

EvoGRAD: A Dynamic Take on the Winograd Schema Challenge with Human Adversaries

Jing Han Sun, Ali Emami

University of Montreal/Mila, Brock University
Montreal, Canada, Saint Catharines, Canada
jing.han.sun@umontreal.ca, aemami@brocku.ca

Abstract

While Large Language Models (LLMs) excel at the Winograd Schema Challenge (WSC), a coreference resolution task testing common-sense reasoning through pronoun disambiguation, they struggle with instances that feature minor alterations or rewording. To address this, we introduce EvoGRAD, an open-source platform that harnesses a human-in-the-loop approach to create a dynamic dataset tailored to such altered WSC instances. Leveraging ChatGPT’s capabilities, we expand our task instances from 182 to 3691, setting a new benchmark for diverse common-sense reasoning datasets. Additionally, we introduce the *error depth* metric, assessing model stability in dynamic tasks. Our results emphasize the challenge posed by EvoGRAD: Even the best performing LLM, GPT-3.5, achieves an accuracy of 65.0% with an average error depth of 7.2, a stark contrast to human performance of 92.8% accuracy without perturbation errors. This highlights ongoing model limitations and the value of dynamic datasets in uncovering them.

Keywords: Winograd Schema Challenge, Common-sense Reasoning, Large Language Models

1. Introduction

The Winograd Schema Challenge (WSC), a coreference resolution task, was developed to gauge the common-sense reasoning of automated systems (Winograd, 1972; Levesque et al., 2011). Given subtly varying sentence pairs, the task is to correctly associate a pronoun with a noun, as illustrated below:

- (1) a. Tom told Ralph, “Check,” as he moved his bishop. (Answer: Tom)
- b. Tom told Ralph, “Check,” as he took his bishop. (Answer: Ralph)

In these examples, chess knowledge informs our interpretation of the pronoun *his*—either referring to *Tom* or *Ralph*—based on the action performed, either a *move* or *take*. While humans find such tasks intuitive, they pose a challenge for statistical models, especially when lacking exposure to basic rules or common knowledge. Yet, recent developments of extensive common-sense reasoning datasets and benchmarks have allowed LLMs to achieve near-human performance on WSC variants (Brown et al., 2020; Sakaguchi et al., 2020). This impressive accomplishment raises the question: has the WSC, seen as a definitive alternative to the Turing Test, been definitively “defeated” (Kocijan et al., 2022)?

At the same time, evidence suggests that even slight alterations to a WSC task can significantly undermine a model’s performance (Jia and Liang, 2017; Trichelair et al., 2018, 2019; Balasubramanian et al., 2020; Lin et al., 2020a; Elazar et al.,

2021a). This instability may reflect a discrepancy between current supervision paradigms and the dynamic nature of common sense acquisition. It suggests the potential value of exploring various approaches, including a human-and-model-in-the-loop concept, as part of a broader strategy to address these challenges (Nie et al., 2020; Kiela et al., 2021; Lu et al., 2022).

Existing datasets, often curated by select scientific communities or crowdsourcing platforms, may also unintentionally bias models toward certain knowledge instances or values, which may not be universally shared. This consideration underscores the need for diverse, dynamic, and inclusive benchmarks in the journey towards systems equipped with generalized common sense.

Consider the chess example mentioned earlier. While the original WSC sentences test the model’s understanding of the game’s basic rules, perturbations can further probe deeper nuances and potential biases:

- (2) a. Maria told Jane, “Your move,” as she adjusted her queen. (Answer: Maria)
- b. Maria told Jane, “Your move,” as she glanced at her clock. (Answer: Jane)

In these variations, the emphasis shifts from the action performed on a chess piece to the broader context of a timed chess match. Slight word changes can dramatically alter the correct answer, exposing potential model biases or gaps in understanding. Such perturbations, especially when generated by diverse human contributors, ensure a broader and more comprehensive test of a model’s common-sense reasoning capabilities.

Figure 1: Interface of EvoGRAD at <https://evograd.com>

In this paper, we propose a revisit to the WSC within the framework of human-and-model-in-the-loop. We introduce EvoGRAD, an open-source, user-centric platform dedicated to the active generation and expansion of nuanced examples for the WSC through human-in-the-loop interactions. Our work contributes three primary advancements:

A novel data construction mechanism: We enhance the WSC with our unique approach to human-adversarial perturbations, combining human creativity with the efficiency of ChatGPT. This innovative union, along with our use of Wordnet for synonym-based variation, led to a dataset expansion from 182 to 3691 instances, setting a new standard for dynamic, diverse, and high-quality common-sense reasoning datasets. Notably, our evaluations highlight the challenging nature of EvoGRAD, revealing significant gaps in model abilities when compared to human benchmarks.

A new metric for model stability: In response to the instability of transformer-based models on WSC-like tasks (Abdou et al., 2020), we introduce a metric termed *error depth*. This measure, derived from our data construction process, offers a quantifiable assessment of model stability. We advocate for its inclusion in evaluation reports alongside accuracy-based metrics, which could discourage the development of models that achieve high scores due to incorrect reasoning.

Online platform for user contributions: Available at <https://evograd.com>¹, our platform

¹All aspects of the website remain anonymous during the submission and review process to maintain the integrity of the user-contributed data and ensure unbiased evaluation.

encourages public participation in the continuous expansion of the dataset. Users can modify existing task instances and observe the predictions of a chosen LLM, fostering a more participatory and immersive data construction process (Figure 1).

2. Related Work

2.1. WSC-based Datasets

The Winograd Schema Challenge (WSC) (Levesque et al., 2011) inspired various datasets for pronominal coreference resolution, each tackling specific challenges in the WSC or model evaluations. Datasets like Winogrande (Sakaguchi et al., 2020) and KnowRef (Emami et al., 2019) address the WSC’s size constraints. WinoGender (Rudinger et al., 2018), WinoBias (Zhao et al., 2018), and KnowRef-60k (Emami et al., 2020) focus on model biases, while WinoWhy (Zhang et al., 2020) and WinoLogic (He et al., 2021) target common sense deficiencies in models. Some research efforts enhanced the original WSC task (Wang et al., 2018; Trichelair et al., 2018; Kocijan et al., 2019; Elazar et al., 2021a; Zahraei and Emami, 2024) and utilized crowd-sourcing for task development (Isaak and Michael, 2019; Sakaguchi et al., 2020). While these static datasets each offer distinct strengths, they often introduce challenges that necessitate prolonged research and iterations. EvoGrad, on the other hand, adopts a dynamic framework, allowing for swift adjustments and refinements in response to emerging challenges.

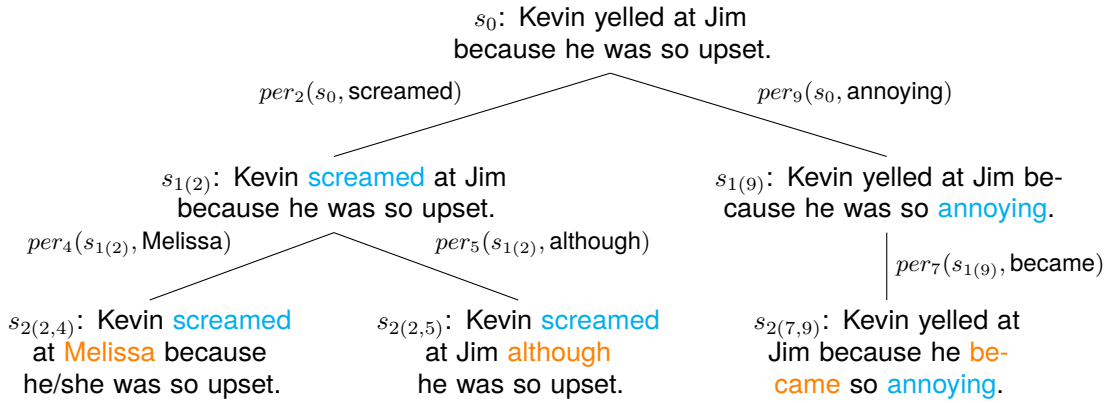


Figure 2: Evolution figure of the sentence “Kevin yelled at Jim because he was so upset.” up to depth level 2.

2.2. Dynamic Datasets

Dynamic datasets, updated over time to present new challenges, have been developed for various tasks (Zellers et al., 2019; Lin et al., 2020b). Adversarial frameworks, as seen in Adversarial SQuAD, SWAG, HellaSWAG, CODAH and ANLI, exemplify this approach (Jia and Liang, 2017; Zellers et al., 2018, 2019; Chen et al., 2019; Nie et al., 2020). Techniques such as AFLite address biases through adversarial filtering (Le Bras et al., 2020), while other methods use continuous learning or a human-model collaborative process (Lan et al., 2017; Yang et al., 2018; Wallace et al., 2019; Dinan et al., 2019; Nie et al., 2020; Xu et al., 2021; Kiela et al., 2021). ANLI and Dynabench are notable for their multi-round adversarial data collection (Nie et al., 2020; Kiela et al., 2021). EvoGRAD, while aligning with the dynamic dataset philosophy, specifically targets WSC-based tasks. It merges human-and-model collaboration, continuous learning, and domain-specific insights for evolutionary data creation, amplifying the depth and relevance of WSC challenges to shed light on common-sense reasoning.

2.3. Data Augmentation Methods in NLP

Data augmentation techniques in NLP create new examples from existing ones, obviating the need for novel data collection (Shi et al., 2021; Feng et al., 2021). These methods include token-level manipulation, text generation restricted, soft data enhancement, and structure-aware data augmentation (Wang and Yang, 2015; Bergmanis et al., 2017; Zhang et al., 2018; Xu et al., 2016). Our approach, mainly a token-level manipulation technique, extends beyond the substitution of words to include the addition and removal of tokens, allowing more significant sentence transformations (Zmigrod et al., 2019; Lu et al., 2020; Shi et al., 2018). We also measure the *depth* of changes (Section 3.6) relative to the original sentence, providing insights into

model stability as a function of perturbations.

2.4. Large Language Models in Data Augmentation and Annotation

Large language models have emerged as effective tools for NLP data augmentation and annotation, often exceeding the performance of crowdworkers in terms of efficiency and cost (Gilardi et al., 2023). These models are effective in tasks such as zero-shot gender identification and providing explanations for implicit hate speech (Kuzman et al., 2023; Huang et al., 2023). AugGPT, for instance, outperforms traditional text augmentation methods in few-shot learning scenarios by rephrasing sentences (Dai et al., 2023). Similarly, ChatGPT has shown the potential to simplify social computing tasks by replicating human-like annotations (Zhu et al., 2023). Building on these insights, we introduce an enhanced data augmentation method that encompasses token substitutions, additions, and removals, aiming to address common-sense reasoning deficiencies in the WSC and related tasks.

3. EvoGRAD

3.1. Dataset Evolution by Perturbation

We adopt an evolutionary approach to dataset expansion, initiating the process with randomly selected instances from the original Winograd Schema Challenge (WSC273) (Levesque et al., 2011) and Winogrande (Sakaguchi et al., 2020), which are correctly resolved by all evaluated models.

Our method introduces a one-word perturbation to each sentence, effectively mutating it via substitution. We define a perturbation function $per_j(s, w)$ that replaces the token at index j in sentence s with the token w . Though primarily substitution-based, this function can also facilitate the addition or re-

moval of words, denoted as $per_j(s, w_j + w)$ and $per_j(s, \epsilon)$ respectively, with ϵ symbolizing an empty string.

The function is generalized as follows:

$$per_j(s_{k(i_1, \dots, i_k)}, w) = s_{(k+1)(i_1, \dots, i_j, \dots, i_{k+1})} \quad j \notin \{i_1, \dots, i_k\} \ \& \ i_1 < \dots < i_{k+1} \quad (1)$$

In this equation, $s_{k(i_1, \dots, i_k)}$ signifies the k th perturbation on the base sentence s_0 , wherein tokens at indices i_1, \dots, i_k have been modified from s_0 (Equation 1). The term k denotes the ‘depth’ or generation of the sentence.

The conditions set for j and indices i_1, \dots, i_{k+1} ensure that a depth increment corresponds solely to the perturbation of a token distinct from those previously perturbed (i.e., i_1, \dots, i_k). Although repeated modifications at the same token position are not prohibited, such sentences maintain their original depths. This approach follows our depth interpretation, emphasizing model stability against sentences that are increasingly divergent from the original. This methodological choice facilitates the systematic generation of progressively varied sentences, thereby enriching the dataset.

The perturbation function is applied iteratively, generating a cascade of output instances from each input instance. This process is illustrated in Figure 2 by the sentence ‘Kevin yelled at Jim because he was so upset.’ Through several iterations of the perturbation function, we generate a wide spectrum of sentences, each incrementally divergent from the original.

3.2. Scaling with ChatGPT

Beyond user contributions, we strategically employed ChatGPT² to vastly expand our dataset. We initialized the process with 14 seed sentences (7 from WSC273 and 7 from Winogrande-valid) and expanded them to 182 instances via human annotation via our website. Then, we designed an elaborate prompt that enabled ChatGPT to act as an ‘expert human annotator’. This step was not only intended to emulate human expertise but also to demonstrate the potential scalability of our dataset. By guiding ChatGPT through a structured decision-making process, similar to that of a human annotator, we aimed to show how LLMs can be used to augment and accelerate data generation and assessment. The prompts were meticulously crafted to guide the model generation process via demonstrative examples and called for frequent self-reflection to ensure the quality of the output. One unique aspect of these prompts was the incorporation of a segmented generation process, interspersed with feedback to ensure quality

²<https://chat.openai.com/>

control and continuous self-assessment. For each instance, we verified semantic coherence and implemented a validation step to ensure pronouns and co-references matched commonly accepted or typical human readings. An illustrative dialogue sample can be found in the Appendix in section A.1.

This rigorous approach to prompt engineering culminated in the generation of approximately 100 new instances per seed sentence. We further diversified these generated sentences by modifying words, altering the correct antecedent, and varying the total perturbation depth from the original sentences. This strategy effectively harnessed the power of human creativity and the scalability of the model to significantly expand our dataset. As a result, we managed to augment our initial 182-instance dataset to a much more extensive collection of 1,414 sentences, thereby facilitating a more comprehensive evaluation of model performance on dynamic WSC tasks.

3.3. Scaling with Wordnet

To increase the diversity of our dataset, we utilized Wordnet (Fellbaum, 2010), a lexical database, to augment the 1,414 sentences obtained from our *ChatGPT Scaling* stage. This process enabled us to nearly triple our dataset size to a final count of 3,691 sentences.

Our strategy was to introduce variability while preserving the context of the sentence and grammatical accuracy. We achieved this by iterating over each sentence and randomly selecting a word—excluding stop words and named entities—for replacement. Once a word was selected, a random synonym from Wordnet was chosen as its substitute. In cases where the chosen word was a verb, we ensured that the replacement synonym matched the tense of the original verb.

This approach allowed us to maintain the integrity of our original dataset while significantly enhancing its size and complexity. The resulting sentences provided a rich basis for model testing, aiding in the generation of a more diverse and nuanced set of pronoun disambiguation scenarios.

3.4. The Dataset

Table 1 outlines the construction and allocation process for our datasets, specifically EvoGRAD-small (S), EvoGRAD-medium (M) and EvoGRAD-large (L). The initial dataset, EvoGRAD-S, comprised 182 instances, all of which were adaptations induced by humans from an original set of 14 sentences.

Subsequently, we generated the EvoGRAD-M dataset, which was divided into three distinct subsets: ‘train’, ‘val’, and ‘test’. These subsets were

Dataset	Sub	Size	Method
EvoGRAD-S	-	182	Human (14 orig.)
EvoGRAD-M	Train	1010	ChatGPT (1-10)
	Val	202	ChatGPT (11-12)
	Test	202	ChatGPT (13-14)
EvoGRAD-L	Train	2963	WordNet (M Train)
	Val	526	WordNet (M Val)
	Test	202	ChatGPT (13-14)

Table 1: Summary of EvoGRAD Allocation

created by perturbing the original sentences using ChatGPT, resulting in a total of 1,414 instances.

Finally, our most extensive dataset, EvoGRAD-L, was constructed by augmenting both the ‘train’ and ‘val’ subsets of EvoGRAD-M using Wordnet, leading to an overall count of 3,691 instances. The ‘test’ subset was retained from the EvoGRAD-M ‘test’ dataset and was generated through further perturbation of EvoGRAD-S sentences via ChatGPT. To illustrate the range of perturbations and their sources, we provide sample instances in Table 2 derived from an original WSC sentence.

3.5. The Platform

To foster collaborative development of EvoGRAD, we have developed an interactive platform, accessible at <https://evograd.com>. Here, global users can actively contribute to the dataset’s evolution by modifying existing sentences.

In the *Build dataset* page, users can select an original or perturbed sentence from a drop-down menu labeled *Original Sentence*. They are then guided to input a modified version of this sentence, replacing the target pronoun with an underscore, in the *New Sentence* field. Following the Winogrande format (Sakaguchi et al., 2020), users also provide the two potential noun antecedents in the *Option 1* and *Option 2* fields, specifying the correct answer.

To enhance user engagement, our platform offers immediate feedback. Users can choose an LLM from a list - including BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and Albert (Lan et al., 2020)—and observe the model’s live prediction. By clicking *Submit*, this prediction is generated, and the newly provided data is incorporated into the dataset.

We prioritize transparency by allowing the dataset, stored as a CSV file, to be downloaded and inspected directly from the platform. To ensure the quality and appropriateness of the submissions, we manually validate all entries. Users are further supported with examples and guidelines. A glimpse of the platform’s interface is depicted in Figure 1.

3.6. Error Depth

Given our dataset construction methodology, we propose the *error depth* (ED) metric to evaluate model stability. While accuracy is a widely used metric to gauge model performance on prediction tasks such as the WSC, it might not effectively capture a model’s resilience against instances that progressively deviate from the original.

There are scenarios where models predict correctly but possibly for the wrong reasons. Sole reliance on accuracy can obscure these nuances. Ideally, a model should demonstrate stability against token substitutions. Although, in the context of the WSC, a token change can alter the answer label, a truly robust model should not be overly sensitive to such modifications.

The error depth metric quantifies a model’s performance on sentences that increasingly diverge from a correctly understood original. Specifically, the error depth denotes the number of perturbations made to the original sentence before the model produces its first incorrect prediction. In other words, it represents the number of modifications a sentence can undergo before the model fails to predict correctly on it.

For clarity, let’s define the symbols:

- s_0 : The original seed sentence.
- $\text{label}(s)$: The true label of sentence s .
- $\text{pred}(s)$: The model’s predicted label for sentence s .
- n_{wrong} : The number of incorrect predictions made by the model on perturbed versions of the original sentence.

With these definitions, the error depth (ED) is formulated as:

$$\overline{ED} \stackrel{\text{def}}{=} \frac{1}{n_{\text{wrong}}} \sum_k^{n_{\text{wrong}}} k \quad (2)$$

if $\text{label}(s_0) = \text{pred}(s_0)$ and

$\text{label}(s_{k(i_1, \dots, i_k)}) \neq \text{pred}(s_{k(i_1, \dots, i_k)})$

Refer to Table 3 for an application of the metric to perturbations of a sentence. In this demonstration, the model mispredicts three sentences: two after five perturbations and one after six. Thus, $\overline{ED} = (5 + 5 + 6)/3 = 5.333$, which means that on average, the model will fail on instances 5.333 modifications away from the original sentence. The error depth functions as an instance-level metric, assessing a model’s stability for individual sentences. Averaging over all instances yields \overline{ED} , which, when paired with accuracy, offers a comprehensive assessment of a model’s performance on tasks like the WSC.

Source	Sentence	Answer	Depth
Original (WSC)	I poured water from the bottle into the cup until _ was full.	cup	0
Human-perturbed	I poured water from the bottle into the cup because _ was empty.	cup	2
ChatGPT-scaled	I poured water from the bottle, filling the cup until _ was empty.	bottle	4
Wordnet-scaled	I decanted water from the feeding bottle into the cup until _ was empty.	feeding bottle	4

Table 2: Sample instances of EvoGRAD derived from an original WSC sentence, showcasing the different methods of sentence generation and perturbation.

Original sentence: Although she was being prosecuted, Monica was welcomed into the sanctuary of the church by Samantha because _ was a sinful criminal.			
Perturbed Sentence	Prediction	True Label	Depth
Although she was being prosecuted, Monica was welcomed into the sanctuary of the church by Samantha because _ was a guilty criminal.	Monica	Monica	1 ✓
Although she was being prosecuted, Monica was welcomed into the sanctuary of the church by Samantha because _ was a compassionate person.	Samantha	Samantha	2 ✓
Even though she was being prosecuted, Monica was guided into the safe haven of the church by Samantha because _ was a virtuous person.	Monica	Samantha	5 ✗
While under prosecution, Monica was brought into the spiritual refuge of the church by Samantha because _ was a good-natured woman.	Monica	Samantha	6 ✗
While being prosecuted, Monica was welcomed into the church’s refuge by Samantha because _ was a law-abiding person.	Monica	Samantha	5 ✗

Table 3: Sample of perturbations constructed from Eq.1 on a Winogrande example, with predictions corresponding to RoBERTa fine-tuned on Winogrande-XL. The model’s incorrect predictions occur at depths 5,6 and 5, respectively, corresponding to the number of modified tokens from the original. Therefore, this sample of 5 perturbed instances has an average error depth (ED) of 5.333.

3.7. Human Performance

Three English-proficient annotators reviewed EvoGRAD-M Val and EvoGRAD-L Val, achieving mean accuracies of 95.2% and 92.8%, respectively. Importantly, they did not exhibit an average error depth, effectively handling perturbations to the full depth of the dataset. A high inter-annotator agreement was recorded with a Fleiss’ Kappa of $\kappa = 0.914$.

4. Experiments and Results

4.1. Model Setup

We evaluated three primary transformer-based models that are masked language models: BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and ALBERT (Lan et al., 2020). These models have been recognized for their strong performance on the WSC and have led the benchmark results. Each

of the models was fine-tuned on the Winogrande-XL dataset (Sakaguchi et al., 2020), which contains approximately 40,000 task instances and is designed to reduce potential annotation biases.

Additionally, we evaluated two left-to-right language models, specifically GPT-3 (text-davinci-003) (Brown et al., 2020) and GPT-3.5 (gpt-3.5-turbo-0613), on the Winogrande-XL and EvoGRAD datasets.

For BERT and RoBERTa, we first aimed to replicate top-performing models from existing literature. Using Huggingface’s package (Wolf et al., 2020), we achieved validation accuracies of 62.75% for BERT-large-uncased and 76.09% for RoBERTa-large. Although slightly below the reported accuracies in (Sakaguchi et al., 2020), variations in hyperparameter tuning may account for the differences. A similar approach was taken for ALBERT-large-v2, with a resulting accuracy of 64.64%.

Hyperparameters for BERT, RoBERTa, and AL-

Model	Tuning	Wino-valid	EvoGRAD-M-val	EvoGRAD-L-val
BERT	EvoGRAD-M	-	60.4 (6.913)	-
	EvoGRAD-L	-	-	54.9 (6.867)
	Wino	62.75	— (7.302)	— (7.258)
	Wino + EvoGRAD-M	63.06	— (7.308)	-
	Wino + EvoGRAD-L	62.98	-	— (7.232)
RoBERTa	EvoGRAD-M	-	58.4 (6.762)	-
	EvoGRAD-L	-	-	60.3 (6.727)
	Wino	76.09	— (6.286)	6.393
	Wino + EvoGRAD-M	76.09	— (6.286)	-
	Wino + EvoGRAD-L	76.64	-	6.652
ALBERT	EvoGRAD-M	-	55.4 (6.989)	-
	EvoGRAD-L	-	-	57.2 (6.853)
	Wino	64.64	— (7.971)	— (7.670)
	Wino + EvoGRAD-M	64.48	— (8.000)	-
	Wino + EvoGRAD-L	64.64	-	— (7.694)
GPT-3*	EvoGRAD-M	-	59.41 (7.122)	-
	EvoGRAD-L	-	-	56.08 (6.753)
GPT-3.5*	EvoGRAD-M	-	67.33 (7.061)	-
	EvoGRAD-L	-	-	65.02 (7.245)

Table 4: Accuracy (and error depth) results of models on Winogrande-valid and EvoGRAD-val sets after training on Winogrande-XL and/or EvoGRAD-train. **Bold** values represent the highest accuracy and underlined values represent the highest error depth for each model in each dataset. A single dash (-) denotes that the model was not tuned on that specific dataset variant, hence was not tested. Dashed (—) values indicate that accuracy was not tested due to potential contamination from EvoGRAD’s seed examples being taken from Winogrande, though error depth was still evaluated. Models marked with an asterisk (*) were evaluated using few-shot learning rather than fine-tuning.

BERT were selected from:

- Learning rates: $1e-5$, $3e-5$, $5e-5$
- Epochs: 3, 4, 5, 8
- Batch sizes: 8, 16

For training on EvoGRAD-train (both medium and large versions), given its resemblance but smaller size to Winogrande, we experimented with:

- Learning rates: $1e-5$, $3e-5$, $5e-5$
- Epochs: 1, 2, 4, 8
- Batch sizes: 8, 16, 32, 64

For evaluations using GPT-based models, we adopted a few-shot learning approach. Each instance was evaluated using an instruction-based prompt consisting of 30 random instances from the respective training set.

4.2. Results

Our evaluation results, as shown in Tables 4 and Figure 3, offer insight into model performance under different training conditions. We trained models exclusively on EvoGRAD-train, on Winogrande-XL (denoted as *Wino*), or sequentially on both

Winogrande and EvoGRAD-train (denoted as *Wino + EvoGRAD*). This approach allowed us to understand how different training datasets influence model robustness and stability.

Table 4 displays the models’ accuracies on the Winogrande-valid dataset alongside their average error depth on the EvoGRAD datasets. The error depth indicates the perturbative distance at which a model starts to fail, providing insights into model stability. While accuracy is the main metric, error depth (shown in parentheses) gives a complementary view of model performance. Due to the potential overlap between EvoGRAD and Winogrande, we have omitted the accuracy scores for Winogrande-trained models in EvoGRAD. GPT-based models were only evaluated on EvoGRAD instances as they are evaluated through few-shot learning.

Figure 3 visualizes the three most frequent perturbation types that lead to incorrect predictions by the models. Each perturbation is categorized by its effect on parts of speech. For instance, “+NN (150)” indicates a noun was added in 150 of the incorrect predictions. A comprehensive breakdown of the perturbation counts and their types, spanning all parts of speech observed, is provided in Table 5.




Figure 3: Distribution of the top three perturbations for models trained on Winogrande + EvoGrad-L. From left to right: BERT, RoBERTa, and ALBERT. The segments represent the relative frequency of each perturbation type: ‘-NN’ (noun removal), ‘+NN’ (noun addition), and either ‘-JJ’ (adjective removal) or ‘-IN’ (preposition removal).

Model	Trained on	EvoGRAD-M-val	EvoGRAD-L-val
BERT	EvoGRAD-M	+NN (150), -NN (148), -JJ (105)	-
	EvoGRAD-L	-	-NN (578), +NN (471), -JJ (342)
	Wino	-NN (92), -JJ (62), +NN (61)	-NN (365), +NN (294), -JJ (228)
	Wino + EvoGRAD-M	-NN (108), +NN (90), -JJ (78)	-
	Wino + EvoGRAD-L	-	-NN (373), +NN (303), -JJ (233)
RoBERTa	EvoGRAD-M	-NN (170), +NN (146), -JJ (120)	-
	EvoGRAD-L	-	-NN (494), +NN (416), -JJ (283)
	Wino	-NN (17), +NN (12), -JJ (11)	-NN (76), +NN (54), -JJ (41)
	Wino + EvoGRAD-M	-NN (17), +NN (12), -JJ (11)	-
	Wino + EvoGRAD-L	-	-NN (61), +NN (43), -IN (32)
ALBERT	EvoGRAD-M	-NN (189), +NN (161), -JJ (131)	-
	EvoGRAD-L	-	-NN (542), +NN (479), -JJ (316)
	Wino	-NN (92), -JJ (62), +NN (61)	-NN (272), +NN (208), -JJ (169)
	Wino + EvoGRAD-M	-NN (92), -JJ (62), +NN (61)	-
	Wino + EvoGRAD-L	-	-NN (294), +NN (220), -JJ (183)
GPT-3	EvoGRAD-M	-NN (173), +NN (144), -JJ (118)	-
	EvoGRAD-L	-	-NN (505), +NN (448), -JJ (306)
GPT-3.5	EvoGRAD-M	-NN (161), -JJ (115), +NN (111)	-
	EvoGRAD-L	-	-NN (464), +NN (364), -JJ (290)

Table 5: Top 3 perturbations and their count on incorrect predictions on EvoGRAD-val sets after fine-tuning on Winogrande-XL and EvoGRAD-train.

5. Discussion

Influence of EvoGRAD on Language Model Performance Table 4 illustrates the varied impacts of EvoGRAD on Transformer models, leading to several key insights:

- BERT’s improved performance post-EvoGRAD training underscores its ability to integrate the dataset’s specific perturbations effectively. This adaptability implies that BERT may be particularly effective for tasks requiring deeper linguistic insight or sensitivity to subtle contextual changes.
- RoBERTa consistently performs well both before and after training EvoGRAD, showcasing its robustness. However, its lower error depth

compared to its accuracy points to a potential trade-off between performance and stability. This observation underscores the need to balance generalization with stability to perturbations.

- The negligible change in ALBERT’s performance across various training regimes raises questions regarding the model’s saturation point and its alignment with the dataset. This warrants further investigation of the limits of adaptability for certain models.
- While GPT-based models, especially GPT-3.5, demonstrate competitive performance, their error depths highlight challenges related to stability. This trend suggests that some of the newer models might prioritize adaptability at

the expense of robustness.

Figure 3 sheds light on the areas where language models are most vulnerable, particularly in handling noun and adjective modifications. Addressing these specific challenges is imperative for the enhancement of common-sense reasoning in future model iterations.

Robustness and Adaptability to New Tasks

One of the challenges in deep learning is ensuring that the models remain adaptable and robust when exposed to new tasks or datasets. Whether through fine-tuning or few-shot learning, a model's ability to incorporate new information without significant detriment to its original capabilities is vital. In our experiments, the transformer models exhibited this adaptability, particularly when introduced to EvoGRAD. For instance, when models were fine-tuned on EvoGRAD, their performance on the Winogrande validation set generally improved or remained consistent (Table 4), indicating that they did not lose their grasp of previously acquired knowledge. However, GPT-based models, through few-shot learning, demonstrated their versatility in quickly adapting to new tasks without the need for extensive retraining. These observations underscore the potential of current architectures in handling evolving datasets and tasks, highlighting their robustness in diverse learning scenarios.

Evolution and Community Involvement with EvoGRAD

The current rendition of EvoGRAD represents only the first phase in a series of envisioned enhancements. As the platform matures, our goal is to achieve multiple cycles of data augmentation, model training, and fine-tuning, striving to foster a greater social impact in the AI domain. In making EvoGRAD accessible to a diverse audience, including those new to WSC-style challenges, we have incorporated clear prompts and guidelines, drawing inspiration from our initial work with the 182 instances in EvoGRAD-small.

Looking ahead, we are also planning to expand the platform to incorporate other foundational NLP tasks by integrating datasets such as OntoNotes 5.0 for Named Entity Recognition (NER) (Weischedel et al., 2012), Natural Questions (NQ) (Kwiatkowski et al., 2019) for Question Answering (QA), and the SemEval tasks for Sentiment Analysis, thereby broadening the scope and utility of EvoGRAD.

Recognizing the scale at which EvoGRAD could grow, we understand the crucial role of user-driven validation. While our dedicated team of in-house researchers currently curates the dataset to ensure its quality, we're eager to transition this role to our users in the near future. This strategy not only offloads the validation responsibility but

also promises a more dynamic, participatory, and community-centric approach to refining LLMs.

6. Conclusion

In this work, we introduced **EvoGRAD**, a dynamic platform that extends the Winograd Schema Challenge with a human-and-model-in-the-loop methodology. The dataset, enriched through our platform, incorporates contributions from human experts, language model expansions, and lexical resource utilization. We also introduced the "error depth" metric as a novel means to assess model stability in evolving tasks. While our evaluations showed potential benefits of using the augmented data from EvoGRAD across different training regimes, the disparity between human and machine performance on this task underlines its complexity and the ongoing challenges in enhancing common-sense reasoning in LLMs.

Ethics Statement

We are presenting our publicly-accessible platform to those outside the scientific and crowd-sourcing communities. However, our platform is still limited to those in society who have access to a mobile device/personal computer and internet access; a large but underrepresented group of people in the world do not. We therefore use our platform as only a first step towards more inclusiveness, which we open to people outside the small community of science and crowd-sourcing, but wish to be involved in efforts that will include the underrepresented mentioned.

We also cannot assume that everyone's foremost priorities involve contributing towards such endeavours as ours – many members of society are currently in turmoils of war, famine, or even indifference or aversion towards AI, which all may amount to their non-involvement in projects related to ours. Accordingly, the direction towards progress is best achieved *outside the laboratory*; after all, if diversity and community involvement in the development of tasks such as ours is as correlated to positive results in AI, our efforts as researchers should also extend towards the education, well being, and thriving of members in society, without which our goal of a global task is never truly realized.

Acknowledgements

This work was supported by the Fonds de recherche du Québec – Nature et technologies (FRQNT), the Natural Sciences and Engineering Research Council of Canada (NSERC), and the New Frontiers in Research Fund (NFRF).

7. Bibliographical 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. Association for Computational Linguistics.
- Alfred V. Aho and Jeffrey D. Ullman. 1972. *The Theory of Parsing, Translation and Compiling*, volume 1. Prentice-Hall, Englewood Cliffs, NJ.
- American Psychological Association. 1983. *Publications Manual*. American Psychological Association, Washington, DC.
- David P Anderson, Jeff Cobb, Eric Korpela, Matt Lebofsky, and Dan Werthimer. 2002. Seti@home: an experiment in public-resource computing. *Communications of the ACM*, 45(11):56–61.
- Rie Kubota Ando and Tong Zhang. 2005. [A framework for learning predictive structures from multiple tasks and unlabeled data](#). *Journal of Machine Learning Research*, 6:1817–1853.
- Galen Andrew and Jianfeng Gao. 2007. [Scalable training of \$L_1\$ -regularized log-linear models](#). In *Proceedings of the 24th International Conference on Machine Learning*, pages 33–40.
- Sriram Balasubramanian, Naman Jain, Gaurav Jindal, Abhijeet Awasthi, and Sunita Sarawagi. 2020. [What’s in a name? are BERT named entity representations just as good for any other name?](#) In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 205–214, Online. Association for Computational Linguistics.
- Toms Bergmanis, Katharina Kann, Hinrich Schütze, and Sharon Goldwater. 2017. [Training data augmentation for low-resource morphological inflection](#). In *Proceedings of the CoNLL SIGMORPHON 2017 Shared Task: Universal Morphological Reinflection*, pages 31–39, Vancouver. Association for Computational Linguistics.
- Edward Loper Bird, Steven and Ewan Klein. 2009. *Natural Language Processing with Python*. O’Reilly Media Inc.
- Daren C. Brabham. 2013. *Crowdsourcing*. MIT Press.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *Advances in neural information processing systems*, 33:1877–1901.
- Ashok K. Chandra, Dexter C. Kozen, and Larry J. Stockmeyer. 1981. [Alternation](#). *Journal of the Association for Computing Machinery*, 28(1):114–133.
- Michael Chen, Mike D’Arcy, Alisa Liu, Jared Fernandez, and Doug Downey. 2019. [CODAH: An adversarially-authored question answering dataset for common sense](#). In *Proceedings of the 3rd Workshop on Evaluating Vector Space Representations for NLP*, pages 63–69, Minneapolis, USA. Association for Computational Linguistics.
- Zhiyuan Chen and Bing Liu. 2018. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3):1–207.
- James W. Cooley and John W. Tukey. 1965. [An algorithm for the machine calculation of complex Fourier series](#). *Mathematics of Computation*, 19(90):297–301.
- Kevin Crowston. 2012. Amazon mechanical turk: A research tool for organizations and information systems scholars. In *Shaping the future of ict research. methods and approaches*, pages 210–221. Springer.
- Haixing Dai, Zhengliang Liu, Wenxiong Liao, Xiaoke Huang, Zihao Wu, Lin Zhao, Wei Liu, Ninghao Liu, Sheng Li, Dajiang Zhu, et al. 2023. Chataug: Leveraging chatgpt for text data augmentation. *arXiv preprint arXiv:2302.13007*.
- 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. Association for Computational Linguistics.
- Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. 2019. [Build it break it fix it for dialogue safety: Robustness from adversarial human attack](#). 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 4537–4546, Hong Kong, China. Association for Computational Linguistics.
- Yanai Elazar, Hongming Zhang, Yoav Goldberg, and Dan Roth. 2021a. [Back to square one: Artifact detection, training and commonsense disentanglement in the Winograd schema](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10486–10500, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yanai Elazar, Hongming Zhang, Yoav Goldberg, and Dan Roth. 2021b. [Back to square one: Artifact detection, training and commonsense disentanglement in the winograd schema](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10486–10500.
- Ali Emami, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2020. [An analysis of dataset overlap on winograd-style tasks](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5855–5865.
- 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. Association for Computational Linguistics.
- Christiane Fellbaum. 2010. Wordnet. In *Theory and applications of ontology: computer applications*, pages 231–243. Springer.
- Steven Y. Feng, Varun Gangal, Jason Wei, Sarath Chandar, Soroush Vosoughi, Teruko Mitamura, and Eduard Hovy. 2021. [A survey of data augmentation approaches for NLP](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 968–988, Online. Association for Computational Linguistics.
- Robert M French. 1999. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135.
- Alexander Gepperth and Barbara Hammer. 2016. Incremental learning algorithms and applications. In *European symposium on artificial neural networks (ESANN)*.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. [Chatgpt outperforms crowdworkers for text-annotation tasks](#). *arXiv preprint arXiv:2303.15056*.
- Dan Gusfield. 1997. *Algorithms on Strings, Trees and Sequences*. Cambridge University Press, Cambridge, UK.
- Weinan He, Canming Huang, Yongmei Liu, and Xiaodan Zhu. 2021. [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*, pages 3779–3789, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mokter Hossain and Ilkka Kauranen. 2015. [Crowdsourcing: a comprehensive literature review](#). *Strategic Outsourcing: An International Journal*.
- Fan Huang, Haewoon Kwak, and Jisun An. 2023. [Is chatgpt better than human annotators? potential and limitations of chatgpt in explaining implicit hate speech](#). In *Companion Proceedings of the ACM Web Conference 2023*, pages 294–297.
- Nicos Isaak and Loizos Michael. 2019. [Winoflexi: A crowdsourcing platform for the development of winograd schemas](#). In *AI 2019: Advances in Artificial Intelligence*, pages 289–302, Cham. Springer International Publishing.
- Nicos Isaak and Loizos Michael. 2020. [Blending nlp and machine learning for the development of winograd schemas](#). In *Agents and Artificial Intelligence: 12th International Conference, ICAART 2020, Valletta, Malta, February 22–24, 2020, Revised Selected Papers*, page 188–214, Berlin, Heidelberg. Springer-Verlag.
- Robin Jia and Percy Liang. 2017. [Adversarial examples for evaluating reading comprehension systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.
- Divyansh Kaushik, Eduard Hovy, and Zachary C Lipton. 2020. [Learning the difference that makes a difference with counterfactually augmented data](#). *International Conference on Learning Representations (ICLR)*.
- Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. 2021. [Dynabench: Rethinking](#)

- benchmarking in NLP. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4110–4124, Online. Association for Computational Linguistics.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2017. [Overcoming catastrophic forgetting in neural networks](#). *Proceedings of the National Academy of Sciences*, 114(13):3521–3526.
- 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. Association for Computational Linguistics.
- Vid Kocijan, Ernest Davis, Thomas Lukasiewicz, Gary Marcus, and Leora Morgenstern. 2022. The defeat of the winograd schema challenge. *arXiv preprint arXiv:2201.02387*.
- Bernhard Kratzwald and Stefan Feuerriegel. 2019. [Learning from on-line user feedback in neural question answering on the web](#). In *The World Wide Web Conference, WWW '19*, page 906–916, New York, NY, USA. Association for Computing Machinery.
- Taja Kuzman, Igor Mozetic, and Nikola Ljubešić. 2023. Chatgpt: Beginning of an end of manual linguistic data annotation? use case of automatic genre identification. *ArXiv, abs/2303.03953*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Wuwei Lan, Siyu Qiu, Hua He, and Wei Xu. 2017. [A continuously growing dataset of sentential phrases](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1224–1234, Copenhagen, Denmark. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.
- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew Peters, Ashish Sabharwal, and Yejin Choi. 2020. Adversarial filters of dataset biases. In *International Conference on Machine Learning*, pages 1078–1088. PMLR.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2011. The winograd schema challenge. In *AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning*.
- Bill Yuchen Lin, Seyeon Lee, Rahul Khanna, and Xiang Ren. 2020a. [Birds have four legs?! NumerSense: Probing Numerical Commonsense Knowledge of Pre-Trained Language Models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6862–6868, Online. Association for Computational Linguistics.
- Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. 2020b. [CommonGen: A constrained text generation challenge for generative commonsense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1823–1840, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Jesús Lovón-Melgarejo, Laure Soulier, Karen Pinel-Sauvagnat, and Lynda Tamine. 2021. Studying catastrophic forgetting in neural ranking models. In *European Conference on Information Retrieval*, pages 375–390. Springer.
- Jinghui Lu, Linyi Yang, Brian Namee, and Yue Zhang. 2022. [A rationale-centric framework for human-in-the-loop machine learning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6986–6996, Dublin, Ireland. Association for Computational Linguistics.
- Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. 2020. Gender bias in neural natural language processing. In *Logic, Language, and Security*, pages 189–202. Springer.
- Gabriela Melo, Vinicius Imaizumi, and Fábio Cozman. 2019. [Winograd schemas in portuguese](#). In

- Proceedings of 16th National Meeting on Artificial and Computational Intelligence*, pages 787–798, Porto Alegre, RS, Brasil. SBC.
- Tom Mitchell, William Cohen, Estevam Hruschka, Partha Talukdar, Justin Betteridge, Andrew Carlson, Bhavana Dalvi Mishra, Matthew Gardner, Bryan Kisiel, Jayant Krishnamurthy, Ni Lao, Kathryn Mazaitis, Thahir Mohamed, Ndapa Nakashole, Emmanouil Platanios, Alan Ritter, Mehdi Samadi, Burr Settles, Richard Wang, Derry Wijaya, Abhinav Gupta, Xinlei Chen, Abulhair Saparov, Malcolm Greaves, and Joel Welling. 2015. [Never-ending learning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 29(1).
- Leora Morgenstern, Ernest Davis, and Charles L. Ortiz. 2016. [Planning, executing, and evaluating the winograd schema challenge](#). *AI Magazine*, 37(1):50–54.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Mohammad Sadegh Rasooli and Joel R. Tetreault. 2015. [Yara parser: A fast and accurate dependency parser](#). *Computing Research Repository*, arXiv:1503.06733. Version 2.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. [Gender bias in coreference resolution](#). 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 8–14, New Orleans, Louisiana. Association for Computational Linguistics.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Winogrande: An adversarial winograd schema challenge at scale. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8732–8740.
- Haoyue Shi, Karen Livescu, and Kevin Gimpel. 2021. [Substructure substitution: Structured data augmentation for NLP](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3494–3508, Online. Association for Computational Linguistics.
- Haoyue Shi, Jiayuan Mao, Tete Xiao, Yuning Jiang, and Jian Sun. 2018. [Learning visually-grounded semantics from contrastive adversarial samples](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3715–3727, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Daniel L Silver, Qiang Yang, and Lianghao Li. 2013. Lifelong machine learning systems: Beyond learning algorithms. In *2013 AAAI spring symposium series*.
- Paul Trichelair, Ali Emami, Jackie Chi Kit Cheung, Adam Trischler, Kaheer Suleman, and Fernando Diaz. 2018. [On the evaluation of common-sense reasoning in natural language understanding](#). In *Critiquing and Correcting Trends in Machine Learning NeurIPS 2018 Workshop*.
- 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. Association for Computational Linguistics.
- Alan M. Turing. 1950. Computing machinery and intelligence. *Mind*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Eric Wallace, Pedro Rodriguez, Shi Feng, Ikuya Yamada, and Jordan Boyd-Graber. 2019. [Trick me if you can: Human-in-the-loop generation of adversarial examples for question answering](#). *Transactions of the Association for Computational Linguistics*, 7:387–401.
- 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. Association for Computational Linguistics.
- William Yang Wang and Diyi Yang. 2015. [That's so annoying!!!: A lexical and frame-semantic embedding based data augmentation approach to automatic categorization of annoying behaviors using #petpeeve tweets](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2557–2563,

- Lisbon, Portugal. Association for Computational Linguistics.
- R Weischedel, S Pradhan, L Ramshaw, J Kaufman, M Franchini, M El-Bachouti, N Xue, M Palmer, JD Hwang, C Bonial, et al. 2012. Ontonotes release 5.0. linguistic data consortium. Technical report, Philadelphia, Technical Report.
- Terry Winograd. 1972. [Understanding natural language](#). *Cognitive Psychology*, 3(1):1–191.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Jing Xu, Da Ju, Margaret Li, Y-Lan Boureau, Jason Weston, and Emily Dinan. 2021. [Bot-adversarial dialogue for safe conversational agents](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2950–2968, Online. Association for Computational Linguistics.
- Yan Xu, Ran Jia, Lili Mou, Ge Li, Yunchuan Chen, Yangyang Lu, and Zhi Jin. 2016. [Improved relation classification by deep recurrent neural networks with data augmentation](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1461–1470, Osaka, Japan. The COLING 2016 Organizing Committee.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Pardis Zahraei and Ali Emami. 2024. [WSC+: Enhancing the Winograd schema challenge using tree-of-experts](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1650–1671, St. Julian’s, Malta. Association for Computational Linguistics.
- 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. 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. Association for Computational Linguistics.
- Hongming Zhang, Xinran Zhao, and Yangqiu Song. 2020. [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*, pages 5736–5745, Online. Association for Computational Linguistics.
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. 2018. [mixup: Beyond empirical risk minimization](#). In *International Conference on Learning Representations*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. [Gender bias in coreference resolution: Evaluation and debiasing methods](#). 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 15–20, New Orleans, Louisiana. Association for Computational Linguistics.
- Yiming Zhu, Peixian Zhang, Ehsan-UI Haq, Pan Hui, and Gareth Tyson. 2023. [Can chatgpt reproduce human-generated labels? a study of social computing tasks](#). *arXiv preprint arXiv:2304.10145*.
- Ran Zmigrod, Sabrina J. Mielke, Hanna Wallach, and Ryan Cotterell. 2019. [Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1651–1661, Florence, Italy. Association for Computational Linguistics.

A. Appendix

A.1. ChatGPT Dialogue During Scaling Process

Author: You are an expert linguist, human annotator and are familiar with the Winograd Schema Challenge. Consider the sentence: Although they ran at about the same speed, Sue beat Sally because _ had such a good start. This sentence is in the format of a Winograd schema, where two antecedents, in this case "Sue" and "Sally" are competing for reference in the position, _. Common sense allows us to conclude that one of the two antecedents is most likely, in this case, Sue.

I am intending on building a dataset composed of modifications to the above sentence. Based on this original sentence, I need you to generate 100 new sentences that are modified by either adding, replacing, or substituting words in the original sentence. They should be organized in of the following format, copied from a csv file :

index sentence option1 option2 answer distance

where the original sentence has the following form: 0 Although they ran at about the same speed, Sue beat Sally because _ had such a good start. Sue Sally 1 0

Here, index refers to the index number for the sentence, sentence refers to the sentence itself (notice the necessity of the _ as the position for one of the two antecedents), option1 refers to one antecedent, option2 refers to the other, answer refers to the correct antecedent that belongs in _ – "1" corresponds to option1, and "2" corresponds to option2. Notice, for this original sentence, the answer is "1" because having a good start is a likely explanation for someone (i.e., Sue) beating another in a race. Finally, "distance" corresponds to the number of modifications made on the original sentence (whose distance is, by definition, 0). It is equivalent to the "token edit distance".

I want you to create a diverse range of new entries, that vary in terms of the modifications, the total distance, as well as the correct answer.

For example, the sentence "Although they sprinted at about the same speed, Sue beat Sally because _ had such a good start." would have an answer of 1 and distance of 1 (because ran was changed to sprinted, but the correct answer still remains option1, which in this case is Sue), and the sentence "Although they sprinted at about the same speed, Sue beat Sally although _ had such a good start." would have the answer of 2 and distance of 2, because the correct answer is now "Sally", and there were two word modifications – the second one here was "although" instead of "because", which flips the answer. Notice that with the right change in words, the answer can flip – I want to encourage you to find such minimal number of changes when you're creating new entries. You are also allowed to change the words in the options too (e.g., Sue and Sally). I would like the dataset to have a similar number of answer=1 entries than answer=2 entries.

Please generate the first 10, formatted as a table here. First, you may ask any clarification questions to make sure you understand the task.

ChatGPT: Thank you for the instructions. Here are the first 10 entries: [...]

Author: Excellent work. Please generate 30 more, but try to keep the distance lower. Don't forget you can change the names too!

ChatGPT: Sure! Here are the next 30 entries with smaller distances and changes in names: [...]

Author: These sentences didn't actually have the distances you claimed they had. These should correspond to token edit distances. Please be careful when generating sentences to actually correspond to distances – it's okay if the distance is a bit higher. I just want you to be accurate. You also forgot to change the names (as possible modifications) at one point. Please generate 10 for me to see if it's okay.

ChatGPT: I apologize for the inaccuracies in the distances and missing name changes in the previous examples. [...] Here are the next 10 sentences: [...]

Author: These were much better. Please generate 30 more. Pay attention to not miscalculate the distance.

ChatGPT: I appreciate your feedback. Here are the next 30 sentences, with the correct token edit distances and a careful attention to detail: [...]

Author: Great job. Please generate 30 more. Again, the token distance should be correctly calculated. For example, in your sentence:

Even though they raced at the same speed, Sue beat Sally although _ had a powerful start.

You said the distance is 3. It's actually 5 here, because Although was dropped and replaced with Even (1), then "though" was added (2), then raced was replaced by ran (3), then about was dropped (4), then good was replaced by powerful (5).

ChatGPT: I understand your concern and I apologize for the mistake. I'll ensure more accuracy in calculating the token edit distance. Here are the next 30 sentences: [...]

Author: Great job. Please generate 20 more. Remember that changing the names is valid but counts as a modification (increments distance).

ChatGPT: Thank you for your feedback. Here are the next 20 sentences, this time including changes in names: [...]