

What Will it Take to Fix Benchmarking in Natural Language Understanding?

Samuel R. Bowman
New York University
bowman@nyu.edu

George E. Dahl
Google Research, Brain Team
gdahl@google.com

Abstract

Evaluation for many natural language understanding (NLU) tasks is broken: Unreliable and biased systems score so highly on standard benchmarks that there is little room for researchers who develop better systems to demonstrate their improvements. The recent trend to abandon IID benchmarks in favor of adversarially-constructed, out-of-distribution test sets ensures that current models will perform poorly, but ultimately only obscures the abilities that we want our benchmarks to measure. In this position paper, we lay out four criteria that we argue NLU benchmarks should meet. We argue most current benchmarks fail at these criteria, and that adversarial data collection does not meaningfully address the causes of these failures. Instead, restoring a healthy evaluation ecosystem will require significant progress in the design of benchmark datasets, the reliability with which they are annotated, their size, and the ways they handle social bias.

1 Introduction

A large and impactful thread of research on natural language understanding (NLU) has focused on improving results on benchmark datasets that feature roughly independent and identically distributed (IID) training, validation, and testing sections, drawn from data that were collected or annotated by crowdsourcing (Maas et al., 2011; Bowman et al., 2015; Rajpurkar et al., 2016; Wang et al., 2019b). Recent methodological progress combined with longstanding issues in crowdsourced data quality has made it so state-of-the-art systems are nearing the maximum achievable values on most of these benchmarks and thus are unlikely to be able to measure further improvements (Devlin et al., 2019; Raffel et al., 2020). At the same time, these apparently high-performing systems have serious known issues and have *not* achieved human-level competence at their tasks (Ribeiro et al., 2020).

1. Good performance on the benchmark should imply robust in-domain performance on the task.
↪ *We need more work on dataset design and data collection methods.*
2. Benchmark examples should be accurately and unambiguously annotated.
↪ *Test examples should be validated thoroughly enough to remove erroneous examples and to properly handle ambiguous ones.*
3. Benchmarks should offer adequate statistical power.
↪ *Benchmark datasets need to be much harder and/or much larger.*
4. Benchmarks should reveal plausibly harmful social biases in systems, and should not incentivize the creation of biased systems.
↪ *We need to better encourage the development and use auxiliary bias evaluation metrics.*

Figure 1: A summary of the criteria we propose.

Progress suffers in the absence of a trustworthy metric for benchmark-driven work: Newcomers and non-specialists are discouraged from trying to contribute, and specialists are given significant freedom to cherry-pick ad-hoc evaluation settings that mask a lack of progress (Church and Hestness, 2019).

The plight of benchmark-driven NLU research has prompted widespread concern about the assumptions underlying standard benchmarks and widespread interest in alternative models of evaluation. As an especially clear example, the documentation for the recent DynaBench benchmark suite argues that “benchmarks saturate”, “benchmarks have artifacts”, “researchers overfit on benchmarks”, and “benchmarks can be deceiving” and use these claims to motivate abandoning the IID paradigm in favor of benchmark data that is collected adversarially by asking a broad population of annotators to try to *fool* some reference neural network model.¹

¹<https://dynabench.org/about>

The DynaBench approach falls into the broader category of *adversarial filtering* (Paperno et al., 2016; Zellers et al., 2018; Nie et al., 2020; Le Bras et al., 2020). Adversarial filtering starts with a pipeline that produces candidate examples for the task, often through crowdsourcing, and then constructs a dataset by selecting those examples from the pipeline where one or more machine learning models fails to predict the correct label. This approach is appealing in that it guarantees that, at least in the short term, existing approaches to dataset construction can be patched to keep producing data that will challenge current systems.

However, collecting examples on which current models fail is neither necessary nor sufficient to create a useful benchmark. Among other points of concern, this approach can create a counterproductive incentive for researchers to develop models that are different without being better, since a model can top the leaderboard either by producing fewer errors than the adversary or by simply producing *different* errors, because the examples on which these new errors would be tested will not appear in the evaluation set. One could attempt to do this by, for example, pretraining new models that deliberately avoid any data that was used to pretrain the original adversary model, in order to minimize the degree to which the idiosyncratic mistakes of the new model line up with those of the old one. This incentive can slow progress and contribute to spurious claims of discovery.

This position paper argues that concerns about standard benchmarks that motivate methods like adversarial filtering are justified, but that they can and should be addressed directly, and that it is possible and reasonable to do so in the context of static, IID evaluation. We propose four criteria that adequate benchmarks should satisfy: benchmarks should offer a valid test of the full set of relevant language phenomena, they should be built around consistently-labeled data, they should offer adequate statistical power, and they should disincentivize the use of systems with potentially harmful biases. We then briefly survey some ongoing or promising research directions that could enable us to meet these challenges, including hybrid data collection protocols involving both crowdworkers and domain experts, larger-scale data validation, and auxiliary bias metric datasets attached to benchmarks.

2 Background

The Problem Performance on popular benchmarks is extremely high, but experts can easily find issues with high-scoring models. The GLUE benchmark (Wang et al., 2019b; Nangia and Bowman, 2019), a compilation of NLU evaluation tasks, has seen performance on its leaderboard approach or exceed human performance on all nine of its tasks. The follow-up SuperGLUE benchmark project (Wang et al., 2019a) solicited dataset submissions from the NLP research community in 2019, but wound up needing to exclude the large majority of the submitted tasks from the leaderboard because the BERT model (Devlin et al., 2019) was *already* showing performance at or above that of a majority vote of human crowdworkers. Of the eight tasks for which BERT did poorly enough to leave clear headroom for further progress, all are now effectively saturated (Raffel et al., 2020; He et al., 2020). State-of-the-art performance on the highly popular SQuAD 2 English reading-comprehension leaderboard (Rajpurkar et al., 2018) has long exceeded that of human annotators.

Ample evidence has emerged that the systems that have topped these leaderboards can fail dramatically on simple test cases that are meant to test the very skills that the leaderboards focus on (McCoy et al., 2019; Ribeiro et al., 2020). This result makes it clear that our systems have significant room to improve. However, we have no guarantee that our benchmarks will detect these needed improvements when they’re made. Most were collected by crowdsourcing with relatively limited quality control, such that we have no reason to expect that perfect performance on their metrics is achievable or that the benchmark will meaningfully distinguish between systems with superhuman metric performance. While the true upper bound on performance for any task (Bayes error) is not measurable, the fact that our systems have exceeded serious estimates of human performance leaves us with no reason to expect there to be much more headroom.

In addition, many of our best models display socially-relevant biases that render them inappropriate for deployment in many applications.² Our best current benchmarks do little or nothing to dis-

²The state-of-the-art T5 model, for example, shows far more sensitivity to irrelevant gender information than humans do when making coreference judgments, according to results on the SuperGLUE leaderboard with the DNC Winogender dataset (Rudinger et al., 2018; Poliak et al., 2018).

courage harmful biases and, by building largely on crowdsourced or naturally-occurring text data, they likely *incentivize* the development of models that reproduce problematic biases, at least to some degree.

The Goal This paper lays out four criteria that we would like our benchmarks to satisfy in order to facilitate further progress toward a primarily scientific goal: *building machines that can demonstrate a comprehensive and reliable understanding of everyday natural language text in the context of some specific well-posed task, language variety, and topic domain.* Among language understanding tasks, we focus on those that use labeled data and that are designed to test relatively general language understanding skills, for which the design of benchmarks can be especially difficult.

We distinguish between a *task* and a *benchmark*: A *task*, in our terms, is a language-related skill or competency that we want a model to demonstrate in the context of a specific input–output format. A *benchmark* attempts to evaluate performance on a task by grounding it to a text domain and instantiating it with a concrete dataset and evaluation metric. As a rough example, multiple-choice reading-comprehension question answering is a task, which the Cosmos benchmark (Huang et al., 2019) attempts to test using an accuracy metric over a specific sample of passages and questions from the English personal narrative domain. There is no general way to prove that a concrete benchmark faithfully measures performance on an abstract task. Nevertheless, since we can only evaluate models on concrete benchmarks, we have no choice but to strengthen the correspondence between the two as best we can.

We set aside the evaluation of computational efficiency and data efficiency, despite its relevance to many specific *applications* of language technology. We will *not* fully set aside issues of social bias. Even though it is possible for the same system to demonstrate both adept language understanding and harmful social prejudices,³ ethical concerns prompt us to argue that community-wide benchmarks should identify and disincentivize potentially harmful biases in models. The widespread sharing of trained models among NLU researchers and en-

³The performance of models like RoBERTa (Liu et al., 2019) or T5 (Raffel et al., 2020) on benchmarks like SuperGLUE that include some coverage of social bias is a good example of this, and typical human behavior is an even better example.

gineers and the fast pace of NLP R&D work mean that it is easy for systems designed with scientific goals in mind to be deployed in settings where their biases can cause real harm. While recent initiatives around data documentation should reduce the accidental deployment of models built on inappropriate data (Bender and Friedman, 2018; Gebru et al., 2018), we see room to do more.

We will also set aside few-shot learning, in which tasks are made artificially difficult by training models only on small subsets of the available training data (as was prominently used for GPT-3 by Brown et al., 2020). This paper focuses instead on the case where one is interested in reaching excellent performance on some language task and is willing to collect data or otherwise expend resources to make that possible. While few-shot learning represents a potentially impactful direction for engineering research, and *success* on some task in a few-shot setting is clear evidence of success more generally, artificial constraints on the use of training data do not fit the broad goals laid out above and do not fit many applied settings.

3 Four Challenges

This paper focuses on four criteria, outlined in Figure 1, that we argue effective future benchmarks for NLU tasks should satisfy. We believe that no current benchmark for any difficult broad-domain NLU task satisfies all four:

3.1 Validity

If one system significantly outperforms another on some benchmark, then that result should be strong evidence that the higher-scoring system is actually better at the task tested by the benchmark. In other words, benchmarks are only useful for language understanding research if they evaluate language understanding. General-purpose benchmarks that are designed to cover tasks like paragraph reading comprehension over Wikipedia are only effective if they test the full range of skills that are required to understand and reason about paragraphs from Wikipedia.

This criterion is difficult to fully formalize, and we know of no simple test that would allow one to determine if a benchmark presents a valid measure of model ability. Minimally, though, it requires the following:

- An evaluation dataset should reflect the full range of linguistic variation—including words

and higher-level constructions—that is used in the relevant domain, context, and language variety.

- An evaluation dataset should have a plausible means by which it tests all of the language-related behaviors that we expect the model to show in the context of the task.
- An evaluation dataset should be sufficiently free of annotation artifacts (as in Si et al., 2019; Sugawara et al., 2020b; Niven and Kao, 2019) that a system cannot reach near-human levels of performance by any means other than demonstrating the required language-related behaviors.

If a benchmark fully meets this challenge, we should expect any clear improvement on the benchmark to translate to similar improvements on *any other* valid and reasonable evaluation data for the same task and language domain.⁴

The rest of this section surveys common paradigms for constructing a benchmark dataset, and points to reasons that none offers a straightforward way to satisfy this criterion:

Naturally-Occurring Examples It is intuitively appealing to, where possible, build benchmark datasets based on naturally-occurring data distributions. This minimizes our effort in creating benchmarks and minimizes the risk that the benchmark is somehow skewed in a way that omits important phenomena. However, this is often not viable.

For tasks like reading comprehension or natural language inference that require multiple related texts (such as a passage and a question) as input, there is often no natural distribution that efficiently isolates the relevant task behaviors. One can find naturally-occurring distributions over questions, like those used to construct Natural Questions (Kwiatkowski et al., 2019), but these will generally be tied to the use contexts of a specific NLP product and will thus be limited by users’ perceptions of the current abilities of that product.

Even for single-input tasks like coreference resolution or Cloze, for which any text corpus can be the basis for a benchmark, naturalistic distributions do nothing to separate skills of interest from factual world knowledge and can be overwhelmingly dominated by the latter, making them poor

⁴Though, of course, any model with non-zero test error could be presented with a potentially-*un*reasonable benchmark entirely consisting of its own test errors.

metrics for incremental progress on NLU. Credible existing NLU-oriented benchmarks for such tasks are generally heavily curated (Paperno et al., 2016; Levesque et al., 2012; Sakaguchi et al., 2019).

Expert-Authored Examples Expert-constructed datasets for language understanding like FraCaS (Cooper et al., 1996) and the Winograd Schema Challenge (Levesque et al., 2012) have been crucial for defining several new tasks and introducing them as objects of study. However, expert example construction isn’t desirable for the creation of benchmarks for the use cases we focus on here.

Setting aside the logistical challenges of creating sufficiently large and diverse datasets by expert labor alone, expert authorship generally gives members of the research community direct, fine-grained control over the data on which their systems will be evaluated. Intentionally or unintentionally, this can produce data that is oriented toward linguistic phenomena that are widely studied and widely known to be important to the task at hand. While this can be helpful when building diagnostic datasets that focus on specific types of model failure (Cooper et al., 1996; Naik et al., 2018; Wang et al., 2019b), it is counterproductive when our goal is to build a broad-coverage benchmark dataset to set priorities and guide progress toward the solution of some task.

Dunietz et al. (2020) and Sugawara et al. (2020a) work around this issue by leaning on taxonomies of required phenomena from outside NLP. This is a direction worth pursuing, but it is not clear that appropriate taxonomies will be available for most NLU tasks of interest, or that these taxonomies will be broad and thorough enough to be straightforwardly implemented as datasets.

Crowdsourcing Most recent benchmarks for language understanding have been collected, at least in part, through crowdsourcing example construction, where non-expert annotators are given some freedom to construct examples based on a simple set of guidelines. This has an obvious appeal: Using non-expert annotators significantly lowers costs and using simple guidelines significantly reduces the risk that the resulting data will be skewed artificially toward phenomena of interest to experts.

However, straightforward standard practice, as was used to collect datasets like SNLI (Bowman et al., 2015) and SQuAD, seem to be relatively poor

at producing difficult datasets that test the intended phenomena. Existing datasets focus heavily on repetitive, easy cases and often fail to isolate key behaviors (Jia and Liang, 2017; Tsuchiya, 2018; McCoy et al., 2019).

Adversarial Filtering Given a source of examples and a model, adversarial-filtering-style approaches build a benchmark based on samples from that source for which the model fails. Adversarial filtering can remove examples that are easy due to trivial artifacts, but it does not ensure that the resulting dataset supports a valid test of model ability, and it can systematically eliminate coverage of linguistic phenomena or skills that are necessary for the task but already well-solved by the adversary model. This mode-seeking (as opposed to mass covering) behavior by adversarial filtering, if left unchecked, tends to reduce dataset diversity and thus make validity harder to achieve.

In contrast with this benchmark data collection setting, adversarial *competitions*, in which one compares the difficulty of collecting valid task examples that are adversarial to each of several systems, could be part of a healthy evaluation ecosystem. Such an ecosystem might involve frequent formative evaluations on a conventional non-adversarial benchmark in conjunction with periodic organized evaluations in an adversarial setting.

3.2 Reliable Annotation

For our benchmarks to incentivize the development of sound new methods, the labels for their test examples should be reliably correct. This means avoiding three failure cases: (i) examples that are carelessly mislabeled, (ii) examples that have no clear correct label due to unclear or underspecified task guidelines, and (iii) examples that have no clear correct label under the relevant metric due to legitimate disagreements in interpretation among annotators. The first two cases straightforwardly compromise the validity of the benchmark, but the third is somewhat subtler.

Legitimate disagreement emerges when an example can be labeled in multiple ways depending on an annotator’s choice between reasonable interpretations of the text of an example. Such disagreements might stem from dialectal variants in the interpretation of words or constructions or different reasonable interpretations of the actual state of the world. As a toy example, consider the question: Does *Ed ate a burrito* entail *Ed ate a sandwich*?

While most US English speakers would likely answer *no*, many pedants and regulatory officials have argued for *yes* (Florestall, 2008).

When a benchmark contains many instances of this kind of legitimate disagreement, a machine learning model will be able to study a benchmark dataset’s training set for clues about typical human behavior that might allow it to perform *better* than any single human annotator. This effect could contribute to misleading reports of super-human performance on such benchmarks, where *human performance* reflects the behavior of humans who are reporting their own judgments, rather than attempting to predict the most frequently assigned label, as the model does. We observe evidence of this kind of ambiguity in existing benchmarks: For example, Pavlick and Kwiatkowski (2019) find that 20% of examples across several textual entailment datasets are significantly ambiguous, and Kwiatkowski et al. (2019) show that 36% of short answer annotations in Natural Questions differ significantly from the majority answer.

3.3 Statistical Power

Benchmark evaluation datasets should be large and discriminative enough to detect any qualitatively relevant performance difference between two models. This criterion introduces a trade-off: If we can create benchmark datasets that are both reliable and highly difficult for the systems that we want to evaluate, then moderate dataset sizes will suffice. However, if our benchmark datasets contain many examples that are easy for current or near-future systems, then we will need dramatically larger evaluation sets to reach adequate power.

In the context of a reliable dataset that is difficult for current systems, a 1% absolute accuracy improvement, such as that from 80% to 81%, may be an acceptable minimum detectable effect. In this case, an evaluation set of a few thousand examples would suffice under typical conditions seen in NLU (Card et al., 2020). Many, though not all, popular benchmark datasets satisfy this size threshold.

Since our systems continue to improve rapidly, though, we should expect to be spending more time in the long tail of our data difficulty distributions: If we build reliable datasets, much of their future value may lie in their ability to measure improvements in accuracy among highly accurate systems. For example, an improvement from 98% accuracy to 98.1% represents the same 5% *relative* improve-

ment as we saw from 80% to 81%. To reliably detect this smaller *absolute* improvement, though, requires two orders of magnitude more evaluation data (Card et al., 2020).

3.4 Disincentives for Biased Models

A benchmark should, in general, favor a model without socially-relevant biases over an otherwise equivalent model with such biases. Many current benchmarks fail this test. Because benchmarks are often built around naturally-occurring or crowd-sourced text, it is often the case that a system can improve its performance by adopting heuristics that reproduce potentially-harmful biases (Rudinger et al., 2017). Developing adequate methods to minimize this effect will be challenging, both because of deep issues with both the precise specification of what constitutes harmful bias and because of the limited set of tools that we have available to us.

There is no precise enumeration of social biases that will be broadly satisfactory across applications and cultural contexts. This can be most easily illustrated with the example of biased associations between word representations for US English (as in Bolukbasi et al., 2016). Associations between race or gender and occupation are generally considered to be undesirable and potentially harmful in most contexts, and are something that benchmarks for word representations should discourage, or at least carefully avoid rewarding. If a set of word representations encodes typically Black female names like *Keisha* as being less similar to professional occupation terms like *lawyer* or *doctor* than typically white male names like *Scott* are, then a model using those representations is likely to reinforce harmful race or gender biases in any downstream content moderation systems or predictive text systems it gets used in.

Adequately enumerating the social attributes for which we might want to evaluate bias in some context can be difficult. For example, Indian castes, like racial categories in the United States, are often signaled by names and are an axis on which managers sometimes discriminate in hiring. Caste is a salient category of social bias in India that is subject to legal and institutional recognition. However, this bias also arises in some cases within the United States, where it has no such recognition (Tiku, 2020), and where it could be easily overlooked by non-specialist bias researchers.

Furthermore, building such a list of attributes

is also deeply political. Within living memory, popular and legal attitudes have changed significantly in the United States about attributes like race, gender, gender expression, sexual orientation, and disability. Attitudes on these issues continue to change, and new categories can gain recognition and protection over time. In many cases, this means that choosing whether to include some attribute in a computational metric of bias means choosing which group of people to align oneself with on a political issue. While there are clear ethical rules of thumb to follow when doing so,⁵ making any particular choice is nonetheless likely to put researchers in conflict with established institutions in ways that can change quickly. Any strategy for handling bias in the context of NLP benchmarks will have to grapple with this difficult reality.

4 Sketching a Solution

Building new benchmarks that improve upon our four axes is likely to be quite difficult. Below we attempt to sketch out some possible directions for improvement along each axis.

4.1 Improving Validity

Building valid benchmarks will require significant new research into data collection methods, at least some of which will be specific to the task under study. We suspect that much of this work will involve improvements in crowdsourcing and the use of non-experts, as most of the annotation behind the tasks we discuss requires no expertise other than fluent knowledge of the language variety under study.

One promising direction involves methods that start from relatively high-quality crowdsourced datasets, then use expert effort to augment them in ways that mitigate annotation artifacts. The Build-it-Break-it challenge (Ettinger et al., 2017), the Open Reading Benchmark (Dua et al., 2019), and the Gardner et al. (2020) contrast sets, among their other features, allow expert annotators to add examples to a test set to fill perceived gaps in coverage or correct perceived artifacts in a starting set of crowdsourced examples. To the extent that crowdsourcing with non-experts can produce data that has broad coverage and high difficulty but retains some measurable artifacts or flaws, this compro-

⁵The ACM code of ethics states, “when the interests of multiple groups conflict, the needs of those less advantaged should be given increased attention and priority.”

mise approach may help to create usable benchmark datasets out of the results.

Another approach brings computational linguists directly into the crowdsourcing process. This was recently demonstrated at a small scale by [Hu et al. \(2020\)](#) with OCNLI: They show that it is possible to significantly improve data quality issues by making small interventions during the crowdsourcing process—like offering additional bonus payments for examples that avoid overused words and constructions—without significantly limiting annotators’ freedom to independently construct creative examples.

Of course, implementing interventions like these in a way that offers convincing evidence of validity will be difficult.

4.2 Improving Handling of Annotation Errors and Disagreements

The use of standard techniques from crowdsourcing—generally involving multiple redundant annotations for each example—can largely resolve the issue of mistaken annotations. Careful planning and pilot work before data collection can largely resolve the issue of ambiguous annotation guidelines. Handling legitimate annotator disagreements can take two fairly different approaches, depending on the goals of the benchmark.

The simplest approach treats ambiguously labeled examples in the same way as mislabeled examples, and systematically identifies and discards them during a validation phase. For some tasks, it may still be possible to test models’ handling of fundamentally ambiguous linguistic phenomena or domains using *unambiguous* examples: In the case of multiple-choice question answering, for example, one can construct examples where one answer candidate is only debatably correct, but all other candidates are unequivocally wrong. Any sound model would then be expected to select the debatable choice.

Alternately, one can decline to assign single, discrete labels to ambiguous examples. This can involve asking models to predict the empirical *distribution* of labels that trustworthy annotators assign ([Pavlick and Kwiatkowski, 2019](#); [Poesio et al., 2019](#)), or allowing models to predict any of several answer choices that are supported by trustworthy annotators (as in the SQuAD benchmark). This comes at the cost, though, of requiring many more

annotator judgments per evaluation example.

4.3 Improving Statistical Power

In principle, achieving adequate statistical power is straightforward: we simply estimate the number of examples required to reach the desired statistical power for any plausible short-to-medium term system evaluation for the task, and collect that number of examples. In practice, however, costs can become prohibitive.

For a relatively simple task like NLI, labeling an existing example likely requires a bare minimum of 45 seconds ([Vania et al., 2020](#)), and creating a new example requires at least one minute ([Bowman et al., 2020](#)). Even if we use these very optimistic numbers to estimate annotation speed, a ten-way-annotated dataset of 500,000 examples will still cost over \$1 million at a \$15/hr pay rate.⁶ Recruiting more experienced annotators or encouraging annotators to work more carefully could increase this figure dramatically. While such an amount of money is not completely out of reach in a well-funded field like NLP,⁷ investments of this kind will inevitably be rare enough that they help reinforce the field’s concentration of data and effort on a few high-resource languages and tasks.

For settings in which large datasets are necessary, we see no clear way to avoid high costs. Gamification, in the style of the ESP game or ZombiLingo ([Von Ahn and Dabbish, 2004](#); [Fort et al., 2014](#)), promises to offer free human labor, but at the cost of the expert time needed to refine the task definition into a game that is widely enjoyable. This approach also introduces severe constraints on the kinds of data collection protocols that can be used and raises tricky new ethical issues ([Morschheuser and Hamari, 2019](#)). Ultimately, the community needs to compare the cost of making serious investments in better benchmarks to the cost of wasting researcher time and computational resources due to our inability to measure progress.

4.4 Disincentives for Biased Models

Because there is no one-size-fits-all definition of harmful social bias, there is little prospect of creat-

⁶This figure ignores platform fees and makes the additional optimistic assumption that only 10% of fully-annotated examples will be discarded because of annotator disagreement.

⁷To put this number in context, public estimates of the cost of OpenAI’s GPT-3 ([Brown et al., 2020](#)) exceed \$10M ([Wiggers, 2020](#)), and in machine translation, [Meng et al. \(2019\)](#)’s use of 512 Nvidia V100 GPUs for three months would have cost over \$1M USD on commodity cloud infrastructure.

ing a benchmark for language understanding that is guaranteed to never reward the development of harmfully biased models. **This is not a compelling reason to accept the status quo**, and we nonetheless have a clear opportunity to mitigate some of the potential harms caused by applied NLP systems before those systems are even developed. Opting not to test models for some plausible and potentially-harmful social bias is, intentionally or not, a political choice.

While it would be appealing to try to guarantee that our evaluation data does not itself demonstrate evidence of bias, we are aware of no robust strategy for reliably accomplishing this, and work on the closely-related problem of *model* bias mitigation has been fraught with false starts and overly optimistic claims (Gonen and Goldberg, 2019).

A viable alternate approach could involve the expanded use of auxiliary metrics: Rather than trying to fully mitigate bias within a single general dataset and metric for some task, benchmark creators can introduce a family of additional expert-constructed test datasets and metrics that each isolate and measure a specific type of bias. Any time a model is evaluated on the primary task test set in this setting, it would be evaluated in parallel on these additional bias test sets. This would not prevent the primary metric from unintentionally and subtly rewarding biased models, but it would combat this effect by more directly highlighting and penalizing bias in models. In addition, the fact that these metrics would target specific types of biases would make it easier for benchmark maintainers to adapt as changing norms or changing downstream applications demand coverage of additional potential harms.

For several tasks, metrics like this already exist, at least for gender in English, in the form of auxiliary test sets meant to be combined with a pre-existing training set (Rudinger et al., 2018; Webster et al., 2018; Kiritchenko and Mohammad, 2018; Li et al., 2020). Even so, refining these metrics and developing new ones will likely require us to face many of the same challenges that we highlight in this paper for benchmark design more generally.

The larger challenge in implementing this approach, however, is a matter of community structure and incentive design. Methods papers dealing with tasks for which metrics already exist rarely report numbers on these metrics. Even for the SuperGLUE benchmark, which *requires* users to com-

pute test set metrics on the DNC Winogender test set in order to reveal test set results for any other target task, a large majority of papers that report test set numbers omit this metric and decline to report potentially unflattering bias numbers (Raffel et al., 2020; Pruksachatkun et al., 2020; Schick and Schütze, 2020; He et al., 2020).

The difficulty, then, is in developing community infrastructure to encourage the widespread reporting of metrics that address the full range of relevant likely harms. This could plausibly involve peer review norms, explicit publication venue policies, stricter versions of the SuperGLUE approach for which users can *only* retrieve aggregate performance numbers, without a precise separation of the primary and bias-oriented metrics, or even the introduction of professional licensing standards.

Of course, ensuring that bias is measured and reported is not enough to prevent bias-related harms from emerging in practice: It is also necessary to ensure that those who build and deploy NLP products will take these metrics seriously and respond to them appropriately. And, of course, even if a system encodes no social bias at all, it can still be deployed in ways that produce unfair or unjust outcomes. These difficult issues are beyond the scope of a paper on benchmark design.

5 Related Work

The NLP and ML research communities are increasingly interested in issues surrounding data and evaluation. This section surveys relevant positions and issues that don't quite fit our schema.

Welty et al. (2019) advocate for the more precise reporting of the focus and abilities of test sets and metrics in ML broadly, with a focus on issues surrounding statistical power. Bender and Friedman (2018) and Gebu et al. (2018) advocate for explicit freestanding datasheets documenting dataset releases of all kinds, with a focus on making potential harmful mismatches between data and application visible, and Hutchinson et al. (2021) argue along similar lines for a broader program of transparency and stakeholder engagement in data creation. Dodge et al. (2019) lay out a set of best practices for results reporting, with a focus on the impact of hyperparameter tuning on model comparison. Ethayarajh and Jurafsky (2020) advocate for the inclusion of efficiency considerations in leaderboard design. Boyd-Graber and Börschinger (2020) describe ways that trivia competitions can

provide a model for carefully-considered dataset design.

Church and Hestness (2019) revisit the arguments that motivated the NLP community’s shift toward quantitative benchmarking in the early 1990s and warn that the overwhelming success of this shift has indirectly laid the groundwork for the widespread use of poor-quality benchmarks. Blodgett et al. (2020) challenge researchers working on social bias in NLP to focus more precisely on specific types of harm to specific populations of users, a challenge that our broad position piece does not fully meet.

NLP has had longstanding debates over the types of tasks that best test substantial language understanding skills. Many task-specific papers contribute to this debate, as does a prominent recent thread advocating for an increased focus on grounding of various kinds by Bender and Koller (2020), Bisk et al. (2020), Zellers et al. (2020), and others.

6 Conclusion

Benchmarking for NLU is broken. We lay out four major criteria that benchmarks should fulfill to offer faithful, useful, and responsible measures of language ability. We argue that departing from IID evaluation (as is seen with benchmark datasets collected by adversarial filtering) does not help to address these criteria, but lay out in broad strokes how each criterion might be addressed directly.

Nonetheless, important open research questions remain. Most centrally, it is still unclear how best to integrate expert effort into crowdsourced data collection, and we do not yet see a clear institutional model by which to ensure that bias metrics are built and used when they are most needed.

Ethical Considerations

This paper advocates for reforms to a set of benchmarking practices that have so far largely failed to address issues of social bias, and that have thereby helped create a false sense of security among those building applied systems. While this paper offers no complete and satisfactory solutions, it proposes measures that should contribute to harm reduction.

Acknowledgments

We thank Emily Bender, Iacer Calixto, Haokin Liu, Kyunghyun Cho, Will Huang, Jamie Kiros, and audiences at CMU, Google, and Apple for feedback on these ideas.

This project has benefited from financial support to SB by Eric and Wendy Schmidt (made by recommendation of the Schmidt Futures program), Samsung Research (under the project *Improving Deep Learning using Latent Structure*), and Intuit. This material is based upon work supported by the National Science Foundation under Grant No. 1922658. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- Emily M. Bender and Batya Friedman. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. [Experience grounds language](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8718–8735, 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.
- Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. [Man is to computer programmer as woman is to homemaker? Debiasing word embeddings](#). In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, page 4356–4364, Red Hook, NY, USA. Curran Associates Inc.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.

- Samuel R. Bowman, Jennimaria Palomaki, Livio Baldini Soares, and Emily Pitler. 2020. [New protocols and negative results for textual entailment data collection](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8203–8214, Online. Association for Computational Linguistics.
- Jordan Boyd-Graber and Benjamin Börschinger. 2020. [What question answering can learn from trivia nerds](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7422–7435, Online. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris 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*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Dallas Card, Peter Henderson, Urvashi Khandelwal, Robin Jia, Kyle Mahowald, and Dan Jurafsky. 2020. [With little power comes great responsibility](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9263–9274, Online. Association for Computational Linguistics.
- Kenneth Ward Church and Joel Hestness. 2019. [A survey of 25 years of evaluation](#). *Natural Language Engineering*, 25(6):753–767.
- Robin Cooper, Dick Crouch, Jan Van Eijck, Chris Fox, Johan Van Genabith, Jan Jaspars, Hans Kamp, David Milward, Manfred Pinkal, Massimo Poesio, et al. 1996. Using the framework. Technical report, LRE 62-051 D-16, The FraCaS Consortium.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jesse Dodge, Suchin Gururangan, Dallas Card, Roy Schwartz, and Noah A. Smith. 2019. [Show your work: Improved reporting of experimental results](#). 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 2185–2194, Hong Kong, China. Association for Computational Linguistics.
- Dheeru Dua, Ananth Gottumukkala, Alon Talmor, Sameer Singh, and Matt Gardner. 2019. [ORB: An open reading benchmark for comprehensive evaluation of machine reading for comprehension](#). *arXiv preprint 1912.12598*.
- Jesse Dunietz, Greg Burnham, Akash Bharadwaj, Owen Rambow, Jennifer Chu-Carroll, and Dave Ferrucci. 2020. [To test machine comprehension, start by defining comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7839–7859, Online. Association for Computational Linguistics.
- Kawin Ethayarajh and Dan Jurafsky. 2020. [Utility is in the eye of the user: A critique of NLP leaderboards](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4846–4853, Online. Association for Computational Linguistics.
- Allyson Ettinger, Sudha Rao, Hal Daumé III, and Emily M. Bender. 2017. [Towards linguistically generalizable NLP systems: A workshop and shared task](#). In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*, pages 1–10, Copenhagen, Denmark. Association for Computational Linguistics.
- Marjorie Florestall. 2008. Is a burrito a sandwich? Exploring race, class, and culture in contracts. *Michigan Journal of Race and Language*, 14:1.
- Karën Fort, Bruno Guillaume, and Hadrien Chastant. 2014. Creating Zombilingo, a game with a purpose for dependency syntax annotation. In *Proceedings of the First International Workshop on Gamification for Information Retrieval*, pages 2–6.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. [Evaluating models’ local decision boundaries via contrast sets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 1307–1323, Online. Association for Computational Linguistics.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2018. [Datasheets for datasets](#). *arXiv preprint 1803.09010*.
- Hila Gonen and Yoav Goldberg. 2019. [Lipstick on a pig: Debiasing methods cover up systematic gender](#)

- biases in word embeddings but do not remove them. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 609–614, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. DeBERTa: Decoding-enhanced BERT with disentangled attention. *arXiv preprint 2006.03654*.
- Hai Hu, Kyle Richardson, Liang Xu, Lu Li, Sandra Kübler, and Lawrence Moss. 2020. OCNLI: Original Chinese Natural Language Inference. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3512–3526, Online. Association for Computational Linguistics.
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. Cosmos QA: Machine reading comprehension with contextual commonsense reasoning. 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 2391–2401, Hong Kong, China. Association for Computational Linguistics.
- Ben Hutchinson, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, and Margaret Mitchell. 2021. Towards accountability for machine learning datasets: Practices from software engineering and infrastructure. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 560–575, New York, NY, USA. Association for Computing Machinery.
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.
- Svetlana Kiritchenko and Saif Mohammad. 2018. Examining gender and race bias in two hundred sentiment analysis systems. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 43–53, New Orleans, Louisiana. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew E Peters, Ashish Sabharwal, and Yejin Choi. 2020. Adversarial filters of dataset biases. *arXiv preprint 2002.04108*.
- Hector J. Levesque, Ernest Davis, and Leora Morgenstern. 2012. The Winograd schema challenge. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning, KR'12*, pages 552–561. AAAI Press.
- Tao Li, Daniel Khashabi, Tushar Khot, Ashish Sabharwal, and Vivek Srikumar. 2020. UNCOVERING stereotyping biases via underspecified questions. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3475–3489, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint 1907.11692*.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Yuxian Meng, Xiangyuan Ren, Zijun Sun, Xiaoya Li, Arianna Yuan, Fei Wu, and Jiwei Li. 2019. Large-scale pretraining for neural machine translation with tens of billions of sentence pairs. *arXiv preprint 1909.11861*.
- Benedikt Morschheuser and Juho Hamari. 2019. The gamification of work: Lessons from crowdsourcing. *Journal of Management Inquiry*, 28(2):145–148.
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. Stress test evaluation for natural language inference. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Nikita Nangia and Samuel Bowman. 2019. Human vs. muppet: A conservative estimate of human performance on the glue benchmark. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4566–4575.

- 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.
- Timothy Niven and Hung-Yu Kao. 2019. [Probing neural network comprehension of natural language arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. 2016. [The LAMBADA dataset: Word prediction requiring a broad discourse context](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1525–1534, Berlin, Germany. Association for Computational Linguistics.
- Ellie Pavlick and Tom Kwiatkowski. 2019. [Inherent disagreements in human textual inferences](#). *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Massimo Poesio, Jon Chamberlain, Silviu Paun, Juntao Yu, Alexandra Uma, and Udo Kruschwitz. 2019. [A crowdsourced corpus of multiple judgments and disagreement on anaphoric interpretation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1778–1789, Minneapolis, Minnesota. Association for Computational Linguistics.
- Adam Poliak, Aparajita Haldar, Rachel Rudinger, J. Edward Hu, Ellie Pavlick, Aaron Steven White, and Benjamin Van Durme. 2018. [Collecting diverse natural language inference problems for sentence representation evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 67–81, Brussels, Belgium. Association for Computational Linguistics.
- Yada Pruksachatkun, Jason Phang, Haokun Liu, Phu Mon Htut, Xiaoyi Zhang, Richard Yuanzhe Pang, Clara Vania, Katharina Kann, and Samuel R. Bowman. 2020. [Intermediate-task transfer learning with pretrained language models: When and why does it work?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5231–5247, Online. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Rachel Rudinger, Chandler May, and Benjamin Van Durme. 2017. [Social bias in elicited natural language inferences](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 74–79, Valencia, Spain. Association for Computational Linguistics.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. [Gender bias in coreference resolution](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 8–14, New Orleans, Louisiana. Association for Computational Linguistics.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [Winogrande: An adversarial winograd schema challenge at scale](#). *arXiv preprint 1907.10641*.
- Timo Schick and Hinrich Schütze. 2020. [It’s not just size that matters: Small language models are also few-shot learners](#). *arXiv preprint 2009.07118*.
- Chenglei Si, Shuohang Wang, Min-Yen Kan, and Jing Jiang. 2019. [What does BERT learn from multiple-choice reading comprehension datasets?](#) *arXiv preprint 1910.12391*.
- Saku Sugawara, Pontus Stenetorp, and Akiko Aizawa. 2020a. [Benchmarking machine reading comprehension: A psychological perspective](#). *arXiv preprint 2004.01912*.
- Saku Sugawara, Pontus Stenetorp, Kentaro Inui, and A. Aizawa. 2020b. [Assessing the benchmarking capacity of machine reading comprehension datasets](#). In *AAAI*.
- Nitasha Tiku. 2020. [India’s engineers have thrived in silicon valley. so has its caste system](#). *The Washington Post*.

- Masatoshi Tsuchiya. 2018. [Performance impact caused by hidden bias of training data for recognizing textual entailment](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Clara Vania, Ruijie Chen, and Samuel R. Bowman. 2020. [Asking Crowdworkers to Write Entailment Examples: The Best of Bad options](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 672–686, Suzhou, China. Association for Computational Linguistics.
- Luis Von Ahn and Laura Dabbish. 2004. Labeling images with a computer game. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 319–326.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019a. SuperGLUE: A stickier benchmark for general-purpose language understanding systems. In *Advances in Neural Information Processing Systems*, pages 3266–3280.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019b. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the International Conference on Learning Representations*.
- Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. 2018. Mind the gap: A balanced corpus of gendered ambiguous pronouns. *Transactions of the Association for Computational Linguistics*, 6:605–617.
- Chris Welty, Praveen Paritosh, and Lora Aroyo. 2019. Metrology for AI: From benchmarks to instruments. *arXiv preprint 1911.01875*.
- Kyle Wiggers. 2020. [OpenAI launches an API to commercialize its research](#). *VentureBeat*.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. [SWAG: A large-scale adversarial dataset for grounded commonsense inference](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Elizabeth Clark, Lianhui Qin, Ali Farhadi, and Yejin Choi. 2020. Evaluating machines by their real-world language use. *arXiv preprint 2004.03607*.