

Cascading Biases: Investigating the Effect of Heuristic Annotation Strategies on Data and Models

Chaitanya Malaviya and Sudeep Bhatia and Mark Yatskar

University of Pennsylvania

{cmalaviy, bhatiasu, myatskar}@upenn.edu

Abstract

Cognitive psychologists have documented that humans use cognitive heuristics, or mental shortcuts, to make quick decisions while expending less effort. While performing annotation work on crowdsourcing platforms, we hypothesize that such heuristic use among annotators cascades on to data quality and model robustness. In this work, we study cognitive heuristic use in the context of annotating multiple-choice reading comprehension datasets. We propose tracking *annotator heuristic traces*, where we tangibly measure low-effort annotation strategies that could indicate usage of various cognitive heuristics. We find evidence that annotators might be using multiple such heuristics, based on correlations with a battery of psychological tests. Importantly, heuristic use among annotators determines data quality along several dimensions: (1) known biased models, such as partial input models, more easily solve examples authored by annotators that rate highly on heuristic use, (2) models trained on annotators scoring highly on heuristic use don't generalize as well, and (3) heuristic-seeking annotators tend to create qualitatively less challenging examples. Our findings suggest that tracking heuristic usage among annotators can potentially help with collecting challenging datasets and diagnosing model biases.

1 Introduction

While crowdsourcing is an effective and widely-used data collection method in NLP, it comes with caveats. Crowdsourced datasets have been found to contain artifacts from the annotation process, and models trained on such data, can be brittle and fail to generalize to distribution shifts (Gururangan et al., 2018; Kaushik and Lipton, 2018; McCoy et al., 2019). In this work, we ask whether systematic patterns in annotator behavior influence the quality of collected data.

We hypothesize that usage of *cognitive heuristics*, which are mental shortcuts that humans employ in everyday life, can cascade on to data quality and model robustness. For example, an annotator asked to write a question based on a passage might not read the entire passage or might use just one sentence to frame a question. Annotators may seek shortcuts to economize on the amount of time and effort they put into a task. This behavior in annotators, characterized by examples that are acceptable but not high-quality, can be problematic.

We analyze the extent to which annotators engage in various low-effort strategies, akin to cognitive heuristics, by tracking indicative features from their annotation data in the form of **annotator heuristic traces**. First, we crowdsource reading comprehension questions where we instruct workers to write hard questions. Inspired by research on human cognition (Simon, 1956; Tversky and Kahneman, 1974), we identify several heuristics that could be employed by annotators for our task, such as satisficing (Simon, 1956), availability (Tversky and Kahneman, 1973) and representativeness (Kahneman and Tversky, 1972). We measure their potential usage by featurizing the collected data and annotation metadata (e.g., time spent and keystrokes entered) (§4). Further, we identify instantiations of these heuristics that correlate well with psychological tests measuring heuristic thinking tendencies in humans, such as the cognitive reflection test (Frederick, 2005; Toplak et al., 2014; Sirota et al., 2021). Our psychologically plausible measures of heuristic use during annotation can be aggregated per annotator, forming a holistic summary of the data they produce.

Based on these statistics, we analyze differences between examples created by annotators who engage in different levels of heuristics use. Our first finding is that examples created by strongly heuristic-seeking annotators are also easier for models to solve using heuristics (§5). We eval-

uate models that exploit a few known biases and find that examples from annotators who use cognitive heuristics are more easily solvable by biased models. We also examine what impact heuristics have on trained models. Previous work (Geva et al., 2019) shows that models generalize poorly when datasets are split randomly by annotators, likely due to the existence of artifacts. We replicate this result and find that models generalize even worse when trained on examples from heuristic-seeking annotators.

To understand which parts of the annotation pipeline contribute to heuristic-seeking behavior in annotators, we also tease apart the effect of components inherent to the task (e.g., passage difficulty) as opposed to the annotators themselves (e.g., annotator fatigue) (§6). Unfortunately, we don’t discover simple predictors (i.e., passage difficulty) of when annotators are likely to use heuristics.

A qualitative analysis of the collected data reveals that heuristic-seeking annotators are more likely to create examples that are not valid, and require simpler word-matching on explicitly stated information (§7). Crucially, this suggests that measurements of heuristic usage, such as those examined in this paper, can provide a general method to find unreliable examples in crowdsourced data, and direct our search for discovering artifacts in the data. Because we implicate heuristic use in terms of robustness and data quality, we suggest future dataset creators track similar features and evaluate model sensitivity to annotator heuristic use.¹

2 Background and Related Work

Cognitive Heuristics. The study of heuristics in human judgment, decision making, and reasoning is a popular and influential topic of research (Simon, 1956; Tversky and Kahneman, 1974). Heuristics can be defined as mental shortcuts, that we use in everyday tasks for fast decision-making. For example, Tversky and Kahneman (1974) asked participants whether more English words begin with the letter *K* or contain *K* as the 3rd letter, and more than 70% participants chose the former because words that begin with *K* are easier to recall, although that is incorrect. This is an example of the availability heuristic. Systematic use of such heuristics can lead to cognitive biases, which are irrational patterns in our thinking.

¹Our code and collected data is available at <https://github.com/chaitanyamalaviya/annotator-heuristics>.

At first glance, it may seem that heuristics are always suboptimal, but previous work has argued that heuristics can lead to accurate inferences under uncertainty, compared to optimization (Gigerenzer and Gaissmaier, 2011). We hypothesize that heuristics can play a considerable role in determining data quality and their impact depends on the exact nature of the heuristic. Previous work has shown that crowdworkers are susceptible to cognitive biases in a relevance judgement task (Eickhoff, 2018), and has provided a checklist to combat these biases (Draws et al., 2021). In contrast, our work focuses on how potential use of such heuristics can be measured in a writing task, and provides evidence that heuristic use is linked to model brittleness.

Features of annotator behavior have previously been useful in estimating annotator task accuracies (Rzeszotarski and Kittur, 2011; Goyal et al., 2018). Annotator identities have also been found to influence their annotations (Hube et al., 2019; Sap et al., 2022). Our work builds on these results and estimates heuristic use with features to capture implicit clues about data quality.

Mitigating and discovering biases. The presence of artifacts or biases in datasets is well-documented in NLP, in tasks such as natural language inference, question answering and argument comprehension (Gururangan et al., 2018; McCoy et al., 2019; Niven and Kao, 2019, *inter alia*). These artifacts allow models to solve NLP problems using unreliable shortcuts (Geirhos et al., 2020). Several researchers have proposed approaches to achieve robustness against known biases. We refer the reader to Wang et al. (2022) for a comprehensive review of these methods. Targeting biases that are unknown continues to be a challenge, and our work can help find examples which are likely to contain artifacts, by identifying heuristic-seeking annotators.

Prior work has proposed methods to discover shortcuts using explanations of model predictions (Lertvittayakumjorn and Toni, 2021), including sample-based explanations (Han et al., 2020) and input feature attributions (Bastings et al., 2021; Pezeshkpour et al., 2022). Other techniques that can be helpful in diagnosing model biases include building a checklist of test cases (Ribeiro et al., 2020; Ribeiro and Lundberg, 2022), constructing contrastive (Gardner et al., 2020) or counterfactual (Wu et al., 2021) examples and statistical tests (Gururangan et al., 2018; Gardner et al., 2021). Our

work is complementary to these approaches, as we provide an alternative approach to bias discovery that is tied to annotators.

Improved crowdsourcing. A related line of work has studied modifications to crowdsourcing protocols to improve data quality (Bowman et al., 2020; Nangia et al., 2021). In addition, model-in-the-loop crowdsourcing methods such as adversarial data collection (Nie et al., 2020) and the use of generative models (Bartolo et al., 2022; Liu et al., 2022) have been shown to be helpful in creating more challenging examples. We believe that tracking annotator heuristics use can help make informed adjustments to crowdsourcing protocols.

3 Annotation Protocol

We consider multiple-choice reading comprehension as our crowdsourcing task, because of the richness of responses and interaction we can get from annotators, which allows us to explore a range of hypothetical heuristics. We describe here the methodology for our data collection.

We provided annotators on Amazon Mechanical Turk with passages and ask them to write a multiple-choice question with four options. We used the first paragraphs of ‘vital articles’ from the English Wikipedia², and ensured that passages are at least 50 words long and at most 250 words long. Passages spanned 11 genres including arts, history, physical sciences, and others, and passages were randomly sampled from the 10K passages. Annotators were asked to write challenging questions that cannot be answered by reading just the question or passage alone, and have a single correct answer. Further, they were asked to ensure that passages provided sufficient information to answer the question while allowing questions to require basic inferences using commonsense or causality.

Annotators were first qualified to avoid spamming behavior. This qualification checked for spamming behavior in the form of invalid questions, and not example quality. Annotators were then asked to write a multiple-choice question to 4 passages in a single HIT on MTurk. Annotators were asked to not work on more than 8 HITs. We collected 1225 multiple-choice question-answer pairs from 73 annotators. In addition, we also logged their keystrokes and the time taken to complete an

²Wikipedia Level 4 vital articles: https://en.wikipedia.org/wiki/Wikipedia:Vital_articles/Level/4

example (ensuring that time away from the screen was not counted). Our annotation interface was built upon Nangia et al. (2021). For other details about our annotation protocol, please refer to Appendix A.

4 Cognitive Heuristics in Crowdsourcing

Cognitive heuristics are mental shortcuts, that humans employ in problem-solving tasks to make quick judgments (Simon, 1956; Tversky and Kahneman, 1974). Annotators, tasked with authoring natural language examples, are not infallible to using such heuristics. We hypothesize that, in writing tasks, reliance on heuristics is a traceable indicator of poor data quality. In this section, we identify several heuristics, their consequences in annotator behavior, and features to track them. Later, we also show they are predictors of qualitatively important dimensions of data.

4.1 Methodology

To test the above hypothesis, we consider several known cognitive heuristics which could be relevant for our task. This list is not comprehensive, and we refer the readers to prior work for a thorough overview of cognitive biases (Shah and Oppenheimer, 2008; Draws et al., 2021). To tangibly measure the potential usage of a heuristic, we featurize each heuristic into a measurable quantity that can be computed automatically for an example (see Table 1). While we do not conclusively determine that an annotator is using a heuristic, we explore various featurizations that align with the intuition behind each heuristic. These featurizations can sometimes be mapped to multiple heuristics that interact together, but for ease of presentation, we list them under the most related cognitive heuristic. These help us create *annotator heuristic traces*, which contain average heuristic values across all of an annotator’s examples.

To verify if our instantiation of a heuristic aligns with heuristic-seeking tendencies in annotators, we measure correlations of heuristic values with annotator performances on a battery of psychological tests (Frederick, 2005; Toplak et al., 2014; Sirota et al., 2021), described in §4.4.

4.2 Heuristics Studied

Satisficing: Satisficing is a cognitive heuristic that involves making a satisfactory choice, rather than an optimal one (Simon, 1956). In terms of

Consequence of cognitive heuristic	Featurization
Satisficing (lowtime)	(1) time, (2) log (time), (3) time / doc length, (4) log (time / doc length)
Satisficing (loweffort)	(1) question length, (2) keystroke length, (3) question+ops length, (4) question+ops length / keystroke length
Availability (first option bias)	First option is marked as correct answer
Availability (serial position)	Correct answer matches span in first or last sentence of passage
Representativeness (word overlap)	Average word overlap in all pairs of examples by annotator
Representativeness (copying)	(1) Length of longest common subsequence (lcs) b/w doc & question, (2) Max of normalized length of lcs between doc & {question, options}, (3) Normalized avg of length of lcs between doc & {question, options}

Table 1: Consequences of cognitive heuristics and featurizations for multiple-choice reading comprehension data.

mental process, strong satisficing can involve inattention to information and lack of information synthesis. In social cognition, [Krosnick \(1991\)](#) described how satisficing can manifest in various patterns in survey responses. For example, survey-takers might pick the same response to several questions in sequence, pick a random response, or use the acquiescence bias (where they always choose to *agree* with the given statement). A potential outcome of satisficing in our task is low time spent on the task and low effort put into forming a question.

Assuming the working time is t and number of tokens in a passage d is l_d , we consider the following *lowtime* featurizations: (1) t , (2) $\log t$, (3) t/l_d , (4) $\log(t/l_d)$.³

We estimate an annotator’s amount of effort through their responses. An annotator who is consistently editing their work or writing long questions might be attempting to thoughtfully draft their question. While this may not always be true (for eg, a worker might spend time thinking about their question and only start writing later), we hypothesize that often, short responses can be indicators of satisficing. Given the number of words found in a stream of keystrokes, k , the question q , and all options o_i is l_k , l_q and l_o , we consider these *loweffort* featurizations: (1) l_q , (2) l_k , (3) $l_q + l_o$, (4) $(l_q + l_o)/l_k$.

Availability heuristic: The tendency to rely upon information that is more readily retrievable from our memory is the availability heuristic ([Tversky and Kahneman, 1973](#)). For example, after hearing about a plane crash on the news, people may overstate the dangers of flying. For our task, once an annotator has read a passage and formulated a question, the question and the correct answer are likely to be readily available in their mind. This could cause them to write that information before

³Previous work shows that taking the logarithm normalizes the response time distribution ([Whelan, 2008](#)).

any of the distractor options. Therefore, we check whether the first option specified for an example is also the correct answer (*first option bias*).

Another consequence of this heuristic is the serial-position effect. When presented with a series of items like a list of words or items in a grocery list, people recall the first and last few items from the series better than the middle ones ([Murdock Jr, 1962](#); [Ebbinghaus, 1964](#)) because of their easier availability. This effect can also be explained as a combination of the primacy effect and recency effect. To test if an annotator anchors their questions on the first or last sentence of the passage due to this heuristic, we check if the correct answer marked for an example matches a span in the first or last sentence of the passage (*serial position*).

Representativeness heuristic: The representativeness heuristic is our tendency to use the similarity of items to make decisions ([Kahneman and Tversky, 1972](#)). For example, if a person is picking a movie to watch, they might think of movies they previously liked and look for those attributes in a new movie. Similarly, an annotator may repeat the same construction in their questions to ease decision-making (e.g., "which of the following is true?" or "what year did [event] happen?"). This could either mean that they are not fully engaged, or, they found a writing strategy that works well and they choose to stick to it. We measure this tendency by computing the average *word overlap* across all pairs of questions from an annotator.

A different manner in which this heuristic can manifest is using similarity with the provided context, i.e., through *copying*. Copying, or imitation, is a common building block that guides human behavior and decision making. In deciding what clothes to buy or which book to read, humans use imitation-of-the-majority to make quicker inferences with lesser cognitive effort ([Garcia-Retamero et al., 2009](#); [Gigerenzer and Gaissmaier,](#)

2011). Similarly, annotators can have tendencies to copy text word-for-word from the context they are primed with, to reduce their cognitive load. Assuming LCS is a function that computes the length of the longest common subsequence between two sequences, we consider these featurizations for copying: (1) $LCS(d, q)$, (2) $\max(LCS(d, q), LCS(d, o))$ and (3) $\text{avg}(LCS(d, q), LCS(d, o))$.

4.3 Annotator Heuristic Traces

The consequences of heuristics we compute, as summarized in Table 1, may not in themselves be problematic per example. However, we claim that annotators who consistently rely on such heuristics may impart larger, harder-to-detect, undesirable regularities in data.

Annotator heuristic traces capture global behavioral trends per annotator. For each annotator and heuristic, we average the heuristic values across all of the annotator’s examples, forming a matrix of annotators and their average heuristic values.

Principal components of heuristics: We also evaluate if a low-dimensional representation of an annotator’s heuristic trace is useful for predicting data quality. We compute the first principal component of this matrix to simultaneously consider multiple heuristic indicators.

4.4 Cognitive Reflection Test

Although we cannot determine whether an annotator is definitively using a heuristic, we can probe if our features correlate with heuristic-seeking tendencies in annotators. Previous work in cognitive psychology has designed tests measuring such tendencies. These help us validate the psychological plausibility of our features, ensuring they are generally applicable.

Perhaps the best known test of heuristic-seeking tendencies is the Cognitive Reflection Test (CRT) (Frederick, 2005). The test has 3 questions, but we instead use the 7-item CRT from Toplak et al. (2014) to find more variance among annotators. The numerical CRT requires mathematical reasoning and previous work has highlighted that its results might be conflated with mathematical reasoning capabilities. Further, since our task requires writing, we also perform the verbal CRT (Sirota et al., 2021). This test has 9 items⁴, and is known to correlate well with the numerical CRT, and other

⁴We exclude a question that requires cultural knowledge.




Figure 1: Correlations of annotator scores on the CRT and their average features values for each heuristic. Feature names, left, correspond to feature names from Table 1. The CRT3 includes the original questions from Frederick (2005) and CRT7 includes 4 more questions from Toplak et al. (2014). The black boxes indicate the featurization with the highest average correlations for the heuristic. * indicates $p < 0.01$ and ^ indicates $p < 0.1$.

indicators of cognitive capabilities. The questions in these tests are provided in Appendix B.⁵

We asked annotators who completed at least 5 question writing examples to do two surveys asking logical questions (CRT-7 and Verbal CRT). 49 of 59 annotators completed the surveys. We then compute Pearson correlations between annotator accuracies on the three versions of the CRT, and values in their heuristic traces, shown in Figure 1. The results indicate that our featurizations have significant, medium correlations with the CRTs, and the PCA projection, which captures multiple heuristics, has the highest correlations. For the sake of further analysis, for each feature group, we use the feature that has the highest average correlations with the CRT tests (enclosed in black boxes in Figure 1).

⁵The use of the CRT has issues due to repeated exposure (Stieger and Reips, 2016; Haigh, 2016), so we ensured that the names/quantities are different from the ones used in the original questions. We also emphasize that the CRT does not provide interpretability into the mechanism of the heuristics, whereas our individual heuristic features do.




Figure 2: Precision of labeling heuristic examples H_k as solvable by biased models, when the set H_k is formed by examples from the k^{th} percentile of heuristic-seeking annotators.




Figure 3: Precision of labeling heuristic examples H_k as solvable by biased models on data from Sugawara et al. (2022).

5 Biased Model Solvability

Annotator heuristic traces are cognitively plausible measures that we hypothesize are indicators of large, potentially undesirable patterns annotators impart on data. To verify this, we test if examples created by heuristic-seeking annotators are more easily solvable by biased models.

We consider heuristic examples as examples from those annotators who score highly on our heuristic indicator features. Given the initial set of examples D , we distinguish a subset as heuristic, H_k , formed by all examples from annotators in the top $k\%$ of average heuristic use across all annotators. We form such a subset independently for all heuristic indicator features we consider.⁶ When H_k is formed from the top quartile ($k=25$), 68% of annotators have examples included in at least 1 heuristic set, and 14% in for at least 4/6 heuristic sets. We find that few annotators never use heuristics.

Next, we evaluate how well biased models perform on heuristic subsets (H_k) compared to the remaining examples, $D \setminus H_k$. We evaluate a few biased models, trained to use unreliable heuristics, on examples created with or without heuristics. Below we describe the biased models we use. In all cases,

⁶We exclude those annotators who wrote less than 5 examples and exclude all invalid examples.

we train or finetune models on QA data from Nangia et al. (2021) and evaluate them on our data. For hyperparameter settings, please see Appendix C.

Lexical Overlap Model (overlap). We train a logistic regression classifier by building upon features from the bias-only model from Clark et al. (2019). Assuming the concatenated passage and question are the context for each option, we use the following features: 1) is the option a subsequence of context, 2) do all words in the option exist in context, 3) the fraction of words in the option that exist in context, 4) the log of length difference between the context and the option, 5) the average and maximum of minimum distance between each context word with each option word using 300-dimensional fastText embeddings (Joulin et al., 2017). We then pick the option with the highest probability as the model prediction. The model achieves an accuracy of 42.27% on D .

Partial Input Models. As a benchmark for diagnosing the collected data, we consider several partial input models. These include no passage (**no_passage**), no question (**no_ques**), first & last sentence of passage only (**fl_passage**). We use a RoBERTa-Large (Liu et al., 2019) model initially finetuned on RACE (Lai et al., 2017) and further trained on the baseline data from Nangia

et al. (2021). These models achieve accuracies of 42.44%, 59.49%, and 55.98% on D , respectively, demonstrating better than random performance.

Human heuristic solvability (human_biased).

In addition to biased models, we also consider an implicit notion of example difficulty from a biased human. Specifically, we evaluate whether a human can answer an example just by skimming the passage. We use an interface where a passage is only visible for 30 seconds, after which, a human needs to answer the question. One of the authors conducted this annotation for the collected examples and achieved an accuracy of 79.79%.

Results. Figure 2 shows the precision of H_k being solvable by biased models, as k is varied. As we can see from the plots, there is a downward trend as the percentile is increased for all heuristics. The features for the availability and representative heuristic, and the PCA projection are particularly effective. This suggests that strongly heuristic-seeking annotators are more likely to create examples solvable by biased models.

Other non-Wikipedia domains. To test if the heuristics we considered are indicative of solvability by biased models in domains other than Wikipedia, we repeated our analysis on 1,982 examples from the standard data collection setting in Sugawara et al. (2022), who collected questions for passages from many different sources. The precision plot is shown in Figure 3. With the exception of serial-position, heuristic-seeking features identify annotators that create examples more easily solvable by biased models in these domains too.⁷

As a predictor of bias. In addition to evaluating the predictiveness of annotator heuristic features at the extreme, we also evaluated whether heuristic features are predictive of solvability by biased models across annotators. Specifically, we calculated Pearson correlations between annotators’ average heuristic values, and the accuracies of biased models on their examples, in Figure 4. These correlations are not strong for the satisficing heuristics, but we do notice some significant, medium correlations for the other heuristics we studied. Importantly, we contrast this with the same correlations measured over the entire pool of data (without averaging per annotator). Those correlations, shown in Figure 6

⁷The serial-position feature might be more effective for Wikipedia because they tend to be more factual, which makes it easier to form questions using the first or last sentence.




Figure 4: Pearson correlations of annotators’ average heuristic values and accuracies with biased models on their annotated examples. * and ^ indicate $p < 0.01$ and $p < 0.1$.

in the Appendix, are much weaker showing the value of our annotator-level measures.

Model generalization across annotators. Previous work showed that models do not generalize well to annotator-based random splits of crowd-sourced datasets, suggesting models might be learning annotator-specific biases (Geva et al., 2019). We suspect that generalization might deteriorate when models are trained on heuristic-seeking annotators, as models could more easily specialize to their examples. Hence, we ask whether heuristic-based splits (**heuristic**) lead to worse performance than random annotator splits (**random**).

While controlling the number of training examples, we trained models on examples from heuristic-seeking annotators or a random set of annotators, and test on the remaining examples. For heuristic-based splits, we train on examples in H_{33} , the top 33% of heuristic-seeking annotators for a heuristic indicator. For random annotator splits, we resampled splits with 3 random seeds and report means. In addition, we trained models on random splits of the same training size (**random-pooled**), where data is not split by annotator. The accuracies on these splits are shown in Table 2. We find generalization is poorer for almost all of the heuristic-based annotator splits compared to random annotator splits. This suggests that heuristic-based splits can serve as natural challenge sets and inadvertently sampling heuristic-seeking annotators for training may not generalize well.




Figure 5: Percentage difference of examples in heuristic set, H_{25} , and the remaining examples, $D \setminus H_{25}$, labeled as having a qualitative property. Examples in the heuristic set are less valid & require more word matching based on explicitly stated information.

Heuristic	heuristic	random	random-pooled	n
lowtime	84.27	87.42 \pm 1.17	87.14 \pm 0.64	394
loweffort	85.06	88.56 \pm 1.01	86.27 \pm 1.75	415
first_option	87.33	87.92 \pm 0.85	88.12 \pm 1.63	341
serial_position	84.78	88.07 \pm 0.84	88.43 \pm 0.6	389
word_overlap	86.86	86.64 \pm 2.16	87.50 \pm 0.67	362
copying	83.75	86.94 \pm 1.78	87.47 \pm 1.02	332
pca	83.40	86.94 \pm 0.77	87.52 \pm 0.81	385

Table 2: Performance on heuristic-based, random annotator splits and random splits with the same training set size. We performed 3 runs on the randomly sampled splits, and report means and standard deviations.

6 Influencers of heuristic behavior

Next, we aim to understand what role the annotation pipeline plays in influencing heuristic use among annotators. Various factors have been shown to determine example quality in crowdsourcing. These include task difficulty, incentives, annotator ability, motivation and fatigue (Krosnick, 1991; Yan et al., 2010). We looked at how such markers influence heuristic use among annotators.

We considered two types of measures that could indicate difficulty: *passage length* (number of tokens) and inverse *entity* count (doc length / number of named entities) in the passage. Longer documents, with fewer named entities, might provide context that is harder to form questions about. Further, having completed more examples could make an annotator fatigued and/or gain expertise at the task. Hence, we also used the sequence *index* of each example for an annotator. We computed Pearson correlations between these indicators and the heuristic values for each annotator, and averaged the correlations across annotators. Our results are summarized in Table 3. We find that neither of

Heuristic	passage length	entity	index
lowtime	0.12	0.16	-0.09
loweffort	0.09	-0.03	-0.01
first_option	0.08	0.00	0.02
serial_position	-0.09	-0.05	-0.03
copying	0.13	0.06	-0.05

Table 3: Correlations between heuristic values and factors, averaged across annotators.

these factors show significant correlations with heuristic features among annotators.

7 Qualitative Analysis

To better understand the differences in data produced by heuristic-seeking annotators, and otherwise, we conducted a comprehensive qualitative analysis of all our data. We annotate questions with properties inspired from previous work (Lai et al., 2017; Trischler et al., 2017; Sugawara et al., 2018, 2022) along the following dimensions: **validity** (is the question answerable given the context in the passage), **context** (how much context from the passage is needed to answer the question), and **comprehension type** (what kinds of comprehension are needed to answer the question). Each question can have multiple labels. For a detailed description of these labels, please refer to Appendix E.

Results. Figure 5 presents the results of our annotation. We show the differences in the percentage of examples in the heuristic set, H_{25} , and the remaining examples, $D \setminus H_{25}$. First, examples in the non-heuristic set are more likely to be valid, and less likely to be unsolvable compared to the heuristic set. Further, we find that examples in the heuristic set often require simple word matching and para-

phrasing, while the ones in the non-heuristic set, are more likely to require multi-sentence reasoning. In terms of comprehension type, we find that heuristic examples are more likely to be answerable using information explicitly stated in the passage. On the other hand, non-heuristic examples are more likely to require implicit inference. These results suggest there are significant qualitative differences in examples from heuristic-seeking annotators.

8 Discussion

Our work measures the implications of annotators' potential use of cognitive heuristics in data quality. The analyses we present suggest that models are indirectly influenced by heuristic use and that previous observations, such as the success of partial input models, is a consequence. While many such consequences of heuristic use appear to be negative, we believe that this judgement should be left up to practical applications that use the data. We propose a fruitful direction for characterizing what models learn from data by considering annotator behaviors.

Practically, it is an open question as to how we can control downstream data using annotator heuristic traces. Instead, we propose that future annotation efforts minimally track indicators of heuristic usage, using task-specific features, in an effort to document how they are reflected in the collected data and trained models.

9 Limitations

One limitation of our study is that we analyze the implications of heuristic-seeking behavior in annotators for one task. Future work could consider extending this methodology to many annotation tasks. For example, in sentence-pair classification tasks such as textual entailment, or in annotation of machine translation or summarization datasets, annotator heuristics could be useful in determining the quality of data and the biases embedded in them. To find stronger signals in the annotator heuristic traces, future work could consider training models to featurize heuristics.

Acknowledgements

This work is supported in part by an Allen Institute for AI Young Investigator award and National Science Foundation grant SES-1847794. We thank the anonymous reviewers, and the annotators who

participated in our study. We also thank the University of Pennsylvania NLP group for helpful discussions.

References

- Max Bartolo, Tristan Thrush, Sebastian Riedel, Pontus Stenetorp, Robin Jia, and Douwe Kiela. 2022. [Models in the loop: Aiding crowdworkers with generative annotation assistants](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3754–3767, Seattle, United States. Association for Computational Linguistics.
- Jasmijn Bastings, Sebastian Ebert, Polina Zablotskaia, Anders Sandholm, and Katja Filippova. 2021. ["will you find these shortcuts?" a protocol for evaluating the faithfulness of input salience methods for text classification](#). *arXiv preprint arXiv:2111.07367*.
- Samuel R. Bowman, Jennimaria Palomaki, Livio Baldini Soares, and Emily Pitler. 2020. [New protocols and negative results for textual entailment data collection](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8203–8214, Online. Association for Computational Linguistics.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. [Don't take the easy way out: Ensemble based methods for avoiding known dataset biases](#). 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 4069–4082, Hong Kong, China. Association for Computational Linguistics.
- Tim Draws, Alisa Rieger, Oana Inel, Ujwal Gadiraju, and Nava Tintarev. 2021. [A checklist to combat cognitive biases in crowdsourcing](#). In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 9, pages 48–59.
- Hermann Ebbinghaus. 1964. [Memory: A contribution to experimental psychology](#). *New York, Dover Publications*.
- Carsten Eickhoff. 2018. [Cognitive biases in crowdsourcing](#). In *Proceedings of the eleventh ACM international conference on web search and data mining*, pages 162–170.
- Shane Frederick. 2005. [Cognitive reflection and decision making](#). *Journal of Economic perspectives*, 19(4):25–42.
- Rocio Garcia-Retamero, Masanori Takezawa, and Gerd Gigerenzer. 2009. [Does imitation benefit cue order learning?](#) *Experimental Psychology*, 56(5):307–320.

- 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.
- Matt Gardner, William Merrill, Jesse Dodge, Matthew Peters, Alexis Ross, Sameer Singh, and Noah A. Smith. 2021. [Competency problems: On finding and removing artifacts in language data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1801–1813, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. 2020. [Shortcut learning in deep neural networks](#). *Nature Machine Intelligence*, 2(11):665–673.
- 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.
- Gerd Gigerenzer and Wolfgang Gaissmaier. 2011. [Heuristic decision making](#). *Annual review of psychology*, 62:451–482.
- Tanya Goyal, Tyler McDonnell, Mucahid Kutlu, Tamer Elsayed, and Matthew Lease. 2018. [Your behavior signals your reliability: Modeling crowd behavioral traces to ensure quality relevance annotations](#). In *Sixth AAAI Conference on Human Computation and Crowdsourcing*.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Matthew Haigh. 2016. [Has the standard cognitive reflection test become a victim of its own success?](#) *Advances in cognitive psychology*, 12(3):145.
- Xiaochuang Han, Byron C. Wallace, and Yulia Tsvetkov. 2020. [Explaining black box predictions and unveiling data artifacts through influence functions](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5553–5563, Online. Association for Computational Linguistics.
- Christoph Hube, Besnik Fetahu, and Ujwal Gadiraju. 2019. [Understanding and mitigating worker biases in the crowdsourced collection of subjective judgments](#). In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–12.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. [Bag of tricks for efficient text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Daniel Kahneman and Amos Tversky. 1972. [Subjective probability: A judgment of representativeness](#). *Cognitive psychology*, 3(3):430–454.
- Divyansh Kaushik and Zachary C. Lipton. 2018. [How much reading does reading comprehension require? a critical investigation of popular benchmarks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5010–5015, Brussels, Belgium. Association for Computational Linguistics.
- Jon A Krosnick. 1991. [Response strategies for coping with the cognitive demands of attitude measures in surveys](#). *Applied cognitive psychology*, 5(3):213–236.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale ReAding comprehension dataset from examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Piyawat Lertvittayakumjorn and Francesca Toni. 2021. [Explanation-based human debugging of nlp models: A survey](#). *Transactions of the Association for Computational Linguistics*, 9:1508–1528.
- Alisa Liu, Swabha Swayamdipta, Noah A Smith, and Yejin Choi. 2022. [Wanli: Worker and ai collaboration for natural language inference dataset creation](#). *arXiv preprint arXiv:2201.05955*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of*

- the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Bennet B Murdock Jr. 1962. [The serial position effect of free recall](#). *Journal of experimental psychology*, 64(5):482.
- Nikita Nangia, Saku Sugawara, Harsh Trivedi, Alex Warstadt, Clara Vania, and Samuel R. Bowman. 2021. [What ingredients make for an effective crowdsourcing protocol for difficult NLU data collection tasks?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1221–1235, Online. Association for Computational Linguistics.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Timothy Niven and Hung-Yu Kao. 2019. [Probing neural network comprehension of natural language arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. [Scikit-learn: Machine learning in Python](#). *Journal of Machine Learning Research*, 12:2825–2830.
- Pouya Pezeshkpour, Sarthak Jain, Sameer Singh, and Byron Wallace. 2022. [Combining feature and instance attribution to detect artifacts](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1934–1946, Dublin, Ireland. Association for Computational Linguistics.
- Marco Tulio Ribeiro and Scott Lundberg. 2022. [Adaptive testing and debugging of NLP models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3253–3267, Dublin, Ireland. 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.
- Jeffrey M Rzeszutarski and Aniket Kittur. 2011. [Instrumenting the crowd: using implicit behavioral measures to predict task performance](#). In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 13–22.
- Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi, and Noah A. Smith. 2022. [Annotators with attitudes: How annotator beliefs and identities bias toxic language detection](#). In *Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL)*, Seattle, USA.
- Anuj K Shah and Daniel M Oppenheimer. 2008. [Heuristics made easy: an effort-reduction framework](#). *Psychological bulletin*, 134(2):207.
- Herbert A Simon. 1956. [Rational choice and the structure of the environment](#). *Psychological review*, 63(2):129.
- Miroslav Sirota, Chris Dewberry, Marie Juanchich, Lenka Valuš, and Amanda C Marshall. 2021. [Measuring cognitive reflection without maths: Development and validation of the verbal cognitive reflection test](#). *Journal of Behavioral Decision Making*, 34(3):322–343.
- Stefan Stieger and Ulf-Dietrich Reips. 2016. [A limitation of the cognitive reflection test: familiarity](#). *PeerJ*, 4:e2395.
- Saku Sugawara, Kentaro Inui, Satoshi Sekine, and Akiko Aizawa. 2018. [What makes reading comprehension questions easier?](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4208–4219, Brussels, Belgium. Association for Computational Linguistics.
- Saku Sugawara, Nikita Nangia, Alex Warstadt, and Samuel Bowman. 2022. [What makes reading comprehension questions difficult?](#) In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6951–6971, Dublin, Ireland. Association for Computational Linguistics.
- Maggie E Toplak, Richard F West, and Keith E Stanovich. 2014. [Assessing miserly information processing: An expansion of the cognitive reflection test](#). *Thinking & Reasoning*, 20(2):147–168.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordoni, Philip Bachman, and Kaheer Suleman. 2017. [NewsQA: A machine comprehension dataset](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 191–200, Vancouver, Canada. Association for Computational Linguistics.
- Amos Tversky and Daniel Kahneman. 1973. [Availability: A heuristic for judging frequency and probability](#). *Cognitive psychology*, 5(2):207–232.

- Amos Tversky and Daniel Kahneman. 1974. [Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty.](#) *science*, 185(4157):1124–1131.
- Xuezhi Wang, Haohan Wang, and Diyi Yang. 2022. [Measure and improve robustness in nlp models: A survey.](#) *Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Robert Whelan. 2008. [Effective analysis of reaction time data.](#) *The Psychological Record*, 58(3):475–482.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel Weld. 2021. [Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models.](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6707–6723, Online. Association for Computational Linguistics.
- Yan Yan, Rómer Rosales, Glenn Fung, Mark Schmidt, Gerardo Hermosillo, Luca Bogoni, Linda Moy, and Jennifer Dy. 2010. [Modeling annotator expertise: Learning when everybody knows a bit of something.](#) In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 932–939. JMLR Workshop and Conference Proceedings.

A Crowdsourcing setup

For annotators to participate in our task, they needed to have an acceptance rate greater than or equal to 98% and have at least 1000 approved HITs. In addition, we required annotators to be located in US, UK or Canada. We estimated each HIT to take approximately 15 minutes, and paid \$4 per HIT (\$15 / hr). Figure 7 shows the interface presented to the annotators for data collection.

B Cognitive Reflection Tests

We list the questions used in the numerical CRT and the verbal CRT in Table 4 and Table 5 respectively. The first 3 questions in Table 4 correspond to the original CRT from Frederick (2005).

C Hyperparameter Settings

Lexical overlap model. The logistic regression was trained with $C=100$ and a maximum of 100 iterations for convergence with the scikit-learn library (Pedregosa et al., 2011).

Partial input models. The partial input models were trained with a learning rate of $1e-4$ and batch size of 1, for 4 epochs and all the default hyperparameters in the multiple-choice QA example in the Transformers library (Wolf et al., 2020). These experiments took approximately a week of compute time on a single Quadro RTX 6000 GPU.

D Correlations with pooled data

In Figure 6, we show correlations between heuristic features and biased model accuracies when all examples are pooled together. We contrast this with the annotator-wise plots shown in Figure 4.

E Question Annotation Scheme

We describe the annotation scheme used to label examples for the analysis in section 7. In addition, we show a breakdown of those results across all heuristic features in Figure 8.

Validity. We annotate whether examples are answerable or not using the following labels:

1. **Unsolvable:** It is not possible to answer the question given the context in the passage and question, or the question is underspecified or incoherent.
2. **Incorrect:** The answer is marked incorrectly.




Figure 6: Pearson correlations of average heuristic values and biased model solvability, pooled across all examples. We exclude *word overlap* since it is computed across all examples of an annotator and is not a sample-level measure. * indicates $p < 0.01$ and ^ indicates $p < 0.1$.

3. **Ambiguous:** The question does not have a unique correct answer.
4. **Valid:** The question can be reasonably answered from the passage.

Context. To understand how much context from the passage is needed to answer the question, we label questions using the following labels:

1. **Word matching:** The question matches a span in the passage, and the answer is easily extractable by matching spans.
2. **Paraphrasing:** The question paraphrases information in exactly one sentence in the passage, and the answer can be retrieved from it.
3. **Single-sentence reasoning:** The question can be answered by exactly one sentence in the passage, but requires a conceptual overlap, or performing some other form of inference.
4. **Multi-sentence reasoning:** The question can only be answered by synthesizing information from multiple sentences in the passage. This excludes just performing coreference.
5. **Coreferential Reasoning:** The question requires performing coreference.

▼ Passage 1 / 4

Question Writing → Completed

The Asante Empire (Asante Twi: Asanteman) was an Akan empire and kingdom from 1701 to 1901, in what is now modern-day Ghana. It expanded from the Ashanti Region to include the Brong-Ahafo Region, Central Region, Eastern Region and Western Region of present-day Ghana as well as some parts of Ivory Coast and Togo. Due to the empire's military prowess, wealth, architecture, sophisticated hierarchy and culture, the Ashanti Empire has been extensively studied and has more historic records written by European, primarily British authors than any other indigenous culture of Sub-Saharan Africa. Starting in the late 17th century, the Ashanti king Osei Tutu (c. 1695 – 1717) and his adviser Okomfo Anokye established the Ashanti Kingdom, with the Golden Stool of Asante as a sole unifying symbol. Osei Tutu oversaw a massive Ashanti territorial expansion, building up the army by introducing new organisation and turning a disciplined royal and paramilitary army into an effective fighting machine. In 1701, the Ashanti army conquered Denkyira, giving the Ashanti access to the Gulf of Guinea and the Atlantic Ocean coastal trade with Europeans, notably the Dutch. The economy of the Ashanti Empire was mainly based on the trade of gold and slaves. The army served as the effective tool to procure captives. The Ashanti Empire fought several wars with neighboring kingdoms and lesser organized groups such as the Fante.

Question

input goes here

Options

- 1. input goes here
- 2. input goes here
- 3. input goes here
- 4. input goes here

Write a natural & difficult question!

Submit the question

Figure 7: Annotation interface used for data collection.

Question	Intuitive Answer	Correct Answer
1) A carpet and a lamp cost \$450 in total. The carpet costs \$400 more than the lamp. How much does the lamp cost?	\$50	\$25
2) It takes 10 computers 10 minutes to run 10 programs. How many minutes does it take 500 computers to run 500 programs?	500	10
3) There is a patch of lily pads in a pond. The patch doubles in size every day. If it takes 100 days for the patch to cover the entire pond, how many days would it take to cover half the pond?	50	99
4) If Jason can drink one barrel of water in 6 days, and Jen can drink one barrel of water in 12 days, how long would it take them to drink one barrel of water together?	9	4
5) Aidan received both the 25th highest and the 25th lowest mark in the class. How many students are in the class?	50	49
6) A farmer buys a sheep for \$500, sells it for \$600, buys it back for \$700, and sells it finally for \$800. How much has he made?	\$100	\$200
7) Ramona decided to invest \$5,000 in the stock market early in 2008. Six months after she invested, on July 17, the stocks she had purchased were down 50%. Fortunately for Ramona, from July 17 to October 17, the stocks she had purchased went up 75%. At this point, does Ramona have a) the same amount of money as when she invested, b) more money, c) less money ? (Respond with a, b or c)	b	c

Table 4: Questions in the numerical CRT.

Question	Intuitive Answer	Correct Answer
1) Angie's father has 5 daughters but no sons—Nana, Nene, Nini, Nono. What is the fifth daughter's name probably?	Nunu	Angie
2) If you were running a race, and you passed the person in 5th place, what place would you be in now?	4th	5th
3) It is a stormy night and a plane takes off from JFK airport in New York. The storm worsens, and the plane crashes - half lands in the United States, the other half lands in Canada. In which country do you bury the survivors?	USA	we do not bury survivors
4) A monkey, a squirrel, and a bird are racing to the top of a coconut tree. Who will get the banana first, the monkey, the squirrel, or the bird?	bird	there is no banana on a coconut tree
5) In a one-storey pink house, there was a pink person, a pink cat, a pink fish, a pink computer, a pink chair, a pink table, a pink telephone, a pink shower—everything was pink! What colour were the stairs probably?	pink	no stairs in a one-storey house
6) The wind blows west. An electric train runs east. In which cardinal direction does the smoke from the locomotive blow?	west	no smoke from an electric train
7) If you have only one match and you walk into a dark room where there is an oil lamp, a newspaper and wood— which thing would you light first?	oil lamp	match
8) Would it be ethical for a man to marry the sister of his widow?	no	not possible
9) Which sentence is correct: (a) 'the yolk of the egg are white' or (b) 'the yolk of the egg is white'?	b	the yolk is yellow

Table 5: Questions in the verbal CRT.

Comprehension Type. To determine what kinds of comprehension are required to answer the question, we label examples with the following labels:

1. **Math/numerical:** Questions that require mathematical or numerical reasoning.
2. **Whole:** Questions that require a complete understanding of the passage or ask about the author's opinion on the passage.
3. **Factuality:** Questions asking about truthfulness of the statements presented in the question or the options (e.g., questions of the form "which of the following is true / false?", or "is it true/false that ..")
4. **Spatial/temporal:** Requires understanding of location and temporal order of events.
5. **Explicit:** Asks about information (facts, events or entities) stated in the passage explicitly or in a paraphrased manner. These shouldn't require much of a concept jump.
6. **Implicit:** Asks about information not directly stated, but which can be inferred through commonsense, causality, numerical or other types of inference.
7. **Negation:** Questions which are phrased in the form of a negation (for e.g. using keywords like "not" and "without").



Figure 8: Breakdown of the question annotation results from section 7 when percentile $k=25$.