

Towards Zero-shot Commonsense Reasoning with Self-supervised Refinement of Language Models

Tassilo Klein
SAP AI Research
tassilo.klein@sap.com

Moin Nabi
SAP AI Research
m.nabi@sap.com

Abstract

Can we get existing language models and refine them for zero-shot commonsense reasoning? This paper presents an initial study exploring the feasibility of zero-shot commonsense reasoning for the Winograd Schema Challenge by formulating the task as self-supervised refinement of a pre-trained language model. In contrast to previous studies that rely on fine-tuning annotated datasets, we seek to boost conceptualization via loss landscape refinement. To this end, we propose a novel self-supervised learning approach that refines the language model utilizing a set of linguistic perturbations of similar concept relationships. Empirical analysis of our conceptually simple framework demonstrates the viability of zero-shot commonsense reasoning on multiple benchmarks.¹

1 Introduction

Natural language processing has recently experienced unprecedented progress, boosting the performance of many applications to new levels. However, this gain in performance does not equally transfer to applications requiring commonsense reasoning capabilities, which has largely remained an unsolved problem (Marcus, 2020; Kocijan et al., 2020). In order to assess the commonsense reasoning capabilities of automatic systems, several tasks have been devised. Among them is the popular Winograd Schema Challenge (WSC), which frames commonsense reasoning as a pronoun resolution task (Levesque et al., 2012). Although appearing evident and natural to the human mind, modern machine learning methods still struggle to solve this challenge.

Lately, the research community has experienced an abundance of methods proposing utilization of

¹The source code can be found at: <https://github.com/SAP-samples/emnlp2021-contrastive-refinement/>

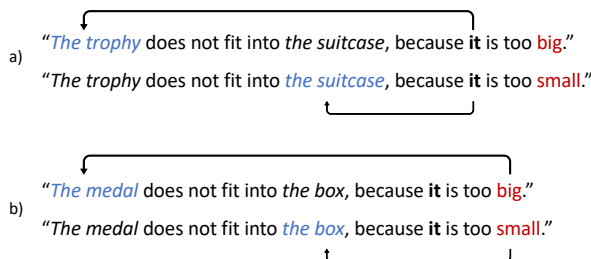


Figure 1: WSC sample: a) original sentence, b) perturbation (noun synonym). Task: resolve **pronoun** with a **candidate**. The **trigger-word** induces an answer flip.

language models (LM) to tackle commonsense reasoning in a two-stage learning pipeline. Starting from an initial self-supervised learned model, commonsense enhanced LMs are obtained in a subsequent fine-tuning (ft) phase. Fine-tuning enforces the LM to solve the downstream WSC task as a plain co-reference resolution task. However, such supervised approaches are prone to leverage statistical data artifacts for reasoning, giving rise to the “Clever Hans” effect (Lapuschkin et al., 2019). As such, instead of truly featuring reasoning capabilities, approaches become very good in faking. On the other hand, the lack of commonsense reasoning capabilities of LMs can be partially attributed to the training corpora itself, as the commonsense knowledge is often not incorporated into the training text due to the assumed triviality (Trichelair et al., 2018; Saba, 2018; Trichelair et al., 2019; Emami et al., 2019; Kavumba et al., 2019; Liu et al., 2020; Cui et al., 2020). We hypothesize that the current self-supervised tasks used in the pre-training phase are insufficient to enforce the model to generalize commonsense concepts (Kejriwal and Shen, 2020). This shortcoming is easily unveiled by the susceptibility of LM to semantic variations. In this regard, it has been shown that LMs are sensitive to linguistic perturbations (Abdou et al., 2020). A case in point is the WSC example in Fig. 1. It shows a pair of sentences sub-

ject to semantic variations establishing the same relationship between entities. This can be defined as the joint concept triplet consisting of two nouns and a verb that determines the relationship between the nouns, e.g., (*container*, *item*, *fit*). Inappropriate semantic sensitivity to semantic variants leads to inadequate “conceptualization” and misconstruction of such triplets. To address this, we propose *self-supervised refinement*. It seeks to achieve generalization through a task agnostic objective.

To this end, we tackle the problem of commonsense reasoning from a zero-shot learning perspective. Leveraging zero-shot models to gauge the intrinsic incorporation of commonsense knowledge suggests being the more valid approach than fine-tuned models. That can be attributed to the exploitation of implicit biases less likely to occur in this setup. Hence, the associated benchmarks constitute a more realistic and reliable benchmark (Elazar et al., 2021). Other zero-shot methods for commonsense reasoning either use large supervised datasets Winogrande (Sakaguchi et al., 2019) or very large LMs such as GPT-3 (Brown et al., 2020). In contrast, the proposed method takes a pre-trained LM as input, which undergoes a refinement step. During refinement, the LM is exposed to semantic variations, aiming at improved concept generalization by making the model more robust w.r.t. perturbations. Motivated by the recent advancements in contrastive representation learning (Chen et al., 2020; He et al., 2020; Jean-Bastien et al., 2020; Klein and Nabi, 2020), we propose refining the LM in a self-supervised contrastive fashion. This entails refinement without the use of any labels and hence with *no gradient update* on the downstream datasets. Consequently, the supervision level is identical to the test time of the Winograd schema challenge.

Our contributions are two-fold: **(i)** we introduce the task of *zero-shot commonsense reasoning* for WSC by reformulating the task as performing self-supervised refinement on a pre-trained language model **(ii)** We propose a self-supervised refinement framework which leverages semantic perturbations to facilitate zero-shot commonsense reasoning.

2 Method

Preliminaries: Transformer-based LMs are based on an encoder-decoder architecture, consisting of a set of encoding layers that process the input it-

eratively. Prior to entering the Transformer stack, the input is pre-processed by a tokenizer that turns the input sentence into a sequence of tokens. Besides tokens arising from the input sentence, there are also auxiliary tokens such as [CLS], [SEP]. In BERT and RoBERTa, these tokens delimit the input from padding for fixed-length sequence processing. Furthermore, there are special tokens that are tailored to frame specific tasks. For example, [MASK] is used to mask out words for learning the masked language model. Instantiation of language models on the tokenized sequence yields a sequence of embedding vectors. To avoid clutter in the notation and subsuming the fact that only fixed-length sequences are encoded, for the following $x \in \mathbb{T}$ will refer to the tensor obtained by stacking the sequence of token embeddings.

2.1 Perturbation Generation Framework

Starting from a pre-trained LM (*init-LM*), we conduct a refinement step exposing the model to semantic variations of Winograd schemas. Given a sentence x and a specific semantic [perturbation token], the LM is trained to generate the embedding \hat{x} of the provided perturbation type. We enforce the generator to estimate the embedding obtained by the LM on the sentence with the actual semantic perturbation as the target. Intuitively speaking, an LM that generates perturbed representations from an unperturbed input is equipped with a generalized view of commonsense concepts. This builds upon the idea that the injection of noise to the input can flatten the loss landscape to promote generalization (Qin et al., 2019; Moosavi-Dezfooli et al., 2019).

To this end, we extend the set of auxiliary tokens with some new tokens referred to as “perturbation tokens”. In the course of training, the perturbation tokens are prepended to the input sentence directly after the [CLS] token. For the following, we let \mathcal{P} denote the set of semantic perturbations. Besides perturbations, \mathcal{P} also includes the identity transformation [IDENTICAL], which implies no semantic change. Figure 1 shows an example of a perturbation induced by the perturbation token [SYNONYM], which entails replacing nouns of the input sentence with synonyms. Following the example from the figure, the LM seeks to map the representation of the (tokenized) sentence **(a)** in conjunction with [SYNONYM] to the representation of **(b)**. To enforce consistency across com-




Figure 2: Schematic illustration of the proposed approach. Two examples x_i and x_j from the WSC dataset, both demonstrating the concept triplet (container, item, fit) and their generated embeddings (dashed outline) for two perturbation types: **top**: [SYNOYM] and **bottom**: [TENSE]. Loss terms defined as *attraction* ($\leftarrow\rightarrow$) and *repulsion* ($\rightarrow\leftarrow$) between embeddings of unperturbed and corresponding generated perturbation, each shown in a different color: **Reconstruction** loss, **Contrastive** loss and **Diversity** loss (best shown in color).

nonsense concepts and semantic perturbations, we embed learning in a contrastive setting.

2.2 Self-supervised Refinement

The method’s core idea is to construct an abstract, generic view of a commonsense concept by exploiting slightly different examples of the same concept (i.e., perturbations). This is achieved by joint optimization of a LM w.r.t. three different loss terms (Reconstruction, Contrastive and Diversity):

$$\min_{\theta_1, \theta_2} \mathcal{L}_R(f_{\theta_1}) + \mathcal{L}_C(f_{\theta_1}) + \mathcal{L}_D(f_{\theta_1}, q_{\theta_2}) \quad (1)$$

Here f denotes the LM, e.g., BERT or RoBERTa parameterized by θ_1 , and $q : \mathbb{T} \rightarrow \mathcal{P}$ denotes a representation discriminator (MLP) parameterized by θ_2 . The functionality of the individual loss terms of Eq. 1 will be explained in the following subsections. Additionally, Fig. 2 shows a schematic illustration of the proposed approach and each loss term.

Optimization of Eq. 1 entails computation of similarities between embeddings, employing a metric $\phi(x, \hat{x}) : \mathbb{T} \times \mathbb{T} \rightarrow \mathbb{R}$. Here, we employ a variant of the BERTscore (Zhang et al., 2020) as a similarity metric. BERTscore computes sentence similarities by matching tokens based on their cosine similarity. Subsequently, the scores for the entire sequence are aggregated. Unlike the original BERTscore, we restrict token matching to each token’s vicinity to accommodate that perturbations typically induce changes only in a small neighborhood. To this end, we restrict token matching by applying a sliding window mechanism centered on each token.

2.2.1 Reconstruction loss

The reconstruction loss’s objective is to regress embeddings by minimizing the distance between the ground-truth and the approximated “perturbation” embedding:

$$\mathcal{L}_R = -\alpha \sum_i^N \sum_{k \in \mathcal{P}} \phi(x_i^{[k]}, \hat{x}_i^{[k]}) \quad (2)$$

2.2.2 Contrastive loss

The objective of the contrastive loss is to preserve the “semantic expressivity of individual samples and prevent the collapse to a singular perturbation representation. This is achieved by pushing apart the embeddings for *different samples of the same perturbation* type.

$$\mathcal{L}_C = \beta \sum_{i, j: i \neq j}^N \sum_{k \in \mathcal{P}} \phi(\hat{x}_i^{[k]}, \hat{x}_j^{[k]}) \quad (3)$$

2.2.3 Diversity loss

The diversity loss term aims to guarantee the discriminativeness of the perturbation embeddings arising from the same sample. As such, it imposes the semantic perturbations for the same sample to be diverse, preventing the collapse of different perturbations to a single embedding. Maximizing diversity entails minimization of cross-entropy w.r.t. perturbations:

$$\mathcal{L}_D = -\gamma \sum_i^N \sum_{k \in \mathcal{P}} \log \frac{q(k | \hat{x}_i^{[k]})}{\sum_{\forall t \in \mathcal{P}: t \neq k} q(t | \hat{x}_i^{[k]})}, \quad (4)$$

Here $q(\cdot) : \mathbb{T} \rightarrow \mathbb{R}$ denotes the likelihood of a classifier w.r.t. embeddings. N denotes the

number of data samples, and $\alpha, \beta, \gamma \in \mathbb{R}$ denote the hyperparameters, balancing the terms in the loss function.

2.2.4 Zero-shot Pronoun Disambiguation

For resolving the WSC we leverage the Transformer Masked Token Prediction following (Kocijan et al., 2019). This entails replacing the [MASK] token with the possible candidates. Given an associated pair of training sentences with $i \in N$, i.e., (s_i^1, s_i^2) , the difference between the sentence pairs is the trigger word(s). Here c_1, c_2 denote the answer candidates, yielding probabilities for the candidates: $p(c_1|s_i^1)$ and $p(c_2|s_i^1)$. The answer prediction corresponds to the candidate with a more significant likelihood. If a candidate consists of several tokens, the probability corresponds to the average of the log probabilities.

3 Experiments and Results

3.1 Setup

We approach training the language model by first training the LM on perturbations on the enhanced-WSC corpus (Abdou et al., 2020). It is a perturbation augmented version of the original WSC dataset. It consists of 285 sample sentences, with up to 10 semantic perturbations per sample. We make use of the following 7 perturbations: tense switch [TENSE], number switch [NUMBER], gender switch [GENDER], voice switch (active to passive or vice versa) [VOICE], relative clause insertion (a relative clause is inserted after the first referent) [RELCLAUSE], adverbial qualification (an adverb is inserted to qualify the main verb of each instance) [ADVERB], synonym/name substitution [SYNONYM].

3.2 Architecture

The proposed approach is applicable to any Transformer architecture. Here, we adopted standard LMs such as BERT and RoBERTa for comparability, without aiming to optimize the results for any downstream dataset/benchmark. Specifically, we employ the Hugging Face (Wolf et al., 2019) implementation of BERT *large-uncased* architecture as well as RoBERTa *large*. The LM is trained for 10 epochs for BERT and 5 for RoBERTa, using a batch size of 10 sentence samples. Each sample was associated with 4 perturbation, yielding an effective batch size of 40. For optimization, we

DPR (Rahman and Ng, 2012)		
Method	BERT	RoBERTa
Baseline (init-LM)	58.50 %	70.39 %
Ours (Zero-shot)	61.35 %	76.95 %
GAP (Webster et al., 2018)		
Method	BERT	RoBERTa
Baseline (init-LM)	58.70 %	58.87 %
Ours (Zero-shot)	58.73 %	59.13 %
KnowRef (Emami et al., 2019)		
Method	BERT	RoBERTa
Baseline (init-LM)	62.36 %	60.42 %
Ours (Zero-shot)	62.44 %	63.97 %
PDP-60 (Davis et al., 2016)		
Method	BERT	RoBERTa
Baseline (init-LM)	60.00 %	50.00 %
Ours (Zero-shot)	58.33 %	55.00 %
WSC-273 (Levesque et al., 2012)		
Method	BERT	RoBERTa
Baseline (init-LM)	62.64 %	67.77 %
Ours (Zero-shot)	61.54 %	71.79 %
WinoGender (Rudinger et al., 2018)		
Method	BERT	RoBERTa
Baseline (init-LM)	62.50 %	61.67 %
Ours (Zero-shot)	62.08 %	69.17 %
WinoGrande (Sakaguchi et al., 2019)		
Method	BERT	RoBERTa
Baseline (init-LM)	51.70 %	53.75 %
Ours (Zero-shot)	52.33 %	55.01 %
WinoBias Anti (Zhao et al., 2018)		
Method	BERT	RoBERTa
Baseline (init-LM)	56.82 %	55.93 %
Ours (Zero-shot)	56.82 %	60.61 %
WinoBias Pro (Zhao et al., 2018)		
Method	BERT	RoBERTa
Baseline (init-LM)	68.43 %	68.43 %
Ours (Zero-shot)	75.12 %	75.76 %

Table 1: Results for zero-shot commonsense reasoning

used a typical setup of AdamW with 500 warmup steps, a learning rate of $5.0 \cdot 10^{-5}$ with $\epsilon = 1.0 \cdot 10^{-8}$ and $\epsilon = 1.0 \cdot 10^{-5}$ for BERT and RoBERTa, respectively. For training BERT, we used $\alpha = 130$, $\beta = 0.5$, $\gamma = 2.5$, for RoBERTa $\alpha = 1.25$, $\beta = 7.25$, $\gamma = 6.255$. For hyperparameter optimization of

α, β, γ we follow a standard greedy heuristic, leveraging a weighted-sum optimization scheme (Jakob and Blume, 2014). From an initial candidate solution set, coarse-grid random search is utilized to explore the neighborhood on a fine grid of a randomly selected candidates.

The representation discriminator q is a MLP consisting of two fully connected layers with Batch-Norm, parametric ReLU (PReLU) activation function and 20% Dropout.

3.3 Results

Given the BERT and RoBERTa language models for comparison, the baseline constitute the initial-LM prior to undergoing refinement. We evaluated our method on nine different benchmarks. Results are reported in Tab. 1. Accuracy gains are significant and consistent with RoBERTa across all benchmarks. On average, the proposed approach increases the accuracy of (+0.8%) with BERT and of (+4.5%) with RoBERTa. The benchmarks and the results are discussed below:

DPR (Rahman and Ng, 2012): a pronoun disambiguation benchmark resembling WSC-273, yet significantly larger. According to (Trichelair et al., 2018), less challenging due to inherent biases. Here the proposed approach outperforms the baseline for both BERT and RoBERTa by a margin of (+2.85%) and (+6.56%), respectively.

GAP (Webster et al., 2018): a gender-balanced co-reference corpus. The proposed approach outperforms the baseline on BERT and RoBERTa with (+0.08%) and (+0.26%).

KnowRef (Emami et al., 2019): a co-reference corpus addressing gender and number bias. The proposed approach outperforms the baseline on BERT and RoBERTa with (+0.08%) and (+3.55%).

PDP-60 (Davis et al., 2016): pronoun disambiguation dataset. Our method outperforms the baseline with RoBERTa with (+5.0%), while on BERT showing a drop of (-1.67%).

WSC-273 (Levesque et al., 2012): a pronoun disambiguation benchmark, known to be more challenging than PDP-60. Our method outperforms the baseline with RoBERTa with (+4.0%), with a drop of (-1.1%) with BERT.

WinoGender (Rudinger et al., 2018): a gender-balanced co-reference corpus. The proposed approach outperforms the baseline on RoBERTa by (+7.6%), observing a drop on BERT (-0.42%).

WinoGrande (W.G.) (Sakaguchi et al., 2019): the

Method	DPR	W.G.
RoBERTa (Liu et al., 2019)	70.39	53.75
Ours ($\mathcal{L}_C + \mathcal{L}_D$)	73.76	53.28
Ours ($\mathcal{L}_R + \mathcal{L}_D$)	65.60	53.59
Ours ($\mathcal{L}_R + \mathcal{L}_C$)	65.07	52.01
Ours ($\mathcal{L}_R + \mathcal{L}_C + \mathcal{L}_D$)	76.95	55.01

Table 2: Ablation study, performance in accuracy (%)

largest dataset for Winograd commonsense reasoning. Our method outperforms the baseline with BERT by (+0.63) and with RoBERTa by (+1.26%).

WinoBias (Rudinger et al., 2018): a gender-balanced co-reference corpus consisting of two types. Type-1 requiring world knowledge, Type-2 requiring syntactic understanding. While on par for the first type in combination with BERT and a margin of (+6.69%), the proposed approach outperforms the baseline with RoBERTa with (+4.68) and (+7.33).

3.3.1 Ablation Study

To assess each loss term’s contribution, we evaluated each component’s performance by removing them individually from the loss. It should be noted that \mathcal{L}_C of Eq. 3 and \mathcal{L}_D of Eq. 4 both interact in a competitive fashion. Hence, only the equilibrium of these terms yields an optimal solution. Changes - such as eliminating a term - have detrimental effects, as they prevent achieving such an equilibrium, resulting in a significant drop in performance. See Tab. 2 for the ablation study on two benchmarks. Best performance is achieved in the presence of all loss terms.

4 Discussion and Conclusion

We introduced a method for self-supervised refinement of LMs. Its conceptual simplicity facilitates generic integration into frameworks tackling commonsense reasoning. A first empirical analysis on multiple benchmarks indicates that the proposed approach consistently outperforming the baselines in terms of standard pre-trained LMs, confirming the fundamental viability. We believe that the performance gain will be more pronounced when leveraging larger perturbation datasets for LM refinement. Hence, future work will focus on the generation of perturbations. This could specifically entail the consideration of sample-specific perturbations.

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.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher 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](#).
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*.
- Leyang Cui, Sijie Cheng, Yu Wu, and Yue Zhang. 2020. Does bert solve commonsense task via commonsense knowledge? *arXiv preprint arXiv:2008.03945*.
- Ernest Davis, Leora Morgenstern, and Charles Ortiz. 2016. Human tests of materials for the winograd schema challenge 2016.
- Yanai Elazar, Hongming Zhang, Yoav Goldberg, and Dan Roth. 2021. [Back to square one: Bias detection, training and commonsense disentanglement in the winograd schema](#).
- 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.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wilfried Jakob and Christian Blume. 2014. [Pareto optimization or cascaded weighted sum: A comparison of concepts](#). *Algorithms*, 7(1):166–185.
- Grill Jean-Bastien, Strub Florian, Alché Florent, Tallec Corentin, Pierre Richemond H., Buchatskaya Elena, Doersch Carl, Bernardo Pires Avila, Zhao-han Guo Daniel, Mohammad Azar Gheshlaghi, Piot Bilal, Kavukcuoglu Koray, Munos Rémi, and Valko Michal. 2020. Bootstrap your own latent - a new approach to self-supervised learning. *NIPS 2020*.
- Pride Kavumba, Naoya Inoue, Benjamin Heinzerling, Keshav Singh, Paul Reiser, and Kentaro Inui. 2019. In *Proceedings of the First Workshop on Commonsense Inference in Natural Language Processing*, pages 33–42, Hong Kong, China. Association for Computational Linguistics. [\[link\]](#).
- Mayank Kejriwal and Ke Shen. 2020. [Do fine-tuned commonsense language models really generalize?](#)
- Tassilo Klein and Moin Nabi. 2020. Contrastive self-supervised learning for commonsense reasoning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7517–7523.
- Vid Kocijan, Ana-Maria Cretu, Oana-Maria Camburu, Yordan Yordanov, and Thomas Lukasiewicz. 2019. A surprisingly robust trick for winograd schema challenge. In *The 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, Florence, Italy.
- Vid Kocijan, Thomas Lukasiewicz, Ernest Davis, Gary Marcus, and Leora Morgenstern. 2020. A review of winograd schema challenge datasets and approaches. *arXiv preprint arXiv:2004.13831*.
- Sebastian Lapuschkin, Stephan Wäldchen, Alexander Binder, Grégoire Montavon, Wojciech Samek, and Klaus-Robert Müller. 2019. [Unmasking clever hans predictors and assessing what machines really learn](#). *Nature Communications*, 10:1096.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*.
- Haokun Liu, William Huang, Dhara Mungra, and Samuel Bowman. 2020. Precise task formalization matters in winograd schema evaluations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8275–8280.
- 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](#).
- Gary Marcus. 2020. The next decade in ai: four steps towards robust artificial intelligence. *arXiv preprint arXiv:2002.06177*.
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Jonathan Uesato, and Pascal Frossard. 2019. Robustness via curvature regularization, and vice versa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

- Chongli Qin, James Martens, Sven Gowal, Dilip Krishnan, Krishnamurthy Dvijotham, Alhussein Fawzi, Soham De, Robert Stanforth, and Pushmeet Kohli. 2019. [Adversarial robustness through local linearization](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Altat Rahman and Vincent Ng. 2012. [Resolving complex cases of definite pronouns: The Winograd schema challenge](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 777–789, Jeju Island, Korea. Association for Computational Linguistics.
- 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.
- Walid S. Saba. 2018. [A simple machine learning method for commonsense reasoning? A short commentary on trinh & le \(2018\)](#). *CoRR*, abs/1810.00521.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [WINOGRANDE: an adversarial winograd schema challenge at scale](#). *CoRR*, abs/1907.10641.
- 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](#). *CoRR*, abs/1811.01778.
- 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.
- Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. 2018. [Mind the GAP: A balanced corpus of gendered ambiguous pronouns](#). *Transactions of the Association for Computational Linguistics*, 6:605–617.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *ArXiv*, abs/1910.03771.
- Varsha Kishore Zhang, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- 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.