

Reliability Testing for Natural Language Processing Systems

Samson Tan^{§†*} Shafiq Joty^{§†} Kathy Baxter[§]
Araz Taeihagh^{◇♣} Gregory A. Bennett[§] Min-Yen Kan[‡]

[§]Salesforce Research

[‡]Nanyang Technological University

[‡]School of Computing, National University of Singapore

[◇]Lee Kuan Yew School of Public Policy, National University of Singapore

[♣]Centre for Trusted Internet and Community, National University of Singapore

Abstract

Questions of fairness, robustness, and transparency are paramount to address before deploying NLP systems. Central to these concerns is the question of reliability: Can NLP systems reliably treat different demographics fairly *and* function correctly in diverse and noisy environments? To address this, we argue for the need for reliability testing and contextualize it among existing work on improving accountability. We show how adversarial attacks can be reframed for this goal, via a framework for developing reliability tests. We argue that reliability testing — with an emphasis on interdisciplinary collaboration — will enable rigorous and targeted testing, and aid in the enactment and enforcement of industry standards.

1 Introduction

Rigorous testing is critical to ensuring a program works as intended (functionality) when used under real-world conditions (reliability). Hence, it is troubling that while natural language technologies are becoming increasingly pervasive in our everyday lives, there is little assurance that these NLP systems will not fail catastrophically or amplify discrimination against minority demographics when exposed to input from outside the training distribution. Recent examples include GPT-3 (Brown et al., 2020) agreeing with suggested suicide (Rousseau et al., 2020), the mistranslation of an innocuous social media post resulting in a minority’s arrest (Hern, 2017), and biased grading algorithms that can negatively impact a minority student’s future (Feathers, 2019). Additionally, a lack of rigorous testing, coupled with machine learning’s (ML) implicit assumption of identical training and testing distributions, may inadvertently result in systems that discriminate against minorities, who are often underrepresented in the training data. This can take

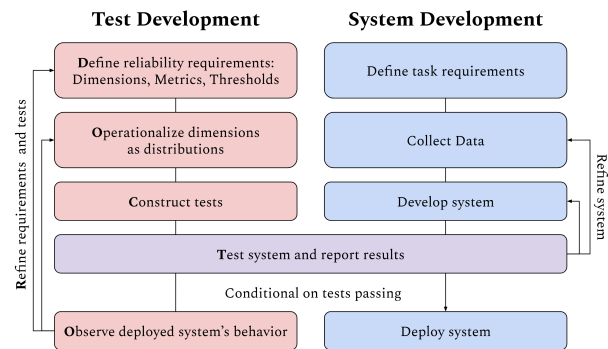


Figure 1: How DOCTOR can integrate with existing system development workflows. Test (left) and system development (right) take place in parallel, separate teams. Reliability tests can thus be constructed independent of the system development team, either by an internal “red team” or by independent auditors.

the form of misrepresentation of or poorer performance for people with disabilities, specific gender, ethnic, age, or linguistic groups (Hovy and Spruit, 2016; Crawford, 2017; Hutchinson et al., 2020).

Amongst claims of NLP systems achieving human parity in challenging tasks such as question answering (Yu et al., 2018), machine translation (Hasan et al., 2018), and commonsense inference (Devlin et al., 2019), research has demonstrated these systems’ fragility to natural and adversarial noise (Goodfellow et al., 2015; Belinkov and Bisk, 2018) and out-of-distribution data (Fisch et al., 2019). It is also still common practice to equate “testing” with “measuring held-out accuracy”, even as datasets are revealed to be harmfully biased (Wagner et al., 2015; Geva et al., 2019; Sap et al., 2019).

Many potential harms can be mitigated by detecting them early and preventing the offending model from being put into production. Hence, in addition to being mindful of the biases in the NLP pipeline (Bender and Friedman, 2018; Mitchell et al., 2019;

*Correspondence to: samson.tan@salesforce.com

Waseem et al., 2021) and holding creators accountable via audits (Raji et al., 2020; Brundage et al., 2020), we argue for the need to evaluate an NLP system’s reliability in diverse operating conditions.

Initial research on evaluating out-of-distribution generalization involved manually-designed challenge sets (Jia and Liang, 2017; Nie et al., 2020; Gardner et al., 2020), counterfactuals (Kaushik et al., 2019; Khashabi et al., 2020; Wu et al., 2021), biased sampling (Søgaard et al., 2021) or toolkits for testing if a system has specific capabilities (Ribeiro et al., 2020) or robustness to distribution shifts (Goel et al., 2021). However, most of these approaches inevitably overestimate a given system’s *worst-case performance* since they do not mimic the NLP system’s adversarial distribution¹.

A promising technique for evaluating worst-case performance is the adversarial attack. However, although some adversarial attacks explicitly focus on specific linguistic levels of analysis (Belinkov and Bisk, 2018; Iyyer et al., 2018; Tan et al., 2020; Eger and Benz, 2020), many often simply rely on word embeddings or language models for perturbation proposal (see §4). While the latter may be useful to evaluate a system’s robustness to malicious actors, they are less useful for dimension-specific testing (e.g., reliability when encountering grammatical variation). This is because they often perturb the input across multiple dimensions at once, which may make the resulting adversaries unnatural.

Hence, in this paper targeted at NLP researchers, practitioners, and policymakers, we make the case for reliability testing and reformulate adversarial attacks as *dimension-specific*, worst-case tests that can be used to approximate real-world variation. We contribute a reliability testing framework — DOCTOR — that translates safety and fairness concerns around NLP systems into quantitative tests. We demonstrate how testing dimensions for DOCTOR can be drafted for a specific use case. Finally, we discuss the policy implications, challenges, and directions for future research on reliability testing.

2 Terminology Definitions

Let’s define key terms to be used in our discussion.

NLP system. The entire text processing pipeline built to solve a specific task; taking raw text as input and producing predictions in the form of labels

¹The distribution of adversarial cases or failure profile.

(classification) or text (generation). We exclude *raw* language models from the discussion since it is unclear how performance, and hence worst-case performance, should be evaluated. We do include NLP systems that use language models internally (e.g., BERT-based classifiers (Devlin et al., 2019)).

Reliability. Defined by IEEE (2017) as the “degree to which a system, product or component performs specified functions under specified conditions for a specified period of time”. We prefer this term over robustness² to challenge the NLP community’s common framing of inputs from outside the training distribution as “noisy”. The notion of reliability requires us to explicitly consider the specific, diverse environments (i.e., communities) a system will operate in. This is crucial to reducing the NLP’s negative impact on the underrepresented.

Dimension. An axis along which variation can occur in the real world, similar to Plank (2016)’s variety space. A taxonomy of possible dimensions can be found in Table 1 (Appendix).

Adversarial attack. A method of perturbing the input to degrade a target model’s accuracy (Goodfellow et al., 2015). In computer vision, this is achieved by adding adversarial noise to the image, optimized to be maximally damaging to the model. §4 describes how this is done in the NLP context.

Stakeholder. A person who is (in-)directly impacted by the NLP system’s predictions.

Actor. Someone who has influence over a) the design of an NLP system and its reliability testing regime; b) whether the system is deployed; and c) who it can interact with. Within the context of our discussion, actors are likely to be regulators, experts, and stakeholder advocates.

Expert. An actor who has specialized knowledge, such as ethicists, linguists, domain experts, social scientists, or NLP practitioners.

3 The Case for Reliability Testing in NLP

The accelerating interest in building NLP-based products that impact many lives has led to urgent questions of fairness, safety, and accountability (Hovy and Spruit, 2016; Bender et al., 2021),

²The “degree to which a system or component can function correctly in the presence of invalid inputs or stressful environmental conditions” (IEEE, 2017).

prompting research into algorithmic bias (Bolukbasi et al., 2016; Blodgett et al., 2020), explainability (Ribeiro et al., 2016; Danilevsky et al., 2020), robustness (Jia and Liang, 2017), etc. Research is also emerging on best practices for productizing ML: from detailed dataset documentation (Bender and Friedman, 2018; Gebru et al., 2018), model documentation for highlighting important but often unreported details such as its training data, intended use, and caveats (Mitchell et al., 2019), and documentation best practices (Partnership on AI, 2019), to institutional mechanisms such as auditing (Raji et al., 2020) to enforce accountability and red-teaming (Brundage et al., 2020) to address developer blind spots, not to mention studies on the impact of organizational structures on responsible AI initiatives (Rakova et al., 2020).

Calls for increased accountability and transparency are gaining traction among governments (116th U.S. Congress, 2019; NIST, 2019; European Commission, 2020; Smith, 2020; California State Legislature, 2020; FDA, 2021) and customers increasingly cite ethical concerns as a reason for not engaging AI service providers (EIU, 2020).

While there has been significant discussion around best practices for dataset and model creation, work to ensure NLP systems are evaluated in a manner representative of their operational conditions has only just begun. Initial work in constructing representative tests focuses on enabling development teams to easily evaluate their models' linguistic capabilities (Ribeiro et al., 2020) and accuracy on subpopulations and distribution shifts (Goel et al., 2021). However, there is a clear need for a paradigm that allows experts and stakeholder advocates to collaboratively develop tests that are representative of the practical and ethical concerns of an NLP system's target demographic. We argue that *reliability testing*, by reframing the concept of adversarial attacks, has the potential to fill this gap.

3.1 What is reliability testing?

Despite the recent advances in neural architectures resulting in breakthrough performance on benchmark datasets, research into adversarial examples and out-of-distribution generalization has found ML systems to be particularly vulnerable to slight perturbations in the input (Goodfellow et al., 2015) and natural distribution shifts (Fisch et al., 2019). While these perturbations are often chosen to max-

imize model failure, they highlight serious reliability issues for putting ML models into production since they show that these models *could* fail catastrophically in naturally noisy, diverse, real-world environments (Saria and Subbaswamy, 2019). Additionally, bias can seep into the system at multiple stages of the NLP lifecycle (Shah et al., 2020), resulting in discrimination against minority groups (O'Neil, 2016). The good news, however, is that rigorous testing can help to highlight potential issues before the systems are deployed.

The need for rigorous testing in NLP is reflected in ACL 2020 giving the Best Paper Award to CheckList (Ribeiro et al., 2020), which applied the idea of behavior testing from software engineering to testing NLP systems. While invaluable as a first step towards the development of comprehensive testing methodology, the current implementation of CheckList may still overestimate the reliability of NLP systems since the individual test examples are largely manually constructed. Importantly, with the complexity and scale of current models, humans cannot accurately determine a model's *adversarial* distribution (i.e., the examples that cause model failure). Consequently, the test examples they construct are unlikely to be the *worst-case examples* for the model. Automated assistance is needed.

Therefore, we propose to perform *reliability testing*, which can be thought of as one component of behavior testing. We categorize reliability tests as *average-case* tests or the *worst-case* tests. As their names suggest, average-case and worst-case tests estimate the expected and lower-bound performance, respectively, when the NLP system is exposed to the phenomena modeled by the tests. Average-case tests are conceptually similar to Wu et al. (2021)'s counterfactuals, which is contemporaneous work, while worst-case tests are most similar to adversarial attacks (§4).

Our approach parallels *boundary value testing* in software engineering: In boundary value testing, tests evaluate a program's ability to handle edge cases using test examples drawn from the extremes of the ranges the program is expected to handle. Similarly, reliability testing aims to quantify the system's reliability under diverse and potentially extreme conditions. This allows teams to perform better quality control of their NLP systems and introduce more nuance into discussions of why and when models fail (§5). Finally, we note that reliabil-

ity testing and standards are established practices in engineering industries (e.g., aerospace (Nelson, 2003; Wilkinson et al., 2016)) and advocate for NL engineering to be at parity with these fields.

3.2 Evaluating worst-case performance in a label-scarce world

A proposed approach for testing robustness to natural and adverse distribution shifts is to construct test sets using data from different domains or writing styles (Miller et al., 2020; Hendrycks et al., 2020), or to use a human vs. model method of constructing challenge sets (Nie et al., 2020; Zhang et al., 2019b). While they are the gold standard, such datasets are expensive to construct,³ making it infeasible to manually create worst-case test examples for each NLP system being evaluated. Consequently, these challenge sets necessarily overestimate each system’s worst-case performance when the inference distribution differs from the training one. Additionally, due to their crowdsourced nature, these challenge sets inevitably introduce distribution shifts across multiple dimensions at once, and even their own biases (Geva et al., 2019), unless explicitly controlled for. Building individual challenge sets for each dimension would be prohibitively expensive due to combinatorial explosion, even before having to account for concept drift (Widmer and Kubat, 1996). This coupling complicates efforts to design a nuanced and comprehensive testing regime. Hence, simulating variation in a controlled manner via reliability tests can be a complementary method of evaluating the system’s out-of-distribution generalization ability.

4 Adversarial Attacks as Reliability Tests

We first give a brief introduction to adversarial attacks in NLP before showing how they can be used for reliability testing. We refer the reader to Zhang et al. (2020b) for a comprehensive survey.

Existing work on NLP adversarial attacks perturbs the input at various levels of linguistic analysis: phonology (Eger and Benz, 2020), orthography (Ebrahimi et al., 2018), morphology (Tan et al., 2020), lexicon (Alzantot et al., 2018; Jin et al., 2020), and syntax (Iyyer et al., 2018).

Early work did not place any constraints on the attacks and merely used the degradation to a tar-

³Dua et al. (2019) reports a cost of 60k USD for 96k question-answer pairs.

Algorithm 1 General Reliability Test

Require: Data distribution $\mathcal{D}_d = \{\mathcal{X}, \mathcal{Y}\}$ modeling the dimension of interest d , NLP system \mathcal{M} , Source dataset $X \sim \mathcal{X}$, Desired labels $Y' \sim \mathcal{Y}$, Scoring function S .
Ensure: Average- or worst-case examples X' , Result r .

```

1:  $X' \leftarrow \{\emptyset\}, r \leftarrow 0$ 
2: for  $x, y'$  in  $X, Y'$  do
3:    $C \leftarrow \text{SAMPLECANDIDATES}(\mathcal{X})$ 
4:   switch TestType do
5:     case AverageCaseTest
6:        $s \leftarrow \text{MEAN}(S(y', \mathcal{M}(C)))$ 
7:        $X' \leftarrow X' \cup C$ 
8:     case WorstCaseTest
9:        $x', s \leftarrow \arg \min_{x_c \in C} S(y', \mathcal{M}(x_c))$ 
10:       $X' \leftarrow X' \cup \{x'\}$ 
11:    $r \leftarrow r + s$ 
12: end for
13:  $r \leftarrow \frac{r}{|X'|}$ 
14: return  $X', r$ 

```

get model’s accuracy as *the* measure of success. However, this often resulted in the semantics and expected prediction changing, leading to an overestimation of the attack’s success. Recent attacks aim to preserve the original input’s semantics. A popular approach has been to substitute words with their synonyms using word embeddings or a language model as a measure of semantic similarity (Alzantot et al., 2018; Ribeiro et al., 2018; Michel et al., 2019; Ren et al., 2019; Zhang et al., 2019a; Li et al., 2019; Jin et al., 2020; Garg and Ramakrishnan, 2020; Li et al., 2020a).

Focusing on maximally degrading model accuracy overlooks the key feature of adversarial attacks: the ability to find the worst-case example for a model from an arbitrary distribution. Many recent attacks perturb the input across multiple dimensions at once, which may make the result unnatural. By constraining our sample perturbations to a distribution modeling a specific dimension of interest, the performance on the generated adversaries is a valid lower bound performance for that dimension. Said another way, adversarial attacks can be reframed as interpretable reliability tests if we constrain them to meaningful distributions.

This is the key element of our approach as detailed in Alg. 1. We specify either an average (Lines 5–7) or worse case test (Lines 8–10), but conditioned on the data distribution \mathcal{D} that models a particular dimension of interest d . The resultant reliability score gauges real-world performance and the worst-case variant returns the adversarial examples that cause worst-case performance. When invariance to input variation is expected, y' is equivalent to the

source label y . Note that by ignoring the average-case test logic and removing d , we recover the general adversarial attack algorithm.

However, the key difference between an adversarial robustness mindset and a testing one is the latter's emphasis on identifying ways in which natural phenomena or ethical concerns can be operationalized as reliability tests. This change in perspective opens up new avenues for interdisciplinary research that will allow researchers and practitioners to have more nuanced discussions about model reliability and can be used to design comprehensive reliability testing regimes. We describe such a framework for interdisciplinary collaboration next.

5 A Framework for Reliability Testing

We introduce and then describe our general framework, DOCTOR, for testing the reliability of NLP systems. DOCTOR comprises six steps:

1. Define reliability requirements
2. Operationalize dimensions as distributions
3. Construct tests
4. Test system and report results
5. Observe deployed system's behavior
6. Refine reliability requirements and tests

Defining reliability requirements. Before any tests are constructed, experts and stakeholder advocates should work together to understand the demographics and values of the communities the NLP system will interact with (Friedman and Hendry, 2019) and the system's impact on their lives. The latter is also known as algorithmic risk assessment (Ada Lovelace Institute and DataKind UK, 2021). There are three critical questions to address: 1) Along what dimensions should the model be tested? 2) What metrics should be used to measure system performance? 3) What are acceptable performance thresholds for each dimension?

Question 1 can be further broken down into: a) general linguistic phenomena, such as alternative spellings or code-mixing; b) task-specific quirks, e.g., an essay grading system should not use text length to predict score; c) sensitive attributes, such as gender, ethnicity, sexual orientation, age, or disability status. This presents an opportunity for interdisciplinary expert collaboration: Linguists are best equipped to contribute to discussions around (a),

domain experts to (b), and ethicists and social scientists to (c). However, we recognize that such collaboration may not be feasible for every NLP system being tested. It is more realistic to expect ethicists to be involved when applying DOCTOR at the company and industry levels, and ethics-trained NLP practitioners to answer these questions within the development team. We provide a taxonomy of potential dimensions in Table 1 (Appendix).

Since it is likely unfeasible to test every possible dimension, stakeholder advocates should be involved to ensure their values and interests are accurately represented and prioritized (Hagerty and Rubinov, 2019), while experts should ensure the dimensions identified can be feasibly tested. A similar approach to that of community juries⁴ may be taken. We recommend using this question to evaluate the feasibility of operationalizing potential dimensions: "What is the system's performance when exposed to variation along dimension d ?" For example, rather than simply "gender", a better-defined dimension would be "gender pronouns". With this understanding, experts and policymakers can then create a set of *reliability requirements*, comprising the testing dimensions, performance metric(s), and passing thresholds.

Next, we recommend using the same metrics for held-out, average-case, and worst-case performance for easy comparison. These often vary from task to task and are still a subject of active research (Novikova et al., 2017; Reiter, 2018; Kryscinski et al., 2019), hence the question of the right metric to use is beyond the scope of this paper. Finally, ethicists, in consultation with the other aforementioned experts and stakeholders, will determine acceptable thresholds for worst-case performance. The system under test must perform above said thresholds when exposed to variation along those dimensions in order to pass. For worst-case performance, we recommend reporting thresholds as relative differences (δ) between the average-case and worst-case performance. These questions may help in applying this step and deciding if specific NLP solutions should even exist (Leins et al., 2020):

- Who will interact with the NLP system, in what context, and using which language varieties?
- What are the distinguishing features of these varieties compared to those used for training?

⁴docs.microsoft.com/en-us/azure/.../community-jury

- What is the (short- and long-term) impact on the community’s most underrepresented members if the system performs more poorly for them?

We note that our framework is general enough to be applied at various levels of organization: within the development team, within the company (compliance team, internal auditor), and within the industry (self-regulation or independent regulator). However, we expect the exact set of dimensions, metrics and acceptable thresholds defined in Step 1 to vary depending on the reliability concerns of the actors at each level. For example, independent regulators will be most concerned with establishing minimum safety and fairness standards that all NLP systems used in their industries must meet, while compliance teams may wish to have stricter and more comprehensive standards for brand reasons. Developers can use DOCTOR to meet the other two levels of requirements and understand their system’s behaviour better with targeted testing.

Operationalizing dimensions. While the abstractness of dimensions allows people who are not NLP practitioners to participate in drafting the set of reliability requirements, there is no way to test NLP systems using fuzzy concepts. Therefore, every dimension the system is to be tested along must be operationalizable as a distribution from which perturbed examples can be sampled in order for NLP practitioners to realize them as tests.

Since average-case tests attempt to estimate a system’s expected performance in its deployed environment, the availability of datasets that reflect real-world distributions is paramount to ensure that the tests themselves are unbiased. This is less of an issue for worst-case tests; the tests only need to know which perturbations that are possible, but not how frequently they occur in the real world. Figuring out key dimensions for different classes of NLP tasks and exploring ways of operationalizing them as reliability tests are also promising directions for future research. Such research would help NLP practitioners and policymakers define reliability requirements that can be feasibly implemented.

Constructing tests. Next, average- and worst-case tests are constructed (Alg. 1). Average-case tests can be data-driven and could take the form of manually curated datasets or model-based perturbation generation (e.g., PolyJuice (Wu et al., 2021)), while worst-case tests can be rule-based

(e.g., Morpheus (Tan et al., 2020)) or model-based (e.g., BERT-Attack (Li et al., 2020a)). We recommend constructing tests that do not require access to the NLP model’s parameters (black-box assumption); this not only yields more system-agnostic tests, but also allows for (some) tests to be created independently from the system development team. If the black-box assumption proves limiting, the community can establish a standard set of items an NLP system should export for testing purposes, e.g., network gradients if the system uses a neural model. Regardless of assumption, keeping the regulators’ test implementations separate and hidden from the system developers is critical for stakeholders and regulators to trust the results. This separation also reduces overfitting to the test suite.

Testing systems. A possible model for test ownership is to have independently implemented tests at the three levels of organization described above (team, company, industry). At the development team level, reliability tests can be used to diagnose weaknesses with the goal of improving the NLP system for a specific use case and set of target users. Compared to unconstrained adversarial examples, contrasting worst-case examples that have been constrained along specific dimensions with non-worst-case examples will likely yield greater intuition into the model’s inner workings. Studying how modifications (to the architecture, training data and process) affect the system’s reliability on each dimension will also give engineers insight into the factors affecting system reliability. These tests should be executed and updated regularly during development, according to software engineering best practices such as Agile (Beck et al., 2001).

Red teams are company-internal teams tasked with finding security vulnerabilities in their developed software or systems. Brundage et al. (2020) propose to apply the concept of red teaming to surface flaws in an AI system’s safety and security. In companies that maintain multiple NLP systems, we propose employing similar, specialized teams composed of NLP experts to build and maintain reliability tests that ensure their NLP systems adhere to company-level reliability standards. These tests will likely be less task-/domain-specific than those developed by engineering teams due to their wider scope, while the reliability standards may be created and maintained by compliance teams or the red teams themselves. Making these stan-

dards available for public scrutiny and ensuring their products meet them will enable companies to build trust with their users. To ensure all NLP systems meet the company’s reliability standards, these reliability tests should be executed as a part of regular internal audits (Raji et al., 2020), investigative audits after incidents, and before major releases (especially if it is the system’s first release or if it received a major update). They may also be regularly executed on randomly chosen production systems and trigger an alert upon failure.

At the independent regulator level, reliability tests would likely be carried out during product certification (e.g., ANSI/ISO certification) and external audits. These industry-level reliability standards and tests may be developed in a similar manner to the company-level ones. However, we expect them to be more general and less comprehensive than the latter, analogous to minimum safety standards such as IEC 60335-1 (IEC, 2020). Naturally, high risk applications and NLP systems used in regulated industries should comply with more stringent requirements (European Commission, 2021).

Our proposed framework is also highly compatible with the use of model cards (Mitchell et al., 2019) for auditing and transparent reporting (Raji et al., 2020). In addition to performance on task-related metrics, model cards surface information and assumptions about a machine learning system and training process that may not be readily available otherwise. When a system has passed all tests and is ready to be deployed, its average- and worst-case performance on all tested dimensions can be included as an extra section on the accompanying model card. In addition, the perturbed examples generated during testing and their labels (x' , y') can be stored for audit purposes or examined to ensure that the tests are performing as expected.

Observing and Refining requirements. It is crucial to regularly monitor the systems’ impact post-launch and add, update, or re-prioritize dimensions and thresholds accordingly. Monitoring large-scale deployments can be done via community juries, in which stakeholders who will be likely impacted (or their advocates) give feedback on their pain points and raise concerns about potential negative effects. Smaller teams without the resources to organize community juries can set up avenues (e.g., online forms) for affected stakeholders to give feedback, raise concerns, and seek remediation.

6 From Concerns to Dimensions

We now illustrate how reliability concerns can be converted into concrete testing dimensions (Step 1) by considering the scenario of applying automated text scoring to short answers and essays from students in the multilingual population of Singapore. We study a second scenario in Appendix A. **Automated Text Scoring (ATS)** systems are increasingly used to grade tests and essays (Markoff, 2013; Feathers, 2019). While they can provide instant feedback and help teachers and test agencies cope with large loads, studies have shown that they often exhibit demographic and language biases, such as scoring African- and Indian-American males lower on the GRE Argument task compared to human graders (Bridgeman et al., 2012; Ramineni and Williamson, 2018). Since the results of some tests will affect the futures of the test takers (Salaky, 2018), the scoring algorithms used must be sufficiently reliable. Hence, let us imagine that Singapore’s education ministry has decided to create a standard set of reliability requirements that all ATS systems used in education must adhere to.

Linguistic landscape. A mix of language varieties are used in Singapore: a prestige English variety, a colloquial English variety, three other official languages (Chinese, Malay, and Tamil), and a large number of other languages. English is the *lingua franca*, with fluency in the prestige variety correlating with socioeconomic status (Vaish and Tan, 2008). A significant portion of the population does not speak English at home. Subjects other than languages are taught in English.

Stakeholder impact. The key stakeholders affected by ATS systems would be students in schools and universities. The consequences of lower scores could be life-altering for the student who is unable to enroll in the major of their choice. At the population level, biases in an ATS system trained on normally sampled data would unfairly discriminate against already underrepresented groups. Additionally, biases against disfluent or ungrammatical text when they are not the tested attributes would result in discrimination against students with a lower socioeconomic status or for whom English is a second language.

Finally, NLP systems have also been known to be overly sensitive to alternative spellings (Belinkov and Bisk, 2018). When used to score subject tests, this could result in the ATS system unfairly penaliz-

ing dyslexic students (Coleman et al., 2009). Since education is often credited with enabling social mobility,⁵ unfair grading may perpetuate systemic discrimination and increase social inequality.

Dimension. We can generally categorize written tests into those that test for content correctness (e.g., essay questions in a history test), and those that test for language skills (e.g., proper use of grammar). While there are tests that simultaneously assess both aspects, modern ATS systems often grade them separately (Ke and Ng, 2019). We treat each aspect as a separate test here.

When grading students on content correctness, we would expect the ATS system to ignore linguistic variation and sensitive attributes as long as they do not affect the answer’s validity. Hence, we would expect variation in these dimensions to have no effect on scores: answer length, language/vocabulary simplicity, alternative spellings/misspellings of non-keywords, grammatical variation, syntactic variation (especially those resembling transfer from a first language), and proxies for sensitive attributes. On the other hand, the system should be able to differentiate proper answers from those aimed at gaming the test (Chin, 2020; Ding et al., 2020).

When grading students on language skills, however, we would expect ATS systems to be only sensitive to the relevant skill. For example, when assessing grammar use, we would expect the system to be sensitive to grammatical errors (from the perspective of the language variety the student is expected to use), but not to the other dimensions mentioned above (e.g., misspellings).

Actors. Relevant experts include teachers of the subjects where the ATS systems will be deployed, linguists, and computer scientists. The stakeholders (students) may be represented by student unions (at the university level) or focus groups comprising a representative sample of the student population.

7 Implications for Policy

There is a mounting effort to increase accountability and transparency around the development and use of NLP systems to prevent them from amplifying societal biases. DOCTOR is highly complementary to the model card approach increasingly adopted⁶ to surface oft hidden details about NLP

models: Developers simply need to list the tested dimensions, metrics, and score on each dimension in the model card. Crucially, reliability tests can be used to highlight fairness issues in NLP systems by including sensitive attributes for the target population, but it is paramount these requirements reflect *local concerns* rather than any prescriptivist perspective (Sambasivan et al., 2021).

At the same time, the ability to conduct quantitative, targeted reliability testing along specifiable dimensions paves the way for reliability standards to be established, with varying levels of stringency and rigor for different use cases and industries. We envision minimum safety and fairness standards being established for applications that are non-sensitive, not safety-critical, and used in unregulated industries, analogous to standards for household appliances. Naturally, applications at greater risks (Li et al., 2020b) of causing harm upon failure should be held to stricter standards. Policymakers are starting to propose and implement regulations to enforce transparency and accountability in the use of AI systems. For example, the European Union’s General Data Protection Regulation grants data subjects the right to obtain “meaningful information about the logic involved” in automated decision systems (EU, 2016). The EU is developing AI-specific regulation (European Commission, 2020): e.g., requiring developers of high-risk AI systems to report their “capabilities and limitations, ... [and] the conditions under which they can be expected to function as intended”. In the U.S., a proposed bill of the state of Washington will require public agencies to report “any potential impacts of the automated decision system on civil rights and liberties and potential disparate impacts on marginalized communities” before using automated decision systems (Washington State Legislature, 2021).

One may note that language in the proposed regulation is intentionally vague. There are many ways to measure bias and fairness, depending on the type of model, context of use, and goal of the system. Today, companies developing AI systems employ the definitions they believe most reasonable (or perhaps easiest to implement), but regulation will need to be more specific for there to be meaningful compliance. DOCTOR’s requirement to explicitly define specific dimensions instead of a vague notion of reliability will help policymakers in this

⁵www.encyclopedia.com/.../education-and-mobility

⁶huggingface.co/models;

github.com/ivylee/model-cards-and-datasheets;

blog.einstein.ai/model-cards-for-ai-model-transparency

regard, and can inform the ongoing development of national (NIST, 2019) and international standards⁷.

While external algorithm audits are becoming popular, testing remains a challenge since companies wishing to protect their intellectual property may be resistant to sharing their code (Johnson, 2021), and implementing custom tests for each system is unscalable. Our approach to reliability testing offers a potential solution to this conundrum by treating NLP systems as black boxes. If reliability tests become a legal requirement, regulatory authorities will be able to mandate independently conducted reliability tests for transparency. Such standards, combined with certification programs (e.g., IEEE’s Ethics Certification Program for Autonomous and Intelligent Systems⁸), will further incentivize the development of responsible NLP, as the companies purchasing NLP systems will insist on certified systems to protect them from both legal and brand risk. To avoid confusion, we expect certification to occur for individual NLP systems (e.g., an end-to-end question answering system for customer enquiries), rather than for general purpose language models that will be further trained to perform some specific NLP task. While concrete standards and certification programs that can serve this purpose do not yet exist, we believe that they eventually will and hope our paper will inform their development. This multi-pronged approach can help to mitigate NLP’s potential harms while increasing public trust in language technology.

8 Challenges and Future Directions

While DOCTOR is a useful starting point to implement reliability testing for NLP systems, we observe key challenges to its widespread adoption. First, identifying and prioritizing the dimensions that can attest a system’s reliability and fairness. The former is relatively straightforward and can be achieved via collaboration with experts (e.g., as part of the U.S. NIST’s future AI standards (NIST, 2019)). The latter, however, is a question of values and power (Noble, 2018; Mohamed et al., 2020; Leins et al., 2020), and should be addressed via a code of ethics and ensuring that all stakeholders are adequately represented at the decision table.

Second, our proposed method of reliability testing may suffer from similar issues plaguing automatic

evaluation metrics for natural language generation (Novikova et al., 2017; Reiter, 2018; Kryscinski et al., 2019): due to the tests’ synthetic nature they may not fully capture the nuances of reality. For example, if a test’s objective were to test an NLP system’s reliability when interacting with African American English (AAE) speakers, would it be possible to guarantee (in practice) that all generated examples fall within the distribution of AAE texts? Potential research directions would be to design adversary generation techniques that can offer such guarantees or incorporate human feedback (Nguyen et al., 2017; Kreutzer et al., 2018; Stiennon et al., 2020).

9 Conclusion

Once language technologies leave the lab and start impacting real lives, concerns around safety, fairness, and accountability cease to be thought experiments. While it is clear that NLP can have a positive impact on our lives, from typing auto-completion to revitalizing endangered languages (Zhang et al., 2020a), it also has the potential to perpetuate harmful stereotypes (Bolukbasi et al., 2016; Sap et al., 2019), perform disproportionately poorly for underrepresented groups (Hern, 2017; Bridgeman et al., 2012), and even erase already marginalized communities (Bender et al., 2021).

Trust in our tools stems from an assurance that stakeholders will remain unharmed, even in the worst-case scenario. In many mature industries, this takes the form of reliability standards. However, for standards to be enacted and enforced, we must first operationalize “reliability”. Hence, we argue for the need for reliability testing (especially worst-case testing) in NLP by contextualizing it among existing work on promoting accountability and improving generalization beyond the training distribution. Next, we showed how adversarial attacks can be reframed as worst-case tests. Finally, we proposed a possible paradigm, DOCTOR, for how reliability concerns can be realized as quantitative tests, and discussed how this framework can be used at different levels of organization or industry.

Acknowledgements

Samson is supported by Salesforce and Singapore’s Economic Development Board under the Industrial Postgraduate Programme. Araz is supported by the NUS Centre for Trusted Internet and Community through project CTIC-RP-20-02.

⁷ethicsstandards.org/p7000

⁸standards.ieee.org/industry-connections/ecpais.html

Broader Impact

Much like how we expect to not be exposed to harmful electric shocks when using electrical appliances, we should expect some minimum levels of safety and fairness for the NLP systems we interact with in our everyday lives. As mentioned in §1, §3, and §7, standards and regulations for AI systems are in the process of being developed for this purpose, especially for applications deemed “high-risk”, e.g., healthcare (European Commission, 2020). Reliability testing, and our proposed framework, is one way to approach the problem of enacting enforceable standards and regulations.

However, the flip side of heavily regulating every single application of NLP is that it may slow down innovation. Therefore, it is important that the level of regulation for a particular application is proportionate to its potential for harm (Daten Ethik Kommission, 2019). Our framework can be adapted to different levels of risk by scaling down the implementation of some steps (e.g., the method and depth in which stakeholder consultation happens or the comprehensiveness of the set of testing dimensions) for low-risk applications.

Finally, it is important to ensure that any tests, standards, or regulations developed adequately represents the needs of the most vulnerable stakeholders, instead of constructing them in a prescriptivist manner (Hagerty and Rubinov, 2019). Hence, DDC-TOR places a strong emphasis on involving stakeholder advocates and analyzing the impact of an application of NLP on the target community.

References

116th U.S. Congress. 2019. Algorithmic Accountability Act of 2019.

Ada Lovelace Institute and DataKind UK. 2021. *Examining the black box: Tools for assessing algorithmic systems*. Technical report.

Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. *Generating natural language adversarial examples*. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Brussels, Belgium. Association for Computational Linguistics.

Amnesty International. 2018. *Toxic Twitter - triggers of violence and abuse against women on Twitter*.

Kent Beck, Mike Beedle, Arie van Bennekum, Alistair Cockburn, Ward Cunningham, Martin Fowler, James

Grønning, Jim Highsmith, Andrew Hunt, Ron Jeffries, Jon Kern, Brian Marick, Robert C. Martin, Steve Mellor, Ken Schwaber, Jeff Sutherland, and Dave Thomas. 2001. *Manifesto for Agile Software Development*.

Yonatan Belinkov and Yonatan Bisk. 2018. *Synthetic and natural noise both break neural machine translation*. In *6th International Conference on Learning Representations*, Vancouver, BC, Canada.

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, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. *On the dangers of stochastic parrots: Can language models be too big?* In *Proceedings of the Conference on Fairness, Accountability, and Transparency*.

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 Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. *Man is to computer programmer as woman is to homemaker? debiasing word embeddings*. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 4349–4357. Curran Associates, Inc.

Brent Bridgeman, Catherine Trapani, and Yigal Attali. 2012. *Comparison of human and machine scoring of essays: Differences by gender, ethnicity, and country*. *Applied Measurement in Education*, 25(1):27–40.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems 33*.

Miles Brundage, Shahar Avin, Jasmine Wang, Haydn Belfield, Gretchen Krueger, Gillian Hadfield, Heidy Khlaaf, Jingying Yang, Helen Toner, Ruth Fong, et al. 2020. *Toward trustworthy AI development: mechanisms for supporting verifiable claims*. *arXiv preprint arXiv:2004.07213*.

California State Legislature. 2020. *California Privacy Rights Act*.

Minhao Cheng, Wei Wei, and Cho-Jui Hsieh. 2019.

- Evaluating and enhancing the robustness of dialogue systems: A case study on a negotiation agent. 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 3325–3335, Minneapolis, Minnesota. Association for Computational Linguistics.
- Monica Chin. 2020. [These students figured out their tests were graded by AI — and the easy way to cheat.](#) *The Verge*.
- Katie Cohen, Fredrik Johansson, Lisa Kaati, and Jonas Clausen Mork. 2014. Detecting linguistic markers for radical violence in social media. *Terrorism and Political Violence*, 26(1):246–256.
- Chris Coleman, Noël Gregg, Lisa McLain, and Leslie W Bellair. 2009. A comparison of spelling performance across young adults with and without dyslexia. *Assessment for effective intervention*, 34(2):94–105.
- Kate Crawford. 2017. The trouble with bias (keynote). *Advances in Neural Information Processing Systems* 30.
- Marina Danilevsky, Kun Qian, Ranit Aharonov, Yannis Katsis, Ban Kawas, and Prithviraj Sen. 2020. [A survey of the state of explainable AI for natural language processing.](#) 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 447–459, Suzhou, China. Association for Computational Linguistics.
- Daten Ethik Kommission. 2019. [Opinion of the data ethics commission.](#) Technical report, Data Ethics Commission of the Federal Government (Germany).
- Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. 2019. [Racial bias in hate speech and abusive language detection datasets.](#) In *Proceedings of the Third Workshop on Abusive Language Online*, pages 25–35, Florence, Italy. Association for Computational Linguistics.
- 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.
- Yuning Ding, Brian Riordan, Andrea Horbach, Aoife Cahill, and Torsten Zesch. 2020. [Don’t take “nswvt-nvakgxp” for an answer —the surprising vulnerability of automatic content scoring systems to adversarial input.](#) In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 882–892, Barcelona, Spain (Online). International Committee on 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.
- Michael Dunn. 2014. Gender determined dialect variation. In *The expression of gender*, pages 39–68. De Gruyter.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. [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.
- Steffen Eger and Yannik Benz. 2020. [From hero to zéro: A benchmark of low-level adversarial attacks.](#) 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 786–803, Suzhou, China. Association for Computational Linguistics.
- EIU. 2020. [Staying ahead of the curve: The business case for responsible AI.](#) Technical report, The Economist Intelligence Unit.
- EU. 2016. [General data protection regulation.](#)
- European Commission. 2020. [On artificial intelligence - a European approach to excellence and trust.](#) Technical report, European Commission.
- European Commission. 2021. [Proposal for a regulation laying down harmonised rules on artificial intelligence \(artificial intelligence act\).](#) Technical report, European Commission.
- FDA. 2021. [Artificial intelligence/machine learning \(ai/ml\)-based software as a medical device \(samd\) action plan.](#) Technical report, U.S. Food & Drug Administration.
- Todd Feathers. 2019. [Flawed algorithms are grading millions of students’ essays.](#) *Vice*.
- Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. [MRQA 2019 shared task: Evaluating generalization in reading comprehension.](#) In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 1–13, Hong Kong, China. Association for Computational Linguistics.
- Batya Friedman and David G Hendry. 2019. *Value Sensitive Design: Shaping Technology with Moral Imagination.* MIT Press.
- Bharath Ganesh. 2018. [The ungovernability of digital hate culture.](#) *Columbia Journal of International Affairs*.

- 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 *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics.
- Siddhant Garg and Goutham Ramakrishnan. 2020. [BAE: BERT-based adversarial examples for text classification](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 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 arXiv:1803.09010*.
- Mor Geva, Yoav Goldberg, and Jonathan Berant. 2019. [Are we modeling the task or the annotator? an investigation of annotator bias in natural language understanding datasets](#). 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 1161–1166, Hong Kong, China. Association for Computational Linguistics.
- Karan Goel, Nazneen Rajani, Jesse Vig, Samson Tan, Jason Wu, Stephan Zheng, Caiming Xiong and Mohit Bansal, and Christopher Ré. 2021. [Robustness Gym: Unifying the NLP evaluation landscape](#). *arXiv preprint arXiv:2101.04840*.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and harnessing adversarial examples](#). In *3rd International Conference on Learning Representations*, San Diego, California.
- Alexa Hagerty and Igor Rubinov. 2019. [Global ai ethics: a review of the social impacts and ethical implications of artificial intelligence](#). *arXiv preprint arXiv:1907.07892*.
- Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, et al. 2018. [Achieving human parity on automatic chinese to english news translation](#). *arXiv preprint arXiv:1803.05567*.
- Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzić, Rishabh Krishnan, and Dawn Song. 2020. [Pretrained transformers improve out-of-distribution robustness](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2744–2751, Online. Association for Computational Linguistics.
- Alex Hern. 2017. [Facebook translates 'good morning' into 'attack them', leading to arrest](#). *The Guardian*.
- Dirk Hovy, Federico Bianchi, and Tommaso Fornaciari. 2020. ["you sound just like your father" commercial machine translation systems include stylistic biases](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1686–1690, Online. Association for Computational Linguistics.
- Dirk Hovy and Shannon L Spruit. 2016. The social impact of natural language processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598.
- Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. 2020. [Social biases in NLP models as barriers for persons with disabilities](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5491–5501, Online. Association for Computational Linguistics.
- IEC. 2020. Household and similar electrical appliances – Safety – Part 1: General requirements. *IEC 60335-1:2020*.
- IEEE. 2017. [ISO/IEC/IEEE International standard - systems and software engineering—vocabulary. ISO/IEC/IEEE 24765:2017\(E\)](#), pages 1–541.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. [Adversarial example generation with syntactically controlled paraphrase networks](#). 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 1875–1885, New Orleans, Louisiana. 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.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. [Is BERT really robust? A strong baseline for natural language attack on text classification and entailment](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8018–8025. AAAI Press.
- Khari Johnson. 2021. [What algorithm auditing startups need to succeed](#). *VentureBeat*.
- Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. 2019. [Learning the difference that makes a difference with counterfactually-augmented data](#). In *International Conference on Learning Representations*.
- Zixuan Ke and Vincent Ng. 2019. [Automated essay scoring: A survey of the state of the art](#). In *International Joint Conference on Artificial Intelligence*. Inter-

- national Joint Conferences on Artificial Intelligence Organization.
- Daniel Khashabi, Tushar Khot, and Ashish Sabharwal. 2020. [More bang for your buck: Natural perturbation for robust question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 163–170, Online. Association for Computational Linguistics.
- Kate Klonick. 2018. [The new governors: The people, rules, and processes governing online speech](#). *Harvard Law Review*, 131(6):1598.
- Julia Kreutzer, Shahram Khadivi, Evgeny Matusov, and Stefan Riezler. 2018. Can neural machine translation be improved with user feedback? In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 92–105.
- Wojciech Kryscinski, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Neural text summarization: A critical evaluation](#). 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 540–551, Hong Kong, China. Association for Computational Linguistics.
- Zachary Laub. 2019. [Hate speech on social media: Global comparisons](#). *Council on Foreign Relations*.
- Kobi Leins, Jey Han Lau, and Timothy Baldwin. 2020. [Give me convenience and give her death: Who should decide what uses of NLP are appropriate, and on what basis?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2908–2913, Online. Association for Computational Linguistics.
- Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. 2019. [Textbugger: Generating adversarial text against real-world applications](#). In *26th Annual Network and Distributed System Security Symposium*.
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020a. [Bert-attack: Adversarial attack against BERT using BERT](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, Online. Association for Computational Linguistics.
- Yanwei Li, Araz Taeihagh, Martin de Jong, and Andreas Klinke. 2020b. [Toward a commonly shared public policy perspective for analyzing risk coping strategies](#). *Risk Analysis*.
- John Markoff. 2013. [Essay-grading software offers professors a break](#). *The New York Times*.
- Alexandria Marsters. 2019. [When Hate Speech Leads to Hateful Actions: A Corpus and Discourse Analytic Approach to Linguistic Threat Assessment of Hate Speech](#). Ph.D. thesis, Georgetown University, Washington, D.C.
- Paul Michel, Xian Li, Graham Neubig, and Juan Pino. 2019. [On evaluation of adversarial perturbations for sequence-to-sequence models](#). 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 3103–3114, Minneapolis, Minnesota. Association for Computational Linguistics.
- John Miller, Karl Krauth, Benjamin Recht, and Ludwig Schmidt. 2020. [The effect of natural distribution shift on question answering models](#). *arXiv preprint arXiv:2004.14444*.
- Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. [Model cards for model reporting](#). In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 220–229.
- Shakir Mohamed, Marie-Therese Png, and William Isaac. 2020. [Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence](#). *Philosophy & Technology*, 33(4):659–684.
- Stacy Nelson. 2003. [Certification processes for safety-critical and mission-critical aerospace software](#). Technical report, NASA Technical Reports Server.
- Khanh Nguyen, Hal Daumé III, and Jordan Boyd-Graber. 2017. [Reinforcement learning for bandit neural machine translation with simulated human feedback](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1464–1474.
- 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.
- NIST. 2019. [U.S. leadership in AI: A plan for federal engagement in developing technical standards and related tools](#). Technical report, National Institute of Standards and Technology.
- Safiya Umoja Noble. 2018. [Algorithms of oppression: How search engines reinforce racism](#). NYU Press.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. [Why we need new evaluation metrics for NLG](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2241–2252, Copenhagen, Denmark. Association for Computational Linguistics.
- Cathy O’Neil. 2016. [Weapons of math destruction: How big data increases inequality and threatens democracy](#). Crown.
- Partnership on AI. 2019. [About ML](#). Technical report, Partnership on AI.

- Barbara Plank. 2016. What to do about non-standard (or non-canonical) language in NLP. *Proceedings of the 13th Conference on Natural Language Processing (KONVENS 2016)*.
- Megha Rajagopalan, Lam Thuy Vo, and Aung Naing Soe. 2018. [How Facebook failed the Rohingya in Myanmar](#). *BuzzFeed News*.
- Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 33–44.
- Bogdana Rakova, Jingying Yang, Henriette Cramer, and Rumman Chowdhury. 2020. [Where responsible ai meets reality: Practitioner perspectives on enablers for shifting organizational practices](#). *arXiv preprint arXiv:2006.12358*.
- Chaitanya Ramineni and David Williamson. 2018. Understanding mean score differences between the e-rater® automated scoring engine and humans for demographically based groups in the GRE® general test. *ETS Research Report Series*, 2018(1):1–31.
- Ehud Reiter. 2018. [A structured review of the validity of BLEU](#). *Computational Linguistics*, 44(3):393–401.
- Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. Generating natural language adversarial examples through probability weighted word saliency. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1085–1097.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. ["why should i trust you?": Explaining the predictions of any classifier](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, page 1135–1144, New York, NY, USA. Association for Computing Machinery.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. [Semantically equivalent adversarial rules for debugging NLP models](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 856–865, Melbourne, Australia. 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.
- Anne-Laure Rousseau, Clément Baudelaire, and Kevin Riera. 2020. [Doctor GPT-3: hype or reality?](#) *Nabla Technologies Blog*.
- Kristin Salaky. 2018. [What standardized tests look like in 10 places around the world](#). *INSIDER*.
- Nithya Sambasivan, Erin Arnesen, Ben Hutchinson, Tulsee Doshi, and Vinodkumar Prabhakaran. 2021. [Reimagining algorithmic fairness in india and beyond](#). In *Proceedings of the 2021 Conference on Fairness, Accountability, and Transparency*.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. [The risk of racial bias in hate speech detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence, Italy. Association for Computational Linguistics.
- Suchi Saria and Adarsh Subbaswamy. 2019. [Safe and reliable machine learning \(tutorial\)](#). ACM Conference on Fairness, Accountability, and Transparency.
- Deven Santosh Shah, H. Andrew Schwartz, and Dirk Hovy. 2020. [Predictive biases in natural language processing models: A conceptual framework and overview](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5248–5264, Online. Association for Computational Linguistics.
- Andrew Smith. 2020. [Using artificial intelligence and algorithms](#).
- Anders Søgaard, Sebastian Ebert, Jasmijn Bastings, and Katja Filippova. 2021. [We need to talk about random splits](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1823–1832, Online. Association for Computational Linguistics.
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. [Learning to summarize from human feedback](#). *arXiv preprint arXiv:2009.01325*.
- Samson Tan and Shafiq Joty. 2021. Code-Mixing on Sesame Street: Dawn of the adversarial polyglots. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Online. Association for Computational Linguistics.
- Samson Tan, Shafiq Joty, Min-Yen Kan, and Richard Socher. 2020. [It's morphin' time! Combating linguistic discrimination with inflectional perturbations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2920–2935, Online. Association for Computational Linguistics.
- Deborah Tannen. 1991. *You just don't understand: Women and men in conversation*. Ballantine books New York.
- Deborah Tannen et al. 2005. *Conversational style: Analyzing talk among friends*. Oxford University Press.
- Viniti Vaish and Teck Kiang Tan. 2008. [Language and social class: Linguistic capital in Singapore](#). In *Annual Meeting of the American Educational Research Association*. American Educational Research Association.

- Claudia Wagner, David Garcia, Mohsen Jadidi, and Markus Strohmaier. 2015. It’s a man’s Wikipedia? Assessing gender inequality in an online encyclopedia. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 9.
- Lily Wakefield. 2020. [Queer people are being forced off social media by trolling and online abuse, searingly obvious report confirms](#). *PinkNews*.
- 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.
- Zeerak Waseem, Smarika Lulz, Joachim Bingel, and Isabelle Augenstein. 2021. [Disembodied machine learning: On the illusion of objectivity in NLP](#). *arXiv preprint arXiv:2101.11974*.
- Washington State Legislature. 2021. [Senate bill SB 5116](#).
- Gerhard Widmer and Miroslav Kubat. 1996. Learning in the presence of concept drift and hidden contexts. *Machine Learning*, 23:69–101.
- Chris Wilkinson, Jonathan Lynch, Raj Bharadwaj, and Kurt Woodham. 2016. [Verification of adaptive systems](#). Technical report, Federal Aviation Administration William J. Hughes Technical Center.
- Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel S Weld. 2021. [Polyjuice: Automated, general-purpose counterfactual generation](#). *arXiv preprint arXiv:2101.00288*.
- Adams Wei Yu, David Dohan, Quoc Le, Thang Luong, Rui Zhao, and Kai Chen. 2018. [Fast and accurate reading comprehension by combining self-attention and convolution](#). In *International Conference on Learning Representations*.
- Yuan Zang, Fanchao Qi, Chenghao Yang, Zhiyuan Liu, Meng Zhang, Qun Liu, and Maosong Sun. 2020. Word-level textual adversarial attacking as combinatorial optimization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6066–6080.
- Huangzhao Zhang, Hao Zhou, Ning Miao, and Lei Li. 2019a. [Generating fluent adversarial examples for natural languages](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5564–5569, Florence, Italy. Association for Computational Linguistics.
- Shiyue Zhang, Benjamin Frey, and Mohit Bansal. 2020a. [ChrEn: Cherokee-English machine translation for endangered language revitalization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 577–595, Online. Association for Computational Linguistics.
- Wei Emma Zhang, Quan Z. Sheng, Ahoud Alhazmi, and Chenliang Li. 2020b. [Adversarial attacks on deep-learning models in natural language processing: A survey](#). *ACM Transactions on Intelligent Systems and Technology*, 11(3).
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019b. [PAWS: Paraphrase adversaries from word scrambling](#). 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 1298–1308, Minneapolis, Minnesota. Association for Computational Linguistics.

Appendix

A Testing Dimensions: Detecting Violent Content on Social Media

In this second case study, we apply DOCTOR for measuring the reliability of a violent content detection system for English social media posts. Although we limit this discussion to the U.S., this is a growing global problem (Laub, 2019) that can lead to deadly outcomes (Rajagopalan et al., 2018). In this hypothetical use case, the NLP system may automatically remove violent content or alert content moderators to potential violations of the social media company’s acceptable use policy. Moderators can decide if specific content should be removed, and if necessary, notify law enforcement to avert pending violence (e.g., threats against individuals, planned violent events). As a result of the 1996 Communications Decency Act⁹, social media platforms have broad latitude (Klonick, 2018) to develop their own policies for acceptable content and how they handle it. In this scenario, the compliance officer of the company developing the system is responsible for making sure it does not discriminate against specific user demographics.

Research has shown that hate speech can lead to hateful actions (Marsters, 2019). In many cases, individuals posted their intents online prior to committing violence (Cohen et al., 2014). When identifying content to remove and especially when involving law enforcement, it is important to distinguish between “Hunters” — those who act — and “Howlers” — those who do not (Marsters, 2019). This is to avoid wrongly detaining individuals who have no intention of committing violence, even if their words are indefensible. Between these extremes, posters may harass, stalk, dox, or otherwise abuse victims from a distance, therefore it is still necessary to flag, remove, and potentially track or document violent content.

Linguistic landscape. We focus solely on English speakers, but we acknowledge that the actual linguistic landscape is much more complex (over 350 languages). Posters on social media may speak English as their first language or as a second language and they often code-switch/-mix. Standard American English is used for business purposes in the U.S. but there are other frequently used language varieties including African American English (AAE), Cajun Vernacular English, and three

different Latinx (Hispanic) vernacular Englishes.

Stakeholder Impact. The key stakeholders that will be impacted are those most often facing violent threats online: minorities, women, immigrants, and the LGBTQ community (Amnesty International, 2018; Ganesh, 2018; Davidson et al., 2019; Wakefield, 2020). Additionally, anyone that posts content on the social media site is a stakeholder. Unfortunately, the very communities that are often the target of violent posts are also often wrongly flagged as posting toxic content themselves due to racial biases present in the training data (Sap et al., 2019; Davidson et al., 2019). Given the risk of harm to victims if the system misses violent posts from hunters or misidentifies legitimate content as violent and notifies law enforcement, it is critical the right balance of false positives and false negatives is achieved in flagging content.

Dimensions. There are two tasks under consideration here: identifying violent content and identifying Hunters who “truly intend to use lethal violence” (Marsters, 2019). In the first task, the system is looking for content that negatively targets a socially defined group. Additionally, the content includes not only hate speech (e.g., profanity, epithets, vulgarity) but also content that incites others to hatred or violence. Since content written in AAE has been shown to be flagged as toxic more often (Sap et al., 2019; Davidson et al., 2019), we must ensure that the system is reliable when encountering dialectal variation. Additionally, due to the casual environment of social media, multilingual speakers often code-switch and code-mix. Hence, we expect variation in these dimensions to have no effect on the system’s predictions: alternative spellings, morphosyntactic variation, word choice, code-mixing, idioms, and references to and manifestations of sensitive attributes and their proxies. However, we must expect the system to be sensitive to in-group and out-group usage of reclaimed slurs so that the in-group usage does not result in a flag while out-group usage result in flagged posts.

When identifying hunters, we may expect the system to be sensitive to uses of first person pronouns, certainty adverbs, negative evaluative adjectives, and modifiers (Marsters, 2019). However, in order to avoid unfairly penalizing vernacular English speakers we should expect the system’s predictions to be equally unaffected by variation in the dimensions listed for the first task.

⁹[fcc.gov/general/telecommunications-act-1996](https://www.fcc.gov/general/telecommunications-act-1996)

Linguistic Phenomena	Orthography	Hyphenation Capitalization Punctuation Reduplication of letters Emojis/emoticons Homonyms Disemvoweling (Eger and Benz, 2020) Homophones (e.g., accept vs. except) (Eger and Benz, 2020) Accidental misspellings (Belinkov and Bisk, 2018) Intentional alternative spellings (e.g., Yas, thru, startin) Open compound concatenation (e.g., couch potato/couchpotato) Dialectal differences (e.g., favor vs. favour) (Ribeiro et al., 2018) Mixing writing scripts (Tan and Joty, 2021) Transliteration
	Morphology	Grammatical gender shifts Grammatical category (Tan et al., 2020) Dialectal differences (Tan et al., 2020) Clitics
	Lexicon	Dialectal variation (e.g., fries vs. chips) Synonyms/Sememes (Zang et al., 2020) Vocabulary simplicity/complexity Cross-lingual synonyms (Tan and Joty, 2021) Loanwords
	Semantics	Idioms (e.g., finer than frog hair)
	Syntax	Matching number and tense Word/phrase order (especially for languages without strict word ordering) Prepositional variation (e.g., stand on line vs. stand in line) Syntactic variation (Iyyer et al., 2018) Sentence simplicity/complexity Code-mixing (Tan and Joty, 2021)
	Discourse & Pragmatics	Register (e.g., formality) Conversational style (involvement/considerateness) (Tannen et al., 2005) Discourse markers / connector words Cross-cultural differences Code-switching
Sensitive Attributes	Gender Identity	Gender pronouns Names Reclaimed slurs Genderlects (Tannen, 1991; Dunn, 2014)
	Race	Names Reclaimed slurs Race-aligned language varieties
	Age	Age/generation-aligned language styles (Hovy et al., 2020)
	Religion	Names Reclaimed slurs
	Sexual Orientation	Reclaimed slurs
	Disability status	Associated adjectives (Hutchinson et al., 2020)
	Place of origin	Location names (e.g., cities, countries) Figures of speech
Proxies	Geographic locations (for ethnicity, socioeconomic status)	
Malicious Attacks	Black-box	Rule-based (Alzantot et al., 2018; Jin et al., 2020) Model-based (Garg and Ramakrishnan, 2020; Li et al., 2020a)
	Gradient-based	HotFlip (Ebrahimi et al., 2018), Universal Triggers (Wallace et al., 2019)
	Policy-based	Adversarial negotiation agent (Cheng et al., 2019)

Table 1: Taxonomy of possible dimensions with references to linguistics literature and existing adversarial attacks that could be used as worst-case tests. Linguists are best equipped to decide which linguistic phenomena are high priority for each use case, ethicists for sensitive attributes, and NLP practitioners for malicious attacks.