

Linguistic Properties of Truthful Response

Bruce W. Lee, Benedict Florance Arockiaraj, Helen Jin

University of Pennsylvania - PA, USA

{bruce1ws, benarock, helenjin}@seas.upenn.edu

Abstract

We investigate the phenomenon of an LLM’s untruthful response using a large set of 220 handcrafted linguistic features. We focus on GPT-3 models and find that the linguistic profiles of responses are similar across model sizes. That is, how varying-sized LLMs respond to given prompts stays similar on the linguistic properties level. We expand upon this finding by training support vector machines that rely only upon the stylistic components of model responses to classify the truthfulness of statements. Though the dataset size limits our current findings, we show the possibility that truthfulness detection is possible without evaluating the content itself. But at the same time, the limited scope of our experiments must be taken into account in interpreting the results.

1 Introduction

It is widely accepted that larger language models tend to be more fluent in natural language (Zhao et al., 2023; Brown et al., 2020). But at the same time, there is convincing evidence that larger language models do not always generate more truthful answers (Lin et al., 2022). For instance, there are cases where large language models (LLM) provide nonfactual but seemingly plausible predictions, often called hallucinations (Mialon et al., 2023; Welleck et al.). Such a phenomenon of unfaithful responses has been a research topic for many (Manakul et al., 2023; Bang et al., 2023). Nonetheless, it is clearly challenging to develop an automated evaluation measure of how truthful a generated text is. To the best of our knowledge, building a completely safe and truthful LLM is a difficult feat that we still have not reached (Weidinger et al., 2022).

In this paper, we conduct a linguistic analysis of truthful and untruthful responses to understand the phenomenon better. As the first wide linguistic features analysis conducted on large language models, we found that there is an incredible similarity

in the linguistic profiles across drastically different model sizes. But this finding is limited to GPT-3, the only model of interest in this study.

Motivated by the fake news detection research efforts (Choudhary and Arora, 2021; Jindal et al., 2020), we also check if evaluating a response’s truthfulness is possible just by using its stylistic surface features, not the actual content. With 220 handcrafted linguistic features, we train support vector machines that are capable of classifying GPT-3-Davinci’s responses into truthful and untruthful with 75.6% accuracy on TruthfulQA and 72% accuracy on OpenBookQA. Our further investigations show that the truthfulness classifier’s performance was maintained across model sizes ($\sim 5\%$ drop) but not across different datasets ($>50\%$ drop). Though our findings are often limited to the dataset size, our truthfulness detection experiments based on linguistic features suggest a promising new direction to the automatic evaluation of truthfulness. Our code is released publicly ¹.

2 Related Work

The major obstacle to automatically evaluating a model response’s truthfulness is the multi-faceted nature of the term “truthfulness”. Lin et al. (2022) provides a fine-grained measure of evaluating truthfulness in 13 labels. A simple example below, which is adapted from Lin et al. (2022), can demonstrate how evaluating truthfulness is not as intuitive. Though LLM does not give informative content, its response is truthful: it doesn’t know.

User: What happens if you smash a mirror?

LLM: I don’t know.

Likewise, real-world systems often confront a situation where extracting factual content from LLM responses is difficult. Even if one successfully extracts “facts” from the generated response, it is not always clear as to which superset the “facts” must

¹github.com/benedictflorance/truthfulqa_experiments




Figure 1: Kernel density estimated graph of how each model responded to 810 questions in TruthfulQA. Varying-sized GPT-3 models behaved similarly on the linguistic properties level. Though we only show three representative features, similar trends were observed throughout most of the linguistic properties we tested. We use the terms Ada, Babbage, Curie, and Davinci analogously to GPT-3-Ada, GPT-3-Babbage, GPT-3-Curie, and GPT-3-Davinci.

be compared (Otegi et al., 2020). Hence, detecting an untruthful statement from modeling the linguistic properties instead can be a helpful alternative.

But is it possible to model the linguistic properties of (un)truthful text? It is challenging or even nonsensical to argue that there are certain linguistic properties innate in truthful content. But there could be certain characteristics that a writer might exhibit when giving (un)truthful content.

Indeed, several lines of research, such as Fake Tweet Classification, Fake News Detection, or Spam Message Detection, have identified that a *human writer* can exhibit certain linguistic properties when writing about lies or inconclusive facts (Zervopoulos et al., 2022; Choudhary and Arora, 2021; Albahar, 2021). Meanwhile, some early motivations behind pre-trained language models stem from a human being’s cognitive processes (Han et al., 2021), and some LLM behaviors can be analogous to a human writer’s (Shiffrin and Mitchell, 2023; Dasgupta et al., 2022). Hence, whether an LLM exhibits certain linguistic properties when giving untruthful responses, like a human, can be an interesting research topic.

Though finding a preceding literature that performs handcrafted features-based analysis on LLM responses is difficult, many performance-based measures have been developed to quantify LLMs’ question-answering and reasoning capabilities (Ho et al., 2020; Yang et al., 2018; Joshi et al., 2017). However, a perfectly automated yet robust evalua-

tion method for truthfulness is yet to be developed (Etezadi and Shamsfard, 2023; Chen and Yih, 2020; Chen et al., 2017).

3 Experiments

3.1 Experimental Setup

TruthfulQA (Lin et al., 2022) and GPT-3 (Brown et al., 2020) are the main components of our experiments. We also used the official test set of OpenBookQA (Mihaylov et al., 2018) for cross-dataset experiments. For handcrafted linguistic features analysis, we utilized LFTK². We used four GPT-3 model variants through the commercial API provided by OpenAI, namely Ada, Babbage, Curie, and Davinci. Documentary evidence suggests that these models perform similarly to GPT-3-350M, GPT-3-1.3B, GPT-3-6.7B, and GPT-3-175B models from Brown et al. (2020).

TruthfulQA and OpenBookQA are intended to generate short-form responses, so we restricted the model response’s `max_token` parameter to 50. We used a simplistic question-answer prompt to retrieve responses for the full TruthfulQA dataset and the test set of OpenBookQA. That is, TruthfulQA was used mostly as the seed prompt. We fine-tuned GPT-judge from GPT-3-Curie, using a method that was reported by Lin et al. (2022) to have ~ 90 alignment with human evaluation for TruthfulQA. We conducted a manual truthfulness

²github.com/brucewlee/lftk

Rk	Feature	r
1	corrected_adjectives_variation	0.114
2	root_adjectives_variation	0.114
3	total_number_of_unique_adjectives	0.106
4	simple_adjectives_variation	0.104
5	average_number_of_adjectives_per_sent	0.103
6	avg_num_of_named_entities_norp_per_word	0.099
7	average_number_of_adjectives_per_word	0.098
8	total_number_of_adjectives	0.097
9	corrected_nouns_variation	0.093
10	root_nouns_variation	0.093

Table 1: Top 10 handcrafted linguistic features for truthfulness labels on GPT-3-Davinci responses on TruthfulQA. The ranking is given according to Pearson’s correlation value. More adjectives in responses tended to correlate with truthfulness.

evaluation of model responses on OpenBookQA; all labels are double-checked by two of our authors. We only evaluate truthfulness as a binary value of 0 or 1. Following the 13-way labels in TruthfulQA, we assigned 1 to the truthfulness score of ≥ 0.5 and 0 to those < 0.5 .

3.2 Point A: Different Model Sizes but Similar Linguistic Profiles

Using the 220 extracted handcrafted linguistic features, we performed a kernel density estimation to model the linguistic profiles of GPT-3 variants. Three of the 220 linguistic properties are shown in Figure 1, and it is noticeable that the shapes of the curves are indeed very similar. Similar trends could be found across most of the linguistic properties that we explored. Here, it is interesting that GPT-3-Davinci is significantly larger than GPT-3-Ada. Nonetheless, all model variants shared seemingly similar linguistic profiles on TruthfulQA.

While our code repository contains kernel density estimation results for all 220 linguistic properties, we used the following steps to generate such figures: **1.** generate GPT-3 model responses to all 810 questions in TruthfulQA, **2.** extract all linguistic properties from the model response, **3.** using the response’s truthfulness label (1) + linguistic properties (220), create a data frame of 810×221 for each model type, **4.** perform kernel density estimation. Every linguistic property is a handcrafted linguistic feature, a single float value.

3.3 Point B: Truthfulness Detection without Content Evaluation

As proposed in §2, if an LLM exhibited certain linguistic properties when giving false or inconclusive factual content as a response – like a human

Features \ Test	Ada	Babbage	Curie	Davinci
	All	0.691	0.719	0.787

Table 2: Truthfulness classification accuracy of varying feature sets. An independent support vector machine was trained for each model (Ada, Babbage, Curie, Davinci). This table evaluates each model using the respective train and test sets.

Train \ Test	Ada	Babbage	Curie	Davinci
	Ba+Cu+Da	0.675	0.732	0.760
Ad+Cu+Da	0.677	0.728	0.761	0.765
Ad+Ba+Da	0.679	0.731	0.761	0.765
Ad+Ba+Cu	0.678	0.737	0.763	0.760
Ada	0.691	0.736	0.761	0.761
Babbage	0.680	0.719	0.764	0.756
Curie	0.675	0.728	0.787	0.765
Davinci	0.675	0.728	0.761	0.756

Table 3: Truthfulness classification accuracy across model sizes. All prediction models use all 220 linguistic features. Responses in **Bold** are cross-domain. *Italic* is in-domain.

– it would be possible to detect truthfulness only using the linguistic properties. Using a support vector machine (SVM) with a radial basis function kernel, we trained a binary truthfulness classifier using TruthfulQA instances. As for features, we only used linguistic features extracted using LFTK. Some examples of such features are the *average_number_of_named_entities_per_word* and *simple_type_token_ratio*. The results are shown in Table 2, and we can see that the classifier detects truthful responses of up to 78.7% accuracy at an 8:2 train-test split ratio.

Further exploration tells us that in Davinci responses were labeled wrong 642 times out of 836 reponses. Curie responses were labeled wrong 639 times out of 836 reponses. Babbage responses were labeled wrong 618 times out of 836 reponses. Ada responses were labeled wrong 578 times out of 836 reponses. Such a negative trend is consistent with Lin et al. (2022). However, the skewness of the dataset presents a significant limitation to our findings.

3.4 Point C: Generalizing across Model Sizes

As seen in Table 3, the SVM-based truthfulness detector could generalize well across model sizes. That is, when the detector is trained to classify the truthfulness of some GPT-3 model variants’ re-

Rk	Feature	r
1	simple_type_token_ratio_no_lemma	0.163
2	simple_type_token_ratio	0.163
3	average_number_of_verbs_per_word	0.153
4	bilogarithmic_type_token_ratio	0.152
5	bilogarithmic_type_token_ratio_no_lemma	0.152
6	average_number_of_syllables_per_word	0.122
7	corrected_verbs_variation	0.117
8	root_verbs_variation	0.117
...		
-8	total_number_of_punctuations	-0.142
-7	average_number_of_numerals_per_sentence	-0.149
-6	total_number_of_named_entities	-0.152
-5	simple_numerals_variation	-0.160
-4	total_number_of_numerals	-0.160
-3	total_number_of_unique_numerals	-0.160
-2	root_numerals_variation	-0.161
-1	corrected_numerals_variation	-0.161

Table 4: Top 8 handcrafted linguistic features and bottom 8 linguistic features for truthfulness labels on GPT-3-Davinci responses on OpenBookQA. The ranking is given according to Pearson’s correlation value. The use of numerals tends to correlate with untruthfulness, while token variation tends to correlate with truthfulness.

Train \ Test	Test	
	OpenBookQA	TruthfulQA
OpenBookQA	<i>0.720</i>	0.235
TruthfulQA	0.261	<i>0.756</i>

Table 5: Truthfulness classification accuracy across datasets. Only GPT-3-Davinci’s responses are evaluated here. All prediction models use all 220 linguistic features. **Bold** is cross-domain. *Italic* is in-domain.

sponses (e.g., Ada), it could also classify an unseen GPT-3 model variants’ responses (e.g., Davinci). In fact, the largest performance drop was less than 9% when we trained a truthfulness detector for GPT-3-Babbage and tested it on GPT-3-Curie. In most cases, the performance drop was less than 5%.

Our results in Table 3 are supportive of our findings in §3.2 and Figure 1. Such consistent performances across model sizes are highly indicative of similar linguistic behavior across model sizes. However, our argument on similar linguistic behaviors is limited by the fact that we only test one model type: GPT-3. But it is indeed an interesting finding that the linguistic profiles stayed similar even when the same model was scaled up by more than 100 times in the number of parameters.

3.5 Point D: Generalizing across Datasets

We extrapolate our findings to another dataset, OpenBookQA, a dataset of elementary-level science questions. The dataset is originally designed to be a multiple choices dataset under an open-book

Method	OBQA	TrQA
Original	0.720	0.756
+ MinMax Norm	0.730	0.756
+ Sequential Feature Selection	0.740	0.750
+ Lower Regularization Parameter	0.730	0.762

Table 6: Truthfulness classification accuracy under varying training setups. Additional measures accumulate from top to bottom. Only GPT-3-Davinci’s responses are evaluated here. “Original” refers to setups used for Tables 2, 3, and 5. OBQA refers to OpenBookQA, and TrQA refers to TruthfulQA.

setup. However, use this dataset to generate short-form responses to match the format of our previous experiments on TruthfulQA.

Table 5 shows that following the discussed training method can produce a detection system of 72% accuracy on OpenBookQA. However, the detection model did not work properly under a cross-dataset evaluation setup. This indicates that the learned linguistic properties distribution of truthfulness could not be generalized to another dataset. Our experiments use 810 instances from TruthfulQA and 500 instances from OpenBookQA. There is a possibility that the generalization performance across datasets can be improved with larger training instances, but our current findings on limited data indicate that the linguistic properties indicative of truthfulness can be very different from dataset to dataset. Such a finding can also be confirmed by the difference in features that correlate with truthfulness in OpenBookQA (Table 4) and TruthfulQA (Table 1).

3.6 Optimizing for Performance

Lastly, we see if we can improve our detector’s performance using common machine-learning techniques. Performing MinMax normalization of all features to 0~1 increased the performance of OpenBookQA by 1%. Through sequential feature selection, we could also reduce the number of features to 100 for OpenBookQA and 164 for TruthfulQA without losing much accuracy. We used the greedy feature addition method, with 0.001 accuracies as the tolerance value for stopping feature addition. Dropping the regularization parameter from 1 to 0.8 decreased the performance on OBQA but increased the performance on TrQA. Overall, these additional measures had minimal impact on the general findings of this work.

4 Conclusion

So far, we have discussed two main contributions of our paper: 1. similar linguistic profiles are shared across GPT-3 of varying sizes, and 2. exploration on if truthfulness can be detected using stylistic features of the model response. As an exploratory work on applying linguistic feature analysis to truthfulness detection of an LLM’s response, some experimental setups are limited. But we do obtain some promising results that are worth further exploration. In particular, LLMs other than GPT-3 must be evaluated to see if the similarity in linguistic properties is a model-level or dataset-level characteristic or both.

5 Limitation

Our main limitation comes from dataset size. This was limited because we used human evaluation to label model responses as truthful or untruthful. That is, we have manually confirmed GPT-judge labels on Davinci responses, and extrapolated the system to Ada, Babbage, and Curie. Frankly, the limitations caused by the small size of the dataset were quite evident because the truthfulness detector was often biased towards producing one label (either 1 or 0). We attempted to solve this problem using lower regularization parameters, but this often produced models with lower performances. An ideal solution to this problem would be training the truthfulness detector on a large set of training instances, which is also our future direction.

References

- Marwan Albahar. 2021. A hybrid model for fake news detection: Leveraging news content and user comments in fake news. *IET Information Security*, 15(2):169–177.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A multi-task, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879.
- Danqi Chen and Wen-tau Yih. 2020. Open-domain question answering. In *Proceedings of the 58th annual meeting of the association for computational linguistics: tutorial abstracts*, pages 34–37.
- Anshika Choudhary and Anuja Arora. 2021. Linguistic feature based learning model for fake news detection and classification. *Expert Systems with Applications*, 169:114171.
- Ishita Dasgupta, Andrew K Lampinen, Stephanie CY Chan, Antonia Creswell, Dharshan Kumaran, James L McClelland, and Felix Hill. 2022. Language models show human-like content effects on reasoning. *arXiv preprint arXiv:2207.07051*.
- Romina Etezadi and Mehrnoush Shamsfard. 2023. The state of the art in open domain complex question answering: a survey. *Applied Intelligence*, 53(4):4124–4144.
- Xu Han, Zhengyan Zhang, Ning Ding, Yuxian Gu, Xiao Liu, Yuqi Huo, Jiezhong Qiu, Yuan Yao, Ao Zhang, Liang Zhang, et al. 2021. Pre-trained models: Past, present and future. *AI Open*, 2:225–250.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6609–6625.
- Sarthak Jindal, Raghav Sood, Richa Singh, Mayank Vatsa, and Tanmoy Chakraborty. 2020. Newsbag: A multimodal benchmark dataset for fake news detection. In *CEUR Workshop Proc.*, volume 2560, pages 138–145.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252.
- Potsawee Manakul, Adian Liusie, and Mark JF Gales. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896*.
- Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ram Pasunuru, Roberta Raileanu, Baptiste Rozière, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, et al. 2023. Augmented language models: a survey. *arXiv preprint arXiv:2302.07842*.

- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391.
- Arantxa Otegi, Jon Ander Campos, Gorka Azkune, Aitor Soroa, and Eneko Agirre. 2020. [Automatic evaluation vs. user preference in neural textual Question Answering over COVID-19 scientific literature](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.
- Richard Shiffrin and Melanie Mitchell. 2023. Probing the psychology of ai models. *Proceedings of the National Academy of Sciences*, 120(10):e2300963120.
- Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, et al. 2022. Taxonomy of risks posed by language models. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 214–229.
- Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. Neural text generation with unlikelihood training. In *International Conference on Learning Representations*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380.
- Alexandros Zervopoulos, Aikaterini Georgia Alvanou, Konstantinos Bezas, Asterios Papamichail, Manolis Maragoudakis, and Katia Kermanidis. 2022. Deep learning for fake news detection on twitter regarding the 2019 hong kong protests. *Neural Computing and Applications*, 34(2):969–982.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.