

Pragmatic Constraint on Distributional Semantics

Elizaveta Zhemchuzhina,^{1*} Nikolay Filippov,¹ Ivan P. Yamshchikov^{2,3}

¹LEYA Lab for Natural Language Processing, Higher School of Economics, Yandex, St. Petersburg, Russia

²Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

³CEMAPRE, University of Lisbon, Portugal

ivan@yamshchikov.info

Abstract

This paper studies the limits of language models' statistical learning in the context of Zipf's law. First, we demonstrate that Zipf-law token distribution emerges irrespective of the chosen tokenization. Second, we show that Zipf distribution is characterized by two distinct groups of tokens that differ both in terms of their frequency and their semantics. Namely, the tokens that have a one-to-one correspondence with one semantic concept have different statistical properties than those with semantic ambiguity. Finally, we demonstrate how these properties interfere with statistical learning procedures motivated by distributional semantics.

Introduction

Various modern Natural Language Processing (NLP) models are designed and trained based on the distributional semantics hypotheses – namely, that linguistic items used in similar contexts tend to have similar meanings (Harris 1954). At the same time, Zipf's law states that given some corpus of natural language utterances, the frequency of any word is inversely proportional to its rank in the frequency table (Zipf 1932). As vocabulary size (or amount of data) grows, this rank-frequency distribution provides a "heavy tail" – a significant amount of linguistic items whose frequency is not enough for distributional semantics to be effective with them. Most modern NLP methods cut this "tail" out of the discussion by introducing the notion of a "token" that could be shorter than the word, and restricting the number of tokens in the model's vocabulary. However, the fact that distributional semantics models are trained on the cor-

pora that follow Zipf's law raises a series of interesting questions such as:

- How does a particular tokenization procedure interfere with Zipf's law?
- Are there distinct differences between frequent "head" and infrequent "tail" tokens?
- Could those differences affect the model's performance and be vital for some aspects of NLP tasks that the model can solve?

This paper tries to initiate the discussion of those questions by providing insights into how the empirical fact of the "heavy tail" of Zipf's distribution interferes with the methods based on the idea of distributional semantics.

Since its introduction in the first half of XX century (Zipf 1932), Zipf's law has been extensively studied and applied in various fields: natural languages (Montemurro 2001), random text generation (Ferrer-i Cancho and Solé 2002), information theory (Harremoës and Topsøe 2006), population statistics (Reed 2002), internet traffic statistics (Breslau et al. 1999), infometrics (Egghe 2005), economics (de Wit 2005), ecology (Camacho and Solé 2001), biology (Gamow and Yčas 1955), information security (Wang and Wang 2016) just to name a few.

Many researchers studied the significance of Zipf's law for various NLP problems. For example, (Yang 2013) suggests that Zipf's law facilitates early language acquisition and uses Zipf's law to distinguish human language usage from bioacoustics of other species, such as chimps. (Zhang 2008) discovers that the distribution of lexical tokens in Java source code follows the Zipf's law. (Takahashi and Tanaka-Ishii 2017) argue that a signature of efficient representations is that frequency distributions follow power laws. (Pimentel et al. 2021) demonstrate that natural codes are closer to not being optimized (in the Zipfian sense) than to being maximally compressed. (Cristelli, Batty, and Pietronero 2012) find that many natural systems do not show true power law behavior because they are incomplete or inconsistent with the conditions under which one might expect power laws emergence. (Ferrer-i Cancho and Vitevitch 2018) show that a single assumption on the joint probability of a word and a meaning suffices to infer Zipf's meaning-frequency law and

*This work is an output of a research project implemented as part of the Basic Research Program at the National Research University Higher School of Economics (HSE University). This research was supported in part through computational resources of HPC facilities at HSE University (Kostenetskiy, Chulkevich, and Kozyrev 2021).

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Accepted to Workshop on Knowledge Augmented Methods for Natural Language Processing, in conjunction with AAAI 2023.

argue that this assumption can be justified as the outcome of a biased random walk in the process of mental exploration. Finally, (Nikkarinen et al. 2021) argue that any approach that assigns zero probability to any out-of-vocabulary word form produces negatively biased probabilities for any out-of-vocabulary word, while positively biased probabilities to in-corpora words. The authors make a compelling argument in favor of properly modeling the unigram distribution and claim it should be a central task in NLP.

In this paper, we study the rank-frequency distribution behavior for sets of tokens with different maximum token lengths (derived from vocabulary size and tokenization algorithm) and show that this distribution is well approximated with the superposition of two Zipf’s laws for two distinct, coherent sets of tokens. We infer that subsets of tokens sampled from the head and tail of such distributions have differences both in terms of statistics and semantics. Thus we introduce the terms ”pragma” and ”idea” to distinguish between the two subsets of tokens. We believe that conceptual understanding of these two token categories is fundamental for assessing the limitations of current models based on the so-called distributional hypothesis (Harris 1954).

Experiments

Here we present a series of experiments that provides insights into statistic and semantic properties of various tokens provided by the chosen tokenization.

Zipf’s Law for Different Tokenizations

We experiment with three different tokenization algorithms: Byte Pair Encoding or BPE (Gage 1994), WordPiece (Wu et al. 2016), and Unigram (Kudo 2018). We find no major differences regarding token frequency vs. token rank behavior. Up to some vocabulary size limit, the rank-frequency distribution follows Zipf’s law regardless of the chosen tokenization algorithm.

Figure 1 shows the rank-frequency distributions of tokens for Wikitext-103. dataset with BPE, WordPiece, and Unigram tokenization, respectively, with three different vocabulary sizes. All of them demonstrate Zipf-like behavior. We found BPE to be the most illustrative algorithm for this paper due to the nature of the algorithm. BPE would produce longer tokens as vocabulary size increases since BPE keeps adding the most frequent pair of existing tokens to the vocabulary as a new token. This makes it straightforward: a bigger vocabulary size leads to a bigger maximum token length. From now on, if the tokenization algorithm is not explicitly specified, BPE is used.

Closer Look at Zipf’s Law

We run a series of experiments with different vocabulary sizes, from a few thousand (which gives tokenization on the symbol, N-gram, and word levels) to several million (which inevitably adds tokens consisting of several words and/or whole phrases to the vocabulary). We found that,

Wikitext-103 dataset was used for all experiments reported in the paper. <https://huggingface.co/datasets/wikitext>




Figure 1: Rank-Frequency distribution for BPE, WordPiece and Unigram tokenization algorithms, vocabulary sizes 10 000, 100 000, 1 000 000




Figure 2: Rank-Frequency distribution and token lengths distribution: BPE tokenization algorithm, vocabulary size 1 000 000

with smaller vocabulary sizes, the rank-frequency distribution tends to follow Zipf’s law even on the level of sub-word tokens. As the vocabulary size grows further to include tokens consisting of several words and/or whole phrases in the vocabulary, the rank-frequency diagram does not anymore follow the “pure” power law but rather resembles the superposition of two Zipf’s laws, thus following the “coherence” concept introduced by Matthieu Cristelli, Michael Batty and Luciano Pietronero (Cristelli, Batty, and Pietronero 2012). Once the vocabulary size reaches some point, there appears a threshold near which the rank-frequency diagram experiences a phase transition. The parameter of Zipf’s approximating rank-frequency distribution shifts. We presume that such behavior is caused by the fact that, in the case of larger vocabularies, we deal with two subsets of tokens, each subset being coherent in itself. Figure 2 represents the rank-frequency distribution for vocabulary size 1 million, illustrating the “bend” of Zipfian distribution. This figure represents BPE tokenization algorithm, but we observe the same behavior for other tokenization algorithms as well. For further experimental results, we address the reader to the Appendix .

Semantics and Zipf’s law

As seen from the experiment results (see diagrams in Appendix), the distribution starts behaving as a superposition of two Zipf’s laws at the point where token length starts exceeding some threshold. We presume that this threshold is determined by semantics: the distribution behavior at the head part of such diagrams differs from that at the tail part. The hypothesis is that the head part mostly consists of shorter tokens with possible semantic variations, whereas the tail part mostly consists of longer tokens connected to one specific semantic field. Figure 3 illustrates the token lengths distribution. Both shorter multi-meaning and longer single-meaning tokens are coherent and demonstrate distribution under Zipf’s law if considered separately. This is vis-




Figure 3: Token Lengths frequency: vocabulary size 1 million, logarithmic scale

ible on distribution diagrams with smaller vocabulary sizes that, presumably, consist mostly of shorter tokens (see example at Figure 11 in Appendix). Together these two subsets do not demonstrate pure Zipfian behavior anymore. A “heavy tail” of such distribution is clearly visible at Figure 2. Figure 3 shows that the tail mostly consists of longer tokens which, as shown below, tend to have one meaning rather than many.

To illustrate qualitative differences between these two distributions, we carried out an additional experiment with two subsets of tokens: one from the head of the one-million-tokens vocabulary distribution, and another one from the tail. As expected, the “head” subset consisted of shorter tokens, primarily words, and the “tail” subset consisted of longer tokens, mainly phrases and parts of sentences. Since there are no long tokens neither in the head nor in the middle of the distribution (see Figure 3), and because of the nature of the experiment (we were mainly interested in the audience’s perception of shorter tokens from the head and longer tokens from the tail), we filtered out a few short tokens that might occur in the tail, and left the tail part with long tokens only. We conducted a poll regarding the shuffled sequence of tokens from the two subsets among a group of professionals in linguistics to find out the difference between these subsets based on the opinion of people with relevant professional backgrounds. We asked the following questions regarding every single token X in the set:

1. Can you reformulate X ?
2. How many meanings does X have depending on context?
3. Can you place X into context?

Figures 4 – 5 illustrate one of the poll results: shorter tokens from head part of distribution are characterized with semantic ambiguity and more often can have several mean-

ings depending on context. Longer tokens from tail part of distribution tend to have one or two specific meanings.




Figure 4: Poll results: token length vs. average number of meanings (heatmap) listed by a linguist




Figure 5: Poll results: token length vs. average number of meanings (boxplot) listed by a linguist

Figures 6 – 8 illustrate that the shorter, semantically ambiguous tokens can be easily placed into different contexts, while longer, single-meaning tokens represent context themselves: attempt to place them in context results in changing a few shorter tokens they consist of, leaving the rest unchanged. The difference is clearly visible by the normalized Levenshtein distance between the original token and the token in context. It is also worth mentioning that the normalized Levenshtein distance distribution in the left part of the diagram visually resembles Zipf’s law as well – this may be a topic for a closer look and more detailed study with a much bigger audience.




Figure 6: Poll results: token length vs. percentage of token placed in context




Figure 7: Poll results: token length vs. normalized Levenshtein distance between original token and token placed in context. "Head" tokens can be placed in various contexts that differ significantly. "Tail" tokens tend to be embedded in the similar contexts that are comparable with the size of the token itself.




Figure 8: Poll results: token length vs. Levenshtein distance between original token and token after restatement

The results were clustered in accordance with the subset each token belonged to, and show that, based on the opinion of people with professional linguistic backgrounds:

- the "head" tokens can be easily replaced with synonyms, homonyms, or phrases with synonymous or homonymous meaning, the full content of the token is changed in the course of restatement, along with (possibly) meaning;
- the "tail" tokens are difficult to reformulate, mainly through changing one or more shorter tokens they consist of but leaving the rest unchanged. Regardless of the restatement, the whole token's meaning does not change;
- "head" tokens often have several meanings while "tail" tokens mostly have one or two meanings;
- "head" tokens can be easily placed into context and may have different meanings depending on the context;
- placing the "tail" token into context is not that easy and is done by placing a few shorter tokens into the context of the longer token rather than vice versa, while the whole long token's meaning remains unchanged.

These results illustrate that "head" and "tail" tokens differ both in terms of semantics and in terms of statistics. "Head" tokens have semantic ambiguity, can be easily placed in different contexts with different meanings. "Tail" tokens have a one-to-one correspondence with the specific semantic concept. They are not easy to place in a new context other than the context they define themselves. The tail part of rank-frequency distribution, mostly consisting of such tokens, demonstrates behavior inherent to the rank-frequency distribution of distinct, coherent sets of tokens. Both "head" and "tail" subsets are coherent in terms of (Cristelli, Batty, and Pietronero 2012), and the distribution of each of them seems to follow Zipf's law. In contrast, the distribution of their union looks more like a superposition of two Zipf's laws associated with two different coherent sets of tokens.

Pragma and Idea

We suggest to introduce the following terms that characterize different types of tokens:

Atom – the smallest element of a discrete sequence that cannot be divided further into smaller parts. In written language, the single symbols are considered atoms. The initial set of atoms is believed to be an exhaustive list limited in size.

Pragma – a part of a discrete sequence that consists of atoms and represents an integral part of an idea. Pragma has semantic ambiguity and may have different meanings depending on context. In written language, pragmas are represented by symbol n-grams, words, and word n-grams.

Idea – a part of a discrete sequence that consists of pragmas and has a one-to-one correspondence to some specific concept (a semantic concept in the case of language). In written language, ideas can be represented by semantically distinct sentences, set phrases, and colloquialisms having specific meanings.

We do not mention atoms in this paper, but we introduce them here for integrity. We believe such an atom-pragma-idea granularity structure helps to understand and describe some concepts related to both statistical behavior and semantics of specific sets of tokens. We also believe that such

a three-layer granularity structure may find its application in discrete sequence processing studies related to many fields other than natural language: chemistry, biology, music, etc.

Discussion

Following the experiment results and findings of the experiments, we consider some topics as subjects for further investigation and discussion.

Language models tend to use vocabulary sizes smaller than the ones we used in our experiments. Therefore, they represent ideas as sequences of pragmas rather than forming a distinct, coherent subset of tokens out of ideas. This makes it difficult to register the "phase shift" where the longer sequence of pragmas becomes an idea: it starts to demonstrate one-to-one correspondence to specific semantic concept.

We believe this might be one of the reasons why language models show poor performance when dealing with longer sequences, in contrast to humans who do not have such problems. We suggest addressing this phenomenon as *Pragmatic Constraint* — the capability of statistical learning to operate within pragmatic tokenization that does not leverage the reduction of semantic ambiguity characteristic for the idea-based tokens.

Conclusion

This paper studies the statistical and semantic properties of tokens. Both statistical and semantic differences were found for two distinct subsets of tokens that we called "pragma" and "idea". Pragmas have semantic ambiguity, while ideas have a one-to-one correspondence to the specific semantic concept. Statistically, the set containing both of these subsets does not demonstrate pure Zipfian behavior but rather behaves like a superposition of two Zipf's laws. Results of a poll conducted among a small group of linguists are in line with these findings. This may impose limitations on machine learning methods and techniques based on statistical learning and assumptions about the Zipfian nature of sets of tokens. Such limitations and their impacts are subject to further study.

Limitations

Wikitext-103 dataset was used for this research. We believe the results are reproducible with other datasets in languages other than English, but this has not been proved yet.

For results to be reproducible, the dataset shall be big enough to have at least a million tokens in vocabulary.

The poll was performed among a small group of five professionals since the primary goal was not to gather statistics but rather to get the opinions of people with strong professional background in linguistics. For statistically significant results, a broader audience is required.

GPU is preferable to achieve results within a reasonable time.

References

Breslau, L.; Cao, P.; Fan, L.; Phillips, G.; and Shenker, S. 1999. Web caching and Zipf-like distributions: evidence

and implications. *IEEE INFOCOM '99. Conference on Computer Communications. Proceedings. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. The Future is Now (Cat. No.99CH36320)*, 1: 126–134.

Camacho, J.; and Solé, R. V. 2001. Scaling in ecological size spectra. *EPL*, 55: 774–780.

Cristelli, M.; Batty, M.; and Pietronero, L. 2012. There is More than a Power Law in Zipf. *Scientific reports*, 2: 812.

de Wit, G. 2005. Zipf's Law in Economics. *Scales research reports*.

Egghe, L. 2005. The exact place of Zipf's and Pareto's law amongst the classical informetric laws. *Scientometrics*, 20: 93–106.

Ferrer-i Cancho, R.; and Solé, R. 2002. Zipf's Law and Random Texts. *Adv. Complex Syst.*, 5: 1–6.

Ferrer-i Cancho, R.; and Vitevitich, M. S. 2018. The origins of Zipf's meaning-frequency law. *Journal of the Association for Information Science and Technology*, 69(11): 1369–1379.

Gage, P. 1994. A New Algorithm for Data Compression. *C Users J.*, 12(2): 23–38.

Gamow, G.; and Yčas, M. 1955. Statistical Correlation of Protein and Ribonucleic Acid Composition. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 41, issue 12, 1011–1019.

Harremoës, P.; and Topsøe, F. 2006. Zipf's Law, Hyperbolic Distributions and Entropy Loss. In *Electronic Notes in Discrete Mathematics*, volume 21, 788–792. ISBN 978-3-540-46244-6.

Harris, Z. 1954. Distributional Hypothesis. *Word*, 10(23): 146–162.

Kostenetskiy, P.; Chulkevich, R.; and Kozyrev, V. 2021. HPC Resources of the Higher School of Economics. In *Journal of Physics: Conference Series*, volume 1740, 012050. IOP Publishing.

Kudo, T. 2018. Subword Regularization: Improving Neural Network Translation Models with Multiple Subword Candidates. *CoRR*, abs/1804.10959.

Montemurro, M. 2001. Beyond the Zipf–Mandelbrot law in quantitative linguistics. *Physica A: Statistical Mechanics and its Applications*, 300: 567–578.

Nikkarinen, I.; Pimentel, T.; Blasi, D.; and Cotterell, R. 2021. Modeling the Unigram Distribution. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 3721–3729.

Pimentel, T.; Nikkarinen, I.; Mahowald, K.; Cotterell, R.; and Blasi, D. 2021. How (Non-) Optimal is the Lexicon? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4426–4438.

Reed, W. J. 2002. On the Rank-Size Distribution for Human Settlements. *Environmental Economics*.

Takahashi, S.; and Tanaka-Ishii, K. 2017. Do neural nets learn statistical laws behind natural language? *PLoS one*, 12(12): e0189326.

Wang, D.; and Wang, P. 2016. On the Implications of Zipf's Law in Passwords. In *ESORICS 2016*, volume 9878. ISBN 978-3-319-45743-7.

Wu, Y.; Schuster, M.; Chen, Z.; Le, Q. V.; Norouzi, M.; Macherey, W.; Krikun, M.; Cao, Y.; Gao, Q.; Macherey, K.; et al. 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.

Yang, C. 2013. Who's afraid of George Kingsley Zipf? Or: Do children and chimps have language? *Significance*, 10(6): 29–34.

Zhang, H. 2008. Exploring regularity in source code: Software science and Zipf's law. In *2008 15th Working Conference on Reverse Engineering*, 101–110. IEEE.

Zipf, G. K. 1932. *Selected Studies of the Principle of Relative Frequency in Language*. Cambridge, Massachusetts: Harvard University Press.

Appendix

Below are figures representing experiment results referred to in the article.

Figures 9–10 illustrate Zipf's law behavior for different tokenization algorithms.




Figure 9: Rank-Frequency distribution: WordPiece tokenization, vocabulary size 30 000 (rare tokens excluded)




Figure 10: Rank-Frequency distribution: BPE tokenization, vocabulary size 30 000 (rare tokens excluded)




Figure 12: Rank-Frequency distribution: vocabulary size 10 000

Figures 11–14 illustrate Zipf’s law behavior for different vocabulary sizes and, hence, different maximum token lengths.




Figure 11: Rank-Frequency distribution: vocabulary size 5 000




Figure 13: Rank-Frequency distribution: vocabulary size 30 000




Figure 14: Rank-Frequency distribution: vocabulary size 100 000