

Language Models Can Learn Exceptions to Syntactic Rules

Cara Su-Yi Leong Tal Linzen
New York University
{caraleong, linzen}@nyu.edu

Abstract

Artificial neural networks can generalize productively to novel contexts. Can they also learn exceptions to those productive rules? We explore this question using the case of restrictions on English passivization (e.g., the fact that “The vacation lasted five days” is grammatical, but “*Five days was lasted by the vacation” is not). We collect human acceptability judgments for passive sentences with a range of verbs, and show that the probability distribution defined by GPT-2, a language model, matches the human judgments with high correlation. We also show that the relative acceptability of a verb in the active vs. passive voice is positively correlated with the relative frequency of its occurrence in those voices. These results provide preliminary support for the entrenchment hypothesis, according to which learners track and uses the distributional properties of their input to learn negative exceptions to rules. At the same time, this hypothesis fails to explain the magnitude of unpassivizability demonstrated by certain individual verbs, suggesting that other cues to exceptionality are available in the linguistic input.

1 Introduction

Many studies have demonstrated language models’ ability to extend a generalization from a small set of examples to novel lexical items, structures, and contexts, even if the models do not always do so in a human-like way (Hupkes et al., 2020; Kim and Linzen, 2020; Lake and Baroni, 2018; McCoy et al., 2018). These studies show that models can substitute novel lexical items into rules where those items were previously unseen. At the same time, language models can sometimes *over-generalize*, for instance by producing a literal, compositional translation of idiomatic expressions like *kick the bucket* when humans would not (Dankers et al., 2022). A full evaluation of language models’ generalization abilities should thus not only measure

whether models can generalize when humans do, but also whether models are able to *constrain* their generalizations when humans do.

We address this question by building on a line of work that probes whether human-like acceptability judgments for argument structure alternations can be predicted from the probability distribution that a from language model defines over sentences. This studies have shown, for example, that the GPT-2 language model (Radford et al., 2019) can match human judgments about whether the dative alternation applies to a verb (Hawkins et al., 2020), and that information about which syntactic frames a verb can appear in (e.g. whether a verb participates in the SPRAY/LOAD alternation) can be recovered from the verb’s contextualized representations and from sentence embeddings (Kann et al., 2019).

In this work, we evaluate models’ ability to identify exceptions using the case study of the English passive.¹ The passive voice is highly productive in English; most strikingly, young children exposed to novel verbs in the active voice are able to understand and produce passive constructions using those verbs (Pinker et al., 1987; Brooks and Tomasello, 1999). This suggests that English speakers do not in general conclude that verbs that they have never encountered in the passive voice are unacceptable in that voice. Yet there are limits to the productivity of the English passive; examples such as (1) have been reported to be unacceptable in the passive voice:

- (1) a. The vacation lasted five days.
b. *Five days was lasted by the vacation.

Sentences like (1b) are unlikely to occur productively in natural speech—just like passives of infrequent verbs. Yet even though they do not receive

¹Data and code are available at <https://github.com/craaaa/exceptions>.

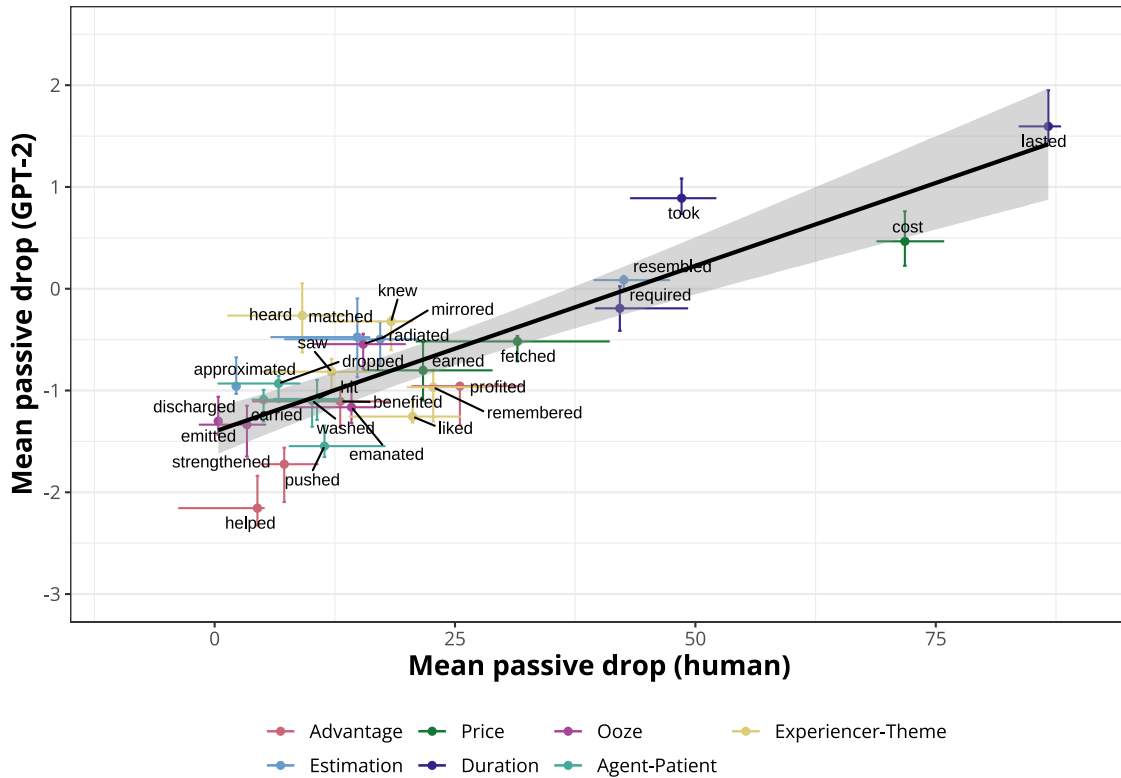


Figure 1: *Passive drop in humans vs. GPT-2* — A GPT-2 model trained on 100M words approximately predicts variable amounts of passive drop equivalent to human judgments. Horizontal and vertical error bars indicate bootstrapped 95% confidence intervals.

explicit evidence that these sentences are unacceptable, rather than simply rare, English speakers nonetheless learn that they constitute exceptions, and do not judge (1b) to be acceptable.

How do humans acquire such exceptions? The **entrenchment hypothesis** suggests that speakers track and use the distributional properties of their input as indirect negative evidence for the existence of an exception (Braine and Brooks, 1995; Regier and Gahl, 2004; Theakston, 2004). For instance, if an English learner never encounters the verb *last* in the passive voice, despite having seen *last* used productively in the active voice, they may conclude that *last* cannot occur in the the passive voice. Are language models—which do not have access to human feedback or syntactic supervision, and are trained solely to perform next-word prediction—attentive to the same information that humans are when determining the extent to which syntactic rules can generalize?

In this paper, we tackle these questions by com-

paring human acceptability judgments on sentences containing verbs that are exceptional in the passive voice, on the one hand, to the probability distribution defined by a GPT-2-like model trained on a 100-million word English corpus. We find that the language model matches human acceptability judgments on active and passive sentences to a large degree (Figure 1), suggesting that language models can constrain their syntactic generalizations in a human-like way. Using our model’s training corpus, we further show that there is a weak but positive correlation between the relative frequency of actives and passives in the input and their relative acceptability. Together, these empirical results suggest that the linguistic input contains useful information from which exceptions to syntactic generalizations can be learned.

2 Restrictions on passivization

Although the English verbal passive is highly productive, not all verbs can occur in the passive. For

Verb class	Active sentence	Passive sentence
Advantage	Your investment ___ the community.	The community was ___ by your investment.
Price	Your book ___ thirty dollars.	Thirty dollars was ___ by your book.
Ooze	That machine ___ a sound.	A sound was ___ by that machine.
Duration	The journey ___ three days.	Three days was ___ by the journey.
Estimation	Your drawing ___ her likeness.	Her likeness was ___ by your drawing.

Table 1: *Example sentence frames* — Each verb in the verb class was substituted into frames specific to the class.

instance, intransitive and middle verbs resist passivization in general (Perlmutter, 1978; Zaenen, 1993). In this paper, we focus on passives of transitive verbs that occur with by-phrases. These long passives are clauses of the form given in (2), which in most cases have an uncontroversially acceptable passive form:

- (2) a. The ball was hit by the boy.

A small list of lexical exceptions have been described for which the passive voice is deemed ungrammatical (Levin, 1993; Postal, 2004). Some of these exceptions can be classed together based on the semantics of the verb or types of arguments the verb takes. For instance, verbs that take measure phrases as their object reportedly do not occur in the passive:

- (3) a. That house costs fifty thousand dollars.
 b. *Fifty thousand dollars is/are cost by that house.

(Hale and Keyser, 1997, 17-8)

Even within a particular verb class, passivizability may also be an idiosyncratic characteristic of individual lexical items (Zwicky, 1987): verbs which can be substituted for each other in any other syntactic context may differ in their ability to passivize. Thus, for instance, although in the active voice *matched*, *mirrored*, *approximated* and *resembled* can occur in the same environment, (4a) is grammatical, while (4b) is not.

- (4) a. Kim is matched/mirrored/approximated by the model in nearly every detail.
 b. *Kim is resembled by the model in nearly every detail. (Zwicky, 1987)

We may thus expect differences in passivizability not only between verbs with different semantics and argument frames, but also among verbs with very similar meaning.

3 Human Acceptability Judgments

In order to test whether language models follow a human-like generalization patterns, we need to first characterize the human judgment pattern, which will serve as the target of modeling. In this section, we report on an acceptability judgment study whose goal was to verify the judgments from the syntax literature and measure any gradient differences in the degree to which different verbs can be passivized.

3.1 Materials

We identified five verb classes containing verbs that have been reported to be unpassivizable (Levin, 1993; Postal, 2004; Zwicky, 1987):

- **Advantage** verbs: *benefit, help, profit, strengthen*
- **Price** verbs: *cost, earn, fetch*
- **Ooze** verbs: *discharge, emanate, emit, radiate*
- **Duration** verbs: *last, require, take*
- **Estimation** verbs: *approximate, match, mirror, resemble*

Each of these class includes verbs with similar semantics that can be substituted into the same position in a sentence in the active voice. While some of these verbs can be used in other senses, we tested the specific sense that was reported in the literature by controlling the sentence frames used. Five past-tense sentence frames were constructed for each verb class (Table 1).

Each of the verbs in the class was substituted into the sentence frame, resulting in 90 total test sentence pairs. Example (5) demonstrates a sentence pair generated from the sentence frame in Table 1 using the verb *matched*:

- (5) a. Your friend matched my brother.
 b. My brother was matched by your friend.

As control verbs, we also selected five agent-patient and five experiencer-theme verbs; we expect these verbs to be passivizable:

- **Agent-Patient:** *hit, push, wash, drop, carry*
- **Experiencer-Theme:** *see, hear, know, like, remember*

Because of the varied semantics of the verbs in these groups, unique sentence pairs were created for each verb, yielding 50 control sentence pairs. An example of a sentence pair for the verb *push* is given in (6):

- (6) a. A boy pushed the cup.
 b. The cup was pushed by a boy.

Each participant only saw either the active or the passive of a sentence pair. The 140 sentence pairs (90 test + 50 control) were divided into two buckets of 70 sentence pairs each such that each bucket contained two or three sentence frames per verb. Each bucket was then further divided into groups of 70 sentences such that the active and passive forms of a sentence pair were in different groups. Each group of sentences contained one quarter of the test and control stimuli (70 sentences).

Presentation order was counterbalanced by making four ordered lists for each group. Each group was organized into two lists such that an item that appeared in the first half of one list appeared in the second half of the other list. The order of items was pseudorandomized within those lists to ensure that not more than two active or passive sentences and no two sentences within the same verb class were seen in succession. These lists were then reversed, so that a total of four ordered sentence lists were made per sentence group.

Additionally, every experimental trial alternated with a filler sentence. Filler sentences were also used as attention checks. We used 24 grammatical and 46 ungrammatical filler sentences: since the passives of control sentences were expected to be acceptable, the greater number of ungrammatical fillers was intended to balance the experimental stimuli. The full set of materials is available in Appendix A.

3.2 Participants

We recruited 84 participants who had IP addresses located in the US and self-reported as native English speakers via the crowdsourcing platform Prolific. Each participant rated 140 sentences (70 test + 70 filler) and was paid US\$3.50. The experiment took approximately 12 minutes to complete.

Participants were asked to rate how acceptable each sentence sounded based on their gut reaction. They were told that there were no right or wrong answers. Participants rated sentences by moving a slider from “Completely unacceptable” to “Completely acceptable”, which corresponded to an integer score (invisible to them) between 0 and 100. They were not able to rate a sentence with a score of 50. Two practice sentences (one ungrammatical, one grammatical) were used to familiarize participants with the paradigm.

Participants were excluded from the results if they answered more than 15 filler questions unexpectedly, either by giving ungrammatical sentences scores above 50 or giving grammatical sentences scores below 50. We excluded 10 participants from analysis for this reason.

3.3 Results

We calculate the **passive drop** of an item as the difference in mean acceptability ratings between its active and passive version. The results are reported in Figure 2; a steeper downward gradient corresponds with a larger passive drop. Since corresponding active and passive sentences contain the same lexical items except for the auxiliary *was* and *by*, which are common across all sentences, directly comparing active and passive sentences isolates the effect of the passivization from lexical effects that might increase the acceptability of sentences with more common verbs like *helped* over low-frequency verbs like *profited*.

Across all verb classes, there was a significant difference between scores given to active and passive sentences. This difference may be accounted for by pragmatic factors: although the passive construction is more pragmatically marked than the active (Comrie, 1988), each sentence in the acceptability judgment task was presented to participants without establishing a relevant context. This setting might have caused participants to rate passive sentences as worse than their active counterparts.

Although the passive drop was positive for all verbs, its **magnitude** differed across verb classes. The *duration* class showed the largest mean passive drop (59.4 points), and the *ooze* class showed the lowest mean passive drop (8.0 points) among the test verb classes.

We fit a linear mixed-effects model to predict SENTENCE SCORE using the *agent-patient* verb class as the baseline. We used SENTENCE TYPE

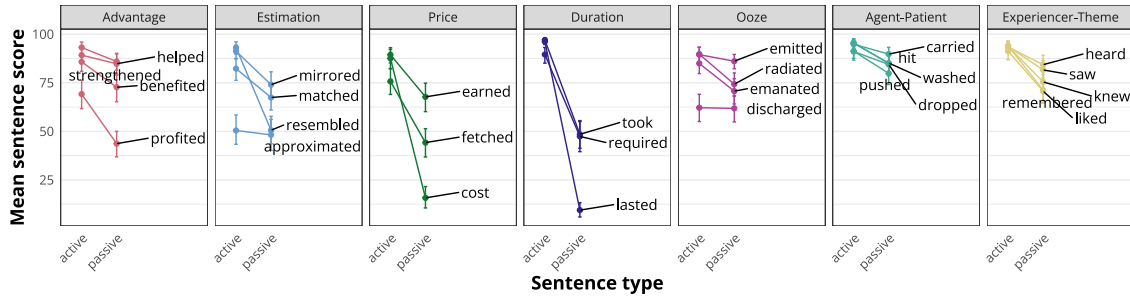


Figure 2: *Passive drop in human acceptability judgments of active and passive sentences by verb* — The steeper the downward gradient between active and passive conditions, the larger the passive drop. Error bars indicate bootstrapped 95% confidence intervals.

and VERB CLASS as well as their interaction as fixed effects and FRAME, PARTICIPANT and VERB as random intercepts. We found a significant difference between **agent-patient** verbs and three other verb classes: **estimation** verbs ($p = 5.74e-06$), **price** verbs ($p < 2e-16$), and **duration** verbs ($p < 2e-16$). On the other hand, there was no significant difference in the sentence scores obtained from **agent-patient** verbs and **ooze** verbs, **advantage** verbs, or **experiencer-theme** verbs as a class.

Within each verb class, some verbs were more passivizable than others. For example, *last* was significantly less passivizable than *took* and *required*, and *cost* was less passivizable than *fetched*. Similarly, while *resembled* had a high passive drop, the remaining verbs in the **estimation** class showed relatively low passive drops. These results validate the claim that some verbs may be more passivizable than others despite sharing similar paradigmatic relationships (Zwicky, 1987).

In summary, the human acceptability judgment experiment demonstrated that some verbs in the verb classes being tested are degraded in the passive voice, and that unacceptability was gradient between verbs. For a model to adequately approximate such behaviour, it must exhibit the following characteristics:

- **Exceptionality:** some verbs (e.g. **duration** verbs) exhibit passive drops that are significantly different from the baseline passive drop expected of the canonically passivizable **agent-patient** verbs.
- **Gradience:** (un)acceptability is gradient, with some verbs on average exhibiting higher passive drop than others.

4 Comparison with Language Models

With the quantitative human acceptability judgment data in hand, we now turn to evaluating language models. If distributional data is sufficient to learn the extent to which verbs are unacceptable in the passive, we expected GPT-2 to be able to match human judgments on both passivizable verbs and unpassivizable verbs. We also expect GPT-2 to be able to match the relative gradience of passive drop that humans display.

4.1 Method

We evaluated GPT-2 (Radford et al., 2019), a Transformer (Vaswani et al., 2017) language model. We tested four different pre-trained GPT-2 models, which differed in their number of parameters and number of layers, but were trained on the same data. Each model was trained on Open-AI’s WebText corpus, which contains 40GB of data — approximately 8B words, assuming each word contains an average of 5 bytes/chars. Pre-trained GPT-2 models have performed well on targeted syntactic evaluations requiring knowledge of argument structure, such as differentiating between verbs that participate in the causative alternation and those that do not (Warstadt et al., 2020).

The GPT-2 models available for download are trained on a much larger corpus than is realistic for any human to be exposed to (Linzen, 2020). English-speaking children are exposed to 2–7M words per year (Gilkerson et al., 2017), or 26M–91M words by the age of 13. Rounding to the nearest order of magnitude, we trained a GPT-2 model on a 100M word subset of the OpenWebText corpus (Gokaslan and Cohen, 2019), an open-source reproduction of the Web Text corpus; this simulates more closely the amount of linguistic input

a human may receive (though not its genre). We trained five iterations of this model, which we call **GPT2-100M**, using different random seeds and report averages of the results obtained from these five models.

We adapted the targeted syntactic evaluation paradigm (Linzen et al., 2016; Lau et al., 2017; Warstadt et al., 2019) to compare the language models to humans. This paradigm involves obtaining model “judgments” for minimal pairs of sentences. For each sentence, a score is obtained by summing the log-probabilities assigned to each token in the sentence, which gives the probability the model assigns to that sentence. We conclude that a model’s distribution is consistent with human judgments if it assigns a higher probability to the acceptable sentence than to the corresponding unacceptable one. Unlike some prior work, we collected numeric scores instead of binary acceptability judgments: we calculated a gradient passive drop of each sentence pair by subtracting the score of the active sentence from the score of its passive counterpart.

Since we compared active sentences to long passives, which contain by-phrases, every passive sentence contained two more words than its active counterpart. A sentence with more tokens will on balance be less probable than a sentence with fewer tokens; we thus normalized each sentence score by dividing it by the number of tokens in the sentence (Lau et al., 2017). Doing so accounts for the effect of sentence length on the sentence score, and also allows us to compare sentences where words are split into separate tokens during GPT-2’s tokenization process, e.g. approximated \rightarrow approx + imated.

4.2 Results

The four pre-trained models as well as the five GPT2-100M models showed positive correlations between mean human passive drop and mean model passive drop, reported in Table 2. For pre-trained GPT-2 models, we calculate mean model passive drop for each verb by averaging over the passive drop of all five sentence frames. For GPT2-100M, we calculate the average passive drop of each verb over all sentence frames across the five versions of the model (trained with different random seeds); we report these results as **GPT2-100M-avg**.

The results were qualitative similar for all models (Figure 3); in what follows, we focus on GPT2-

Model	# parameters	r_s
GPT2-100M-avg	124M	0.709
GPT2	124M	0.659
GPT2-med	345M	0.385
GPT2-large	774M	0.549
GPT2-xl	1558M	0.559

Table 2: *GPT2 model parameters and correlation coefficients* — in all five models, a correlation was found between human passive drop and the model’s passive drop, but it was stronger for smaller models, and strongest for the models trained on only 100M words.

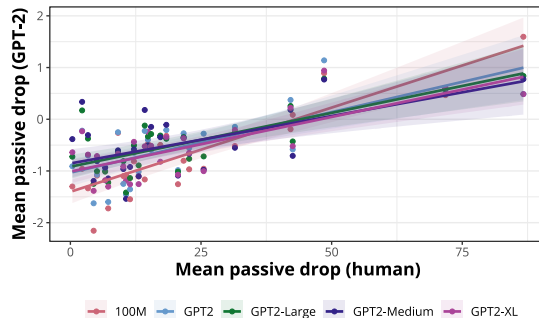


Figure 3: *Passive drop of different-sized GPT-2 models compared to human judgments* — Each point in represents a single verb. Models differed in number of parameters and/or training data, but showed qualitatively similar passive drops.

100M-avg, whose behaviour showed the strongest correlation with human judgments. These models are also trained on the most cognitively realistic corpus.

Figure 1 plots GPT2-100M-avg’s passive drop against the passive drop observed in the human experiment. A strong correlation was found between the passive drop in the models’ sentence scores and human passive drop ($r_s = 0.709$), suggesting that predictions learned from linguistic input match human **gradient** judgments on passivization relatively well.

GPT2-100M-avg also matched humans’ judgments of **exceptionality** within verb classes: among verbs with similar meanings, both humans and the model identified the same verbs as being less passivizable. In verbs for which humans demonstrated low passive drop, such as *strengthened* and *discharged*, close to no passive drop was observed in the model’s predictions. GPT2-100M-avg also predicted high passive drops for verbs like

Frequency of selected verbs in active and passive sentences

Corpus: 100M corpus

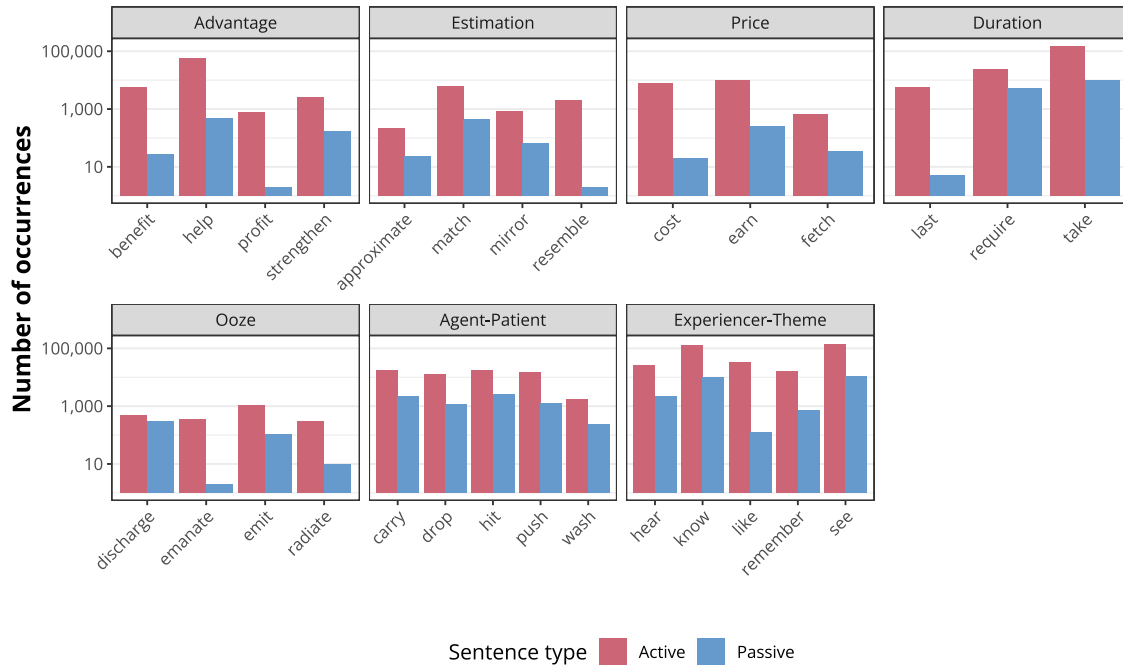


Figure 4: Occurrence of active transitive and passive sentences using test verbs in the training corpus — sentences whose verb had a passive dependent (csubjpass, nsubjpass, or auxpass) were tagged as passive, while all other instances of the verb were tagged as active.

lasted, *resembled* and *cost*, aligning with human judgments that these verbs are unique in their verb class.

5 Does Frequency Explain Passivizability Judgments?

Having established that a language model can successfully model humans’ gradient passivizability judgments, we now examine the extent to which GPT2-100M’s passivization judgments correlate with the distributional properties of its training data. Specifically, we explore the utility of the entrenchment hypothesis in explaining GPT2-100M’s gradient judgments of passivization. Recall that this hypothesis argues that learners conclude that a verb cannot appear in a particular context if it appears in many other contexts but systematically fails to appear in the context in question.

Here, we consider a weaker version of the entrenchment hypothesis, which does not presuppose that exceptions *never* occur in the learner’s input. Instead, we hypothesize that the less frequently a verb is used in the passive voice relative to the active voice, the less acceptable passive constructions

using that verb will be.

5.1 Method

We conducted a corpus study on the data that GPT2-100M was trained on. We processed each document in the corpus using the spaCy Transformer-based lemmatizer, POS tokenizer and dependency parser (Honnibal et al., 2020) and extracted all sentences that contained a verbal lemma corresponding to the test and control verbs. Sentences that contained the verbs in question and had a dependency edge to a passive auxiliary (auxpass), a passive nominal subject (nsubjpass) or a passive clausal subject (csubjpass) were classified as passive sentences, while all other sentences containing the verb were classified as active sentences. We hand-checked a 1000 sentence subset of the training data to verify the accuracy of the tagging process. No sentences were incorrectly tagged in the manually verified subset, although the corpus did contain instances of typos such as (7) (tagged as passive):

- (7) It was fun while it was lasted.

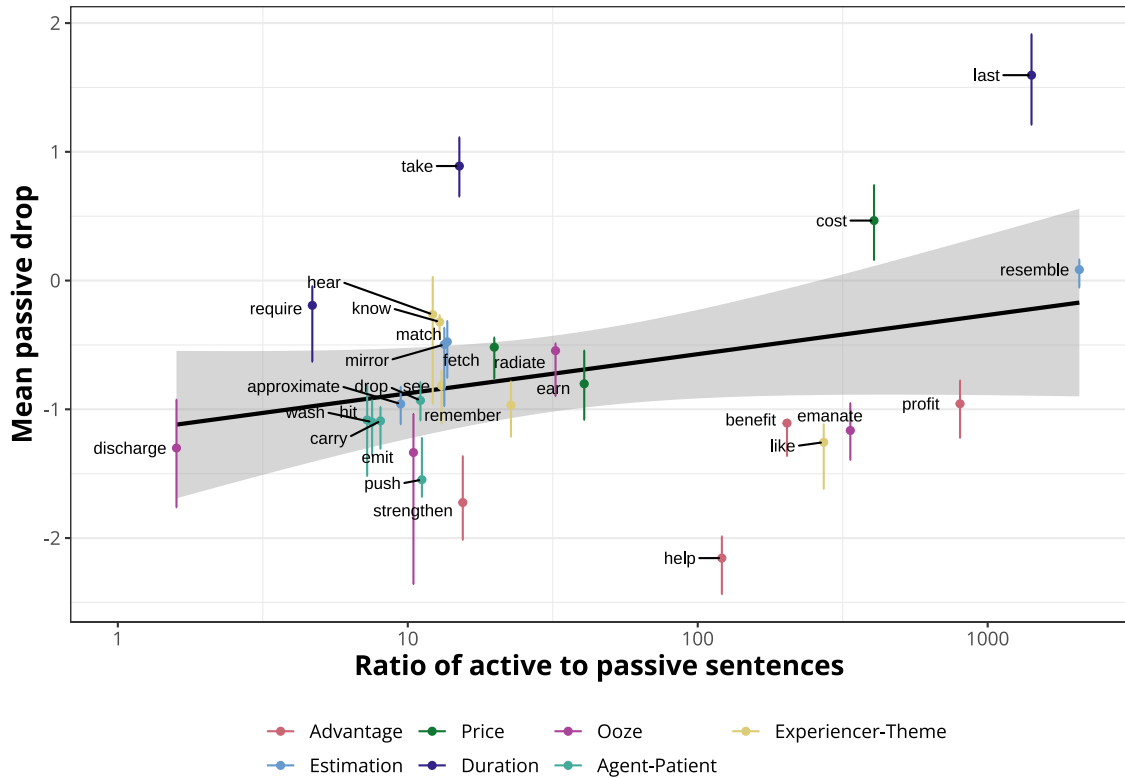


Figure 5: *GPT2-100M*'s passive drop against the ratio of active to passive sentences in its training corpus. Error bars indicate bootstrapped 95% confidence intervals across sentence frames.

5.2 Results

Figure 4 shows the number of active and passive sentences in *GPT2-100M*'s training corpus.

Not all verbs appear in the same ratios in the active and passive voice. Agent-patient verbs consistently appeared in approximately 10 times as many active sentences as passive sentences, matching estimates from previous corpus studies (Roland et al., 2007). On the other hand, test verbs appeared in varying amounts in the active and passive. For instance, *last* appeared 5666 times in the active and four times in the passive in the 100M word corpus, while *cost* appeared 7706 times in the active and 19 times in the passive. This result suggests that the test verbs differ from canonically passivizable control verbs in their distribution.

Figure 5 graphs the correlation between the ratio of active to passive sentences for a given verb, on the one hand, and that verb's mean passive drop on the other hand. We find a weak but positive correlation between the two variables ($r_s = 0.212$).

Two key outliers that are not well accounted

for by this measure of relative frequency are *last* and *cost*. In both humans and model judgments, these verbs demonstrated high passive drops; yet, they are similar in relative frequency of active and passive to verbs like *emanate*, *profit* and *resemble*, whose passive drops are lower. While frequency seems to predict some amount of unpassivizability, then, it cannot account for the full magnitude of the passive drop displayed by these particular verbs.

Furthermore, entire verb classes are systematically over- or under-predicted in Figure 5. The *duration* verb class on the whole has a high passive drop relative to its frequency in the corpus, while frequency over-predicts the passive drop expected for the *advantage* verb class. We thus conclude that while the relative frequency of active and passive voice sentences positively correlates with passive drop, other factors are likely to also be relevant on a verb-class level.

Although *take* appears to be an outlier in Figure 5, with an active to passive ratio similar to that of the *agent-patient* and *experiencer-theme* con-

trol verbs, the measure of frequency we used does not take into account the fact that *take* has multiple senses. If a different sense than the one being tested is heavily represented by passive sentences, the number of passives counted may be overestimated. For example, although we only test the duration sense of *take*, as given in (8a), the sense used in (8b) may be more prevalent in the corpus:

- (8) a. *Two days was taken by the meeting.
b. The photo was taken by the boy.

These differences in verb sense are not accounted for in the current corpus study; future work should make use of word sense disambiguation to conduct more targeted corpus analyses. Additionally, the issue of differentiating verb senses in polysemous verbs is one that both human and machine learners face, raising the question of the extent to which learners differentiate between verb senses that are more or less difficult to passivize.

Overall, while the relative frequency of a verb’s occurrence in the active and passive does positively correlate with its unpassivizability, it does not account for crucial verb-level differences in the magnitude of passive drop demonstrated by GPT2-100M-avg.

6 Discussion

The goal of this study was to explore whether a language model can identify exceptions to a productive syntactic rule in a human-like way. We compared human acceptability judgments to sentence scores produced by a GPT-2 model trained on the amount of linguistic input that a human can plausibly be exposed to, and found that the model displayed human-like exceptionality and gradience in its judgments of passive sentences. The results of our study suggest that language models are able to refrain from over-generalizing to exceptions. Our results suggest that future empirical inroads may be made towards understanding the mechanisms and input required to overcome the projection problem (Baker, 1979), i.e. the problem of acquiring arbitrary negative exceptions, using language models as experimental subjects.

We took a first step in this direction by showing a positive correlation between the relative frequency of active and passive sentences containing a given verb and the difference between that verb’s acceptability in the active and passive voice (i.e.

its passive drop) in GPT-2. Although our results lend some credence to the entrenchment hypothesis, they suggest that additional factors must be recruited to explain the full magnitude of exceptionality displayed by highly unpassivizable verbs such as *last* and *cost*.

Moreover, although we demonstrated that the relative frequency of a verb’s occurrence in the active and passive is correlated with its passive drop, a causal relationship between the two cannot be established from our data. A single underlying factor, such as verbal semantics, may affect both the frequency of a verb in the passive in relation to the passive *and* its acceptability in the passive construction.

Future research should test the causal impact of a verb’s absolute and relative frequency in the training corpus on its predicted passivizability. Following Wei et al. (2021), we plan to create an altered training dataset where we match the frequency of active and passive sentences containing passivizable verbs like *drop* to the absolute frequency of sentences containing highly unpassivizable verbs, such as *last*. Comparing models trained on this dataset against GPT2-100M will allow us to move beyond a correlational analysis and explore whether altering the frequency of a verb in the active and passive voice in a model’s training data has a causal effect on the model’s predictions of that verb’s passivizability.

7 Conclusion

In this paper, we explored whether a language model trained on a human-scale amount of linguistic input is able to learn lexical exceptions to a productive syntactic generalization in English. We showed that it was able to match humans’ reported judgments on unpassivizable verbs like *last*, showing both the ability to identify exceptions as well as to identify the magnitude of an exception. We also demonstrated a weak correlation between the degree to which a model prefers active over passive sentences using a given verb, on the one hand, and the ratio between the frequencies with which sentences containing that verb occur in the active and passive voice, on the other hand. Together, these results suggest that distributional information plays a role in learning exceptions to syntactic rules.

Acknowledgments

We would like to thank Alec Marantz and Gary Thoms for valuable comments and suggestions related to this paper. This material is based upon work supported by the National Science Foundation (NSF) under Grant No. BCS-2114505. This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise.

References

- C. L. Baker. 1979. [Syntactic Theory and the Projection Problem](#). *Linguistic Inquiry*, 10(4):533–581.
- Martin D S Braine and Patricia J Brooks. 1995. Verb Argument Structure and the Problem of Avoiding an Overgeneral Grammar. In Michael Tomasello and William Edward Merriman, editors, *Beyond Names for Things: Young Children’s Acquisition of Verbs*. L. Erlbaum, Hillsdale, N.J.
- Patricia J. Brooks and Michael Tomasello. 1999. [Young children learn to produce passives with nonce verbs](#). *Developmental Psychology*, 35(1):29.
- Bernard Comrie. 1988. [Passive and voice](#). In Masayoshi Shibatani, editor, *Passive and Voice*, Typological Studies in Language, pages 9–24. John Benjamins Publishing Company.
- Verna Dankers, Christopher Lucas, and Ivan Titov. 2022. [Can Transformer be Too Compositional? Analysing Idiom Processing in Neural Machine Translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3608–3626, Dublin, Ireland. Association for Computational Linguistics.
- Jill Gilkerson, Jeffrey A. Richards, Steven F. Warren, Judith K. Montgomery, Charles R. Greenwood, Oller D. Kimbrough, John H. L. Hansen, and Terrance D. Paul. 2017. [Mapping the Early Language Environment Using All-Day Recordings and Automated Analysis](#). *American Journal of Speech-Language Pathology*, 26(2):248–265.
- Aaron Gokaslan and Vanya Cohen. 2019. OpenWeb-Text corpus.
- Ken Hale and Samuel Jay Keyser. 1997. Adjectives, other stative predicates and the roots of stativity.
- Robert Hawkins, Takateru Yamakoshi, Thomas Griffiths, and Adele Goldberg. 2020. [Investigating representations of verb bias in neural language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4653–4663, Online. Association for Computational Linguistics.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spaCy: Industrial-strength natural language processing in python.
- Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. 2020. [Compositionality Decomposed: How do Neural Networks Generalise? \(Extended Abstract\)](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pages 5065–5069, Yokohama, Japan. International Joint Conferences on Artificial Intelligence Organization.
- Katharina Kann, Alex Warstadt, Adina Williams, and Samuel R. Bowman. 2019. [Verb argument structure alternations in word and sentence embeddings](#). In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*, pages 287–297.
- Najoung Kim and Tal Linzen. 2020. [COGS: A compositional generalization challenge based on semantic interpretation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9087–9105, Online. Association for Computational Linguistics.
- Brenden M. Lake and Marco Baroni. 2018. Generalization without Systematicity: On the Compositional Skills of Sequence-to-Sequence Recurrent Networks. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmmsäsan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 2879–2888. PMLR.
- Jey Han Lau, Alexander Clark, and Shalom Lappin. 2017. [Grammaticality, Acceptability, and Probability: A Probabilistic View of Linguistic Knowledge](#). *Cognitive Science*, 41(5):1202–1241.
- Beth Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press, Chicago.
- Tal Linzen. 2020. [How Can We Accelerate Progress Towards Human-like Linguistic Generalization?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217, Online. Association for Computational Linguistics.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the Ability of LSTMs to Learn Syntax-Sensitive Dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- R. Thomas McCoy, Robert Frank, and Tal Linzen. 2018. Revisiting the poverty of the stimulus: hierarchical generalization without a hierarchical bias in recurrent neural networks. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society, CogSci 2018*, Proceedings of the 40th Annual Meeting of the Cognitive Science Society, CogSci 2018, pages 2096–2101. The Cognitive Science Society.

- David M. Perlmutter. 1978. *Impersonal Passives and the Unaccusative Hypothesis*. *Annual Meeting of the Berkeley Linguistics Society*, 4(0):157–190.
- Steven Pinker, David S. Lebeaux, and Loren Ann Frost. 1987. *Productivity and constraints in the acquisition of the passive*. *Cognition*, 26(3):195–267.
- Paul Martin Postal. 2004. *Skeptical Linguistic Essays*. Oxford University Press, Oxford ; New York.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. *Language Models are Unsupervised Multitask Learners*. Technical report, OpenAI.
- Terry Regier and Susanne Gahl. 2004. *Learning the unlearnable: The role of missing evidence*. *Cognition*, 93(2):147–155.
- Douglas Roland, Frederic Dick, and Jeffrey L. Elman. 2007. *Frequency of basic English grammatical structures: A corpus analysis*. *Journal of Memory and Language*, 57(3):348–379.
- Anna L Theakston. 2004. *The role of entrenchment in children’s and adults’ performance on grammaticality judgment tasks*. *Cognitive Development*, 19(1):15–34.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. *Attention is All you Need*. In *Advances in Neural Information Processing Systems (NIPS 2017)*, volume 30. Curran Associates, Inc.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. *BLiMP: The Benchmark of Linguistic Minimal Pairs for English*. *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. *Neural network acceptability judgments*. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Jason Wei, Dan Garrette, Tal Linzen, and Ellie Pavlick. 2021. *Frequency effects on syntactic rule learning in transformers*. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 932–948, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Annie Zaenen. 1993. *Unaccusativity in Dutch: Integrating Syntax and Lexical Semantics*. In James Pustejovsky, editor, *Semantics and the Lexicon*, Studies in Linguistics and Philosophy, pages 129–161. Springer Netherlands, Dordrecht.
- Arnold M. Zwicky. 1987. *Slashes in the passive*. *Linguistics*, 25(4).

A Stimuli

A.1 Test sentence frames

Verb class	Sentence frame
Advantage	Your investment ___ the community.
	The exercise ___ his fitness.
	Our friendship ___ my life.
	The law ___ these workers.
Price	The treaty ___ both countries.
	Your dish ___ ninety dollars.
	The painting ___ a fortune.
	The tickets ___ a lot of money.
Ooze	Your book ___ thirty dollars.
	His actions ___ the medal.
	My friend ___ confidence.
	The lightbulb ___ some light.
Estimation	That machine ___ a sound.
	The teacher ___ wisdom.
	The trash ___ an odor.
	Your drawing ___ her likeness.
Duration	Your friend ___ my brother.
	The character ___ the author.
	Her son ___ her father.
	The copy ___ the original.
Duration	The journey ___ three days.
	My meeting ___ two hours.
	The interview ___ some time.
	Her speech ___ seventeen minutes.
Duration	His trek ___ a month.

A.2 Agent-patient sentences

Verb	Active sentence
hit	My brother hit your friend.
	Your sister hit the target.
	The child hit the ball.
	A boy hit my bag.
	A monkey hit the toy.
pushed	My brother pushed a child.
	The mother pushed my toy.
	A boy pushed the cup.
	A child pushed the bag.
	Your sister pushed your friend.
washed	A boy washed the cup.
	A child washed the bag.
	My sister washed a towel.
	My brother washed my plate.
	Your mother washed my toy.
dropped	My brother dropped my plate.
	The mother dropped my toy.
	A boy dropped the cup.
	A child dropped the bag.
	Your sister dropped a book.
carried	A boy carried my bag.
	Your mother carried the child.
	My brother carried your friend.
	The dog carried the toy.
	The donkey carried the load.

A.3 Experiencer-theme sentences

Verb	Active sentence
saw	My brother saw your friend.
	Your dog saw the toy.
	Your sister saw a book.
	A boy saw my bag.
	The child saw a monkey.
heard	A boy heard the sound.
	The child heard the rules.
	My brother heard your friend.
	Your dog heard the toy.
	Your sister heard a squeak.
knew	My brother knew your friend.
	Your dog knew my cat.
	Your sister knew my brother.
	A boy knew my mother.
	The mother knew the dog.
liked	A boy liked the game.
	The child liked a monkey.
	My brother liked your friend.
	Your dog liked the toy.
	Your sister liked a book.
remembered	My brother remembered your friend.
	Your dog remembered my toy.
	Your sister remembered a book.
	A boy remembered the game.
	The child remembered the rules.