The Defeat of the Winograd Schema Challenge

Vid Kocijan[a], Ernest Davis[b], Thomas Lukasiewicz[a,c], Gary Marcus[d],
Leora Morgenstern[e]

[a]*University of Oxford, Department of Computer Science, Oxford, OX1 3QD, UK*

[b]*New York University, Department of Computer Science, 251 Mercer St, NY 10012, United States*

[c]*Alan Turing Institute, 96 Euston Rd, London NW1 2DB*

[d]*Robust AI, 380 Portage Avenue Palo Alto, CA 94306 United States*

[e]*Palo Alto Research Center, 3333 Coyote Hill Rd, Palo Alto, CA 94304, United States*

Abstract

The Winograd Schema Challenge—a set of twin sentences involving pronoun reference disambiguation that seem to require the use of commonsense knowledge—was proposed by Hector Levesque in 2011. By 2019, a number of AI systems, based on large pre-trained transformer-based language models and fine-tuned on these kinds of problems, achieved better than 90% accuracy. In this paper, we review the history of the Winograd Schema Challenge and assess its significance.

Keywords: Commonsense Reasoning, Winograd Schema Challenge

1. Introduction

In 2010, Hector Levesque (Levesque, 2011) proposed a new challenge for artificial intelligence: The Winograd Schema Challenge. The challenge was named after a well-known example in Terry Winograd's 1972 groundbreaking doctoral thesis, *Understanding Natural Language*. The example consists of a pair of sentences:

> The city councilmen refused the demonstrators a permit because they feared violence.
> The city councilmen refused the demonstrators a permit because they advocated violence.

1972: Winograd's (1972) thesis introduces the original example.

2010: Levesque (2011) proposes the Winograd Schema Challenge.

2010–2011: The initial corpus of Winograd schemas is created.

2014: Levesque's Research Excellence talk "On our best behavior"
(Levesque, 2014).

2016: The Winograd Schema Challenge is run at IJCAI-16. No systems do
much better than chance (Davis et al., 2017b).

2018: WNLI is incorporated in the GLUE set of benchmarks. BERT-based
systems do no better than most-frequent-class guessing
(Wang et al., 2019b).

2019, May: Kocijan et al. (2019b) achieve 72.5% accuracy on WSC273 using
pretraining.

2019, June: Liu et al. (2019) achieve 89.0% on WNLI.

2019, November: Sakaguchi et al. (2020) achieve 90.1% on WSC273.

Table 1: Time line of the Winograd Schema Challenge

In the first sentence, the pronoun "they" in the first sentence is naturally interpreted as referring to the city councilmen; in the second, it is naturally interpreted as referring to the demonstrators. The only difference between the two sentences is that the first has "feared" where the second has "advocated". Therefore, the different referents chosen for "they" must somehow reflect that different choice of word. Presumably when humans read the first sentence, their choice of referent for "they" is guided by their knowledge that X fearing that Y would be violent would be a reason for X to deny Y a permit to demonstrate, whereas Y fearing violence would rarely be a reason for X to deny Y such a permit. Humans' choice of referent in the second sentence is likewise guided by their knowledge that Y advocating violence would be a reason to deny Y a permit to demonstrate whereas X advocating violence would not be a good reason for X to deny Y a permit to demonstrate. Knowledge about the attitudes of city councils' vs. demonstrators' stereotypical attitudes toward (non-state-sanctioned) violence no doubt also plays a role in the disambiguation. In 1972, and still in 2011, it seemed reasonable to suppose that an AI that would be able to do this kind of disambiguation would similarly have to draw on a body of commonsense knowledge.

Levesque therefore proposed to use these kinds of sentences as a test for the depth of understanding of AI natural language programs. In particu-

lar, he defined a *Winograd schema* as a pair of sentences,[1] comparable to Winograd's example above, with the following features.

1. The two sentences are identical except for one or two words: "feared" vs. "advocated" in Winograd's example.

2. The two sentences both contain two noun phrases and a pronoun: "the city council", "the demonstrators", and "they" in Winograd's example.

3. The natural readings of the two sentences in isolation would assign different choices of referents for the pronoun.

4. Simple feature matching, known as "selectional restrictions", will not suffice to do the disambiguation. For instance, the pair of sentences

> The women stopped taking the pills because they were [pregnant/carcinogenic].

is disallowed, because pills cannot be pregnant and women cannot be carcinogenic.

5. Matching based on simple frequency of co-occurrence will not suffice to do the disambiguation. For instance, the pair of sentences

> The racecar zoomed by the school bus because it was going so [fast/slow].

would be disallowed, because the words "racecar" and "fast" tend to appear together in text.

6. Both sentences must seem natural and must be easily understood by a human listener or reader; ideally, so much so that, coming across the sentence in some context, the reader would not even notice the potential ambiguity.

Point (6) was taken for granted by Levesque (2011) but was made explicit in (Levesque et al., 2012). To what extent the condition has been met – that is, whether readers ordinarily would not notice that the pronoun is ambiguous, has never been tested, as far as we know. (Reading-time experiments have been carried for similar issues in the naturalness and coherence of texts; see, e.g., (Wolf et al., 2004).) However, as discussed further in Section 2.1.1, the constructors of the dataset kept it very much in mind as a design consideration, and in fact a number of examples that otherwise met the criteria

---

[1]In later work, this was often relaxed, so that each element of the pair could be a two-sentence text.

were rejected, because the dataset's authors could not find any way to phrase them that sounded natural.

The naturalness of the task is, indeed, an aspect of the Winograd Schema Challenge that is particularly striking as compared to many other natural language tasks used as benchmarks and makes it particularly compelling as a powerful test of human understanding. In a sentence from a well-designed schema, human readers carry out the inference *automatically*, in a "System 1" (Kahneman, 2011) process in reading the sentence, *before* they read the question that points out the ambiguity; but this inference seems to require commonsense reasoning of some depth and complexity.

By contrast, characterizing the relation between two sentences as "entailment", "neutral", or "contradiction", as in the RTE task (Candela et al., 2006); or identifying the segment of a text that answers a question, as in SQuAD (Rajpurkar et al., 2016); or word analogies (Mikolov et al., 2013); or predicting a probability distribution over the next word following a specified text, which is the fundamental task in "language modeling" (Devlin et al., 2019), are not ecologically valid tasks that people ordinarily carry out explicitly.

As a contrast in a different direction, Cozman and Munhoz (2020) propose encoding brainteasers as Winograd schemas; for example,

> Miss Marple was looking for the jewel, so she asked the girls about it: Ann said she took the jewel, Bella said Donna was the thief, Carol said she did not even see the jewel, and Donna said Ann (did/didn't) in fact take the jewel. When Miss Marple learned that only one of the girls was telling the truth, she immediately knew who had the jewel, and she smiled to *her*.

It is not clear to us what is gained by casting this in terms of pronoun disambiguation, instead of simply asking the question; and it is certainly not what we had in mind.

Levesque's original 2011 paper included eighteen original Winograd schemas. The first of these, "The trophy doesn't fit in the brown suitcase because it's too [small/large]," has become a standard example, cited hundreds of times in articles about AI language understanding.

Once a reasonable-sized collection of Winograd schemas had been assembled, these could be used as a challenge for AI system as follows. One sentence of each pair would be chosen at random; that sentence would be turned into

a binary-choice question. For instance, the "trophy" schema would yield the following question

> The trophy doesn't fit in the brown suitcase because it is too large.
> What is too large?
> A. The trophy
> B. The suitcase

The advantage of using halves of Winograd schemas rather than simply using sentences with pronoun ambiguity is that one can be sure that the only clue to the correct resolution lies in the choice of the alternating word; any clue that does not depend on the alternating word would point in the same direction for both sentences, which would be the wrong direction for one of the pair. For example, one can be entirely sure that the "recency" heuristic has exactly a 50% chance of working, since it will work on one sentence of the pair, and fail on the other and the sentence in the test set is chosen on random.

Thus, a program that guesses randomly will have an expected accuracy of 50%. Human readers should have accuracy close to 100%. Thus, an AI program that achieves an accuracy that is statistically larger than 50% is accomplishing something significant by way of understanding; an AI cannot claim to have achieved human-level abilities unless its acccuracy is close to 100%.

Levesque argued that this test would be a better test of understanding than the Turing test—trivial to evaluate, difficult to game. He also argued that it would be more robust than some alternative semantic tasks that had been proposed, such as the RTE (recognizing textual entailment) challenge (Candela et al., 2006).

2. Datasets

2.1. The initial Winograd schema dataset

Building on Levesque's initial work, in 2010, Ernest Davis created an additional eighty-nine schemas. Together with four suggested by Ray Jackendoff, these were published on the web in HTML and XML.[2] Since then, the

---

[2]https://cs.nyu.edu/faculty/davise/papers/WinogradSchemas/WSCollection.html and

collections have been amplified with new schemas, by Davis and others; in particular, the collection includes twenty schemas written by David Bender. Currently, the HTML file includes 150 schemas; the XML version includes 285 individual sentences. (Some of the HTML schemas do not fit into the stricter form of the XML framework and are thus not included.) The latter is generally known by the acronym WSC285 in recent articles on the subject; since it has become a standard, it is now fixed and will not be modified further. To ensure consistency with earlier models, several authors often prefer to report the performance on the first 273 examples only. This dataset is known as WSC273.

Levesque, Davis, and Leora Morgenstern 2012 published an expanded version of (Levesque, 2011), expanding the comparison to other proposed tests of AI semantic understanding and of the significance of the test in the general context of AI research, as well as presenting the corpus. Levesque (2014) made the challenge the subject of his speech "On Our Best Behavior" accepting the IJCAI Research Excellence Award.

The spatial and physical knowledge used in disambiguating the trophy example—that proposed content A will not fit inside container B if either A is too large or B is too small—is fairly complex from a logical standpoint, particularly the interpretation of "too". (Davis (2013) includes a logic-based analysis.) Other Winograd schemas likewise draw on fairly complex knowledge and reasoning. "The sack of potatoes had been placed [above/below] the bag of flour, so it had to be moved first" draws on a different body of spatial and physical knowledge. "There is a pillar between me and the stage, and I can't see [around it/it]" draws on knowledge of spatial relations and perception. "Bob paid for Charlie's college education, but now Charlie acts as though it never happened. He is very [hurt/ungrateful]" draws on a body of knowledge of human interactions and emotions. "Alice tried frantically to stop her daughter from [chatting/barking] at the party, leaving us to wonder why she was behaving so strangely" requires reasoning about atypical human behavior; for the variant with "chatting", this requires applying the rather complex rule that "It is generally inappropriate for A to try to prevent B from carrying out an innocuous behavior E". All told, the Winograd schemas draw on a wide range of knowledge, often of fairly complex structure.

The challenge attracted a fair amount of favorable interest from both

---

https://cs.nyu.edu/faculty/davise/papers/WinogradSchemas/WSCollection.xml

the research community and the popular science press (e.g., (Markoff, 2015; Knight, 2016)). The naturalness of the problems made it a fair one for AI systems; the complexity of the inferences involved seemed to put it far beyond the range of current technology.

But even at the start and continuing throughout the early years of the challenge, when there were few signs that AI systems could succeed at the Winogrand Schema Challenge, there were those who warned that it might be a less robust test than its creators hoped and believed. They argued that, as was the case with past challenges, success might come without solving the underlying problem that had motivated the challenge, and moreover, that the very nature of the test invited systems that were tailored to the test rather than the underlying problem. In a prescient email to Davis on Feb. 5, 2011, Doug Hofstadter wrote:

> The problem is that I believe that what will happen is that you will simply wind up spawning a whole host of new and ultra-clever brute-force techniques to solve the "Winograd Challenge" without solving the problem of understanding whatsoever. I always liked Terry Winograd's sample sentence, but I hardly think that it represents the epitome or the essence of what is wrong with today's computer approaches to language. It is just one type of example among thousands of types. It's a great example but it's misleading to focus on it as if it were really the crux of the matter. Getting people to spend huge amounts of time on just one kind of challenge is not going to be helpful. In fact, I fear it will be counterproductive, because I don't think that anyone who will be moved to tackle this particular challenge is likely to take up the deeper and more general challenge of what language understanding really is. People are daunted by that, as well they should be, and no one is going to be motivated by a prize to suddenly tackle that gigantic challenge. Instead, very smart engineering types are going to be motivated to seek clever tricks that will allow computers to solve this very narrow type of linguistic disambiguation problem with a high degree of accuracy.

A similar sentiment was expressed by Craig Boutillier at the KR-2012 conference after presentation of the (Levesque et al., 2012) paper. He pointed out that challenges often do get solved sooner or later; and challenged the authors to a thought experiment. If a system did succeed at the Winograd

Schema Challenge, would they be convinced of its intelligence through this alone? In a conversation with Morgenstern in April 2011, Ed Hovy pointed out that this challenge was very likely the sort of problem that a smart graduate student could figure out how to solve, given the right resources. Indeed, the first author of this survey is a graduate student who came up with the first system that performed much better than chance and achieved what was close to human performance, confirming Hovy's prediction. Dan Weld argued, at Marcus's AAAI-2015 workshop Beyond the Turing Test, that the nature of a narrow test was that researchers would figure out how to game the test without making real progress in the field.

Hofstadter, Boutillier, Hovy, and Weld all turned out to be correct. Solving Winograd schemas is not a surrogate for the ability to do commonsense reasoning, let alone for intelligence.

2.1.1. Manually constructing Winograd Schemas

We did not in (Levesque et al., 2012) or subsequent papers, discuss the actual process of constructing Winograd Schemas, because at that time it seemed of no great interest. But the proliferation of adversarial data sets of all kinds in the last few years, and the attempts to produce larger colllections of Winograd schemas using inexpert labor, make it now worth a brief discussion.

The bulk of the Winograd Schemas in the online collection were constructed by Davis over a month in the summer of 2010. Working an hour a day, he found that he could reliably generate three schemas an hour — rarely fewer, almost never more. In fall of 2015, preparing for the Winograd Schema competition at AAAI-16, he produced another collection of 60, which has been held in reserve and has never been published. This was carried out less systematically, and was somewhat slower, but not enormously so.

The primary difficulty in constructing Winograd schemas is that English usage imposes a variety of formal constraints on how pronouns can be used, beyond purely syntactic rules. These constraints are complex and not fully understood (Rohde, 2018), but can be quite rigid, nonetheless.

For instance, both halves of Levesque's canonical trophy example are fine:

> The trophy doesn't fit inside the suitcase because it is too large.
> The trophy doesn't fit inside the suitcase because it is too small.

Commonsense knowledge about fitting and size require that, in the first sentence, "it" refers to the trophy, and in the second "it" refers to the suitcase.

But you can hardly say

* The trophy doesn't fit in the suitcase because the trophy is an awkward shape and it is too small.

with the intent that "it" should refer to the suitcase.

Likewise, you can reasonably say in English,

Ann has no children, but Barbara has two sons, Carl and David. Her children are in elementary school.

However, you cannot say

* Barbara has two sons, Carl and David, but Ann has no children. Her children are in elementary school.

in the expectation that semantic constraints will force the interpretation of "Her" as "Barbara's". The natural reading of this sentence is to interpret "Her" as meaning "Ann's" and then try to work out the contradiction.

Kehler (2015) similarly gives a number of examples where the form of the sentence either forces a pragmatically implausible reading or conflicts with pragmatics in a way that makes the sentence unintelligible:

(a) The demonstrators were denied a permit by the city council because they feared violence.
(b) The city council denied the demonstrators a permit because they felt strongly that the best way to draw attention to current political issues is to advocate violence.
(c) Norm lent his car to his brother's girlfriend. He doesn't own one.
(d) Margaret Thatcher admires Hillary Clinton, and George W. Bush absolutely worships her.

In (a), readers generally interpret "they" as referring to the demonstrators. In (b), readers "garden path": they initially interpret "they" as meaning "the city council" and then consciously backtrack and correct when they reach the end of the sentence. In (c), readers end up confused rather than deciding that "he" refers to the brother. In (d), readers interpret "her" as meaning Clinton rather than Thatcher, despite the fact that, when the example was devised, Thatcher was much more plausible.

Since the two sentences in a Winograd schema can differ only in a word or two, and both sentences have to be very easily understood by a human reader,

9

and both have to sound natural in their use of pronouns, and neither can be resolvable by selectional restrictions or considerations of frequency, the problem of generating Winograd schemas is substantially overconstrained. The problem of generating a collection of Winograd schemas that span a range of commonsense knowledge and linguistic contructions is still harder. Often it happened that a promising first sentence could not be completed to a full schema. For instance, the text "Tom's books are full of mistakes. Some of them are quite [foolish/worthless]." seemed promising; but "foolish mistake" is a much more common association than "foolish book", and we couldn't find a way to fix this.

Another example: The following sentence is from a published article (Joukovsky and Powell, 2011):

> Late in 2009, the novelist Jim Powell found a cache of letter written by the 19th century novelist George Meredith to his great-great-grandmother Susan Mary Neil.

The pronoun "his" here must refer to Powell, not Meredith; very few people are in a position to write letters to their own great-great-grandmother. It could be turned into a Winograd schema by changing "great-great-grand-mother" to "friend"; in that version, "his" must refer to Meredith, since it would be unusual (though not quite impossible) for a nineteenth-century novelist to write letters to a friend of someone alive in 2009. However, the problem with this example is that the first version is too easily misread. In fact, Davis, reading the article, *did* misinterpret it until he got further in the article and realized that interpreting "his" as "Meredith's" was not only wrong but impossible.

There is a web page[3] with a few more examples of attempts at Winograd schemas that could not be brought to fruition.

Manually creating a large, diverse collection of high-quality Winograd schemas is inherently difficult. As a result, the collection was small; initially a little more than a hundred, and later less than three hundred. This led to some complaints; for instance, Trichelair et al. (2018) wrote that "the main drawback of the Winograd Schema Challenge is its limited size and the absence of training and validation sets for (hyper)parameter tuning. For its size, it turns out that if one were to choose from a set of 10 random binary

---

[3]https://cs.nyu.edu/faculty/davise/papers/WinogradSchemas/WSfailed.html

classifiers, the best based on its performance on the WSC, there is more than a 1-in-3 chance of scoring above 55% accuracy with this chosen classifier." However, it was not our purpose to construct a training set, because there is no point in training an AI program to solve Winograd schemas specifically. The point of the Winograd Schema Challenge is to test programs that claim to have solved the problem of pronoun reference resolution; and for that purpose, what would be relevant would be a training set for pronoun reference resolution generally. As for testing: a sample of 100 would be small for polling voters in an election, where you generally expect percentages between 40% and 60% and you would like to get precision of around 2% or 3%. But for distinguishing between success rates for humans on the Winograd schemas, which are around 95% and those of 2011-level AIs, which were little better than 50%, a sample of 100 is more than adequate. If an AI scores 50% on a sample of 100, then one can reject the hypothesis that its actual accuracy is 95% with a confidence of around $1-10^{-36}$. (Linzen (2020) similarly argues for the value of expert-constructed, comparatively small, "test-only" datasets.)

2.1.2. Subclasses of the original collection

Trichelair et al. (2018) have observed that 37 sentences in the Wsc273 dataset (13.6%) can be easily solved using statistics over simple patterns; that is, they fail condition (5) above. An example of such a sentence is

*In the storm, the tree fell down and crashed through the roof of my house. Now, I have to get **it** [repaired/removed].*

The roof is commonly associated with *being repaired*, while *the tree* is not. They call these examples *associative* and name the rest *non-associative.*

Such patterns can involve multi-word phrases. For instance, the example

*Joe has sold his house and bought a new one a few miles away. He will be moving out of **it** on Thursday.*

can be solved using the fact that *moving out of his old house* is a much more common phrase than *moving out of his new house.* (As of May 13, 2021, a Google search claimed that there were 481,000 results for the first quoted phrase and 2 for the second.) Trichelair et al. (2018) used human annotators to identify the associative sentences; their identifications were largely replicated by Elazar et al. (2021), using an automated measure.

Additionally, 131 sentences (48% of Wsc273) were found to form meaningful examples if the candidates in the sentence are switched. An example of such sentence is

Bob collapsed on the sidewalk. Soon he saw Carl coming to help. **He** *was very [ill/concerned].*

In this sentence, *Bob* and *Carl* can be switched to obtain an equivalent example with the opposite answers. Such sentences were named *switchable*. Trichelair et al. (2018) encourage future researchers to additionally report the consistency on the *switchable* dataset, when the candidates are switched, and when they are not.

The list of associative and switchable examples together with their switched counterparts have been made public[4].

2.1.3. Winograd Schema Challenge in other languages.

While the inspiration and original design of the challenge was in English, translations into other languages exist. (As far as we have found, the WSC is the only commonsense benchmark or dataset that exists in any language other than English.) Amsili and Seminck (2017) translated the collection of 144 Winograd schemas into French, and 285 original Winograd schemas were translated into Portuguese by Melo et al. (2020). Authors of French and Portuguese translation both report having to make some changes to the content to avoid unintended cues, such as grammatical gender. In the case of Portuguese, 8 sentences had to be dropped, as no appropriate translation could be found.

Translations to Japanese[5] and Chinese[6] are available on the official web page of the challenge. A second translation into Chinese, "Mandarinograd",[7] is reported in (Bernard and Han, 2020) together with an account of the difficulties involved in the translation process.

The Chinese Winograd Schema Collection (CLUEWSC2020) is an anaphora/coreference resolution task where the model is asked to decide whether a pronoun and a noun (phrase) in a sentence co-refer (binary classification). The collection contains 1838 questions hand-selected from thirty-six contemporary literary works in Chinese. The anaphora relations have been hand-annotated by linguists. The dataset is part of CLUE (Chinese Language Understanding Evaluation), a collection of Chinese language benchmarks analogous to GLUE (Xu et al., 2020)

---

[4]`https://github.com/ptrichel/How-Reasonable-are-Common-Sense-Reasoning-Tasks`
[5]`http://arakilab.media.eng.hokudai.ac.jp/~kabura/collection_katakana.html`
[6]`https://cs.nyu.edu/faculty/davise/papers/WinogradSchemas/WSChinese.html`
[7]`https://gitlab.com/vanTot/mandarinograd/`

The Indic General Language Understanding Evaluation (IndicGLUE) (Kakwani et al., 2020) includes translations of GLUE's WNLI into Hindi, Marathi, and Gujarathi. Slovene translation of SuperGLUE benchmark includes the translation of WNLI into Slovene (Žagar et al., 2020).

Vered Shwartz has translated the WSC collection into Hebrew.[8]

2.2. Definite Pronoun Resolution Dataset

The Definite Pronoun Resolution (DPR) dataset is an easier variation of the Winograd Schema Challenge (Rahman and Ng, 2012). The constraints on the Winograd schemas have been relaxed, and several examples in the dataset are not *Google-proof*. The dataset consists of 1322 training examples and 564 test examples, constructed manually. 6 examples in the training set reappear in WSC273 in a very similar form. These should be removed when training on DPR and evaluating on WSC273. This dataset is also referred to as WSCR, as named by Opitz and Frank (2018).

An expanded version of this dataset, called WINOCOREF, has been released by Peng et al. (2015), who further annotate all previously ignored mentions (in their work, a mention can be either a pronoun or an entity) in the sentences that were not annotated in the original work. In this way, they add 746 mentions to the dataset, 709 of which are pronouns. Moreover, Peng et al. (2015) argue that accuracy is not an appropriate metric of performance on the WINOCOREF dataset and introduce a new one, called *AntePre*.

They define *AntePre* as follows: Suppose there are k pronouns in the dataset, and each pronoun has $n_1, \ldots, n_k$ antecedents. We can treat finding the correct candidate for each pronoun as a binary classification for each antecedent-pronoun pair. Let m be the number of correct binary decisions. *AntePre* is then computed as $m \big/ \sum_{i=1}^{k} n_i$.

2.3. Pronoun Disambiguation Problem Dataset

The Pronoun Disambiguation Problem (PDP) dataset consists of 122 problems of pronoun disambiguation collected from classic and popular literature, newspapers, and magazines. As discussed by Morgenstern et al. (2016), PDPs were originally conceived as an auxiliary dataset to Winograd schemas

---

[8]https://vered1986.github.io/resources/winograd_he.html
and https://vered1986.github.io/resources/winograd_he.jsonl

primarily because constructing Winograd schemas according to Levesque's original guidelines was a difficult, manual process. Indeed, as of this date, the plurality of such Winograd schemas has been constructed by Davis. Moreover, in the decade since the Winograd Schema Challenge was first published, no one has successfully developed a methodology for constructing large numbers of such schemas. There have been efforts to construct large sets of Winograd schemas, such as WinoGrande (Sakaguchi et al., 2020), but as discussed in section 2.7, the schemas are often flawed and do not meet Levesque's guidelines.

The chief distinction between PDPs and Winograd schemas is the lifting of Levesque's criterion of the "special word" or phrase that when substituted creates a pair of similar sentences with different pronoun referents. Winograd schemas are rarely found "in nature," that is, in novels, short stories, newspaper articles, or other pieces of existing text. In contrast, PDPs are ubiquitous in text. Not every instance of a pronoun in text, of course, is a suitable PDP. For example, the sentence starting Chapter 4 in Maud Hart Lovelace's *Carney's House Party* "During the long trip Carney began to think about her reunion with Bonnie," has the pronoun *her*, but there is no challenge in resolving that pronoun to *C*arney. A suitable PDP should ideally present a slight challenge to a human reader, a challenge that is easily resolvable using commonsense knowledge, but a challenge nonetheless. An example, from Alcott's *L*ittle Women is:

> Mrs. March gave the mother tea and gruel, while she dressed the
> little baby as tenderly as if it had been her own

There are two possible referents for *she*: *Mrs. March* and *the mother*. It is not difficult to reason that *she* must refer to Mrs. March rather than the mother, since it is clear from the remainder of the sentence, especially the phrase "as if it had been her own" that the little baby does not in fact belong to the referent of *she*. Humans are likely to do this reasoning using, among other things, the commonsense knowledge that a baby is understood to belong to its mother.

The aim in collecting and vetting PDPs has been to come up with a set consisting of challenging problems that tests ability to do something akin to commonsense reasoning. They were intended to be used as a gateway set before administration of the Winograd Schema Challenge and in fact were used in this way in the IJCAI-2016 running of the Challenge Davis et al. (2017b). Our aim in construction and vetting was to create a test which was roughly

equivalent in difficulty to a WSC test corpus, though of course, the absence of pairs of sentences distinguished only by a single word or phrase means that one cannot guarantee that the pronoun reference is resolved through commonsense knowledge and reasoning ability, rather than sentence structure.

PDPs are clearly easier to collect and to vet than Winograd schemas are to construct. In addition, there are other reasons that appropriate sets of PDPs are of interest. First, by construction, Winograd schemas are generally quite simple — in general, the more complex the sentence, the more difficult it is to create a sentence identical in all but a word or short phrase — and it may be of interest to test a system's ability to resolve pronoun referents in contexts that are more complex.

Second, the reasoning and the knowledge involved in many PDPs has tended to be deeper and involve more types of commonsense knowledge than that needed for Winograd schemas. An example of a PDP of this sort is taken from *All-of-a-Kind Family*:

> Mama came over and sat down beside Sarah. Gently she stroked her hair and let the child weep.

This example involves relatively profound concepts of parenting and empathy and comforting and soothing that are not, at least at this point, present in existing Winograd schema collections, perhaps partly because it is not straightforward to find a sentence that is nearly identical to the one above, but for which the referent for *she* would be Sarah rather than Mama. Such a range of difficulty and complexity of reasoning in Winograd schemas was a desideratum in Levesque's original paper.

Despite the fact that collecting and vetting PDPs is a much simpler task than constructing Winograd schemas, it is not a trivial task. We focused on children's books partly because they tended to contain a wide range of commonsense reasoning concepts of interest. However, many possible examples did not stand up to further examination.

For example, consider the following from *The Wizard of Oz*:

> The Scarecrow seized the oilcan from Dorothy's basket and oiled the Woodman's jaws, so that after a few moments he could talk as well as before.

There are many interesting commonsense concepts which children are likely to use when resolve the referent of *he*, including the importance of jaw

movement in speaking. However, this example fails one of the tests shared by Winograd schemas and PDPs, that they be Google proof. "Jaws" and "talk" have a higher correlation – of several orders of magnitude – than "seize" or "oil" and talk, making it much more likely that the referent of *he* is the Woodman than the Scarecrow.

Moreover, many examples simply were trivial or did not involve interesting or complex commonsense concepts.

62 examples of PDPs were published before the Winograd Schema Challenge was administered,[9] and 60 PDPs were included in the Winograd Schema Challenge that was administered at IJCAI 2016 (Davis et al., 2017b).[10] A corpus of 400 sentences was collected semi-automatically from online text, with less vetting by Davis and Pan (2015).[11] In addition, in developing this set there was little emphasis on collecting examples that, for pronoun resolution by humans, would involve complex concepts of commonsense reasoning.

2.4. Winograd Natural Language Inference Dataset

The Winograd Natural Language Inference (WNLI) dataset is part of the GLUE benchmark (Wang et al., 2019b) and is a textual entailment variation of the Winograd Schema Challenge. An example from WNLI is given below with the goal to determine whether the hypothesis follows from the premise.

Premise: *The city councilmen refused the demonstrators a permit because they feared violence.*
Hypothesis: *The demonstrators feared violence.*
Answer: true / **false**

The dataset consists of 634 training examples, 70 validation examples, and 145 test examples. Training and validation sets contain a major overlap with the WSC273 dataset, while test samples come from a previously unreleased collection of Winograd schemas. Not all examples in this dataset contain the *special* word and therefore do not come in pairs. Kocijan et al. (2019b) note that examples are much easier to approach if the Winograd schemas are transformed from the textual entailment back into the pronoun resolution problem, and approached as such.

---

[9]`http://commonsensereasoning.org/disambiguation.html`
[10]`https://cs.nyu.edu/faculty/davise/papers/PDPChallenge.xml`
[11]https://cs.nyu.edu/faculty/davise/annotate/corpus.xml

The same collection of examples is used for the SuperGLUE benchmark (Wang et al., 2019a) as a pronoun resolution problem to begin with. For the purpose of this survey paper, Wnli and SuperGlue Wsc are considered the same dataset. They consist of the same examples and all approaches to Wnli described in this paper transform the examples as noted in the previous paragraph.

2.5. WinoGender Dataset

Unlike the previous datasets, WinoGender was created as a diagnostic dataset and is aimed to measure gender bias of the systems for pronoun resolution (Rudinger et al., 2018). WinoGender consists of 120 hand-written sentence templates, together with candidates and pronouns that can be inserted into the templates to create valid sentences.

In each sentence, one of the candidates is an occupation, usually one with a high imbalance in gender ratio (e.g., surgeon). The other candidate is a participant (e.g., patient) or a neutral *someone*. For each sentence, either of the pronouns *he, she*, or *they* can be included to create a valid sentence, as the candidates are gender-neutral. All together, this gives 720 Winograd schemas. An example from the dataset is

*The surgeon operated on the child with great care; [**his/her**] [tumor/affection] had grown over time.*

Note that the gender of the pronoun does not affect the expected answer; however, a biased system that associates the pronoun *his* with the surgeon is likely to answer one of them incorrectly. The aim of this dataset is not to measure model performance, as its data distribution is highly skewed, but to help analyse the models for gender bias.

A Swedish-language version of WinoGender has also been developed (Hansson et al., 2021).

2.6. WinoBias Dataset

WinoBias was created by Zhao et al. (2018), which tries to identify gender bias in pronoun resolution models. WinoBias and WinoGender were created concurrently but independently, despite the same objective. They introduce a dataset with $3,160$ sentences, split equally into development and test. Each sentence contains two candidates that are selected from a list of jobs with highly imbalanced gender ratio.

Two different templates are used to create Winograd schemas. Type 1 sentences follow a structure that does not give away any syntactic cues. The

authors thus estimate these sentences to be more challenging. An example of such a sentence is

*The farmer knows the editor because [**he/she**] [is really famous/likes the book].*

Type 2 sentences can be answered based on the structure of the sentence. The authors thus expect the models to perform better. An example of such a sentence is

*The accountant met the janitor and wished [**her/him**] well.*

Just like in example of type 1, the referent is not affected by the grammatical gender of the pronoun. Its "twin pair" has the candidates swapped. As the structure of the sentence gives the answer away, there is no *special word*.

Moreover, the authors evenly split the whole dataset into *pro-stereotypical* and *anti-stereotypical*, depending on whether the gender of the pronoun matches the most common gender of the referent occupation or not. They observe that publicly available models for co-reference resolution exhibit a major difference (up to 21.1% F_1) in performance on *pro-* and *anti-* subsets of the dataset.

2.7. WinoGrande Dataset

In contrast to the sets of Winograd schemas discussed above, created either entirely by hand or using templates and a fixed set of rules for instantiating the templates, and consisting of anywhere from a few hundred to a few thousand examples, (Sakaguchi et al., 2020) presented a substantially larger (∼44K) set of coreference resolution problem sentences, along with methods to filter these sentences to remove bias. As discussed in (Morgenstern, 2021), Sakaguchi et al. were primarily motivated by the fact that systems already existed that achieved near-human performance on the Winograd Schema challenge, yet none had achieved the ability to do commonsense reasoning, a central aim of posing the WSC. They therefore sought to re-engineer the WSC as a more meaningful benchmark of a system's CSR ability, by introducing new methods of dataset development and adversarial filtering, expressly designed to prevent AI systems from making claims of smashing through benchmarks without making real progress.

The WinoGrande corpus was collected using the Amazon Mechanical Turk (MTurk) crowdsourcing marketplace. To prevent crowdworkers from creating lexically and stylistically repetitive examples, they were primed by a randomly chosen topic from a WikiHow article as a suggestive context.

Additionally, they were instructed to adhere to constraints close to the original WSC constraints, e.g., avoiding writing examples where word association would lead to inferring the correct pronoun referent. Workers were told to focus on the domains of social commonsense and physical commonsense. An additonal set of Mechanical Turkers checked all sentences generated to increase the likelihood that humans could easily infer pronoun referents but that making such inferences required at least some commonsense knowledge.

The authors additionally introduced the AFLITE adversarial filtering algorithm.The idea was to retain only examples that minimize representation bias. Removed pairs include those with data-set specific polarity basis (e.g., advanced rock climbing is more strongly associated with being *strong* than being *weak*). They used a fine-tuned RoBERTa language model (Liu et al., 2019) to gain contextualized embeddings for each instance. Using these embeddings, they iteratively trained an ensemble of linear classifiers, trained on random subsets of the data and discard top-k instances that were correctly resolved by more than 75% of the classifiers. By iteratively applying this algorithm, the authors identified a subset $(12, 282$ instances), called WinoGrande$_{\text{debiased}}$. Finally, they split this dataset into training $(9, 248)$, development $(1, 267)$, and test $(1, 767)$ sets. In many cases, the AFLITe filtering removes only one sentence of a pair. Slightly more than $1/2$ of the questions in WinoGrande$_{\text{debiased}}$ are not twins. [12] They also released the unfiltered training set WinoGrande$_{\text{all}}$ with $40, 938$ examples;(Sakaguchi et al., 2020) these are all twins.

Despite the promise of this approach, we note first, that many of the sentences in the filtered WINOGRANDE corpus do not fit the criteria laid out in the WSC, and therefore are not actually Winograd schemas. Many sentences are easily solvable using word correlation; that is, they are not Google-proof. For example (all examples below are from the debiased, filtered set released in September 2020), consider: *The doctor diagnosed Justin with bipolar and Robert with anxiety. Justin/Robert had terrible nerves recently.* A Google search shows that *anxiety* and *nerves* are correlated at about twice the rate as *bipolar* and *nerves*. Some sentences appear to be poorly written or to contain the answer directly in the sentence, e.g. *The waiter could not cover the round tables with the square tablecloths because the tables/tablecloths were*

---

[12](Sakaguchi et al., 2021) state that 1/3 are not twins, but our own count found that the true figure is 56%.

square. Others are genuinely hard to understand, e.g., *George opted for both of them to use a knife instead of a gun in the duel because the knife/gun could partially injure them.*

We note, second, that the goal of postponing human-level performance on this dataset until systems achieve commonsense reasoning ability was achieved only momentarily. The WinoGrande leaderboard shows that the UNICORN system (Lourie et al., 2021a) achieves 91.2% accuracy, despite little progress in commonsense reasoning.

2.8. WinoFlexi Dataset

Similarly to WinoGrande, Isaak and Michael (2019) aimed to construct a dataset through crowdsourcing. They built their own crowdsourcing interface for generating and evaluating schemas, and they collected 165 candidate Winograd schemas. In our manual review of this, we found that there were 84 distinct valid schemas (168 examples). Unlike workers on WinoGrande, workers on WinoFlexi were not presented with any particular topic and were free to pick it on their own; and they were not instructed to avoid examples that could be solved using word associations.

2.9. Winventor Dataset

The Winventor program (Isaak and Michael, 2020) uses automated methods to generate candidates for pronoun disambiguation problems, which could then be polished or adapted by human post-editors. Using hand-coded techniques, the program collected Wikipedia sentences with a pronoun with two potential referents and formulated a question based on the resolution of that pronoun. (An alternative, less successful, program tried to use a deep learning architecture for the same task.)

For example, starting with the Wikipedia sentence, "As Frederick was rather distant to his family, Eleanor had a great influence on the raising and education of Frederick's children, and she therefore played an important role in the House of Hapsburg's rise to prominence," Winventor proposed the question "Who therefore played an important role in the House of Hapsburg's rise to prominence?" Since, as it stands, "she" can be disambiguated by gender constraints, a human editor changed "Eleanor" to "John" and "she" to "he". The authors compared the performance of human experts in generating PDPs when aided by Winventor as compared to working by themselves; they found that, using Winventor the editors were able to generate examples

1.5 times as fast as without it (10 PDPs in 20 minutes as opposed to 7), and that their examples were much more varied in form.

2.10. Datasets of Explanations

WINOWHY (Zhang et al., 2020) and WINOLOGIC (He et al., 2021) are collections of correct and incorrect natural language explanations of all Winograd Schemas in the WSC273 dataset. The goal of both datasets is to determine whether systems that can correctly answer Winograd Schemas are also capable to identify the correct explanations for their choice. The answer to this question at the time of the writing seems to be negative. The difference between the datasets is in their construction. WinoWhy was constructed through crowdsourcing with no guarantee that the explanations are exhaustive, that is, a correct explanation is not guaranteed to contain all the necessary steps to get the answer. WinoLogic, on the other hand, consists of explanations that were obtained by manually deriving the answers to examples in first-order logic first and then rewriting the derivation into natural language explanations.

2.11. Winograd Schemas for Machine Translation

Machine translation of schemas to languages with a strong presence of grammatical gender can require their implicit resolution. Emelin and Sennrich (2021), e.g., provide an example of the following translation from English to German:

> The gardener used the shovel more than the rake because *it* was poorly made.
>
> Der Gärtner benutzte die Schaufel mehr als den Rechen, weil [*er*/sie] schleht gemacht war.

Die Schaufel and *der Rechen* have feminine and masculine grammatical gender, respectively, and to correctly translate the pronoun *it* into *sie* or *er*, the system must resolve the pronoun first. Alternatively, Stanovsky et al. (2019) note that translating the sentence

> The doctor asked the nurse to help *her* in the procedure.

requires implicit resolution of the pronoun to determine whether to use the feminine or masculine form of *the doctor* (in German, *der Arzt/die Ärztin*).

Stanovsky et al. (2019) create a WINOMT dataset to detect gender bias in machine translation systems by observing how sentences in WinoGender and WinoBias are translated to Spanish, French, Italian, Russian, Ukrainian, Hebrew, Arabic, and German. Kocmi et al. (2020) extend this to Polish and Czech. They find that systems for machine translation regularly assign grammatical gender based on gender stereotypes. Emelin and Sennrich (2021), on the other hand, focus on the subset of WinoGrande that talks about inanimate objects that can be referred to with the pronoun *it*. They create a Wino-X, a diagnostic dataset for machine translation from English to French, German, and Russian. They, too, find that existing machine translation systems do not resolve Winograd Schemas well. Curiously, they find that systems are more likely to (incorrectly) translate an entity into a masculine than feminine form, indicating that gender bias and possible imbalance in the training data affects even the treatment of inanimate objects with feminine grammatical gender.

2.12. An earlier dataset for psycholinguistic experimentation

In psycholinguistic studies of pronoun disambiguation, Kehler et al. (2008) report using experimental material that includes a set of 16 pairs of Winograd schemas following a fixed framework:

> Samuel threatened Justin with a knife, and Erin [blindfolded/stopped] him. Whom did Erin [blindfold/stop]?

> Samuel threatened Justin with a knife, and he [blindfolded Erin/alerted security]. Who [blindfolded Erin/alerted security]?

2.13. Evaluation with Human Subjects

Human subject evaluation has been carried out for a number of data sets (see Table 2). The study reported in (Kehler et al., 2008) were carried out by undergraduates at UCSD, who were self-reported monolingual English speakers. The studies reported in (Davis et al., 2016) and the unpublished studies carried out by Davis were carried out in person by paid volunteers from the NYU community (in the latter case, the volunteers were screened for being native speakers of English). In all the other studies, the subjects were crowd workers recruited on Amazon Mechanical Turk or similar platforms. The details of the methodology used are given in the references. The anomalously low result of 86.5% accuracy on WSC273 is mentioned in passing in a footnote in (Rudinger et al., 2018); the method used is not described.

Dataset	% correct	Reference
Experimental dataset of 32 schemas	93.0%	(Kehler et al., 2008)
143 schemas from WSC273	92.1%	(Bender, 2015)
66 texts with 108 PDPs	90.1%	(Davis et al., 2016)
89 unpublished Winograd schemas	92%	(Davis et al., 2016)
Unpublished collection of 86 texts with 101 PDPs	96%	Unpublished test, (Ernest Davis)
WNLI	96.1%	(Nangia and Bowman, 2019)
WINOGENDER	94.9%	(Rudinger et al., 2018)
WSC273	86.5%	(Rudinger et al., 2018)
WINOGRANDE	94%	(Sakaguchi et al., 2020)

Table 2: Human subject evaluation of datasets

3. Methods

At least three different methods have been used to try to solve the Winograd Schema Challenge. One class of approaches consists of feature-based approaches, typically extracting information such as semantic relations. Additional *commonsense knowledge* is usually included in the form of explicitly written rules from knowledge bases, web searches, or word co-occurrences. The collected information is then used to make a decision, using rule-based systems, various types of logics, or discrete optimization algorithms. We observe that the extraction of relevant information from the sentence is usually the bottleneck of these approaches. Given the nature of the challenge, even the slightest noise in the feature collection can make the problem unsolvable.

The second group of approaches are neural approaches, excluding language-model-based approaches, which we consider as a separate group. Neural-network-based approaches usually read the sentence as a whole, removing the bottleneck of information extraction. To incorporate background information, these networks or their components are usually pre-trained on unstructured data, usually unstructured text, or other datasets for coreference resolution. Common approaches to the tasks in this group take advantage of semantic similarities between word embeddings or use recurrent neural networks to encode the local context. We find this group of approaches to lack reasoning capabilities, as semantic similarity or local context usually do not contain sufficient information to solve Winograd schemas.

The third group includes approaches that make use of large-scale pre-trained language models, trained with deep neural networks, extensively pre-trained on large corpora of text. Some of the approaches then additionally fine-tune the model on Winograd-Schema-Challenge-style data to maximize their performance. Approaches in this group achieve visibly better performance than approaches from the first two groups.

3.1. History of pronoun disambiguation in AI and linguistics

The AI problem of building natural language understanding systems capable of determining the referent of pronouns in text,[13] and the related linguistics problem of characterizing the rules that govern pronoun reference have been studied extensively since the 1970s.

Pronouns in English are generally, though not always, anaphoric—that is, a pronoun generally refers to some entity mentioned by a noun phrase earlier in the text—and most of the AI research on pronoun use has focused on finding anaphoric references in the text, since those are more easily characterized than non-anaphoric references. The problem is thus a subset of the broader problem of anaphoric resolution. Indeed, in general, pronoun resolution is a particularly easy category of anaphoric resolution for computational systems; the problem of determining whether two noun phrases in text refer to the same entity is often much more difficult (Claire Cardie, personal communication). (Extensive surveys of anaphoric reference are given in (Poesio et al., 2011), (Poesio et al., 2016), and (Kehler et al., 2008).)

Early AI work in pronoun resolution mostly focused on the role of commonsense knowledge in disambiguation (e.g., (Charniak, 1972; Wilks, 1975; Schank and Abelson, 1977; Hobbs, 1979; Hobbs et al., 1993)) or on structural characteristics of the text, such as parallelism and focus (Grosz, 1977; Sidner, 1979; Kameyama, 1986). Implementations were almost always proof-of-concept rather than practical systems. Development was guided by the researchers' intuitions and interests, rather than by evaluation. Systematic evaluation of systems was extremely rare (McDermott, 1976).

As we have discussed in Section 2.1.1, it is clear that both types of information are used in human interpretation of natural language. When different constraints conflict, the rules for resolving the conflict are complex:

---

[13]As with many linguistic phenomena, the issue is much less crucial in generation than in interpretation, because a generator generally can avoid difficult cases by limiting its pronoun use to clear-cut cases, with at worst some loss in the naturalness of the expression.

sometimes pragmatic preferences overrule formal preferences, sometimes *vice versa,* sometimes the conflict leads to an unintelligible sentence. This line of work continues among linguists (Kehler et al., 2008), but has largely been abandoned in AI, due to the difficulties of developing such systems at scale.

Statistical approaches: The DARPA-funded conferences MUC-6 and MUC-7 (Message Understanding Conferences) (Grishman and Sundheim, 1996; Chinchor, 1998) included an evaluation of coreference resolution systems. Both conferences involved the construction of large text corpora, drawn from natural sources, and carefully hand-annotated following guidelines that were carefully worked out (though not uncontroversial) and presented in careful detail to the annotators. Systems used the machine learning technology available at the time and were both trained and evaluated on these corpora. Systems were generally trained for a fairly narrowly defined tasks; progress relative to the evaluation metric tended to take place steadily, in small increments. This general approach dominated work in the area for the next decade.

Deep learning systems trained on big data: A second revolution in natural language processing began in the early 2010's, and has dominated the field since around 2016. The technology used is some variant of deep learning; in the last two years, particularly transformer-based technology (Devlin et al., 2019). Much of the most impactful technology is task-independent, particularly enormous language modeling systems and word embeddings; performance on individual tasks can then be enhanced with fine-tuning. Data sets of all kinds—enormous unannotated corpora, datasets annotated or created by crowd-workers, synthetic datasets—have proliferated. Progress has in some cases been astonishing.

In statistical approaches, and still more in approaches based on deep learning, commonsense knowledge has been incorporated into the systems only to the extent that the learning procedure implicitly abstracts it from text corpora.

Though the deep-learning technology had led to incomparably more powerful systems than the knowledge-based approach, the former does not subsume the latter. Many of the issues that were addressed in knowledge-based systems are not reliably solved in 2021 state-of-the-art technology. For instance, the language modeling system GPT-3 failed (Marcus and Davis, 2020) when tested on an example from (Charniak, 1972).

3.2. Feature-based Approaches to the Winograd Schema Challenge

This section covers the approaches that collect knowledge in form of explicit rules from knowledge bases, internet search queries, and use logic-based systems or optimization techniques to deduce the answer. Results of methods that rely on search engines, such as Google, can be irreproducible, as they strongly depend on the search results.

The majority of the early approaches to the Winograd Schema Challenge were feature-based or logic-based. While powerful in theory, their main bottleneck was usually the processing of the input data or relevant background knowledge. This is best demonstrated by Sharma (2019) who designs a reasoning algorithm that solves 84.2% of examples in $WSC285$ if the input processing and background knowledge acquisition are done manually. The same algorithm solves fewer than 50% of examples when data processing is done with K-Parser (Sharma et al., 2015).

At the same time, the methods in this group are the hardest to compare one to another. The models are often evaluated on specific subsets of Wsc273 or other datasets, making direct comparison hard. We thus chose not to summarize the results in a table like we do for neural and language-model-based approaches.

Approaches in this work usually work in three steps: First, the input sentence is processed to extract keywords and semantics. In the second step, relevant background knowledge is extracted from various sources, most commonly knowledge bases and internet search queries. Finally, all this information is combined used the input to the *reasoning algorithm*, which gives the final answer. The third step is where approaches differ the most one from another. Approaches used were, for example, SVM-rankers (Rahman and Ng, 2012), integer linear programming (Peng et al., 2015), answer set programming (Sharma et al., 2015), message passing on a graph (Fähndrich et al., 2018), and formal logic (Isaak and Michael, 2016).

While many of these models achieve decent results on the data they have been evaluated on, their ability to generalize can be contested. Sharma et al. (2015), for example, report achieving 69% on a selected subset of Wsc285, but when re-evaluated on a different subset of the data, the same model only achieved chance-level performance (Zhang and Song, 2018).

We highlight Emami et al. (2018) as the first model to achieve a better-than-chance accuracy (57.1%) performance on the entire Wsc273. Their system is completely rule-based and focuses on high-quality knowledge hunting, rather than reasoning, showing the importance of the former. Unlike

26

	training resource	Wsc273	Pdp
Liu et al. (2017a)	cause-effect pairs	70%†	–
Liu et al. (2017b)	cause-effect pairs, Ontonotes	52.8%	58.3%
Zhang and Song (2018)	Wikipedia, dependency parser	60.33%†	–
Opitz and Frank (2018)	DPR, InferSent	56%	–
Wang et al. (2019c)	Gutenberg, 1 Billion Word	62.4%	78.3%

Table 3: Resources, both data and tools, used by different neural approaches, ordered by the time of publication.
† denotes results that were obtained on modified version of the Wsc273 dataset, usually a hand-picked subset. Such results are not guaranteed to generalize to full dataset.

neural approaches from later sections, this model is not negatively affected by switching candidates.

With the increase in the popularity of large language models, interest in feature and reasoning-based methods decreased, however, Hong and Bennett (2020) note that hybrid approaches can be useful when solving schemas of a very specific type, using a language model when the symbolic approach does not yield an answer.

3.3. Neural Approaches

This section contains approaches that rely on neural networks and deep learning, but do not use pre-trained language models. Models in this section are usually designed, built, and trained from scratch in contrast to models that use language models and are usually built on top of an off-the-shelf pre-trained neural network. We find that several ideas introduced in this section are later adjusted and scaled to language models; see Section 3.4. Note that each work comes with a collection of model-specific architecture designs that are not covered in detail.

We summarize the resources and approaches of neural methods in Table 3 and give main observations in the following paragraph. Most of the methods in this section attempt to incorporate *semantic features* into the unsupervised training of the models. The listed pieces of work usually introduce tricks that aim to force a model to pick up the relevant training signals, either through specific architectures, or through appropriate formulation of the training data. Success is often limited and many models are thus only evaluated on subsets of the dataset that suit the model better, making a direct comparison of the results non-trivial. An evident trend through

	Language Model	fine-tuning or external data
Trinh and Le (2018)	custom LSTM	–
Radford et al. (2019)	GPT-2	–
Klein and Nabi (2019)	BERT	–
Prakash et al. (2019)	custom LSTM	internet querying
Kocijan et al. (2019b)	BERT	MaskedWiki, DPR
Kocijan et al. (2019a)	BERT	WikiCREM, DPR, GAP
Ruan et al. (2019)	BERT	DPR
He et al. (2019)	BERT	DPR
Ye et al. (2019)	BERT	ConceptNet, DPR
Sakaguchi et al. (2020)	RoBERTa	WinoGrande
Melo et al. (2020)	custom LSTM	–
Brown et al. (2020)	GPT-3	–
Yang et al. (2020)	RoBERTa	WinoGrande and generated data
Lin et al. (2020)	T5 (3B)	WinoGrande
Khashabi et al. (2020)	T5	WinoGrande and QA tasks
Lourie et al. (2021b)	T5	WinoGrande and RAINBOW

Table 4: Resources used by different language-model based approaches, ordered by the time of publication. With time, ever larger language models and more additional fine-tuning data was used.

time is the movement from methods based on *(contextual) word embeddings* (e.g., Zhang and Song (2018)) into end-to-end systems with a more complex set of objectives Wang et al. (2019c).

3.4. Language Model Approaches

This section covers the approaches that use neural language models to tackle the Winograd Schema Challenge. Most of them use one or more language models that were trained on a large corpus of text. Several authors use large pre-trained language models, such as BERT (Devlin et al., 2019), and have to tailor their approach accordingly. Such works thus focus on better fine-tuning of such language models instead of inventing new architectures.

We summarize the resources and results of all language-model based approaches in Tables 4 and 5, respectively. A careful comparison of the approaches in the listed works reveals multiple changes of trends over time. We list and discuss them below.

- Used language models consistently rise in complexity. First approaches

	Wsc273	Wnli	Pdp	WinoGrande
Trinh and Le (2018)	63.7%	–	70%	–
Radford et al. (2019)	70.7%	–	–	–
Klein and Nabi (2019)	60.3%	–	–	–
Prakash et al. (2019)	70.17%	–	–	–
Kocijan et al. (2019b)	72.5%	74.7%	–	–
Kocijan et al. (2019a)	71.8%	74.7%	86.7%	–
Ruan et al. (2019)	71.1%	–	–	–
He et al. (2019)	75.1%	89%	90.0%	–
Ye et al. (2019)	75.5%	83.6%	–	–
Sakaguchi et al. (2020)	90.1%	85.6%	87.5%	79.1%
Brown et al. (2020)	88.3%	–	–	77.7%
Yang et al. (2020)	–	–	–	80.0%
Lin et al. (2020)	–	–	–	84.6%
Khashabi et al. (2020)	–	–	–	89.4%
Lourie et al. (2021b)	–	–	–	91.3%

Table 5: Results of language-model based approaches on the most commonly used datasets. Less commonly used evaluation sets include Dpr and Wsc285. The only non-English result by Melo et al. (2020) who evaluate their model on a set of Portugese Winograd Schemas is not included.

use custom-made LSTMs, trained specifically for the task, which are replaced with large-scale language models such as BERT, in turn replaced with even larger T5 and GPT-3. This trend is not exclusive to the Winograd Schema Challenge, and it is a general ongoing trend in NLP.

- Appropriate fine-tuning grows in importance, as it turns out to be the most efficient way of improving performance. Approaches that use language models without fine-tuning (e.g., Trinh and Le (2018) and Klein and Nabi (2019)) are outperformed by approaches that fine-tune the same models on WSC-like data (e.g., DPR). To increase the impact of fine-tuning, language models are often fine-tuned on more than one dataset. Earlier approaches usually use synthetic datasets (Kocijan et al., 2019b; Ye et al., 2019; Yang et al., 2020), while later approaches resort to multi-task learning on related tasks (Khashabi et al., 2020; Lourie et al., 2021b).

- WNLI only becomes commonly used after the reformulation trick described in Section 2.4 is introduced by Kocijan et al. (2019b). We note that many models were evaluated on this dataset as part of the GLUE benchmark (Wang et al., 2019b), which we do not include in this survey, because they do not introduce any ideas specific to the task. Their novelty is in better language modelling, which is not the focus of this paper.

- The introduction of WinoGrande dataset made most other smaller datasets less interesting to researchers, even if these small datasets are higher in quality. The majority of the papers published after the release of WinoGrande (Sakaguchi et al., 2020) only evaluated their work on that dataset.

We highlight that many of the approaches that use large-scale pre-trained language models use various input and output formats, which can affect the performance of the model. For a detailed comparison of different training objectives and their impact on the performance of the pre-trained language model, we refer the reader to the work by Yordanov et al. (2020) and Liu et al. (2020).

3.5. The Winograd Schema Challenge at IJCAI-16

The first and last running of the Winograd Schema Challenge as a competition took place at IJCAI-16 (Davis et al., 2017b). Six systems were entered and tested on a collection of 60 PDPs. The most successful of these was that of Liu et al. (2017b) discussed in Section 3.3 above, which achieved a score of 58% on the test collection. The other contestants for the most part used knowledge-based technology, and submitted systems that incorporated only the specific knowledge that was required for the examples that had been published. Not surprisingly, these systems did not do better than chance when tested on new examples.

4. The Problem of Commonsense Reasoning Remains

The Winograd Schema Challenge as originally formulated has largely been overcome. However, this accomplishment may in part reflect flaws in its formulation and execution. Indeed, Elazar et al. (2021) argue that the success of existing models at solving WSC may be largely artifactual. They write

> We provide three explanations for the perceived progress on the WS task: (1) lax evaluation criteria, (2) artifacts in the datasets that remain despite efforts to remove them, and (3) knowledge and reasoning leakage from large training data.

In their experiments, they determined that, when the form of the task, the training regime, the training set, and the evaluation measure were modified to correct for these, the performance of existing language models dropped significantly.

There is no question that the problem of commonsense reasoning still stands as one of the major challenges facing AI (Davis and Marcus, 2015). Large language models spew nonsense Marcus and Davis (2020), as do question-answering systems. Speech transcription and auto-correct systems of every kind regularly turn sensible speech and text into gibberish, by turns confusing, humorous, or embarrassing. Vision systems do not ensure that their interpretations are physically or even geometrically coherent. Video interpretation systems have the barest idea of what might be going on. Robotic systems are extremely limited in their understanding of the world they are interacting with.

This gap has recently come increasingly to the fore of the general discussion of AI challenges; indeed DARPA created a large project, "Machines Common Sense" to fund research in the area (Gunning, 2018) and "commonsense" has become a hot buzzword. However, it still remains the case that research into commonsense reasoning constitutes a tiny fraction of AI research—indeed, a rapidly shrinking fraction, given the extraordinary explosion of AI research—and that, almost without exception, successful AI applications avoid the issue.

Meanwhile, the need for commonsense understanding in AI systems remains as crucial as ever. The facts that, if a small object O does not fits into a container C, then O will not fit into a container smaller than C, and an object larger than O will not fit into C[14] are the same whether the object and container are a trophy in a suitcase, a picture on a printed page, or a quart of milk in a bottle. They are the same whether they are being used to disambiguate the "trophy" sentence; to understand a narrative or video in which a container is being packed; to plan how to load a container; or to write a program for a robot to load containers. They can be used metaphorically in reasoning about fitting data into computer memory or fitting content into an article with word limits. They are part of a rich, general understanding of space and motion and of the constraints on how various kinds of physical objects can be moved and manipulated (Davis et al., 2017a; Lake and Murphy, 2020). The fact that large language models with one hundred billion parameters trained on half a trillion words and fine-tuned on the pronoun disambiguation task can learn enough linguistic patterns that they can disambiguate the pronouns in the "trophy" sentence does not solve the larger problem; it is not even progress toward solving the larger problem.

The WSC and similar datasets do not even explore all aspects of the use of commonsense knowledge in resolving pronoun references. For one thing, these datasets are limited to referents explicitly named in the same or previous sentence. In ordinary usage, the referent of a pronoun may be implicit from the context but not previously named in any noun phrase; for example, the referent of *it* in *I hiked for an hour and it tired me out* or the referent of *they* in *I went to the hospital but they sent me home.*

Furthermore, as argued in (Elazar et al., 2021), the success at a particular

---

[14]Depending on the exact interpretation of "larger" and "fit", this can be either a reliable inference or a plausible inference.

task, such as pronoun disambiguation of a model that has been fine-tuned to that task is not at all a reliable measure of the degree to which the model has learned commmonsense knowledge broadly, or even to which it has learned the commonsense knowledge needed for language understanding. For one thing, pronoun disambiguation may only probe a limited subset of the commonsense knowledge needed in natural language understanding. For another, the fine-tuned model may only able to use its commonsense knowledge for this single task but not for other understanding tasks. For instance, it may "know" that small objects fit better in large containers than vice versa for the purpose of pronoun reference resolution, but not for any other kind of disambiguation or interpretation.

The goal of the WSC was to present a challenge that (a) was clearly "commonsensical" and easily carried out by humans; (b) was easy to evaluate; (c) was an adequate test of commonsense reasoning abilities. It succeeded at (a) and (b), but not at (c). At this point, it seems increasingly doubtful that any benchmark or challenge set that consists of many bite-sized instances of a single simple tasks can adequately test for commonsense reasoning. It may always be possible to construct a training set comparable to the dataset and to train a system that can exploit task-specific features while still avoiding confronting the commonsense knowledge that are involved when humans do the task—especially if, as is general practice, the training set and the test set are random samples of the same corpus.

Rather, it may be necessary to look at larger tasks with less tightly-structured measures of success, such as reading an extended document and giving free-form answers to free-form questions, or carrying out a new task from verbal instructions. An alternative approach would be to interactively probe the AI system's depth of understanding. In designing such a measure, it is always important to keep in mind that, at least at present, AI tends to stumble over basic concrete realities much more than over abstractions; it is more likely to fail conspicuously in discussing a pot being knocked off a counter than in discussing whether automata can be truly conscious (Davis, 2021).

Acknowledgments

OUCS/EPSRC-NPIF/VK/1123106.

References

P. Amsili, O. Seminck, A Google-proof collection of French Winograd schemas, in: Proc. 2nd CORBON Workshop, 2017.

D. Bender, Establishing a human baseline for the Winograd Schema Challenge, 2015.

T. Bernard, T. Han, Mandarinograd: A Chinese collection of Winograd schemas, in: Proceedings of the 12th Language Resources and Evaluation Conference, European Language Resources Association, 2020.

T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, in: Advances in Neural Information Processing Systems, 2020.

J. Q. Candela, I. Dagan, B. Magnini, F. d'Alché-Buc (Eds.), Machine Learning Challenges, Evaluating Predictive Uncertainty, Visual Object Classification and Recognizing Textual Entailment, First PASCAL Machine Learning Challenges Workshop, MLCW 2005, Southampton, UK, April 11-13, 2005, Revised Selected Papers, Springer, 2006.

E. Charniak, Toward a Model of Children's Story Comprehension, Ph.D. thesis, Massachusetts Institute of Technology, 1972.

N. A. Chinchor, Overview of MUC-7/MET-2, Technical Report, Science Applications International Corp., 1998.

F. Cozman, H. Munhoz, The Winograd schemas from hell, in: Anais do XVII Encontro Nacional de Inteligência Artificial e Computacional, SBC, 2020, pp. 531–542.

E. Davis, Qualitative spatial reasoning in interpreting text and narrative. spatial cognition and computation. forthcoming, 2013.

E. Davis, Using human skills taxonomies and tests in as measures of artificial intelligence, in: S. Elliott (Ed.), Skills and Tests for Assessing AI and Robotics, OECD Publishing, 2021.

E. Davis, G. Marcus, Commonsense reasoning and commonsense knowledge in artificial intelligence, Communications of the ACM 58 (2015) 92–103.

E. Davis, G. Marcus, N. Frazier-Logue, Commonsense reasoning about containers using radically incomplete information, Artificial intelligence 248 (2017a) 46–84.

E. Davis, L. Morgenstern, C. Ortiz, Human tests of material for the Winograd Schema Challenge, 2016. URL: https://cs.nyu.edu/faculty/davise/papers/WS2016SubjectTests.pdf.

E. Davis, L. Morgenstern, C. L. Ortiz, The first Winograd Schema Challenge at IJCAI-16, AI Magazine (2017b).

E. Davis, X. Pan, A corpus of challenging pronoun disambiguation problems, adapted from children's books, 2015. URL: https://cs.nyu.edu/faculty/davise/annotate/corpus.pdf.

J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in: Proc. NAACL, 2019.

Y. Elazar, H. Zhang, Y. Goldberg, D. Roth, Back to square one: Bias detection, training and commonsense disentanglement in the Winograd schema, arXiv preprint arXiv:2104.08161 (2021).

A. Emami, N. De La Cruz, A. Trischler, K. Suleman, J. C. K. Cheung, A knowledge hunting framework for common sense reasoning, in: Proc. EMNLP, Brussels, Belgium, 2018.

D. Emelin, R. Sennrich, Wino-X: Multilingual Winograd schemas for commonsense reasoning and coreference resolution, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2021, pp. 8517–8532. URL: https://aclanthology.org/2021.emnlp-main.670.

J. Fähndrich, S. Weber, H. Kanthak, A marker passing approach to Winograd schemas, in: Semantic Technology, Springer, 2018.

R. Grishman, B. Sundheim, Design of the MUC-6 evaluation, Technical Report, New York University, Dept. of Computer Science, 1996.

B. J. Grosz, The representation and use of focus in a system for understanding dialogs, in: IJCAI, 1977, pp. 67–76.

D. Gunning, Machine common sense concept paper, arXiv preprint arXiv:1810.07528 (2018).

S. Hansson, C. Mavromatakis, Y. Adesam, G. Bouma, D. Dannells, The Swedish Winogender Database, in: Proceedings of the 23rd Nordic Conference on Computational Linguistics, 2021, pp. 452–457.

P. He, X. Liu, W. Chen, J. Gao, A hybrid neural network model for commonsense reasoning, in: Proc. 1st Workshop on Commonsense Inference in NLP, 2019.

W. He, C. Huang, Y. Liu, X. Zhu, WinoLogic: A zero-shot logic-based diagnostic dataset for Winograd Schema Challenge, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2021, pp. 3779–3789. URL: https://aclanthology.org/2021.emnlp-main.307.

J. R. Hobbs, Coherence and coreference, Cognitive Science 3 (1979) 67–90.

J. R. Hobbs, M. E. Stickel, D. E. Appelt, P. Martin, Interpretation as abduction, Artificial intelligence 63 (1993) 69–142.

S. J. Hong, B. Bennett, Tackling domain-specific Winograd schemas with knowledge-based reasoning and machine learning, arXiv preprint arXiv:2011.12081 (2020).

N. Isaak, L. Michael, Tackling the Winograd Schema Challenge through machine logical inferences., in: STAIRS, 2016.

N. Isaak, L. Michael, Winoflexi: A crowdsourcing platform for the development of Winograd schemas, in: Proc. AI, 2019.

N. Isaak, L. Michael, Winventor: A machine-driven approach for the development of Winograd schemas., in: ICAART (2), 2020, pp. 26–35.

N. Joukovsky, J. Powell, A peacock in the attic, Times Literary Supplement (2011).

D. Kahneman, Thinking, fast and slow, Farrar, Straus and Giroux, 2011.

D. Kakwani, A. Kunchukuttan, S. Golla, N. Gokul, A. Bhattacharyya, M. M. Khapra, P. Kumar, inlpsuite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for indian languages, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, 2020, pp. 4948–4961.

M. Kameyama, A property-sharing constraint in centering, in: 24th Annual Meeting of the Association for Computational Linguistics, 1986, pp. 200–206.

A. Kehler, Testing for common sense: Thoughts on pronoun interpretation and the Winograd schema challenge, 2015. Talk presented at the Workshop on Language & Common Sense: Integrating Across Psychology, Linguistics, and Computer Science, CogSci-2015.

A. Kehler, L. Kertz, H. Rohde, J. L. Elman, Coherence and coreference revisited, Journal of semantics 25 (2008) 1–44.

D. Khashabi, S. Min, T. Khot, A. Sabharwal, O. Tafjord, P. Clark, H. Hajishirzi, UnifiedQA: Crossing format boundaries with a single QA system, arXiv preprint arXiv:2005.00700 (2020).

T. Klein, M. Nabi, Attention is (not) all you need for commonsense reasoning, in: Proc. ACL, 2019.

W. Knight, Tougher turing test exposes chatbots' stupidity, Technology Review (2016).

V. Kocijan, O.-M. Camburu, A.-M. Cretu, Y. Yordanov, P. Blunsom, T. Lukasiewicz, WikiCREM: a large unsupervised corpus for coreference resolution, in: Proc. EMNLP, 2019a.

V. Kocijan, A.-M. Cretu, O.-M. Camburu, Y. Yordanov, T. Lukasiewicz, A surprisingly robust trick for Winograd Schema Challenge, in: Proc. ACL, 2019b.

T. Kocmi, T. Limisiewicz, G. Stanovsky, Gender coreference and bias evaluation at WMT 2020, in: Fifth Conference on Machine Translation - Proceedings of the Conference, Association for Computational Linguistics, Stroudsburg, PA, USA, 2020, pp. 357–364.

B. M. Lake, G. L. Murphy, Word meaning in minds and machines, arXiv preprint arXiv:2008.01766 (2020).

H. J. Levesque, The Winograd Schema Challenge, AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning (2011).

H. J. Levesque, On our best behaviour, Artif. Intell. (2014).

H. J. Levesque, E. Davis, L. Morgenstern, The Winograd Schema Challenge, in: Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning, 2012.

S.-C. Lin, J.-H. Yang, R. Nogueira, M.-F. Tsai, C.-J. Wang, J. Lin, TTTT-Tackling WinoGrande schemas, arXiv preprint arXiv:2003.08380 (2020).

T. Linzen, How can we accelerate progress towards human-like linguistic generalization?, arXiv preprint arXiv:2005.00955 (2020).

H. Liu, W. Huang, D. Mungra, S. R. Bowman, Precise task formalization matters in Winograd schema evaluations, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), Association for Computational Linguistics, 2020.

Q. Liu, H. Jiang, A. Evdokimov, Z.-H. Ling, X. Zhu, S. Wei, Y. Hu, Cause-effect knowledge acquisition and neural association model for solving a set of Winograd Schema Problems, in: Proc. IJCAI-17, 2017a.

Q. Liu, H. Jiang, Z.-H. Ling, X. Zhu, S. Wei, Y. Hu, Combing context and commonsense knowledge through neural networks for solving Winograd Schema Problems (2017b).

Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, RoBERTa: A robustly optimized BERT pretraining approach, arXiv:1907.11692 (2019).

N. Lourie, R. L. Bras, C. Bhagavatula, Y. Choi, UNICORN on RAIN-BOW: A universal commonsense reasoning model on a new multitask benchmark, in: Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021, AAAI Press, 2021a, pp. 13480–13488. URL: `https://ojs.aaai.org/index.php/AAAI/article/view/17590`.

N. Lourie, R. L. Bras, C. Bhagavatula, Y. Choi, UNICORN on RAINBOW: A universal commonsense reasoning model on a new multitask benchmark, in: Proc. AAAI, 2021b.

G. Marcus, E. Davis, GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about, Technology Review (2020).

J. Markoff, Software is smart enough for SAT, but still far from intelligent, The New York Times (2015).

D. McDermott, Artificial intelligence meets natural stupidity, ACM Sigart Bulletin (1976) 4–9.

G. Melo, V. Imaizumi, F. Cozman, Esquemas de Winograd em português, in: Anais do XVI Encontro Nacional de Inteligência Artificial e Computacional, 2020.

T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, in: Advances in NIPS 26, 2013.

L. Morgenstern, Technical perspective: The importance of WINOGRANDE, Commun. ACM 64 (2021) 98. URL: `https://doi.org/10.1145/3474378`. doi:`10.1145/3474378`.

L. Morgenstern, E. Davis, C. L. O. Jr., Planning, executing, and evaluating the winograd schema challenge, AI Mag.

37 (2016) 50–54. URL: https://doi.org/10.1609/aimag.v37i1.2639. doi:10.1609/aimag.v37i1.2639.

N. Nangia, S. R. Bowman, Human vs. Muppet: A conservative estimate of human performance on the GLUE benchmark, arXiv preprint arXiv:1905.10425 (2019).

J. Opitz, A. Frank, Addressing the Winograd Schema Challenge as a sequence ranking task, in: Proc. 1st International Workshop on Language Cognition and Computational Models, ACL, 2018.

H. Peng, D. Khashabi, D. Roth, Solving hard co-reference problems, in: Proc. NAACL, 2015.

M. Poesio, S. Ponzetto, Y. Versley, Computational models of anaphora resolution: A survey, 2011. URL: http://web.informatik.uni-mannheim.de/ponzetto/pubs/poesio10a.pdf.

M. Poesio, R. Stuckardt, Y. Versley, Anaphora resolution, Springer, 2016.

A. Prakash, A. Sharma, A. Mitra, C. Baral, Combining knowledge hunting and neural language models to solve the Winograd schema challenge, in: Proc. ACL, 2019.

A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, Language models are unsupervised multitask learners (2019).

A. Rahman, V. Ng, Resolving complex cases of definite pronouns: The Winograd Schema Challenge, in: Proc. EMNLP, 2012.

P. Rajpurkar, J. Zhang, K. Lopyrev, P. Liang, Squad: 100,000+ questions for machine comprehension of text, arXiv preprint arXiv:1606.05250 (2016).

H. Rohde, Pronoun interpretation and production, in: C. Cummins, N. Katsos (Eds.), The Oxford Handbook of Experimental Semantics and Pragmatics, Oxford University Press, 2018.

Y. Ruan, X. Zhu, Z. Ling, Z. Shi, Q. Liu, S. Wei, Exploring unsupervised pretraining and sentence structure modelling for Winograd Schema Challenge, arXiv:1904.09705 (2019).

R. Rudinger, J. Naradowsky, B. Leonard, B. Van Durme, Gender bias in coreference resolution, in: Proc. NAACL, 2018.

K. Sakaguchi, R. L. Bras, C. Bhagavatula, Y. Choi, WINOGRANDE: An adversarial Winograd Schema Challenge at scale, in: Proc. AAAI, 2020.

K. Sakaguchi, R. L. Bras, C. Bhagavatula, Y. Choi, Winogrande: an adversarial winograd schema challenge at scale, Commun. ACM 64 (2021) 99–106. URL: https://doi.org/10.1145/3474381. doi:10.1145/3474381.

R. C. Schank, R. P. Abelson, Scripts, plans, goals, and understanding: An inquiry into human knowledge structures, Lawrence Erlbaum Associates, 1977.

A. Sharma, Using answer set programming for commonsense reasoning in the Winograd schema challenge, arXiv:1907.11112 (2019).

A. Sharma, N. Vo, S. Aditya, C. Baral, Towards addressing the Winograd Schema Challenge — Building and using a semantic parser and a knowledge hunting module, in: Proc. IJCAI, 2015.

C. Sidner, Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse, Ph.D. thesis, Massachusetts Institute of Technology, 1979.

G. Stanovsky, N. A. Smith, L. Zettlemoyer, Evaluating gender bias in machine translation, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, 2019.

P. Trichelair, A. Emami, J. C. K. Cheung, A. Trischler, K. Suleman, F. Diaz, On the evaluation of common-sense reasoning in natural language understanding, in: Proc. NeurIPS Workshop on Critiquing and Correcting Trends in Machine Learning, 2018.

T. H. Trinh, Q. V. Le, A simple method for commonsense reasoning, arXiv:1806.02847 (2018).

A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, S. R. Bowman, SuperGLUE: A stickier benchmark for general-purpose language understanding systems, arXiv preprint 1905.00537 (2019a).

A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, S. R. Bowman, GLUE: A multi-task benchmark and analysis platform for natural language understanding, in: Proc. ICLR., 2019b.

S. Wang, S. Zhang, Y. Shen, X. Liu, J. Liu, J. Gao, J. Jiang, Unsupervised deep structured semantic models for commonsense reasoning, in: Proc. NAACL, 2019c.

Y. Wilks, An intelligent analyzer and understander of english, Communications of the ACM 18 (1975) 264–274.

T. Winograd, Understanding Natural Language, Academic Press, Inc., Orlando, FL, USA, 1972.

F. Wolf, E. Gibson, T. Desmet, Discourse coherence and pronoun resolution, Language and cognitive processes 19 (2004) 665–675.

L. Xu, X. Zhang, L. Li, H. Hu, C. Cao, W. Liu, J. Li, Y. Li, K. Sun, Y. Xu, et al., Clue: A Chinese language understanding evaluation benchmark, arXiv preprint arXiv:2004.05986 (2020).

Y. Yang, C. Malaviya, J. Fernandez, S. Swayamdipta, R. Le Bras, J.-P. Wang, C. Bhagavatula, Y. Choi, D. Downey, Generative data augmentation for commonsense reasoning, in: Findings of the Association for Computational Linguistics: EMNLP 2020, Association for Computational Linguistics, 2020.

Z.-X. Ye, Q. Chen, W. Wang, Z.-H. Ling, Align, mask and select: A simple method for incorporating commonsense knowledge into language representation models, arXiv:1908.06725 (2019).

Y. Yordanov, O.-M. Camburu, V. Kocijan, T. Lukasiewicz, Does the objective matter? comparing training objectives for pronoun resolution, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, November 16–20, 2020, 2020.

A. Žagar, M. Robnik-Šikonja, T. Goli, Š. Arhar Holdt, Slovene translation of SuperGLUE, 2020. URL: http://hdl.handle.net/11356/1380, slovenian language resource repository CLARIN.SI.

H. Zhang, Y. Song, A distributed solution for Winograd schema challenge, in: ACM International Conference Proceeding Series, ACM, 2018.

H. Zhang, X. Zhao, Y. Song, WinoWhy: A deep diagnosis of essential commonsense knowledge for answering Winograd schema challenge, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, 2020.

J. Zhao, T. Wang, M. Yatskar, V. Ordonez, K.-W. Chang, Gender bias in coreference resolution: Evaluation and debiasing methods, in: Proc. NAACL 2018, 2018.