

WINODICT: Probing language models for in-context word acquisition

Julian Martin Eisenschlos¹, Jeremy R. Cole¹, Fangyu Liu², William W. Cohen¹

¹ Google Research ² University of Cambridge

{eisenjulian, jrcole, wcohen}@google.com fl399@cam.ac.uk

Abstract

We introduce a new in-context learning paradigm to measure Large Language Models’ (LLMs) ability to learn novel words during inference. In particular, we rewrite Winograd-style co-reference resolution problems by replacing the key concept word with a synthetic but plausible word that the model must understand to complete the task. Solving this task requires the model to make use of the dictionary definition of the new word given in the prompt. This benchmark addresses word acquisition, one important aspect of the diachronic degradation known to afflict LLMs. As LLMs are frozen in time at the moment they are trained, they are normally unable to reflect the way language changes over time. We show that the accuracy of LLMs compared to the original Winograd tasks decreases radically in our benchmark, thus identifying a limitation of current models and providing a benchmark to measure future improvements in LLMs ability to do in-context learning.

1 Introduction

Large Language Models (LLMs) such as GPT-3 (Brown et al., 2020) and PALM (Chowdhery et al., 2022) can only learn from information that is in their training corpus. However, this is naturally limiting because the training corpus itself is bounded in time to the point of its collection. As a result, recent work has studied how to adapt such models to new data without an expensive re-training phase. Methods range from using semi-parametric methods with access to external memory (e.g., Guu et al. 2020; Lewis et al. 2020), to continual learning (e.g., Dhingra et al. 2022; Lazaridou et al. 2021), to parameter efficient fine-tuning (e.g., Ben Zaken et al. 2022; Pfeiffer et al. 2021).

Much of this work concerns factual knowledge or task distribution shifts. However, language also changes subtly: for instance, the popularity or meaning of individual words can change over

time. In fact, such shifts also cause a consistent decrease in model performance for downstream tasks (Huang and Paul, 2018; Jaidka et al., 2018; Lukes and Sjøgaard, 2018; Florio et al., 2020).

Acquiring new words through either examples or definitions is therefore an important test of LLMs’ ability to overcome diachronic degradation. With in-context learning having emerged as the primary way to interact with LLMs (Brown et al., 2020), we propose to study LLMs capability of acquiring new vocabulary via prompting.

We propose WINODICT, a novel benchmark for word acquisition for LLMs. Word acquisition is challenging to study in a realistic setting as it is hard to know which terms a model has already been exposed to. To overcome this, we rely on a heuristic method to introduce newly invented words and define them in terms of existing concepts. Following previous work (Chakrabarty et al., 2022), we incorporate the required knowledge into the prompt. We then ask models to perform tasks that require successfully interpreting the invented words.

We consider the English co-reference resolution datasets Winograd Schema Challenge (Levesque et al., 2012) and WinoGrande (Sakaguchi et al., 2020). The examples are built in pairs with minimal changes, which allow the identification of the key concept that must be understood to solve the example. An example of WINODICT can be seen in Figure 1. Our contributions are the following:

- (a) We propose WINODICT, a method and dataset to test models for word acquisition skills.
- (b) We benchmark the performance of several state-of-the-art models across scale and number of shots.
- (c) We analyze the effect of prompt, POS tags, word likelihood and similarity for ease of acquisition.

These results help us understand the challenges for incorporating new concepts into LLMs. The code to build the dataset has been open-sourced.¹

¹https://github.com/google-research/language/tree/master/language/wino_dict

WINOGRAD	
The <u>city councilmen</u> refused the demonstrators a permit because <i>they</i> feared violence.	The city councilmen refused the <u>demonstrators</u> a permit because <i>they</i> advocated violence.
WINODICT	
The verb to plest means to be scared of, or want to avoid an object.	The verb to sparn means to to publicly recommend or support.
The <u>city councilmen</u> refused the demonstrators a permit because <i>they</i> plested violence.	The city councilmen refused the <u>demonstrators</u> a permit because <i>they</i> sparned violence.

Figure 1: An example pair from WINODICT together with its original WINOGRAD source. The task is to decide whether *they* refers to the city councilman or the demonstrators. Here, the correct answer is shown in blue and the incorrect answer in **red**. Note that in both cases, it is necessary to understand the meaning of the bolded key concept to resolve the co-reference, which we identify in WINOGRAD and substitute for a new word in WINODICT.

2 Methods

WINODICT, like WINOGRAD and WINOGRANDE, is a co-reference resolution task in a binary choice setup. A model is given two alternative noun phrases, and has to decide which one is more likely to correspond to a highlighted pronoun or blank.

2.1 Dataset Construction

To build WINODICT, we rely on the fact that the examples from WINOGRAD and WINOGRANDE are constructed from contrasting pairs (Gardner et al., 2020; Kaushik et al., 2020). Each instance differs in a minimal way from its counterpart with the true label reversed. This allows the identification of the key concept that needs to be parsed in order to resolve the task. In Figure 1 for instance, the verbs *fear* and *advocate* correspond to the key concepts.

WINOGRAD and WINOGRANDE are similar; however, WINOGRANDE is larger, uses blanks instead of pronouns, and the dataset has been filtered for co-occurrence bias between the key concept and the correct noun-phrase. This results in some examples that do not have a corresponding paired example with a different key concept.

To create our examples, we first recover the pairing between the examples, dropping those with no pairing. Secondly, we identify the key concept tokens that change from one example to the other, dropping examples where the key concept consists of multiple tokens. Finally we run the sentence through the spaCy² syntactic analyzer and fetch WordNet³ definitions of the key concepts’ lemmas. In the next section we show how the key concept

²<https://spacy.io>

³<https://wordnet.princeton.edu> (Miller, 1995)

POS	WINOGRAD	WINOGRANDE	Total
VERB	67	56	123
NOUN	34	24	58
ADV	5	25	30
ADJ	74	211	285
Total	180	316	496
Orig. Size	273	12,282	12,555
Sent. Len	16.34	18.93	17.99
Def. Len	14.07	14.3	14.22

Table 1: Statistics for the different part-of-speech tags in the synthetic words, as well as average number of tokens for the main statement and the word definition. WINODICT consists of 496 examples.

tokens are replaced by synthetic words. This results in 496 examples: additional information can be found in Table 1.

2.2 New Word Creation

Our goal is to create plausible synthetic words. We create plausible words using a simple probabilistic model of every one-, two-, and three-letter sequence that is trained on the vocabulary of English words⁴. These three-letter sequences are then sampled and combined to form new synthetic words. We filter any words that have a three letter sequence that does not occur in any other English word. We then sample the words based on their log probability, placing them into five buckets and keeping around 500 for each bucket.

The morphology for each word is created by aggregating over a sample of proposed synthetic word morphologies. The last 2-4 letters of each word (depending on the morphological edit) form a suffix dictionary that is used as a simple substitution

⁴<https://pypi.org/project/english-words>

dictionary for the remaining words: failures are dropped. This produces a combination of regular and irregular conjugations over the new words.

2.3 Answer Scoring

Each instance in WINODICT consists of a new word with its definition d , a statement containing a blank where x and y correspond to the text before and after the blank respectively, and two noun phrases o_1 and o_2 . The task consists of identifying which of the noun-phrases better fits the blank.

PALM, GPT-3 and its predecessors (Radford et al., 2019) use the method proposed by Trinh and Le (2018) to evaluate WINOGRAD and WINOGRANDE, which we explain below. A prediction score is obtained through comparing the log likelihood of the same continuation y of two possible prefix texts ($x : o_1$ and $x : o_2$) where the co-reference pronoun or blank marker has been replaced. It is correct if it scores the suffix higher for the prefix with the correct interpretation of the co-reference problem.

$$\ln P_{\Theta}(y|x : o_1) - \ln P_{\Theta}(y|x : o_2)$$

$$= \sum_{i=0}^n (\ln P_{\Theta}(y_i|y_{<i} : x : o_1) - \ln P_{\Theta}(y_i|y_{<i} : x : o_2))$$

where $:$ denotes concatenation and variables map to:

$x =$ "The city councilmen refused the demonstrators a permit because"

$o_1 =$ "the city councilmen"

$o_2 =$ "the demonstrators"

$\{y_i\}_{i=1}^n = y =$ "feared violence."

In our setup we add the definition of the new concept as a suffix to the shared term y , thus replacing it with $y : d$. This achieves higher accuracy than the alternatives. Note that this means that the model is *scoring* the definition rather than conditioning on it. See Section 4 and Table 4 for a discussion of other variants of the setup, including adding the definition as a prefix.

3 Experiments

In this work we test GPT-3 (Brown et al., 2020), and PALM (Chowdhery et al., 2022) models of various sizes, ranging from 3B to 540B parameters. Appendix A has more details on the model sizes.

As in the original in-context learning evaluations, we try 0, 1, and 5-shot experiments, using random examples to build the prompt. We compare to both

a zero-shot human evaluation as well as the original source datasets with only our filtered examples.

The main experimental results are shown in Table 2. We observe a consistent gap of 18 or more points between WINODICT examples and their original counterparts. Similar to trends observed in other datasets (Chowdhery et al., 2022), scaling the number of shots and model size consistently improves accuracy. The three smaller versions of GPT-3 and PALM-8B all perform close to random.

We verify that omitting any information of the new word yields random results for even the best PALM-540B model. We discuss this and other prompting strategies in more detail in Appendix B.

3.1 Human Evaluation

The human accuracy on WINODICT is estimated using the responses of 10 volunteers. No native English proficiency was required for participation. Participants were told that the aim of the research is to study how to use words based on their definition. They were presented with 15 sentences that included a pronoun / blank and asked to select one of two noun phrases it most likely refers to.

3.2 Foreign Inspired Words

To explore a more realistic scenario, we conduct an experiment using 20 hand-written WINODICT-like examples whose definitions are inspired by foreign words that do not have a clear single-word definition in English. For instance, "estrenar" refers to wearing a piece of clothing for the first time, which does not have a clear English word equivalent. We can then create an example that requires knowing this definition, such as "I really [love | hate] my new dress. I can't wait to <word> it.

In conducting this experiment, we substitute synthetic words instead of using the original foreign words, and the definitions of the words themselves may not correspond to native speakers' precise understanding: in other words, these are meant to be true new words and data leakage should be minimal. We run the same experiment on these examples. Results are in Table 3 and full details in Appendix C. Overall, the numbers are comparable to the WINODICT results, suggesting that models are unlikely to be solving the task using a reverse dictionary.

4 Prompt Analysis

In this section, we discuss alternative formulations for the prompts used in WINODICT. We focus on

Shots	WINOGRAD						WINOGRANDE					
	WINODICT (Ours)			Original			WINODICT (Ours)			Original		
	0	1	5	0	1	5	0	1	5	0	1	5
PALM 8B	59.2 \pm 1.6	57.1 \pm 2.1	59.1 \pm 1.6	83.3	83.3	87.2	51.8 \pm 1.6	54.2 \pm 0.4	52.4 \pm 1.1	69.3	65.5	67.4
PALM 62B	62.2 \pm 0.6	65.9 \pm 3.6	70.3 \pm 1.3	91.1	90.0	92.2	56.7 \pm 1.1	58.2 \pm 1.0	59.7 \pm 1.1	76.6	77.8	78.2
PALM 540B	65.9 \pm 2.5	75.4 \pm 1.3	78.6 \pm 0.6	92.8	92.2	95.6	60.3 \pm 1.4	63.9 \pm 2.3	68.5 \pm 1.9	80.1	81.3	85.8
GPT-3 Ada	51.9 \pm 2.2	50.9 \pm 1.7	50.2 \pm 4.3	60.0	57.8	61.7	52.2 \pm 1.2	52.0 \pm 3.6	49.4 \pm 1.7	48.1	53.8	53.2
GPT-3 Babbage	51.8 \pm 0.8	52.8 \pm 2.0	54.4 \pm 2.3	75.6	71.7	65.6	50.8 \pm 1.7	52.3 \pm 1.0	52.2 \pm 0.8	52.8	55.1	56.6
GPT-3 Curie	54.2 \pm 1.6	54.6 \pm 2.4	59.9 \pm 1.5	85.0	81.7	82.8	50.2 \pm 1.5	50.6 \pm 1.6	52.2 \pm 1.0	62.0	61.1	60.8
GPT-3 Davinci	60.3 \pm 1.3	63.6 \pm 2.3	72.9 \pm 0.5	88.3	85.0	91.1	55.0 \pm 1.1	55.7 \pm 1.4	61.3 \pm 1.4	71.8	69.6	72.5
Human	91.7			96.5*			83.3			94.0*		

Table 2: Binary classification accuracy on WINODICT vs. the original datasets using average and standard deviation across 5 sets of new words. Original results may differ from the ones reported by Chowdhery et al. (2022) since only a subset of the examples are used. A consistent gap of 18+ points appears when comparing against the original sets. The original human evaluation numbers denoted with * are taken from Sakaguchi et al. (2020).

Shots	0	1	5
PALM 540B	68.7	73.0	76.0
GPT-3 Davinci	61.0	56.0	68.0

Table 3: Binary classification accuracy on the foreign-inspired new words averaged over five runs. Overall, accuracy is comparable to the original dataset.

Word Type	Prompt	WINOGRAD	WINOGRANDE
Synthetic	Def Prefix	72.2	62.7
	Def Suffix*	78.6	68.5
	Syn Prefix	74.1	60.5
	Syn Suffix	88.4	78.2
	Empty	52.0	51.9
Original	Def Prefix	85.5	74.0
	Def Suffix	93.8	84.4
	Syn Prefix	87.2	74.3
	Syn Suffix	91.6	83.2
	Empty*	95.6	85.8
Meaning shift	Def Prefix	66.1	60.8
	Def Suffix	75.6	60.4
	Syn Prefix	69.4	60.1
	Syn Suffix	83.3	74.7
	Empty	51.1	49.7

Table 4: Analysis of different prompts. We show the results on the synthetic words, original words, and existing words but assigned to a new meaning (“Meaning shift”). Prefix/Suffix correspond to the location of the definition, Syn/Def corresponds to using the definition or synonyms of the synthetic word. Empty means neither (should be random for synthetic words). Providing synonyms yields the best results. All results are on PALM-540B 5-shot. The lines marked with * correspond to the experiments in Table 2.

the best-performing PALM-540B model using a 5-shot setup. See Table 4 for the full results.

Concretely, we vary the prompts along a few axes. First, we test whether the definition should be

part of the prefix, where the model would condition on it, or the suffix, where the model would score it. Note that in all setups, putting the definition in the suffix works consistently better.

Additionally, we test whether the task is made easier by using synonyms instead of definitions. This task indeed appears to be easier, potentially because the model needs to learn only a simple substitution between the new word and the provided synonym, whose definition it knows. We focus on definitions in this work as exact synonyms would rarely be available for novel words.

As a baseline, we also examine the “Empty” setup, where the model is provided no information about the new word. We observe that PALM approximates random guessing without being given the definition, showing that the task remains roughly unbiased.

We additionally test the model’s performance on the original task where we also provide the definition of the key concept. Note that the “Empty” case here corresponds precisely to the original task. Interestingly, the definition seems to serve as a slight distraction, especially as a prefix, though accuracy is still well above the model’s performance on the synthetic words.

Finally, in the “Meaning shift” scenario, we map new definitions to already known words. This task appears to be even more difficult than the standard WINODICT setup, implying that the model is distracted by the surface forms of the words.

5 New Word Analysis

Several factors can affect the capabilities for word acquisition of LLMs. We investigate several attributes, split into quartiles, using PALM-540B with 5-shots, which is the best model from Table 2.

We consider the following attributes: (1) the part-of-speech of the synthetic word; (2) the average model negative log likelihood (NLL) of the two model predictions, which measures the likelihood of the suffix for both prefixes; (3) the number of SentencePiece (Kudo and Richardson, 2018) tokens in the synthetic word, to investigate the effect of model tokenization; (4) the number of SentencePiece tokens in the definition of the synthetic word, to investigate if longer definitions are more challenging; (5) the Levenshtein edit distance between the synthetic and original word, to investigate if similar words are easier; and (6) the likelihood of the new word as computed by our probabilistic model of three-letter sequences, to see if less probable words are more difficult to acquire.

Of the six attributes, the two most correlated with accuracy are (4) the definition length and (2) the average NLL. We observe no clear pattern in the other four attributes. In Figure 2 we show their effect in each quartile. The effect of definition length indicated that the 25% longest definitions are the hardest to acquire by a significant margin (12% relative drop for WINOGRAD, 5% for WINOGRANDE). The relative accuracy drop for the largest quartile of the NLL average is 13% for WINOGRAD and 4% for WINOGRANDE. The drop in NLL suggests that when models assigns low probabilities to answers, they make more mistakes: the low probability may indicate the model has a poor understanding of the prefix so scores the suffix randomly.

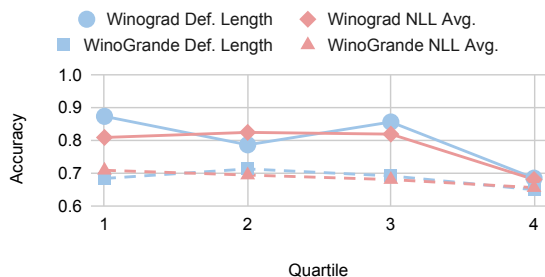


Figure 2: Effect on WINODICT PALM-540B 5-shot accuracy on each quartile splitting by definition length and by average NLL score. Longer definitions and higher NLL correlate with lower accuracy.

6 Related Work

Word acquisition for LLMs. Inspired by developmental linguistics (Carey and Bartlett, 1978), Radford et al. (2019) succeeded to prompt GPT-3 to generate plausible example sentences based on

definitions of synthetic words. Unlike WINODICT, the evaluation was purely qualitative.

Common sense. Li et al. (2021) study how prompt structures and scoring methods affect the performance of LLMs on common sense tasks including WINOGRANDE, where they observe the least variation. The format from WINOGRAD has been subsequently used to probe models for other phenomena such as explanations (Zhang et al., 2020) and gender bias (Zhao et al., 2018).

Benchmarks for lexical knowledge. Schick and Schütze (2020) introduce a benchmark for probing a model’s knowledge of the properties of rare words. Hill et al. (2016) train models to match word and definition representations, which they apply to a reverse dictionary task.

7 Conclusion

In this work, we study the question of in-context word acquisition by large language models. While non-trivial to measure, the ability to incorporate knowledge about new words in-context may be useful to decrease the effect of diachronic degradation. We design a mechanism to transform Winograd-style tasks into challenging probes for reasoning on the meaning assigned to synthetic words, allowing for a more objective measurement of word acquisition. We study the results of models of multiple sizes and families and conclude that while the problem becomes easier with scale, there is still a substantial gap with human performance and the original WINOGRAD and WINOGRANDE tasks, demonstrating the difficulty of the proposed task. Finally, we show that acquiring novel definitions is of similar difficulty, indicating the task is realistic.

Limitations

The task described in this work is synthetic and thus an imperfect measure of the phenomena under study. The words in WINODICT are synthetic words with definitions copied from existing concepts; the model could thus solve WINODICT with a reduction to a reverse dictionary task. To partially address this, we conducted pilot experiments described in Section 3.2.

Additionally, the choice of prompts for LLMs has been shown to have a large influence on the resulting accuracy (Min et al., 2022; Lu et al., 2022). While we tried multiple templates, it is possible that substantially better prompts exist for this task.

Acknowledgements

We thank Yasemin Altun, Iulia-Maria Comşa and Srini Narayanan, as well as our anonymous reviewers, for their valuable feedback.

Ethics Statement

All annotations for the human evaluation were done by English speaking volunteers using a multiple choice survey tool. No personal identifiable information was collected.

References

- Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. 2022. [BitFit: Simple parameter-efficient fine-tuning for transformer-based masked language-models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1–9, Dublin, Ireland. Association for Computational Linguistics.
- 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](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Susan Carey and E. Bartlett. 1978. [Acquiring a single new word](#). *Proceedings of the Stanford Child Language Conference*, 15:17–29.
- Tuhin Chakrabarty, Yejin Choi, and Vered Shwartz. 2022. [It’s not rocket science: Interpreting figurative language in narratives](#). *Transactions of the Association for Computational Linguistics*, 10:589–606.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. [Palm: Scaling language modeling with pathways](#). *arXiv preprint arXiv:2204.02311*.
- Bhuwan Dhingra, Jeremy Cole, Julian Eisenschlos, Daniel Gillick, Jacob Eisenstein, and William Cohen. 2022. [Time-aware language models as temporal knowledge bases](#). *Transactions of the Association for Computational Linguistics*, 10(0):257–273.
- Komal Florio, Valerio Basile, Marco Polignano, Pierpaolo Basile, and Viviana Patti. 2020. [Time of your hate: The challenge of time in hate speech detection on social media](#). *Applied Sciences*, 10(12):4180.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. [Evaluating models’ local decision boundaries via contrast sets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pappas, and Mingwei Chang. 2020. [Retrieval augmented language model pre-training](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3929–3938. PMLR.
- Felix Hill, Kyunghyun Cho, Anna Korhonen, and Yoshua Bengio. 2016. [Learning to understand phrases by embedding the dictionary](#). *Transactions of the Association for Computational Linguistics*, 4:17–30.
- Xiaolei Huang and Michael J. Paul. 2018. [Examining temporality in document classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 694–699, Melbourne, Australia. Association for Computational Linguistics.
- Kokil Jaidka, Niyati Chhaya, and Lyle Ungar. 2018. [Diachronic degradation of language models: Insights from social media](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 195–200, Melbourne, Australia. Association for Computational Linguistics.
- Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. 2020. [Learning the difference that makes a difference with counterfactually-augmented data](#). In *International Conference on Learning Representations*.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Angeliki Lazaridou, Adhi Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Tomas Kocisky, Sebastian Ruder, et al. 2021. [Mind the gap: Assessing temporal generalization in neural language models](#). *Advances in Neural Information Processing Systems*, 34:29348–29363.
- Hector J. Levesque, Ernest Davis, and Leora Morgenstern. 2012. [The winograd schema challenge](#). In

- Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning*, KR'12, page 552–561. AAAI Press.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Xiang Lorraine Li, Adhi Kuncoro, Cyprien de Masson d’Autume, Phil Blunsom, and Aida Nematzadeh. 2021. [A systematic investigation of commonsense understanding in large language models](#). *arXiv preprint arXiv:2111.00607*.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022. [Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, Dublin, Ireland. Association for Computational Linguistics.
- Jan Lukes and Anders Søgaard. 2018. [Sentiment analysis under temporal shift](#). In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 65–71, Brussels, Belgium. Association for Computational Linguistics.
- George A. Miller. 1995. [Wordnet: A lexical database for english](#). *Commun. ACM*, 38(11):39–41.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. [AdapterFusion: Non-destructive task composition for transfer learning](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 487–503, Online. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *CoRR*.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [Winogrande: An adversarial winograd schema challenge at scale](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8732–8740.
- Timo Schick and Hinrich Schütze. 2020. [Rare words: A major problem for contextualized embeddings and how to fix it by attentive mimicking](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8766–8774.
- Trieu H. Trinh and Quoc V. Le. 2018. [A simple method for commonsense reasoning](#). *CoRR*.
- 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.
- 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.

A Model Sizes

While OpenAI does not officially disclose the size of their four models Davinci, Curie, Babbage and Ada, we use the numbers approximated in a blog-post as estimates.⁵ Table 5 contains the number of parameters for the models used in our experiments.

Model	# Parameters
GPT-3-Ada	350M
GPT-3-Babbage	1.3B
GPT-3-Curie	6.7B
GPT-3-Davinci	175B
PaLM-8B	8B
PaLM-62B	62B
PaLM-540B	540B

Table 5: Number of parameters of the reported models.

B Prompts

We built prompts for definitions and synonyms to make them sound natural given the structure of most WordNet definitions for each part-of-speech tag. Table 6 shows the different prompt templates in each case.

Type	Prompt
Synonym	The meaning of {lemma} is similar to {synonym}.
Verb definition	The verb to {lemma} means to {definition}.
Noun definition	The word {lemma} refers to {definition}.
Adj. Definition	The meaning of {lemma} is definition.
Adv. Definition	The word {lemma} means {definition}.

Table 6: Templates used to integrate the definition into the prompt for each part-of-speech tag.

C Foreign Inspired Words

In Table 7 we list the word, approximate definition, and WINODICT-like example. Note that these examples are handwritten and did not go through a debiasing process like WINOGRANDE. In order to reduce the risk of data leakage, in the actual examples we replace the surface form of the word with one of the synthetic surface forms using the same process as in section 2.

⁵<https://blog.eleuther.ai/gpt3-model-sizes>

Example	Definition
John frequently goes backpacking and Jake never does because [<u>Jake</u> John] disdains the feeling of waldeinsamkeit .	the feeling of solitude and connectedness to nature when being alone in the woods
After returning from backpacking, John thought he would go again [<u>frequently</u> infrequently]. John likely appreciates the feeling of waldeinsamkeit .	the feeling of solitude and connectedness to nature when being alone in the woods
Mary loves going to antique stores and Ashley never does because [<u>Mary</u> Ashley] wabi-sabis old things.	finding beauty in imperfections
Mary loves going to [<u>antique</u> modern] stores because she wabi-sabis old things.	finding beauty in imperfections
Pierre is from France and John is from Ireland. Pierre and John like to go to Irish bars and talk about [<u>Pierre</u> John]'s feeling of depaysement there.	the feeling that comes from not being in one's home country; being a foreigner
Pierre has lived in France all his life. When he's in [<u>Ireland</u> France], Pierre frequently talks about his feeling of depaysement .	the feeling that comes from not being in one's home country; being a foreigner
Jake and Ashley plan to get married, Ashley's parents are happy, but Jake's parents don't like it because a friend said they had bad yuanfen . [<u>Jake</u> Ashley]'s parents are more likely to go to a fortune teller.	the fate between two people
Jake and Ashley plan to get married. Ashley's parents are very practical while Jake's parents believe in destiny. When an advisor said Jake and Ashley had bad yuanfen , [<u>Jake</u> Ashley] wanted to call it off.	the fate between two people
Theresa doesn't get why Martha thinks the statue in the museum was so duende that [<u>Martha</u> Theresa] spent a lot of time looking at it.	a work of art's mysterious power to deeply move a person
Martha spends a lot of times in museums while Theresa spends little. [<u>Martha</u> Theresa] finds art duende .	a work of art's mysterious power to deeply move a person
After losing his [<u>religion</u> job], John fell into a sense of toska .	a sensation of great spiritual anguish, often without a specific cause; a longing with nothing to long for
John kept yelling at Joey for not doing chores, but Joey wouldn't even respond. [<u>Joey</u> John] really seems tosked .	a sensation of great spiritual anguish, often without a specific cause; a longing with nothing to long for
Because he [<u>loves</u> hates] reptiles, John found seeing that group of lizards very gigil .	a situation of such extreme cuteness it's overwhelming or the irresistible urge to hug something cute
John only keeps salamanders as pets and Joey likes more traditional ones, so [<u>John</u> Joey] found seeing the group of lizards very gigil .	a situation of such extreme cuteness it's overwhelming or the irresistible urge to hug something cute
John thought his marriage with Joey was shougani , so he wanted to hire a [<u>lawyer</u> therapist].	a situation that can't be helped, or an act of resignation
John thought his marriage with Joey was shougani but Joey disagreed, so [<u>John</u> Joey] decided to hire a lawyer.	a situation that can't be helped, or an act of resignation
Joey still can't get over when John drunkenly called him Mark at his wedding, and now whenever they see each other, [<u>Joey</u> John] tartles .	a moment of hesitation when introducing someone because you can't remember their name
I really [<u>love</u> hate] my new dress. I can't wait to estrene it.	wearing something for the very first time
Mary and Sue went dress shopping together. Mary hates her dress while Sue loves hers. [<u>Sue</u> Mary] can't wait to estrene it.	wearing something for the very first time
After a long day of work, James xinkued the job John did. John was [<u>grateful</u> upset].	acknowledging someone's effort for working hard or doing you a favor

Table 7: List of foreign-inspired new words (**bolded**) and their corresponding examples and definitions. The possible choices for the example are shown, with the correct choice underlined. The definition is shown on the right. These definitions may or may not be idiosyncratic to a native speaker; however, the actual examples use a synthetic word to more closely resemble new word acquisition and minimize the risk of data leakage.