

Dynabench: Rethinking Benchmarking in NLP

Douwe Kiela[†], Max Bartolo[‡], Yixin Nie^{*}, Divyansh Kaushik[§], Atticus Geiger[¶],

Zhengxuan Wu[¶], Bertie Vidgen^{||}, Grusha Prasad^{**}, Amanpreet Singh[†], Pratik Ringshia[†],

Zhiyi Ma[†], Tristan Thrush[†], Sebastian Riedel^{†‡}, Zeerak Waseem^{††}, Pontus Stenetorp[‡],

Robin Jia[†], Mohit Bansal^{*}, Christopher Potts[¶] and Adina Williams[†]

[†] Facebook AI Research; [‡] UCL; ^{*} UNC Chapel Hill; [§] CMU; [¶] Stanford University

^{||} Alan Turing Institute; ^{**} JHU; ^{††} Simon Fraser University

dynabench@fb.com

Abstract

We introduce Dynabench, an open-source platform for dynamic dataset creation and model benchmarking. Dynabench runs in a web browser and supports human-and-model-in-the-loop dataset creation: annotators seek to create examples that a target model will misclassify, but that another person will not. In this paper, we argue that Dynabench addresses a critical need in our community: contemporary models quickly achieve outstanding performance on benchmark tasks but nonetheless fail on simple challenge examples and falter in real-world scenarios. With Dynabench, dataset creation, model development, and model assessment can directly inform each other, leading to more robust and informative benchmarks. We report on four initial NLP tasks, illustrating these concepts and highlighting the promise of the platform, and address potential objections to dynamic benchmarking as a new standard for the field.

1 Introduction

While it used to take decades for machine learning models to surpass estimates of human performance on benchmark tasks, that milestone is now routinely reached within just a few years for newer datasets (see Figure 1). As with the rest of AI, NLP has advanced rapidly thanks to improvements in computational power, as well as algorithmic breakthroughs, ranging from attention mechanisms (Bahdanau et al., 2014; Luong et al., 2015), to Transformers (Vaswani et al., 2017), to pre-trained language models (Howard and Ruder, 2018; Devlin et al., 2019; Liu et al., 2019b; Radford et al., 2019; Brown et al., 2020). Equally important has been the rise of benchmarks that support the development of ambitious new data-driven models and that encourage apples-to-apples model comparisons. Benchmarks provide a north star goal for researchers, and




Figure 1: Benchmark saturation over time for popular benchmarks, normalized with initial performance at minus one and human performance at zero.

are part of the reason we can confidently say we have made great strides in our field.

In light of these developments, one might be forgiven for thinking that NLP has created models with human-like language capabilities. Practitioners know that, despite our progress, we are actually far from this goal. Models that achieve super-human performance on benchmark tasks (according to the narrow criteria used to define human performance) nonetheless fail on simple challenge examples and falter in real-world scenarios. A substantial part of the problem is that our benchmark tasks are not adequate proxies for the sophisticated and wide-ranging capabilities we are targeting: they contain inadvertent and unwanted statistical and social biases that make them artificially easy and misaligned with our true goals.

We believe the time is ripe to radically rethink benchmarking. In this paper, which both takes a position and seeks to offer a partial solution, we introduce Dynabench, an open-source, web-based research platform for dynamic data collection and model benchmarking. The guiding hypothesis be-

hind Dynabench is that we can make even faster progress if we evaluate models and collect data dynamically, with humans and models in the loop, rather than the traditional static way.

Concretely, Dynabench hosts tasks for which we dynamically collect data against state-of-the-art models in the loop, over multiple rounds. The stronger the models are and the fewer weaknesses they have, the lower their error rate will be when interacting with humans, giving us a concrete metric—i.e., how well do AI systems perform when interacting with humans? This reveals the shortcomings of state-of-the-art models, and it yields valuable training and assessment data which the community can use to develop even stronger models.

In this paper, we first document the background that led us to propose this platform. We then describe the platform in technical detail, report on findings for four initial tasks, and address possible objections. We finish with a discussion of future plans and next steps.

2 Background

Progress in NLP has traditionally been measured through a selection of task-level datasets that gradually became accepted benchmarks (Marcus et al., 1993; Pradhan et al., 2012). Recent well-known examples include the Stanford Sentiment Treebank (Socher et al., 2013), SQuAD (Rajpurkar et al., 2016, 2018), SNLI (Bowman et al., 2015), and MultiNLI (Williams et al., 2018). More recently, multi-task benchmarks such as SentEval (Conneau and Kiela, 2018), DecaNLP (McCann et al., 2018), GLUE (Wang et al., 2018), and SuperGLUE (Wang et al., 2019) were proposed with the aim of measuring general progress across several tasks. When the GLUE dataset was introduced, “solving GLUE” was deemed “beyond the capability of current transfer learning methods” (Wang et al., 2018). However, GLUE saturated within a year and its successor, SuperGLUE, already has models rather than humans at the top of its leaderboard. These are remarkable achievements, but there is an extensive body of evidence indicating that these models do not in fact have the human-level natural language capabilities one might be lead to believe.

2.1 Challenge Sets and Adversarial Settings

Whether our models have learned to solve tasks in robust and generalizable ways has been a topic

of much recent interest. Challenging test sets have shown that many state-of-the-art NLP models struggle with compositionality (Nie et al., 2019; Kim and Linzen, 2020; Yu and Ettinger, 2020; White et al., 2020), and find it difficult to pass the myriad stress tests for social (Rudinger et al., 2018; May et al., 2019; Nangia et al., 2020) and/or linguistic competencies (Geiger et al., 2018; Naik et al., 2018; Glockner et al., 2018; White et al., 2018; Warstadt et al., 2019; Gauthier et al., 2020; Hossain et al., 2020; Jeretic et al., 2020; Lewis et al., 2020; Saha et al., 2020; Schuster et al., 2020; Sugawara et al., 2020; Warstadt et al., 2020). Yet, challenge sets may suffer from performance instability (Liu et al., 2019a; Rozen et al., 2019; Zhou et al., 2020) and often lack sufficient statistical power (Card et al., 2020), suggesting that, although they may be valuable assessment tools, they are not sufficient for ensuring that our models have achieved the learning targets we set for them.

Models are susceptible to adversarial attacks, and despite impressive task-level performance, state-of-the-art systems still struggle to learn robust representations of linguistic knowledge (Ettinger et al., 2017), as also shown by work analyzing model diagnostics (Ettinger, 2020; Ribeiro et al., 2020). For example, question answering models can be fooled by simply adding a relevant sentence to the passage (Jia and Liang, 2017).

Text classification models have been shown to be sensitive to single input character change (Ebrahimi et al., 2018b) and first-order logic inconsistencies (Minervini and Riedel, 2018). Similarly, machine translation systems have been found susceptible to character-level perturbations (Ebrahimi et al., 2018a) and synthetic and natural noise (Belinkov and Bisk, 2018; Khayrallah and Koehn, 2018). Natural language inference models can be fooled by simple syntactic heuristics or hypothesis-only biases (Gururangan et al., 2018; Poliak et al., 2018; Tsuchiya, 2018; Belinkov et al., 2019; McCoy et al., 2019). Dialogue models may ignore perturbations of dialogue history (Sankar et al., 2019). More generally, Wallace et al. (2019) find universal adversarial perturbations forcing targeted model errors across a range of tasks. Recent work has also focused on evaluating model diagnostics through counterfactual augmentation (Kaushik et al., 2020), decision boundary analysis (Gardner et al., 2020; Swayamdipta et al., 2020), and behavioural testing (Ribeiro et al., 2020).

2.2 Adversarial Training and Testing

Research progress has traditionally been driven by a cyclical process of resource collection and architectural improvements. Similar to Dynabench, recent work seeks to embrace this phenomenon, addressing many of the previously mentioned issues through an iterative human-and-model-in-the-loop annotation process (Yang et al., 2017; Dinan et al., 2019; Chen et al., 2019; Bartolo et al., 2020; Nie et al., 2020), to find “unknown unknowns” (Attenberg et al., 2015) or in a never-ending or life-long learning setting (Silver et al., 2013; Mitchell et al., 2018). The Adversarial NLI (ANLI) dataset (Nie et al., 2020), for example, was collected with an adversarial setting over multiple rounds to yield “a ‘moving post’ dynamic target for NLU systems, rather than a static benchmark that will eventually saturate”. In its few-shot learning mode, GPT-3 barely shows “signs of life” (Brown et al., 2020) (i.e., it is barely above random) on ANLI, which is evidence that we are still far away from human performance on that task.

2.3 Other Related Work

While crowdsourcing has been a boon for large-scale NLP dataset creation (Snow et al., 2008; Munro et al., 2010), we ultimately want NLP systems to handle “natural” data (Kwiatkowski et al., 2019) and be “ecologically valid” (de Vries et al., 2020). Ethayarajh and Jurafsky (2020) analyze the distinction between what leaderboards incentivize and “what is useful in practice” through the lens of microeconomics. A natural setting for exploring these ideas might be dialogue (Hancock et al., 2019; Shuster et al., 2020). Other works have pointed out misalignments between maximum-likelihood training on i.i.d. train/test splits and human language (Linzen, 2020; Stiennon et al., 2020).

We think there is widespread agreement that something has to change about our standard evaluation paradigm and that we need to explore alternatives. The persistent misalignment between benchmark performance and performance on challenge and adversarial test sets reveals that standard evaluation paradigms overstate the ability of our models to perform the tasks we have set for them. Dynabench offers one path forward from here, by allowing researchers to combine model development with the stress-testing that needs to be done to achieve true robustness and generalization.

3 Dynabench

Dynabench is a platform that encompasses different *tasks*. Data for each task is collected over multiple *rounds*, each starting from the current state of the art. In every round, we have one or more *target models* “in the loop.” These models interact with humans, be they expert linguists or crowdworkers, who are in a position to identify models’ shortcomings by providing *examples* for an optional *context*. Examples that models get wrong, or struggle with, can be validated by other humans to ensure their correctness. The data collected through this process can be used to evaluate state-of-the-art models, and to train even stronger ones, hopefully creating a virtuous cycle that helps drive progress in the field. Figure 2 provides a sense of what the example creation interface looks like.

As a large-scale collaborative effort, the platform is meant to be a platform technology for human-and-model-in-the-loop evaluation that belongs to the entire community. In the current iteration, the platform is set up for dynamic *adversarial* data collection, where humans can attempt to find model-fooling examples. This design choice is due to the fact that the *average* case, as measured by maximum likelihood training on i.i.d. datasets, is much less interesting than the *worst* (i.e., adversarial) case, which is what we want our systems to be able to handle if they are put in critical systems where they interact with humans in real-world settings.

However, Dynabench is not limited to the adversarial setting, and one can imagine scenarios where humans are rewarded not for fooling a model or ensemble of models, but for finding examples that models, even if they are right, are very uncertain about, perhaps in an active learning setting. Similarly, the paradigm is perfectly compatible with collaborative settings that utilize human feedback, or even negotiation. The crucial aspect of this proposal is the fact that models and humans interact live “in the loop” for evaluation and data collection.

One of the aims of this platform is to put expert linguists center stage. Creating model-fooling examples is not as easy as it used to be, and finding interesting examples is rapidly becoming a less trivial task. In ANLI, the verified model error rate for crowd workers in the later rounds went below 1-in-10 (Nie et al., 2020), while in “Beat the AI”, human performance decreased while time per valid adversarial example went up with stronger models in the loop (Bartolo et al., 2020). For expert linguists, we

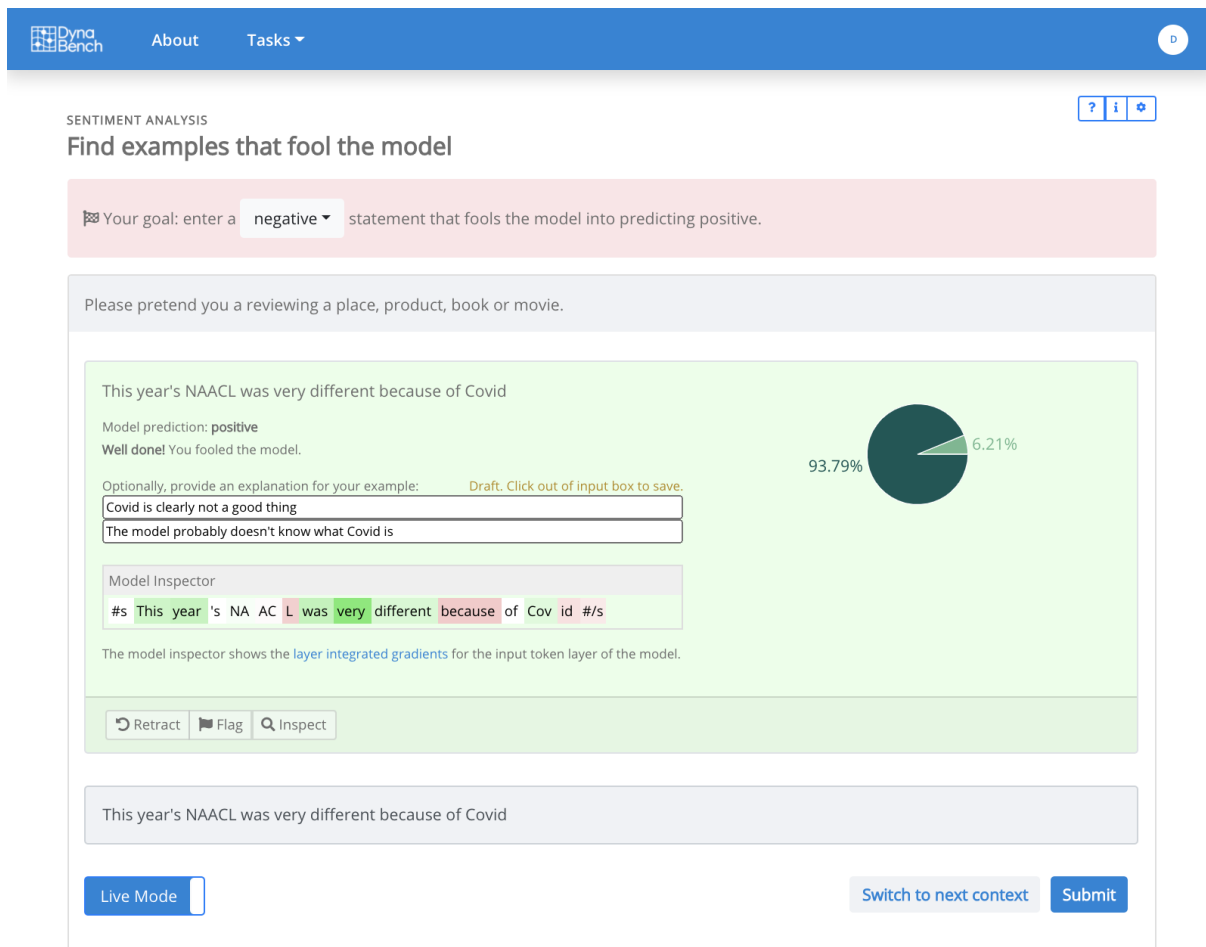


Figure 2: The Dynabench example creation interface for sentiment analysis with illustrative example.

expect the model error to be much higher, but if the platform actually lives up to its virtuous cycle promise, that error rate will go down quickly. Thus, we predict that linguists with expertise in exploring the decision boundaries of machine learning models will become essential.

While we are primarily motivated by evaluating progress, both ANLI and “Beat the AI” show that models can overcome some of their existing blind spots through adversarial training. They also find that best model performance is still quite far from that of humans, suggesting that while the collected data appears to lie closer to the model decision boundaries, there still exist adversarial examples beyond the remit of current model capabilities.

3.1 Features and Implementation Details

Dynabench offers low-latency, real-time feedback on the behavior of state-of-the-art NLP models. The technology stack is based on PyTorch (Paszke et al., 2019), with models served via TorchServe.¹

¹<https://pytorch.org/serve>

The platform not only displays prediction probabilities, but through an “inspect model” functionality, allows the user to examine the token-level layer integrated gradients (Sundararajan et al., 2017), obtained via the Captum interpretability library.²

For each example, we allow the user to *explain* what the correct label is, as well as why they think it fooled a model if the model got it wrong; or why the model might have been fooled if it wasn’t. All collected model-fooling (or, depending on the task, even non-model-fooling) examples are verified by other humans to ensure their validity.

Task owners can collect examples through the web interface, by engaging with the community, or through Mephisto,³ which makes it easy to connect, e.g., Mechanical Turk workers to the exact same backend. All collected data will be open sourced, in an anonymized fashion.

In its current mode, Dynabench could be de-

²<https://captum.ai/>

³<https://github.com/facebookresearch/Mephisto>

scribed as a fairly conservative departure from the status quo. It is being used to develop datasets that support the same metrics that drive existing benchmarks. The crucial change is that the datasets are now dynamically created, allowing for more kinds of evaluation—e.g., tracking progress through rounds and across different conditions.

3.2 Initial Tasks

We have selected four official tasks as a starting point, which we believe represent an appropriate cross-section of the field at this point in time. Natural Language Inference (NLI) and Question Answering (QA) are canonical tasks in the field. Sentiment analysis is a task that some consider “solved” (and is definitely treated as such, with all kinds of ethically problematic repercussions), which we show is not the case. Hate speech is very important as it can inflict harm on people, yet classifying it remains challenging for NLP.

Natural language inference. Built upon the semantic foundation of natural logic (Sánchez Valencia, 1991, i.a.) and hailing back much further (van Benthem, 2008), NLI is one of the quintessential natural language understanding tasks. NLI, also known as ‘recognizing textual entailment’ (Dagan et al., 2006), is often formulated as a 3-way classification problem where the input is a context sentence paired with a hypothesis, and the output is a label (entailment, contradiction, or neutral) indicating the relation between the pair.

We build on the ANLI dataset (Nie et al., 2020) and its three rounds to seed the Dynabench NLI task. During the ANLI data collection process, the annotators were presented with a context (extracted from a pre-selected corpus) and a desired target label, and asked to provide a hypothesis that fools the target model adversary into misclassifying the example. If the target model is fooled, the annotator was invited to speculate about why, or motivate why their example was right. The target model of the first round (R1) was a single BERT-Large model fine-tuned on SNLI and MNLI, while the target model of the second and third rounds (R2, R3) was an ensemble of RoBERTa-Large models fine-tuned on SNLI, MNLI, FEVER (Thorne et al., 2018) recast as NLI, and all of the ANLI data collected prior to the corresponding round. The contexts for Round 1 and Round 2 were Wikipedia passages curated in Yang et al. (2018) and the contexts for Round 3 were from various domains. Results indi-

cate that state-of-the-art models (which can obtain 90%+ accuracy on SNLI and MNLI) cannot exceed 50% accuracy on rounds 2 and 3.

With the launch of Dynabench, we have started collection of a fourth round, which has several innovations: not only do we select candidate contexts from a more diverse set of Wikipedia featured articles but we also use an ensemble of two different models with different architectures as target adversaries to increase diversity and robustness. Moreover, the ensemble of adversaries will help mitigate issues with creating a dataset whose distribution is too closely aligned to a particular target model or architecture. Additionally, we are collecting two types of natural language explanations: why an example is correct and why a target model might be wrong. We hope that disentangling this information will yield an additional layer of interpretability and yield models that are as least as explainable as they are robust.

Question answering. The QA task takes the same format as SQuAD1.1 (Rajpurkar et al., 2016), i.e., given a context and a question, extract an answer from the context as a continuous span of text. The first round of adversarial QA (AQA) data comes from “Beat the AI” (Bartolo et al., 2020). During annotation, crowd workers were presented with a context sourced from Wikipedia, identical to those in SQuAD1.1, and asked to write a question and select an answer. The annotated answer was compared to the model prediction using a word-overlap F_1 threshold and, if sufficiently different, considered to have fooled the model. The target models in round 1 were BiDAF (Seo et al., 2017), BERT-Large, and RoBERTa-Large.

The model in the loop for the current round is RoBERTa trained on the examples from the first round combined with SQuAD1.1. Despite the super-human performance achieved on SQuAD1.1, machine performance is still far from humans on the current leaderboard. In the current phase, we seek to collect rich and diverse examples, focusing on improving model robustness through generative data augmentation, to provide more challenging model adversaries in this constrained task setting. We should emphasize that we don’t consider this task structure representative of the broader definition even of closed-domain QA, and are looking to expand this to include unanswerable questions (Rajpurkar et al., 2018), longer and more complex passages, Yes/No questions and multi-span

answers (Kwiatkowski et al., 2019), and numbers, dates and spans from the question (Dua et al., 2019) as model performance progresses.

Sentiment analysis. The sentiment analysis project is a multi-pronged effort to create a dynamic benchmark for sentiment analysis and to evaluate some of the core hypotheses behind Dynabench. Potts et al. (2020) provide an initial report and the first two rounds of this dataset.

The task is structured as a 3-way classification problem: positive, negative, and neutral. The motivation for using a simple positive/negative dichotomy is to show that there are still very challenging phenomena in this traditional sentiment space. The neutral category was added to avoid (and helped trained models avoid) the false presupposition that every text conveys sentiment information (Pang and Lee, 2008). In future iterations, we plan to consider additional dimensions of sentiment and emotional expression (Alm et al., 2005; Neviarouskaya et al., 2010; Wiebe et al., 2005; Liu et al., 2003; Sudhof et al., 2014).

In this first phase, we examined the question of how best to elicit examples from workers that are diverse, creative, and naturalistic. In the “prompt” condition, we provide workers with an actual sentence from an existing product or service review and ask them to edit it so that it fools the model. In the “no prompt” condition, workers try to write original sentences that fool the model. We find that the “prompt” condition is superior: workers generally make substantial edits, and the resulting sentences are more linguistically diverse than those in the “no prompt” condition.

In a parallel effort, we also collected and validated hard sentiment examples from existing corpora, which will enable another set of comparisons that will help us to refine the Dynabench protocols and interfaces. We plan for the dataset to continue to grow, probably mixing attested examples with those created on Dynabench with the help of prompts. With these diverse rounds, we can address a wide range of question pertaining to dataset artifacts, domain transfer, and overall robustness of sentiment analysis systems.

Hate speech detection. The hate speech task classifies whether a statement expresses hate against a protected characteristic or not. Detecting hate is notoriously difficult given the important role played by context and speaker (Leader May-

nard and Benesch, 2016) and the variety of ways in which hate can be expressed (Waseem et al., 2017). Few high-quality, varied and large training datasets are available for training hate detection systems (Vidgen and Derczynski, 2020; Poletto et al., 2020; Vidgen et al., 2019).

We organised four rounds of data collection and model training, with preliminary results reported in Vidgen et al. (2020). In each round, annotators are tasked with entering content that tricks the model into giving an incorrect classification. The content is created by the annotators and as such is synthetic in nature. At the end of each round the model is retrained and the process is repeated. For the first round, we trained a RoBERTa model on 470,000 hateful and abusive statements⁴. For subsequent rounds the model was trained on the original data plus content from the prior rounds. Due to the complexity of online hate, we hired and trained analysts rather than paying for crowd-sourced annotations. Each analyst was given training, support, and feedback throughout their work.

In all rounds annotators provided a label for whether content is hateful or not. In rounds 2, 3 and 4, they also gave labels for the target (i.e., which group has been attacked) and type of statement (e.g., derogatory remarks, dehumanization, or threatening language). These granular labels help to investigate model errors and improve performance, as well as directing the identification of new data for future entry. For approximately half of entries in rounds 2, 3 and 4, annotators created “perturbations” where the text is minimally adjusted so as to flip the label (Gardner et al., 2020; Kaushik et al., 2020). This helps to identify decision boundaries within the model, and minimizes the risk of overfitting given the small pool of annotators.

Over the four rounds, content becomes increasingly adversarial (shown by the fact that target models have lower performance on later rounds’ data) and models improve (shown by the fact that the model error rate declines and the later rounds’ models have the highest accuracy on each round). We externally validate performance using the HATECHECK suite of diagnostic tests from Röttger et al. (2020). We show substantial improvement over the four rounds, and our final round target model achieves 94% on HATECHECK, outperforming the models presented by the original authors.

⁴Derived from <https://hatespeechdata.com>, in anonymized form.

Task	Rounds	Examples	vMER
NLI	4	170,294	33.24%
QA	2	36,406	33.74%
Sentiment	3	19,975	35.00%
Hate speech	4	41,255	43.90%

Table 1: Statistics for the initial four official tasks.

3.3 Dynabenchmarking NLP

Table 1 shows an overview of the current situation for the four tasks. Some tasks are further along in their data collection efforts than others. As we can see, the validated model error rate (vMER; the number of human-validated model errors divided by the total number of examples—note that the error rates are not necessarily comparable across tasks, since the interfaces and in-the-loop models are not identical) is still very high across all tasks, clearly demonstrating that NLP is far from solved.

4 Caveats and Objections

There are several obvious and valid objections one can raise. We do not have all the answers, but we can try to address some common concerns.

Won’t this lead to unnatural distributions and distributional shift? Yes, that is a real risk. First, we acknowledge that crowdsourced texts are likely to have unnatural qualities: the setting itself is artificial from the perspective of genuine communication, and crowdworkers are not representative of the general population. Dynabench could exacerbate this, but it also has features that can help alleviate it. For instance, as we discussed earlier, the sentiment analysis project is using naturalistic prompt sentences to try to help workers create more diverse and naturalistic data.

Second, if we rely solely on dynamic *adversarial* collection, then we increase the risks of creating unnatural datasets. For instance, [Bartolo et al. \(2020\)](#) show that training solely on adversarially-collected data for QA was detrimental to performance on non-adversarially collected data. However, they also show that models are capable of simultaneously learning both distributions when trained on the combined data, retaining if not slightly improving performance on the original distribution (of course, this may not hold if we have many more examples of one particular kind). Ideally, we would combine adversarially collected data with

non-adversarial—preferably naturally collected—data, so as to capture both the average and worst case scenarios in our evaluation.

Finally, we note that Dynabench could enable the community to explore the kinds of distributional shift that are characteristic of natural languages. Words and phrases change their meanings over time, between different domains, and even between different interlocutors. Dynabench could be a tool for studying such shifts and finding models that can succeed on such phenomena.

What if annotators “overfit” on models? A potential risk is cyclical “progress,” where improved models forget things that were relevant in earlier rounds because annotators focus too much on a particular weakness. Continual learning is an exciting research direction here: we should try to understand distributional shift better, as well as how to characterize how data shifts over time might impact learning, and how any adverse effects might be overcome. Because of how most of us have been trained, it is natural to assume that the *last* round is automatically the best evaluation round, but that does not mean that it should be the only round: in fact, most likely, the best way to evaluate progress is to evaluate on *all* rounds as well as *any* high-quality static test set that exists, possibly with a recency-based discount factor. To make an analogy with software testing, similar to checklists ([Ribeiro et al., 2020](#)), it would be a bad idea to throw away old tests just because you’ve written some new ones. As long as we factor in previous rounds, Dynabench’s dynamic nature offers a way out from forgetting and cyclical issues: any model biases will be fixed in the limit by annotators exploiting vulnerabilities.

Another risk is that the data distribution might be too heavily dependent on the target model in the loop. When this becomes an issue, it can be mitigated by using ensembles of many different architectures in the loop, for example the top current state-of-the-art ones, with multiple seeds.⁵

How do we account for future, not-yet-in-the-loop models? Obviously, we can’t—so this is a very valid criticism. However, we can assume that an ensemble of model architectures is a reasonable approximation, if and only if the models are not too bad at their task. This latter point is crucial: we

⁵ANLI does not show dramatically different results across models, suggesting that this is not necessarily a big problem yet, but it shows in R2 and R3 that ensembles are possible.

take the stance that models by now, especially in aggregate, are probably good enough to be reasonably close enough to the decision boundaries—but it is definitely true that we have no guarantees that this is the case.

How do we compare results if the benchmark keeps changing? This is probably the main hurdle from a community adoption standpoint. But if we consider, e.g., the multiple iterations of SemEval or WMT datasets over the years, we’ve already been handling this quite well—we accept that a model’s BLEU score on WMT16 is not comparable to WMT14. That is, it is perfectly natural for benchmark datasets to evolve as the community makes progress. The only thing Dynabench does differently is that it *anticipates* dataset saturation and embraces the loop so that we can make faster and more sustained progress.

What about generative tasks? For now Dynabench focuses on classification or span extraction tasks where it is relatively straightforward to establish whether a model was wrong. If instead the evaluation metric is something like ROUGE or BLEU and we are interested in generation, we need a way to discretize an answer to determine correctness, since we wouldn’t have ground truth annotations; which makes determining whether a model was successfully fooled less straightforward. However, we could discretize generation by re-framing it as multiple choice with hard negatives, or simply by asking the annotator if the generation is good enough. In short, going beyond classification will require further research, but is definitely doable.

Do we need models in the loop for good data? The potential usefulness of adversarial examples can be explained at least in part by the fact that having an annotation partner (so far, a model) simply provides better incentives for generating quality annotation. Having the model in the loop is obviously useful for evaluation, but it’s less clear if the resultant data is necessarily also useful in general for training. So far, there is evidence that adversarially collected data provides performance gains irrespective of the model in the loop (Nie et al., 2020; Dinan et al., 2019; Bartolo et al., 2020). For example, ANLI shows that replacing equal amounts of “normally collected” SNLI and MNLI training data with ANLI data improves model performance, especially when training size is small (Nie et al., 2020), suggesting higher data efficiency. How-

ever, it has also been found that model-in-the-loop counterfactually-augmented training data does not necessarily lead to better generalization (Huang et al., 2020). Given the distributional shift induced by adversarial settings, it would probably be wisest to combine adversarially collected data with non-adversarial data during training (ANLI takes this approach), and to also test models in both scenarios. To get the most useful training and testing data, it seems the focus should be on collecting adversarial data with the best available model(s), preferably with a wide range of expertise, as that will likely be beneficial to future models also. That said, we expect this to be both task and model dependent. Much more research is required, and we encourage the community to explore these topics.

Is it expensive? Dynamic benchmarking is indeed expensive, but it is worth putting the numbers in context, as all data collection efforts are expensive when done at the scale of our current benchmark tasks. For instance, SNLI has 20K examples that were separately validated, and each one of these examples cost approximately \$0.50 to obtain and validate (personal communication with SNLI authors). Similarly, the 40K validated examples in MultiNLI cost \$0.64 each (p.c., MultiNLI authors). By comparison, the average cost of creation and validation for ANLI examples is closer to \$1.00 (p.c., ANLI authors). This is a substantial increase at scale. However, dynamic adversarial datasets may also last longer as benchmarks. If true, then the increased costs could turn out to be a bargain.

We should acknowledge, though, that dynamic benchmarks will tend to be more expensive than regular benchmarks for comparable tasks, because not every annotation attempt will be model-fooling and validation is required. Such expenses are likely to increase through successive rounds, as the models become more robust to workers’ adversarial attacks. The research bet is that each example obtained this way is actually worth more to the community and thus worth the expense.

In addition, we hope that language enthusiasts and other non-crowdworker model breakers will appreciate the honor that comes with being high up on the user leaderboard for breaking models. We are working on making the tool useful for education, as well as gamifying the interface to make it (even) more fun to try to fool models, as a “game with a purpose” (Von Ahn and Dabbish, 2008), for example through the ability to earn badges.

5 Conclusion and Outlook

We introduced Dynabench, a research platform for dynamic benchmarking. Dynabench opens up exciting new research directions, such as investigating the effects of ensembles in the loop, distributional shift characterisation, exploring annotator efficiency, investigating the effects of annotator expertise, and improving model robustness to targeted adversarial attacks in an interactive setting. It also facilitates further study in dynamic data collection, and more general cross-task analyses of human-and-machine interaction. The current iteration of the platform is only just the beginning of a longer journey. In the immediate future, we aim to achieve the following goals:

Anyone can run a task. Having created a tool that allows for human-in-the-loop model evaluation and data collection, we aim to make it possible for anyone to run their own task. To get started, only three things are needed: a target model, a (set of) context(s), and a pool of annotators.

Multilinguality and multimodality. As of now, Dynabench is text-only and focuses on English, but we hope to change that soon.

Live model evaluation. Model evaluation should not be about one single number on some test set. If models are uploaded through a standard interface, they can be scored automatically along many dimensions. We would be able to capture not only accuracy, for example, but also usage of computational resources, inference time, fairness, and many other relevant dimensions. This will in turn enable dynamic leaderboards, for example based on utility (Ethayarajh and Jurafsky, 2020). This would also allow for backward-compatible comparisons, not having to worry about the benchmark changing, and automatically putting new state of the art models in the loop, addressing some of the main objections.

One can easily imagine a future where, in order to fulfill reproducibility requirements, authors do not only link to their open source codebase but also to their model inference point so others can “talk with” their model. This will help drive progress, as it will allow others to examine models’ capabilities and identify failures to address with newer even better models. If we cannot always democratize the *training* of state-of-the-art AI models, at the very least we can democratize their *evaluation*.

Acknowledgements

We would like to thank Jason Weston, Emily Dinan and Kyunghyun Cho for their input on this project, and Sonia Kris for her support. ZW has been supported in part by the Canada 150 Research Chair program and the UK-Canada AI Artificial Intelligence Initiative. YN and MB have been supported in part by DARPA MCS N66001-19-2-4031, DARPA YFA17-D17AP00022, and ONR N00014-18-1-2871. CP has been supported in part by grants from Facebook, Google, and by Stanford’s Institute for Human-Centered AI.

References

- Cecilia Ovesdotter Alm, Dan Roth, and Richard Sproat. 2005. [Emotions from text: Machine learning for text-based emotion prediction](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 579–586, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Joshua Attenberg, Panos Ipeirotis, and Foster Provost. 2015. [Beat the machine: Challenging humans to find a predictive model’s “unknown unknowns”](#). *J. Data and Information Quality*, 6(1).
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). *arXiv preprint arXiv:1409.0473*.
- Max Bartolo, Alastair Roberts, Johannes Welbl, Sebastian Riedel, and Pontus Stenetorp. 2020. [Beat the ai: Investigating adversarial human annotation for reading comprehension](#). *Transactions of the Association for Computational Linguistics*, 8:662–678.
- Yonatan Belinkov and Yonatan Bisk. 2018. [Synthetic and natural noise both break neural machine translation](#). In *International Conference on Learning Representations*.
- Yonatan Belinkov, Adam Poliak, Stuart Shieber, Benjamin Van Durme, and Alexander Rush. 2019. [Don’t take the premise for granted: Mitigating artifacts in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 877–891, Florence, Italy. Association for Computational Linguistics.
- 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.

- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *arXiv preprint arXiv:2005.14165*.
- 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.
- Michael Chen, Mike D’Arcy, Alisa Liu, Jared Fernandez, and Doug Downey. 2019. [CODAH: An adversarially-authored question answering dataset for common sense](#). In *Proceedings of the 3rd Workshop on Evaluating Vector Space Representations for NLP*, pages 63–69, Minneapolis, USA. Association for Computational Linguistics.
- Alexis Conneau and Douwe Kiela. 2018. [SentEval: An evaluation toolkit for universal sentence representations](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. [The PASCAL recognising textual entailment challenge](#). In *Machine learning challenges. evaluating predictive uncertainty, visual object classification, and recognising textual entailment*, pages 177–190. Springer.
- Harm de Vries, Dzmitry Bahdanau, and Christopher Manning. 2020. [Towards ecologically valid research on language user interfaces](#). *arXiv preprint arXiv:2007.14435*.
- 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.
- Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. 2019. [Build it break it fix it for dialogue safety: Robustness from adversarial human attack](#). 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 4537–4546, Hong Kong, China. Association for Computational Linguistics.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. [DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs](#). 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 2368–2378, Minneapolis, Minnesota. Association for Computational Linguistics.
- Javid Ebrahimi, Daniel Lowd, and Dejing Dou. 2018a. [On adversarial examples for character-level neural machine translation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 653–663, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018b. [HotFlip: White-box adversarial examples for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia. Association for Computational Linguistics.
- Kawin Ethayarajh and Dan Jurafsky. 2020. [Utility is in the eye of the user: A critique of NLP leaderboard design](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4846–4853, Online. Association for Computational Linguistics.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- 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.
- 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.
- Jon Gauthier, Jennifer Hu, Ethan Wilcox, Peng Qian, and Roger Levy. 2020. [SyntaxGym: An online platform for targeted evaluation of language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 70–76, Online. Association for Computational Linguistics.
- Atticus Geiger, Ignacio Cases, Lauri Karttunen, and Christopher Potts. 2018. [Stress-testing neural models of natural language inference with](#)

- multiply-quantified sentences. *arXiv preprint arXiv:1810.13033*.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. [Breaking NLI systems with sentences that require simple lexical inferences](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–655, Melbourne, Australia. Association for Computational Linguistics.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). 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 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazare, and Jason Weston. 2019. [Learning from dialogue after deployment: Feed yourself, chatbot!](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3667–3684, Florence, Italy. Association for Computational Linguistics.
- Md Mosharaf Hossain, Venelin Kovatchev, Pranoy Dutta, Tiffany Kao, Elizabeth Wei, and Eduardo Blanco. 2020. [An analysis of natural language inference benchmarks through the lens of negation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9106–9118, Online. Association for Computational Linguistics.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- William Huang, Haokun Liu, and Samuel R Bowman. 2020. [Counterfactually-augmented snli training data does not yield better generalization than un-augmented data](#). *arXiv preprint arXiv:2010.04762*.
- Paloma Jeretic, Alex Warstadt, Suvrat Bhooshan, and Adina Williams. 2020. [Are natural language inference models IMPPRESSive? Learning IMpliciture and PRESupposition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8690–8705, Online. Association for Computational Linguistics.
- 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.
- Divyansh Kaushik, Eduard Hovy, and Zachary C Lipton. 2020. [Learning the difference that makes a difference with counterfactually-augmented data](#). *International Conference on Learning Representations (ICLR)*.
- Huda Khayrallah and Philipp Koehn. 2018. [On the impact of various types of noise on neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 74–83, Melbourne, Australia. Association for Computational Linguistics.
- Najoung Kim and Tal Linzen. 2020. [COGS: A compositional generalization challenge based on semantic interpretation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9087–9105, Online. 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.
- Jonathan Leader Maynard and Susan Benesch. 2016. [Dangerous Speech and Dangerous Ideology: An Integrated Model for Monitoring and Prevention](#). *Genocide Studies and Prevention*, 9(3):70–95.
- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2020. [Question and answer test-train overlap in open-domain question answering datasets](#). *arXiv preprint arXiv:2008.02637*.
- Tal Linzen. 2020. [How can we accelerate progress towards human-like linguistic generalization?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217, Online. Association for Computational Linguistics.
- Hugo Liu, Henry Lieberman, and Ted Selker. 2003. [A model of textual affect sensing using real-world knowledge](#). In *Proceedings of Intelligent User Interfaces (IUI)*, pages 125–132.
- Nelson F. Liu, Roy Schwartz, and Noah A. Smith. 2019a. [Inoculation by fine-tuning: A method for analyzing challenge datasets](#). 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 2171–2179, Minneapolis, Minnesota. 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. 2019b. [RoBERTa: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.

- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a large annotated corpus of English: The Penn Treebank](#). *Computational Linguistics*, 19(2):313–330.
- Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 2019. [On measuring social biases in sentence encoders](#). 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 622–628, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. 2018. [The natural language deathlon: Multitask learning as question answering](#). *arXiv preprint arXiv:1806.08730*.
- 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.
- Pasquale Minervini and Sebastian Riedel. 2018. [Adversarially regularising neural NLI models to integrate logical background knowledge](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 65–74, Brussels, Belgium. Association for Computational Linguistics.
- Tom Mitchell, William Cohen, Estevam Hruschka, Partha Talukdar, Bishan Yang, Justin Betteridge, Andrew Carlson, Bhavana Dalvi, Matt Gardner, Bryan Kisiel, et al. 2018. [Never-ending learning](#). *Communications of the ACM*, 61(5):103–115.
- Robert Munro, Steven Bethard, Victor Kuperman, Vicky Tzuyin Lai, Robin Melnick, Christopher Potts, Tyler Schnoebelen, and Harry Tily. 2010. [Crowdsourcing and language studies: the new generation of linguistic data](#). In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*, pages 122–130, Los Angeles. Association for Computational Linguistics.
- 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, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. [CrowS-pairs: A challenge dataset for measuring social biases in masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. 2010. [Recognition of affect, judgment, and appreciation in text](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 806–814, Beijing, China. Coling 2010 Organizing Committee.
- Yixin Nie, Yicheng Wang, and Mohit Bansal. 2019. [Analyzing compositionality-sensitivity of nli models](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6867–6874.
- 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.
- Bo Pang and Lillian Lee. 2008. [Opinion mining and sentiment analysis](#). *Foundations and Trends in Information Retrieval*, 2(1):1–135.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems*, volume 32, pages 8026–8037. Curran Associates, Inc.
- Fabio Poletto, Valerio Basile, Manuela Sanguinetti, Cristina Bosco, and Viviana Patti. 2020. [Resources and benchmark corpora for hate speech detection: a systematic review](#). *Language Resources and Evaluation*.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. [Hypothesis only baselines in natural language inference](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.
- Christopher Potts, Zhengxuan Wu, Atticus Geiger, and Douwe Kiela. 2020. [DynaSent: A dynamic benchmark for sentiment analysis](#). *arXiv preprint arXiv:2012.15349*.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. [CoNLL-](#)

- 2012 shared task: [Modeling multilingual unrestricted coreference in OntoNotes](#). In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8):9.
- 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.
- Ohad Rozen, Vered Shwartz, Roei Aharoni, and Ido Dagan. 2019. [Diversify your datasets: Analyzing generalization via controlled variance in adversarial datasets](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 196–205, Hong Kong, China. 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.
- Paul Röttger, Bertram Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet Pierrehumbert. 2020. [Hatecheck: Functional tests for hate speech detection models](#).
- Swarnadeep Saha, Yixin Nie, and Mohit Bansal. 2020. [ConjNLI: Natural language inference over conjunctive sentences](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8240–8252, Online. Association for Computational Linguistics.
- Victor Sánchez Valencia. 1991. *Studies on natural logic and categorial grammar*, University of Amsterdam Ph. D. Ph.D. thesis, thesis.
- Chinnadhurai Sankar, Sandeep Subramanian, Chris Pal, Sarath Chandar, and Yoshua Bengio. 2019. [Do neural dialog systems use the conversation history effectively? an empirical study](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 32–37, Florence, Italy. Association for Computational Linguistics.
- Sebastian Schuster, Yuxing Chen, and Judith Degen. 2020. [Harnessing the linguistic signal to predict scalar inferences](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5387–5403, Online. Association for Computational Linguistics.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. [Bidirectional attention flow for machine comprehension](#). In *The International Conference on Learning Representations (ICLR)*.
- Kurt Shuster, Da Ju, Stephen Roller, Emily Dinan, Y-Lan Boureau, and Jason Weston. 2020. [The dialogue dodecaathlon: Open-domain knowledge and image grounded conversational agents](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2453–2470, Online. Association for Computational Linguistics.
- Daniel L Silver, Qiang Yang, and Lianghao Li. 2013. [Lifelong machine learning systems: Beyond learning algorithms](#). In *2013 AAAI spring symposium series*.
- Rion Snow, Brendan O’Connor, Daniel Jurafsky, and Andrew Ng. 2008. [Cheap and fast – but is it good? evaluating non-expert annotations for natural language tasks](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 254–263, Honolulu, Hawaii. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. [Learning to summarize with human feedback](#). *Advances in Neural Information Processing Systems*, 33.
- Moritz Sudhof, Andrés Gómez Emilsson, Andrew L. Maas, and Christopher Potts. 2014. [Sentiment expression conditioned by affective transitions and social forces](#). In *Proceedings of 20th Conference on Knowledge Discovery and Data Mining*, pages 1136–1145, New York. ACM.

- Saku Sugawara, Pontus Stenetorp, Kentaro Inui, and Akiko Aizawa. 2020. Assessing the benchmarking capacity of machine reading comprehension datasets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8918–8927.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. [Axiomatic attribution for deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 3319–3328. PMLR.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. [Dataset cartography: Mapping and diagnosing datasets with training dynamics](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293, Online. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- 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).
- Johan van Benthem. 2008. A brief history of natural logic. In *Logic, Navya-Nyaya and Applications: Homage to Bimal Matilal*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Bertie Vidgen and Leon Derczynski. 2020. [Directions in Abusive Language Training Data: Garbage In, Garbage Out](#). *arXiv:2004.01670*, pages 1–26.
- Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019. [Challenges and frontiers in abusive content detection](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 80–93, Florence, Italy. Association for Computational Linguistics.
- Bertie Vidgen, Tristan Thrush, Zeerak Waseem, and Douwe Kiela. 2020. [Learning from the worst: Dynamically generated datasets to improve online hate detection](#). *arXiv preprint arXiv:2012.15761*.
- Luis Von Ahn and Laura Dabbish. 2008. [Designing games with a purpose](#). *Communications of the ACM*, 51(8):58–67.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. [Universal adversarial triggers for attacking and analyzing NLP](#). 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 2153–2162, Hong Kong, China. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems*, volume 32, pages 3266–3280. Curran Associates, Inc.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Alex Warstadt, Yu Cao, Ioana Grosu, Wei Peng, Hagen Blix, Yining Nie, Anna Alsop, Shikha Bordia, Haokun Liu, Alicia Parrish, Sheng-Fu Wang, Jason Phang, Anhad Mohananey, Phu Mon Htut, Paloma Jeretic, and Samuel R. Bowman. 2019. [Investigating BERT’s knowledge of language: Five analysis methods with NPIs](#). 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 2877–2887, Hong Kong, China. Association for Computational Linguistics.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: A benchmark of linguistic minimal pairs for English](#). In *Proceedings of the Society for Computation in Linguistics 2020*, pages 409–410, New York, New York. Association for Computational Linguistics.
- Zeerak Waseem, Thomas Davidson, Dana Warmusley, and Ingmar Weber. 2017. [Understanding Abuse: A Typology of Abusive Language Detection Subtasks](#). In *Proceedings of the First Workshop on Abusive Language Online*, pages 78–84.
- Aaron Steven White, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2018. [Lexicosyntactic](#)

- [inference in neural models](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4717–4724, Brussels, Belgium. Association for Computational Linguistics.
- Aaron Steven White, Elias Stengel-Eskin, Siddharth Vashishtha, Venkata Subrahmanyam Govindarajan, Dee Ann Reisinger, Tim Vieira, Keisuke Sakaguchi, Sheng Zhang, Francis Ferraro, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2020. [The universal decompositional semantics dataset and decomp toolkit](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5698–5707, Marseille, France. European Language Resources Association.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. [Annotating expressions of opinions and emotions in language](#). *Language Resources and Evaluation*, 39(2–3):165–210.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Zhilin Yang, Saizheng Zhang, Jack Urbanek, Will Feng, Alexander H Miller, Arthur Szlam, Douwe Kiela, and Jason Weston. 2017. [Mastering the dungeon: Grounded language learning by mechanical turker descent](#). *arXiv preprint arXiv:1711.07950*.
- Lang Yu and Allyson Ettinger. 2020. [Assessing phrasal representation and composition in transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4896–4907, Online. Association for Computational Linguistics.
- Xiang Zhou, Yixin Nie, Hao Tan, and Mohit Bansal. 2020. [The curse of performance instability in analysis datasets: Consequences, source, and suggestions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8215–8228, Online. Association for Computational Linguistics.