in NLP conferences, journals and workshops over the past 15 years that mention IRB-related terms. The intensity of color is proportional to the % values. The boxes with a gray hatch reflect the years when that particular conference was not held.

Figure 2 shows the percentage(%) of papers in each NLP conference iteration that mention IRB-related terms. It is immediately obvious that for almost all the conferences only a minuscule number of papers mention IRB approvals. However, it is heartening to see that the number of mentions is increasing in recent years. **LREC** and **WS** particularly stand out among the other conferences for having at least some mentions of IRB approvals in every iteration. For **LREC**, it is understandable, since the theme of the conference revolves around data and resource generation. In the case of **WS**, IRB mentions are consistently increasing over the years. We observed that this is mostly due to the diverse nature of workshops some of which are on resource generation ([Popescu-Belis et al., 2019](#)) or cover sensitive topics ([Niederhoffer et al., 2019](#); [Yannakoudakis et al., 2019](#)). Journals such as **CL** and **TACL** have very few papers in each iteration, so even one IRB mention appears to be a lot. It should be noted that there is a possibility that a research study has obtained IRB approval but has not disclosed it in their paper. However, based on the authors’ experience (and anecdotes from personal conversations), it is highly unlikely that anyone who has been through the IRB approval process will fail to mention it.

3.2 What kinds of research seek IRB approvals?

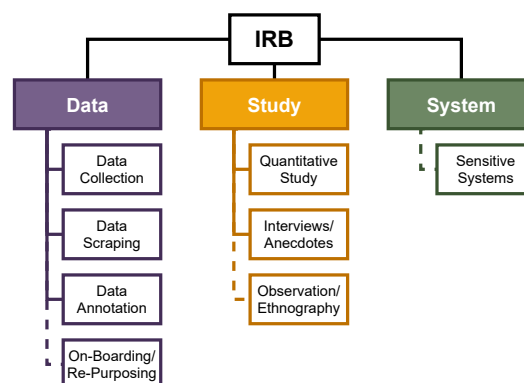


Figure 3: Different aspects of research for which IRB approvals were sought in the papers that we manually analyzed.

We manually go through each of the 58 NLP papers (which excludes **WS**) to derive the aspects and understand the context in which IRB approvals are sought and build a taxonomy of the broad topics covered (Figure 3). We see that most of them (24 papers; 41.3%) take IRB approvals for collection of data which can often involve human subjects directly. It is followed by the annotation of data with 20 papers (34%) taking IRB approvals. A meager 7 papers (12%) in our set take IRB approvals for scraping data which is the automatic collection of data from web pages or social media posts without explicit consent from the users. We see that only one paper takes IRB approval for re-purposing and further annotation of data ([Rogers et al., 2018](#)). One of the core concerns of GDPR is the usage of personal data collected by media platforms for a purpose different than what the user consented to and hence such re-purposing of data should ideally undergo IRB approvals. 12 papers (20%) take approvals for conducting user studies of both qualitative (survey, interview) and quantitative nature (semantic edits). Interestingly, we see that only one paper has taken IRB approval for the whole system owing to its sensitive nature ([Cao et al., 2018](#)).

We also look at the nature (or topics) of the work for which the IRB approvals are taken. We observe that 48.4% of papers that mention IRB have sensitive topics (such as mental health, hate speech, clinical/medical NLP), 20.3% of the papers are for collection of eye movement, EEG and audio/video recordings of human subjects, and rest of them are generic data collection or user study. To further understand the trends, we look at certain tracks of **ACL 2020** which deal with sensitive topics or

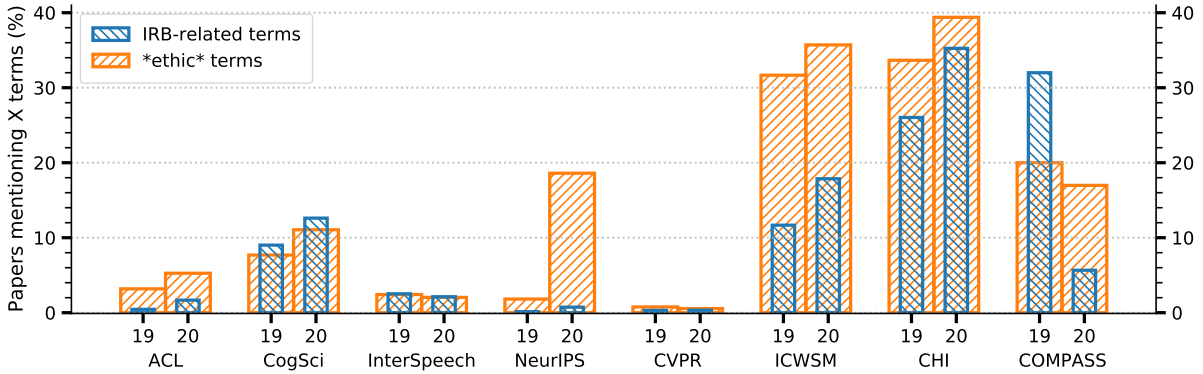


Figure 4: Percentage(%) of papers mentioning IRB-related and *ethic* terms in related conferences.

data collection. We notice that only 1/42 *Resource and Evaluation* papers and 3/24 *Computation Social Science* papers have taken IRB approvals. It is worth noting that 29/210 papers we have annotated only used an informed consent form without explicitly mentioning whether an ethical review board was involved in the process.

3.3 What is the distribution of IRB approvals by country and industry/academia?

	Total Papers	Percent Papers
Countries		
USA	47/5368	0.88%
Canada	5/358	1.40%
Germany	5/850	0.59%
UK	5/1088	0.46%
Netherlands	3/226	1.33%
Sweden	2/100	2.00%
South Korea	2/151	1.32%
China	2/2350	0.09%
Affiliation Types		
University	52/7730	0.67%
Industry	1/841	0.12%
National Lab	1/182	0.55%
Joint/Collaboration	11/2651	0.41%

Table 1: Distribution of % IRB-related term mentions among countries and different types of affiliations for NLP conferences (excluding **LREC** and **WS**) from 2012 to 2020. ⁶

Table 1 shows the distribution of papers which mention IRB approvals along two dimensions: countries and types of institutions. As can be observed, most of the listed countries are WEIRD⁷ societies. When it comes to the type of institution, we find that universities account for the vast majority

⁶Country and Affiliation data obtained from https://github.com/marekrei/ml_nlp_paper_data/

⁷Western, Educated, Industrialized, Rich and Democratic

of papers seeking IRB approvals, followed by joint collaborations. This trend can be counter-intuitive as an industry is more likely to be regulated and accountable for the ethical and legal concerns of their work. One possibility is that industries perhaps do not engage in external data collection/annotation work or conduct user studies as much as academic institutions do. Alternatively, it is possible that the data collection/annotation process is a completely independent pipeline that is not specific to the research paper in which it is used and thus is not reported.

3.4 How do the IRB trends in NLP research compare with those in related fields?

We look at the following conferences for our analysis: **CogSci** in cognitive science, **InterSpeech** in speech processing, **NeurIPS** in machine learning, **CVPR** in computer vision, **ICWSM** in social media mining, **CHI** in human-computer interaction, and **COMPASS** in computing systems deployment. We specifically analyze these conferences for 2019 and 2020 iterations as there have been significant changes made in the conferences during this period in terms of reporting the ethical ramifications of their research. Figure 4 shows the % of papers mentioning *ethic* and IRB-related terms for each conference iteration. We calculate for *ethic* to understand how aware and concerned each field/conference is towards the ethical implications of the research they conduct.

It is not surprising that IRB mentions for **CHI** are so high (~ 35%) given that more than 65% percent of CHI papers include at least one user study (Koe-man, 2020). **ICWSM** works with datasets and systems related to web and social media analytics and hence would need to undergo IRB approvals. This is apparent in the relatively high number of IRB

mentions in both 2019 and 2020. Unlike the other conferences we choose for our analysis, **CogSci** is a non-computing conference. Linguistics work is frequently found in **CogSci**, which often makes use of human subjects. We observe that it has the most consistent representation of both *ethic* and IRB-related term mentions among the years. As previously discussed, one of the concerns is that sensitive systems are seldom taking IRB approvals. On the contrary, we notice that **COMPASS**, a conference largely focused on deploying computing systems, is prevalent in taking IRB approvals. **InterSpeech** and **CVPR** have significantly fewer papers with IRB mentions ($< 0.35\%$ and $< 2.5\%$, respectively) and the trends have hardly changed over the years, despite the fact that they conduct research with speech, multimodal, and vision data that may have been collected from human subjects. Among these, there is a ray of hope for **ACL** which is showing a significant positive trend in both *ethic* and IRB mentions without any external reinforcement, thanks to the increasing awareness in the field. **NeurIPS**, on the other hand, has seen a meteoric rise in their *ethic* mentions, which, on manual inspection reveals, is due to their mandatory inclusion of broader impact statements. There has also been a slight increase in their IRB mentions, which could be attributed to this, indicating that broader impact statements might help researchers be more cautious when proposing their research to the larger community. This quantitative testimony from **NeurIPS** shows that **ACL** and other *CL conferences are moving in the right direction with their inclusion of stringent ethics reviews for their papers.

4 Way Forward

In this paper, we conduct a survey of IRB approvals in NLP research. The two key observations we make are as follows. First, very few papers ($< 0.8\%$ of all papers published) since 2006 have sought an IRB approval; though we do observe a rise in numbers ($< 1.3\%$ of all papers published) since 2016. This is much smaller compared to the numbers we observe for other conferences such as **CogSci**, **CHI**, **ICWSM** or **COMPASS**. Second, the majority of the IRB approvals were obtained for data collection or annotation that directly involved users, with only a few studies seeking approvals for data scraping or re-purposing. Such approvals are even more scarce for sensitive systems where

we seldom see any paper taking IRB approvals solely for the system. The number of papers creating new datasets is expected to be greater than 1% of all NLP papers⁸; the number of papers that re-purpose an existing dataset is expected to be even greater than this. Therefore, clearly not all papers creating datasets, and almost no paper re-purposing datasets take approvals from IRB. As such, re-purposing data collected from human subjects without their explicit consent on how the data will be re-used is potentially dangerous and may even have legal repercussions. Furthermore, with the exception of a couple of papers, to date, there is no practice or trend of taking IRB approval for designing, developing, and deploying systems. This is in stark contrast to the practice in other related fields/conferences such as **COMPASS**. Much of the harm caused by a system could actually come from its design or style of training or deployment, rather than the underlying datasets.

We see that the broad impact statements have helped conferences such as **NeurIPS** which were traditionally oblivious to ethical issues (Nanayakkara et al., 2021). We believe that, in a similar way, the impact statements introduced in **NAACL**⁹ and **ACL 2021**, with specific clauses for seeking IRB, will be highly beneficial in limiting the aforementioned potential risks by increasing the awareness amongst researchers of broader ethical repercussions of their research. It will be interesting to conduct a similar study a few years down the line and contrast with the findings of the current study.

Acknowledgements

We thank Prof. Emily M. Bender (UW) for providing useful insights and inspiration at the early stages of this work. We also thank Dr. Mary L. Gray and Ms. Anu Kannepalli (Microsoft) for their many insightful inputs, Prof. Yoav Goldberg (Allen AI) for setting up the SPIKE instance for us, and Prof. Katharina Reinecke (UW) and Dr. Prasanta Bhattacharya (A*STAR) for the broad discussions around IRBs. We also thank anonymous reviewers from ACL and NLP4PI workshop for their thoughtful suggestions and feedback.

⁸As a crude statistics, consider the fact that the number of accepted papers in the *Resource and Evaluation* track in **ACL 2020** was 5.4% , whereas only 1.7% of all the papers in that year sought for IRB approvals.

⁹Report: <https://2021.naacl.org/blog/ethics-review-process-report-back/>

References

- Mark Alfano, Dirk Hovy, Margaret Mitchell, and Michael Strube, editors. 2018. *Proceedings of the Second ACL Workshop on Ethics in Natural Language Processing*. Association for Computational Linguistics, New Orleans, Louisiana, USA.
- Elizabeth A Bankert and Robert J Amdur. 2006. *Institutional review board: Management and function*. Jones & Bartlett Learning.
- Emily M. Bender and Batya Friedman. 2018. [Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Adrian Benton, Glen Coppersmith, and Mark Dredze. 2017. [Ethical Research Protocols for Social Media Health Research](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 94–102, Valencia, Spain. Association for Computational Linguistics.
- Elizabeth A. Buchanan and Charles M. Ess. 2009. [Internet Research Ethics and the Institutional Review Board: Current Practices and Issues](#). *SIGCAS Comput. Soc.*, 39(3):43–49.
- Wesley G. Byerly. 2009. [Working with the institutional review board](#). *American Journal of Health-System Pharmacy*, 66(2):176–184.
- Xuân-Nga Cao, Cyrille Dakhliya, Patricia Del Carmen, Mohamed-Amine Jaouani, Malik Ould-Arbi, and Emmanuel Dupoux. 2018. [BabyCloud, a Technological Platform for Parents and Researchers](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Mike Conway and Daniel O’Connor. 2016. [Social media, big data, and mental health: current advances and ethical implications](#). *Current Opinion in Psychology*, 9:77–82.
- Alain Couillault, Karèn Fort, Gilles Adda, and Hugues de Mazancourt. 2014. [Evaluating corpora documentation with regards to the ethics and big data charter](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 4225–4229, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Dirk Hovy, Shannon Spruit, Margaret Mitchell, Emily M. Bender, Michael Strube, and Hanna Wallach, editors. 2017. *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*. Association for Computational Linguistics, Valencia, Spain.
- Dirk Hovy and Shannon L. Spruit. 2016. [The Social Impact of Natural Language Processing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598, Berlin, Germany. Association for Computational Linguistics.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The State and Fate of Linguistic Diversity and Inclusion in the NLP World](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Robert Klitzman. 2012. [Institutional Review Board Community Members: who are they, what do they do, and whom do they represent?](#) *Academic Medicine*, 87(7):975–981.
- Lisa Koeman. 2020. [HCI/UX research: what methods do we use?](#)
- Brian Larson. 2017. [Gender as a Variable in Natural-Language Processing: Ethical Considerations](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 1–11, Valencia, Spain. Association for Computational Linguistics.
- Jochen L. Leidner and Vassilis Plachouras. 2017. [Ethical by Design: Ethics Best Practices for Natural Language Processing](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 30–40, Valencia, Spain. Association for Computational Linguistics.
- Jude Mikal, Samantha Hurst, and Mike Conway. 2016. [Ethical issues in using twitter for population-level depression monitoring: a qualitative study](#). *BMC Medical Ethics*, 17(1).
- Priyanka Nanayakkara, Jessica Hullman, and Nicholas Diakopoulos. 2021. [Unpacking the Expressed Consequences of AI Research in Broader Impact Statements](#). In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’21, New York, NY, USA. Association for Computing Machinery.
- Kate Niederhoffer, Kristy Hollingshead, Philip Resnik, Rebecca Resnik, and Kate Loveys, editors. 2019. *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*. Association for Computational Linguistics, Minneapolis, Minnesota.
- Andrei Popescu-Belis, Sharid Loáiciga, Christian Hardmeier, and Deyi Xiong, editors. 2019. *Proceedings of the Fourth Workshop on Discourse in Machine Translation (DiscoMT 2019)*. Association for Computational Linguistics, Hong Kong, China.
- Shrimai Prabhumoye, Elijah Mayfield, and Alan W Black. 2019. [Principled Frameworks for Evaluating Ethics in NLP Systems](#). In *Proceedings of the 2019 Workshop on Widening NLP*, pages 118–121,

- Florence, Italy. Association for Computational Linguistics.
- Anna Rogers, Alexey Romanov, Anna Rumshisky, Svitlana Volkova, Mikhail Gronas, and Alex Gribov. 2018. [RuSentiment: An enriched sentiment analysis dataset for social media in Russian](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 755–763, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. [The Risk of Racial Bias in Hate Speech Detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence, Italy. Association for Computational Linguistics.
- Tyler Schnoebelen. 2017. [Goal-Oriented Design for Ethical Machine Learning and NLP](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 88–93, Valencia, Spain. Association for Computational Linguistics.
- Micah Shlain, Hillel Taub-Tabib, Shoval Sadde, and Yoav Goldberg. 2020. [Syntactic search by example](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 17–23, Online. Association for Computational Linguistics.
- Boaz Shmueli, Jan Fell, Soumya Ray, and Lun-Wei Ku. 2021. [Beyond Fair Pay: Ethical Implications of NLP Crowdsourcing](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3758–3769, Online. Association for Computational Linguistics.
- Simon Šuster, Stéphan Tulkens, and Walter Daelemans. 2017. [A Short Review of Ethical Challenges in Clinical Natural Language Processing](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 80–87, Valencia, Spain. Association for Computational Linguistics.
- Hillel Taub Tabib, Micah Shlain, Shoval Sadde, Dan Lahav, Matan Eyal, Yaara Cohen, and Yoav Goldberg. 2020. [Interactive extractive search over biomedical corpora](#). In *Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing*, pages 28–37, Online. Association for Computational Linguistics.
- Yulia Tsvetkov, Vinodkumar Prabhakaran, and Rob Voigt. 2018. [Socially Responsible NLP](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorial Abstracts*, pages 24–26, New Orleans, Louisiana. Association for Computational Linguistics.
- Jessica Vitak, Nicholas Proferes, Katie Shilton, and Zahra Ashktorab. 2017. [Ethics regulation in social computing research: Examining the role of institutional review boards](#). *Journal of Empirical Research on Human Research Ethics*, 12(5):372–382.
- Helen Yannakoudakis, Ekaterina Kochmar, Claudia Leacock, Nitin Madnani, Ildikó Pilán, and Torsten Zesch, editors. 2019. *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*. Association for Computational Linguistics, Florence, Italy.