

Machine Reading, Fast and Slow: When Do Models “Understand” Language?

Sagnik Ray Choudhury^{1*,2}, Anna Rogers², and Isabelle Augenstein²

¹University of Michigan

²University of Copenhagen

sagnikrayc@gmail.com, arogers@sodas.ku.dk, augenstein@di.ku.dk

Abstract

Two of the most fundamental challenges in Natural Language Understanding (NLU) at present are: (a) how to establish whether deep learning-based models score highly on NLU benchmarks for the ‘right’ reasons; and (b) to understand what those reasons would even be. We investigate the behavior of reading comprehension models with respect to two linguistic ‘skills’: coreference resolution and comparison. We propose a definition for the reasoning steps expected from a system that would be ‘reading slowly’, and compare that with the behavior of five models of the BERT family of various sizes, observed through saliency scores and counterfactual explanations. We find that for comparison (but not coreference) the systems based on larger encoders are more likely to rely on the ‘right’ information, but even they struggle with generalization, suggesting that they still learn specific lexical patterns rather than the general principles of comparison.

1 Introduction

Generally, human decisions may be based on deliberate, careful reasoning (‘slow thinking’) or quick heuristics (‘fast thinking’) (Kahneman, 2011). These two processes have parallels in the realm of reading comprehension (RC): a human reader would ideally fully process the text to answer questions, but in practice, we may deliberately skim rather than read to save effort. Even capable students may be misled by superficial cues (Ackerman et al., 2013).

The previous generations of NLP models have already achieved high performance on many RC benchmarks, but they were found to often ‘read fast’, i.e. rely on shallow patterns (Chen et al., 2016; Jia and Liang, 2017; Rychalska et al., 2018). Fine-tuned Transformer-based models (Devlin et al., 2019) still have similar shortcomings

* Work done while employed at the University of Copenhagen

Context: Leo Strauss was a political philosopher and classicist. He was born in Germany ... Thoughts on Machiavelli is a book by Leo Strauss ...
Question: Where was the author of Thoughts of Machiavelli born ?
Answer: Germany

Figure 1: A sample question from the SQuAD (Rajpurkar et al., 2016) dataset. green tokens are the words that a reader relying on coreference resolution would take into account, and red tokens are the words that could be used to answer the question with entity type matching.

(Sugawara et al., 2020; Rogers et al., 2020; Sen and Saffari, 2020; Kassner and Schütze, 2020, inter alia) in RC, as well as other tasks (McCoy et al., 2019; Jin et al., 2020).

Consider the example in Figure 1. A human reader would ideally construct the coreference chain resolving the pronoun ‘he’ to ‘Leo Strauss’. A possible heuristic-based solution is entity type matching (Jia and Liang, 2017): a model could observe that a ‘where’ question can only be answered by a ‘location’ and among two such entities (‘Germany’ and ‘United States’) the correct answer (‘Germany’) is closer to ‘born’. Such heuristic reasoners will not generalize to unseen examples. Thus a key challenge in building trustworthy and explainable RC systems is to make sure their decisions are based on valid reasoning steps. However, it is difficult to establish: (a) what that reasoning should be; and (b) whether a blackbox system adheres to it.

The present study proposes a framework for the analysis of RC models that includes: (a) defining the expected reasoning; (b) analysing model performance using explainability techniques. In particular, we contribute a case study for RC questions involving coreference resolution and comparison: we define the expected ‘reasoning’ for them (§2) and use a combination of saliency-based and coun-

terfactual explanations (§3) to analyze RC systems based on BERT and RoBERTa encoders of various sizes (§4). Overall, we find that the larger models are more likely to rely on the ‘right’ information, but even they seem to learn specific lexical patterns rather than underlying linguistic phenomena.

2 When do RC Model ‘Understand’ A Text?

2.1 Understanding in Humans

The phenomenon of ‘natural language understanding’ is not yet sufficiently well defined even for human speakers, although it is pursued by at least three different fields: philosophy of mind (e.g. Grimm, 2021; Dellsén, 2020), psychology (e.g. Christianson, 2016; Zwaan, 2016), and pedagogy (e.g. Lander, 2010; Duffin and Simpson, 2000). We cannot do this topic justice within the scope of this paper, but let us briefly outline the key premises about human understanding that we rely on in our work:

- Understanding is not truth-connected: it is “a merely psychological state” (Grimm, 2012);
- Its objects are something like ‘connections’ or ‘relations’ of the phenomenon X to other phenomena (Grimm, 2021);
- It is not binary: teachers routinely talk of ‘levels of understanding’, ‘continuum of understanding’ or ‘partial understanding’ (Nurhuda et al., 2017);
- It is different from ‘knowledge’, i.a. since it is “not transmissible¹ in the same sense as knowledge is” (Burnyeat and Barnes, 1980).

If human understanding is about establishing connections between new and existing conceptualizations, its success depends on the pre-existence of a suitable set of conceptualizations, to which the connections can be established (this is why e.g. algebra is taught in schools before differential calculus). The set of conceptualizations that each of us possesses is unique, since it depends on our experience of the world (cf. Fillmore’s ‘semantics of understanding’ (Fillmore, 1985)). This, together with other factors like level of motivation, attention etc., explains the variation in human understanding: we may grasp different sets of possible connections

¹This is why, as any teacher knows from practice, simply presenting the students with definitions or principles does not necessarily result in understanding of those principles or definitions.

between different aspects of the new phenomenon and our pre-existing worldview.

2.2 ‘Understanding’ in Machines

Much research on human understanding focuses on mechanisms that fundamentally do not apply to current NLP systems, such as the distinction between ‘knowledge’ and ‘understanding’ or the fact that humans will fail to understand if they don’t have suitable pre-existing conceptualizations (while an encoder will encode text even if its weights are random). Since the mechanism (and its results) is so fundamentally different, terms like ‘natural language understanding’ or ‘reading comprehension’² for the current NLP systems are arguably misleading. It would be more accurate to talk instead of ‘natural language processing’ and ‘information retrieval’.

While terms like ‘understanding’ are widely (mis)applied to models in AI research (Mitchell, 2021), their definitions are scarce. Turing famously posited that the question “can machines think?” is too ill-defined to deserve serious consideration, and replaced it with a behavioral test (conversation with a human judge) for when we would say that thinking occurs (Turing, 1950). Conceptually, this is still the idea underlying the ‘NLU’ benchmarks used today: we assume that for models to perform well on collections of tests such as GLUE (Wang et al., 2018, 2019), some capacity for language understanding is required, and hence if our systems get increasingly higher scores on such behavioral tests, this would mean progress on ‘NLU’. However, just like the Turing test itself turned out to be “highly gameable” (Marcus et al., 2016), so are our tests³ (Sugawara et al., 2020; Rogers et al., 2020; Sen and Saffari, 2020; Kassner and Schütze, 2020; McCoy et al., 2019; Jin et al., 2020, inter alia).

All this suggests that, at the very least, we need a better specification for the success criteria for such behavioral tests. Instead of asking “Does my RC model “understand” language?” we could ask: “Does my RC model produce its output based

²Marcus and Davis (2019) dispute even the applicability of the term “reading”, declaring the current QA/RC systems “functionally illiterate” since they cannot draw the implicit inferences crucial for human reading.

³In fact, the larger the dataset, the more of likely spurious patterns are to occur (Gardner et al., 2021). This presents a fundamental problem for data-hungry deep learning systems: “the models, unable to discern the intentions of the data set’s designers, happily recapitulate any statistical patterns they find in the training data” (Linzen, 2020).

on valid information retrieval and inference strategies?” Then the next question is to specify what strategies would be valid and acceptable, which is possible to do on case-by-case basis.

With respect to ‘machine reading comprehension’, a recent proposal by [Dunietz et al. \(2020\)](#) is based on whether a model can extract certain information that should be salient for a human reader (e.g. spatial, temporal, causal relations in a story). However, a model can extract such ‘right’ information through a ‘wrong’ process, e.g. some shallow heuristic. Hence the definition of ‘NLU’ needs at least two components: (a) the specific information that the model is expected to be ‘extract’ from the text; and (b) a valid process with which such ‘extraction’ is performed. And this would still not be enough: the model could have simply memorized *both* the right answer and the strategy to find it for some limited set of examples. We argue that the third key prerequisite is the ability to generalize: to *consistently* use the ‘right’ information-seeking strategy in novel contexts.⁴

Thus we propose the following general success criteria for NLP systems:

Definition 1 *A NLP system has human-level competence with respect to its task X iff:*

- (a) *it is able to correctly perform the task X (identify the target information in QA, correctly classify texts, generate an appropriate translation etc.);*
- (b) *it does so by relying predominantly on information that a competent human speaker would also find relevant⁵;*
- (c) *it does so consistently under distribution shifts that do not pose challenges to competent human speakers.*

2.3 Reasoning an RC Model *should* Perform

The second principle in our Def. 1 is that the model should rely on the ‘right’ information. While models can discover patterns unknown to humans, a competent human reader should at least find such patterns relevant post-factum.

What information-seeking strategy is needed depends on the type of question and the context. [Rogers et al. \(2022\)](#) propose a classification of RC ‘skills’ into five main groups: situation/world modeling, different types of inference/logical reasoning,

⁴This does not preclude errors (humans make them too).

⁵Note that this leaves room for NLP systems to rely on patterns humans may not be even aware of, as long as such patterns are valid. E.g. if a system learned to make health outcome predictions based on latent information about unknown drug interactions, that would be the discovery of new knowledge that the experts would then accept – but not if its predictions were based on a spurious correlation with Marvel movie release dates.

the ability to combine information in multi-step reasoning, knowing what kind of information is needed and where to find it, and interpreting/manipulating linguistic input. A single question may require the competency of several types of ‘skills’.

This study contributes an empirical investigation on two RC ‘skills’ in the broad category of ‘interpreting/manipulating linguistic input’: coreference resolution and comparison. Both of them rely on the contextual information and linguistic competence. Assuming that a human reader would first read the question and then read the context in order to find the answer, they would need to perform roughly three steps: (a) to interpret the ‘question’ (akin to its transformation to a formal semantic representation or a query); (b) to identify the relevant information in the context through establishing the referential equality between expressions in the question and in the context; (c) to use that information to perform the operation of comparison or coreference resolution (see [Table 1](#)).⁶

2.4 Reasoning an RC Model *does* Perform

Having established what reasoning steps an RC model *should* perform, the next step would be to ascertain whether that is the case for specific models. But generally, the interpretability of DL models is an actively developed research area ([Belinkov and Glass, 2019](#); [Molnar, 2022](#)). In this study, we rely on a combination of two popular post-hoc explanation techniques, but we also discuss their limitations, and expect that new methods could soon be developed and used in the overall paradigm for the analysis of RC models that we propose.

Attribution/saliency-based methods [Li et al. \(2016\)](#); [Sundararajan et al. \(2017\)](#) provide a saliency score for each token in the input, which shows how ‘important’ a given token is for the model decision in this instance. [Figure 2](#) illustrates that such scores may not necessarily map onto human rationales.

To establish whether a model performs a given reasoning step (see [Table 1](#)), we define the following partition of the token space: the tokens the model *should* find important (positive) vs the ones it *should not* (negative). For example, to know if the model ‘attends’ to the entities being compared, we can define the positive partition as {blind,

⁶This definition could be developed further for more complex cases of coreference and comparison, or to model other variations of the human reading process, but this approximation suffices for our purposes and our RC data (see §3.1).

	Example	Step	Relevant Spans
Comparison	<p>Context: Blind Shaft is a 2003 film about a pair of brutal con artists operating in the illegal coal mines of present day northern China. The Mask Of Fu Manchu is a 1932 pre-Code adventure film directed by Charles Brabin.</p> <p>Question: Which film came out earlier, Blind Shaft or The Mask Of Fu Manchu?</p> <p>Answer: The Mask Of Fu Manchu</p>	Interpreting the question	<p><i>came out</i> relation: <film, release date></p> <p>film entities: <i>Blind Shaft</i>, <i>The Mask Of Fu Manchu</i></p> <p><i>earlier</i>: date comparison</p> <p>target: $\min(\text{release date}_{\text{Blind Shaft}}, \text{release date}_{\text{The Mask of Fu Manchu}})$</p>
		Identifying relevant information through referential equality	<p>$\text{Blind Shaft}_q := \text{Blind Shaft}_c$</p> <p>$\text{The Mask Of Fu Manchu}_q := \text{The Mask Of Fu Manchu}_c$.</p> <p><i>came out</i>_q := <date, film> construction_c</p> <p>release dates: <<i>Blind Shaft</i>, 2003>, <<i>The Mask Of Fu Manchu</i>, 1932></p>
		Value comparison	<p>solution: $\text{earlier}_q := \min_c$</p> <p>$\min(1932, 2003) = 1932$</p>
Coreference	<p>Context: Barack Obama was the 44th president of the US. He was born in Hawaii.</p> <p>Question: Who was born in Hawaii?</p> <p>Answer: Barack Obama.</p>	Interpreting the question	<p><i>born</i> relation: <person, location></p> <p><i>Hawaii</i>: location</p> <p>target: <i>born</i>: <<i>Hawaii</i>, UNK></p>
		Identifying relevant information through referential equality	<p>$\text{Hawaii}_q := \text{Hawaii}_c$</p> <p><i>born</i> relation: <<i>he</i>, <i>Hawaii</i>></p>
		Coref. resolution	<p><<i>Barack Obama</i>, <i>he</i>></p> <p>solution: <i>born</i> <<i>Barack Obama</i>, <i>Hawaii</i>></p>

Table 1: The basic reasoning steps for answering comparison and coreference questions.

[CLS] which film came out earlier , blind shaft or the mask of fu manchu ? [SEP] blind shaft is a 2003 film about a pair of brutal con artists operating in the illegal coal mines of present day northern china . the mask of fu manchu is a 1932 pre - code adventure film directed by charles bra ##bin . [SEP]

Figure 2: IG saliency scores example. Green/red denotes positive/negative scores.

shaft, mask, of, fu, munchu} and the negative partition as {northern, china}. If the model consistently follows this strategy, the average score should be higher for the positive rather than negative partition.

A limitation of saliency explanations is that they are not always faithful, i.e., do not reflect a model’s true decision process (Atanasova et al., 2020, 2022a; Ye et al., 2021). Also, even when they are faithful, i.e., when we can reliably say that a model places more ‘importance’ on token i than token j in an instance, this does not imply that a set of tokens I is more salient than another set J .

Counterfactual explanations have the form: “had X not occurred, Y would not have occurred” (Molnar, 2022). In NLP, they are based on input perturbations (Kaushik et al., 2020; Gardner et al., 2020; Sen et al., 2021; Atanasova et al., 2022b). In our case, it translates to “had the model not relied on information X, it could not have answered both

the original and the perturbed instance correctly”. Thus the perturbation has to change the correct label, unlike for contrast sets (Gardner et al., 2020).

Counterfactual (CF) explanations are considered to be more faithful, since they identify input features that impact predictions. However, they typically have to be manually generated (Kaushik et al., 2020), which makes large-scale CF generation prohibitively expensive (Khashabi et al., 2020).

We rely on both types of explanations as parallel sources of evidence about RC model reasoning, and define their alignment as follows:

Definition 2 Explanation Alignment. A CF and saliency-based explanation align when: (a) both the original and the counterfactually modified instance are answered correctly,⁷; and (b) the positive partition has a statistical significantly higher average saliency score than the negative partition.

We define the alignment score as follows:

Definition 3 Alignment Score: The Alignment Score for a <dataset, model, reasoning step> triple is the proportion of instances in that dataset for which different kinds of explanations align (according to our Def. 2).

We interpret a high alignment score as evidence that both kinds of explanations are faithful, and the model indeed performs the expected reasoning steps.

⁷i.e. there is an exact match between the predicted and the correct answer.

3 Methodology

3.1 Datasets and Models

For **Coreference**, we use the Quoref (Dasigi et al., 2019) dataset (20K training and 2.4K validation instances) where the annotators were asked to *design* questions for a given text so that answering those would require resolving anaphora. For **Comparison**, we sample questions from HotpotQA (Yang et al., 2018) and 2WikiMultiHopQA (Ho et al., 2020): two datasets with questions manually annotated with their reasoning type (bridge or comparison). We select the ‘comparison’ questions containing comparative adjectives or adverbs in them (23K training, 3K validation instances). These resources are based on Wikipedia and have multiple passages as contexts, but the sentences (typically 2-3) necessary to answer a question are marked as ‘supporting facts’. Since we are not focusing on the multi-hop information retrieval skill, we limit the contexts to these sentences.

We experiment with five pre-trained Transformer-based encoders of the BERT family: RoBERTa_{large} (Liu et al., 2019), BERT_{large-cased}, BERT_{base-cased} (Devlin et al., 2019), BERT_{medium}, and BERT_{small} (Turc et al., 2019; Bhargava et al., 2021). These BERT models differ mainly in the structure of architecture blocks and the number of parameters, while RoBERTa also has a different training corpus and optimization. Since larger models were shown to generalize better for some use cases (Hendrycks et al., 2020; Bhargava et al., 2021), we investigate whether they also are more likely to be right for the right reasons.

We fine-tune each encoder using the architecture in Devlin et al. (2019) (see the appendix for details) and evaluate them on the validation set (as the test sets are not public). We use the standard evaluation metrics in extractive QA: **F1-Score** (the percentage of token overlap between predicted and ‘gold’ answers, averaged over all data points), and **Exact-match** (the number of data points where the predicted answer matches the ‘gold’ answer).

3.2 Counterfactual Explanations

Our formulation of reasoning (Table 1) consists of three basic steps for both coreference and comparison: interpreting the question, identifying the relevant information through referential equality, and the target operation on the identified information (coreference resolution or value comparison). We focus on the final step, since: (a) it implicitly

relies on correct semantic parsing of the question and the context; (b) referential equality in our data is in large part trivial: most entities have the same surface form in the question and the text.

An obvious semantically valid perturbation that should change the prediction (and thus test for the model’s understanding of the comparison operation) is to replace the comparative adjectives with their antonyms (Figure 3d). Since our sample only contains 6 tokens used as comparison operators, we define appropriate replacements manually.⁸

For coreference questions, a competent RC model would at least resolve the coreference chain for the target entity. A context can have many coreference clusters, so we need to identify the relevant one. In the Quoref dataset, we use the instances where the relevant cluster itself contains the answer entity⁹ (see Figure 3a), and therefore, can be extracted automatically. This leaves us with 55%(1329/2418) of the validation instances. These are further subsampled to manually create 100 CF instances by inserting a new sentence, which includes the new and excludes the old answer (see Figure 3b). Similarly to the comparison questions, the original answer entity remains in the context. If the model uses the ‘shortcut’ of choosing the most frequent entity in the context (Wu et al., 2021), it should not be able to answer both the original and the perturbed instance correctly.

3.3 Saliency-based Explanations

We obtain token saliency scores from two families of attribution/saliency methods: Occlusion (DeYoung et al., 2020), a method based on perturbations, and Integrated Gradients (IG, Sundararajan et al. (2017)), a method based on gradients.¹⁰

Design decisions: RC models typically predict two scores (t_s, t_e) for each token t : the probability of t being the start and the end of the answer span. Any attribution method produces two scores (A_{start}^t, A_{end}^t) for each token t , indicating how ‘im-

⁸earlier \leftrightarrow later, first \rightarrow later, more recently \rightarrow earlier, older \leftrightarrow younger.

⁹We extract the clusters using an off-the-shelf coreference resolver (Clark and Manning, 2016) implemented in Spacy.

¹⁰Atanasova et al. (2020) shows that for Transformer based architectures, Occlusion is the best perturbation method by two evaluation criteria: agreement with human rationale and faithfulness. A recent paper by Ye et al. (2021) finds IG to be one of the most faithful gradient-based methods for extractive QA, only outperformed by Layerwise Attention Attribution (LAA), a method proposed in the paper itself. We leave LAA and other popular explainability methods such as LIME (Ribeiro et al., 2016) for future work.

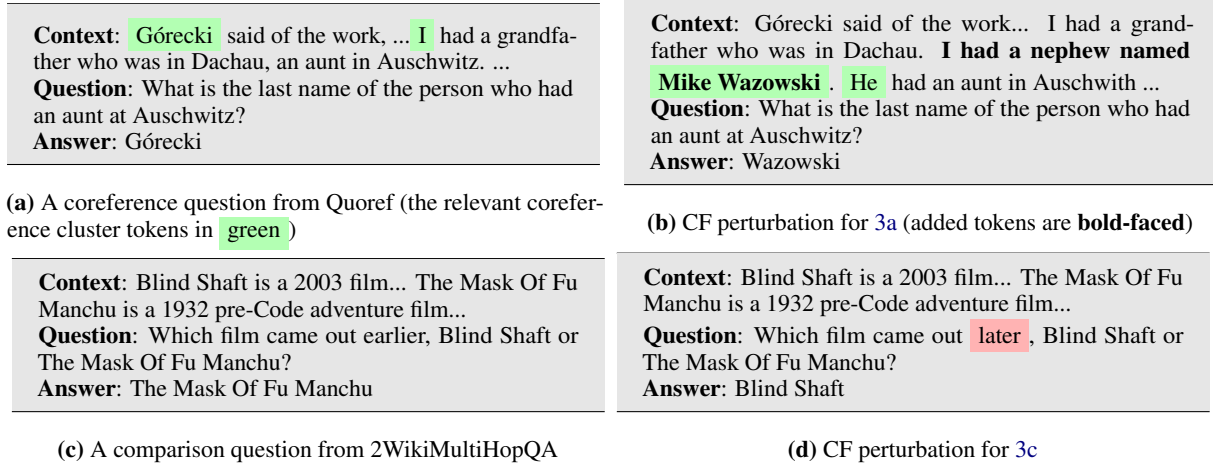


Figure 3: Examples of CF perturbations used in this study.




Figure 4: Positive and negative partitions for saliency explanations.

portant’ t is for predicting the start/end of the answer span. Following Kokhlikyan et al. (2020), we use A_{start} in all our saliency experiments.¹¹

For Occlusion, we calculate A_{start}^t by replacing t in the input with a baseline token (MASK) and measuring the change in t_s . DNNs map an input vector to a scalar value (loss/ class probability). Gradient-based methods measure A_{start}^t using the gradient of the token t w.r.t. this scalar function (we use $\text{argmax}(t_s)$). IG sums these gradient values along a linear path from a baseline to the current instance. Both Occlusion and IG need a baseline token, which for us is the MASK token.

Gradient-based methods in NLP do not produce a scalar saliency score, i.e., A_{start}^t is a vector because the input is an embedding *matrix* and not a vector. Two common ways to summarize this vector to a scalar are: (a) scalar product between the input and the gradient vector (Han et al., 2020); or (b) l_p norm, where $p \in 1, 2$ (Atanasova et al., 2020). We use l_2 norm (see the discussion in §4.3).

Token partitions: Figure 4 shows the token

¹¹We also briefly experimented with $(A_{start} + A_{end})/2$ for Occlusion but it yielded very similar saliency ranking of the tokens on a 100 sample subset of the Comparison dataset.

partitions used for the same reasoning steps (comparison and coreference resolution) that we also target with the CF perturbations. For comparison the positive partition consists of the **question** token(s) expressing the comparison operation (e.g. ‘more recently’). The negative partition consists of question tokens that are not in the set of entities or values that need to be compared, or in the set of verbs (which could capture the relation between the entities and their values). For coreference resolution, the positive partition is the **context** tokens in the relevant coreference cluster (§3.2). The negative partition is the set of context tokens that are not in: (a) the positive partition; and (b) match the question tokens.

4 Results & Analysis

4.1 Base Model Performance

As a sanity check, we fine-tune all models on the data described in §3.1 (Table 2). For coreference, the F1-Score of our best model (RoBERTa_{large}) is slightly better (82.10) than the previously reported score (79.64, Wu et al. (2021)). The comparison instances are sampled from *parts* of two datasets, and so a direct comparison is not possible.¹²

The size and the model family matter: RoBERTa performs better than BERT for two models of the same size, and the larger models do better. Interestingly, the difference is more pronounced for the Quoref dataset, where the instances have longer contexts and the questions are more complex.

¹²The best model (RoBERTa_{large}) has an F1-Score of 92%, slightly better than the highest score reported on the HotpotQA leaderboard (89.14%) and much better than the baseline model for the 2WikiMultiHopQA dataset (65.02, (Ho et al., 2020)).

	Comparison		Coreference	
	F1	EM	F1	EM
RoBERTa _{large}	92.08	91.07	82.10	79.39
BERT _{large-cased}	89.23	88.57	71.91	68.47
BERT _{base-cased}	89.38	88.37	64.62	59.38
BERT _{medium}	86.45	85.96	60.16	54.82
BERT _{small}	71.44	69.87	50.94	43.39

Table 2: Average (3 runs) results of different models on Comparison and Coreference datasets. The STD varies between 0.01 – 0.72%. **Green** indicates the best scores.

	Coreference		Comparison	
	og	cf	og	cf
RoBERTa _{large}	92.0	70.7	99.4	98.9
BERT _{large-cased}	86.2	50.8	98.9	93.1
BERT _{base-cased}	82.5	39.2	98.4	91.8
BERT _{medium}	74.0	35.8	97.4	96.5
BERT _{small}	67.2	29.4	68.2	45.3

Table 3: F1-Score for the original and the CF perturbations. **Red** denotes significant drop.

4.2 Counterfactual Explanations

Table 3 compares the F1-Score of the original (‘og’) vs counterfactual (‘cf’) instances. For the comparison questions, the performance on the original and CF instances are very close for all models except BERT_{small}. Bigger models consistently perform better, but in most cases the difference with the next larger model is relatively small.

For coreference questions, CF instances are much more difficult for all models. Even the best model RoBERTa_{large} experiences a 24% drop. All BERT models perform poorly: even the larger ones have a 40% performance drop (BERT_{large-cased}). Thus, the CF tests show that the models are more likely to follow the expected reasoning strategy for comparison, but not for the coreference questions.

4.3 Alignment Score

For statistical significance testing in ‘Expectation Alignment’ (Def. 2), we use a one-tailed independent t_{test} ($p = 0.05$) with the null hypothesis that *the positive partition does not have a higher average saliency score*. Table 4 shows the ‘Alignment Score’ (Def. 3) results for comparison and coreference resolution (§3.3), using saliency scores from IG and Occlusion.

Ideally, for a random partition of tokens in any instance, the positive and the negative partitions should have similar saliency scores. For a dataset, they should be *significantly* different in $\approx 0\%$

	Coreference		Comparison	
	IG	Occ	IG	Occ
RoBERTa _{large}	33.3	69.7	33.8	67.0
BERT _{large-cased}	12.5	58.3	34.1	65.9
BERT _{base-cased}	21.4	42.9	83.8	69.0
BERT _{medium}	81.8	36.4	82.2	42.0
BERT _{small}	83.3	33.3	86.3	16.3

Table 4: Alignment score between counterfactual explanations vs IG (Integrated Gradients) or Occ (Occlusion). **Green** indicates methods with $> 80\%$ alignment.

cases.¹³ For Occlusion, the saliency scores are significantly different in only 5.6 – 8.2% instances for a random partition. Recall that in §3.3 we discussed 3 summarizers for IG. Among all of them, the l_2 norm is the only one where this happens in 5.2 – 7.3% cases, for the other two the numbers are between 11.3 – 28.9%.

Table 4 shows that, counter-intuitively, for both comparison and coreference questions the larger models overall have *lower* IG alignment scores, meaning that they do not pay as much ‘attention’ to the tokens we defined as important. This is despite the fact that for comparison the above CF experiment suggests that the models do perform the expected reasoning operations. One possible explanation is that IG simply does not reliably capture the model’s reasoning process, and Occlusion does better at that because its trend in alignment is the opposite of IG: bigger models tend to have significantly higher alignment scores.¹⁴

Another possible explanation is that IG explanations are in fact faithful, but, having more ‘attention’ to the tokens we defined as important is counter-productive. Consider that the BERT_{small} model achieves an Exact-match of 87% on the original questions containing the comparative tokens ‘earlier’, ‘first’ and ‘older’ (which are 2.1 times more frequent in the training data than all others), and an Exact-match of 28% on the other original questions. Yet overall the model performs poorly, and thus the reliance on these highly frequent comparative adjectives could be a bug rather than a feature. As this hypothesis brings into question the overall utility of saliency-based explanations for testing for the ‘correct’ reasoning steps, we hope it will be investigated in more depth in future work.

¹³Aggregation of local explanations such as saliency scores are not *guaranteed* to produce faithful global explanations (Setzu et al., 2021), but this is a convincing evidence.

¹⁴The lack of alignment between the two techniques is consistent with the findings of Atanasova et al. (2020).

	Supporting Facts		Paragraphs			
	OG	CF	CF-ood	OG	CF	CF-ood
RoBERTa _{large}	99.4	98.9	77.2	98.7	96.4	74.8
BERT _{large-cased}	98.9	93.1	68.7	98.0	90.8	67.5
BERT _{base-cased}	98.4	91.8	58.1	97.0	86.8	59.9
BERT _{medium}	97.4	96.5	64.4	96.2	86.3	66.3
BERT _{small}	68.2	45.3	57.1	68.3	47.6	58.8

Table 5: F1-Score for the original (OG) comparison questions and their counterfactual perturbations in (CF) and out (CF-ood) of the training distribution. The models are provided either a smaller context of supporting facts or full paragraphs.

Red indicates a significant drop in performance.

4.4 Generalization Tests

Table 3 shows that when measured with CF tests, most models do not follow the expected coreference resolution strategy, but they do so for comparison. Still, based on our success criteria (Def. 1), we cannot yet conclude that they ‘understand’ comparison. A human would be able to disassociate the logical operation of comparison from the surface realizations, i.e., they would be able to answer a question correctly with either of the surface forms ‘younger’ and ‘more junior’.

For the CF experiments reported up until this point the perturbations were in-distribution, i.e., the training data had both the original question “who is younger” and the CF “who is older”. Now we replace the comparative adjectives with antonyms that are not in the training data (see the appendix for details). We also increase the context size by using full paragraphs instead of just the sentences marked as ‘supporting facts’, to see if the models would be ‘distracted’ by more information.

Table 5 shows a considerable drop in performance for CF-ood condition for all models. The larger models generalize better: RoBERTa_{large} and BERT_{large-cased} perform 2% and 8% worse for CF questions, whereas BERT_{small} exhibits a 29% reduction. The ‘supporting facts only’ condition is overall easier than the ‘paragraphs’ condition.

4.5 Heuristics for Coreference Questions

Since the CF tests (§4.2) do not show that BERT models can cope with the altered coreference chains, we have to conclude that they do not follow the expected reasoning steps. Though given the above-chance performance they must follow some other strategy. We test the hypothesis that many of

	Coreference		SQuAD	
	F1-Score	EM	F1-Score	EM
Token overlap	21.5	12.9	26.68	21.64
LCS	17.2	12.9	19.59	15.97
Position	12.3	7.9	21.62	16.32
Sentence encoder	20.43	9.67	25.91	20.61

Table 6: Results for different heuristic methods on the coreference and SQuAD datasets. Green indicates the best score.

the coreference questions can be answered by simple heuristics and that the models resort to those. Specifically, we define an unsupervised dataset-independent heuristic method consisting of two steps: sentence selection and phrase extraction.

Sentence Selection: Among all the context sentences $\{c_i\}$, select the one that is the ‘closest’ to the question q . We experiment with 4 options for similarity: **token-overlap** (number of common tokens in q and c_i), **sentence encoder** (cosine similarity between the sentence embeddings of q and c_i created by a sentence encoder (Reimers and Gurevych, 2019)), **LCS** (number of tokens in the Longest Common Subsequence between q and c_i), and **position** (simply taking the first sentence in the context following Ko et al. (2020)).

Phrase Extraction: We assume that the model would also learn to look for a named entity in the selected sentence. The question dictates the *type* of this entity (e.g. ‘where’ \rightarrow location, ‘who’ \rightarrow person name). The type could be determined by a simple mapping between ‘wh’ question words and entity types, but this can fail (e.g. for the question “Who won the World Cup in 2002?” the expected answer is a location, not a person). Therefore, we fine-tune a Transformer model to predict the answer type from the question.¹⁵

Table 6 shows the best heuristic has an F1-Score of 21.5% on the coreference dataset, and 26.68% on SQuAD (Rajpurkar et al., 2016), which we use for validation. The SQuAD score is comparable to the previously reported result of 26.7% in Sen and Saffari (2020) for an algorithm predicting entity types heuristically, and choosing the entity from the whole context instead of the best possible sentence. Ray Choudhury et al. (2022) uses a similar approach to find Quoref questions that can be answered heuristically, but our algorithm has

¹⁵The accuracy for this model is 85.7%. See the appendix for results from multiple models and loss functions.

more sentence selection strategies, and unlike ours, Ray Choudhury et al. (2022) only uses one loss function in the phrase extraction model.

Nevertheless, the best heuristic algorithm performs considerably worse than the smallest BERT_{small} model (51%, Table 2). Performance alone cannot reveal whether this strategy is used in the instances where it *would* be sufficient, but this result shows that even the smaller models must either rely on a more successful (but still imperfect) strategy, or at least rely on more than one heuristic. The problem with discovering potential ‘shortcuts’ in low-performing models is complicated as these strategies are not necessarily human-interpretable: González et al. (2021) show that humans struggle to predict the answer chosen by poorly performing RC models, even when the saliency explanations for that answer is shown, because these answers simply do not align with human RC strategies.

5 Discussion and Related Work

Our work continues the emerging trend of research on being ‘right for the right reasons’ (McCoy et al., 2019; Chen and Durrett, 2019; Min et al., 2019; Atanasova et al., 2022b, inter alia). We contribute stricter success criteria for behavioral tests of NLP models (Def. 1), and, for the RC task, develop the methodology of: (a) defining what information the model should rely on for a given linguistic, logical, or world knowledge ‘skill’; (b) systematically testing the behavior of RC models with interpretability techniques for whether they rely on that information. This is most closely related to the work on ‘defining comprehension’ by Dunietz et al. (2020), though their testing is limited to probing the models with RC questions. Another related study is the QED framework (Lamm et al., 2021), annotating Natural Questions (Kwiatkowski et al., 2019) with ‘explanations’ of the expected reasoning process. Their expected reasoning process also contains 3 steps, partly similar to ours: selecting a relevant sentence, referential equality, and deciding on whether this sentence entails the predicate in the question. However, the goals of QED are to: (a) predict both the answer and the explanation for a question; and (b) understand if explanations help QA models. Such explanation annotations are unavailable for most datasets, and few QA models produce explanations. Therefore, our approach of: (a) defining expected reasoning steps; and (b) using model interpretations to validate such steps applies

to a broader class of models.

This study is also related to the overall efforts to define what kinds of ‘skills’ RC models can be expected to exhibit (Sugawara et al., 2018; Schlegel et al., 2020; Rogers et al., 2022). While these works focus on the high-level taxonomies of ‘skills’, we contribute practical definitions for two linguistic ‘skills’ (comparison and coreference resolution) which could be used for analyzing model performance. Implicitly, research proposing RC resources that target various specific ‘skills’ (e.g. TempQuestions (Jia et al., 2018) for temporal order, MathQA (Amini et al., 2019) for numerical reasoning, etc.) also contributes to this area, but they typically rely on broad linguistic definitions rather than on steps for machine reasoning.

The saliency techniques we rely on have previously been used for extractive QA (Madsen et al., 2021), but we are among the first (Ye et al., 2021) to investigate their correlation with counterfactual explanations. For counterfactual perturbations, we also ensure that the perturbations are human-interpretable and change the prediction, which is not the case for adding incomprehensible text (Kaushik and Lipton, 2018), removing words from questions, shuffling the context (Sen and Saffari, 2020), or replacing context tokens with random tokens (Sugawara et al., 2020).

6 Conclusion

Making progress towards trustworthy NLP models requires specific definitions for the behavior expected of these models in different situations. We propose a framework for RC model analysis that involves: (a) the definition of the expected ‘reasoning’ steps; (b) analysis of model behavior. We contribute such definitions for two linguistic ‘skills’ (comparison and coreference resolution), and use parallel explainability techniques to investigate whether RC models based on BERT family encoders answer such questions correctly for the right reasons. We find that to be the case for comparison, but not for coreference. Moreover, we find that, even for comparison, the models ‘break’ when encountering out-of-distribution counterfactual perturbations, suggesting that they memorize specific lexical patterns rather than learn more general reasoning ‘skills’. As such, more research is needed on developing definitions and tests for specific ‘skills’ expected of NLU models, as well as on more faithful interpretability techniques.

7 Acknowledgements

We would like to thank Pepa Atanasova, Gary Marcus, Mark Steedman, and Bonnie Webber for the discussion of various aspects of this work. We also thank the anonymous reviewers for their time and insightful comments.

References

- Rakefet Ackerman, David Leiser, and Maya Shpigelman. 2013. [Is Comprehension of Problem Solutions Resistant to Misleading Heuristic Cues?](#) *Acta Psychologica*, 143(1):105–112.
- Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. 2019. [MathQA: Towards Interpretable Math Word Problem Solving with Operation-Based Formalisms](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2357–2367, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. [A Diagnostic Study of Explainability Techniques for Text Classification](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 3256–3274. Association for Computational Linguistics.
- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2022a. [Diagnostics-Guided Explanation Generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):10445–10453.
- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2022b. [Fact Checking with Insufficient Evidence](#). *Transactions of the Association for Computational Linguistics*, 10:746–763.
- Yonatan Belinkov and James Glass. 2019. [Analysis Methods in Neural Language Processing: A Survey](#). *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Prajjwal Bhargava, Aleksandr Drozd, and Anna Rogers. 2021. [Generalization in NLI: Ways \(Not\) To Go Beyond Simple Heuristics](#). In *Proceedings of the Second Workshop on Insights from Negative Results in NLP*, pages 125–135, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- M. F. Burnyeat and Jonathan Barnes. 1980. [Socrates and the Jury: Paradoxes in Plato’s Distinction between Knowledge and True Belief](#). *Proceedings of the Aristotelian Society, Supplementary Volumes*, 54:173–206.
- Danqi Chen, Jason Bolton, and Christopher D. Manning. 2016. [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. Association for Computational Linguistics.
- Jifan Chen and Greg Durrett. 2019. [Understanding Dataset Design Choices for Multi-hop Reasoning](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4026–4032. Association for Computational Linguistics.
- Kiel Christianson. 2016. [When Language Comprehension Goes Wrong for the Right Reasons: Good-enough, Underspecified, or Shallow Language Processing](#). *Quarterly Journal of Experimental Psychology*, 69(5):817–828.
- Kevin Clark and Christopher D. Manning. 2016. [Deep Reinforcement Learning for Mention-Ranking Coreference Models](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2256–2262. The Association for Computational Linguistics.
- Pradeep Dasigi, Nelson F. Liu, Ana Marasovic, Noah A. Smith, and Matt Gardner. 2019. [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 2019, Hong Kong, China, November 3-7, 2019*, pages 5924–5931. Association for Computational Linguistics.
- Finnur Dellsén. 2020. [Beyond Explanation: Understanding as Dependency Modelling](#). *The British Journal for the Philosophy of Science*, 71(4):1261–1286.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. 2020. [ERASER: A Benchmark to Evaluate Rationalized NLP Models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 4443–4458. Association for Computational Linguistics.

- Janet M. Duffin and Adrian P. Simpson. 2000. [A Search for Understanding](#). *The Journal of Mathematical Behavior*, 18(4):415–427.
- Jesse Dunietz, Gregory Burnham, Akash Bharadwaj, Owen Rambow, Jennifer Chu-Carroll, and David A. Ferrucci. 2020. [To Test Machine Comprehension, Start by Defining Comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7839–7859. Association for Computational Linguistics.
- Charles J Fillmore. 1985. Frames and the Semantics of Understanding. *Quaderni di Semantica*, 6(2):222–254.
- 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 E. 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, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 1801–1813. Association for Computational Linguistics.
- Ana Valeria González, Anna Rogers, and Anders Søgaard. 2021. [On the Interaction of Belief Bias and Explanations](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2930–2942, Online. Association for Computational Linguistics.
- Stephen Grimm. 2012. [The Value of Understanding](#). *Philosophy Compass*, 7(2):103–117.
- Stephen Grimm. 2021. [Understanding](#). In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, summer 2021 edition. Metaphysics Research Lab, Stanford University.
- 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, ACL 2020, Online, July 5-10, 2020*, pages 5553–5563. Association for Computational Linguistics.
- Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzic, Rishabh Krishnan, and Dawn Song. 2020. [Pretrained Transformers Improve Out-of-Distribution Robustness](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2744–2751, Online. Association for Computational Linguistics.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. [Constructing A Multi-hop QA Dataset for Comprehensive Evaluation of Reasoning Steps](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 6609–6625. International Committee on Computational Linguistics.
- Robin Jia and Percy Liang. 2017. [Adversarial Examples for Evaluating Reading Comprehension Systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2021–2031. Association for Computational Linguistics.
- Zhen Jia, Abdalghani Abujabal, Rishiraj Saha Roy, Jan-nik Strötgen, and Gerhard Weikum. 2018. [TempQuestions: A Benchmark for Temporal Question Answering](#). In *Companion of the The Web Conference 2018 on The Web Conference 2018 - WWW '18*, pages 1057–1062, Lyon, France. ACM Press.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. [Is BERT Really Robust? A Strong Baseline for Natural Language Attack on Text Classification and Entailment](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8018–8025. AAAI Press.
- Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York.
- Nora Kassner and Hinrich Schütze. 2020. [Negated and Misprimed Probes for Pretrained Language Models: Birds Can Talk, But Cannot Fly](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7811–7818. Association for Computational Linguistics.
- Divyansh Kaushik, Eduard H. Hovy, and Zachary Chase Lipton. 2020. [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.
- 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, Brussels, Belgium, October 31 - November 4, 2018*, pages 5010–5015. Association for Computational Linguistics.
- Daniel Khashabi, Tushar Khot, and Ashish Sabharwal. 2020. [More Bang for Your Buck: Natural Perturbation for Robust Question Answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 163–170. Association for Computational Linguistics.
- Yoon Kim. 2014. [Convolutional Neural Networks for Sentence Classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A Method for Stochastic Optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Miyoung Ko, Jinhyuk Lee, Hyunjae Kim, Gangwoo Kim, and Jaewoo Kang. 2020. [Look at the First Sentence: Position Bias in Question Answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 1109–1121. Association for Computational Linguistics.
- Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, and Orion Reblitz-Richardson. 2020. [Captum: A unified and generic model interpretability library for PyTorch](#). *CoRR*, abs/2009.07896.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: a benchmark for question answering research](#). *Trans. Assoc. Comput. Linguistics*, 7:452–466.
- Matthew Lamm, Jennimaria Palomaki, Chris Alberti, Daniel Andor, Eunsol Choi, Livio Baldini Soares, and Michael Collins. 2021. [QED: A Framework and Dataset for Explanations in Question Answering](#). *Trans. Assoc. Comput. Linguistics*, 9:790–806.
- Arthur D. Lander. 2010. [The Edges of Understanding](#). *BMC Biology*, 8(1):40.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016. [Understanding Neural Networks through Representation Erasure](#). *CoRR*, abs/1612.08220.
- Tal Linzen. 2020. [How Can We Accelerate Progress Towards Human-like Linguistic Generalization?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 5210–5217. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#). *CoRR*, abs/1907.11692.
- Andreas Madsen, Nicholas Meade, Vaibhav Adlakha, and Siva Reddy. 2021. [Evaluating the Faithfulness of Importance Measures in NLP by Recursively Masking Allegedly Important Tokens and Retraining](#). *CoRR*, abs/2110.08412.
- Gary Marcus and Ernest Davis. 2019. *Rebooting AI: Building Artificial Intelligence We Can Trust*. Knopf Doubleday Publishing Group.
- Gary Marcus, Francesca Rossi, and Manuela Veloso. 2016. [Beyond the Turing Test](#). *AI Magazine*, 37(1):3–4.
- 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 Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3428–3448. Association for Computational Linguistics.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient Estimation of Word Representations in Vector Space](#). In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Sewon Min, Eric Wallace, Sameer Singh, Matt Gardner, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019. [Compositional Questions Do Not Necessitate Multi-hop Reasoning](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4249–4257. Association for Computational Linguistics.
- Melanie Mitchell. 2021. [Why AI is Harder than We Think](#). In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '21*, page 3, New York, NY, USA. Association for Computing Machinery.
- Christoph Molnar. 2022. *Interpretable Machine Learning*, 2 edition. LeanPub.
- T Nurhuda, D Rusdiana, and W Setiawan. 2017. [Analyzing Students' Level of Understanding on Kinetic Theory of Gases](#). *Journal of Physics: Conference Series*, 812:012105.

- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. 2013. [Towards Robust Linguistic Analysis using OntoNotes](#). In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning, CoNLL 2013, Sofia, Bulgaria, August 8-9, 2013*, pages 143–152. ACL.
- Yada Pruksachatkun, Phil Yeres, Haokun Liu, Jason Phang, Phu Mon Htut, Alex Wang, Ian Tenney, and Samuel R. Bowman. 2020. [jiant: A Software Toolkit for Research on General-Purpose Text Understanding Models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 109–117, Online. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100, 000+ Questions for Machine Comprehension of Text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392. The Association for Computational Linguistics.
- Sagnik Ray Choudhury, Nikita Bhutani, and Isabelle Augenstein. 2022. [Can Edge Probing Tests Reveal Linguistic Knowledge in QA Models?](#) In *Proceedings of the 29th International Conference on Computational Linguistics*. International Committee on Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [SentenceBERT: Sentence Embeddings using Siamese BERT-Networks](#). 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. ["Why Should I Trust You?": Explaining the Predictions of Any Classifier](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, pages 1135–1144. ACM.
- Anna Rogers, Matt Gardner, and Isabelle Augenstein. 2022. [QA Dataset Explosion: A Taxonomy of NLP Resources for Question Answering and Reading Comprehension](#). *Computing Surveys (CSUR)*, to appear.
- Anna Rogers, Olga Kovaleva, Matthew Downey, and Anna Rumshisky. 2020. [Getting Closer to AI Complete Question Answering: A Set of Prerequisite Real Tasks](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8722–8731. AAAI Press.
- Barbara Rychalska, Dominika Basaj, Anna Wróblewska, and Przemyslaw Biecek. 2018. [Does It Care What You Asked? Understanding Importance of Verbs in Deep Learning QA System](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 322–324. Association for Computational Linguistics.
- Viktor Schlegel, Marco Valentino, André Freitas, Goran Nenadic, and Riza Batista-Navarro. 2020. [A Framework for Evaluation of Machine Reading Comprehension Gold Standards](#). In *Proceedings of The 12th Language Resources and Evaluation Conference, LREC 2020, Marseille, France, May 11-16, 2020*, pages 5359–5369. European Language Resources Association.
- Indira Sen, Mattia Samory, Fabian Flöck, Claudia Wagner, and Isabelle Augenstein. 2021. [How Does Counterfactually Augmented Data Impact Models for Social Computing Constructs?](#) In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 325–344, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Priyanka Sen and Amir Saffari. 2020. [What do Models Learn from Question Answering Datasets?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2429–2438. Association for Computational Linguistics.
- Mattia Setzu, Riccardo Guidotti, Anna Monreale, Franco Turini, Dino Pedreschi, and Fosca Giannotti. 2021. [GLocalX - From Local to Global Explanations of Black Box AI Models](#). *Artif. Intell.*, 294:103457.
- 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, Brussels, Belgium, October 31 - November 4, 2018*, pages 4208–4219. Association for Computational Linguistics.
- Saku Sugawara, Pontus Stenetorp, Kentaro Inui, and Akiko Aizawa. 2020. [Assessing the Benchmarking Capacity of Machine Reading Comprehension Datasets](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8918–8927. AAAI Press.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. [Axiomatic Attribution for Deep Networks](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 3319–3328. PMLR.

- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Well-Read Students Learn Better: The Impact of Student Initialization on Knowledge Distillation](#). *CoRR*, abs/1908.08962.
- A. M. Turing. 1950. [Computing Machinery and Intelligence](#). *Mind*, 59(236):433–460.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 3261–3275.
- Alex Wang, Amapreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2018. [GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Mingzhu Wu, Nafise Sadat Moosavi, Dan Roth, and Iryna Gurevych. 2021. [Coreference Reasoning in Machine Reading Comprehension](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 5768–5781. Association for Computational Linguistics.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2369–2380. Association for Computational Linguistics.
- Xi Ye, Rohan Nair, and Greg Durrett. 2021. [Connecting Attributions and QA Model Behavior on Realistic Counterfactuals](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 5496–5512. Association for Computational Linguistics.
- Matthew D. Zeiler. 2012. [ADADELTA: An Adaptive Learning Rate Method](#). *CoRR*, abs/1212.5701.
- Rolf A. Zwaan. 2016. [Situation Models, Mental Simulations, and Abstract Concepts in Discourse Comprehension](#). *Psychonomic Bulletin & Review*, 23(4):1028–1034.

A Appendix

A.1 QA Model Training

For training the QA models in §3.1 The questions and contexts are concatenated, and a linear layer on top of the encoder is used to predict the probability of a context token i being the start ($P_{i,s}$) or end ($P_{i,e}$) of an answer. The score ($S_{i,j}$) for a span with start token i and end token j is computed as $P_{i,s} + P_{j,e}$. For all valid combination of i and j , the span with the highest score is chosen as the answer. A cross entropy loss between the actual and predicted start/end positions is minimized.

The models were trained for 10 epochs with a batch size of 16 using the Adam optimizer (Kingma and Ba, 2015) ($\beta_1 = 0.9, \beta_2 = 0.99, \epsilon = 1e-8, \text{weight_decay} = 0.01$) and gradient clipping. The learning rate (LR) was kept at $1e-05$ with a linear warm-up schedule (starting LR=0). The models were evaluated on a subset of the validation data every 500 mini-batches with early stopping on 100 evaluations (Pruksachatkun et al., 2020). The LR and batch size was determined by a small grid search on the coreference dataset: LR= $\{1e-05, 1e-04, 1e-03\}$, batch size = $\{8, 16, 32\}$.

A.2 Antonym Replacements for CF Generation

The antonym replacements for the generalization test (§4.4) are described below:

- first → less recently
- older → less old, more junior, less mature, less grown-up
- earlier → subsequently, thereafter, less recently
- later → less recently
- younger → more old, less junior, more mature, more grown-up
- more recently → less recently, longer ago

A.3 Supervised Entity Type Predictor

Our goal is to build a classifier to predict the answer entity type from the question (§4.5). A sample data point is shown in Figure 5. The entity types are defined in the Ontonotes-5 dataset (Pradhan et al., 2013). The answer entity type is detected

Text: What is the full name of the person who is the television reporter that brings in a priest versed in Catholic exorcism rites?
Label: PER

Figure 5: A sample instance for answer entity type classifier.

from the context using an off-the-shelf entity detector implemented in Spacy.¹⁶ When the answer is not a named entity, or the entity detector fails to determine its type, that question is discarded.

The classification models are trained on the **training** portion of Quoref and SQuAD which is further divided into train/dev/test (70/20/10) split for training and evaluation. The distribution of the class labels is very skewed.

Models: We use two types of models: 1) a fine-tuned 12 layer 768 dimensional BERT_{base-cased} model; and 2) a popular word convolutional model for sentence classification (Kim, 2014) using three parallel filters (size 3, 4, and 5) and 300 dimensional Google News Word2Vec representations (Mikolov et al., 2013).

BERT model: This model is trained for 5 epochs, with Adam optimizer (Kingma and Ba, 2015) with a weight decay of $1.0e-08$ and a learning rate of $1.0e-05$. The sequence max length is kept at 128. We search for two hyperparameters: 1) number of epochs: 3-7, increasing by 1; and 2) learning rate: $1.0e-05, 5.0e-05, 1.0e-04$.

WordConv model: This model is trained for 40 epochs, with Adadelta optimizer (Zeiler, 2012) with a learning rate of $1.0e-05$. The sequence max length is again kept at 128.

For both models, accuracy was used as the early stopping metric. We minimized the cross entropy (CE) loss in general, but for the WordConv model, a weighted CE loss was also implemented to account for the training data class-imbalance in Quoref. That did not improve the results significantly and was not used in the BERT_{base-cased} model. Table 7 shows the detailed results. Finally, we choose the fine-tuned BERT_{base-cased} model as the entity detector as it performs the best. Ray Choudhury et al. (2022) also proposes a model to determine the answer entity type from a question, but the major difference is the label space. The model in Ray Choudhury et al. (2022) is trained to predict a label of “UNKNOWN_ENTITY” when the an-

¹⁶<https://spacy.io>

Dataset	Model	Accuracy	Macro F1
SQuAD	BERT _{base-cased}	76.4	56.2
	WordConv	72.4	44.9
Coref	BERT _{base-cased}	85.7	73.9
	WordConv	85.0	67.6
	Weighted BCE	85.3	69.7

Table 7: Models for supervised entity type selection. Green indicates the best results.

swer span is a) not a named entity or b) the entity detector can not find its type. However, an “UNKNOWN_ENTITY” label does not help the final algorithm (heuristic answer selection) to find the correct answer span. Therefore, our model never predicts this label, and consequently, has a better accuracy than [Ray Choudhury et al. \(2022\)](#). It potentially makes a mistake on the test data points that fall in the previous two categories, but the final algorithm is no worse than [Ray Choudhury et al. \(2022\)](#).