

---

# CommonsenseQA 2.0: Exposing the Limits of AI through Gamification

---

Alon Talmor<sup>1</sup> Ori Yoran<sup>1,2</sup> Ronan Le Bras<sup>1</sup> Chandra Bhagavatula<sup>1</sup>  
Yoav Goldberg<sup>1</sup> Yejin Choi<sup>1,3</sup> Jonathan Berant<sup>1,2</sup>

<sup>1</sup>The Allen Institute for AI, <sup>2</sup>Tel-Aviv University, <sup>3</sup>University of Washington  
{ronanlb,chandrab,yoavg,yejinc,jonathan}@allenai.org  
{alontalmor,oriy}@mail.tau.ac.il

## Abstract

Constructing benchmarks that test the abilities of modern natural language understanding models is difficult – pre-trained language models exploit artifacts in benchmarks to achieve human parity, but still fail on adversarial examples and make errors that demonstrate a lack of common sense. In this work, we propose gamification as a framework for data construction. The goal of players in the game is to compose questions that mislead a rival AI, while using specific phrases for extra points. The game environment leads to enhanced user engagement and simultaneously gives the game designer control over the collected data, allowing us to collect high-quality data at scale. Using our method we create CommonsenseQA 2.0, which includes 14,343 yes/no questions, and demonstrate its difficulty for models that are orders-of-magnitude larger than the AI used in the game itself. Our best baseline, the T5-based UNICORN with 11B parameters achieves an accuracy of 70.2%, substantially higher than GPT-3 (52.9%) in a few-shot inference setup. Both score well below human performance which is at 94.1%.

## 1 Introduction

Evaluating models for natural language understanding (NLU) has become a frustrating endeavour. On the one hand, state-of-the-art models achieve human parity on a wide range of NLU benchmarks, which are now generated at unprecedented rate. On the other hand, recurring findings reveal that such models are brittle when provided with out-of-domain [1, 2, 3] or adversarial examples [4, 5, 6, 7, 8, 9, 10], and make silly mistakes that contradict basic notions of human common sense. It is almost as if NLU models learn to solve datasets without solving the underlying tasks. This status quo calls for new innovations in the area of benchmark design.

To generate more challenging datasets, recent research proposed algorithmic methods for bias reduction, such as adversarial filtering [11, 12, 13, 14, 15]. While helpful in mitigating bias and reducing dataset artifacts in data that has already been generated, such approaches are fundamentally post-hoc, that is, they can only remove examples, but do not direct the example generation pipeline to create difficult and diverse examples. Indeed, recent research highlights the importance of diverse and hard instances for robust out-of-domain generalization [16, 17]. Thus, a more fundamental solution should be integrated directly into the data creation process. Unfortunately, the vast majority of existing benchmarks are based on crowdsourcing, which are by and large static and non-interactive, leading to datasets that are plagued with shortcuts and artifacts [18, 19, 20, 4]. At the same time, humans can still reliably compose new examples that demonstrate the failing of current AI systems.

In this paper, we propose gamification as a framework for data creation that can tap into human intelligence more effectively. Our framework shares the spirit of model-and-human-in-the-loop

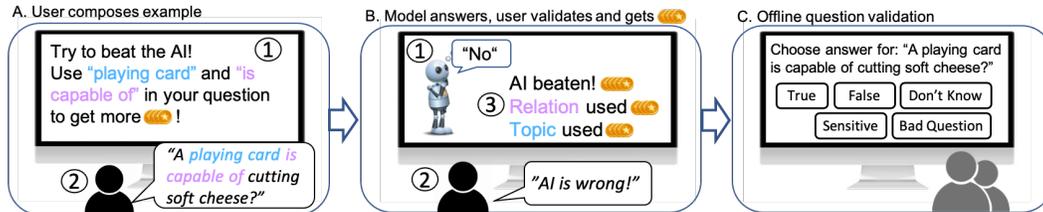


Figure 1: An overview of our approach for data collection through gamification.

approaches, which seek dynamic benchmark creation [21], but provides a game environment that enhances user engagement, which in turn leads to high-quality data and low costs.

Fig. 1 shows a high-level overview of our game, which focuses on collecting yes/no questions or assertions. At a high-level, a player is asked to author a yes/no question, is then shown the answer from the AI, and then marks whether the AI was correct or not. The goal of the player is to earn points, which are used as a flexible vehicle for steering the behaviour of the player. First, points are given for beating the AI, that is, authoring questions where the AI is incorrect. This incentivizes the player to ask difficult questions, conditioned on its understanding of the AI capabilities. Second, the player gets points for using particular phrases in the question. This provides the game designer control to skew the distribution of questions towards topics or other phenomena they are interested in. Last, questions are validated by humans, and points are deducted for questions that do not pass validation. This pushes players to author questions with broad agreement among people. Similar to prior work [21, 22, 23], a model is trained on collected examples, which prevents players from using the same types of questions repeatedly, as the re-trained model quickly adapts to such questions.

Using our method, we collected COMMONSENSEQA 2.0 (CSQA2) containing 14,343 yes/no questions. We extensively evaluate models on CSQA2, experimenting with pre-trained models, fine-tuned models, and reading comprehension (RC) models that utilize web snippets extracted from Google search on top of the question itself.

Despite using RoBERTA-Large [24] as the model in the loop, the resulting dataset proves to be challenging for state-of-the-art models that are orders of magnitude larger. Compared to human performance of 94.1%, the best performance the models achieve is 70.2% by UNICORN-11B [25], substantially better than 67.8% by T5-11B [26]. Notably, GPT-3 [1] achieves only 52.9%, despite being much larger (but without fine-tuning), even after extensive prompt engineering.

Comprehensive analyses offer new insights into the limitations of current neural models; GPT-3 is relatively weaker in causal reasoning, while fine-tuned UNICORN is weaker in size comparisons. In addition, we find that none of the models are robust against relatively simple linguistic or logical variations of the same question.

To summarize, our contributions are:

1. CSQA2: A challenging and diverse new QA dataset, containing 14,343 examples.
2. A new framework for data creation via gamification alongside a model-in-the-loop approach.
3. An empirical evaluation of state-of-the-art pre-trained language models on CSQA2, showing that humans substantially outperform current models.

Our dataset and code are available at <http://allenai.github.io/csqa2>.

## 2 Data Collection through Gamification

We now describe our game environment for eliciting hard and diverse yes/no questions (or assertions). We first describe the game itself (§2.1), and then detail the dataset construction procedure including measures for quality assurance (§2.2).

### 2.1 The Game

Our premise is that a game environment can create an enjoyable experience for players, while enabling control over the quality and type of data obtained through the game’s point system. We now describe

the details of the different parts of the game, which are illustrated in Fig. 1: (a) controlled question generation, (b) model-in-the-loop, and (c) question validation. Full print-screens from the actual game are presented in Fig. 2.

**Controlled question generation** Examples generated in this work are yes/no questions (or assertions) with their correct answer. We use questions since they are a natural medium for players to test capabilities of the AI, but other formats are also possible.

In this step, players are asked to author questions, and get points if the AI errs on the authored questions. This pushes players to generate difficult questions conditioned on their understanding of the abilities of the AI. To give the game designer control over the topic of the question and the reasoning skills required for answering it, we present the player with two phrases, a *topic prompt* and a *relational prompt*, and points are given for incorporating these phrases in the question. For example, in Fig. 1 the topic prompt is “*playing card*”, and the relational prompt is “*is capable of*”, and indeed the player uses them in “*A playing card is capable of cutting soft cheese*”. The prompts not only provide control, but also spark the creativity of the player, which can author the question around tangible concepts.

For topic prompts, we use CONCEPTNET [27], a popular knowledge-graph containing information on commonsense concepts and relations. We choose the 1,875 highest ranking concepts in ConceptNet, and in each question sample one concept uniformly at random. For relational prompts, we use a manually-constructed list of 33 phrases, designed to target a wide set of commonsense reasoning skills, such as understanding meronymy, causality, plausibility, and more. Table 1 shows the full list of relational prompts, their categories and their frequency. While our choice for prompts steers collected data towards commonsense questions, other prompts can be used to get a different distribution.

**Model-in-the-loop** After a player formulates a question (*A2* in Fig. 1), e.g., “*A playing card is capable of cutting soft cheese*”, the question is answered by the model (*B1*), and the user gets points if the model is wrong. A potential issue is that players identify a weakness in the current model and only ask questions that target that particular weakness. Thus, similar to past work [28, 22, 21], we use a model-in-the-loop and re-train the model during data collection. Thus, if players generate many questions exploiting a particular model blind spot, the blind spot will be fixed when the model is re-trained. This leads to more diversity in the collected data.

In our game, the initial model is ROBERTA-LARGE [24], fine-tuned on two datasets of binary (yes/no) question, using standard multi-task training: (a) 50K examples from TWENTY QUESTIONS (20Q),<sup>1</sup> a question answering dataset which focuses on commonsense questions such as “*Does an aircraft fly?*” (true) and “*Do everyone have an alarm?*” (false); and 73K automatically-generated yes/no questions from CONCEPTNET triplets, where we transform triplets to assertions with a manually-constructed template for each CONCEPTNET relation. For example, for the triplet (“*wheel*”, “*part-of*”, “*car*”) we generate “*a wheel is part of a car*” (true).

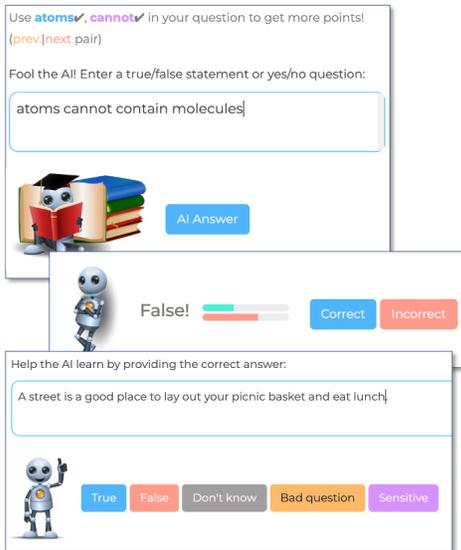


Figure 2: Print-screens of the different game parts.

Type	RELATIONAL PROMPTS	%
Taxonomy / other	is, part of, has, have, is a	20.0
Capable of	is capable of, can, cannot	14.9
Causality	before, after, because, causes	11.1
Plausibility	all, some, at least one, at least two,	
	most, none, exactly, few	10.7
Always-Never	always, almost always, sometimes,	
	almost never, never	10.1
Sizes	larger than, smaller than, same size as	8.9
Conditional	if, only if	6.4
Sequence	done in this order, ordered like this	2.8
Prompt not used		14.4

Table 1: Relational prompts and their frequencies.

<sup>1</sup><https://github.com/allenai/twentyquestions>

We re-trained the model five times when the number of (unvalidated) questions collected reached 1,000, 2,000, 5,000, 10,000 and 20,000. The user then gives immediate feedback to the model on whether the model was correct or incorrect, and receives points accordingly (*B2*, *B3* in Fig. 1).

**Question validation** The last step ensures that composed questions are ones that most people would agree on. Validators answer questions independent of players, and if their answer is different from the answer provided by the question author, then points are deducted from the player, pushing players to ask good questions that can also be answered by other people. This step is also crucial since we train models on collected data and need to make sure examples are not noisy.

In more detail, validators receive a proposed question (*C* in Fig. 1), and choose a *validation label*. Validation labels indicate if the answer is “*True*” or “*False*”. Other validation labels include “*Don’t know*” to indicate they were not able to find an answer, “*Bad question*” for questions that do not make sense, and “*Sensitive*” for questions that involve sensitive topics. Questions are validated multiple times, and a model is trained to determine the final *gold label* based on the validators’ selections. We provide more detail on this model in §2.2. To incentivize players to provide truthful answers, we manually label 2,753 questions, and test the users on these questions in 10% of the validations. If players provides an incorrect answer for these questions, they lose 1 point, otherwise they gain 2 points for each validation performed in the game.

**Details of point system** Beating the AI (asking a question the AI gets wrong) grants the players 5 points plus 4 points if the relational prompt was used and an additional 4 points if the topic prompt was used. If the AI is correct, the user receives a default of 3 points for the effort invested. If a question is discarded after quality assurance, 3 points are deducted from the player who authored the question, otherwise, if the final answer is different than the one provided by the player (*B2*), 2 points are deducted. This prevents players from stating that the model is incorrect to earn more points.

## 2.2 Quality Assurance and Dataset Construction

We now describe the details of collecting CSQA2 and the measures taken for quality assurance.

**Automatic question verification** To automatically filter noisy examples, we manually label 2,753 with the labels “*True*”, “*False*”, and “*Bad Question*” (collapsing “*Bad Question*”, “*Don’t know*” and “*Sensitive*” above). We then train a linear model to predict the gold label. Features in this linear model include: (a) *Validator features*: The conjunction of validation label, validator accuracy, and validator experience (number of validated questions). For example the feature `Label:True, Acc:High, Exp:Low` indicates that a validator that has high accuracy w.r.t the gold labels, but little experience, indicated that the answer to the assertion is “*True*”; (b) *Player features*: Since players also serve as validators (see below), we featurize players in a similar manner to validators but also includes the prediction of the model-in-the-loop (yes or no).

After training the question validation model, we discard any example that is classified as “*Bad Question*” or questions where the confidence of the verification model is low. Then, we label the question according to the prediction of the question verification model. Out of the questions that pass this verification step, 92% are labeled correctly in the development set of the verification model. Further steps for improving quality assurance are described below.

**Crowdsourcing questions** We used Amazon Mechanical Turk (AMT) workers to compose and validate questions. Workers were designated as *players* 70% of the time, and *validators* the rest of the time. Players receive 4.4\$ when they reached 300 points. To increase the quality of the dataset, we only keep questions composed by workers that have a validation accuracy of at least 60% on the expert-labeled examples, and that less than 30% of the questions they author are discarded by the automatic question verification model.

**Adding Google snippets** A possible strategy of malicious players is to find long-tail knowledge facts on the web and ask the model if the fact is correct or not. For example, “*Lobsters taste with their feet*”. To circumvent this, we issue a web query to Google search for every question and take the top-100 returned snippets, as well as the featured Google snippet, if it exists. We then filter out questions for which there is a snippet with a contiguous span that has very low edit-distance to the authored question. This process leads to the filtering of 421 questions from the final dataset. Aside

for quality assurance, in §5 we also investigate whether using Google snippets as additional context can improve the performance of models on CSQA2.

Overall, from the 144,682 questions that were authored, we selected a total of 14,343 examples. The average total cost per question is \$0.17 (compared to \$0.33 in CommonsenseQA 1.0 [29]).

### 3 Game Analysis

To increase the motivation of players to beat the AI, we experimented with various forms of user feedback. We now analyze how such feedback affects the collected data.

**Feedback on performance** We experimented with showing players a bar-chart indicating (a) the fraction of authored question where the player beat the AI (AI beat-rate), (b) the fraction of questions that passed the verification model, and (c) the fraction of correct validations w.r.t the expert-labeled set described above (Fig. 3, left): under 15% in red, under 30% in yellow, and above 30% in green. Interestingly, the average AI beat-rate increased from 26.8% in the 10 days before the introduction of the feedback feature, to 35% in the 10 days immediately after the change.

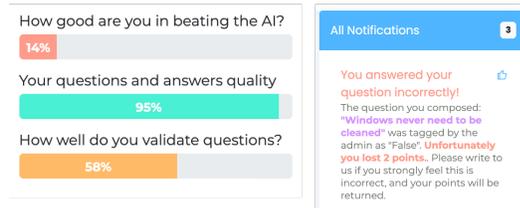


Figure 3: Feedback given to users during the game. Left: metrics on average daily player performance. Right: notification on a bad question that leads to point deduction.

**Validation feedback** We provided daily notifications to players, where we showed them questions where validators changed the answer, or questions that were discarded because they were predicted by the automatic question verification model to be a “Bad Question” (Fig. 3, right). Moreover, as mentioned in §2.1, we introduced a 3-point penalty for such cases. Notifications led to a 4.8% absolute increase in average positive validations (i.e., the validators agree with the yes/no answer given by the player). However, this also led to questions that are easier for the model – Fig. 4 (dashed black line) shows that after notifications were introduced, the average accuracy of models on questions increased, for example, UNICORN-11B [25] improved from 0.66 average accuracy to almost 0.70. This is since players are more hesitant to author questions that may be too difficult for validators. Our analysis emphasizes how subtle design choices in the game can lead to significant ramifications on players’ motivation and data quality.

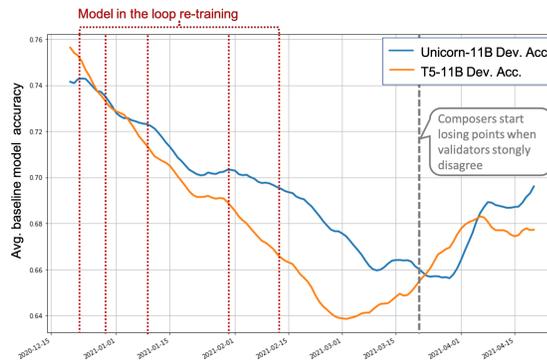


Figure 4: Baselines acc. over question composition time.

**Model re-training** We re-trained the model-in-the-loop five times during the data collection period. Fig. 4 shows in red the dates in which the model-in-the-loop was trained, as well as the performance (Y-axis) of our best performing baselines, UNICORN [25] and T5 [26], on development set questions during those dates. Performance drops significantly and consistently from 0.74 to 0.66 as the model-in-the-loop, RoBERTa-Large, is re-trained, indicating that a stronger model drives players to compose more challenging questions, and that the rise in difficulty is shared across pre-trained models.

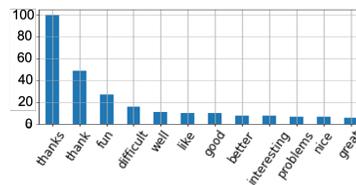


Figure 5: Sentiment words in comments.

**The fun factor** At the end of each playing session, players were encouraged to leave feedback. We selected sentiment words out of the 100 most frequently used words in the comments, shown in Fig. 5. We find that users enjoy the game and mostly use positive sentiment words. Fun is an important factor in encouraging high engagement, allowing us to select annotators that are better at beating the AI, while maintaining a low average cost.



<i>Skill</i>	<i>Description (Example)</i>	<i>%</i>
Capable of	Whether an object is capable of performing an action (“A watch is capable of telling the past time”)	24.5
Long-tail knowledge	The question contains factual long-tail information (“Washington DC is located further south than Washington State”)	24.5
Plausibility	Quantifiers or always-never relations (“The peak of a mountain almost always reaches above the tree line”)	23.6
Comparison	Comparison between two objects (“The end of a baseball bat is larger than the handle”)	16.4
Physical	Physical commonsense (“Do you build the walls on a house before putting on the roof?”)	13.6
Causality	Cause and effect relations (“If you get into an accident because you have been drinking alcohol you will be arrested?”)	13.6
Temporal	Temporal understanding (“None had ever reached the top of Mount Everest before 1977?”)	10.0
Negation	The question includes a negation phrase (“A mock trial is something with no legal consequence”)	9.1
Strategy	Reasoning steps are implicit and should be inferred using a strategy (“Blood banks almost never take cash or checks as deposits”)	9.1
Event chain	Question is about order of events (“Putting on shoes is done in this order normally: person ties shoelaces then slips shoes onto feet”)	3.6

Table 3: Skills and their frequency in the analyzed data (each example can be annotated with multiple skills).

**Reasoning skills** To analyze the types of reasoning skills needed to correctly answer questions in CSQA2, we randomly sampled 110 examples from the development set and performed the following analysis. For each question, we manually annotated the types of commonsense skills a human would potentially use to correctly answer the question. We annotate multiple skills per question, such that the average question includes 1.48 skills. When annotating the skills, we focus on phenomena that are frequent in our questions, such as plausibility, causality, and negation. Table 3 presents the skill categories used, their frequency in the analyzed examples, a description, and an example.

## 5 Experimental Evaluation

**Experimental setup** We split the data into a training/development/test set containing 9, 282/2, 544/2, 517 examples, respectively. We preform a *topic prompt split* where the training-set is disjoint from the development and test sets in terms of topic prompts.

**Human evaluation** To test human accuracy, we created a batch of questions for validation (as explained in §2.1), and asked players to validate questions they did not previously see without knowing the feedback given by the player nor the model prediction. Humans obtained an average accuracy of 90.3% per question, and when using the majority vote per question reached an accuracy of 94.1% when comparing to the final answers of the questions (Table 4).

To empirically evaluate that CSQA2 is indeed challenging for state-of-the-art models, we experiment with multiple baselines, as described below.

### 5.1 Baselines

**T5** [26] is a text-to-text model built on top of the transformer architecture [30] pretrained using masked language modeling [31] on a large collection of NLP benchmarks. We conduct experiments with two model sizes, namely 770 million parameters (T5-LARGE) and 11 billion parameters (T5-11B), fine-tuned on the training set of CSQA2.

**UNICORN** [25] is a pretrained commonsense reasoning model obtained by taking T5 and multi-task training on RAINBOW, a suite of commonsense benchmarks. RAINBOW combines commonsense reasoning about social situations and physical interactions, as well as reasoning about the most plausible explanation or continuation of everyday narratives. Pre-trained on the RAINBOW benchmark, UNICORN offers an efficient commonsense reasoning model ready to be fine-tuned on other downstream commonsense tasks. It was shown to reach state-of-the-art performance on 9 different commonsense datasets, including COMMONSENSEQA 1.0. Similar to T5, we report results for UNICORN-LARGE and UNICORN-11B, with 770 million and 11 billion parameters, respectively, fine-tuning models on the training set.

**GPT-3** [1] is a 175B-parameter autoregressive language model trained on a large corpus of web text. While we fine-tune the previous baselines on the full CSQA2 training set, we evaluate GPT-3 in a few-shot inference setting. Namely, we provide a prompt that contains example questions with their respective answers, followed by a question of interest. We then ask GPT-3 to generate

	<i>Dev</i>	<i>Test</i>
GPT-3	58.4	52.9
T5-LARGE	53.8	54.6
UNICORN-LARGE	56.4	54.9
T5-11B	68.5	67.8
UNICORN-11B	<b>69.9</b>	<b>70.2</b>
Human	94.1	-

Table 4: Development and test accuracies (%) on CSQA2.

a text completion that corresponds to the answer of the question. In §5.2, we report experiments where the prompt contains  $K = 5$  *randomly-selected* (from the training set) example questions. We experimented with multiple values of  $K$ , as well as a variety of other heuristics for choosing the few-shot examples to be included in the prompt. We provide full details of the prompt format, and the results of these experiments in the supplementary material.

## 5.2 Experimental Results

Table 4 shows results of all models on the development and test sets of CSQA2. T5-LARGE and T5-11B achieve a test accuracy of 54.6% and 67.8%, respectively. UNICORN-LARGE and UNICORN-11B obtain a test accuracy of 54.9% and 70.2%, respectively. This illustrates how increasing model size leads to dramatic improvements in commonsense capabilities, and that fine-tuning on RAINBOW leads to a sizable increase for the 11B models – from 67.8% to 70.2%. GPT-3 achieves a test accuracy of 52.9%, well below the T5 and UNICORN baselines, despite being much larger (but without fine-tuning). For all models, performance is significantly lower than human performance, which is at 94.1%, leaving ample room for future model improvements.

**Augmenting questions with Google snippets** In the results above, knowledge needed for answering questions must be encoded in the model weights. We now test whether adding Google search results as additional context improves the performance of our T5 and UNICORN baselines. Specifically, we issue a Google web query for each question, and add the top- $k$  snippets as additional context before the question. Table 5 shows the results of this experiment. We see that snippets from Google search queries consistently improve the quality of the predictions of the T5 and UNICORN models, from 67.8 to 72.5 for T5-11B and from 70.2 to 73.3 for UNICORN-11B. This shows that some knowledge on the web is useful for answering questions from CSQA2.

	<i>Dev</i>	<i>Test</i>
T5-11B ( $k = 1$ )	73.5	<b>73.9</b>
UNICORN-11B ( $k = 1$ )	73.9	73.4
T5-11B ( $k = 5$ )	<b>74.1</b>	72.5
UNICORN-11B ( $k = 5$ )	74.0	73.3

Table 5: Development and test accuracies (%) on CSQA2, when including Google search snippets in the input, where  $k$  indicates the number of snippets prepended before the question.

## 6 Model Analysis

**Baseline analysis** To analyze baselines on different skills, we examine their performance on different sets of relational prompts that correspond to similar skills (Table 6). For example the “*has*”, “*have*”, “*part of*” and “*is a*” relational prompts correspond to meronymy/hypernymy, “*because*” and “*causes*” correspond to causality, and “*smaller than*”, “*larger than*” and “*same size as*”, correspond to size comparison. We only consider cases where the player used the prompt when composing the question. For skills that cannot be inferred from the relational prompt, such as whether the question contains factual long-tail knowledge, we manually annotate questions in which the skill exists.

We find that T5 and UNICORN do well on questions that include meronymy or hypernymy (Table 6). T5-11B and UNICORN-11B outperform GPT-3 by 19.4 and 18.0 points without Google snippets, and 24.4 and 25.0 with 5 Google snippets, respectively. Accuracy for questions that require causality are higher by 6 points for UNICORN-11B (79.3) compared to T5-11B (73.3), outperforming GPT-3 (49.1) by more than 30 points. Accuracy for examples that require size comparison and long-tail factual knowledge is lower for T5 and UNICORN, such that accuracy for UNICORN-11B is 63.1 and 60.8, respectively. Adding Google snippets increases performance on questions that require long-tail factoid knowledge by as much as 22.5 points.

Category	Examples	%	T5-11B	Uni-11B	T5-11B ( $k = 5$ )	Uni-11B ( $k = 5$ )	GPT-3
Meronymy/ Hypernymy	Q: Scales are an important part of music and fish (A: <i>yes</i> )	14.2	<b>75.0</b>	<b>73.6</b>	80.0	80.6	55.6
	Q: The Earth is a planet that is made primarily of air and helium? (A: <i>no</i> )						
Causality	Q: Saltwater is a chemistry solution because it combines water and salt (A: <i>yes</i> )	4.6	<b>73.3</b>	<b>79.3</b>	82.8	85.3	<b>49.1</b>
	Q: Drinking milk causes bones to become weaker (A: <i>no</i> )						
Long-tail factoid knowledge	Q: Independence is a place in California and Utah (A: <i>yes</i> )	24.5	61.8	<b>60.8</b>	77.5	<b>83.3</b>	56.9
	Q: A male seahorse cannot give birth (A: <i>no</i> )						
Size comparison	Q: The outmost layer of an onion is larger than the inner layers are? (A: <i>yes</i> )	9.2	<b>60.1</b>	<b>63.1</b>	64.4	65.2	59.2
	Q: Electrons are smaller than mesons (A: <i>no</i> )						

Table 6: Qualitative error analysis. For each category, we provide two examples and the accuracy.

**Contrast sets** To examine the consistency of our baselines, we created contrast sets [8, 9], that is, we took 60 questions from the development set and authored additional questions that are minor perturbations of the original question. For example, for the original question “*A bird has 3 wings.*”, we add “*A bird has wings*” and “*A bird has only one wing*”, testing if the model understand the exact number of wings birds have.

	Avg.	EM
UNICORN-11B	<b>68.8</b>	<b>18.3</b>
T5-11B	<b>68.8</b>	11.7
GPT-3	53.6	3.3

Table 7: Diagnostic-set acc.

Overall, we authored 276 questions (4.6 additional examples per original question). Table 7 presents model performance on the contrast set with two evaluation metrics: the average score on all questions (including the original ones), and a *consistency* metric, which checks if a model answers correctly all the questions in each contrast set. All models perform poorly in the consistency setup, UNICORN-11B achieving a score of 18.3, outperforming T5-11B (11.7). GPT-3 scores a close-to-zero result of 3.3.

**GPT-3 analysis** To better understand GPT-3’s behaviour, we replace our yes/no prompts with questions that have free-form answers. We then manually label whether GPT-3’s free form answers (FF) reflect understanding of the question and answer (see examples in Table 8). We find that in 10% of the cases GPT-3 answers “*Unknown*”. Of the remaining questions, only 61% of the answers reflect agreement with the original model trained with yes/no prompts. Of the questions where GPT-3’s answer agrees with its original prediction, we add the word “*why*” as a prefix to the question and find that only 36.4% of the explanations make sense. This hints to the lack of consistency in the representation of knowledge of GPT-3.

Q: Question.	FF: GPT-3 free-form prediction.	Why: GPT-3 generative prediction when adding the word “ <i>why</i> ” as a prefix to the question.
Q: “ <i>The office is capable of making people laugh?</i> ”	FF: “ <i>Yes</i> ” ✓	Why: “ <i>The office is capable of making people laugh because it is a comedy.</i> ” ✓
Q: “ <i>Some people don’t seem to mind not breathing for an hour?</i> ”	FF: “ <i>They are called free divers.</i> ” ✗	Why: “ <i>Because they are dead.</i> ” ✗
Q: “ <i>When water freezes, does it get hard?</i> ”	FF: “ <i>No, water expands when it freezes.</i> ” ✗	Why: “ <i>Water expands as it freezes, so it doesn’t get hard.</i> ” ✗

Table 8: Analysis of GPT-3 predictions.

**Google snippets analysis** We analyzed GPT-3 performance on questions whose string was found in a Google web snippet, but for which the answer is “*no*”. Interestingly, we find that in many of these cases GPT-3 does not distinguish between imaginary or counterfactual quotes and real-life facts, answering “*yes*” to questions such as “*Is it true that all the leaves are brown and the sky is gray?*”, “*Is it true that shadows have shadows?*” and “*Is it true that You’ve Got to Hide Your Love Away?*”. This alludes to the challenges involved in separating factual from imaginary content on the web.

## 7 Related Work

Our work joins a line of previously introduced benchmarks targeting commonsense reasoning [12, 13, 29, 15] as well as probing for other reasoning abilities [32, 33, 34, 35, 36, 37, 38]. An emerging trend in NLP dataset creation is the use of a model-in-the-loop when humans compose examples: a model is used either as a post-hoc filter or during annotation. Examples for this include [39, 40, 41, 12, 42, 34, 43, 38, 44], with recent works featuring re-training during dataset collection (Beat-the-AI [22], StrategyQA [23] and Dynabench [21]). This approach makes the resulting dataset more challenging to current models, driving research to tackle new problems. Gamification, the collection of data through a fun interactive design, was attempted by prior work for multi-choice question answering [45], image labeling [46] and protein folding [47]. Contemporarily, Fool Me Twice [28] introduced a game for collecting entailment data through multi-player interaction. Our work focuses on gamifying the interaction between a human and a model.

## 8 Conclusion

In this work, we propose gamification as a general framework for creating diverse and challenging NLU benchmarks. We use this framework to collect CSQA2, a new benchmark that contains 14,343 yes/no questions. We perform a detailed analysis of CSQA2, which elucidates the unique properties of our dataset, and thoroughly evaluate on a strong suite of baselines. We find that the best model, UNICORN-11B, achieves an accuracy of 70.2%, dozens of points lower than human accuracy. We argue that gamification is a promising approach for creating challenge sets that expose the weaknesses of current state-of-the-art models.

## 9 Acknowledgments

We thank our colleagues at The Allen Institute of AI. This research was partially supported by The Yandex Initiative for Machine Learning and the European Union’s Seventh Framework Programme (FP7) under grant agreement no. 802800-DELPHI.

## References

- [1] 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. Language models are few-shot learners. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [2] Alon Talmor and Jonathan Berant. MultiQA: An empirical investigation of generalization and transfer in reading comprehension. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4911–4921, Florence, Italy, 2019. Association for Computational Linguistics.
- [3] Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 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, 2019. Association for Computational Linguistics.
- [4] Robin Jia and Percy Liang. 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, 2017. Association for Computational Linguistics.
- [5] Pranav Rajpurkar, Robin Jia, and Percy Liang. 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, 2018. Association for Computational Linguistics.
- [6] Xiaoyong Yuan, Pan He, Qile Zhu, and Xiaolin Li. Adversarial examples: Attacks and defenses for deep learning. *IEEE transactions on neural networks and learning systems*, 30(9):2805–2824, 2019.
- [7] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 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, 2018. Association for Computational Linguistics.
- [8] 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. Evaluating models’ local decision boundaries via contrast sets. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online, 2020. Association for Computational Linguistics.
- [9] Divyansh Kaushik, Eduard H. Hovy, and Zachary Chase Lipton. Learning the difference that makes A difference with counterfactually-augmented data. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [10] Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Madry. Adversarially robust generalization requires more data. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors,

*Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 5019–5031, 2018.

- [11] Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew E. Peters, Ashish Sabharwal, and Yejin Choi. Adversarial filters of dataset biases. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 1078–1088. PMLR, 2020.
- [12] Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. SWAG: A large-scale adversarial dataset for grounded commonsense inference. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104, Brussels, Belgium, 2018. Association for Computational Linguistics.
- [13] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy, 2019. Association for Computational Linguistics.
- [14] Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. Social IQa: Commonsense reasoning about social interactions. 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 4463–4473, Hong Kong, China, 2019. Association for Computational Linguistics.
- [15] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8732–8740, 2020.
- [16] Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 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, 2020. Association for Computational Linguistics.
- [17] Lifu Tu, Garima Lalwani, Spandana Gella, and He He. An empirical study on robustness to spurious correlations using pre-trained language models. *Transactions of the Association for Computational Linguistics*, 8:621–633, 2020.
- [18] Danqi Chen, Jason Bolton, and Christopher D. Manning. A thorough examination of the CNN/Daily Mail reading comprehension task. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2358–2367, Berlin, Germany, 2016. Association for Computational Linguistics.
- [19] Sewon Min, Eric Wallace, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi, and Luke Zettlemoyer. Compositional questions do not necessitate multi-hop reasoning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4249–4257, Florence, Italy, 2019. Association for Computational Linguistics.
- [20] Karl Moritz Hermann, Tomáš Kociský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. Teaching machines to read and comprehend. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1693–1701, 2015.
- [21] 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. Dynabench: Rethinking benchmarking in NLP. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4110–4124, Online, 2021. Association for Computational Linguistics.

- [22] Max Bartolo, Alastair Roberts, Johannes Welbl, Sebastian Riedel, and Pontus Stenetorp. Beat the AI: Investigating adversarial human annotation for reading comprehension. *Transactions of the Association for Computational Linguistics*, 8:662–678, 2020.
- [23] Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *arXiv preprint arXiv:2101.02235*, 2021.
- [24] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [25] Nicholas Lourie, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Unicorn on rainbow: A universal commonsense reasoning model on a new multitask benchmark. *AAAI*, 2021.
- [26] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020.
- [27] Robyn Speer, Joshua Chin, and Catherine Havasi. Conceptnet 5.5: An open multilingual graph of general knowledge. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4444–4451. AAAI Press, 2017.
- [28] Julian Eisenschlos, Bhuwan Dhingra, Jannis Bulian, Benjamin Börschinger, and Jordan Boyd-Graber. Fool me twice: Entailment from Wikipedia gamification. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 352–365, Online, 2021. Association for Computational Linguistics.
- [29] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. CommonsenseQA: A question answering challenge targeting commonsense knowledge. 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 4149–4158, Minneapolis, Minnesota, 2019. Association for Computational Linguistics.
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017.
- [31] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, 2019. Association for Computational Linguistics.
- [32] Alon Talmor, Yanai Elazar, Yoav Goldberg, and Jonathan Berant. oLMPics-on what language model pre-training captures. *Transactions of the Association for Computational Linguistics*, 8:743–758, 2020.
- [33] A. Talmor and J. Berant. The web as knowledge-base for answering complex questions. In *North American Association for Computational Linguistics (NAACL)*, 2018.
- [34] Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 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, 2019. Association for Computational Linguistics.

- [35] Alon Talmor, Ori Yoran, Amnon Catav, Dan Lahav, Yizhong Wang, Akari Asai, Gabriel Ilharco, Hannaneh Hajishirzi, and Jonathan Berant. Multimodalqa: Complex question answering over text, tables and images. *arXiv preprint arXiv:2104.06039*, 2021.
- [36] Alon Talmor, Oyvind Tafjord, Peter Clark, Yoav Goldberg, and Jonathan Berant. Leap-of-thought: Teaching pre-trained models to systematically reason over implicit knowledge. *Advances in Neural Information Processing Systems*, 33, 2020.
- [37] Kevin Lin, Oyvind Tafjord, Peter Clark, and Matt Gardner. Reasoning over paragraph effects in situations. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 58–62, Hong Kong, China, 2019. Association for Computational Linguistics.
- [38] Pradeep Dasigi, Nelson F. Liu, Ana Marasović, Noah A. Smith, and Matt Gardner. Quoref: A reading comprehension dataset with questions requiring coreferential reasoning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5925–5932, Hong Kong, China, 2019. Association for Computational Linguistics.
- [39] Joshua Attenberg, Panos Ipeirotis, and Foster Provost. Beat the machine: Challenging humans to find a predictive model’s “unknown unknowns”. *Journal of Data and Information Quality (JDIQ)*, 6(1):1–17, 2015.
- [40] Allyson Ettinger, Sudha Rao, Hal Daumé III, and Emily M. Bender. 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, 2017. Association for Computational Linguistics.
- [41] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 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, 2018. Association for Computational Linguistics.
- [42] Zhilin Yang, Saizheng Zhang, Jack Urbanek, Will Feng, Alexander H. Miller, Arthur Szlam, Douwe Kiela, and Jason Weston. Mastering the dungeon: Grounded language learning by mechanical turker descent. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [43] Michael Chen, Mike D’Arcy, Alisa Liu, Jared Fernandez, and Doug Downey. 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, 2019. Association for Computational Linguistics.
- [44] Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 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, 2020. Association for Computational Linguistics.
- [45] Panagiotis G. Ipeirotis and Evgeniy Gabilovich. Quizz: targeted crowdsourcing with a billion (potential) users. In Chin-Wan Chung, Andrei Z. Broder, Kyuseok Shim, and Torsten Suel, editors, *23rd International World Wide Web Conference, WWW ’14, Seoul, Republic of Korea, April 7-11, 2014*, pages 143–154. ACM, 2014.
- [46] Luis Von Ahn and Laura Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 319–326, 2004.
- [47] Seth Cooper, Firas Khatib, Adrien Treuille, Janos Barbero, Jeehyung Lee, Michael Beenen, Andrew Leaver-Fay, David Baker, Zoran Popović, et al. Predicting protein structures with a multiplayer online game. *Nature*, 466(7307):756–760, 2010.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
  - (b) Did you describe the limitations of your work? [Yes]
  - (c) Did you discuss any potential negative societal impacts of your work? [Yes]
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
  - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [Yes]
  - (b) Did you mention the license of the assets? [Yes]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes]
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [Yes]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [Yes]