

# Beyond Fair Pay: Ethical Implications of NLP Crowdsourcing

Boaz Shmueli<sup>1,2,3,\*</sup>, Jan Fell<sup>2</sup>, Soumya Ray<sup>2</sup>, and Lun-Wei Ku<sup>3</sup>

<sup>1</sup>Social Networks and Human-Centered Computing, TIGP, Academia Sinica

<sup>2</sup>Institute of Service Science, National Tsing Hua University

<sup>3</sup>Institute of Information Science, Academia Sinica

## Abstract

The use of crowdworkers in NLP research is growing rapidly, in tandem with the exponential increase in research production in machine learning and AI. Ethical discussion regarding the use of crowdworkers within the NLP research community is typically confined in scope to issues related to labor conditions such as fair pay. We draw attention to the lack of ethical considerations related to the various tasks performed by workers, including labeling, evaluation, and production. We find that the *Final Rule*, the common ethical framework used by researchers, did not anticipate the use of online crowdsourcing platforms for data collection, resulting in gaps between the spirit and practice of human-subjects ethics in NLP research. We enumerate common scenarios where crowdworkers performing NLP tasks are at risk of harm. We thus recommend that researchers evaluate these risks by considering the three ethical principles set up by the Belmont Report. We also clarify some common misconceptions regarding the Institutional Review Board (IRB) application. We hope this paper will serve to reopen the discussion within our community regarding the ethical use of crowdworkers.

## 1 Revisiting the Ethics of Crowdsourcing

The information age brought with it the internet, big data, smartphones, AI, and along with these, a plethora of complex ethical challenges. As a result, there is growing concern and discussion on ethics within the research community at large, including the NLP community. This is manifested in new ethics-focused workshops, ethics conference panels and relevant updates to peer review forms.

While ethics in NLP has multiple aspects, most recent attention focuses on pressing issues related to the societal impact of NLP. These include discrimination, exclusion, over-generalization, bias,

and fairness (Hovy and Spruit, 2016; Leidner and Plachouras, 2017). Other works are concerned with the ethical implications of NLP shared tasks (Parra Escartín et al., 2017), and introducing ethics into the NLP curriculum (Bender et al., 2020).

A substantial amount of NLP research now takes advantage of *crowdworkers* — workers on crowdsourcing platforms such as Amazon Mechanical Turk (known also as AMT or MTurk), Figure Eight<sup>1</sup>, Appen, Upwork, Prolific, Hybrid, Tencent Questionnaire, and Baidu Zhongbao, as well as internal crowdsourcing platforms in companies such as Microsoft and Apple. Workers are recruited to label, evaluate, and produce data. In the pre-internet era, such tasks (e.g. part-of-speech (POS) tagging) were done by hiring expert annotators or linguistics students. However, these are now mostly replaced by crowdworkers due to lower costs, convenience, speed, and scalability.

Overall, the general consensus in the literature is that as long as the pay to the crowdworkers is “fair” (minimum hourly wage or above), there are no further ethical concerns, and there is no need for approval by an Institutional Review Board<sup>2</sup> (with some exceptions). For example, Hovy and Spruit (2016) mention that “[w]ork on existing corpora is unlikely to raise any flags that would require an IRB approval”, with a footnote that there are “a few exceptions”. Fort et al. (2011) mention that only “[a] small number of universities have insisted on institutional review board approval for MTurk experiments”. As another example, NLP students are being taught that “paid labeling does not require IRB approval” since “[i]t’s not an experiment

<sup>1</sup>Previously CrowdFlower; acquired by Appen in 2019.

<sup>2</sup>Institutional Review Boards (IRBs) are university-level, multi-stakeholder committees that review the methods proposed for research involving human subjects to ensure that they conform to ethical principles. IRBs are also known by various other names, such as Research Ethics Boards (REBs) and Research Ethics Committees (RECs). Non-academic organizations may employ similar committees.

\*Corresponding author: [shmueli@iis.sinica.edu.tw](mailto:shmueli@iis.sinica.edu.tw)

Year	Accepted Papers				Papers Using Crowdsourcing	Payment Mentioned	IRB Mentioned
	ACL	EMNLP	NAACL	All			
2015	318	312	186	816	59 (7%)	4 (7%)	0
2016	328	264	182	774	82 (11%)	15 (18%)	0
2017	302	323	—	625	57 (9%)	12 (21%)	3
2018	381	549	332	1262	136 (11%)	17 (13%)	1
2019	660	683	423	1766	189 (11%)	32 (17%)	5
2020	779	754	—	1533	180 (12%)	42 (23%)	5
Total	2768	2885	1123	6776	703 (10%)	122 (17%)	14

Table 1: Papers using crowdsourced tasks in top NLP conferences, 2015–2020. The columns show, from left to right: conference year; number of accepted papers at ACL, EMNLP, and NAACL; total number of accepted papers; number of accepted papers using crowdsourced tasks (percentage of papers using crowdsourced tasks); number of papers using crowdsourced tasks that mention payment (percentage of papers using crowdsourced tasks that mention payment); number of papers using crowdsourced tasks that mention IRB review or exemption.

with human subjects” (Carnegie-Mellon University, Language Technologies Institute, 2020). Indeed, NLP papers that involve crowdsourced work rarely mention a review by an ethics board.

In this work, we wish to revisit the ethical issues of crowdsourcing in the NLP context, highlight several issues of concern, and suggest ways forward. From our survey of top NLP conferences, we find that crowdsourcing tasks are growing rapidly within the community. We therefore establish a common understanding of research ethics and how it relates to crowdsourcing. We demonstrate that the existing ethical framework is often inadequate and does not seek to protect crowdsourced workers. We then dispel common misunderstandings regarding the IRB process that NLP researchers might harbor. And finally, we outline how to apply ethical values as guidelines to minimize potential harms, and conclude with recommendations.

## 2 The Rise and Rise of NLP Crowdsourcing

To get a sense of the extent and growth of crowdsourced tasks within the NLP research community, we analyzed the proceedings of three top NLP conferences: ACL, EMNLP, and NAACL (also known as NAACL-HLT)<sup>3</sup>. We scanned the annual proceedings of these conferences in the six years from 2015 to 2020, looking for papers that mention direct employment of crowdsourced workers. All together,

<sup>3</sup>ACL is the Annual Meeting of the Association for Computational Linguistics, EMNLP is the Conference on Empirical Methods in Natural Language Processing, and NAACL is the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.

6776 papers were accepted for publication in these 16 conferences<sup>4</sup>. In total, we identified 703 papers that use crowdworkers as part of their research.<sup>5</sup> The results are summarized in Table 1.

**Remuneration** For each paper that uses crowdsourced labor, we checked whether the authors discuss payment or labor-related issues. Out of the 703 papers, 122 (17%) discuss payment, either by detailing the amount paid per task, the worker’s hourly wages, or declaring that wages were paid ethically. While in some cases researchers emphasize fair payment (e.g., Nangia et al. (2020) ensured “a pay rate of at least \$15/hour”), many other papers are more concerned about the cost of dataset acquisition, and thus mention the cost per task or the total dataset cost, but not the hourly compensation.

**IRB Review** Finally, we also checked whether authors mention a review by an IRB (or equivalent body) for their research. We found very few papers that mention an IRB approval or exemption — a total of 14 papers — which make up only 2% of the works that use crowdsourcing.

**Growth** We see that research papers using crowdsourced tasks have made up a relatively constant 11-12% of all research papers in the last three years. As research production grows exponentially, we expect a corresponding increase in the number of crowdsourced tasks.<sup>6</sup>

<sup>4</sup>NAACL was not held in 2017 and 2020; it is skipped every three years.

<sup>5</sup>MTurk is used in 80% of tasks.

<sup>6</sup>ACL, EMNLP, and NAACL are only three of many other publication venues in the NLP community.

## 2.1 Categorizing Crowdsourced Tasks

To understand the many nuanced ways in which researchers presently use crowdworkers in NLP tasks, we examined the tasks performed in each of the papers that use crowdsourcing. We found that NLP crowdsourced tasks generally fall into one of three categories, which we designate as *labeling*, *evaluation*, and *production*. We found that the categories account for 34%, 43%, and 23% of crowdsourcing tasks respectively. The results are summarized in Table 2, which also lists common action verbs that researchers use to describe the work performed by crowdworkers.

Task Category	Pct.	Action Verbs
Labeling	34%	annotate, label, identify, indicate, check
Evaluation	43%	rate, decide, score, judge, rank, choose
Production	23%	write, modify, produce, explain, paraphrase

Table 2: Tasks in the papers that use crowdsourcing, broken down by category. We also list typical action verbs used by the authors to describe the task.

**Labeling** entails the processing of existing data by the crowdworker and then the selection or composition of a label or labels for that data<sup>7</sup>. Labeling tasks augment the data with human-supplied labels. The augmented data are often used for training machine learning models. We further divide labeling tasks into two: objective and subjective labeling. In objective labeling, the desired label is factual, and does not depend on the worker. Classical examples are the tagging of sentences with named entities and part-of-speech (POS) labeling, and text transcription. In contrast, subjective labeling comprises of tasks where labels may depend on the worker’s personality, cultural background, opinion, or affective state. Examples include emotion labeling, detecting sarcasm in a tweet, or deciding whether a text constitutes hate speech or not.

In **evaluation** tasks, the worker is presented with data, for example a sentence, tweet, paragraph, or dialogue, and then requested to evaluate and score the data — mostly text — according to pre-defined criteria, such as fluency, coherence, originality, or structure. These tasks are often used by researchers to evaluate natural language generation (NLG) models. Similar to subjective labeling,

<sup>7</sup>We refrain from using the terms “annotation” and “annotators” as these terms are overloaded and often used for non-annotation work.

scores given by the workers may reflect the interpretation, values, or beliefs of the worker.

Finally, in **production** tasks, workers are asked to produce their own data, rather than label or evaluate existing data. In the NLP context, this often amounts to text elicitation or text generation. Examples include captioning a photo or video clip, writing a story given a sequence of images, or composing questions and answers. In this category we also include text translation. The produced data is often used for model training or evaluation.

While the majority of the studies use only one type of task — labeling, evaluation, or production — we found that in 10% of the papers that use crowdsourcing, researchers used two or more types of tasks in the same study. The combination of production and evaluation is particularly common; researchers often ask workers to generate data, which in turn is used to train a model; they then use workers to evaluate the model’s performance.

## 2.2 Surveys and Gamification

Although not common, some papers also collect personal information from workers. For example, [Yang et al. \(2015\)](#) and [Ding and Pan \(2016\)](#) conduct personality surveys among its crowdworkers. [Pérez-Rosas and Mihalcea \(2015\)](#) collect demographic data from the workers which included “their gender, age, country of origin, and education level”. Finally, we also found a few papers that add elements of gaming to their crowdsourced tasks, e.g. [Niculae and Danescu-Niculescu-Mizil \(2016\)](#) and [Urbanek et al. \(2019\)](#).

## 3 The Rules and Institutions of Research Ethics

Given the increasing use of crowdsourced NLP tasks, how can researchers ensure ethical concerns are reasonably addressed? Should a researcher make a judgement call and decide which tasks pose a risk of harm to the worker, and which are benign? To answer such questions, we will first explore the existing ethical framework used by researchers in the biomedical, social, and behavioral sciences.

### 3.1 The Genesis of Modern Research Ethics

The roots of contemporary research ethics originate in the 19th century, when researchers made unparalleled discoveries, but also engaged in hazardous, and frequently deadly, experimentation without great concern for the human subjects involved as

long as these trials advanced the boundaries of scientific knowledge (Ivy, 1948). The dominant ethics paradigm at the time was largely devoid of now-common principles surrounding therapeutic benefits, scientific validity, full knowledge, or subject consent (Lederer, 1995). Examples include researchers infecting intellectually disabled orphans with gonorrhea, or puncturing a healthy and unaware woman with the nodules of a leper patient to observe the clinical course of these diseases (Shamoo and Resnik, 2009).

Such incidents were common before the 1940s, and academia and public discourse were generally ignorant of them and of research ethics in general (Rothman, 1987). The revelation of the Nazi concentration camp experiments at the end of World War II was a watershed moment (Gillespie, 1989) and led to an early precursor of contemporary research ethics, namely the *Nuremberg Code* of 1947 (Jonsen, 1998). Not long after, the fallout of the Tuskegee Syphilis Study in the US prompted the formalization of research ethics at American universities (Caplan, 1992). In the study, which took place between 1932 and 1972, a total of 399 socio-economically disadvantaged African-American males with latent syphilis infections were recruited through the false promise of “free health-care”. Yet, these subjects were actually left without therapy even as effective treatment became available, with the objective to observe the clinical course of the disease (Brandt, 1978).

### 3.2 The Belmont Principles and IRBs

From thereon, beginning with biomedical research, the notion of research ethics has been institutionalized in the US at the university level through institutional review boards (IRBs), as well as national legislation (Israel, 2015). Gradually, research ethics have become a concern also in the social and behavioral sciences, as demonstrated by the 1978 *Belmont Report* created by the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research (Jonsen, 2005). This report became the basis of the 1991 Federal Policy for the Protection of Human Subjects in the United States, more commonly known as the *Common Rule* (Owen, 2006) and superseded by the *Final Rule* (2018). Chiefly, the *Final Rule* aims to ensure that the following three basic principles listed in the Belmont Report (National Commission for the Protection of Human Subjects of Biomedical and

*Behavioral Research*, 1978) are met:

1. **Respect for persons**, which includes “the requirement to acknowledge autonomy and the requirement to protect those with diminished autonomy”;
2. **Beneficence**, which mandates considering whether the benefits resulting from the research can outweigh the risks; and
3. **Justice**, requiring that the burden — and benefits — of the research are equally shared among potential subjects.

The *Final Rule* is codified in Title 45, *Code of Federal Regulations*, Part 46 and applies to all government-funded research in the United States (Israel, 2015). Virtually all universities in the United States apply this regulation to human subjects research projects irrespective of funding source (Klitzman, 2015). Specifically, the *Final Rule* requires that most *research involving human subjects* receives approval from an IRB. The IRB is a special university-level committee that reviews research proposals to verify that they comply with ethical standards. While it is difficult to assess the effectiveness of IRBs, and the process of ethical review is sometimes criticized as overly bureaucratic which may hamper low-risk social science (Resnik, 2018; Schrag, 2010), glaring ethical lapses have been rare after 1974 (Klitzman, 2015).

### 3.3 Are IRBs Universal?

The three ethical principles outlined by the Belmont Report — Respect for persons, Beneficence, and Justice — are also the stated principles guiding the actions of more recently formed ethics boards around the world, as well as the underlying principles of relevant policies of intergovernmental organizations, including the Universal Declaration on Bioethics and Human Rights (UNESCO, 2006) and the International Ethical Guidelines for Biomedical Research Involving Human Subjects (Council for International Organizations of Medical Sciences, 2002). Consequently, many countries worldwide have modeled national regulations after the *Final Rule* or its predecessor, the *Common Rule* (Capron, 2008). Both regulations have also influenced editorial policies of academic journals (e.g., the *Committee on Publication Ethics*, 2020). The *Final Rule* can thus be considered a cross-disciplinary de-facto standard for research ethics in Western/US-



influenced academic settings (Gontcharov, 2018), including India, Japan, Korea, and Taiwan.

Though we find that a global agreement on human-subjects ethics is emerging, countries still vary in the extent to which relevant policies are accepted, framed, implemented, or enforced.

#### 4 Do NLP Tasks Constitute Research Involving Human Subjects?

These formal rules and institutions have been established to protect the rights and interests of humans subjects involved as participants in scientific research. Thus, we must determine whether the rules and institutions of research ethics are even applicable to crowdsourced studies. At the core of this determination are two fundamental questions:

1. Are crowdsourcing tasks *research*?
2. Are crowdworkers *human subjects*?

In the following, we address these two questions.

##### 4.1 Are Crowdsourcing Tasks Research?

The *Final Rule* defines *research* as follows:

**Research** means a systematic investigation, including research development, testing, and evaluation, designed to develop or contribute to generalizable knowledge. Activities that meet this definition constitute research for purposes of this policy, whether or not they are conducted or supported under a program that is considered research for other purposes. (45 *CFR* 46.102(l), *Final Rule*, 2018)

From this definition it is evident that rather than a concrete research behavior on part of the NLP researcher, it is the *purpose* of the research behavior that classifies said behavior as *research* under the *Final Rule*. In other words, all categories of crowdsourced tasks summarized in Section 2, i.e., labeling, evaluation, and production, may be considered part of *research* so long as the intended outcome is to create generalizable knowledge. Typically, this encompasses academic settings where research behavior takes place (course assignments by students being a prominent exception), but does not include research conducted in industry settings (Meyer, 2020; Jackman and Kanerva, 2015).

##### 4.2 Are Crowdworkers Human Subjects?

The *Final Rule* defines *human subjects* as follows:

**Human subject** means a living individual *about whom* an investigator (whether professional or student) conducting research:

- (i) Obtains information or biospecimens through intervention or interaction with the individual, and, uses, studies, or analyzes the information or biospecimens; or
- (ii) Obtains, uses, studies, analyzes, or generates identifiable private information or identifiable biospecimens.

(45 *CFR* 46.102(e)(1), *Final Rule*, 2018)

Clearly, if the researcher obtains identifiable private information (IPI) as part of the crowdsourced task — e.g. name, date of birth, email address, national identity number, or any other information that identifies the worker — then (ii) holds and the worker is considered a human subject.

Even if the researcher does not obtain any IPI, the worker may still be considered a human subject under (i) in certain cases. This is because NLP researchers *interact* with crowdworkers through MTurk or analogous platforms when they publish the task, and obtain information through this interaction. It is also evident that academic NLP researchers make use of this information as they conduct a given study, and hence it is “used, studied, or analyzed”. If the information is *about the crowdworker* then they are considered human subjects and (i) is met. However, *Final Rule* does not expand on what constitutes *information about the individual*. According to University of Washington, Office of Research (2020), for example, *about whom* means that the “data or information relates to the person. Asking what [crowdworkers] think about something, how they do something, or similar questions usually pertain to the individuals. This is in contrast to questions about factual information not related to the person.”

Whether the information obtained in an NLP task is about the worker can initially seem like an easy-to-answer question. For example, Benton et al. (2017) write:

[R]esearch that requires the annotation of corpora for training models involves hu-

man annotators. But since the research does not study the actions of those annotators, the research does not involve human subjects. By contrast, if the goal of the research was to study how humans annotate data, such as to learn about how humans interpret language, then the research may constitute human subjects research.

However, we believe that this is not so clear-cut. First, one might argue that although in a POS labeling task we do not obtain information about the worker, other labeling tasks might be harder to classify. For example, when researchers ask a worker to compose a story given a sequence of photos, do they obtain information about the worker? And if so, what kind of information? Similar questions can be asked about tasks related to emotion classification (which might reveal a worker’s personality or mood), composing questions and answers (which point to areas of interest and cultural background), or identifying hate speech (which can indicate political orientation).

Second, platforms might *automatically* provide information that can be considered by some to be about the individual. Even in the most benign and “objective” tasks such as POS tagging, MTurk supplies researchers with information on the amount of time taken to complete each task. This information is sometimes collected and used by NLP researchers (e.g., [Sen et al., 2020](#)).

In summary, we have shown that in an academic context, NLP crowdsourcing tasks are *research*, but that the categorization of crowdworkers as *human subjects* can, in some cases, be a gray area that is open to interpretation. The *Final Rule* was designed to address ethical issues in medical research, and later in behavioral sciences; lawmakers and experts involved did not anticipate its use in new domains such as crowdsourcing. Therefore, its application to online data collection, and crowdsourcing in particular, can be ambiguous and unsatisfactory.<sup>8</sup> Thus, while in some cases the protections and procedures mandated under the *Final Rule* apply, in others they might not. As a consequence, some NLP crowdsourcing tasks may not require an IRB application, and this may happen even if crowdworkers are at risk.

---

<sup>8</sup>Some universities employ the *Precautionary Principle* and require all crowdsourced-enabled research to go through an IRB application.

## 5 Dispelling IRB Misconceptions

As we now see, if workers constitute human subjects, an IRB application is required. To clarify any other misconceptions that might be present within the community, we list key points related to the IRB process and dispel misconceptions around them. While not exhaustive, the list can serve both researchers and reviewers.

### 5.1 Researchers Cannot Exempt Themselves from an IRB Application

The *Final Rule* includes provisions for IRB exemptions, and we expect the vast majority of crowdsourced NLP tasks to fall into that category. However, it is crucial to understand that granting a research project an IRB exemption is not the prerogative of the researcher; it is only IRB that hands out exemptions following an initial review:

[T]he determination of exempt status (and the type of review that applies) rests with the IRB or with an administration official named by your institution. *The determination does not rest with the investigator.* Therefore, all projects must be submitted to the IRB for initial review. ([American Association for Public Opinion Research, 2020](#))

### 5.2 Worker IDs Constitute IPI

Researchers often obtain and store worker IDs — unique and persistent identifiers assigned to each worker by the crowdsourcing platform — even when workers are “anonymous”. MTurk, for example, assigns each worker a fixed 14-digit string, which is provided to the researcher with completed tasks. A worker ID is part of the worker’s account, and is therefore linked to their personal details, including full name, email address, and bank account number. As a consequence, the worker ID constitutes IPI (identifiable private information). If the worker ID is obtained by the researcher, the research mandates an initial IRB review.

To avoid obtaining this IPI, researchers can create and store pseudonymized worker IDs, provided that these IDs cannot be mapped back to the original worker IDs.

### 5.3 Obtaining Anonymous Data Does Not Automatically Absolve from IRB Review

Even if IPI is not obtained, and only *anonymous* and non-identifiable data is collected, we have

shown in Section 4.2 that crowdsourced NLP tasks involve interaction between researcher and participant through which the data about the worker may be collected, and thus often require an initial review by an IRB.

#### 5.4 Payment to Crowdworkers Does Not Exempt Researchers from IRB Review

Remuneration of human subjects does not change their status to independent contractors beyond the scope of research ethics. In fact, compensation of human subjects for the time and inconvenience involved in participating is a standard practice “especially for research that poses little or no direct benefit for the subject” and at the same time “should not constitute undue inducement to participate”, as the [University of Virginia, Human Research Protection Program \(2020\)](#) points out.

#### 5.5 Non-Published Research Also Requires IRB Review

Some researchers believe that research that will not be published is not subject to an IRB review. For example, [Carnegie-Mellon University, Language Technologies Institute \(2020\)](#) teaches students that “Paid labeling does not require IRB approval... [b]ut sometimes you want to discuss results in papers, so consider IRB approval”.<sup>9</sup>

The definition of *research* in the *Final Rule* is not contingent on subsequent publication of the results. Given the uncertainties of the peer-review process, whether research eventually finds its way into a publication can often be ascertained only *ex post*. An important exception not requiring IRB approval is student work as part of course assignments. However, subsequent use of research data collected originally as part of a student assignment is not mentioned in the *Final Rule*. Consequently, universities handle this situation differently. For example, University of Michigan allows retroactive IRB approval:

Class assignments may become subject to this policy... if the faculty member or the students change their plans... application to the IRB for permission to use the data is required. ([University of Michigan, Research Ethics and Compliance, 2021](#))

while Winthrop University does not:

<sup>9</sup>Moreover, whether the labeling is paid or unpaid is irrelevant; see Section 5.4.

IRB approval cannot be granted retroactively, and data may need to be re-collected for the project. ([University of Winthrop, The Office of Grants and Sponsored Research Development, 2021](#))

## 6 Risks and Harms for Crowdworkers

Previous work on the ethics of crowdsourcing focused on labor conditions, such as fair pay, and on privacy issues ([Fort et al., 2011](#); [Gray and Suri, 2019](#)). However, even when payment is adequate and the privacy of the workers is preserved, there are additional ethical considerations that are often not taken into account, and might put the worker at risk. We propose using the three ethical principles outlined by the Belmont Report — Respect for Persons, Beneficence, and Justice — to guide the action of researchers. We outline some of the specific risks and harms that might befall NLP task crowdworkers in light of these principles. While the list is not comprehensive, it can serve as a starting point to be used by researchers when planning their crowdsourced task, as well as by reviewers examining the ethical implications of a manuscript or research proposal.

### 6.1 Inducing Psychological Harms

NLP researchers are increasingly cognizant that texts can potentially harm readers, as evident by *trigger warnings* they add to their own papers (e.g., [Sap et al., 2020](#); [Nangia et al., 2020](#); [Sharma et al., 2020](#); [Han and Tsvetkov, 2020](#)). Moreover, researchers have long known that annotation work may be psychologically harmful. The Linguistic Data Consortium (LDC), for example, arranged stress-relieving activities for annotators of broadcast news data, following reports of “negative psychological impact” such as intense irritation, overwhelmed feelings, and task-related nightmares ([Strassel et al., 2000](#)). Although literature on the emotional toll on crowdworkers is still scant ([Huws, 2015](#)), there is growing literature on the psychological cost of work done by commercial content moderators ([Steiger et al., 2021](#)). Crowdworkers deserve similar consideration: while NLP tasks can be as benign as the POS tagging of a children’s poem, they may also involve exposure to disturbing textual or visual content.

In 45 *CFR* 46.110 the *Final Rule* allows expedited IRB review for research posing *no more than*

*minimal risk* to the human subjects involved. According to 45 *CFR* 46.102(j) *minimal risk* “means that the probability and magnitude of harm or discomfort anticipated in the research are not greater in and of themselves than those ordinarily encountered in daily life or during the performance of routine [...] psychological examinations or tests.” Exposing crowdworkers to sensitive content may exceed the threshold of *minimal risk* if the data that requires labeling or evaluation might be psychologically harmful.

The amount of harm (if any) depends on the sensitivity of the worker to the specific type of content. The risk is generally higher in data labeling or text evaluation tasks, where workers might be repeatedly asked to categorize offensive tweets, transcribe violent texts, evaluate hateful text that may expose them to emotional stimuli, or, depending on the content and worker, traumatize them. In some cases, sexual material can be highly offending or shocking and cause an emotional disturbance. Although less likely, hazards can also occur when workers are asked to produce text, since workers are elicited to produce texts based on given input. For example, when users are asked to compose a story based on images, certain images might trigger a harmful response in the worker.

## 6.2 Exposing Sensitive Information of Workers

A crowdworker might inadvertently or subconsciously expose sensitive information about themselves to the researcher. This is more pronounced in text production, where the responses produced by such tasks reveal as much about the individual workers as they do about the produced text. However, workers also reveal information about themselves when evaluating or labeling text, especially when subjective labeling is in place. Moreover, even seemingly trivial data — for example, the elapsed time taken by the worker to label, evaluate, or produce text — may contain valuable information about the worker (and this information is automatically captured by MTurk). Table 3 shows the risk level for the different task categories.

Moreover, researchers can obtain sensitive information about workers because the crowdsourcing platforms allow screening of workers using built-in qualification attributes including age, financial situation, physical fitness, gender, employment status, purchasing habits, political affiliation, handedness,

Task Category	Risk of Exposure of Workers’ Sensitive Information
Objective Labeling	Low
Subjective Labeling	Medium
Evaluation	Medium
Production	High

Table 3: Potential risk level by task category.

marital status, and education level (Amazon Mechanical Turk, 2020).

Researchers may also obtain other types of sensitive information by creating their own, arbitrary qualification tests as a quality control measure (Daniel et al., 2018).

In summary, the information obtained by researchers may reveal as much about the individual crowdworkers as they do about the data being labeled, evaluated, or produced.

## 6.3 Unwittingly Including Vulnerable Populations

Research on crowdsourcing platforms such as MTurk is inherently prone to the inclusion of *vulnerable* populations. In 45 *CFR* 46.111(b) the *Final Rule* (2018) non-exhaustively lists “children, prisoners, individuals with impaired decision-making capacity, or economically or educationally disadvantaged persons” as vulnerable groups. The *Final Rule* requires that additional safeguards are implemented to protect these vulnerable populations and makes IRB approval contingent on these safeguards.

A great proportion of MTurk crowdworkers are located in developing countries, such as India or Bangladesh, which makes MTurk an attractive proposition to those offering a task (Gray and Suri, 2019), but increases the risk of including economically-disadvantaged persons. Furthermore, it is difficult to ensure that MTurk crowdworkers are above the age of majority, or fall into any other of the defined vulnerable populations (Mason and Suri, 2011).

Moreover, given the power imbalances between researchers in industrialized countries and crowdworkers in developing countries, ethical consideration should occur regardless of whether the jurisdiction in which the crowdworkers are located even has a legal framework of research ethics in place, or whether such a local framework meets the standard of the Belmont Report or *Final Rule*.



## 6.4 Breaching Anonymity and Privacy

There is a perception among researchers that crowdworkers are anonymous and thus the issue of privacy is not a concern. This is not the case. [Lease et al. \(2013\)](#), for example, discuss a vulnerability that can expose the identity of an Amazon Mechanical Turk worker using their worker ID — a string of 14 letters and digits — because the same worker ID is used also as the identifier of the crowdworker’s account on other Amazon assets and properties. As a result, a Google search for the worker ID can lead to personal information such as product reviews written by the crowdworker on Amazon.com, which in turn can disclose the worker’s identity. Researchers might be unaware of these issues when they make worker IDs publicly available in papers or in datasets. For example, [Gao et al. \(2015\)](#) rank their crowdworkers using MTurk worker IDs in one of the figures.

Moreover, breaches of privacy can also occur unintentionally. For example, workers on MTurk are provided with an option to contact the researcher. In this case, their email address will be sent to the researcher, who is inadvertently exposed to further identifiable private information (IPI). We maintain that the anonymity of crowdworkers cannot be automatically assumed or guaranteed, as this is not a premise of the crowdsourcing platform.

## 6.5 Triggering Addictive Behaviour

[Graber and Graber \(2013\)](#) identify another source for harmful effects, which ties in with the risk of psychological harm, and is specific to *gamified* crowdsourced tasks: a possibility of addiction caused by dopamine release following a reward given during the gamified task. Gamification techniques can be added to data labeling, evaluation, and production. Indeed, some NLP work is using gamification, mostly for data collection (e.g., [Kumaran et al., 2014](#); [Ogawa et al., 2020](#); [Öhman et al., 2018](#)).

Moreover, the crowdsourcing platform may add elements of gamification over which the researcher has no control. For example, MTurk recently introduced a “Daily Goals Dashboard”, where the worker can set game-like “HITs Goal” and “Reward Goal”, as shown in Figure 1.




Figure 1: The “Daily Goals Dashboard” on MTurk

## 7 Ways Forward

The use of crowdworkers is growing within the NLP community, but the ethical framework set in place to guarantee their ethical treatment (whether *de jure* or *de facto*) did not anticipate the emergence of crowdsourcing platforms. In most crowdsourced NLP tasks, researchers do not intend to gather information about the worker. However, the crowdsourcing platform often autonomously collects such information. As a result, it is often difficult to determine whether crowdworkers constitute human subjects — which hinges on whether the researcher collects information about the worker. However, a determination that all crowdworkers are human subjects — and thus mandate an IRB approval for government-supported institutions — might create a “chilling effect” and disadvantage university researchers compared to industry-affiliated researchers. The effect is exacerbated in institutions where the ethics committee is heavily bureaucratic and does not offer a streamlined, expedited exemption process for low-to-no risk studies.

Whether their crowdsourced task requires IRB application or not, we recommend that the ethics-aware researcher should carefully examine their study in light of the three principles set up by the Belmont Report: Respect for persons, Beneficence, and Justice. And while this mandates fair pay, it is important to note that this is just one of the implications. There are other ethical considerations that are often overlooked — in particular risk assessment of causing psychological harm and exposure of sensitive information. Thus, we recommend increasing awareness of the potential ethical implications of crowdsourced NLP tasks. As NLP researchers are now encouraged to add an “ethical considerations” section to their papers ([NAACL, 2020](#)), they should also be encouraged to carefully weigh potential benefits against risks related to the crowdsourced task.

We also propose increasing awareness by disseminating relevant knowledge and information. An educational ethics resource created using a community effort could serve as a beneficial first step. Such a resource can include guidelines, checklists, and case studies that are specific to the ethical challenges of crowdsourced tasks in the context of NLP research. We believe that the creation of such a resource can serve as a springboard for a necessary nuanced conversation regarding the ethical use of crowdworkers in the NLP community.

## Acknowledgements

We thank the anonymous reviewers for their valuable comments and suggestions which helped improve the paper. This research was partially supported by the Ministry of Science and Technology in Taiwan under grants MOST 108-2221-E-001-012-MY3 and MOST 109-2221-E-001-015- [sic].

## References

- Amazon Mechanical Turk. 2020. Amazon Mechanical Turk Pricing. <https://web.archive.org/web/20201126061244/https://requester.mturk.com/pricing>. Accessed: 2020-11-26.
- American Association for Public Opinion Research. 2020. IRB FAQs for Survey Researchers. <https://www.aapor.org/Standards-Ethics/Institutional-Review-Boards/IRB-FAQs-for-Survey-Researchers.aspx>. Accessed: 2020-11-11.
- Emily M. Bender, Dirk Hovy, and Alexandra Schofield. 2020. Integrating ethics into the NLP curriculum. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 6–9, Online. Association for Computational Linguistics.
- 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.
- Allan M. Brandt. 1978. Racism and research: The case of the Tuskegee syphilis study. *The Hastings Center Report*, 8(6):21–29.
- Arthur L. Caplan. 1992. When evil intrudes. *The Hastings Center Report*, 22(6):29–32.
- Alexander M. Capron. 2008. Legal and regulatory standards of informed consent in research. In Ezekiel J. Emanuel, Christine Grady, Robert A. Crouch, Reidar K. Lie, Franklin G. Miller, and David Wendler, editors, *The Oxford Textbook of Clinical Research Ethics*, pages 613–632. Oxford University Press, New York, NY.
- Carnegie-Mellon University, Language Technologies Institute. 2020. Human subjects research determination worksheet. <https://web.archive.org/web/20201111132614/http://demo.clab.cs.cmu.edu/11737fa20/slides/multilingual-21-annotation.pdf>. Accessed: 2020-11-11.
- Committee on Publication Ethics. 2020. <https://publicationethics.org>. Accessed: 2020-11-22.
- Council for International Organizations of Medical Sciences. 2002. International Ethical Guidelines for Biomedical Research Involving Human Subjects. [https://cioms.ch/wp-content/uploads/2016/08/International\\_Ethical\\_Guidelines\\_for\\_Biomedical\\_Research\\_Involving\\_Human\\_Subjects.pdf](https://cioms.ch/wp-content/uploads/2016/08/International_Ethical_Guidelines_for_Biomedical_Research_Involving_Human_Subjects.pdf). Accessed: 2021-04-06.
- Florian Daniel, Pavel Kucherbaev, Cinzia Cappiello, Boualem Benatallah, and Mohammad Allahbakhsh. 2018. Quality control in crowdsourcing: A survey of quality attributes, assessment techniques, and assurance actions. *ACM Comput. Surv.*, 51(1).
- Tao Ding and Shimei Pan. 2016. Personalized emphasis framing for persuasive message generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1432–1441, Austin, Texas. Association for Computational Linguistics.
- Final Rule. 2018. Protection of Human Subjects, Part 46. In *US Department of Health and Human Services, Code of Federal Regulations (CFR) Title 45*.
- Karën Fort, Gilles Adda, and K Bretonnel Cohen. 2011. Amazon mechanical turk: Gold mine or coal mine? *Computational Linguistics*, 37(2):413–420.
- Mingkun Gao, Wei Xu, and Chris Callison-Burch. 2015. Cost optimization in crowdsourcing translation: Low cost translations made even cheaper. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 705–713, Denver, Colorado. Association for Computational Linguistics.
- Richard Gillespie. 1989. Research on human subjects: An historical overview. *Bioethics News*, 8(2):s4–s15.
- Igor Gontcharov. 2018. Qualitative ethics in a positivist frame: The Canadian experience 1998–2014. In *The SAGE Handbook of Qualitative Research Ethics*, pages 231–246. SAGE Publications Ltd.
- Mark A Graber and Abraham Graber. 2013. Internet-based crowdsourcing and research ethics: the case for IRB review. *Journal of medical ethics*, 39(2):115–118.
- Mary L. Gray and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt, Boston, MA.
- Xiaochuang Han and Yulia Tsvetkov. 2020. Fortifying toxic speech detectors against veiled toxicity. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7732–7739, Online. Association for Computational Linguistics.

- 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.
- Ursula Huws. 2015. *Online labour exchanges or 'crowdsourcing': Implications for occupational safety and health*. European Agency for Safety and Health at Work (EU-OSHA).
- Mark Israel. 2015. *Research Ethics and Integrity for Social Scientists: Beyond Regulatory Compliance*, 2nd edition. SAGE Publications, Thousand Oaks, CA.
- Andrew Conway Ivy. 1948. [The history and ethics of the use of human subjects in medical experiments](#). *Science*, 108(2792):1–5.
- Molly Jackman and Lauri Kanerva. 2015. Evolving the IRB: Building robust review for industry research. *Wash. & Lee L. Rev. Online*, 72:445.
- Albert R. Jonsen. 1998. The ethics of research with human subjects: A short history. In Albert R. Jonsen, Robert M. Veatch, and LeRoy Walters, editors, *Source Book in Bioethics: A Documentary History*, pages 3–10. Georgetown University Press, Washington, DC.
- Albert R. Jonsen. 2005. On the origins and future of the Belmont report. In James F. Childress, Eric M. Meslin, and Harold T. Shapiro, editors, *Belmont Revisited: Ethical Principles for Research with Human Subjects*, pages 3–11. Georgetown University Press, Washington, DC.
- Robert L. Klitzman. 2015. *The Ethics Police?: The Struggle to Make Human Research Safe*. Oxford University Press, Oxford.
- A Kumaran, Melissa Densmore, and Shaishav Kumar. 2014. Online gaming for crowd-sourcing phrase-equivalents. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1238–1247.
- Matthew Lease, Jessica Hullman, Jeffrey Bigham, Michael Bernstein, Juho Kim, Walter Lasecki, Saeid Bakhshi, Tanushree Mitra, and Robert Miller. 2013. Mechanical Turk is not anonymous. *Available at SSRN 2228728*.
- Susan E. Lederer. 1995. *Subjected to Science: Human Experimentation in America Before the Second World War*. Johns Hopkins University Press, Baltimore, MD.
- 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.
- Winter Mason and Siddharth Suri. 2011. [Conducting behavioral research on Amazon's Mechanical Turk](#). *Behavior Research Methods*, 44(1):1–23.
- Michelle N. Meyer. 2020. [There oughta be a law: When does\(n't\) the U.S. Common Rule apply?](#) *The Journal of Law, Medicine & Ethics*, 48(1\_suppl):60–73. PMID: 32342740.
- NAACL. 2020. Ethics FAQ. <https://web.archive.org/web/20201129202126/https://2021.naacl.org/ethics/faq/>. Accessed: 2020-11-29.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. [CrowS-pairs: A challenge dataset for measuring social biases in masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. 1978. *The Belmont report: ethical principles and guidelines for the protection of human subjects of research*. The Commission, Bethesda, MD.
- Vlad Niculae and Cristian Danescu-Niculescu-Mizil. 2016. [Conversational markers of constructive discussions](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 568–578, San Diego, California. Association for Computational Linguistics.
- Haruna Ogawa, Hitoshi Nishikawa, Takenobu Tokunaga, and Hikaru Yokono. 2020. [Gamification platform for collecting task-oriented dialogue data](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 7084–7093, Marseille, France. European Language Resources Association.
- Emily Öhman, Kaisla Kajava, Jörg Tiedemann, and Timo Honkela. 2018. [Creating a dataset for multilingual fine-grained emotion-detection using gamification-based annotation](#). In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 24–30, Brussels, Belgium. Association for Computational Linguistics.
- Michael Owen. 2006. Ethical review of social and behavioral science research. In Elliott C. Kulakowski and Lynne U. Chronister, editors, *Research Administration and Management*, pages 543–556. Jones and Bartlett Publishers, Sudbury, MA.
- Carla Parra Escartín, Wessel Reijers, Teresa Lynn, Joss Moorkens, Andy Way, and Chao-Hong Liu. 2017. [Ethical considerations in NLP shared tasks](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 66–73, Valencia, Spain. Association for Computational Linguistics.



- Verónica Pérez-Rosas and Rada Mihalcea. 2015. [Experiments in open domain deception detection](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1120–1125, Lisbon, Portugal. Association for Computational Linguistics.
- David B Resnik. 2018. *The ethics of research with human subjects: Protecting people, advancing science, promoting trust*, volume 74. Springer, Cham.
- David J. Rothman. 1987. [Ethics and human experimentation](#). *New England Journal of Medicine*, 317(19):1195–1199.
- Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A. Smith, and Yejin Choi. 2020. [Social bias frames: Reasoning about social and power implications of language](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5477–5490, Online. Association for Computational Linguistics.
- Zachary M Schrag. 2010. *Ethical imperialism: Institutional review boards and the social sciences, 1965–2009*. The Johns Hopkins University Press, Baltimore, MD.
- Cansu Sen, Thomas Hartvigsen, Biao Yin, Xiangnan Kong, and Elke Rundensteiner. 2020. [Human attention maps for text classification: Do humans and neural networks focus on the same words?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4596–4608, Online. Association for Computational Linguistics.
- Adil E. Shamoo and David E. Resnik. 2009. *Responsible Conduct of Research*, 2 edition. Oxford University Press, New York, NY.
- Ashish Sharma, Adam Miner, David Atkins, and Tim Althoff. 2020. [A computational approach to understanding empathy expressed in text-based mental health support](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5263–5276, Online. Association for Computational Linguistics.
- Miriah Steiger, Timir J Bharucha, Sukrit Venkatagiri, Martin J Riedl, and Matthew Lease. 2021. [The psychological well-being of content moderators](#). In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, CHI '21*, New York, NY, USA. Association for Computing Machinery.
- Stephanie Strassel, David Graff, Nii Martey, and Christopher Cieri. 2000. [Quality control in large annotation projects involving multiple judges: The case of the TDT corpora](#). In *Proceedings of the Second International Conference on Language Resources and Evaluation (LREC'00)*, Athens, Greece. European Language Resources Association (ELRA).
- UNESCO. 2006. Universal Declaration on Bioethics and Human Rights. <https://unesdoc.unesco.org/ark:/48223/pf0000146180>. Accessed: 2021-04-06.
- University of Michigan, Research Ethics and Compliance. 2021. Class assignments & IRB approval. <https://web.archive.org/web/20210325081932/https://research-compliance.umich.edu/human-subjects/human-research-protection-program-hrpp/hrpp-policies/class-assignments-irb-approval>. Accessed: 2021-03-25.
- University of Virginia, Human Research Protection Program. 2020. Subject compensation. <https://web.archive.org/web/20200611193644/https://research.virginia.edu/irb-hsr/subject-compensation>. Accessed: 2020-11-19.
- University of Washington, Office of Research. 2020. Human subjects research determination worksheet. [https://web.archive.org/web/20201111071801/https://www.washington.edu/research/wp/wp-content/uploads/WORKSHEET\\_Human\\_Subjects\\_Research\\_Determination\\_v2.30\\_2020.11.02.pdf](https://web.archive.org/web/20201111071801/https://www.washington.edu/research/wp/wp-content/uploads/WORKSHEET_Human_Subjects_Research_Determination_v2.30_2020.11.02.pdf). Accessed: 2020-11-11.
- University of Winthrop, The Office of Grants and Sponsored Research Development. 2021. Research v/s course assignment. <https://web.archive.org/web/20210325082350/https://www.winthrop.edu/uploadedFiles/grants/Research-vs-Course-Assignment.pdf>. Accessed: 2021-03-25.
- Jack Urbanek, Angela Fan, Siddharth Karamcheti, Saachi Jain, Samuel Humeau, Emily Dinan, Tim Rocktäschel, Douwe Kiela, Arthur Szlam, and Jason Weston. 2019. [Learning to speak and act in a fantasy text adventure game](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 673–683, Hong Kong, China. Association for Computational Linguistics.
- Chao Yang, Shimei Pan, Jalal Mahmud, Huahai Yang, and Padmini Srinivasan. 2015. [Using personal traits for brand preference prediction](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 86–96, Lisbon, Portugal. Association for Computational Linguistics.