

# An Analysis of Dataset Overlap on Winograd-Style Tasks

**Ali Emami**

Mila/McGill University  
ali.emami@mail.mcgill.ca

**Adam Trischler, Kaheer Suleman**

Microsoft Research Montreal  
{adam.trischler, kasulema}@microsoft.com

**Jackie Chi Kit Cheung**

Mila/McGill University  
jcheung@cs.mcgill.ca

## Abstract

The Winograd Schema Challenge (WSC) and variants inspired by it have become important benchmarks for common-sense reasoning (CSR). Model performance on the WSC has quickly progressed from chance-level to near-human using neural language models trained on massive corpora. In this paper, we analyze the effects of varying degrees of overlap between these training corpora and the test instances in WSC-style tasks. We find that a large number of test instances overlap considerably with the corpora on which state-of-the-art models are (pre)trained, and that a significant drop in classification accuracy occurs when we evaluate models on instances with minimal overlap. Based on these results, we develop the KNOWREF-60K dataset, which consists of over 60k pronoun disambiguation problems scraped from web data. KNOWREF-60K is the largest corpus to date for WSC-style common-sense reasoning and exhibits a significantly lower proportion of overlaps with current pretraining corpora.

## 1 Introduction

The original purpose of the Winograd Schema Challenge was to serve as an alternative Turing test to evaluate an automatic system’s capacity for common-sense inference (Levesque et al., 2011). As an example:

- (1) a. Jim yelled at Kevin because *he* was so upset. (Answer: Jim)
- b. Jim comforted Kevin because *he* was so upset. (Answer: Kevin)

For a number of years, models struggled to exceed chance-level performance (Kruengkrai et al., 2014; Sharma et al., 2015; Peng et al., 2015; Liu et al., 2016). The WSC task is carefully controlled, such that heuristics involving syntactic and semantic cues were ineffective, and the common-sense knowledge required to correctly resolve its test instances make it particularly difficult for statistical systems to model. More recently, however, the advent of deep bidirectional transformers (e.g., BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019)) pretrained on massive amounts of data has led to near-human-level performance (Kocijan et al., 2019; Ye et al., 2019; Ruan et al., 2019). Various works have lately re-examined the challenges of the WSC, leading to the proposal of more difficult, larger variants, data and model debiasing methods, and evaluation protocols that clarify which types of instances models excel on and which they struggle with (Trichelair et al., 2019; Emami et al., 2019; Sakaguchi et al., 2020; Abdou et al., 2020). However, little attention has been paid to studying the effects and influence of pretraining data points. While recent work has included some analysis on the effects of 13-gram overlaps between pretraining and test instances for the WSC (Brown et al., 2020), a deeper look into how the *degree* of overlap (and how this can be defined) affects language models’ performance is critical to revealing models’ reasoning

---

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

and inference functions. For example, in 1b), useful knowledge instances in the pretraining corpora may occur as:

- (2) George comforted Melissa because she was very upset. (high overlap)
- (3) And he comforted me because I was so upset by the whole event. (lower overlap)

Studying how models make use of these training instances, in correlation with their overlap with test instances, can provide insight on the roles and downstream influence of exact duplicates (which may be useful if memorized) or highly relevant but distinctly expressed knowledge (useful by retrieval and analogy). In turn, this insight could be used to improve training approaches for models meant to exhibit common-sense reasoning.

**Contributions:** In this work, we address the above issues in CSR modeling by devising a mechanism to score train-test overlap according to a schematization based on BM25, a popular information retrieval function for text matching (Amati, 2009). We use the mechanism to sub-divide test-set instances according to these overlaps. We find that a significant drop in classification accuracy occurs when models are evaluated on the subset with no overlap (we see drops between 3% and 10%, depending on the model, test set, and degree of overlap). Based on this result, we develop the KNOWREF-60K dataset, consisting of 64,301 difficult pronoun disambiguation problems. It is the largest corpus to date for WSC-style common sense reasoning and exhibits a significantly lower proportion of overlaps with current pretraining corpora.<sup>1</sup>

## 2 Related Work

Previous work on the difficulty of instances in the WSC and its variants includes the study by Trichelair et al. (2019), who classified data points into various meaningful subsets. They showed that the success of a then-state-of-the-art LM ensemble (Trinh and Le, 2018) resulted mainly from improvements on simpler “associative” instances. Similarly, experiments by Abdou et al. (2020) show that models are sensitive to linguistic perturbations of Winograd-style examples. New datasets have been proposed to circumvent issues of unintentionally easy test instances, including Winogrande (Sakaguchi et al., 2020), a scaled WSC-variant debiased against RoBERTa, and KnowRef (Emami et al., 2019), which consists of naturally occurring sentences that are free of WSC-specific stylistic quirks.

Given the recent popularity of large, internet-scale datasets for pretraining neural language models, there is an increasing concern that test instances in a downstream task may inadvertently appear in the pretraining corpus. This is a form of data contamination. One of the earliest works that trained a language model on Common Crawl data identified and removed a training documents that overlapped with one of their evaluation datasets (Trinh and Le, 2018). Other work, such as GPT-2 (Radford et al., 2019) and GPT-3 (Brown et al., 2020), conducted post-hoc overlap analysis on CSR benchmarks based on a conservative threshold for contamination—specifically, instance pairs that have a 13-gram overlap. They found the effects of this 13-gram contamination to be negligible. On the other hand, a recent work in computer vision found a significant effect of *near*-duplicates on test performance in an important benchmark, leading to the proposal of a duplicate-free and demonstrably more difficult dataset (Barz and Denzler, 2020). To our knowledge, no work has investigated the effect of varying degrees of overlap between pretraining and CSR test instances for the state-of-the-art transformer-based models (BERT and RoBERTa). Any such investigation must include formulating a more precise definition of contamination.

Methods for purging easy instances from CSR benchmarks have been developed recently: for example, the algorithmic bias reduction of test sets proposed by (Sakaguchi et al., 2020) removes instances from the test set with exploitable annotation artifacts. These techniques depend on pre-computed neural network embeddings of a particular model, and so may be difficult for that model alone but not for previous or up-and-coming models. As the work of Zellers et al. (2018) and the follow-up by Zellers et al. (2019) have shown, adversarial filtering must be iteratively re-adapted to newer models that may be

---

<sup>1</sup>The corpus, the code to scrape the sentences from the source texts, as well as the code to reproduce all experimental results will be publicly available at <https://github.com/aemami1/KnowRef60k>.

immune to previous filtering. This may be costly. Adversarial filtering and related debiasing techniques also do not provide much insight on why certain test instances are filtered out. Our proposed method for data purging is interpretable and model-independent, and can be further supplemented with existing debiasing algorithms like AFLite (Sakaguchi et al., 2020) to ensure that benchmarks remain challenging.

### 3 Hunting for Overlaps

Our procedure for identifying train-test overlaps consists of three main steps: (1) parsing a test instance into its core components, (2) formulating a query using a schema derived from the parse, and (3) quantifying the degree of overlap between a train-test pair using an overlap scoring mechanism.

#### 3.1 Skeletal Representation

We first perform a partial parse of each test instance into a general skeleton of each of the important semantic components, in the order that they appear. We use rules related to the syntactic parse of the sentence implemented by Stanford CoreNLP (Manning et al., 2014).

We use the notation in Emami et al. (2018) to separate the components of WSC-like instances; that is, instances can be divided into a *context* clause, which introduces the two competing antecedents, and a *query* clause, which contains the target pronoun to be resolved:

$E_1, E_2$	the candidate antecedents
$Pred_C$	the context predicate
+	discourse connective
$P$	the target pronoun
$Pred_Q$	the query predicate

$E_1$  and  $E_2$  are noun phrases in the context clause. In the WSC, these two are specified without ambiguity.  $Pred_C$  is the context predicate composed of the verb phrase that relates both antecedents to some event. The context contains  $E_1$ ,  $E_2$ , and the context predicate  $Pred_C$ . The context and the query clauses are often connected by a discourse connective, +. The query contains the target pronoun,  $P$ , which is also specified unambiguously. Preceding or succeeding  $P$  is the query predicate,  $Pred_Q$ , a verb phrase involving the target pronoun. In our case, we will treat  $Pred_C$  and  $Pred_Q$  distinctly, and group all other components ( $E_1$ ,  $E_2$ ,  $P$ , +) together as *content words* in the set  $C$ . Table 1 shows some examples of WSC instances and Table 2 shows sentence pairs in terms of each of these components.

1 a)	The man couldn't lift his son because he was so <u>weak</u> . (Answer: the man)
1 b)	The man couldn't lift his son because he was so <u>heavy</u> . (Answer: son)
2 a)	The older students were bullying the younger ones, so we <u>punished</u> them. (Answer: the older students)
2 b)	The older students were bullying the younger ones, so we <u>rescued</u> them. (Answer: the younger ones)
3 a)	Sam tried to paint a picture of shepherds with sheep, but they ended up looking more like <u>golfers</u> . (Answer: shepherds)
3 b)	Sam tried to paint a picture of shepherds with sheep, but they ended up looking more like <u>dogs</u> . (Answer: sheep)

#### 3.2 Query Schematization

Table 1: Examples of Winograd instances.

We use the above analysis of an instance to formulate a query used to retrieve similar instances in a text corpus. In particular, the per-instance query schema that we propose is:

$$(\text{Phrase}(Pred_c, Pred_Q, 10)) \cap \bigcup_C c_i,$$

where  $\text{Phrase}(Pred_c, Pred_Q, 10)$  denotes that the two predicates must occur in the same order within a distance of 10 tokens to each other, and  $c_i$  are content words in  $C$  that may appear in any order in the sentence.

Predicates		Content Words			
$Pred_C$	$Pred_Q$	$E_1$	$E_2$	$P$	$+$
couldn't lift	was so heavy	the man	his son	he	because
were bullying	punished	the older students	the younger ones	them	so
tried to paint	ended up looking more like	shepherds	sheep	they	but

Table 2: Winograd sentence pairs from Table 1, parsed into the representation schema.

The choice of this schematization stems from the idea that the predicates are the most salient components of WSC-style problem instances. Instances of common-sense knowledge in corpora that support the resolution of a corresponding WSC instance often exhibit only these two components: for example, for 1 a) in Table 1, a possible supporting instance is: *John couldn't lift Melissa and she was so heavy*, although it only shares predicates (underlined) with 1 a). Nevertheless, content words may still contribute informatively and are included as optional components in the query. See Table 3 for the query extracted for the running example.

---

**Sentence:** The man couldn't lift his son because he was so heavy.

---

*Phrase*("couldn't lift", "was so heavy", 10)  $\cap$  (the man  $\cup$  his son  $\cup$  because)

---

Table 3: Query for an example Winograd sentence

### 3.3 Overlap scoring

A retrieval function takes a query related to a given sentence, as formulated above, and estimates its relevance to a given document. In our case, "documents" correspond to individual sentences in the pretraining corpora. One popular retrieval function is BM25 (Amati, 2009), which is a bag-of-words-based function with various components and parameters.

Specifically, given a query  $Q$  containing keywords  $q_1, q_2, \dots, q_n$ , the BM25 score of a document  $D$  is:

$$\text{score}(D, Q) = \sum_{i=1}^n \text{IDF}(q_i) \cdot \frac{f(q_i, D) \cdot (k_1 + 1)}{f(q_i, D) + k_1 \cdot (1 - b + b \cdot \frac{|D|}{\text{avgdl}})},$$

where  $\text{IDF}(q_i) = \log \frac{N - n(q_i) + 0.5}{n(q_i) + 0.5}$ . Here,  $f(q_i, D)$  is  $q_i$ 's term frequency in document  $D$  in words, and  $\text{avgdl}$  is the average document length in the text collection from which documents are drawn. Parameters  $k_1$  and  $b$  are free and usually chosen, in absence of a hyperoptimization, in the range [1.2, 2.0] and as 0.75, respectively.

We use the BM25 score in one of two ways:

1. As a heuristic measure for the degree of overlap or relevance of a sentence in a pretraining corpus, with respect to a given test instance.
2. As a cut-off criterion for sub-dividing a given CSR test set into overlapping and non-overlapping subsets.

We use the Python package *Whoosh* (Chaput, 2017), which has methods to index pretraining corpora, to generate customized queries, and to score these based on the BM25 retrieval function. When the queries are customized in terms of logical operators as in our case, we employ a filter to remove sentences that do not meet the criteria. For example, a document that would otherwise have yielded a high relevancy score to a given query, but whose predicate words do not occur within the 10 token limit, would not be scored at all.

WSC Instance:	The man couldn't lift his son because he was so heavy. Answer: son
Retrieved Sentences & BM25 Scores:	"Nope , our driver had a steel plate in his back and couldn't lift anything ( although he was able to open the truck and put the rather heavy ramp in place , so I am not sure if he was unable or just lazy ) ." → 18.9 "Then I came across a box that weighed a ton - I couldn't even lift it it was so heavy → 26.5 "I man stopped to get it but he couldn't lift it because it was so heavy" → 36.1 "The man couldn't lift his son because he was so heavy" → 43.0 (exact copy)
WSC Instance:	Paul tried to call George on the phone, but he wasn't available. Answer: George
Retrieved Sentences & BM25 Scores:	"A couple of days later we tried to call him at home but his wife told us he wasn't available." → 25.3 "I also tried to call the district attorney, but, unsurprisingly, he wasn't available. → 32.9 "Have a go and check if you're as intelligent as a human: Paul tried to call George on the phone, but he wasn't [successful/available]" → 33.4 (near exact copy) "Paul tried to call George on the phone, but he wasn't available" → 40.3 (exact copy)

Table 4: Example Sentences Retrieved and Corresponing BM25 scores for two WSC problems.

In Table 4, we provide examples of sentences retrieved from the pretraining corpora for a given test instance for various BM25 scores. Qualitatively, a trend appears towards an increased BM25 score with increasing relevance/degree of overlap between a test and pretraining instance. In the case where there was an exact copy found, the score was always significantly higher than when there was not. This suggests that these two steps of query schematization and BM25-based retrieval form an adequate (although by no means perfect) automatic heuristic for ranking the relevance and potential usefulness of pretraining instances.

## 4 Experiments

### 4.1 Existing Benchmarks

**WSC (Levesque et al., 2011)** The original pronoun disambiguation challenge, consisting of 273 problems. Each problem instance is manually crafted by an expert to avoid word association bias, although Trichelair et al. (2019) later report that 13.5% of the questions may still exhibit such bias.

**DPR (Rahman and Ng, 2012)** DPR (Definite Pronoun Resolution Dataset) provides 1,886 additional WSC-style problems authored by undergraduate students. Trichelair et al. (2019) observe that the dataset is likely less challenging than the original WSC because of an increased level of linguistic and dataset-specific biases.

**KnowRef (Emami et al., 2019)** KnowRef introduces over 8k WSC-style coreference resolution problems extracted and filtered using heuristic rules from 100 million web sentences (from Reddit, Wikipedia, and OpenSubtitles).

**Winogrande (Sakaguchi et al., 2020)** Winogrande is a large-scale dataset of 44k WSC-like problems, inspired by the original WSC but adjusted to improve both the scale and the difficulty of the dataset. The key steps of dataset construction are (1) a carefully designed crowdsourcing procedure, followed by (2) systematic bias reduction against a finetuned RoBERTa model by adversarial filtering.

### 4.2 Models

**BERT** BERT (Devlin et al., 2019) is a pretrained neural language model with bidirectional paths and sentence representations in consecutive hidden layers. We finetune BERT by splitting the input sentence into a context and an option component using the candidate answer as delimiter as prescribed in (Devlin et al., 2019). We used grid-search for hyper-parameter tuning: learning rate  $\{1e-5, 3e-5, 5e-5\}$ , number of epochs  $\{3, 4, 5, 8\}$ , batch-size  $\{8, 16\}$  with three different random seeds as in (Sakaguchi et al., 2020). The pretraining corpora are BooksCorpus (800M words) (Zhu et al., 2015) and English Wikipedia (2,500M words). We fine-tune BERT models on DPR-Train for the purpose of comparability

with the state-of-the-art in Kocijan et al. (2019), and we include that corpus as an additional source for querying overlaps.

**RoBERTa** RoBERTa (Liu et al., 2019) is an improved variant of BERT that adds more training data with larger batch sizes and longer training, as well as other refinements like dynamic masking. RoBERTa performs consistently better than BERT across many benchmarks. The pretraining corpora include those used for BERT and three more: CC-News (Nagel, 2016), Openwebtext (Gokaslan and Cohen, 2019), and the Stories Corpus (Trinh and Le, 2018). We fine-tune RoBERTa models on the WNLI-train dataset (Wang et al., 2018) for comparability with the state-of-the-art model in Liu et al. (2019), including the corpus as an additional source for querying potential overlaps.

### 4.3 Results

In the following section, we report the performance of state-of-the-art models on subsets of the CSR test sets for which at least one overlapping instance was retrieved from the pretraining corpora – that is, where the BM25 score between a train-test sentence pair is  $> 0$ . In addition, we investigate how performance changes as we increase the BM25 score cut-off, and use this to assess the relationship between each test set as a whole and the pretraining corpora. Our motivation for these experiments is to gain insight as to the roles that increasingly relevant (and potentially duplicate) pretraining instances play in model performances on CSR benchmarks.

Test Set (size)	WSC Test (273)	DPR-Test (564)	Winogrande-Dev (1,267)
<b>BERT, fine-tuned on DPR-Train</b> (Kocijan et al., 2019)			
Overall Accuracy	71.4%	84.9%	54.6 %
Overlapping Subset Acc., BM25 $> 0$ (subset size)	<b>79.2%</b> (53)	<b>86.9%</b> (246)	<b>59.1%</b> (220)
Non Overlapping Subset Acc., BM25 = 0 (subset size)	69.5% (220)	83.9% (318)	53.7% (1,047)
<b>RoBERTa, fine-tuned on WNLI</b> (Liu et al., 2019)			
Overall Accuracy	91.2%	93.1%	51.9 %
Overlapping Subset Acc., BM25 $> 0$ (subset size)	<b>91.7%</b> (252)	<b>94.1%</b> (306)	<b>55.8%</b> (317)
Non Overlapping Subset Acc., BM25 = 0 (subset size)	85.7% (21)	91.8% (258)	50.6% (950)

Table 5: Performances of models on test set subsets with and without potential overlaps<sup>2</sup> (BM25 score  $> 0$ ). Significant performance differences ( $p < 0.05$ ) according to a chi-squared test are in bold.

Table 5 shows an increase in accuracy between 3% and 10% for both models on subsets of the test data for which overlap scores are greater than 0, suggesting that models tend towards better performance on instances for which there exist similar/overlapping pretraining instances.

Next, we consider subsets of the test set corresponding to more stringent overlapping criteria, that is, for cut-offs of BM25 score  $> 25$  and  $> 35$  (Table 6). We chose these two cut-offs because they corresponded to sharp decreases in numbers of overlapping sentences retrieved<sup>3</sup> (see Figure 1). An increased cut-off score reduces the size of the resulting subset substantially, but in many cases, significantly increased the performance difference. For the original WSC dataset, however, this trend did not hold when the cut-off increased considerably to  $> 35$ ; in fact, the performance difference dropped to negative. This may be explained by the fact that an exact copy of a test instance appearing in the pretraining corpus does not confer knowledge useful at test time, since these copies are by definition ambiguous. Consider retrieving *The man could not lift the boy because he is so heavy*. This does not help resolution in the way that the less similar instance *Tom could not lift Melissa because she was so heavy* does. Indeed, upon further investigation, we found that the WSC contains many exact overlaps with the pretraining corpora (26/29 of its instances with  $\text{BM25} > 35$  corresponded to exact copies.) This suggests that for CSR-based

<sup>2</sup>Choice of fine-tuned models and training set come from the state-of-the-art in Kocijan et al. (2019) and Liu et al. (2019).

<sup>3</sup>While other choices are possible, the results do not vary dramatically within that window.

pronoun disambiguation tasks, current state-of-the-art models tend more towards the retrieval of highly relevant/similar sentences than they do towards memorizing exact duplicates of the test instances.

Test Set (size)	WSC Test (273)	DPR-Test (564)	Winogrande-Dev (1,267)
<b>BERT, fine-tuned on DPR-Train</b> (Kocijan et al., 2019)	71.4%	84.9%	54.6%
Performance Difference (+)			
BM25 > 0 (subset size)	<b>9.7%</b> (53)	2.2% (246)	<b>5.4%</b> (220)
BM25 > 25 (subset size)	<b>8.81%</b> (29)	<b>4.46%</b> (121)	<b>12.6%</b> (92)
BM25 > 35 (subset size)	28.9% (6)	8.6% (15)	5.4% (9)
<b>RoBERTa, fine-tuned on WNLI</b> (Liu et al., 2019)	91.2%	93.1%	51.9%
Performance Difference (+)			
BM25 > 0 (subset size)	<b>6.0%</b> (252)	2.5% (306)	<b>5.2%</b> (317)
BM25 > 25 (subset size)	<b>5.3%</b> (249)	<b>4.1%</b> (186)	0.82% (190)
BM25 > 35 (subset size)	-2.6% (236)	7.3% (32)	4.1% (25)

Table 6: Performance diff. of models between overlapping & non-overlapping test sets

Finally, we graphed the proportion of the original dataset with detected overlaps as a function of the BM25 score cut-off and use this to analyze the overlapping tendencies of each benchmark (Figure 1). Our findings demonstrate that a significant proportion of all datasets have overlaps receiving a BM25 score of at least 20 (ranging between 25% to 68%). We observe a decline as the threshold increases gradually to 40. In the case of most of these test sets, the decline asymptotically approaches 0; notably, however, WSC and KnowRef show no such trend. In the case of the former, the examination mentioned earlier yielded many examples where the test instance was referred to, *verbatim*, in another context, often as a reference to the WSC itself. In the case of KnowRef, the pretraining corpus of Wikipedia was precisely the same corpus used to collect the KnowRef instances, yielding a significant number of highly relevant overlaps. This suggests that various test sets may be subject to leakage/community overfitting or contain test instances for which supporting knowledge may not be as long-tailed as anticipated.

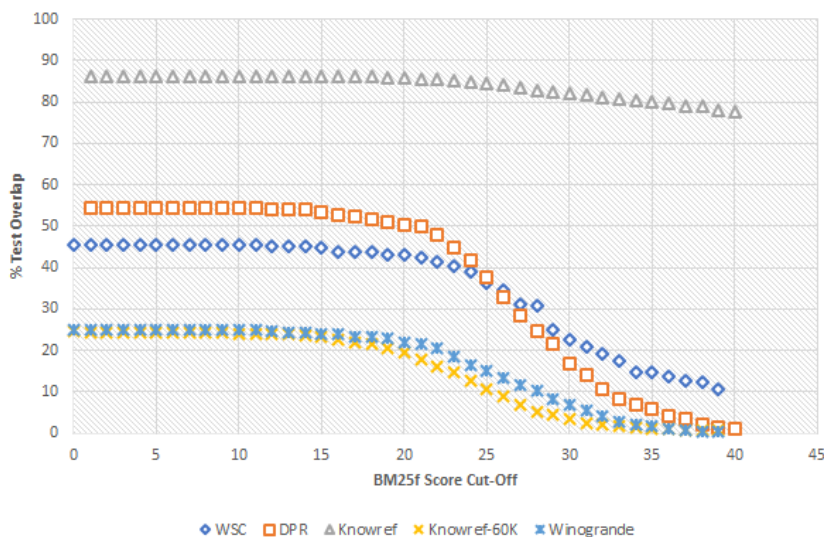


Figure 1: % Test set overlap as a function of BM25 score cut-off

## 5 KNOWREF-60K

### 5.1 Limitations of previous test sets

The results of the previous section suggest that overlapping test instances are, in general, less difficult for models to resolve, and this effect strengthens as the degree of overlap increases. In addition, current

CSR test sets have a large proportion of overlapping instances with the pre-training corpora (BM25 Score > 0), and while this proportion decreases as the degree of overlap increases, certain datasets have a considerable number of highly overlapping instances (BM25 Score > 40). These results suggest two significant limitations of CSR benchmarks.

**Highly Overlapping Instances:** As shown in Figure 1, WSC and KnowRef contain a considerable proportion of instances that overlap significantly with those in the pretraining corpora (either as exact copies or very similar sentences). In the case of KnowRef, almost all instances correspond closely to a sentence in English Wikipedia since this corpus was used specifically for data collection.

**Predictable Structure:** A limitation common to WSC, DPR, and Winogrande is that they are based largely on the structural specifications of the WSC and some of its most famously cited instances (e.g. *The trophy does not fit in the suitcase because it is too large*); that is, instances are often composed only of two clauses connected by a single causal discourse connective, like *because*. For example, during the crowdsourcing protocol for Winogrande, annotators are first primed by classical examples of WSC sentences that may influence their creative process. Test instances that are structurally similar to the original WSC instances seem more likely to overlap with training instances, and it is known that this kind of crowd-sourcing protocol engenders annotation artifacts (Gururangan et al., 2018) that are particularly problematic when they do not corresponding to real-word data (He et al., 2019).

One can find more elaborate, real-word coreference examples that circumvent the above issues. For example, consider this sentence, taken *as is* from Reddit:

- (4) “Forbes wrote that Edison can’t be held accountable because his assistant willingly submitted to the trials and that the dangers of radiation poisoning were not well known.”

Here, despite the instance being a valid pronoun disambiguation problem akin to the WSC instances, there are multiple discourse connectives and more than two clauses, plus distractor content words that contribute variably to the correct resolution. All this renders it much more complex than most, if not all, of the WSC, DPR, and Winogrande instances. The idea that drives our corpus construction process is to identify and collect binary pronoun disambiguation problems as they occur naturally (and potentially, more intricately) in written text, while ensuring that the source of text is not contained in popular pretraining corpora.

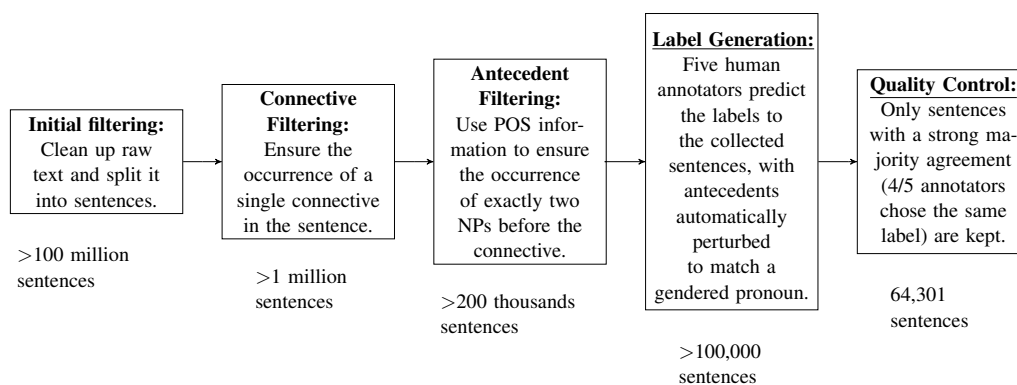


Figure 2: KNOWREF-60K’s construction process. Differences from Emami et al. (2019) are underlined.

## 5.2 Corpus Construction

Motivated by the above limitations, we modify the corpus construction process of Emami et al. (2019) by scraping candidate sentences only from text documents not contained in the pretraining corpora of current models and using human annotators to resolve and label extracted sentences. Specifically, we scrape text samples from Reddit comments dating from 2006–2019. We filter this text through a multi-stage process to ensure quality and diversity as depicted in Figure 2, ultimately yielding 64,301 complex pronoun disambiguation problems scraped from written text. We compile and release these in a coreference task we call KNOWREF-60K. Examples of its instances are shown in Table 7.



Ex. 1:	{Steven} certainly manipulates {Gregory}, but [he] also has the best interest of the world at heart. (Steven)
Ex. 2:	{Gordon} is a better baseball player than {Joseph}, but no one gives [him] any credit. (Gordon)
Ex. 3:	{Vernon} was somewhat insulted and called {Gary} to see if [he] was goofing on him. (Gary)

Table 7: Examples of KNOWREF-60K instances.

### 5.3 Overlap Statistics

We perform experiments to tabulate overlap statistics for KNOWREF-60K and include the results in Figure 1 and Table 8. As seen in Figure 1, KNOWREF-60K’s test set, without applying any explicit debiasing algorithm, yields the smallest proportion of test set overlaps. At the same time, it is nearly three times as large as the largest test set, Winogrande-dev. Table 8 demonstrates also that, according to our retrieval technique, the subset of the test set with overlapping instances yields the lowest increase in performance (just above 2%) for BERT and no change in performance for RoBERTa. Both BERT and RoBERTa achieve their lowest accuracy on the KNOWREF-60K test set (excepting Winogrande-dev, which was specifically debiased against these models to push their performance to chance).

	KNOWREF-60K
<b>BERT, fine-tuned on DPR-Train</b> (Kocijan et al., 2019)	
Overall Accuracy	67.9%
Overlapping Subset Acc., BM25 > 0 (subset size)	69.6% (745)
Non Overlapping Subset Acc., BM25 = 0 (subset size)	67.4% (2,316)
<b>RoBERTa, finetuned on WNLI</b> (Liu et al., 2019)	
Overall Accuracy	76.6%
Overlapping Subset Acc., BM25 > 0 (subset size)	76.1% (770)
Non Overlapping Subset Acc., BM25 = 0 (subset size)	76.8% (2,291)

Table 8: Performances of models on KNOWREF-60K with and without potential overlaps.

## 6 Conclusion

We proposed an automatic method of scoring the degree of overlap between test-train instances in CSR benchmarks and demonstrated that models generally perform better on test instances with high degrees of overlap in pretraining corpora. In response to our findings, we released a more difficult and largest-to-date WSC-style test set called KNOWREF-60K. We ensured that its overlaps with pretraining data are minimal by using a text source not contained in the suite of common pretraining corpora and by basing its construction on naturally-occurring sentences with no direct influence from WSC-like patterns. Our findings suggest that, for better or worse, highly similar pretraining instances have a significant influence on the performance of state-of-the-art transformer-based architectures. Coupled with the large fraction of exact copies or highly-overlapping instances that currently exist in CSR test sets, this effect may bias the evaluation and development of deep learning approaches for common-sense reasoning.

On a positive note, models still performed significantly better than random on non-overlapping test instances, and their relative rankings did not change. This suggests that the community’s efforts do not seem to have overfit to the presence of overlaps yet: transformer-based language models still provide higher accuracy on the subset of non-overlapping instances as well as on our KNOWREF-60K dataset.

We therefore encourage researchers to be cognizant of such overlaps as important factors affecting the performance of CSR-based models, and to use this knowledge to have a clearer picture regarding the true capabilities of machine commonsense across these benchmarks. An important limitation of our analysis is that there is no guarantee that the overlapping subset is drawn from the same distribution as the original dataset, meaning that it is entirely possible that an emergent statistical bias (and neither retrieval nor memorization) caused the overlapping subset to be easier. Accordingly, our work also raises some open questions worthy of pursuit; firstly, how do we more precisely identify the cases of memorization/retrieval? Secondly, what constraints does the retrieval and/or memorization of models as a means of acquiring common sense present in terms of their capabilities/robustness?

## Acknowledgements

This work was supported by the Natural Sciences and Engineering Research Council of Canada and by Microsoft Research. Jackie Chi Kit Cheung is supported by the Canada CIFAR AI Chair program.

## References

- Mostafa Abdou, Vinit Ravishankar, Maria Barrett, Yonatan Belinkov, Desmond Elliott, and Anders Søgaard. 2020. The sensitivity of language models and humans to Winograd schema perturbations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7590–7604, Online, July. Association for Computational Linguistics.
- Giambattista Amati, 2009. *BM25*, pages 257–260. Springer US, Boston, MA.
- Björn Barz and Joachim Denzler. 2020. Do we train on test data? purging cifar of near-duplicates. *Journal of Imaging*, 6(6):41.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Nee-lakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Matt Chapat. 2017. Whoosh: Fast pure-python full text indexing, search and spell checking library.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Ali Emami, Noelia De La Cruz, Adam Trischler, Kaheer Suleman, and Jackie Chi Kit Cheung. 2018. A knowledge hunting framework for common sense reasoning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1949–1958, Brussels, Belgium, October-November. Association for Computational Linguistics.
- Ali Emami, Paul Trichelair, Adam Trischler, Kaheer Suleman, Hannes Schulz, and Jackie Chi Kit Cheung. 2019. The KnowRef coreference corpus: Removing gender and number cues for difficult pronominal anaphora resolution. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3952–3961, Florence, Italy, July. Association for Computational Linguistics.
- Aaron Gokaslan and Vanya Cohen. 2019. Openwebtext corpus. <http://Skylion007.github.io/OpenWebTextCorpus>.
- 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, June. Association for Computational Linguistics.
- He He, Sheng Zha, and Haohan Wang. 2019. Unlearn dataset bias in natural language inference by fitting the residual. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 132–142, Hong Kong, China, November. Association for Computational Linguistics.
- Vid Kocijan, Ana-Maria Cretu, Oana-Maria Camburu, Yordan Yordanov, and Thomas Lukasiewicz. 2019. A surprisingly robust trick for the Winograd schema challenge. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4837–4842, Florence, Italy, July. Association for Computational Linguistics.
- Canasai Kruengkrai, Naoya Inoue, Jun Sugiura, and Kentaro Inui. 2014. An example-based approach to difficult pronoun resolution. In *Proceedings of the 28th Pacific Asia Conference on Language, Information and Computing*, pages 358–367, Phuket, Thailand, December. Department of Linguistics, Chulalongkorn University.
- Hector J Levesque, Ernest Davis, and Leora Morgenstern. 2011. The winograd schema challenge. In *AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning*.
- Quan Liu, Hui Jiang, Andrew Evdokimov, Zhen-Hua Ling, Xiaodan Zhu, Si Wei, and Yu Hu. 2016. Probabilistic reasoning via deep learning: Neural association models. *arXiv preprint arXiv:1603.07704*.

- 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*.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*.
- Sebastian Nagel. 2016. Cc-news.
- Haoruo Peng, Daniel Khashabi, and Dan Roth. 2015. Solving hard coreference problems. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 809–819, Denver, Colorado, May–June. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Altaf Rahman and Vincent Ng. 2012. Resolving complex cases of definite pronouns: The Winograd schema challenge. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 777–789, Jeju Island, Korea, July. Association for Computational Linguistics.
- Yu-Ping Ruan, Xiaodan Zhu, Zhen-Hua Ling, Zhan Shi, Quan Liu, and Si Wei. 2019. Exploring unsupervised pretraining and sentence structure modelling for winograd schema challenge. *arXiv preprint arXiv:1904.09705*.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Winogrande: An adversarial winograd schema challenge at scale. In *AAAI*.
- Arpit Sharma, Nguyen H Vo, Somak Aditya, and Chitta Baral. 2015. Towards addressing the winograd schema challenge—building and using a semantic parser and a knowledge hunting module. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, pages 1319—1325.
- Paul Trichelair, Ali Emami, Adam Trischler, Kaheer Suleman, and Jackie Chi Kit Cheung. 2019. How reasonable are common-sense reasoning tasks: A case-study on the Winograd schema challenge and SWAG. 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 3382–3387, Hong Kong, China, November. Association for Computational Linguistics.
- Trieu H Trinh and Quoc V Le. 2018. A simple method for commonsense reasoning. *arXiv preprint arXiv:1806.02847*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium, November. Association for Computational Linguistics.
- Zhi-Xiu Ye, Qian Chen, Wen Wang, and Zhen-Hua Ling. 2019. Align, mask and select: A simple method for incorporating commonsense knowledge into language representation models. *arXiv preprint arXiv:1908.06725*.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. SWAG: A large-scale adversarial dataset for grounded commonsense inference. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy, July. Association for Computational Linguistics.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pages 19–27.