

Unsupervised Pronoun Resolution via Masked Noun-Phrase Prediction

Ming Shen* Pratyay Banerjee* Chitta Baral

Arizona State University

mshen16, pbanerj6, chitta@asu.edu

Abstract

In this work, we propose Masked Noun-Phrase Prediction (MNPP), a pre-training strategy to tackle pronoun resolution in a fully unsupervised setting. Firstly, We evaluate our pre-trained model on various pronoun resolution datasets without any finetuning. Our method outperforms all previous unsupervised methods on all datasets by large margins. Secondly, we proceed to a few-shot setting where we finetune our pre-trained model on WinoGrande-S and XS separately. Our method outperforms RoBERTa-large baseline with large margins, meanwhile, achieving a higher AUC score after further finetuning on the remaining three official splits of WinoGrande.

1 Introduction

Co-reference resolution is an important NLP task that aims to find all expressions that refer to the same entity in a text. The resolution of an ambiguous pronoun, known as pronoun resolution, is a longstanding challenge for the NLU community and an essential step for various high-level NLP tasks such as natural language inference (Bowman et al., 2015; Williams et al., 2018), question answering (Rajpurkar et al., 2016), and relation extraction (Zhang et al., 2017).

The most successful approach to pronoun resolution is first fine-tuning a large pre-trained language model such as BERT (Devlin et al., 2019) or RoBERTa (Liu et al., 2019) on a human-labeled pronoun resolution dataset such as Definite Pronoun Resolution Dataset (DPR) (Rahman and Ng, 2012) or WinoGrande (WG) (Sakaguchi et al., 2020), and then either directly transferring to a smaller dataset such as Winograd Schema Challenge (WSC) (Levesque et al., 2012) or Pronoun Disambiguation Problems (PDP) (Morgenstern

WSC Sentences	Candidate Choices
The trophy doesn't fit in the suitcase because it is too <u>small</u> .	A. the trophy B. the suitcase
The trophy doesn't fit in the suitcase because it is too <u>big</u> .	A. the trophy B. the suitcase

Table 1: Above are two WSC examples. A system is required to resolve the bold pronoun “it” to “the suitcase” in the first sentence and to “the trophy” in the second sentence.

et al., 2016) or further finetuning on a downstream dataset such as SuperGLUE-WSC (Wang et al., 2019a). However, all the pipelines above can not avoid the phase of pre-training on a large human-labeled pronoun resolution dataset. Crowd-sourced “unbiased” labels that do not introduce annotation-artifacts (Gururangan et al., 2018) are shown to be costly and challenging to collect, requiring a well-designed annotation interface and dedicated annotators. To this end, we propose the unsupervised Masked Noun-Phrase Prediction task to pre-train a language model without any pronoun resolution training signal and directly transfer the pre-trained model to downstream datasets such as WSC.¹ Two examples of WSC are listed in Table 1. Our work improves on all previous unsupervised methods by large margins and even outperforms several strong supervised methods on all datasets we study.

We then proceed to the few-shot setting where we finetune our best zero-shot model on WinoGrande-S and XS respectively. MNPP gives a large margin of improvements over strong baselines including CSS (Klein and Nabi, 2020), RoBERTa-large (Sakaguchi et al., 2020), and UnifiedQA-BART-large (Khashabi et al., 2020). We further finetune on the remaining three data splits and achieve a higher AUC score on all five splits of WinoGrande over RoBERTa-large baseline.

¹We refer to unsupervised or zero-shot transfer as without training on any pronoun resolution dataset.

*Equal Contribution

In summary, our main contributions in this work are threefold.

- **First**, we propose the MNPP pre-training task and study how different synthetic dataset properties affect zero-shot performances.
- **Second**, we show MNPP outperforms all previous fully unsupervised methods and even several strong supervised baselines on all pronoun resolution datasets we study.
- **Finally**, we show that under few-shot settings, MNPP pre-training gives a significant performance boost on WinoGrande-S and XS and furthermore achieves a higher AUC score over all five splits of WinoGrande.

2 Related Works

In this work, we mainly compare with unsupervised methods.² On WSC, Zhang and Song (2018) propose the first unsupervised model where they modify Skip-Gram (Mikolov et al., 2013) objective to predict semantic dependencies then use this additional information during testing. Wang et al. (2019b) propose Unsupervised Deep Structured Semantic Models (UDSSM), which utilizes BiLSTM (Hochreiter and Schmidhuber, 1997) to compute contextual word embedding and uses models ensemble. Klein and Nabi (2019) directly explore the inner attention layers of BERT. Ye et al. (2019) adapt a masking and predicting strategy, called align, mask, and select (AMS), where entities that are connected with ConceptNet (Speer and Havasi, 2012) are masked and the model is required to select from a given list of candidate entities. An ensemble of large pre-trained models is first utilized by Trinh and Le (2018). GPT-2 is directly evaluated on WSC in Radford et al. (2019). Prakash et al. (2019) extend a language model with a knowledge hunting strategy. Kocijan et al. (2019b) and Kocijan et al. (2019a) are the most similar works to us and we will discuss the details in Section 3.1. Most recently, Klein and Nabi (2020) study a contrastive self-supervised learning approach (CSS) for WSC and DPR and also establish the first unsupervised baseline for KnowRef (Emami et al., 2019). On WinoGrande, knowledge hunting (Prakash et al., 2019) and language models ensemble (Sakaguchi et al., 2020) have been studied.

²Please refer to supplemental materials for more details on supervised methods.

3 Masked Noun-Phrase Prediction

We treat MNPP as a binary classification task. Given the sentence: “*She put the cup on the chair, but he knocked over the chair, and the cup fell.*”, the underlined “*the chair*” will be masked and a pair of replacement phrases for this masked position is given as {“*the cup*”, “*the chair*”}. One of the candidates is the masked phrase, “*the chair*”, and the other candidate is a different phrase in the sentence, “*the cup*” extracted from “*She put the cup on the chair*”. The constraint we impose is that both the ground-truth noun-phrase and the alternative candidate need to appear before the masked phrase location, which mimics the pronoun resolution task. We sample sentences following the above constraint to create our synthetic datasets for pre-training.

We convert the sentence into the format of {[CLS] *first_half* **option** *second_half* [SEP]} where *first_half* refers to “*She put the cup on the chair but he knocked over*” and *second_half* refers to “*, and the cup fell.*”. The **option** is replaced by candidates, “*the cup*” or “*the chair*”. We compute $P(\text{the chair}|\text{sentence}, \theta)$ and $P(\text{the cup}|\text{sentence}, \theta)$ and optimize θ , the parameters of the model, using cross-entropy loss. We use the final layer [CLS] vector from transformer-based language models and pass it through a single layer feed-forward network to calculate the logits.

3.1 Discussion

The intuition behind MNPP is that given sufficient samples that mimic pronoun resolution task, the model can learn rich knowledge to perform well on human-annotated pronoun resolution datasets. Such idea is also in-line with recent advances in unsupervised QA (Lewis et al., 2019; Li et al., 2020; Banerjee and Baral, 2020; Banerjee et al., 2020, 2021), where synthetic QA datasets are created from unannotated corpora to perform unsupervised pre-training. Strictly speaking, MNPP is even more unsupervised since our synthetic datasets are not created with true pronoun resolution signals, whereas synthetic QA datasets in works cited above contain true question-answer pairs.

As mentioned in previous Section 2, similar to our work, Kocijan et al. (2019b) studied such pre-training strategy by constructing a synthetic dataset, called MaskedWiki, which is crawled from English Wikipedia. However, our work is significantly different from theirs in the following ways. First, their

Dataset \ Source	CNN	QUOREF	Gutenberg	Knowledge	Total
Hybrid Source	100,556	51,451	6,381	-	158,388
Hybrid Source w/ Knowledge	189,376	98,844	19,424	75,993	383,637

Table 2: Number of instances from each source of two hybrid-source synthetic datasets in the first group.

Synth. Dataset \ Downstream	WinoGrande (AUC)	WSC	DPR	KnowRef	COPA
Hybrid Source (160k)	58.08 (0.6961)	79.48	82.27	79.83	71.29
Hybrid Source w/ Know. (380k)	58.56 (0.6821)	78.39	83.88	79.04	73.27
Gutenberg-10k	57.93 (-)	75.09	81.21	77.15	79.21
Gutenberg-50k	57.40 (-)	76.19	77.84	75.10	74.26
Gutenberg-100k	58.56 (-)	72.53	75.00	74.40	75.25
Gutenberg-300k	57.38 (-)	75.82	81.56	76.44	78.22
Gutenberg-500k	59.19 (0.6748)	76.56	80.50	79.12	85.51
Gutenberg-Easy (33k)	56.43 (-)	69.60	70.92	75.10	77.23
Gutenberg-Medium (33k)	57.00 (-)	75.10	80.32	78.17	79.21
Gutenberg-Hard (33k)	57.54 (-)	75.82	80.67	79.98	74.36

Table 3: Zero-shot transfer performances (%) on downstream datasets. AUC scores of WinoGrande are calculated after finetuning on all 5 splits of WinoGrande training sets. Difficulty level is decided using cosine similarity between the two candidate word vectors. Hard samples are the top 33% of samples when they are sorted in descending order using similarity score. Easy are bottom 33%, with Medium in-between.

pipeline requires further finetuning on another pronoun resolution task before transferring to downstream datasets, whereas our method can be directly evaluated on downstream datasets. Second, the size of MaskedWiki is 2.4 millions, which is 15 times the size of our best performing synthetic dataset. Third, we study how different properties of synthetic datasets affect zero-shot performances. Finally, they use a masked token prediction loss, and we model it as a classification task. Kocijan et al. (2019a) also construct another synthetic dataset called WikiCREM following the same masking principle but with only personal names masked.

4 Experiments and Results

4.1 Synthetic Dataset

We study three properties of synthetic dataset: source style, size, and difficulty level. The sources we choose include various styles of texts, including CNN stories (See et al., 2017), Wikipedia, and PG-19 language modeling benchmark (Rae et al., 2020). We study 3 groups and a total of 10 different synthetic datasets. The first group contains two synthetic datasets collected from all sources with and without knowledge hunting strategy (Prakash et al., 2019). The second group contains five synthetic datasets collected only from PG-19 but with varying sizes from 10k to 500k. The third group contains three synthetic datasets collected from PG-19 but with easy, medium, and hard samples with

the same size of 33k each.³ Datasets’ names are listed in the first column of Table 3 and statistics of the first group are described in Table 2.

4.2 Unsupervised Pronoun Resolution

The downstream datasets we test on are the WinoGrande test set (17k instances), DPR test set (564 instances), KnowRef test set (12k instances), and COPA validation set (101 instances). Although COPA (Wang et al., 2019a) is a cause and effect identification dataset, Sakaguchi et al. (2020) show that directly transferring from a WinoGrande-finetuned RoBERTa-large model to COPA already achieves a good performance, indicating that finetuning on WinoGrande can serve as a resource for common sense knowledge. We also investigate whether learning through MNPP can serve as a resource for common sense. Note that we also provide evaluation on the GAP dataset (Webster et al., 2018) in Table 5 for reference although the authors of GAP explicitly mention in their paper that they urge the community to not treat GAP as a Winograd-style task but a co-reference resolution task without gold mention provided.

4.2.1 Results

We report our experiment results in Table 3 and Table 4. Table 3 shows that different downstream

³Please refer to supplemental materials for details on synthetic datasets constructions.

WSC (Levesque et al., 2012)	
Bi-LSTM-DPR (2018)	56.0
BERT_NSP-DPR (2019)	71.1
CorefBERT _{LARGE} (2020)	71.4
BERT-WIKICREM-DPR (2019a)	71.8
BERT-MASKEDWIKI-DPR (2019b)	72.5
UDSSM-MASKEDWIKI-DPR (2019)	75.1
AMS-CSQA-DPR (2019)	75.5
RoBERTa-DPR (2020)	83.1
CorefRoBERTa _{LARGE} (Ye et al., 2020)	83.2
RoBERTa-WG (2020)	90.1
Modified Skip-Gram (2018)	60.3
BERT Inner Attention (2019)	60.3
BERT-MASKEDWIKI (2019b)	61.9
UDSSM (2019b)	62.4
BERT-WIKICREAM (2019a)	63.4
Ensemble LMs (2018)	63.7
CSS (2020)	69.6
GPT-2 (2019)	70.7
WSC Know. Hunting (2019)	71.1
MNPP (this work)	79.5

WinoGrande (Sakaguchi et al., 2020)	AUC
RoBERTa (local context) (2020)	50.0 -
BERT-DPR (2020)	51.0 -
BERT (local context) (2020)	51.9 -
RoBERTa-DPR (2020)	58.9 -
BERT (2020)	64.9 <u>0.5289</u>
CSS (2020)	65.0 <u>0.6046</u>
UnifiedQA-Bart-large (2020)	73.3 <u>0.6358</u>
CorefRoBERTa _{LARGE} (2020)	77.9 -
RoBERTa-large (2020)	79.1 <u>0.6641</u>
CorefBERT _{LARGE} (2020)	80.8 -
TTTTT (2020)	84.6 0.7673
UnifiedQA-T5-11B (2020)	89.4 0.8571
Wino Know. Hunting (2020)	49.6 -
Ensemble LMs (2020)	50.9 -
MNPP (this work)	59.2 0.6706

DPR (Rahman and Ng, 2012)	
Bi-LSTM (2018)	63.0
FeatureEng+Ranking (2012)	73.0
BERT-WIKICREM-DPR (2019a)	80.0
BERT-DPR (2019a)	83.3
BERT-MASKEDWIKI-DPR (2019b)	84.8
BERT-WG (2020)	84.9
CorefBERT _{LARGE} (Ye et al., 2020)	85.1
RoBERTa-DPR (2020)	91.7
CorefRoBERTa _{LARGE} (Ye et al., 2020)	92.2
RoBERTa-WG (2020)	92.5
RoBERTa-WG-DPR (2020)	93.1
BERT-WIKICREAM (2019a)	67.4
CSS (2020)	80.1
MNPP (this work)	83.9

KnowRef (Emami et al., 2019)	
E2E-CoNLL (2019)	60.0
E2E-KnowRef (2019)	61.0
BERT (2019)	65.0
E2E-KnowRef+CoNLL (2019)	65.0
RoBERTa-DPR (2020)	84.2
RoBERTa-WG (2020)	85.6
CSS (2020)	65.5
MNPP (this work)	80.0

COPA (Wang et al., 2019a)	
RoBERTa-WG (2020)	84.4
MNPP (this work)	85.5

Table 4: Comparisons of zero-shot transfer performance (%) among baselines and MNPP. Works highlighted with gray are supervised methods either directly finetuned on downstream datasets or additionally finetuned on another pronoun resolution dataset. Works highlighted with cyan are fully unsupervised methods. Best performances are in bold. We also underline supervised methods that our method outperforms. Note that AUC score for MNPP is obtained after finetuning on all WinoGrande data splits. (Model-A-B stands for model finetuned on A and B sequentially.)

dataset benefits from different property of the synthetic dataset. The hybrid-source synthetic dataset of size 160k outperforms PG-500k by a large margin on both WSC and DPR. It shows that pre-training on text of various styles instead of larger size is probably a better guarantee for better zero-shot performance on WSC and DPR. However, on WinoGrande and KnowRef, text style and dataset size both seem to impact zero-shot performance. On WinoGrande, larger size matters slightly more, whereas on KnowRef, synthetic dataset with various styles of texts gives better performance. On COPA, it is clear that using books as the source and with larger size at the same time is the key, probably because fictional event descriptions describing day-to-day activities in books contain more common sense, whereas CNN or Wikipedia articles contain precise, factual, non-fictional event descriptions. Finally, pre-training on more challenging examples helps on all tasks except COPA.

Compared with previous methods in Table 4, MNPP outperforms all unsupervised methods on all datasets and is comparable with several strong supervised methods. Current best unsupervised methods on WinoGrande is either random guess or below it, however, MNPP outperforms all of them by a margin of at least 8%. Even compared with a supervised baseline where BERT is first finetuned on DPR, our method outperforms it by 8%. On WSC, MNPP also outperforms all SOTA unsupervised methods by more than 8% and outperforms most supervised methods by at least 4% except RoBERTa-large finetuned on another pronoun resolution dataset. On DPR, our method outperforms the SOTA unsupervised baseline over 3% and also achieves only 1% behind the strong supervised baseline that finetunes BERT on MaskedWiki and DPR sequentially or only on WinoGrande. On KnowRef, MNPP outperforms the only unsuper-

	M	F	B	O
BERT (Kocijan et al., 2019a)	75.3	75.1	1.00	75.2
CorefBERT _{LARGE} (Ye et al., 2020)	-	-	-	76.8
BERT-WIKICREM-GAP (Kocijan et al., 2019a)	76.4	78.4	1.03	77.4
CorefRoBERTa _{LARGE} (Ye et al., 2020)	-	-	-	77.8
BERT-WIKICREM-ALL-GAP (Kocijan et al., 2019a)	76.7	79.4	1.04	78.0
BERT-WIKICREM (Kocijan et al., 2019a)	60.5	57.5	0.95	59.0
MNPP (this work)	71.3	75.2	1.05	73.3

Table 5: Performance comparisons among previous works and MNPP on GAP measured in F1. M stands for male, F stands for female, B stands for bias, and O stands for overall. Works highlighted with lightgray are supervised methods and works highlighted with cyan are fully un-supervised methods.

vised baseline by nearly 15% and achieves only 5% behind SOTA supervised model. Finally, on COPA, we show that MNPP gives models better common sense knowledge than finetuning on WinoGrande.

Meanwhile, we are not surprised that SOTA supervised methods still outperform unsupervised methods, including ours, considering the supervision itself and huge models with billions of parameters such as T5-11B.

4.3 Few-Shot Pronoun Resolution

We further proceed to the few-shot setting on WinoGrande-S and XS. We take the top three performance zero-shot models on WinoGrande development set and finetune them on WinoGrande-XS (160 instances) and S (640 instances) separately. After few-shot evaluation, we also finetune on the remaining three data splits, which are WinoGrande-M, L, and XL. Best performances on all 5 data splits are reported in Fig. 1 and AUC scores are reported in third column of WinoGrande section in Table 4.

4.3.1 Results

As indicated in Figure 1, MNPP outperforms CCS, UnifiedQA-BART-large, and RoBERTa-large on WinoGrande-S and XS with a large margin, and more importantly, achieves a higher AUC score as indicated in Table 4. It is clear that MNPP pre-training gives the model crucial additional information in the few-shot setting where only minimal data is available. We also notice that in the AUC column of Table 3, there is a negative correlation between zero-shot performance and AUC score, which means higher zero-shot performance does

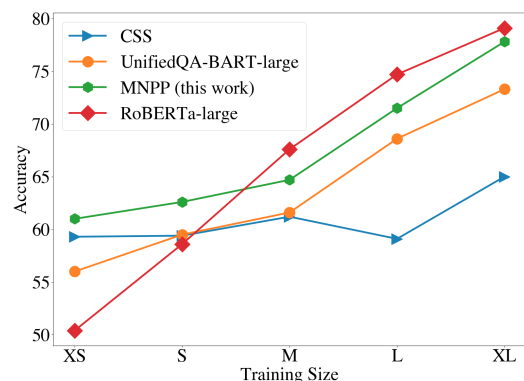


Figure 1: Performances (%) on WinoGrande test set after finetuning on 5 sizes of WinoGrande training set.

not guarantee better finetuning results.

Again we need to mention that we are not comparing with SOTA performances from billions-parameters models such as UnifiedQA-T5-11B from Khashabi et al. (2020) or T5-3B from Lin et al. (2020).

5 Conclusion

In this work, we propose MNPP pre-training to tackle unsupervised pronoun resolution and study how different properties of the synthetic pre-training dataset impact zero-shot performance on downstream datasets. Without finetuning on any pronoun resolution signal, MNPP outperforms all previous fully unsupervised methods on all tasks we study and even several strong supervised baselines. In the few-shot case where we finetune the zero-shot transfer model on WinoGrande-S and XS respectively, our model outperforms baselines by large margins, and further achieves a higher AUC score.

This work shows the effectiveness of unsupervised task definitions on text-based pronoun-resolution and common sense reasoning tasks. It would be interesting to design such tasks for multi-modal common sense reasoning (Zellers et al., 2019; Fang et al., 2020).

Acknowledgements

The authors acknowledge support from the DARPA SAIL-ON program W911NF2020006, ONR award N00014-20-1-2332, and NSF grant 1816039; and thank Yulong Chen for proofreading and the anonymous reviewers for their insightful discussion.

References

- Pratyay Banerjee and Chitta Baral. 2020. [Self-supervised knowledge triplet learning for zero-shot question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 151–162, Online. Association for Computational Linguistics.
- Pratyay Banerjee, Tejas Gokhale, and Chitta Baral. 2021. [Self-supervised test-time learning for reading comprehension](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1200–1211, Online. Association for Computational Linguistics.
- Pratyay Banerjee, Tejas Gokhale, Yezhou Yang, and Chitta Baral. 2020. [Self-supervised vqa: Answering visual questions using images and captions](#). *arXiv preprint arXiv:2012.02356*.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Pradeep Dasigi, Nelson F. Liu, Ana Marasović, Noah A. Smith, and Matt Gardner. 2019. [Quoref: A reading comprehension dataset with questions requiring coreferential reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5925–5932, Hong Kong, China. Association for Computational Linguistics.
- 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.
- 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.
- Zhiyuan Fang, Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. 2020. [Video2Commonsense: Generating commonsense descriptions to enrich video captioning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 840–860, Online. Association for Computational Linguistics.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Pengcheng He, Xiaodong Liu, Weizhu Chen, and Jianfeng Gao. 2019. [A hybrid neural network model for commonsense reasoning](#). *arXiv preprint arXiv:1907.11983*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural computation*, 9(8):1735–1780.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hananeh Hajishirzi. 2020. [UNIFIEDQA: Crossing format boundaries with a single QA system](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online. Association for Computational Linguistics.
- Tushar Khot, Peter Clark, Michal Guerquin, Peter Jansen, and Ashish Sabharwal. 2020. [Qasc: A dataset for question answering via sentence composition](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8082–8090.
- Tassilo Klein and Moin Nabi. 2019. [Attention is \(not\) all you need for commonsense reasoning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4831–4836, Florence, Italy. Association for Computational Linguistics.
- 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, Online. Association for Computational Linguistics.
- Vid Kocijan, Oana-Maria Camburu, Ana-Maria Cretu, Yordan Yordanov, Phil Blunsom, and Thomas Lukasiewicz. 2019a. [WikiCREM: A large unsupervised corpus for coreference resolution](#). 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 4303–4312, Hong Kong, China. Association for Computational Linguistics.
- Vid Kocijan, Ana-Maria Cretu, Oana-Maria Camburu, Yordan Yordanov, and Thomas Lukasiewicz. 2019b. [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, Thomas Lukasiewicz, Ernest Davis, Gary Marcus, and Leora Morgenstern. 2020. [A review of winograd schema challenge datasets and approaches](#). *arXiv preprint arXiv:2004.13831*.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale ReAding comprehension dataset from examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. [Higher-order coreference resolution with coarse-to-fine inference](#). 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 687–692, New Orleans, Louisiana. Association for Computational Linguistics.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. [The winograd schema challenge](#). In *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*. Citeseer.
- Patrick Lewis, Ludovic Denoyer, and Sebastian Riedel. 2019. [Unsupervised question answering by cloze translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4896–4910, Florence, Italy. Association for Computational Linguistics.
- Zhongli Li, Wenhui Wang, Li Dong, Furu Wei, and Ke Xu. 2020. [Harvesting and refining question-answer pairs for unsupervised QA](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6719–6728, Online. Association for Computational Linguistics.
- Sheng-Chieh Lin, Jheng-Hong Yang, Rodrigo Nogueira, Ming-Feng Tsai, Chuan-Ju Wang, and Jimmy Lin. 2020. [Ttttackling winogrande schemas](#). *arXiv preprint arXiv:2003.08380*.
- 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*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). *arXiv preprint arXiv:1301.3781*.
- Leora Morgenstern, Ernest Davis, and Charles L Ortiz. 2016. [Planning, executing, and evaluating the winograd schema challenge](#). *AI Magazine*, 37(1):50–54.
- Juri Opitz and Anette Frank. 2018. [Addressing the Winograd schema challenge as a sequence ranking task](#). In *Proceedings of the First International Workshop on Language Cognition and Computational Models*, pages 41–52, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. [CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes](#). In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea. Association for Computational Linguistics.
- Ashok Prakash, Arpit Sharma, Arindam Mitra, and Chitta Baral. 2019. [Combining knowledge hunting and neural language models to solve the Winograd schema challenge](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6110–6119, Florence, Italy. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8):9.
- Jack W. Rae, Anna Potapenko, Siddhant M. Jayakumar, Chloe Hillier, and Timothy P. Lillicrap. 2020. [Compressive transformers for long-range sequence modelling](#). In *International Conference on Learning Representations*.
- Altaf Rahman and Vincent Ng. 2012. [Resolving complex cases of definite pronouns: The Winograd schema challenge](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 777–789, Jeju Island, Korea. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Yu-Ping Ruan, Xiaodan Zhu, Zhen-Hua Ling, Zhan Shi, Quan Liu, and Si Wei. 2019. [Exploring unsupervised pretraining and sentence structure modelling for winograd schema challenge](#). *arXiv preprint arXiv:1904.09705*.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [Winogrande: An adversarial winograd schema challenge at scale](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8732–8740.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational*

- Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Robert Speer and Catherine Havasi. 2012. [Representing general relational knowledge in conceptnet 5](#). In *LREC*, pages 3679–3686.
- Trieu H Trinh and Quoc V Le. 2018. [A simple method for commonsense reasoning](#). *arXiv preprint arXiv:1806.02847*.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019a. [Super-glue: A stickier benchmark for general-purpose language understanding systems](#). *arXiv preprint arXiv:1905.00537*.
- Shuohang Wang, Sheng Zhang, Yelong Shen, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, and Jing Jiang. 2019b. [Unsupervised deep structured semantic models for commonsense reasoning](#). 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 882–891, Minneapolis, Minnesota. 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.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Deming Ye, Yankai Lin, Jiaju Du, Zhenghao Liu, Peng Li, Maosong Sun, and Zhiyuan Liu. 2020. [Coreferential Reasoning Learning for Language Representation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7170–7186, Online. Association for Computational Linguistics.
- Zhi-Xiu Ye, Qian Chen, Wen Wang, and Zhen-Hua Ling. 2019. [Align, mask and select: A simple method for incorporating commonsense knowledge into language representation models](#). *arXiv preprint arXiv:1908.06725*.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [From recognition to cognition: Visual commonsense reasoning](#). *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6713–6724.
- Hongming Zhang and Yangqiu Song. 2018. [A distributed solution for winograd schema challenge](#). In *Proceedings of the 2018 10th International Conference on Machine Learning and Computing*, pages 322–326.
- Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D. Manning. 2017. [Position-aware attention and supervised data improve slot filling](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 35–45, Copenhagen, Denmark. Association for Computational Linguistics.

A Related Work on Supervised Methods

WSC & DPR. Opitz and Frank (2018) is the first work to propose transfer learning from another pronoun resolution dataset such as DPR to WSC. He et al. (2019) use a hybrid model of Wang et al. (2019b) and Kocijan et al. (2019b). Ruan et al. (2019) explore BERT’s next sentence prediction with finetuning on DPR. Ye et al. (2020) finetune a new language representation model called CorefBERT, which is trained with a novel task to strengthen the co-referential reasoning ability of BERT, on DPR and then test on DPR and WSC. The SOTA supervised performance is provided by Sakaguchi et al. (2020) where they finetune a RoBERTa-large model on WinoGrande or DPR and evaluate on WSC and DPR without and with further finetuning. A detailed review of WSC and WSC related dataset can be found at Kocijan et al. (2020).

KnowRef. In Emami et al. (2019), an end-to-end neural system (Lee et al., 2018) is trained on CoNLL2012 shared task (Pradhan et al., 2012) and then tested under three settings: directly applying to KnowRef test set, retraining on KnowRef, and retraining on KnowRef plus CoNLL2012. Sakaguchi et al. (2020) transfer a WinoGrande-finetuned RoBERTa-large model and DPR-finetuned RoBERTa-large model to KnowRef test set respectively.

WinoGrande. The authors of WinoGrande finetune a RoBERTa-large on WinoGrande training set and evaluate on the test set in standard supervised setting, and Lin et al. (2020) finetune a T5-3B model instead. Sakaguchi et al. (2020) also study finetuning BERT and RoBERTa with only local context (only tokens near the pronoun location are available instead of the whole sentence). Ye et al. (2020) finetune WinoGrande using CorefBERT. Klein and Nabi (2020) finetune their unsupervised CSS model. Finally, UnifiedQA (Khashabi et al., 2020), which is pre-trained on eight seed QA datasets spanning four different formats in a unified way, is finetuned on WinoGrande.

B Synthetic Datasets Construction

For the first synthetic dataset in the first group, we choose 5000 stories in CNN stories, a small portion of Gutenberg books, and the whole training set of QUOREF (Dasigi et al., 2019), which is a reading comprehension dataset that requires resolving co-

reference among entities crawled from Wikipedia, and these sources result in the size of 160k. The second synthetic dataset in the first group comprises the same sources as above plus extra knowledge crawled by Google query using the knowledge hunting strategy introduced in Prakash et al. (2019). Following their strategy, we scrap 6531 and 69462 knowledge sentences for WSC and WinoGrande respectively. We relax the filtering process to allow longer sentences than those in the first synthetic dataset and lead to 380k samples in total. We then fix the text style and study the influence of data size on pre-training. We use 2000 books from PG-19 as the source and create five synthetic datasets with size of 500k, 300k, 100k, 50k, and 10k as the second group. We further study how difficulty levels of samples affect the downstream zero-shot performance. We select 100k samples from the PG-19 books described above and evenly split them into three synthetic datasets with low, medium, and high similarity scores between candidate choices as the third group. As a result, we create 3 groups of synthetic datasets with ten synthetic datasets in total. We used spaCy⁴ to pre-process raw text, including removing blank spaces, special characters, sentences that are too short or too long, and extracting noun-phrases.

C Zero-shot Experiment Details

Recent study (Khot et al., 2020) has shown that finetuning a RACE-finetuned (Lai et al., 2017) RoBERTa model as a start point is much more stable than directly finetuning a RoBERTa model from scratch, we follow the same strategy to start finetuning a RACE-finetuned RoBERTa-large model on all synthetic datasets. We use Hugging Face Transformers⁵ as our codebase. We set Adam optimizer with an initial learning rate of $1e - 5$ and epsilon of $1e - 8$, and without weight decaying for all settings. For a synthetic dataset whose size is larger or equal to 100k, we choose the batch size of 32 and train for 20 epochs, otherwise, we choose the batch size of 16 and train for 50 epochs. We checkpoint every X steps, with X in [50,500].

D Few-shot Experiment Details

We set Adam optimizer with an initial learning rate of $1e - 5$ and epsilon of $1e - 8$, without weight decaying, and batch size between 16 and 32 for all

⁴<https://spacy.io/>

⁵<https://github.com/huggingface/>

sizes. We finetune 20 epochs for WinoGrande-XL, L, and M, 40 epochs for S, and 160 epochs for XS. We checkpoint every X steps, with X in $[50,500]$.