

How Good Is NLP? A Sober Look at NLP Tasks through the Lens of Social Impact

Conference Paper**Author(s):**

Jin, Zhijing; Chauhan, Geeticka; Tse, Brian; Sachan, Mrinmaya; Mihalcea, Rada

Publication date:

2021-08

Permanent link:

<https://doi.org/10.3929/ethz-b-000527311>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

<https://doi.org/10.18653/v1/2021.findings-acl.273>

How *Good* Is NLP? A Sober Look at NLP Tasks through the Lens of Social Impact

Zhijing Jin

Max Planck Institute for Intelligent Systems
Tübingen, Germany
zjin@tue.mpg.de

Geeticka Chauhan

MIT
geeticka@mit.edu

Brian Tse

Oxford
briantsemanhei@gmail.com

Mrinmaya Sachan

ETH Zürich
msachan@ethz.ch

Rada Mihalcea

University of Michigan
mihalcea@umich.edu

Abstract

Recent years have seen many breakthroughs in natural language processing (NLP), transitioning it from a mostly theoretical field to one with many real-world applications. Noting the rising number of applications of other machine learning and AI techniques with pervasive societal impact, we anticipate the rising importance of developing NLP technologies for social good. Inspired by theories in moral philosophy and global priorities research, we aim to promote a guideline for social good in the context of NLP. We lay the foundations via the moral philosophy definition of social good, propose a framework to evaluate the direct and indirect real-world impact of NLP tasks, and adopt the methodology of global priorities research to identify priority causes for NLP research. Finally, we use our theoretical framework to provide some practical guidelines for future NLP research for social good.¹

1 Introduction

Advances on multiple NLP fronts have given rise to a plethora of applications that are now integrated into our daily lives. NLP-based intelligent agents like Amazon Echo and Google Home have entered millions of households (Voicebot, 2020). NLP tools are now prevalent on phones, in cars, and in many daily services such as Google search and electronic health record analysis (Townsend, 2013).

In the current COVID-19 context, NLP has already had important positive social impact in the face of a public health crisis. When the pandemic broke out, Allen AI collected the COVID-19 dataset (Wang et al., 2020) with the goal of helping public health experts efficiently sift through the myriad of COVID-19 research papers that emerged in a short time period. Subsequently, NLP services such as

Amazon Kendra were deployed to help organize the research knowledge around COVID-19 (Bhatia et al., 2020). The NLP research community worked on several problems like the question-answering and summarization system CAiRE-COVID (Su et al., 2020), the expressive interviewing conversational system (Welch et al., 2020) and annotation schemas to help fight COVID-19 misinformation online (Alam et al., 2020; Hossain et al., 2020).

As NLP transits from theory into practice and into daily lives, unintended negative consequences that early theoretical researchers did not anticipate have also emerged, from the toxic language of Microsoft’s Twitter bot Tay (Shah and Chokkattu, 2016), to the leak of privacy of Amazon Alexa (Chung et al., 2017). A current highly-debated topic in NLP ethics is GPT-3 (Brown et al., 2020), whose risks and harms include encoding gender and racist biases (Bender et al., 2021).

It is now evident that we must consider the negative and positive impacts of NLP as two sides of the same coin, a consequence of how NLP and more generally AI pervade our daily lives. The consideration of the negative impacts of AI has engendered the recent and popular interdisciplinary field of AI ethics, which puts forth issues such as algorithmic bias, fairness, transparency and equity with an aim to provide recommendations for ethical development of algorithms.

Highly influential works in AI ethics include (Buolamwini and Gebru, 2018; Mitchell et al., 2019; Raji et al., 2020; Chen et al., 2019; Blodgett et al., 2020). AI for social good (AI4SG) (Tomašev et al., 2020) is a related sub-field that benefits from results of AI ethics and while keeping ethical principles as a pre-requisite, has the goal of creating positive impact and addressing society’s biggest challenges. Work in this space includes (Wang et al., 2020; Bhatia et al., 2020; Killian et al., 2019; Lampos et al., 2020).

¹Our data and code are available at http://github.com/zhijing-jin/nlp4sg_acl2021.

Active conversations about ethics and social good have expanded broadly, in the NLP community as well as the broader AI and ML communities. Starting with early discussions in works such as (Hovy and Spruit, 2016; Leidner and Plachouras, 2017), the communities introduced the first workshop on ethics in NLP (Hovy et al., 2017) and the AI for social good workshop (Luck et al., 2018), which inspired various follow-up workshops at venues like ICML and ICLR. The upcoming NLP for Positive Impact Workshop (Field et al., 2020) finds inspiration from these early papers and workshops. In 2020, NeurIPS required all research papers to submit broader impact statements (Castelvecchi, 2020; Gibney, 2020). NLP conferences followed suit and introduced optional ethical and impact statements, starting with ACL in 2021 (Association for Computational Linguistics, 2021).

With the growing impact of our models in daily lives, we need comprehensive guidelines for following ethical standards to result in positive impact and prevent unnecessary societal harm. (Tomašev et al., 2020) provide general guidelines for successful AI4SG collaborations through the lens of United Nations (UN) sustainable development goals (SDGs) (United Nations, 2015) and (Hovy and Spruit, 2016; Leidner and Plachouras, 2017) begin the ethics discussions in NLP. However, there is room for iteration in terms of presenting a comprehensive picture of NLP for social good, with an evaluation framework and guidelines. At the moment, researchers eager to make a beneficial contribution need to base their research agenda on intuition and word of mouth recommendations, rather than a scientific evaluation framework.

To this end, our paper presents a modest effort to the understanding of social good, and sketches thinking guidelines and heuristics for NLP for social good. Our main goal is to answer the question:

Given a specific researcher or team with skills s , and the set of NLP technologies T they can work on, what is the best technology $t \in T$ for them to optimize the social good impact I ?

In order to answer this overall question, we take a multidisciplinary approach in our paper:

- §2 relies on theories in moral philosophy to approach what is social good versus bad (i.e., the sign and rough magnitude of impact I for a *direct* act a);

- §3 relies on causal structure models as a framework to estimate I for $t \in T$, considering that t can be an *indirect* cause of impact;
- §4 relies on concepts from global priorities research and economics to introduce a high-level framework to choose a technology t that optimizes the social impact I ;
- §5 applies the above tools to analyze several example NLP directions, and provides a practical guide on how to reflect on the social impact of NLP.

We acknowledge the iterative nature of a newly emerging field in NLP for social good, requiring continuing discussions on definitions and the development of ethical frameworks and guidelines. Echoing the history of scientific development (Kuhn, 2012), the goal of our work is not to provide a perfect, quantitative, and deterministic answer about how to maximize social good with our NLP applications. The scope of our work is to take one step closer to a comprehensive understanding, through high-level philosophies, thinking frameworks, together with heuristics and examples.

2 What is social good?

Defining social good can be controversial. For example, if we define saving energy as social good, then what about people who get sick because of not turning on the air-conditioner on a cold day? Therefore, social good is context-dependent, relevant to *people, times, and states of nature* (Broome, 2017). This section is to provide a theoretical framework about the social impact I for a direct act a .

2.1 Moral philosophy theories

We can observe that for some acts, it is relatively certain to judge whether the impact is positive or negative. For example, solving global hunger is in general a positive act. Such judgement is called intuitionism (Sidgwick, 1874), a school of moral philosophy.

There are many areas of social impact that cannot receive consensus by intuitions. To find analytical solutions to these debatable topics, several moral philosophies have been proposed. We introduce below three categories of philosophical perspectives to judge moral laws (Kagan, 2018), and provide the percentage of professional philosophers who support the theory (Bourget and Chalmers, 2014):

1. Deontology: emphasizes duties or rules, endorsed by 25.9% philosophers;
2. Consequentialism: emphasizes consequences of acts, endorsed by 23.6% philosophers;
3. Virtue ethics: emphasizes virtues and moral character, endorsed by 18.2% philosophers.

Note that the above three schools, deontology, consequentialism, and virtue ethics, follows the standard textbook introductions for normative ethics in the analytic philosophy tradition. It is also possible for future research to consider different perspectives while defining social good.

A practical guide for using these philosophies.

The three perspectives provide us dimensions to think about the impact I of an act a , so that the final decision is (hopefully) more reliable than one single thought which is subject to biases. Such decomposition practices are often used in highly complicated analyses (e.g., business decisions), such as radar charts to rate a decision/candidate or SMART goals.

A practical guide for using moral philosophies to judge an act a is to think along each of the three perspectives, collect estimations of how good the act a is from the three dimensions, and merge them. For example, using NLP for healthcare to save lives can be good from all three perspectives, and thus it is an overall social good act.

When merging judgements from the above philosophical views, there can be tradeoffs, such as sacrificing one life for five lives in the Trolley problem (Thomson, 1976), which scores high on consequentialism but low on deontology and virtue ethics. One solution by the moral uncertainty theory (MacAskill et al., 2020) is to favor acts with more balanced judgements on all criteria, and reject acts that are completely unacceptable on any criterion.

2.2 Principles for future AI

Many agencies from academia, government, and industries have proposed principles for future AI (Jobin et al., 2019), which can be regarded as a practical guide by *deontology*. Zeng et al. (2019) surveyed the principles of the governance of AI proposed by 27 agencies. The main areas are as follows (with keywords):

- Humanity: beneficial, well-being, human right, dignity, freedom, education, human-friendly.

- Privacy: personal information, data protection, explicit confirmation, control of the data, notice and consent.
- Security: cybersecurity, hack, confidential.
- Fairness: justice, bias, discrimination.
- Safety: validation, test, controllability.
- Accountability: responsibility.
- Transparency: explainable, predictable, intelligible.
- Collaboration: partnership, dialog.
- Share: share, equal.
- AGI: superintelligence.

3 Evaluating the indirect impact of NLP

Given the general moral guide to judge an act with *direct* impacts, we now step towards the second stage – understanding the downstream impact of scientific research which typically has *indirect* impacts. For example, it is not easily tractable to estimate the impact of some linguistic theories. To sketch a solution, this section will first classify NLP tasks by the dimension of theory→application, and then provide an evaluation framework for I of a technology t that may have indirect real-life impacts.

3.1 Classifying tasks from upstream to downstream

To evaluate each NLP research topic, we propose four stages in the *theory*→*application* development, as shown in Figure 1, and categorize the 570 long papers from ACL 2020² according to the four stages in Figure 2. Details of the annotation are in Appendix A. The four stages are as follows.

Stage 1. Fundamental theories. Fundamental theories are the foundations of knowledge, such as linguistic theories by Noam Chomsky. In ACL 2020, the most prevalent topic for papers in Stage 1 is linguistics theory in Figure 2.

Stage 2. Building block tools. Moving one step from theory towards applications is the research on building block tools, which serves as important building blocks and toolboxes for downstream technologies. The most frequently researched Stage-2 topics at ACL 2020 are information extraction, model design, and interpretability (in Figure 2).

²<https://www.aclweb.org/anthology/events/acl-2020/#2020-acl-main>

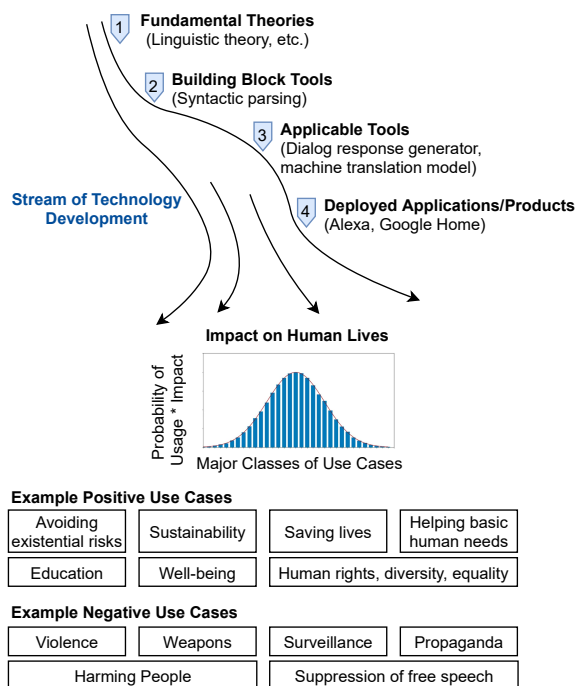


Figure 1: Stream of technology development from theory to application with end impacts. The end impacts are a distribution of use cases and their corresponding weighted impacts.

Stage 3. Applicable tools. Applicable tools are pre-commercialized NLP systems which can serve as the backbones of real-world applications. This category includes NLP tasks such as dialog response generation, question answering, and machine translation. The most common research topics in this category are dialog, machine translation, and question answering (in Figure 2).

Stage 4. Deployed applications/products. Deployed applications often build upon tools in Stage 3, and wrap them with user interfaces, customer services, and business models. Typical examples of Stage-4 technologies include Amazon Echo, Google Translate, and so on. The top three topics of ACL 2020 papers in this category are ways to address misinformation (e.g., a fact checker for news bias), dialog, and NLP for healthcare.

3.2 Estimating impact

Direct impacts of Stage-4 technologies. A direct impact of NLP development is allowing users more free time. This is evident in automatic machine translation, which saves the effort and time of human translators, or in NLP for healthcare, which allows doctors to more quickly sift through patient history. Automatic fake news detection frees up

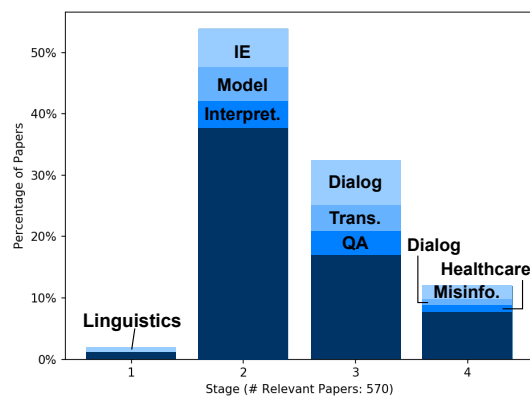


Figure 2: Distribution of ACL 2020 papers by the four stages. For each stage, we highlight the top several topics of the papers. We only list the top one topic for Stage 1 due to visual space limit. Abbreviations of technologies include Information Extraction (IE), Interpretability (Interpret.), machine translation (Trans.), question answering (QA), and misinformation (Misinfo.).

time for human fact-checkers, to aid them in more quickly detecting fake news through the increasing number of digital news articles being published.

The impact of more user free time is varied. In the case of healthcare, NLP can free up time for more personalized patient care, or allow free time for activities of choice, such as spending time on passion projects or more time with family. We recognize these varied impacts of NLP deployment, and recommend user productivity as one way to measure it.

Note that there can be positive as well as negative impact associated with rising productivity, and the polarity can be decided according to Section 2.1. Typical positive impacts of NLP technology include better healthcare and well-being, and in some cases it indirectly helps with avoiding existential risks, sustainability, and so on. Typical negative impacts include more prevalent surveillance, propaganda, breach of privacy, and so on. For example, intelligent bots can improve efficiency at work (to benefit economics), and bring generally better well-being for households, but they might leak user privacy (Chung et al., 2017).

Thus, estimating the overall end impact of a technology t in the Stage 4 needs to accumulate over a set of aspects \mathcal{AS} :

$$I(t) = \sum_{as \in \mathcal{AS}} \text{scale}_{as}(t) \cdot \text{impact}_{as}(t), \quad (1)$$

where $\text{scale}_{as}(t)$ is the usage scale of applications of technology t used in the aspect as , and




Figure 3: Viewing the theory→application process with a structural causal model.

$\text{impact}_{as}(t)$ is the impact of t in this aspect.

Indirect impacts of early stage technologies.

Although the direct impact of Stage-4 technologies can be estimated by Eq. (1), it is difficult to calculate the impact of a technology in earlier stages (i.e., stage 1-3).

We can approach the calculation of indirect impacts I of an early-stage technology t by a structural causal model. As shown in the causal graph \mathcal{G} in Figure 3, each technology t is in a causal chain from its parent vertex set $\text{PA}(t)$ (i.e., upstream technologies that directly causes the invention of t), to its children vertex set (i.e., downstream technologies directly resulting from t). Formally, we denote a directed (causal) path in \mathcal{G} as a sequence of distinct vertices (t_1, t_2, \dots, t_n) such that $t_{i+1} \in \text{CH}(t_i)$ for all $i = 1, \dots, n-1$. We call t_n a descendant of t_1 . After enumerating all paths, we denote the set of all descendants of t as $\text{DE}(t)$. Specifically, we denote all descendant nodes in Stage 4 as Stage-4 $\text{DE}(t)$.

Hence, the impact of any technology t is the sum of impact of all its descendants in Stage 4:

$$I(t) = \sum_{x \in \text{Stage-4 DE}(t)} p(x) \cdot c_x(t) \cdot I(x), \quad (2)$$

where $p(x)$ is the probability that the descendent technology x can be successfully developed, $c_x(t)$ is the contribution of t to x , and $I(x)$ can be calculated by Eq. (1). This formula can also be interpreted from the light of do-calculus (Pearl, 1995) as $P(X|\text{do}(t)) - P(X)$, for $X \in \text{Stage-4 DE}(t)$, which means the effect of intervention $\text{do}(t)$ on Stage 4 descendants.

Note that Eq. (1) and (2) are meta frameworks, and we leave it to future work to utilize these for assessing the social impact of their work.

3.3 Takeaways for NLP tasks

With the growing interest of AI and NLP publication venues (e.g., NeurIPS, ACL) in ethical and

broader impact statements, it will be useful and important for researchers to have practical guidelines on evaluating the impact of their NLP tasks.

We first introduce some thinking steps to estimate the impact of research on an NLP task t :

- (S1) Classify the NLP task t into one of the four stages (§3.1)
- (S2) If t is in Stage 4, think of the set of aspects \mathcal{AS} that t will impact, the scale of applications, and aspect-specific impact magnitude. Finally, estimate impact using Eq. (1).
- (S2') If t is in Stage 1-3, think of its descendant technologies, their success rate, and the contribution of t to them. Finally, estimate impact using Eq. (1) and (2).

Next, we introduce some high-level heuristics to facilitate fast decisions:

- (H1) For earlier stages (i.e., Stage 1-2), it is challenging to quantify the exact social impact. Their overall impact tends to lean towards positive as they create more knowledge that benefits future technology development.
- (H2) Developers of Stage-4 technologies should be the most careful about ethical concerns. Enumerate the use cases, and estimate the scale of each usage by thinking of the stakeholders, economic impact, and users in the market. Finally, evaluate the final impact before proceeding. (E.g., if the final impact is very negative, then abandon or do it with restrictions).
- (H3) For Stage-3 technologies, if their Stage-4 descendants are tractable to enumerate and estimate for their impacts, then aggregate the descendants' impacts by Eq. 2. Otherwise, treat them like (H1).

4 Deciding research priority

There are many directions for expansion of our efforts for social good; however, due to limited resources and availability of support for each researcher, we provide a research priority list. In this section, we are effectively trying to answer the overall question proposed in Section 1. Specifically, we adopt the practice in the research field *global priorities* (GP) (MacAskill, 2015; Greaves and McAskill, 2017). We first introduce the high-level decision-making framework in Section 4.1, and then formulate these principles using technical terms in Section 4.2.

4.1 Important/Neglected/Tractable (INT) framework

Our thinking framework to address the research priority follows the practice of existing cost-benefit analysis in GP (MacAskill, 2015; Greaves and McAskill, 2017), which aligns with the norms in established fields such as development economics, welfare economics, and public policy.

We draw an analogy between the existing GP research and NLP for social good. Basically, GP addresses the following problem: given, for example, 500 billion US dollars (which is the annual worldwide expenditure on social good), what priority areas should we spend on? Inspired by this practical setting, we form an analogy to NLP research efforts, namely to answer the question proposed in Section 1 about how to attribute resources and efforts on NLP research for social good.

The high-level intuitions are drawn from the Important/Neglected/Tractable (INT) framework (MacAskill, 2015), a commonly adopted framework in global priorities research on social good. Assume each agent has something to contribute (e.g., money, effort, etc.). It is generally effective to contribute to important, neglected, and tractable areas.

4.2 Calculation of priority

Although the INT framework is commonly used in practice of many philanthropy organizations (MacAskill, 2015), it will be more helpful to formulate it using mathematical terms and economic concepts. Note that the terms we formulate in this section can be regarded as elements in our proposed thinking framework, but they are not directly calculable.³

Our end goal is to estimate the cost-effectiveness of contributing a unit time and effort of a certain researcher or team to research on the technology t . So far we have a meta framework to estimate the impacts I brought by successful development of a technology t . And we introduce the notations in Table 1.

³We adapted these terms from GP. Such terms to estimate priority has been successfully used by real-world social good organizations, e.g., GiveWell, Global Priorities Institute, the Open Philanthropy Project (a foundation with over 10 billion USD investment), ReThink Priorities, 80,000 Hours Organization. In the long run, the NLP community may potentially benefit from aligning with GP’s terminology. Still, we do not recommend applying our framework in high-stake settings yet, since it serves only as a starting point currently.

Notation	Meaning
r	An NLP researcher or research group
$\mathbf{T}(r)$	The set of NLP topics that the researcher can pursue (limited by skills, resources, and passion)
t	An NLP technology
$I(t)$	Social impacts brought by successful development of t
$\text{prog}(t)$	The current progress of t
$p(t; r)$	Probability that research in t succeeds based on the skills of the researcher r
$p(t; r)I(t)$	Expected social impact of the researcher r ’s work on t
$\Delta t(r)$	Improvement of t per unit resource (incl. time, effort, money, etc.) of the researcher r

Table 1: Notations and their corresponding meanings used for cost-effectiveness calculation.

For a researcher r , the action set per unit resource is $\{\Delta t | t \in \mathbf{T}(r)\}$. Equivalently speaking, they can intervene at a node t by the amount of $\Delta t(r)$ in the structured causal graph \mathcal{G} in Figure 3.

The first useful concept is $p(t; r)I(t)$, **the expected social impact** of research on a technology t . Here the success rate $p(t; r)$ is crucial because most research does not necessarily produce the expected outcome. However, if the impact of a technology can be extremely large (for example, prevention of extinction has impact near positive infinity), then even with a very little success rate, we should still devote considerable efforts into it.

The second concept that is worth attention is **the marginal impact** (Pindyck et al., 1995) of one more unit of resources of the researcher r into the technology t , calculated as

$$\Delta I(t; r) := I(\text{prog}(t) + \Delta t(r)) - I(\text{prog}(t)). \quad (3)$$

For example, if the field associated with the technology is almost saturated, or if many other researchers working on this field are highly competent, then, for a certain research group, blindly devoting time to the field may have little marginal impact. However, on the other hand, if a field is important but neglected, the marginal impact of pushing it forward can be large. This also explains why researchers are passionate about creating a new research field.

The third useful concept is **the opportunity cost** (Palmer and Raftery, 1999) to devote researcher r ’s resources into the technology t instead of a possibly more optimal technology t^* . Formally,

the opportunity cost is calculated as

$$t^*(r) := \arg \max_x \Delta I(x(r)), \quad (4)$$

$$\text{Cost}(t; r) := \Delta I(t^*(r); r) - \Delta I(t; r), \quad (5)$$

where t^* is the optimal technology that can bring the largest expected improvement of social impact. The opportunity cost conveys the important message that we should not just do *good*, but do the *best*, because the difference from good to best can be a large loss.

Estimating the variables. Note that the frameworks we have proposed so far are at the meta level, useful for guiding thought experiments, and future research. Exact calculations are not possible with the current state of research in NLP for social good, although achievable in the future.

A practical insight is that NLP researchers estimate the impact of their research via qualitative explanations (natural language) or rough quantitative ones. For example, the introduction section of most NLP papers or funding proposals is a natural language-based estimation of the impact of the research. Such estimations can be useful to some extent (Hubbard and Drummond, 2011), although precise indicators of impact can motivate the work more strongly.

We can also borrow some criteria from effective altruism, a global movement that establishes a philosophical framework, and also statistical calculations of social good. One of the established metrics for calculating impact is called the “quality-adjusted life years” (QALYs) proposed by MacAskill (2015). QALYs count the number of life years (calibrated by life quality such as health conditions) that an act helps to increase.

5 Evaluating NLP tasks

In this section, we will first try to categorize the current state of NLP research for social good based on ACL 2020 papers, and then highlight NLP topics that are aligned with the UN’s SDGs. We will conclude with a practical checklist and case studies of common NLP tasks using this checklist.

5.1 Current state of NLP research for social good – ACL 2020 as a case study

We want to compare the ideal priority list with the current distribution of NLP papers for social good. As a case study of the current research frontier, we plot the topic distribution of the 89 ACL 2020




Figure 4: Social good topics at ACL 2020 by countries.

papers that are related to NLP for social good in Figure 4. We also show the portion of papers by the 10 countries with the most social-good papers. Our annotation details are in Appendix A.

Illustrated in Figure 4, most social-good papers work on interpretability, tackling misinformation (e.g., fact-checking for news), and healthcare (e.g., to increase the capacity of doctors). In terms of countries, the US has the most papers on interpretability, and no papers on NLP for education, NLP for legal applications, and some other topics. China has few papers on interpretability, although interpretability is the largest topic. India has no papers on fighting misinformation, although it is the second largest topic. Only 5 countries have publications across more than two social good topics. Please refer to Appendix B for more analyses such as social-good papers by academia vs. industries.

However, compared with the UN’s SDGs (United Nations, 2015), the current NLP research (at least in the scope of ACL conference submissions) lacks attention to other important cause areas such as tackling global hunger, extreme poverty, clean water and sanitation, and clean energy. There are also too few research papers on NLP for education, although education is the 4th most important area in SDGs.

One cause of this difference is *value misalignment*. Most NLP research is supported by stakeholders and funding agencies, which have a large impact on the current research trends or preferences in the NLP community. The perspective from social good with a framework to calculate the priority list has still not reached many in the NLP community.

Although we do not have data on expenditure

Cause	Annual Spending (USD)
Global R&D	1.5 trillion (UNESCO, 2017)
Luxury Goods	1.3 trillion (D’arpizio et al., 2015)
US Social Welfare	900 billion (Ferrara, 2011)
Climate Change	>300 billion (Buchner et al., 2014)
Global Poverty	>250 billion (Todd, 2017)
Nuclear Security	1-10 billion (Todd, 2017)
Pandemic Prevention	1 billion (Todd, 2017)
AI Safety Research	10 million (Todd, 2017)

Table 2: Annual spending of the cause areas.

Priority	Example NLP research topics
Poverty	<ul style="list-style-type: none"> Predicting poverty by geo-located Wikipedia articles (Sheehan et al., 2019) Parsing fund applicant profiles (proposed)
Hunger	<ul style="list-style-type: none"> NLP for agriculture (Yunpeng et al., 2019) NLP for food allocation (proposed)
Health & Well-being	<ul style="list-style-type: none"> NLP to analyze clinical notes (Dernoncourt et al., 2017a,b; Luo et al., 2018; Gopinath et al., 2020; Leiter et al., 2020a,b) NLP for psychotherapy and counseling (Biester et al., 2020; Xu et al., 2020; Pérez-Rosas et al., 2019) NLP for happiness (Asai et al., 2018; Evensen et al., 2019) Assistive speech generation (proposed)
Education	<ul style="list-style-type: none"> NLP for educational question answering (Atapattu et al., 2015; Lende and Raghuvanshi, 2016) Improving textbooks (Agrawal et al., 2010) Automated grading (Madnani and Cahill, 2018; Taghipour and Ng, 2016) Plagiarism detection (Chong et al., 2010) Tools for learners with disabilities (proposed)
Equality	<ul style="list-style-type: none"> Interpretability (Köhn, 2015; Belinkov et al., 2017; Nie et al., 2020) Ethics of NLP (Hovy and Spruit, 2016; Stanovsky et al., 2019; Sap et al., 2019) NLP for low-resource languages (Zoph et al., 2016; Kim et al., 2017) NLP on resource-limited devices (Sun et al., 2020) NLP tools that signal bias in human language and speech (proposed)
Clean water	<ul style="list-style-type: none"> Raising public awareness of water sanitation (proposed)
Clean energy	<ul style="list-style-type: none"> Green NLP (Strubell et al., 2019; Schwartz et al., 2020) NLP to analyze cultural values regarding climate change (Jiang et al., 2017; Koenecke and Feliu-Fabà, 2019) Cross-cultural models of climate change perceptions (proposed)

Table 3: Top priorities and some NLP research related to each of them. This list may not be exhaustive. We also propose a high-impact research problem in each of the areas which has received less attention so far.

in each NLP subarea, we can get a glimpse of the value misalignment in general. Table 2 shows the annual spending of some cause areas. Note that the ranking of the expenditure does not align with

our priority list for social good. For example, luxury goods are not as important as global poverty, but luxury goods cost 1.3 trillion USD each year, almost five times the expenditure in global poverty.

5.2 Aligning NLP with social good

In this subsection, we list the top priorities according to UN’s SDGs (United Nations, 2015). For each goal, in Table 3 we include examples of existing NLP research, and suggest potential NLP tasks that can be developed (labeled as (proposed)).

5.3 Checklist

As a practical guide, we compile the takeaways of this paper into a list of heuristics that might be helpful for future practitioners of NLP for social good. To inspect the social goodness of an NLP research direction (especially in Stage 3-4), the potential list of questions to answer is as follows:

- (Q1) What kind of people/process will be benefited from the technology?
- (Q2) Does it enforce traditional structure of beneficiaries? I.e., what groups of underprivileged people can be benefited? (e.g., by gender, demographics, socio-economic status, country, native languages, disability type)
- (Q3) Does it contribute to SDG priority goals such as poverty, hunger, health, education, equality, clean water, and clean energy?
- (Q4) Can it directly improve quality of lives? E.g., how many QALYs might it result in?
- (Q5) Does it count as (a) mitigating problems brought by NLP, or (b) proactively helping out-of-NLP social problems?

5.4 Case studies by the checklist

We conduct some case studies of NLP technologies using the checklist.

Low-resource NLP & machine translation.

This category includes NLP on low-resource languages, such as NLP for Filipino (Sagum et al., 2019; Cruz et al., 2020), and machine translation in general. Because this direction expands the users of NLP technologies from English-speaking people to other languages, it benefits people speaking these languages (Q1), and helps to narrow the gap between English-speaking and non-English speaking end users (Q2), although it is still likely that people who can afford intelligent devices will benefit more than those who cannot. This category can contribute directly to goals such as equality

and education, and indirectly to other goals because translation of documents in general helps the sharing of information and knowledge (Q3). It directly improves quality of lives, for example, for immigrants who may have difficulties with the local language (Q4). Thus, it counts as social good category (b) by (Q5).

Transparency, interpretability, algorithmic fairness and bias. Research in this direction can impact users who need more reliable decision-making NLP, such as the selection process for loans, jobs, criminal judgements, and medical treatments (Q1). It can shorten the waiting time of candidates and still make fair decisions regardless of spurious correlations (Q2) (Q4). It reduces inequality raised by AI, but not increasing equality over man-made decisions, at least by the current technology (Q2). Thus, it is social good category (a) by (Q5).

Green NLP. Green NLP reduces the energy consumption of large-scale NLP models. Although it works towards the goal of affordable and clean energy (Q3), it just neutralizes the negative impact of training NLP models, but not impacting out-of-NLP energy problems. Green NLP belongs to social good category (a) by (Q5). It does not have large impacts directly targeted at (Q1), (Q2) and (Q4).

QA & dialog. People who can afford devices embedded with intelligent agents can use it, which is about 48.46% of the global population (BankMyCell, 2021) (Q1). So this benefits people with higher socio-economic status, and benefits English speaking people more than others, not to mention job replacements for labor-intensive service positions (Q2). It does not contribute to priority goals except for education and healthcare for people who can afford intelligent devices (Q3). Nonetheless, it can improve the quality of lives for its user group (Q4). It can be regarded as social good of category (b) by (Q5).

Information extraction, NLP-powered search engine & summarization. This direction speeds up the information compilation process, which can increase the productivity in many areas. About 50% of the world population have access to the Internet and thus can use it (Meeker, 2019) (Q1) (Q2). This category indirectly helps education, and the information compilation process of other goals

(Q3). It can largely improve the lives of its user group because people gather information very frequently (e.g., do at least one Google search every day) (Q4). Thus, it belongs to social good category (b) by (Q5).

NLP for social media. This category is beneficial for social scientists interested in the trends and culture, or companies and organizations interested in market analyses (Q1) (Q2). The benefits to social science can indirectly help evidence-based policy makers, despite the caveat of breach of user privacy, and information manipulation (Q4). This direction has limited contribution to priorities such as poverty and hunger, because extremely poor population do not own devices to post on social media (Q3). It is social good category (b) by (Q5).

6 Conclusion

This paper presented a meta framework to evaluate NLP tasks in the light of social good, and proposed a practical guide for practitioners in NLP. We call for more attention towards awareness and categorization of social impact of NLP research, and we envision future NLP research taking on an important social role and contributing to multiple priority areas. We also acknowledge the iterative nature of this emerging field, requiring continuing discussions, improvements to our thinking framework and different ways to implement it in practice. We highlight that the goal of our work is to take one step closer to a comprehensive understanding of social good rather than introducing a deterministic answer about how to maximize social good with NLP applications.

Acknowledgments

We thank Bernhard Schoelkopf, Kevin Jin, and Qipeng Guo for insightful discussions on the main ideas and methodology of the paper. We thank Osmond Wang for checking the economic concepts in the paper. We also thank Chris Brockett for checking many details in the paper. We thank the labmates in the LIT lab at University of Michigan, especially Laura Biester, Ian Stewart, Ashkan Kazemi, and Andrew Lee for constructive feedback. We also thank labmates at the MIT MEDG group, especially William Boag and Peter Szolovits for their constructive feedback. We thank Jingwei Ni for the help with the annotation of the country and affiliation of the ACL 2020 papers.

Ethical and societal implications

Our paper establishes a framework to better understand the definition of social good in the context of NLP research, and lays out a recommended direction on how to achieve it. The contributions of our paper could benefit a focused, organized and accountable development of NLP for social good. The data used in our work is public, and without privacy concerns.

References

- Rakesh Agrawal, Sreenivas Gollapudi, Krishnaram Kenthapadi, Nitish Srivastava, and Raja Velu. 2010. Enriching textbooks through data mining. In *Proceedings of the First ACM Symposium on Computing for Development*, pages 1–9.
- Firoj Alam, Fahim Dalvi, Shaden Shaar, Nadir Durrani, Hamdy Mubarak, Alex Nikolov, Giovanni Da San Martino, Ahmed Abdelali, Hassan Sajjad, Kareem Darwish, et al. 2020. Fighting the covid-19 infodemic in social media: a holistic perspective and a call to arms. *arXiv preprint arXiv:2007.07996*.
- Akari Asai, Sara Eversen, Behzad Golshan, Alon Y. Halevy, Vivian Li, Andrei Lopatenko, Daniela Stepanov, Yoshihiko Suhara, Wang-Chiew Tan, and Yinzhan Xu. 2018. [HappyDB: A corpus of 100, 000 crowdsourced happy moments](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC 2018, Miyazaki, Japan, May 7-12, 2018*. European Language Resources Association (ELRA).
- Association for Computational Linguistics. 2021. [Ethics FAQ at ACL-IJCNLP 2021](#).
- Thushari Atapattu, Katrina Falkner, and Nickolas Falkner. 2015. Educational question answering motivated by question-specific concept maps. In *International Conference on Artificial Intelligence in Education*, pages 13–22. Springer.
- BankMyCell. 2021. [How many smartphones are in the world?](#)
- Yonatan Belinkov, Lluís Màrquez, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James Glass. 2017. [Evaluating layers of representation in neural machine translation on part-of-speech and semantic tagging tasks](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1–10, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big. In *Proceedings of the 2021 Conference on Fairness, Accountability, and Transparency*.
- Parminder Bhatia, Kristjan Arumae, Nima Pourdamghani, Suyog Deshpande, Ben Snively, Mona Mona, Colby Wise, George Price, Shyam Ramaswamy, and Taha Kass-Hout. 2020. AWS CORD19-search: A scientific literature search engine for COVID-19. *arXiv preprint arXiv:2007.09186*.
- Laura Biester, Katie Matton, Janarathanan Rajendran, Emily Mower Provost, and Rada Mihalcea. 2020. [Quantifying the effects of COVID-19 on mental health support forums](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- David Bourget and David J Chalmers. 2014. What do philosophers believe? *Philosophical studies*, 170(3):465–500.
- John Broome. 2017. *Weighing goods: Equality, uncertainty and time*. John Wiley & Sons.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Barbara Buchner, Morgan Herve-Mignucci, Chiara Trabacchi, Jane Wilkinson, Martin Stadelmann, Rodney Boyd, Federico Mazza, Angela Falconer, and Valerio Micale. 2014. Global landscape of climate finance 2015. *Climate Policy Initiative*, 32.
- Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR.
- Davide Castelvecchi. 2020. Prestigious ai meeting takes steps to improve ethics of research. *Nature*.
- Irene Y Chen, Peter Szolovits, and Marzyeh Ghassemi. 2019. Can ai help reduce disparities in general medical and mental health care? *AMA journal of ethics*, 21(2):167–179.

- Miranda Chong, Lucia Specia, and Ruslan Mitkov. 2010. Using natural language processing for automatic detection of plagiarism. In *Proceedings of the 4th International Plagiarism Conference (IPC-2010)*.
- Hyunji Chung, Michaela Iorga, Jeffrey Voas, and Sangjin Lee. 2017. Alexa, can i trust you? *Computer*, 50(9):100–104.
- Jan Christian Blaise Cruz, Jose Kristian Resabal, James Lin, Dan John Velasco, and Charibeth Cheng. 2020. [Investigating the true performance of transformers in low-resource languages: A case study in automatic corpus creation](#). *CoRR*, abs/2010.11574.
- Claudia D’arpizio, Federica Levato, Daniele Zito, and Joelle de Montgolfier. 2015. Luxury goods worldwide market study. *Bain & Company’s report*.
- Franck Deroncourt, Ji Young Lee, and Peter Szolovits. 2017a. [NeuroNER: an easy-to-use program for named-entity recognition based on neural networks](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 97–102. Association for Computational Linguistics.
- Franck Deroncourt, Ji Young Lee, Ozlem Uzuner, and Peter Szolovits. 2017b. De-identification of patient notes with recurrent neural networks. *Journal of the American Medical Informatics Association*, 24(3):596–606.
- Sara Evensen, Yoshihiko Suhara, Alon Y. Halevy, Vivian Li, Wang-Chiew Tan, and Saran Mumick. 2019. [Happiness entailment: Automating suggestions for well-being](#). In *8th International Conference on Affective Computing and Intelligent Interaction, ACII 2019, Cambridge, United Kingdom, September 3-6, 2019*, pages 62–68. IEEE.
- Peter Ferrara. 2011. America’s ever expanding welfare empire. *Forbes*, April, 22.
- Anjalie Field, Shrimai Prabhumoye, Maarten Sap, Zhi-jing Jin, Jieyu Zhao, and Chris Brockett. 2020. [1st workshop on nlp for positive impact](#).
- Elizabeth Gibney. 2020. The battle for ethical ai at the world’s biggest machine-learning conference. *Nature*, 577(7792):609–609.
- Divya Gopinath, Monica Agrawal, Luke Murray, Steven Horng, David Karger, and David Sontag. 2020. Fast, structured clinical documentation via contextual autocomplete. In *Machine Learning for Healthcare Conference*, pages 842–870. PMLR.
- Hilary Greaves and William McAskill. 2017. A research agenda for the Global Priorities Institute. *University of Oxford, London*.
- Tamanna Hossain, Robert L. Logan IV, Arjuna Ugarte, Yoshitomo Matsubara, Sean Young, and Sameer Singh. 2020. [COVIDLies: Detecting COVID-19 misinformation on social media](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.
- 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.
- Douglas W Hubbard and David Drummond. 2011. *How to measure anything*. Wiley Online Library.
- Ye Jiang, Xingyi Song, Jackie Harrison, Shaun Quegan, and Diana Maynard. 2017. [Comparing attitudes to climate change in the media using sentiment analysis based on latent dirichlet allocation](#). In *Proceedings of the 2017 Workshop: Natural Language Processing meets Journalism, NLPmJ@EMNLP, Copenhagen, Denmark, September 7, 2017*, pages 25–30. Association for Computational Linguistics.
- Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The global landscape of ai ethics guidelines. *Nature Machine Intelligence*, 1(9):389–399.
- Shelly Kagan. 2018. *Normative ethics*. Routledge.
- Jackson A Killian, Bryan Wilder, Amit Sharma, Vinod Choudhary, Bistra Dilkina, and Milind Tambe. 2019. Learning to prescribe interventions for tuberculosis patients using digital adherence data. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2430–2438.
- Joo-Kyung Kim, Young-Bum Kim, Ruhi Sarikaya, and Eric Fosler-Lussier. 2017. Cross-lingual transfer learning for pos tagging without cross-lingual resources. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 2832–2838.
- Allison Koenecke and Jordi Feliu-Fabà. 2019. [Learning twitter user sentiments on climate change with limited labeled data](#). *CoRR*, abs/1904.07342.
- Arne Köhn. 2015. What’s in an embedding? analyzing word embeddings through multilingual evaluation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2067–2073.
- Thomas S Kuhn. 2012. *The structure of scientific revolutions*. University of Chicago press.

- Vasileios Lampos, Simon Moura, Elad Yom-Tov, Ingemar J Cox, Rachel McKendry, and Michael Edelstein. 2020. Tracking COVID-19 using online search. *arXiv preprint arXiv:2003.08086*.
- 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.
- Richard Leiter, Enrico Santus, Zhijing Jin, Katherine Lee, Miryam Yusufov, Edward Moseley, Yujie Qian, Jiang Guo, and Charlotta Lindvall. 2020a. An artificial intelligence algorithm to identify documented symptoms in patients with heart failure who received cardiac resynchronization therapy (s717). *Journal of Pain and Symptom Management*, 59(2):537–538.
- Richard E Leiter, Enrico Santus, Zhijing Jin, Katherine C Lee, Miryam Yusufov, Isabel Chien, Ashwin Ramaswamy, Edward T Moseley, Yujie Qian, Deborah Schrag, et al. 2020b. Deep natural language processing to identify symptom documentation in clinical notes for patients with heart failure undergoing cardiac resynchronization therapy. *Journal of Pain and Symptom Management*, 60(5):948–958.
- Sweta P Lende and MM Raghuvanshi. 2016. Question answering system on education acts using nlp techniques. In *2016 world conference on futuristic trends in research and innovation for social welfare (Startup Conclave)*, pages 1–6. IEEE.
- Margaux Luck, Jonnie Penn, Tristan Sylvain, Mark Crowley, Joseph Paul Cohen, Margaux Luck, Sean McGregor, Myriam Côté, Valentine Goddard, Margaux Luck, Kenny Chen, Arsene Fansi Tchango, Joseph Paul Cohen, and Yoshua Bengio. 2018. [AI for social good workshop](#).
- Yuan Luo, Yu Cheng, Özlem Uzuner, Peter Szolovits, and Justin Starren. 2018. Segment convolutional neural networks (seg-cnns) for classifying relations in clinical notes. *Journal of the American Medical Informatics Association*, 25(1):93–98.
- Michael MacAskill, Krister Bykvist, and Toby Ord. 2020. *Moral uncertainty*. Oxford University Press.
- William MacAskill. 2015. *Doing good better: Effective altruism and a radical new way to make a difference*. Guardian Faber Publishing.
- Nitin Madnani and Aoife Cahill. 2018. Automated scoring: Beyond natural language processing. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1099–1109.
- Mary Meeker. 2019. [Internet trends report; 2019](#).
- Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 220–229.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Stephen Palmer and James Raftery. 1999. Opportunity cost. *Bmj*, 318(7197):1551–1552.
- Judea Pearl. 1995. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688.
- Verónica Pérez-Rosas, Xinyi Wu, Kenneth Resnicow, and Rada Mihalcea. 2019. [What makes a good counselor? learning to distinguish between high-quality and low-quality counseling conversations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 926–935, Florence, Italy. Association for Computational Linguistics.
- Robert S Pindyck, Daniel L Rubinfeld, and Prem L Mehta. 1995. *Microeconomics*, volume 4. Prentice Hall Englewood Cliffs, NJ.
- Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the ai accountability gap: defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 33–44.
- Ria Ambrocio Sagum, Aldrin D Ramos, and Monique T Llanes. 2019. Ficobu: Filipino wordnet construction using decision tree and language modeling. *International Journal of Machine Learning and Computing*, 9(1):103–107.
- 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.
- Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2020. [Green ai](#). *Commun. ACM*, 63(12):54–63.
- Saqib Shah and Julian Chokkattu. 2016. [Microsoft kills AI chatbot Tay \(twice\) after it goes full Nazi](#).
- Evan Sheehan, Chenlin Meng, Matthew Tan, Burak Uzkent, Neal Jean, Marshall Burke, David B. Lobell, and Stefano Ermon. 2019. [Predicting economic development using geolocated wikipedia articles](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*, pages 2698–2706. ACM.
- Henry Sidgwick. 1874. *The methods of ethics*. Macmillan and Co.

- Gabriel Stanovsky, Noah A. Smith, and Luke Zettlemoyer. 2019. [Evaluating gender bias in machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy. Association for Computational Linguistics.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and policy considerations for deep learning in nlp. *arXiv preprint arXiv:1906.02243*.
- Dan Su, Yan Xu, Tiezheng Yu, Farhad Bin Siddique, Elham Barezi, and Pascale Fung. 2020. [CAiRE-COVID: A question answering and query-focused multi-document summarization system for COVID-19 scholarly information management](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. [Mobilebert: a compact task-agnostic BERT for resource-limited devices](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2158–2170. Association for Computational Linguistics.
- Kaveh Taghipour and Hwee Tou Ng. 2016. A neural approach to automated essay scoring. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 1882–1891.
- Judith Jarvis Thomson. 1976. Killing, letting die, and the trolley problem. *The Monist*, 59(2):204–217.
- Benjamin Todd. 2017. [The case for reducing existential risks](#).
- Nenad Tomašev, Julien Cornebise, Frank Hutter, Shakir Mohamed, Angela Picciariello, Bec Connelly, Danielle CM Belgrave, Daphne Ezer, Fanny Cachat van der Haert, Frank Mugisha, et al. 2020. Ai for social good: unlocking the opportunity for positive impact. *Nature Communications*, 11(1):1–6.
- Hilary Townsend. 2013. Natural language processing and clinical outcomes: the promise and progress of nlp for improved care. *Journal of AHIMA*, 84(2):44–45.
- UNESCO. 2017. [Facts and figures: R&d expenditure](#).
- United Nations. 2015. [UN Sustainable Development Goals](#).
- Voicebot. 2020. [Amazon Echo & Alexa stats](#).
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Darrin Eide, Kathryn Funk, Rodney Kinney, Ziyang Liu, William Merrill, et al. 2020. Cord-19: The COVID-19 open research dataset. *ArXiv*.
- Charles Welch, Allison Lahnal, Veronica Perez-Rosas, Siqi Shen, Sarah Seraj, Larry An, Kenneth Resnicow, James Pennebaker, and Rada Mihalcea. 2020. [Expressive interviewing: A conversational system for coping with COVID-19](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.
- Zhentao Xu, Verónica Pérez-Rosas, and Rada Mihalcea. 2020. [Inferring social media users’ mental health status from multimodal information](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6292–6299, Marseille, France. European Language Resources Association.
- Cui Yunpeng, Wang Jian, and Liu Juan. 2019. [The development of deep learning based natural language processing \(nlp\) technology and applications in agriculture](#). *Journal of Agricultural Big Data*, 1(1):38.
- Yi Zeng, Enmeng Lu, and Cunqing Huangfu. 2019. [Linking artificial intelligence principles](#). In *Workshop on Artificial Intelligence Safety 2019 co-located with the Thirty-Third AAAI Conference on Artificial Intelligence 2019 (AAAI-19), Honolulu, Hawaii, January 27, 2019*, volume 2301 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. [Transfer learning for low-resource neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575, Austin, Texas. Association for Computational Linguistics.

A ACL 2020 paper annotations

For the case study on ACL 2020 papers, such as Figure 2 and 4, we collect the 570 long papers at ACL 2020. An NLP researcher with four years of research experience conducted the entire annotation, so that the categorization is consistent across all papers.⁴

The first annotation task is to categorize all papers into one of the four stages in the theory→application development. We showed the annotator the description of the four stages in Section 3.1. Next, provided with the title, abstract, and PDF of each paper, the annotator was asked to annotate which of the four stages each paper belongs to. The annotator had passed a test batch before starting the large-scale annotation.

The second annotation task is to annotate the research topics of the papers related to social good at ACL 2020. If the paper has a clear social good impact (89 out of 570 papers), the annotator needs to classify the topic of the paper into one of the given categories: bias mitigation, education, equality, fighting misinformation, green NLP, healthcare, interpretability, legal applications, low-resource language, mental healthcare, robustness, science literature parsing, and others. For the other meta information such as countries, or academia vs. industry, we decide based on the information of the leading first author.

B More statistics about ACL 2020 papers

For the case study on ACL 2020 papers, we further investigate the following statistics.

Stage 1-4 by countries. Recall that in Figure 2 of the main paper, we plot the distributions of papers by the four stages, and highlight the most frequent topics in each stage. Additionally, it is also interesting to explore the distribution of stages for different countries. In Figure 5, we have the following observations:

China does not have Stage-1 papers (i.e., fundamental theories), although it has the second largest total number of papers. The reason might be that there are not many Chinese researchers on linguistic theories who publish at English conferences.

Most countries' number of papers in the four stages follows the overall trend (i.e., Stage-2 papers > Stage-3 papers > Stage-4 papers > Stage-1

⁴The annotation file has been uploaded to the softconf system.

papers), with a few exceptions. For example, China has almost the same number of papers in Stage 2 and 3, Germany has more papers in Stage 4 (i.e., deployed applications) than in Stage 3, and Canada has the most papers in Stage 3.




Figure 5: Stage 1-4 of ACL 2020 papers by countries.

Social good topics by academia vs. industry.

As we call for more research attention to NLP for social good, it is important to understand the affiliations behind the current social good papers. A coarse way is to look at the affiliation of the first author, and inspect whether the main work of the paper is done by people from academia or industry.




Figure 6: Social good topics at ACL 2020 by affiliations.

As in Figure 6, overall academia publishes several times more papers on social good than the industry. This ratio is higher than the average ratio of papers from academia out of all ACL 2020 papers (389 from academia out of 570). Industry does not have ACL 2020 papers on topics such as

NLP ethics. Note that using statistics from ACL papers alone could be limiting because researchers in academia typically present almost all research achievements through publications, but many industry researchers do not publish in public venues such as ACL, although their research may impact various products.