

Verbing Weirds Language (Models): Evaluation of English Zero-Derivation in Five LLMs

David R. Mortensen,¹ Valentina Izrailevitch,² Yunze Xiao,¹
Hinrich Schütze,³ Leonie Weissweiler³

¹Carnegie Mellon University, ²TU Munich, ³LMU Munich & MCML

Abstract

Lexical-syntactic flexibility, in the form of conversion (or zero-derivation) is a hallmark of English morphology. In conversion, a word with one part of speech is placed in a non-prototypical context, where it is coerced to behave as if it had a different part of speech. However, while this process affects a large part of the English lexicon, little work has been done to establish the degree to which language models capture this type of generalization. This paper reports the first study on the behavior of large language models with reference to conversion. We design a task for testing lexical-syntactic flexibility—the degree to which models can generalize over words in a construction with a non-prototypical part of speech. This task is situated within a natural language inference paradigm. We test the abilities of five language models—two proprietary models (GPT-3.5 and GPT-4), three open-source models (Mistral 7B, Falcon 40B, and Llama 2 70B). We find that GPT-4 performs best on the task, followed by GPT-3.5, but that the open source language models are also able to perform it and that the 7B parameter Mistral displays as little difference between its baseline performance on the natural language inference task and the non-prototypical syntactic category task, as the massive GPT-4.

Keywords: morphology, syntax, large language models, open source

1. Introduction

English displays a relatively high degree of flexibility regarding the syntactic category of lexical items. This tendency is so pervasive that it is even commented on in popular culture, as in the American comic strip Calvin and Hobbes, in which one of the main characters proclaims “I like to verb words” (Figure 1). In this case *verb*—whose prototypical function is as a noun—functions as a verb. It is coerced into acting as a verb (specifically, an infinitive) by being placed as the head of a verb phrase, a position a noun could never occupy. In fact, some linguists have quipped, in only slightly hyperbolic fashion, that—in English—you can verb anything.

Other examples include:

Adjective His hair has begun to *gray*.

Mass Noun If you don’t want to *water* the plants, please *coffee* the graduate students instead.

Count Noun The fascist tried to *knife* me in the back.

English also allows adjectives and verbs to be converted (zero-derived) into nouns:

Intransitive Verb I think I’ll go for a *swim*.

Transitive Verb I sustained a direct *hit*.

Adjective She’s got lots of *green* but she’s not spending any of it on me.

Note that, while *swim*, *hit*, and *green* are treated as nouns (as well as verbs or adjectives) by lexicographers, they are etymologically verbs and adjectives.

There are three different linguistic approaches to this phenomenon. One is to say that conversion (or zero-derivation) is like any other morpho-




Figure 1: Calvin and Hobbes ©1993 Watterson. Reprinted with permission of Andrews McMeel Syndication. All rights reserved.

logical process (except that no overt affix, stress shift, or other formal change is evident) (Marchand, 1969). The principal evidence for this, as pointed out by Beard (2017), is that conversion is subject to blocking effects: if a word of the same meaning as a potential converted word already exists in the lexicon (e.g., due to derivation), it will be used instead of the converted word. The appeal of this approach is that it accounts for the many different semantic relationships that can exist between bases and zero-derived words with the same parts of speech. Contrariwise, it has been proposed that conversion is not morphological at all—that it is, effectively, coining new words (Lieber, 1980). Our study is conducted in the spirit of a third approach, namely that English, Dutch, and other, similar, languages allow flexibility with regard to syntactic categories (parts of speech) when this is (1) allowed by the context semantically and (2) required by the context syntactically (Clark and Clark, 1979). We do not attempt to implement these three approaches computationally, or to distinguish them empirically, but methodologically we manipulate part of speech by manipulating aspects of the syntactic and semantic context, inspired by insights of Clark and Clark (1979).

In this study, we evaluate five large language models of varying sizes, based on their ability to make inferences that require reasoning about words used in non-prototypical grammatical contexts. We investigate four major hypotheses:

- performance on the task with prototypical parts of speech is better than with non-prototypical parts of speech
- non-prototypical parts of speech are associated with better performance than nonce words
- correlation between performance in the prototypical, non-prototypical, and nonce conditions
- differences in model size account for differences in performance

Performance in the prototypical condition is—indeed—the best, but performance in the nonce and non-prototypical conditions are similar. The performance of each model was correlated across conditions. We find that the models vary greatly in their ability in this area, with the very large, commercial models performing best. However, we also show that the number of parameters alone does not predict the performance of models on this task. Instead, a good predictor is the performance of the models on a generic version of the task in which all words are used in prototypical ways. This suggests that most of the variance

in the model scores is variance in the ability to perform the NLI task itself, and that the other differences are of lesser, but still significant, importance.

We make three main contributions:

- A new methodology for investigating conversion in language models
- A test set for systematically applying this methodology to arbitrary models
- The demonstration that lexical-syntactic flexibility does not increase monotonically with model size

2. Related Work

Linguistic work on conversion in English dates back to Sweet (1891). It has been taken up sporadically by researchers since then (Marchand, 1969; Clark and Clark, 1979; Lieber, 1980; Kiparsky, 1982; Don, 1993; Bauer and Hernández, 2005). Much of the literature regarding conversion concerns whether conversion is due to a kind of zero-affixation (Marchand, 1969), a process of coinage (Lieber, 1980), or the flouting of syntactic category constraints when constrained (and allowed) by context (Clark and Clark, 1979). This paper assumes the position of Clark and Clark, namely that languages like English allow syntactic flexibility when licensed by semantics and required by the encompassing constructions. This is analogous, in some ways, to Goldberg (1995)’s analysis of argument structure alternations, where the broader construction coerces, for example, intransitive verbs to function as transitive verbs (in the caused-motion construction). Similarly, we assume that constructional contexts coerce words to function as if they have a non-prototypical part of speech.

While this is the first study modeling conversion computationally, there is a growing body of work addressing related issues for a broader range of phenomena in derivational morphology and neologism. The most relevant to the current work are Hofmann et al. (2021) and Hofmann et al. (2020a), which address derivational morphology in the context of older BERT-like language models (but not contemporary LLMs). Factors in the emergence of new words have been elucidated by Ryskina et al. (2020).

3. Methodology

We sought to design a task that would test whether words from non-prototypical syntactic categories—converted words and nonce words—

Subtask	Prompt	Intended
transitive	If I asked you to X it, do I want you to $\left\{ \begin{array}{l} \emptyset \\ \text{not} \end{array} \right\}$ X it?	✓ ✗
	If I asked you $\left\{ \begin{array}{l} \text{not to} \\ \text{to not} \end{array} \right\}$ X it, do I want you to $\left\{ \begin{array}{l} \emptyset \\ \text{not} \end{array} \right\}$ X it?	✗ ✓ ✓
	If I say, "Don't X me," am I asking you to $\left\{ \begin{array}{l} \text{not} \end{array} \right\}$ X me?	✓ ✗
intransitive	If I $\left\{ \begin{array}{l} \emptyset \\ \text{don't} \end{array} \right\}$ X daily, do I X every day?	✓ ✗
	If I $\left\{ \begin{array}{l} \text{tried} \\ \text{did not try} \end{array} \right\}$ to X , did I attempt to X ?	✓ ✗
count	If I like this X more than the other one, do I prefer $\left\{ \begin{array}{l} \text{this } \mathbf{X} \text{ to the other one?} \\ \text{the other } \mathbf{X} \text{ to this one?} \end{array} \right\}$	✓ ✗
	If I like this X less than the other one, do I prefer $\left\{ \begin{array}{l} \text{this } \mathbf{X} \text{ to the other one?} \\ \text{the other } \mathbf{X} \text{ to this one?} \end{array} \right\}$	✗ ✓
mass	I prefer less X . Do I prefer $\left\{ \begin{array}{l} \text{less} \\ \text{more} \end{array} \right\}$ X ?	✓ ✗
	I prefer more X . Do I prefer $\left\{ \begin{array}{l} \text{less} \\ \text{more} \end{array} \right\}$ X ?	✗ ✓

Table 2: Frames used for generating prompts

affect the ability of text-in-text-out language models to make pragmatic generalizations.

3.1. Materials

We created a set of 3,069 prompts based on the frames in Table 2 and five word lists (Table 1). Nonce words were generated with Unipseudo

Part of Speech	Number	Example
transitive verbs	42	bamboozle
intransitive verbs	42	deign
mass nouns	51	music
count nouns	79	professor
nonce words	49	theord

Table 1: Word lists derived from UniMorph and Unipseudo (New et al., in press) and verified manually

(New et al., in press) based on a list of the 59 most frequent mono-morphemic nouns verbs in English with length 6 (as listed in UniMorph). This list was manually culled to remove words that were (1) too similar to or (2) too distant from any known English words. All lexical sets were manually curated by a native-speaker linguist.

The frames and the wordlists were combined according to principled criteria, yielding 3,069 items. Prompts reflect the format of the following example:

If I asked you to **day** it, do I want you to **day** it?

3.2. Experiments

The prompts were presented to five models, two closed models (GPT-3.5 and GPT-4) and three open models of varying sizes (Mistral 7B, Falcon 40B, and Llama 2 70B). The closed models were prompted via the OpenAI API. The open models were all evaluated (without quantization) using vLLM on a cluster node with 4 A6000 GPUs. The prompts described above were presented to the models with the suffix, "Answer with one word." Answers were automatically identified using regular expressions. Responses starting with "yes", "yeah", "sure", "correct", "right", and "true" were coded as affirmative and those starting with "no", "nope", "wrong", "incorrect", and "false" were coded as negative. Other responses were treated as "null". For each model, four runs (of 3,069 prompts) were made.

4. Results

Overall results are shown in Figure 2. GPT-4 achieves almost perfect results (maximal lexical-syntactic flexibility) across all categories (count noun frames, mass noun frames, and transitive verb frames). The exception is the intransitive verb frames, where its performance, when nulls are removed, is worse than that of Mistral 7B. GPT-3.5 is consistently worse than GPT-4 but is, on balance, a stronger performer than the open-source models. Falcon 40B performs better on the metric than the other open-source models, on the prototypical condition. The glaring exception is the mass noun frames, where Falcon gen-

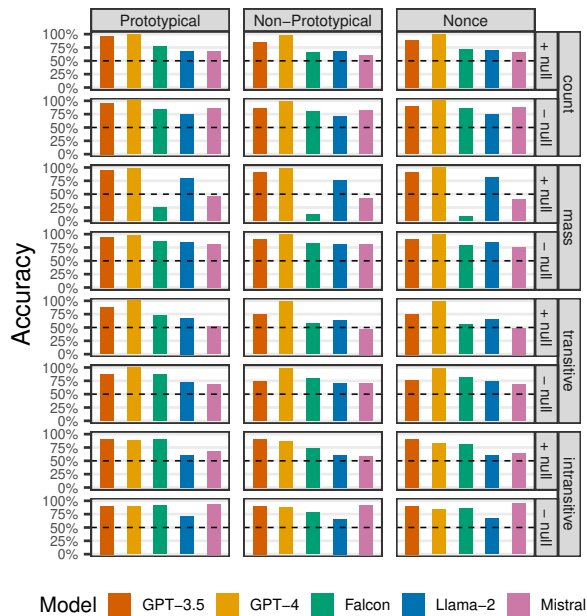


Figure 2: Average accuracy grouped by model and typicality (p- for “prototypical,” np- for “non-prototypical,” and no- for “nonce”)

erated mostly non-sequiturs rather than yes/no responses. When all responses are considered, Mistral is a relatively weak performer. However, when null responses are filtered out, Mistral appears to display greater flexibility than the other open-source models.

In order to separate the ability to carry out the natural language inference task from lexical-syntactic flexibility, we analyzed the difference between the average accuracy on the prototypical condition (expected part of speech) and the non-prototypical condition (unexpected part of speech). These results are shown in Table 3. The inaccuracy associated with non-prototypical contexts was substantially higher for GPT-3.5 and Falcon 40B, but was relatively low for Llama 2 70B and was minimal for Mistral 7B and GPT-4.

model	with nulls	without nulls
gpt-3.5	0.08	0.08
gpt-4	0.01	0.01
falcon	0.15	0.07
llama2	0.03	0.03
mistral	0.07	0.01

Table 3: Difference in average accuracy between the prototypical and non-prototypical condition, both with all records (left, null treated as negative) and with only non-null records (right)

5. Discussion

In order to understand which factors most contributed to performance on the syntactic flexibility task, we fitted a Logistic Regression to these results using the `Logit` function from the Python `statsmodels` library. The features were prototypical part of speech, model type, (proto)typicality of the filler given the frame, and whether the answer was yes, no, or “null”. It showed all factors to be significant predictors of correctness ($p < 0.01$), with answer type (yes, no, or null) as the strongest predictor. This is likely because Mistral and Falcon often fail to give correct responses by generating answers that are coded neither “yes” nor “no”. This is associated with a confounder (these frames, with two sentences, often elicited null responses). Controlling for all of the other factors, model type is also predictive, with GPT-3.5 and GPT-4 most associated with correct responses and Llama 2 least associated with correct responses.

Returning to our hypotheses in Table 4, we find that the models do, almost without exception, perform better under the prototypical condition than the non-prototypical condition. This suggests that conversion is more challenging than the use of unconverted words. However, non-prototypical performance is not significantly different from nonce performance, suggesting that the models are treating converted words as, essentially, nonce words. Scores for all three conditions are generally correlated with one another—models that are good at using words in a prototypical way are also good at using them in non-prototypical ways. Perhaps most surprising, though, is the fact that model size was not a good predictor of performance. The largest of the open models (Llama 2 70B) was in some ways the weakest performer. Mistral 7B, was the smallest, but held its own against Falcon 40B and even the much larger GPT models, and least in certain sub-tasks. This was true in spite of the fact that Mistral and Falcon were generally worse at following instructions.

Investigating the differences between the models in detail, it is clear that GPT-4 displays, far and away, the best scores on our lexical-syntactic flexibility task. It might be tempting to attribute this difference to the number of parameters in the model (since GPT-4 is believed to be a Mixture of Experts of several large models). However, it is not the case that this kind of generalization is necessarily simply a function of model size: the best-performing open-source model, when all else is held equal, is also the smallest (Mistral 7B). The largest of the open-source models—Llama 2 70B—is consistently mediocre in its performance

Hypothesis	Finding
prototypical performance > non-prototypical performance	✓ Supported
non-prototypical performance > nonce performance	✗ Not supported
Correlation between prototypical, non-prototypical, nonce performance	✓ Supported
Difference between model size accounts for difference in performance	✗ Not supported

Table 4: Findings with regard to each major hypothesis

on this task. And Falcon 40B, which is almost six times the size of Mistral, shows impressive abilities at the natural language inference task but lackluster performance at lexical-syntactic flexibility.

Mistral and Falcon’s scores are hurt by the fact that they frequently general null responses, particularly in response to frames eliciting mass nouns (Falcon and Mistral) and transitive verbs (Mistral). The causal mechanism, in these cases, is unclear. The mass noun frames all consist of two sentences: “I prefer more/less X. Do I prefer more/less X?” The other frames have one sentence each. The intransitive frames require the model to reason about semantically related words (though the count noun subtask does as well). The generation of null responses may be due either to these superficial factors or to more basic differences in model behavior with respect to frames that call for mass nouns or intransitive verbs.

GPT-4 also displays a dip in performance on the intransitive subtask, not because it is generating null responses but because it is generating incorrect yes/no responses in a non-trivial number of cases. Again, because the sets of frames are so small and lacking in diversity, it is not possible to construct a valid causal explanation to account for the fact that the models perform differently on some subtasks than others. What is clear, though, is that there are significant differences between the models and these differences are in some way correlated with the subtasks defined here.

6. Conclusion

We have introduced the first experiment testing the lexical-syntactic flexibility of LLMs, finding that language models are challenged by words in syntactically non-prototypical context (when compared to words in syntactically prototypical contexts). However, we did not find that words in syntactically non-prototypical contexts presented challenges to the models that nonce words did not. As we posited, there is a correlation between performance on prototypical and non-prototypical items and the model type was a significant predictor of performance. However, contrary to expectations, the model size was not a

good predictor of lexical-syntactic flexibility. The findings are summarized in Table 4.

With this foundation in place, we plan to investigate lexical-syntactic flexibility more systematically by using a much larger number of frames for each subtask and by testing a larger set of (open and proprietary) models. Now that truly open models like Olmo (Groeneveld et al., 2024) are available, it is possible to know, more precisely, what words have been seen by the model and in what contexts. This will allow us to state unambiguously when models are generalizing old vocabulary to new contexts and when they are directly recapitulating what they have seen in their training data.

7. Bibliographical References

- Laurie Bauer and Salvador Valera Hernández. 2005. *Approaches to conversion/zero-derivation*. Waxmann Verlag.
- Laurie Bauer and Salvador Valera. 2005. Conversion or zero-derivation: An introduction. *Approaches to conversion/zero-derivation*, pages 7–17.
- Robert Beard. 2017. *Derivation*, chapter 2. John Wiley & Sons, Ltd.
- Eve V Clark and Herbert H Clark. 1979. When nouns surface as verbs. *Language*, pages 767–811.
- Jan Don. 1993. *Morphological Conversion*. Ph.D. thesis, University of Utrecht.
- David Francis, Ella Rabinovich, Farhan Samir, David Mortensen, and Suzanne Stevenson. 2021. [Quantifying cognitive factors in lexical decline](#). *Transactions of the Association for Computational Linguistics*, 9:1529–1545.
- Adele E. Goldberg. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. University of Chicago Press, Chicago, IL.

- Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, and Hannaneh Hajishirzi. 2024. [Olmo: Accelerating the science of language models](#).
- Valentin Hofmann, Janet Pierrehumbert, and Hinrich Schütze. 2020a. [DagoBERT: Generating derivational morphology with a pretrained language model](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3848–3861, Online. Association for Computational Linguistics.
- Valentin Hofmann, Janet Pierrehumbert, and Hinrich Schütze. 2021. [Superbizarre is not superb: Derivational morphology improves BERT’s interpretation of complex words](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3594–3608, Online. Association for Computational Linguistics.
- Valentin Hofmann, Hinrich Schütze, and Janet Pierrehumbert. 2020b. [A graph auto-encoder model of derivational morphology](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1127–1138, Online. Association for Computational Linguistics.
- Paul Kiparsky. 1982. From cyclic phonology to lexical phonology. In H. van der Hulst and N. Smith, editors, *The Structure of Phonological Representations*, pages 131–175. Foris, Dordrecht.
- Rochelle Lieber. 1980. *On the organization of the lexicon*. Ph.D. thesis, University of New Hampshire.
- Hans Marchand. 1969. *The Categories and Types of Present-Day English Word-Formation*. C. H. Beck, München.
- Boris New, Christophe Pallier, Jessica Bourgin, and Julien Barra. in press. Unipseudo. <http://www.lexique.org/shiny/unipseudo/>. Accessed: 19.10.2023.
- Maria Ryskina, Ella Rabinovich, Taylor Berg-Kirkpatrick, David R. Mortensen, and Yulia Tsvetkov. 2020. Where new words are born: Distributional semantic analysis of neologisms and their semantic neighborhoods. In *Proceedings of the Society for Computation in Linguistics*, volume 3.
- Henry Sweet. 1891. *A new English grammar, logical and historical. Part I. Introduction, phonology, and accidence*. Clarendon Press, Oxford.
- Leonie Weissweiler, Valentin Hofmann, Anjali Kantharuban, Anna Cai, Ritan Dutt, Amey Hengle, Anubha Kabra, Atharva Kulkarni, Abhishek Vijayakumar, Haofei Yu, Hinrich Schütze, Kemal Oflazer, and David R. Mortensen. 2023. Counting the bugs in ChatGPT’s wugs: A multilingual investigation into the morphological capabilities of a large language model. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online. Association for Computational Linguistics.

8. Language Resource References