

Lost in Space: Finding the Right Tokens for Structured Output

Sil Hamilton and David Mimno
Department of Information Science
Cornell University
{srh255,mimno}@cornell.edu

Abstract

General-purpose language models are trained to produce varied natural language outputs, but for some tasks, like annotation or classification, we need more specific output formats. LLM systems increasingly support structured output, which enforces formats by sampling tokens according to a grammar but also unpredictably reduces downstream performance. Are there systematic differences between grammars that appear semantically (and often visually) similar to humans? To answer this, we test four popular model families with five varying output formats on four common NLP benchmarks. We find all models perform most accurately when guided to use formats respecting convention, such as letters for multiple choice and real numbers for numerical prediction. Performance also improves by 5%-10% when guiding models to return tokens incorporating leading whitespace, with smaller models benefiting the most. We find leading whitespace helps models avoid structural deficiencies in subword token representations. We finally present best practices for researchers using language models as zero-shot classifiers with structured output.

1 Introduction

Language model outputs are flexible, which makes them effective for a wide range of applications, but this can also be a liability. While language models are often used to emit free text, e.g. writing a story, other important applications require their output match specific formats like class labels and domain-specific languages (Wang et al., 2023a; Geng et al., 2023). Rigid formats make it easier to parse responses, helping users avoid wasting queries on potentially invalid stochastic output.

Multiple choice options are the simplest example, but users are increasingly using more flexible (and more failure-prone) output format restrictions. Reinforcement learning from human feedback (RLHF) and instruction tuning can guarantee

Prompt:

Rate the similarity of these words from 1 to 5: "train" and "underground". Use only the following characters:

Tokens decoded as digits:

1 2 3 4 5 → 3

Tokens decoded as digits with whitespace:

1 2 3 4 5 → 5

Llama 3.1 picks different tokens depending on whether given tokens incorporate whitespace.

Figure 1: Llama 3.1 8b provides different answers to a question from the MEN Test Collection when constrained to produce digit tokens *with* and *without* a leading space. Sampled with temperature set to 0.

well-formed output (Ouyang et al., 2022), but these methods are time consuming and expensive to implement, and are impractical for the variable and rapid-turnaround scenarios typical of LLM usage.

An increasingly popular method for ensuring output consistency without fine-tuning is selectively sampling tokens valid according to a formal language (Shin et al., 2021; Wang et al., 2023b). Structured output, powered by grammar-constrained decoding (GCD), is regularly implemented in server software for running LLM inference (Gerganov, 2024; Microsoft, 2024; Rickard, 2024). Some servers allow users to specify grammars with regular languages (Willard and Louf, 2023; Microsoft, 2024), while others allow users to write context-free grammars in Backus–Naur form (Chomsky, 1956; Knuth, 1964; Gerganov, 2024).

But despite this popularity, enabling GCD appears to degrade downstream task accuracy (Willard and Louf, 2023; Park et al., 2024; Beurer-

Kellner et al., 2024; Geng et al., 2023; Le et al., 2024). We hypothesize that the reduced performance of GCD is not due to GCD itself, but to non-optimal grammars. For example, Figure 1 shows Llama 3.1 produces different results when constrained by two different output token sets, one with leading spaces and another without.

We investigate this phenomenon by comparing the performance achieved when using one of five distinct token sets across four different data labeling scenarios with four popular families of pre-trained language models. We present three key contributions. First, although all models vary considerably between formats for a given task, they often agree on a preferred format, such as real numbers for prediction tasks. Second, important distinctions may not be obvious to humans. We find that tokens incorporating leading whitespace consistently improve performance, echoing results for syntactic tokens (Pimentel and Meister, 2024; Oh and Schuler, 2024). We offer initial explanation for this difference using evidence in the embedding weights of our models of interest. We finally propose a set of strategies for selecting target tokens when specifying grammars for GCD libraries. These strategies can help achieve better performance when using structured output, particularly for users of smaller, less computationally intensive models.

2 Related Work

Grammar-constrained decoding failure modes. Grammar-model misalignment is a recognized issue with grammar-constrained decoding algorithms. Existing solutions generally involve modifying GCD algorithms to avoid misalignment by adding additional constraints. Park et al. (2024) propose an adaptive algorithm where beam search discovers valid *and* optimal token sequences according to the grammar. They report some success in generating valid sequences, but note accuracy is not improved in all respects. Beurer-Kellner et al. (2024) propose “DOMINO,” a “minimally invasive” algorithm, pre-computing allowed sequences of tokens before dynamically selecting certain paths according to probability mass. While these approaches show promise, both propose non-trivial adjustments to how GCD is commonly implemented in software. This limits their downstream applicability, and does not provide a generalizable token alignment strategy for common inference servers *as they are presently implemented*.

Non-exchangeable token embeddings. Artifacts in LLM training can result in tokens that appear semantically similar to humans but whose latent representations differ to a great extent. Wang and Mimno (2023) explore trained token representations in the embedding weights of highly multilingual LLMs, observing significant variation in how models represent semantically similar tokens in the embedding space. Zhang et al. (2024) report similar behavior in monolingual LLMs, noting asymmetries in token categories such as pronouns and nouns. Neither study assesses the downstream impact of this behavior.

Leading whitespace tokens. Tokenizers descending from SentencePiece duplicate certain tokens: once with leading whitespace to capture instances where the token begins a word, and again without whitespace (Kudo and Richardson, 2018). For example, the Llama 3 tokenizer can represent “culture” either as the token culture (index key 70905) or `_culture` (7829). These tokens are distinct despite decoding to strings nearly identical to humans.

Pimentel and Meister (2024) note tokenizers making use of LW tokens can complicate computing token surprisal. Oh and Schuler (2024) similarly find models trained with LW tokens are prone to an increased likelihood of failing to learn word boundaries, given their tokenizers explicitly penalize the appearance of space tokens. Neither work estimates the impact tokens may have on downstream tasks, nor whether LW tokens can converge with non-LW tokens in training.

3 Methods

If we can dependably elicit grammar preferences, we can avoid extending existing GCD-enabled inference servers as is routinely recommended by the literature (Geng et al., 2023; Park et al., 2024; Beurer-Kellner et al., 2024). In this section we describe our method for systematically discovering these GCD-optimal token sets.

3.1 Data

Benchmark selection. We assess the viability of certain token formats against four NLP benchmarks involving quantitative outputs.¹ Our first

¹Common multiple choice benchmarks (e.g. MMLU) are subject to token influence (see Table 5), but do not require GCD as responses can be read straight from model logits. Compare with prediction tasks where responses may be com-

Category	Variant	Token Set	Maximum Length	Example
Integer	Integer, <i>numeric</i>	{1, 2, ..., 10}	1 token	7
	Integer, <i>word</i>	{One, Two, ..., Ten}	1 token	“Seven”
Real Number	Real, <i>numeric</i>	{0.1, ..., 1}	≤ 3 token(s)	0.1
	Real, <i>word</i>	{Zero point zero, ..., One}	≤ 4 token(s)	“Zero point one”
Percentage	Percentage, <i>numeric</i>	{0%, ..., 100%}	2 tokens	10%
	Percentage, <i>word</i>	{Zero, ..., One hundred; percent}	≤ 4 token(s)	“Ten percent”
Boolean	Binary, <i>numeric</i>	{1, 0}	1 token	1
	Binary, <i>word</i>	{True, False}	1 token	“True”
Likert	Likert task, <i>numeric</i>	{1, 2, 3, 4, 5}	1 token	4
	Likert task, <i>word</i>	{Strongly disagree, ..., Strongly agree}	≤ 3 token(s)	“Somewhat agree”

Table 1: Target token sets considered in our study, noting the value range and specifying the maximum length of a valid sequence with Llama 3.1 followed by an example.

two benchmarks are the Semantic Textual Similarity Benchmark (Cer et al., 2017) and the MEN Test Collection (Bruni et al., 2014). Both are prediction tasks asking annotators to rate the semantic similarity of two sentences on a scale between two numbers. The last two are *binary detection* tasks asking annotators to assess whether a given phenomenon is present: Quora Question Pairs (DataCanary et al., 2017) and ToxicChat (Lin et al., 2023).

Benchmark preprocessing. We randomly sample 1,000 problems from each benchmarks for a total of 4,000 problems. We normalize each benchmark to a standard format, a JSON structure containing a question and answer in the form of a real number. We then specify a standard task prompt template valid for problems sampled from all four datasets. We present example task prompts for all four benchmarks in Appendix A.

3.2 Evaluation

Baseline measures. We contextualize our LLM results using two baseline measures. Given two of our four benchmarks target semantic similarity as a real number, we make use of BERTScore to provide a baseline BERT-based semantic similarity score for all problems (Zhang et al., 2020). This system was trained to perform semantic similarity scoring, conveniently serving as an upper bound to compare our zero-shot prompting strategies against. We estimate the lower bound by measuring the Levenshtein distance (Levenshtein, 1966) between both choices in all problems, with the expectation these scores will trend lower given all benchmarks appear only lightly correlated with string edit distance.

posed of two to three tokens, e.g. 12% being composed of the tokens [1], [2], [%].

We intend neither measure to be state-of-the-art. Rather, BERTScore represents the performance of a reasonable language model, while the edit distance of each problem is a simple method expected to outperform random guessing by some margin.

Model selection. To assess whether grammar-dependent degradations are consistent between LLMs, we evaluate eight models of varying size.² We select these eight from four popular families of instruction-tuned LLMs: Llama 3.1 8b/70b, Qwen 2.5 7b/72b, Gemma 2 9b/24b and Phi 3 3b/14b (Dubey et al., 2024; Qwen et al., 2025; Team et al., 2024; Abdin et al., 2024). All models were trained in the last two years and feature pre-training sets ranging between 3.3 and 18 trillion tokens. We select these model families for three reasons: their high performance relative to proprietary models, their high compatibility with GCD-enabled inference servers, and their popularity with researchers for certain downstream tasks such as classification and annotation, e.g. Bamman et al. (2024).

Model inference. All models are run with model-default temperature and Top_p parameters. We set context length to 512 tokens in all cases. We obtain models from their official repositories, quantizing each to q8_0 with scripts provided by Gerganov (2024) to maximize local inference speed.³ We host the models with llama.cpp (compiled with CUDA support) and run the server on a system equipped with two Nvidia RTX A6000 GPUs. We

²We do not include proprietary models in our assessment because existing APIs, e.g. that of OpenAI, do not enable defining context-free grammars.

³Prior research suggests reducing weight precision to an average of no fewer than eight bits ensures quantized performance remains comparable to unquantized models (Dettmers et al., 2022).

Benchmark	BERTScore		Levenshtein	
	Corr.	MSE	Corr.	MSE
MEN	0.768	397	-0.055	3,736
QUORA	0.609	2,574	-0.323	8,106
STSB	0.824	379	-0.107	8,076
TOXICCHAT	-0.081	4,248	-0.014	267,889
\bar{x}	0.530	—	-0.125	—

Table 2: Baseline correlations and MSE for each STS benchmark considered in this paper. Note semantic similarity consistently outperforms Levenshtein distance.

use the GBNF GCD implementation included with `llama.cpp` without modification.⁴

Token output formats. Current GCD implementations have users specify output formats with CFGs meaning many possible output formats exist.

We consider five commonly used output formats, shown in Table 1. These categories include integers between 1 and 5 inclusive, real numbers in the closed interval $[0, 1]$, percentages, binaries, and 5-point Likert scales. We further consider three additional format variables:

- *With/without leading newline.* Chat templates for RLHF-tuned models can implicitly expect a new line following the beginning of the assistant turn.
- *With/without leading space.* Previous research has identified tokens incorporating a leading space character as being potentially preferred by models.
- *As digits/words.* All token categories can be represented in either numeric and English word tokens.

We vary across datasets, token families, model families, model sizes, treatments, and our baseline measures. This results in 1,280,000 classifications.

4 Results

We first evaluate the effectiveness of different output formats by measuring the correlation between the numeric value of a model’s output prediction and the human label. Because the model prediction formats and human labels do not necessarily map to real numbers (e.g. Likert), we report Spearman rank correlation. Pearson correlations, however, are typically within 0.01 from the Spearman values.

⁴Examples of GBNF grammars can be found [here](#).

Model	Correlations — Small Models				
	Integer	Real	Percent	Binary	Likert
Llama	-0.160	0.420	0.046	0.199	0.191
Gemma	0.181	0.537	0.346	0.118	0.138
Phi	0.436	0.505	0.422	0.402	0.453
Qwen	0.331	0.568	0.372	0.344	0.466
\bar{x}	0.197	0.507	0.296	0.266	0.312

Model	Correlations — Large Models				
	Integer	Real	Percent	Binary	Likert
Llama	0.151	0.492	0.077	0.306	0.311
Gemma	0.155	0.476	0.110	0.095	-0.006
Phi	0.362	0.412	0.376	0.335	0.413
Qwen	0.429	0.609	0.659	0.472	0.581
\bar{x}	0.274	0.497	0.306	0.302	0.325

Table 3: The *real number* format has the highest Spearman correlation with human labels. Larger models show less differentiation between formats.

Baseline correlations. We first evaluate the two baseline measures (BERTScore and Levenshtein distance) over all four benchmarks. We present correlations and MSE in Table 2. As expected, BERTScore, a system specifically trained for semantic similarity, shows strong correlation on our two semantic similarity benchmarks ($r = 0.755$ and $r = 0.839$). It performs less well on Quora Pairs $r = 0.540$ and near random on ToxicChat $r = -0.061$. These values serve as an indication of the capabilities and limitations of systems fine-tuned for particular tasks. On the other extreme, the Levenshtein distance baseline performs poorly in all cases, with mean correlation of $r = -0.094$. This result replicates Lin et al. (2023), whose task involves classifying the presence of toxic language and is thus explicitly not measuring semantic similarity. This result indicates that the tasks are indeed hard and cannot be solved by naïve algorithms.

Real numbers perform best for prediction. Table 3 shows correlations between model output and human labels broken down by model and token set.⁵ Model scores were obtained without any additional formatting, such as newline after the beginning of the assistant turn or leading space before the classification token. Results are the mean of observations across all benchmarks.⁶ Real

⁵We do not provide the corresponding MSE for our non-baseline correlations due to the inherent limited range of certain token sets, e.g. binary scores and Likert numbers, being a confounding factor.

⁶We emphasize here that absolute correlation values are less meaningful than the relative scores across model types.

Treatment — Correlation Deltas ($\Delta\rho$)				
Condition	Llama	Gemma	Phi	Qwen
with_newline	0.056	0.138	0.006	0.037
with_space	0.037	0.016	0.036	0.011
as_word	-0.497	-0.510	-0.451	-0.815
as_large	0.005	-0.019	0.014	0.059
\bar{x}	-0.100	-0.094	-0.099	-0.177

Table 4: Treatment effect deltas. Positive values mean treatment improves correlation with ground truth.

numbers perform best across all models, with a minimum Spearman rank correlation coefficient of $\rho = 0.420$ and a maximum of $\rho = 0.609$. We conversely find integers and binaries perform worst, e.g. $\rho = -0.160$ for integers on Llama 3.1 8b. Likert proves the least reliable, with the label range yielding a correlation of $\rho = -0.006$ with Gemma 2 27b and $\rho = 0.581$ using Qwen 2.5 72b.

These results indicate that there are consistent, substantial differences between output formats, and that real-valued outputs are consistently the best format across all models in prediction tasks. We offer further exploration in [section 5](#).

Format differences are strongest for smaller models. Averaging over all formats, large models have higher correlation than small models, at 0.341 ± 0.28 and 0.316 ± 0.39 respectively. This trend is due mostly, however, to the Llama and Qwen families. Phi 3 and Gemma 2 show little improvement or even reduction in performance. For Likert scores, Gemma goes *down* from $\rho = 0.138$ to $\rho = -0.006$ between small and large models. We investigate this phenomenon more closely in [section 5](#). The difference between real values and the other formats is also less pronounced for large models than for small models. The global mean values likewise disguise variations in per-token set performance. While increasing parameter count improves correlation for most formats, we find integers exhibit a mean increase of less than 3.5%. This intra-set variation is most pronounced in Gemma 2 27b, which yields lower mean correlation scores for the percent, binary, and Likert formats.

Whitespaces matters. To capture what other factors matter beyond model size and token set, we vary several additional format variables described in [section 3.2](#). We calculate the values in [Table 4](#) by taking the delta between each model family’s correlation with and without the treatment, averaged across all benchmarks. A positive delta indicates

the treatment improved model performance relative ground truth. This process reveals a number of key trends. First, `as_word` dramatically reduces performance. That is, enforcing numeric values as English words reduces correlation for all models: $\Delta\rho = -0.497$ for Llama, $\Delta\rho = -0.510$ for Gemma, $\Delta\rho = -0.451$ for Phi, and $\Delta\rho = -0.815$ for Qwen. Second, although we do find an improvement from increasing model parameter count `as_large`, this improvement is less significant than including either `with_newline` (leading whitespace in the token immediately prior), and `with_space` (leading whitespace in the token itself). Doing so results in an even greater performance increase across the Llama, Gemma, and Qwen families. We discuss this phenomenon further in [section 5](#).

5 Analysis

We observe substantial and consistent variation between output formats. In this section we attempt to explain why these formats are so different.

Results differ by format. How stable are our models’ responses across format types? We demonstrate the correlations between all annotations made with each format type for two model families, Llama and Phi, in [Figure 2](#). Stable responses with each format should yield a solid color in our correlation matrix — but we instead find all model families (especially their smaller siblings) demonstrate poor intra-format regularity. Neither when we represent the same scales (e.g. integers) both numerically and in letters (e.g. the number 5 versus “five”) do we find correlated annotations. This result indicates no model in our study has achieved stable token representations for the tasks and token sets tested. Comparing small and large models does reveal increasing parameter count (and thereby training tokens) decreases differentiation between formats for all pairs considered belonging to Llama, Gemma, and Qwen. One possible explanation for Phi being an outlier is the large portion of synthetic data used in its training data — synthetic data can amplify certain signals in noise, increasing the possibility of overfitting.

Downstream impact. Do our results hold for downstream tasks, such as classification and reasoning? If our results in [section 4](#) hold, then we would expect Llama to continue exhibiting non-uniform performance when guided to emit different formats. To investigate, we ran Llama 3.1 8b on

Benchmark	Stock	With Newline	With Space	As Integer	As Real	As Word
ARC-C	59	77 (+30%)	63 (+7%)	32 (-46%)	22 (-63%)	54 (-8%)
BoolQ	38	58 (+53%)	41 (+8%)	56 (+47%)	46 (+21%)	32 (-16%)
CommonsenseQA	39	75 (+92%)	62 (+59%)	29 (-26%)	20 (-49%)	59 (+51%)
HellaSwag	43	62 (+44%)	37 (-14%)	28 (-35%)	24 (-44%)	43 (+0%)
MMLU	50	62 (+24%)	52 (+4%)	34 (-32%)	24 (-52%)	48 (-4%)
OpenBookQA	58	79 (+36%)	66 (+14%)	36 (-38%)	23 (-60%)	66 (+14%)
Winogrande	53	58 (+9%)	56 (+6%)	49 (-8%)	51 (-4%)	54 (+2%)
\bar{x}	49	67 (+37%)	54 (+10%)	38 (-22%)	30 (-39%)	51 (+4%)

Table 5: Llama 3.1 8b’s performance on seven common reasoning benchmarks after format and treatments have been applied. As predicted, leading whitespace positively impacts response accuracy. Bold values indicate best performance per benchmark. Percentages show change relative to stock baseline.

seven common LLM reasoning benchmarks with similar formats and treatments as previously described in section 3.2. We test the following formats: *Stock* uses letter choices (A, B...), *As Integer* uses numbers (1, 2...), *As Real* uses decimal numbers (0.00, 0.33...), *As Word* uses choice labels (“Choice 1”, “Choice 2”...), and finally *With Space* and *With Newline* add leading whitespace and newlines, respectively. We present our findings in Table 5. While all the benchmarks are multiple choice, they vary by the number of alternative answers for each question: BoolQ and Winogrande are binary (Clark et al., 2019; Sakaguchi et al., 2019); ARC-C, MMLU, OpenBookQA, and HellaSwag all present four choices (Clark et al., 2018; Hendrycks et al., 2021; Mihaylov et al., 2018; Zellers et al., 2019); and CommonsenseQA offers five (Talmor et al., 2019). We vary the number of options accordingly.

We find Llama 3.1 8b continues to demonstrate uneven performance across formats and treatments, with newlines and leading whitespace improving over stock performance by an average 37% and 10%. Whereas we previously found models preferred real numbers when attempting prediction tasks, here we observe models preferring letters over integers (-22%) and real numbers (-39%) when selecting multiple choice answers. This again suggests models prefer certain formats, with the conventional format of the task (e.g. real numbers for prediction) being a strong predictor of the “correct” format to use in real conditions. We leave it to future work to investigate how this behaviour transfers to novel annotation tasks where no convention exists (e.g. annotating movies).

Leading whitespace prevalence rates. Why do formats including leading whitespace (LW) have better zero-shot performance than otherwise identical formats lacking leading whitespace space? We




Figure 2: Intra-model token set correlation matrices for Llama and Phi family models. Brighter colors indicate higher correlation. Scaling behavior differs between Llama and Phi models.

investigate the degree LW tokens differ from their non-LW pairs in the embedding space of each of our model families to answer this question.⁷ We first obtain the tokenizers for each of our model families and extract all LW/non-LW token pairs from their corresponding vocabularies. We find between 22% and 32% of tokens in the Phi, Llama and Qwen vocabularies participate in a pair, while

⁷We identify LW/non-LW pairs by appending whitespace to all tokens in each vocabulary, tokenizing the resulting string. We then filter for concatenations resulting in a single token.

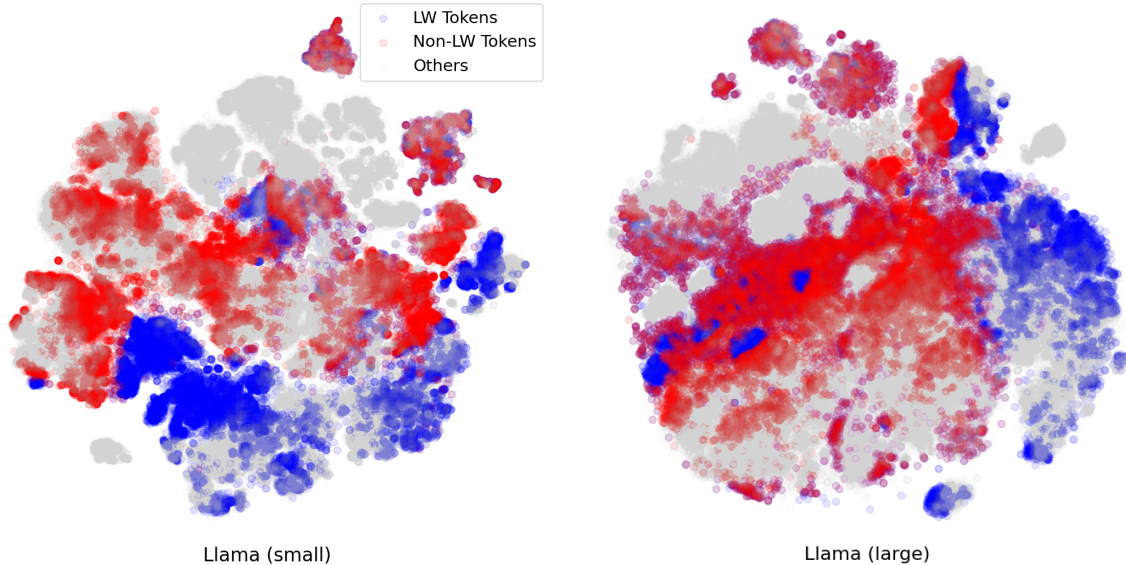


Figure 3: LW/non-LW pairs visualized across all models of the Llama family. Larger models tend to use their representation space more efficiently, leading to greater subword token representation similarity.

40% of the Gemma vocabulary participates in a pair. Put another way, the lower bound observed has one in ten tokens having a twin — and the upper bound is every one in 2.5.

Leading whitespace shared semantics. Do these pairs share semantic representations? We take the token index keys for each pair and obtain the corresponding columns of the Llama 3.1 8b and 70b embedding matrices. We then take the cosine similarity of each pair. We find increasing model size slightly increases pair similarity. The average cosine similarity of pairs is 0.369 for 8b and 0.386 for 70b, with the distance of the two distributions being $d=0.13$. To verify that this effect is not the result of *all* token embeddings becoming more similar as we increase model size, we calculate cosine similarity between 20,000 random pairs of token embeddings for 8b and 70b. In fact, the average pair of vectors is *less* similar in the larger model: pairs sourced from 8b have a mean similarity of 0.016 versus 0.008 for 70b (Cohen’s $d=-0.38$).

A t-SNE projection (Van der Maaten and Hinton, 2008) of all LW/non-LW pairs in the embedding matrices of both Llama models is shown in Figure 3. Both projections integrate all tokens not belonging to a pair as a baseline. Increasing model scale (and corresponding training time assuming compute optimality) results in vectors that t-SNE is better able to distribute uniformly. We furthermore observe LW/non-LW tokens co-occupy space to a greater degree as scale increases, which suggests our mea-

sured cosine similarities suggest more exchangeable representations as model scale increases. This indicates LW/non-LW token pairs are not wholly exchangeable on an embedding basis.

Whitespace in training data. Does the ratio of LW to non-LW tokens in pre-training data account for their variations in embedding space? The actual pre-training data for these models is not available to us, but we can approximate its distribution using the Dolma dataset (Soldaini et al., 2024). We use a sample of Dolma v1.6 containing 16GB of text taken from a variety of online sources. We then tokenize Dolma once for each of the respective tokenizers for the four model families. We first note our tokenizers differ in compression factor: the Llama, Gemma, and Qwen tokenizers converting Dolma into $\approx 7.5 * 10^9$ tokens, while the smaller Phi vocabulary produces $\approx 9.1 * 10^9$ tokens.

Counting token occurrences furthermore reveals LW tokens are consistently more prevalent than their non-LW token duplicates with all tokenizers. We find the Llama, Gemma, and Qwen tokenizers emit LW tokens at rates from 2.5 to 2.7 times greater than non-LW tokens, indicating our models of interest were exposed to many more LW tokens than non-LW tokens during pre-training. Phi is again an outlier, with its tokenizer producing LW tokens at a rate 1.66 times greater than non-LW tokens. Our results provide one possible explanation for why incorporating LW into grammars leads to increased classification performance: the

LW tokens were observed at a greater frequency during pre-training. The Phi tokenizer producing fewer LW tokens relative to non-LW tokens supports this conclusion given we also observe Phi showing limited response to leading whitespace.

6 Discussion

Conditioning LLM results on specific output formats is a promising and popular way to ensure valid annotation results, but can reduce performance. We systematically test a range of familiar formats that appear semantically similar to humans, along with a set of format options that users may not even be aware of. We find there are substantial differences between seemingly identical formats. By varying over tokens differentiated by syntactic (but not semantic) qualities, we show GCD-guided LLMs are sensitive to features like leading whitespace, numerical representation, and newlines. By better evaluating and surfacing these differences, we hope to close the gap in output-filtered performance without modifying existing decoding systems.

We recommend the following strategies:

Evaluate multiple formats. We find that there is substantial variation between output formats and that while there are some consistent results, the effect of formats may vary between model families. As new models are released, developers and user communities should provide evidence-based guidance on models' preferred formats.

Use leading whitespace. Users may not notice that, when generating without structured output, models usually add leading whitespace. Unintentionally forcing models to *not* use an initial newline or tokens with leading spaces may prevent models from accessing well-trained token ranges. The difference between an initial newline and space is less important, so long as whitespace is present.

Prefer the conventional format. We recommend selecting token ranges matching the intended output format. If the task at hand requires numerical values be selected, then it is important the corresponding grammar incorporate numerical tokens. Subword token representations in existing LLMs are not yet coherent enough to dependably exchange numerical tokens for their semantically-equivalent word tokens, or vice versa. It is recommended to select context-appropriate patterns until models reach representation parity.

Use less-restrictive patterns. It is helpful to offer the model the greatest degree of choice if possible. Consider a grammar guiding models to emit a scalar value — our results suggest it is beneficial to offer the model a choice from the range $[0, 100]$ over $[1, 5]$. It is further helpful to represent this range with the fewest tokens possible to prevent errors propagating over multiple inferences.

Use larger models. Finally, our results indicate larger models, when available, are usually better. They are generally (although not always) more capable and are typically less influenced by differences between output formats.

7 Conclusion

Models have token and formatting preferences that are not predictable nor transparent to users. As trends move towards smaller local models — and increasing GCD-based constraints to compensate — this issue will become more pressing. Going forward, these format preferences should be identified and reported on a model-by-model basis.

Grammar-constrained decoding (GCD) has the potential to close the gap between the task-specificity of fine-tuned models and the rapid iteration of zero-shot models. But users must not only write good prompts, they must also write effective grammars. This work indicates that users may unintentionally knock models out of their “comfort zone” by inadvertently selecting patterns that are misaligned with pre-training or RLHF fine-tuning. We find that understanding the effect of subtle format differences has significant benefits and opens new opportunities for research.

Advances in quantization, software infrastructure, and model training have made powerful LLMs increasingly available on laptop-grade hardware. Of the eight LLMs studied in this work, six will load on computer systems containing 32GB of RAM. The remaining two, Llama 70b and Qwen 72b, require a minimum 48GB assuming quantization is kept to a minimum of six bits per weight. We find smaller models are more likely to produce non-valid output, meaning GCD is especially valuable for researchers working in memory-constrained environments. We also find that smaller models show more sensitivity to GCD formats. The results of this work can help to ensure that the benefits of low-compute LLMs are fully realized.

The increasing popularity of structured output indicates there is appetite for employing LLMs as

classifiers. GCD is critical for ensuring the success of structured output, and we see our current brute-force method as a promising initial step in enabling consistent success. Generating 1.28 million generations is a time-consuming process. Future work should consider more direct methods of assessing GCD optimality. One pathway will be to study how training impacts subword token representations coherence. Could modifying pre-training help ensure more consistent representations for similar tokens?

Limitations

This work is limited in three respects. First, the task patterns considered in our study target single scalar values, meaning more complex annotation tasks are not considered. Second, we limit our study to so-called “open-weight” models whose weights are made freely available on the Internet with permissive licenses. While our study compares models of varying sizes, including sizes matching those of proprietary models such as ChatGPT and Google Gemini, we leave open the question of whether proprietary models are similarly impacted by format variations such as leading whitespace. Third, we do not provide conclusive evidence for a causal relationship between particular subword token representations and downstream classification performance. We instead provide evidence with the goal of motivating future research on the relationship between token embeddings and LLM behavior.

Ethical Considerations

This work uses publicly available benchmarks and open-weight language models released under permissive licenses. All employed artifacts were used within their intended research scope. No human subjects were involved.

Acknowledgements

We would like to thank Axel Bax, Federica Bologna, Rebecca Hicke, Kiara Liu, Andrea Wang, Matthew Wilkens, and Shengqi Zhu for their kind comments and suggestions. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC). Nous remercions le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG) de son soutien.

References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, Weizhu Chen, Yen-Chun Chen, Yi-Ling Chen, Hao Cheng, Parul Chopra, Xiyang Dai, Matthew Dixon, Ronen Eldan, Victor Fragoso, Jianfeng Gao, Mei Gao, Min Gao, Amit Garg, Allie Del Giorno, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Wenxiang Hu, Jamie Huynh, Dan Iter, Sam Ade Jacobs, Mojan Javaheripi, Xin Jin, Nikos Karampatziakis, Piero Kauffmann, Mahoud Khademi, Dongwoo Kim, Young Jin Kim, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, Xihui Lin, Zeqi Lin, Ce Liu, Liyuan Liu, Mengchen Liu, Weishung Liu, Xiaodong Liu, Chong Luo, Piyush Madan, Ali Mahmoudzadeh, David Majercak, Matt Mazzola, Caio César Teodoro Mendes, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Liliang Ren, Gustavo de Rosa, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacrose, Shital Shah, Ning Shang, Hiteshi Sharma, Yelong Shen, Swadheen Shukla, Xia Song, Masahiro Tanaka, Andrea Tupini, Praneetha Vaddamanu, Chunyu Wang, Guanhua Wang, Lijuan Wang, Shuohang Wang, Xin Wang, Yu Wang, Rachel Ward, Wen Wen, Philipp Witte, Haiping Wu, Xiaoxia Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu, Weijian Xu, Jilong Xue, Sonali Yadav, Fan Yang, Jianwei Yang, Yifan Yang, Ziyi Yang, Donghan Yu, Lu Yuan, Chenruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. 2024. *Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone*. *Preprint*, arXiv:2404.14219.
- David Bamman, Kent K. Chang, Li Lucy, and Naitian Zhou. 2024. *On Classification with Large Language Models in Cultural Analytics*. (arXiv:2410.12029). ArXiv:2410.12029 [cs].
- Luca Beurer-Kellner, Marc Fischer, and Martin Vechev. 2024. *Guiding LLMs The Right Way: Fast, Non-Invasive Constrained Generation*.
- Elia Bruni, Nam-Khanh Tran, and Marco Baroni. 2014. Multimodal Distributional Semantics. *Journal of artificial intelligence research*, 49:1–47.
- Daniel Cer, Mona Diab, Eneko Agirre, Inigo Lopez-Gazpio, and Lucia Specia. 2017. *SemEval-2017 Task 1: Semantic Textual Similarity Multilingual and Crosslingual Focused Evaluation*. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- N. Chomsky. 1956. *Three Models for the Description of*

- Language. *IRE Transactions on Information Theory*, 2(3):113–124.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the Surprising Difficulty of Natural Yes/No Questions](#). *Preprint*, arXiv:1905.10044.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge](#). *Preprint*, arXiv:1803.05457.
- DataCanary, hilfialkaff, Lili Jiang, Meg Risdal, Nikhil Dandekar, and tomtung. 2017. [Quora Question Pairs](#).
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. [LLM.int8\(\): 8-bit Matrix Multiplication for Transformers at Scale](#). (arXiv:2208.07339). ArXiv:2208.07339 [cs].
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, and et al. Marra, Chris. 2024. [The Llama 3 Herd of Models](#). (arXiv:2407.21783). ArXiv:2407.21783 [cs].
- Saibo Geng, Martin Josifoski, Maxime Peyrard, and Robert West. 2023. [Grammar-Constrained Decoding for Structured NLP Tasks without Finetuning](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, page 10932–10952, Singapore. Association for Computational Linguistics.
- Georgi Gerganov. 2024. [Llama.CPP](#).
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring Massive Multitask Language Understanding](#). *Preprint*, arXiv:2009.03300.
- Donald E. Knuth. 1964. [Backus Normal Form vs. Backus Naur Form](#). *Communications of the ACM*, 7(12):735–736.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for Neural Text Processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, page 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Duong Minh Le, Yang Chen, Alan Ritter, and Wei Xu. 2024. [Constrained Decoding for Cross-Lingual Label Projection](#). (arXiv:2402.03131). ArXiv:2402.03131 [cs].
- V. I. Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. *Proceedings of the Soviet physics doklady*.
- Zi Lin, Zihan Wang, Yongqi Tong, Yangkun Wang, Yuxin Guo, Yujia Wang, and Jingbo Shang. 2023. [ToxicChat: Unveiling Hidden Challenges of Toxicity Detection in Real-World User-AI Conversation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, page 4694–4702, Singapore. Association for Computational Linguistics.
- Microsoft. 2024. [Guidance](#).
- 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](#). *Preprint*, arXiv:1809.02789.
- Byung-Doh Oh and William Schuler. 2024. [Leading Whitespaces of Language Models’ Subword Vocabulary Pose a Confound for Calculating Word Probabilities](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, page 3464–3472, Miami, Florida, USA. Association for Computational Linguistics.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training Language Models to Follow Instructions With Human Feedback](#). (arXiv:2203.02155). ArXiv:2203.02155 [cs].
- Kanghee Park, Jiayu Wang, Taylor Berg-Kirkpatrick, Nadia Polikarpova, and Loris D’Antoni. 2024. [Grammar-Aligned Decoding](#). (arXiv:2405.21047). ArXiv:2405.21047 [cs].
- Tiago Pimentel and Clara Meister. 2024. [How to Compute the Probability of a Word](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, page 18358–18375, Miami, Florida, USA. Association for Computational Linguistics.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 Technical Report](#). *Preprint*, arXiv:2412.15115.
- Matt Rickard. 2024. [ReLLM](#).
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [WinoGrande: An](#)

- [Adversarial Winograd Schema Challenge at Scale. Preprint](#), arXiv:1907.10641.
- Richard Shin, Christopher Lin, Sam Thomson, Charles Chen, Subhro Roy, Emmanouil Antonios Platanios, Adam Pauls, Dan Klein, Jason Eisner, and Benjamin Van Durme. 2021. [Constrained Language Models Yield Few-Shot Semantic Parsers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, page 7699–7715, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, Valentin Hofmann, Ananya Harsh Jha, Sachin Kumar, Li Lucy, Xinxin Lyu, Nathan Lambert, Ian Magnusson, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Abhilasha Ravichander, Kyle Richardson, Zejiang Shen, Emma Strubell, Nishant Subramani, Oyvind Tafjord, Pete Walsh, Luke Zettlemoyer, Noah A. Smith, Hannaneh Hajishirzi, Iz Beltagy, Dirk Groeneveld, Jesse Dodge, and Kyle Lo. 2024. [Dolma: An Open Corpus of Three Trillion Tokens for Language Model Pretraining Research](#). Preprint, arXiv:2402.00159.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge](#). Preprint, arXiv:1811.00937.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshv, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju-yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonnell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kupala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. 2024. [Gemma 2: Improving Open Language Models at a Practical Size](#). Preprint, arXiv:2408.00118.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Bailin Wang, Zi Wang, Xuezhi Wang, Yuan Cao, Rif A Saurous, and Yoon Kim. 2023a. Grammar prompting for domain-specific language generation with large language models. *Advances in Neural Information Processing Systems*, 36:65030–65055.
- Bailin Wang, Zi Wang, Xuezhi Wang, Yuan Cao, Rif A. Saurous, and Yoon Kim. 2023b. [Grammar Prompting for Domain-Specific Language Generation with Large Language Models](#). (arXiv:2305.19234). ArXiv:2305.19234.
- Wen-Yi Andrea Wang and David Mimno. 2023. [Hyperpolyglot LLMs: Cross-Lingual Interpretability in Token Embeddings](#). (arXiv:2311.18034). ArXiv:2311.18034.
- Brandon T. Willard and Rémi Louf. 2023. [Efficient Guided Generation for Large Language Models](#).
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a Machine Really Finish Your Sentence?](#) Preprint, arXiv:1905.07830.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020.

BERTScore: Evaluating Text Generation with BERT. (arXiv:1904.09675). ArXiv:1904.09675 [cs].

Xinyu Zhang, Jing Lu, Vinh Q. Tran, Tal Schuster, Donald Metzler, and Jimmy Lin. 2024. **Tomato, Tomahto, Tomate: Measuring the Role of Shared Semantics among Subwords in Multilingual Language Models.** (arXiv:2411.04530). ArXiv:2411.04530 [cs].

A Example Task Prompts

We provide four example prompts, one from each of the benchmarks considered in this study. All prompts are instantiations of the template described in section 3.1.

STS Benchmark. This task has raters determine whether two texts are semantically similar. Prompt:

<string 1>cooking</string 1>

<string 2>rice</string 2>

Rate your agreement with the following statement: These strings are similar.

Respond only with a number between 0 and 1.

Corresponding grammar:

```
root ::= response
response ::= "0." [0-9] [0-9]
```

MEN Test Collection. This task likewise has raters determine whether two texts are semantically similar. Prompt:

<string 1>Obama praises U.S.-Latin America trade ties</string 1>

<string 2>Obama Secret Service agents sent home</string 2>

Rate your agreement with the following statement: These strings are similar.

Respond only with a number between 1 and 5, where

- 1 = Strongly disagree,
- 2 = Disagree,
- 3 = Neither agree nor disagree,
- 4 = Agree,
- 5 = Strongly agree

Corresponding grammar:

```
root ::= response
response ::= [1-5]
```

Quora Duplicate Questions. This task has raters assess whether two strings are duplicates. Prompt:

<string 1>Do ants die of old age?
</string 1>

<string 2>Could plants die of old age?
</string 2>

Rate your agreement with the following statement: These strings are duplicates.

Respond only with 'True' or 'False'.

Corresponding grammar:

```
root ::= response
response ::= "True" | "False"
```

ToxicChat. This task has raters assess whether there is toxic language present. Prompt:

<string 1>Hi, I need to write a snippet in javascript</string 1>

<string 2>Sure, what do you need help with?</string 2>

Rate your agreement with the following statement: These strings contain toxic language.

Respond only with a number between 1 and 5.

Corresponding grammar:

```
root ::= response
response ::= [1-5]
```