

From Language Models over Tokens to Language Models over Characters

Tim Vieira¹ Benjamin LeBrun² Mario Giulianelli¹ Juan Luis Gastaldi¹ Brian DuSell¹ John Terilla³
Timothy J. O'Donnell^{4,2,5} Ryan Cotterell¹

Abstract

Modern language models are internally—and mathematically—distributions over *token* strings rather than *character* strings, posing numerous challenges for programmers building user applications on top of them. For example, if a prompt is specified as a character string, it must be tokenized before passing it to the token-level language model. Thus, the tokenizer and consequent processing are very sensitive to the specification of the prompt (e.g., whether the prompt ends with a space or not). This paper presents algorithms for converting token-level language models to character-level ones. We present both exact and approximate algorithms. In the empirical portion of the paper, we benchmark the practical runtime and approximation quality. Across four publicly available language models, we find that—even with a small computation budget—our method is able to accurately approximate the character-level distribution at reasonably fast speeds, and that a significant improvement in the language model’s compression rate (bits/byte) is achieved.

 <https://github.com/genlm/genlm-bytes>

1. Introduction

Modern language models are engineered as probability distributions over strings of *tokens* rather than strings of *characters*. However, this leads to a fundamental tension between the users of language models and the engineers who build them. Specifically, token-level models are rife with unintuitive behaviors that—without a technical fix—baffle users. As an illustrative example of a common user complaint, we exhibit the prompt boundary problem (see below). This paper provides a principled solution to

the prompt boundary problem as well as other oddities that make interfacing with token-level language models with character-level prompts hard for users.

Tokenized language models: A brief overview.¹ Let Σ be an alphabet of characters, and let Σ^* denote the set of all strings that can be built from it. Suppose there is a true distribution p_Σ^* over Σ^* that we seek to model. We observe a training corpus of character strings: $\sigma^{(1)}, \dots, \sigma^{(M)} \stackrel{\text{i.i.d.}}{\sim} p_\Sigma^*$. However, rather than estimating a language model that approximates p_Σ^* directly, we employ a (possibly stochastic) tokenizer τ that transforms the training corpus into a corpus of *token* strings. $\delta^{(1)} \sim \tau(\cdot \mid \sigma^{(1)}), \dots, \delta^{(M)} \sim \tau(\cdot \mid \sigma^{(M)})$. Next, we estimate a token-level language model to fit the strings $\delta^{(1)}, \dots, \delta^{(M)}$. Lastly, we use p_Δ to generate character strings through the following generative process: (i) sample $\delta \sim p_\Delta$ and (ii) return $\sigma = \kappa(\delta)$ where κ is a decoding function. Let p_Σ denote the resulting distribution of this process. Practically, we hope that the choice of τ and κ should aid in our ability to estimate p_Σ^* in the sense that $p_\Sigma^* \approx p_\Sigma$. Commonly used tokenizers in the realm of LLMs use tokenizers τ that break long strings into chunks. Intuitively, generating chunks instead of individual characters helps because it effectively shortens strings without obfuscating them using a complicated encoding.

The prompt boundary problem. Consider the case of GPT2 (Radford et al., 2019), which was trained over token strings created from byte-pair encoding (BPE; Sennrich et al. (2016); Gage (1994)). Suppose we wish to generate continuations of the prompt: `"In_the_kingdom_of_the_blind,_the`. Unfortunately, the *interface* to p_Δ is not ideal: it does not accept a *character* string; thus, it is common to encode it as a *token* string with an encoding function τ_{BPE} :^{2,3}

```
 $\tau_{\text{BPE}}(\text{"In\_the\_kingdom\_of\_the\_blind,\_the"})$   
 $= [1, 818, 262, 13239, 286, 262, 7770, 11, 262]$ 
```

¹ETH Zürich ²Mila ³City University of New York ⁴McGill University ⁵Canada CIFAR AI Chair . Correspondence to: Tim Vieira <tim.f.vieira@gmail.com>.

¹We adopt the notation of Gastaldi et al. (2025), who gave a general characterization for when training with tokenization allows for consistent estimation of the true distribution over characters.

²In practice, τ_{BPE} outputs an integer sequence; we provide the substring *gloss* for readability.

³We write $\delta = \tau(\sigma)$ when τ is deterministic (i.e., a function).

If we complete the prompt by taking the most likely next token (*greedy completion*), we generate the following:

```
[_one - eyed _man _is _king ."]
[530 , 12 , 18834 , 582 , 318 , 5822 , 526]
```

Now that we have our generated output, we can apply the decoding function κ that maps token strings to character strings as part of the tokenization protocol:

```
 $\kappa_{\text{BPE}}(\dots) = \text{"In\_the\_kingdom\_of\_the\_blind,\_the\_one-eyed\_man\_is\_king."}$ 
```

This is a good completion, as the string is a well-known proverb. However, if we tweak the prompt *ever so slightly* by inserting a trailing whitespace:

```
 $\tau_{\text{BPE}}(\text{"In\_the\_kingdom\_of\_the\_blind,\_the\_"})$ 
 $= [1, 818, 262, 13239, 286, 262, 7770, 11, 262, 220]$ 
```

greedy completion gives $[\text{ills_of_the_world_are_seen}, \dots]$ because the conditional probability (p_{Δ}) of generating the characters we want becomes highly unlikely, dropping from $0.98 = \vec{p}_{\Delta}(\text{one}_{530} \mid \tau_{\text{BPE}}(\dots \text{the}))$ to $5 \times 10^{-7} = \vec{p}_{\Delta}(\text{one}_{505} \mid \tau_{\text{BPE}}(\dots \text{the_}))$ upon appending the $_$ character. Hence, the trailing white space undesirably impacts the output, which no longer corresponds to the proverb.

We formalize the **prompt boundary problem** as follows: Let $\sigma \in \Sigma^*$ be our initial prompt, such as `"In_the_kingdom_of_the_blind,_the"`. Now, consider a pair of future strings $\sigma' \cdot \alpha$ and $\sigma' \cdot \beta$ with a common prefix σ' (e.g., `_one` and `_ills`). We say a language model interface suffers from the prompt boundary problem if moving the common prefix σ' (e.g., `_`) across the boundary into the prompt does not preserve the relative probability of the pair of future strings, i.e., the relative probability of $\sigma' \cdot \alpha$ and $\sigma' \cdot \beta$ given σ does not equal the relative probability of α and β given $\sigma \cdot \sigma'$. Our example fails this test because the probability of `_one` was much higher than `_ills`, but then, the relative order swaps after moving `_` into the prompt.

Our perspective is that the prompt boundary problem arises from incorrectly conditioning the *token-level* language model on a string of *characters* by using $\tau(\sigma)$ rather than finding token strings that best match the prompt σ .⁴

\hookrightarrow *The token healing heuristic:* [Lundberg & Ribeiro \(2023\)](#) present *token healing* as a heuristic that mitigates the prompt boundary problem; it works as follows: (1) Tokenize (i.e., encode) the prompt. (2) Backup the tokenized prompt to the penultimate token. (3) Generate the next token subject to the constraint that it starts with the unmatched substring at the

end of the prompt. (4) Continue generating as usual. Now, we can see how token healing patches the running example:

```
 $\tau_{\text{BPE}}(\text{"In\_the\_kingdom\_of\_the\_blind,\_the\_"})$ 
 $= [1, 818, 262, 13239, 286, 262, 7770, 11, 262, 220]$ 
```

The most probable next token starting with `_` is `one530`. Now, generating the remaining tokens recovers the desired output, as it matches the prompt before we added the whitespace.

Unfortunately, backing up one token is insufficient for the general case, as the following example will illustrate. Consider generating from GPT2 using the prompt `Hello,_worl`:

```
 $\tau_{\text{BPE}}(\text{Hello,\_worl}) = [\text{Hello ,\_wor l}]$ 
 $[15496, 11, 476, 75]$ 
```

The most likely next *character* ought to be `d` as `Hello,_world` is a common expression (popularized in educational material). However, the most likely *token* results in `Hello,_worldwide`, an apparent misspelling of `worldwide`.⁵ Unfortunately, token healing’s strategy of backing up by a single token cannot salvage the poor tokenization as `l75` is still the most common next token that is consistent with the string `l`. Thus, after generating `l75`, we are back where we started, generating `[Hello ,_wor l wide]`.

\hookrightarrow *Getting it right:* A simple “probability 101” expression tells us the correct solution to the prompt boundary problem. Consider a Δ^* -valued random variable Y , distributed according to p_{Δ} . Then, the correct way to sample from p_{Δ} conditioned on a character string σ is according to

$$p_{\Delta|\Sigma}(\delta \mid \sigma) \stackrel{\text{def}}{=} \mathbb{P}_{Y \sim p_{\Delta}} [Y = \delta \mid \kappa(Y) \succeq \sigma] \quad (1)$$

where we have conditioned on the event that the decoded string $\kappa(Y)$ has σ as a prefix (i.e., $\kappa(Y) \succeq \sigma$). While innovative with respect to the literature, the expression $\mathbb{P}_{Y \sim p_{\Delta}} [Y = \delta \mid \kappa(Y) \succeq \sigma]$ conveys the probability we are interested in precisely and concisely. For the more procedurally minded, this corresponds to the following process:

```
1 def conditional_token_generation( $\sigma$ ):
2   while True:
3     sample  $\delta \sim p_{\Delta}$ 
4     if  $\kappa(\delta) \succeq \sigma$ : return  $\delta$  # accept
```

This is, of course, very inefficient, but fear not—we will provide an equivalent, efficient algorithm.

Our method finds a set of token strings that form a *covering*, a key technical concept we introduce in this paper. We will provide the precise definition in due course; for now, we will illustrate the covering of `Hello,_worl`:

⁴Probabilistic condition preserves the necessary relative probabilities that we specified in the prompt boundary problem.

⁵The misspelling is a testament to the extent to which the tokenized prompt is out-of-distribution.

██████████ (0.9820714) :	[15496, 11, 995]
l(0.0106702) :	[15496, 11, 11621]
l(0.0070749) :	[15496, 11, 8688]
l(0.0000830) :	[15496, 11, 43249]
l(0.0000369) :	[15496, 11, 29081]
l(0.0000225) :	[28254, 78, 11, 995]
l(0.0000179) :	[15496, 11, 476, 75]
:	:

Notice that each token string in the covering has the character string `Hello, _world` as a prefix (after decoding), but it may have some extra characters (marked by underlining) in the partially matched last token—just like token healing. The size of the complete covering may, in general, be exponential in the length of the character string. However, we can enumerate its high-probability members very quickly in practice. Since the worst-case time to find the top- K elements of the covering might be exponential, we provide a more aggressive approximation based on beam search that gives a good approximation even with small beam sizes.

Unlike token healing, the sequence `[15496, 11, 995]` is the highest probability member of the covering; thus, it is not pruned away by the aggressive heuristics based on τ , as in token healing. The covering is defined only in terms of κ , and the pruning is based on the language model probability, which typically prioritizes the token strings that are most reflective of the language model’s training data. We provide an algorithm for correctly conditioning a token-level model on a character string in §3.4.

Character-level model. At the character level, working with the language model is intuitive. Consider, again, our example illustrating the prompt boundary problem. Appending whitespace behaves predictably, as it satisfies the *probabilistic chain rule*:

$$\begin{aligned} \vec{p}_{\Sigma}(_one \mid \text{"In_the_kingdom_of_the_blind,_the"}) \\ = \vec{p}_{\Sigma}(_ \mid \text{"In_the_kingdom_of_the_blind,_the"}) \\ \cdot \vec{p}_{\Sigma}(one \mid \text{"In_the_kingdom_of_the_blind,_the_}) \end{aligned}$$

Here, \vec{p}_{Σ} denotes the character-level model’s conditional distribution. Similarly, recall the `Hello, _world` example above. The character-level model correctly infers that `d` is the most likely next character given `Hello, _worl`. The computation of this conditional probability is simply the total probability of the covering of `Hello, _world` divided by the total probability of the covering of `Hello, _worl`. These quantities are derived from our concept of covering, which directly leads to an algorithm for determining the distribution over possi-

ble next characters. Beyond the prompt boundary problem, computing the conditional probability of a character string given a token-level language model has many applications.

Applications. Aside from making character-level conditioning well-behaved, we highlight a few applications of language models requiring careful character-level reasoning:

↪ *Character-level constraints:* Enforcing character-level constraints on allowed strings is a promising area that has received much recent attention (e.g., Scholak et al., 2021; Poesia et al., 2022; Geng et al., 2023; Microsoft, 2023; Willard & Louf, 2023; Koo et al., 2024; Loula et al., 2025).

↪ *Computational psycholinguistics:* Computing the contextual surprisal (negative log probability) of a character substring to predict human reading times (Hale, 2001; Levy, 2008). Two recent papers (Oh & Schuler, 2024; Pimentel & Meister, 2024) have given algorithms for computing the surprisal of whitespace-separated *words* under a number of strong assumptions. Our algorithms can compute the contextual surprisal of arbitrary character strings. Giulianelli et al. (2024) show experimentally that having the flexibility to compute character substring surprisals leads to a more predictive model of reading time than a fixed notion of a word.

Does it work? In the experimental portion of our paper (§4), we report the empirical runtime of our algorithm for converting token-level language models to character-level ones and quantify its accuracy in estimating the conditional distribution over characters. We find that even with a limited computational budget, our method provides an accurate estimate of the conditional distribution over the next character under four publicly available language models.⁶ We also find that the compression rate (bits/byte) is significantly improved by estimating the probability of the corpus as a character string rather than a canonical token string.

2. Background

2.1. Alphabets and Strings

An **alphabet** Γ is a non-empty, finite set of elements called **symbols**. A **string** γ over alphabet Γ is a finite sequence $\gamma = \gamma_1 \dots \gamma_N$ for some $0 \leq N < \infty$ of symbols where $\gamma_1, \dots, \gamma_N \in \Gamma$. Let $|\gamma|$ denote the string’s length N . We denote the empty string as ε . For any alphabet Γ , let Γ^* denote the set of all strings over Γ , and let Γ^+ denote the set of all non-empty strings over Γ . For any two strings $\gamma', \gamma'' \in \Gamma^*$, we denote their concatenation as $\gamma' \cdot \gamma''$. Additionally, we define $S \cdot S' \stackrel{\text{def}}{=} \{\gamma \cdot \gamma' \mid \gamma \in S, \gamma' \in S'\}$ for any $S, S' \subseteq \Gamma^*$. Given a string γ such that $|\gamma| \geq t$, let $\gamma_{<t}$ denote the string made from the first $t-1$ characters of γ . We write $\gamma \preceq \gamma'$

⁶Specifically, Llama-3.2-1B, Meta-Llama-3.1-8B, DeepSeek-R1-Distill-Llama-8B, and phi-4 (14B).

if γ is a prefix of γ' and $\gamma \prec \gamma'$ if γ is a proper prefix of γ' . The relation \preceq defines a partial order on Γ^* . We write \succeq and \succ to refer to the relations \preceq and \prec with their respective arguments transposed.

2.2. Language Models and Prefix Probability

A **language model** p_Γ is a probability distribution over Γ^* where Γ is an alphabet. Let Y be a Γ^* -valued random variable distributed according to p_Γ and $\gamma \in \Gamma^*$. The **prefix probability** $\vec{p}_\Gamma(\gamma)$ is the probability that Y has prefix γ :

$$\vec{p}_\Gamma(\gamma) \stackrel{\text{def}}{=} \mathbb{P}_{Y \sim p_\Gamma} [Y \succeq \gamma] = \sum_{\gamma' \in \Gamma^*} \mathbb{1}\{\gamma' \succeq \gamma\} p_\Gamma(\gamma') \quad (2)$$

We also define the following shorthand for the **conditional prefix probability** $\vec{p}_\Gamma(\gamma' \mid \gamma)$ as the probability of the event $Y \succeq \gamma \cdot \gamma'$ provided that $Y \succeq \gamma$:

$$\vec{p}_\Gamma(\gamma' \mid \gamma) \stackrel{\text{def}}{=} \mathbb{P}_{Y \sim p_\Gamma} [Y \succeq \gamma \cdot \gamma' \mid Y \succeq \gamma] = \frac{\vec{p}_\Gamma(\gamma \cdot \gamma')}{\vec{p}_\Gamma(\gamma)} \quad (3)$$

We may express the probability of γ as a product of conditional prefix probabilities:

$$p_\Gamma(\gamma) = \vec{p}_\Gamma(\text{EOS} \mid \gamma) \prod_{t=1}^{|\gamma|} \vec{p}_\Gamma(\gamma_t \mid \gamma_{<t}) \quad (4)$$

where each $\vec{p}_\Gamma(\gamma_t \mid \gamma_{<t})$ is an instance of Eq. (3), and

$$\vec{p}_\Gamma(\text{EOS} \mid \gamma) \stackrel{\text{def}}{=} \frac{p_\Gamma(\gamma)}{\vec{p}_\Gamma(\gamma)} \quad (5)$$

Here, EOS is a distinguished **end-of-string symbol** that cannot appear in any alphabet. Of particular interest are the single-symbol conditional prefix distributions $\vec{p}_\Gamma(\cdot \mid \gamma_{<t})$, as they may each be interpreted as a probability distribution over the set $\Gamma \cup \{\text{EOS}\}$ —in fact, modern language models⁷ are *defined* via the product in Eq. (4) where each single-symbol conditional prefix probability comes from the learned parametric model.⁸

2.3. Tokenization

We now discuss our basic formalization for tokenization.

Definition 1. An (exact) **tokenization model** is a tuple $(\Sigma, \Delta, \tau, \kappa)$ where

- Σ is an alphabet of **character** symbols

- Δ is an alphabet of **token** symbols
- τ is a (possibly) **stochastic encoder**: $\tau(\cdot \mid \sigma)$ is a probability distribution over Δ^* for each $\sigma \in \Sigma^*$
- $\kappa: \Delta^* \rightarrow \Sigma^*$ is a **decoder function** satisfying **exactness** $\sum_{\delta \in \Delta^*} \mathbb{1}\{\kappa(\delta) = \sigma\} \tau(\delta \mid \sigma) = 1$, for all $\sigma \in \Sigma^*$

Definition 2. A **tokenized language model** p_Σ is a language model over Σ^* that is parameterized by a language model p_Δ over Δ^* and a decoding function $\kappa: \Delta^* \rightarrow \Sigma^*$. This tokenized language model generates character strings via the following process: (i) $\delta \sim p_\Delta$, (ii) $\sigma \leftarrow \kappa(\delta)$. Thus, the character strings σ generated have the distribution:

$$p_\Sigma(\sigma) \stackrel{\text{def}}{=} \mathbb{P}_{Y \sim p_\Delta} [\kappa(Y) = \sigma] \quad (6)$$

Note that $p_\Sigma(\sigma)$ accounts for the fact that many token strings may be associated with a given character string through κ .⁹ To describe that association, we define $\mathcal{E}(\sigma) \stackrel{\text{def}}{=} \{\delta \in \Delta^* : \sigma = \kappa(\delta)\}$, the set of **encodings** for any character string $\sigma \in \Sigma^*$.¹⁰

What about τ ? The reader may notice that τ does not appear in Eq. (6). Although τ is essential for generating training data, once the model p_Δ has been trained, the information in τ is not of immediate practical use. Moreover, attempts to leverage τ seem to lead to faulty heuristics, as we discussed in the introduction. We note that under exactness, $\tau(\sigma)$ must be present in $\mathcal{E}(\sigma)$. This is because exactness implies that $\mathcal{E}(\sigma) \supseteq \{\delta \in \Delta^* : \tau(\delta \mid \sigma) > 0\}$ for all $\sigma \in \Sigma^*$. In the common case where τ is **deterministic**¹¹ we emphasize that $\mathcal{E}(\sigma) \supseteq \{\tau(\sigma)\}$. The tokenization model would need to be **bijective**¹² for $\mathcal{E}(\sigma) = \{\tau(\sigma)\}$. Unfortunately, common tokenizers (e.g., BPE) are *not* bijective.

The mirage of the canonical tokenization. Consider the case when the encoder τ is deterministic. In that case, we write $\delta = \tau(\sigma)$, and we call this δ the **canonical** tokenization of σ . Note that even if τ is deterministic, there may exist many **noncanonical** tokenizations $\delta' \in \mathcal{E}(\sigma)$ such that $\delta' \neq \tau(\sigma)$ with nonzero probability $p_\Delta(\delta') > 0$. Thus, the character string generation process includes a mix of canonical and noncanonical token strings—making it incorrect to only consider a character string’s canonical tokenization when assessing its probability. In practice, the conditional probability $\mathbb{P}_{Y \sim p_\Delta} [Y = \delta \mid \kappa(Y) = \sigma]$ over the encodings δ of a character string σ tends to be highly

⁷E.g., transformers (Vaswani et al., 2017), RNNs (e.g., Mikolov et al., 2010), and n -gram models (e.g., Shannon, 1948).

⁸The reader may notice that the equations for $p_\Gamma(\gamma)$, $\vec{p}_\Gamma(\gamma' \mid \gamma)$, and $\vec{p}_\Gamma(\text{EOS} \mid \gamma)$ are mutually recursive. Thus, they may appear circular. The key to resolving this concern is to recognize that p_Γ is given as a *base case*. We note, however, that some readers may view Eq. (4) as the *definition* of the language model p_Γ , taking the components on the right-hand side of Eq. (4) as the base case.

⁹Many authors (Cao & Rimell, 2021; Chirkova et al., 2023; Phan et al., 2024) have discussed this particular complication; however, no algorithms for computing the next-character distribution exist.

¹⁰Note that $|\mathcal{E}(\sigma)|$ can be very large, e.g., infinite in the worst case. In the case of BPE, it is exponential in $|\sigma|$.

¹¹I.e., $\tau(\delta \mid \sigma) \in \{0, 1\}$ for all $\delta \in \Delta^*, \sigma \in \Sigma^*$.

¹²I.e., a bijective tokenizer is deterministic, exact, and satisfies $\tau(\kappa(\delta)) = \delta$ for all $\delta \in \Delta^*$.

concentrated on the canonical tokenizations, as illustrated in Example 1 below.

Example 1. Below, we show GPT2’s top encodings from $\mathcal{E}(\text{Hello, world})$, ranked by their conditional probability.

This short string has 78 encodings from numerous partitions of `Hello, world` into substrings of the tokenization alphabet Δ . We see that the probability heavily concentrates on the top string, which is canonical.

(0.9999719)	<div> <div></div> <div> Hello , _world 15496' 11' 995 </div> </div>
l(0.0000229)	<div> <div></div> <div> Hell o , _world 28254' 78' 11' 995 </div> </div>
(0.0000024)	<div> <div></div> <div> Hello , _wor ld 15496' 11' 476' 335 </div> </div>
(0.0000017)	<div> <div></div> <div> He llo , _world 1544' 18798' 11' 995 </div> </div>
(0.0000004)	<div> <div></div> <div> H ell o , _world 39' 695' 78' 11' 995 </div> </div>
(0.0000002)	<div> <div></div> <div> Hello , _w orld 15496' 11' 266' 1764 </div> </div>
(0.0000002)	<div> <div></div> <div> H ello , _world 39' 11109' 11' 995 </div> </div>

A character-level interface. A character-level interface to the token-level language model p_Δ is available in the following equations, which hold $\forall \sigma, \sigma' \in \Sigma^*$:

$$\vec{p}_\Sigma(\sigma) = \mathbb{P}_{Y \sim p_\Delta} [\kappa(Y) \succeq \sigma] \quad (7)$$

$$\vec{p}_\Sigma(\sigma' \mid \sigma) = \frac{\vec{p}_\Sigma(\sigma \cdot \sigma')}{\vec{p}_\Sigma(\sigma)} \quad (8)$$

$$\vec{p}_\Sigma(\text{EOS} \mid \sigma) = \frac{p_\Sigma(\sigma)}{\vec{p}_\Sigma(\sigma)} \quad (9)$$

These equations show that we can have a complete character-level language model derived from the tokenized language model if we can compute—or approximate—the necessary summations implied by Eq. (6) and (7); specifically,

$$p_\Sigma(\sigma) = \sum_{\delta \in \Delta^*} \mathbb{1}\{\kappa(\delta) = \sigma\} p_\Delta(\delta) \quad (10)$$

$$\vec{p}_\Sigma(\sigma) = \sum_{\delta \in \Delta^*} \mathbb{1}\{\kappa(\delta) \succeq \sigma\} p_\Delta(\delta) \quad (11)$$

We will develop effective methods for these summations for the family of strict-prefix monotone decoders κ (described in §2.4) where Eq. (10) and Eq. (11) admit a finite summation.

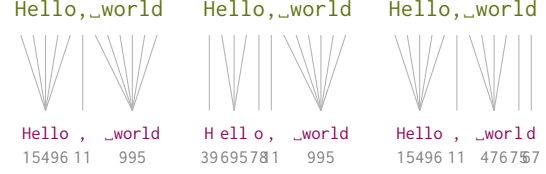
2.4. Key Properties of κ

This section defines the essential properties of κ we require.

Definition 3. We say $\kappa: \Delta^* \rightarrow \Sigma^*$ is

- **prefix monotone** if $\delta \preceq \delta' \implies \kappa(\delta) \preceq \kappa(\delta')$
- **strict-prefix monotone** if $\delta \prec \delta' \implies \kappa(\delta) \prec \kappa(\delta')$

In simpler terms, strict-prefix monotonicity implies that concatenating a token to the encoding necessarily concatenates at least one character to the decoded character string. The diagrams below illustrate how GPT2’s strict-prefix monotone κ gives rise to a certain alignment between three token strings and the character string `Hello, world`:



More formally, every application $\sigma_1 \cdots \sigma_M = \kappa(\delta_1 \cdots \delta_M)$ of a strict-prefix monotone mapping has the following properties. Each token in $\delta_1 \cdots \delta_M$ maps to one or more contiguous characters in $\sigma_1 \cdots \sigma_M$. Moreover, the mappings do not exclude any characters, and no edges of the mapping cross one another. *Strict* prefix monotonicity, in contrast to prefix monotonicity, ensures that there are no deletions of tokens in the mapping, i.e., each token maps to *at least one* character.

Strict-prefix monotonicity is the key structural property required by §3’s algorithms, as it allows us to replace an infinite sum with a finite sum in Proposition 1. We briefly mention an important special case. A decoder κ is **multiplicative** if $\kappa(\delta_1 \cdots \delta_N) = \kappa(\delta_1) \cdots \kappa(\delta_N)$ for all $\delta_1 \cdots \delta_N \in \Delta^*$, and **non-erasing** if $\kappa(\delta) = \varepsilon \implies \delta = \varepsilon$. If κ is multiplicative, it is prefix monotone; if it is also non-erasing, it is strict-prefix monotone. Both BPE and WordPiece are multiplicative and non-erasing.

3. Algorithms

This section gives algorithms for computing $p_\Sigma(\sigma)$, $\vec{p}_\Sigma(\sigma)$, $\vec{p}_\Sigma(\sigma' \mid \sigma)$, $\vec{p}_\Sigma(\text{EOS} \mid \sigma)$, and conditional token generation. We assume throughout that κ is strict-prefix monotone.

3.1. Covering

Eq. (11) shows that we can, in principle, compute the prefix probability $\vec{p}_\Sigma(\sigma)$ by summing over **prefix-encodings** of σ , $\mathcal{P}(\sigma) \stackrel{\text{def}}{=} \{\delta \in \Delta^* : \kappa(\delta) \succeq \sigma\}$. Although $\mathcal{P}(\sigma)$ is infinitely large, we can exploit the prefix monotone structure of κ to find a different way to perform the summation by summing over a *finite* set. Let $\delta \in \Delta^*$, $\delta_1 \cdots \delta_M = \delta$, and $\sigma \in \Sigma^*$. We say that δ **covers** σ if and only if $\kappa(\delta) \succeq \sigma$. Monotonicity ensures that for all $\delta \in \mathcal{P}(\sigma)$, we have that $\forall \delta' \in \Delta^* : \kappa(\delta \cdot \delta') \succeq \sigma$. In other words, any δ that decodes to an extension of σ (i.e., $\kappa(\delta) \succeq \sigma$) will continue to do so if we append tokens to it. Thus, we may additionally qualify the relationship as δ **minimally covers** σ if additionally $\kappa(\delta_1 \cdots \delta_{M-1}) \prec \sigma$. With that in mind, we define $\phi_\sigma(\delta)$ as the *shortest* prefix $\delta' \preceq \delta$ such that $\kappa(\delta') \succeq \sigma$, i.e., ϕ_σ maps any δ that covers σ to a (possibly equal) token string that minimally covers σ . Next, we define the set of *minimal* prefix encodings of σ , which we call the **covering** of σ , $\mathcal{C}(\sigma) \stackrel{\text{def}}{=} \{\phi_\sigma(\delta) \mid \delta \in \mathcal{P}(\sigma)\}$. A more convenient expression for the covering $\mathcal{C}(\sigma)$ of a

string $\sigma \in \Sigma^*$ is equal to the following subset of Δ^* :

$$\mathcal{C}(\sigma) = \begin{cases} \{\varepsilon\} & \text{if } \sigma = \varepsilon \\ \{\delta_1 \dots \delta_M \in \Delta^+ : \kappa(\delta_1 \dots \delta_{M-1}) \prec \sigma \preceq \kappa(\delta_1 \dots \delta_M)\} & \text{otherwise} \end{cases} \quad (12)$$

Example 2.

Recall the covering of the string $\sigma = \text{Hello, world}$ from the introduction. We have repeated it on the right and couched it in our terminology. Note that the complete covering $\mathcal{C}(\text{Hello, world})$ contains 36,608 token strings; we only show the top strings according to their respective \vec{p}_Δ . Below, we list several properties and observations about the structure of the covering:

Hello , _world	[15496:11, 995]
Hello , _worlds	[15496:11, 11621]
Hello , _worldwide	[15496:11, 8688]
Hello , _worldly	[15496:11, 43249]
Hello , _worldview	[15496:11, 29081]
Hell o , _world	[28254:78:11, 995]
Hello , _wor l	[15496:11, 476:75]

- The covering for any given string σ always has the property that each non-empty token string δ in the covering decodes to a string that has σ as a prefix, i.e., $\sigma \preceq \kappa(\delta)$. This is illustrated by the gloss string in the tokenization.
- Note that δ may include a partially matched token at its end (i.e., δ_M in Eq. (12)). We have marked the extra characters by underlining them. We note that the 7th member does not have a partially matched last token.
- Each token string in the covering has at most one partially matched token thanks to the condition $\kappa(\delta) \prec \sigma \preceq \kappa(\delta \cdot \delta')$. The 7th member of the cover has a completely matched last token; hence, there is no underlining.
- We see that if we were to extend any member $\delta \in \mathcal{C}(\sigma)$ with an arbitrary string of additional tokens δ' , it would continue to decode to a string such that $\kappa(\delta \cdot \delta') \succeq \sigma$. Moreover, δ is minimal (i.e., $\phi_\sigma(\delta) = \delta$).

The notion of a covering is used to derive an algorithm for computing character-level probabilities given a token-level language model. We first show how it gives us the prefix probability and subsequently give equations for the remaining quantities of the character-level language model.

Proposition 1. Suppose $(\Sigma, \Delta, \tau, \kappa)$ is a tokenization model where κ is strict-prefix monotone and p_Δ is a token-level language model. Then, the prefix probability $\vec{p}_\Sigma(\sigma)$ for the character-level model Eq. (6) is given by

$$\vec{p}_\Sigma(\sigma) = \sum_{\delta \in \mathcal{C}(\sigma)} \vec{p}_\Delta(\delta), \quad \forall \sigma \in \Sigma^* \quad (13)$$

Proof. See App. B. ■

Eq. (13) is a substantial improvement over Eq. (7) for computing $\vec{p}_\Sigma(\sigma)$. Specifically, we now have a *finite* sum, as

$|\mathcal{C}(\sigma)|$ is finite for all $\sigma \in \Sigma^*$. Bear in mind that the covering's size is likely too large to be practical, as there may still be a large number of summands; however, the set of high-prefix-probability elements of the covering tends to be reasonably small, an observation that we verify in §4, and leverage to develop practical algorithms in §3.

Lastly, we note that the covering contains the set of encodings, i.e., $\mathcal{E}(\sigma) \subseteq \mathcal{C}(\sigma)$; hence, the encodings may be extracted from the covering as follows: $\mathcal{E}(\sigma) = \{\delta \in \mathcal{C}(\sigma) : \kappa(\delta) = \sigma\}$. We may also express the probability of σ in terms of the covering:

$$p_\Sigma(\sigma) = \sum_{\delta \in \mathcal{C}(\sigma)} \mathbb{1}\{\kappa(\delta) = \sigma\} p_\Delta(\delta) \quad (14)$$

3.2. Algorithms for $\vec{p}_\Sigma(\sigma)$ and $p_\Sigma(\sigma)$

The enumeration algorithm will enumerate elements of the covering along with their prefix probability (for convenience). It filters prefixes of token strings that cannot eventually cover the target string σ . The strict-prefix monotonicity property is essential for this filtering.

Our algorithm `enum_cover` performs recursive enumeration of the members of the covering $\mathcal{C}(\sigma)$ along with some metadata. Specifically, the algorithm returns a collection of triples where each triple (p', σ', δ') satisfies $\delta' \in \mathcal{C}(\sigma)$, $p' = \vec{p}_\Delta(\delta')$, and $\sigma' = \kappa(\delta')$.

```

5 def enum_cover( $\sigma_1 \dots \sigma_N$ ):
6   if  $N = 0$ : return [(1,  $\varepsilon$ ,  $\varepsilon$ )] # base case
7   out  $\leftarrow$  []
8   for ( $p', \sigma', \delta'$ ) in enum_cover( $\sigma_1 \dots \sigma_{N-1}$ ):
9     if  $|\sigma'| < N$ : # extend
10      for  $\delta'' \in \Delta$ :
11         $\sigma'' \leftarrow \kappa(\delta' \cdot \delta'')$ 
12        if  $\sigma''_N = \sigma_N$ : # filter
13          out.append(( $p' \cdot \vec{p}_\Delta(\delta'' | \delta'), \sigma'', \delta' \cdot \delta''$ ))
14      elif  $\sigma'_N = \sigma_N$ : # filter character matches
15        out.append(( $p', \sigma', \delta'$ ))
16   return prune( $\sigma_1 \dots \sigma_N$ , out)
```

Note that this method has an additional parameter, the function `prune`, which is used on the last line. This method, as the name suggests, is used to limit the size of the covering to prevent excessive growth. We will discuss this parameter shortly. For now, consider the following definition:

```

17 def prune_nothing( $\sigma_1 \dots \sigma_N$ , out):
18   return out
```

From the output of the enumeration algorithm, we can compute the other key objects and quantities (i.e., $\mathcal{C}(\sigma)$, $\vec{p}_\Sigma(\sigma)$, $\mathcal{E}(\sigma)$, $p_\Sigma(\sigma)$) in the character-level interface.

Time and space complexity. To meaningfully discuss its running time, we assume the following:

- $\kappa(\delta' \cdot \delta'')$ can be evaluated in constant time given $\kappa(\delta')$.
- the cost of evaluating $\vec{p}_{\Delta}(\delta_t \mid \delta_{<t})$ is constant given that $\vec{p}_{\Delta}(\delta_s \mid \delta_{<s})$ has been computed for $0 \leq s < t$.¹³

Under these assumptions, the running time of $\text{enum_cover}(\sigma_1 \dots \sigma_N)$ can be exponential in N when no pruning is used. We provide detailed bounds on the covering's size in App. C. It is straightforward to verify that the space complexity is $\mathcal{O}(|\mathcal{C}(\sigma)|)$, and the running time is $\mathcal{O}(|\Delta| \cdot \sum_{t=1}^{|\sigma|} |\mathcal{C}(\sigma_{<t})|)$ which is dominated by the $|\Delta| \cdot |\mathcal{C}(\sigma)|$ term; thus, $\mathcal{O}(|\Delta| \cdot |\mathcal{C}(\sigma)|)$.

Pruning. We now consider some useful pruning heuristics for the algorithm, which make it an *approximation* but substantially improve its running time. We propose a heuristic based on beam search. This heuristic is very effective: it gives us a linear running time as a function of the character string's length. It has a parameter K that controls the approximation quality. Larger K makes the approximation more accurate, and the approximation becomes exact as K approaches the size of the (largest intermediate) covering. We take K to be a global variable in the pseudocode.

\hookrightarrow *Our pruning heuristic:* Our pruning heuristic enumerates $\leq K$ distinct token strings modulo their last token. This choice allows up to $|\Delta|$ versions of the last token to be enumerated (if it is not completely matched). Thus, the work done at each step is $\mathcal{O}(K \cdot |\Delta|)$, and the size of the pruned list is at most that size. Therefore, the overall running time is $\mathcal{O}(N \cdot K \cdot |\Delta|)$ for a character string of length N .¹⁴

```

19 def prune_top_K_buckets( $\sigma_1 \dots \sigma_N$ , results):
20     buckets  $\leftarrow \{\}$ 
21     for item in results:
22         ( $p, \sigma', \delta_1 \dots \delta_M$ )  $\leftarrow$  item
23         # Exclude a partially matched last
24         key  $\leftarrow \delta_1 \dots \delta_{M-1}$  if  $|\sigma'| > N$  else  $\delta_1 \dots \delta_M$ 
25         buckets[key].append(item)
26     pruned  $\leftarrow []$ 
27     for bucket in buckets.top( $K$ ): # by prob.
28         for item in bucket:
29             pruned.append(item)
30     return pruned
    
```

Bundled beam summing implementation. App. D describes an implementation strategy that improves the constant factors associated with the pseudocode above. The key idea is to group the token sequences that fall into the same bucket in `prune_top_K_buckets` into a *bundle* that represents them compactly. In particular, we can use a trie

¹³In the case of the common transformer language model (Vaswani et al., 2017), this can be achieved with efficient caching and limiting context windows to a constant size.

¹⁴Note: finding the (unordered) set of top- K elements from a set S is possible in $\mathcal{O}(|S|)$ time via the median-of-medians algorithm.

to efficiently filter out the next tokens that disagree with the next character. We can regard the trie as a local language model that generates the next token character-by-character according to the probability assigned by $\vec{p}_{\Delta}(\cdot \mid \delta)$. Each bundle can be unbundled (if necessary) into the respective tuples that the `enum_cover` algorithm maintains.

3.3. Algorithms for $\vec{p}_{\Sigma}(\cdot \mid \sigma)$

This section gives algorithms for computing the character-level conditional prefix probability. Recall the definition of the character-level conditional prefix probability, that is, Eq. (8) and (9), can be computed from a certain ratio of calls to \vec{p}_{Σ} (and p_{Σ} in the case of EOS). From here, Eq. (8) and (9) give a straightforward algorithm for computing the distribution over $\Sigma \cup \{\text{EOS}\}$ given $\sigma \in \Sigma^*$. However, a direct translation would perform duplicate work. Therefore, we provide the following version, which reuses work between the calls to \vec{p}_{Σ} than directly evaluating those equations would do.

```

31 def next_character_probability( $\sigma$ ):
32      $N \leftarrow |\sigma|$ ;  $Z \leftarrow 0$ ;  $\bar{p} \leftarrow \{\sigma' : 0 \text{ for } \sigma' \in \Sigma \cup \{\text{EOS}\}\}$ 
33     for ( $p', \sigma', \delta'$ ) in enum_cover( $\sigma$ ):
34          $Z += p'$ 
35         if  $|\sigma'| = N$ : # i.e.,  $\sigma' = \sigma$ 
36              $\bar{p}(\text{EOS}) += p' \cdot \vec{p}_{\Delta}(\text{EOS} \mid \delta')$ 
37             for  $\delta'' \in \Delta$ : # extend
38                  $\sigma'' \leftarrow \kappa(\delta' \cdot \delta'')$ 
39                  $\bar{p}(\sigma''_{N+1}) += p' \cdot \vec{p}_{\Delta}(\delta'' \mid \delta')$ 
40             else: # i.e.,  $\sigma' \succeq \sigma$ 
41                  $\bar{p}(\sigma'_{N+1}) += p'$ 
42     return  $\bar{p}/Z$  #  $Z = \vec{p}_{\Sigma}(\sigma)$ 
    
```

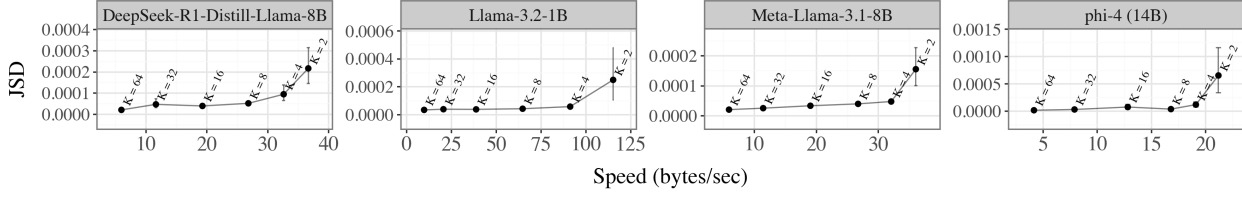
3.4. Conditional Generation $p_{\Delta|\Sigma}(\delta \mid \sigma)$

This section gives a simple algorithm for correctly generating a token string Y that has a given character-level prompt σ as its prefix. This algorithm is equivalent to the algorithm in the introduction but significantly faster.

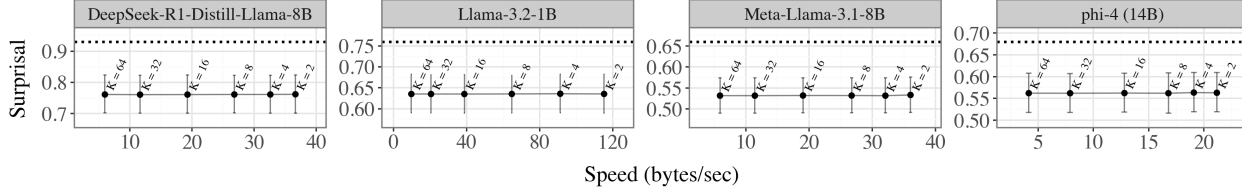
The algorithm works by enumerating the covering $\mathcal{C}(\sigma)$, drawing a token string from it in proportion to its prefix probability, and finishing the token string by sampling a completion, which can be done from the token-level model.

```

43 def conditional_token_generation( $\sigma$ ):
44      $\delta' \sim \text{Categorical}(\{\delta' : p' / \vec{p}_{\Sigma}(\sigma)\})$ 
45     for ( $p', \_, \delta'$ ) in enum_cover( $\sigma$ ):
46         return sample_completion( $\delta'$ )
47 def sample_completion( $\delta'$ ):
48      $\delta'' \leftarrow \epsilon$ 
49     while True:
50          $\delta \sim \vec{p}_{\Delta}(\cdot \mid \delta' \cdot \delta'')$ 
51         if  $\delta = \text{EOS}$ : break
52          $\delta'' \leftarrow \delta'' \cdot \delta$ 
53     return  $\delta' \cdot \delta''$ 
    
```



(a) Error (JSD/byte) vs. speed (bytes/sec). Error is computed between the character-level conditional distributions with beam sizes $K \in \{2, 4, 8, 16, 32, 64\}$ and a reference distribution computed using a much larger value of $K = 128$.



(b) Surprisal (bits/byte) vs. speed (bytes/sec). The dotted line shows the canonically tokenized baseline’s surprisal.

Figure 1: Experimental results. Averages are computed over the first 4000 bytes of the `wikitext-103-v1` test set. Error bars denote bootstrapped 95% confidence intervals. See App. A for the full table of results.

The following proposition establishes correctness:

Proposition 2. $\text{conditional_token_generation}(\sigma)$ generates samples according to $p_{\Delta|\Sigma}(\cdot | \sigma)$ for all $\sigma \in \Sigma^*$.

Proof. See App. E. ■

We also note the following corollary, as it gives an interpretation for the categorical distribution in the efficient $\text{conditional_token_generation}$ algorithm.

Corollary 1. For all $\sigma \in \Sigma^*$, $\delta \in \Delta^*$,

$$\mathbb{P}_{Y \sim p_{\Delta}}[\phi_{\sigma}(Y) = \delta | \kappa(Y) \succeq \sigma] = \frac{\overrightarrow{p_{\Delta}}(\delta)}{\overrightarrow{p_{\Sigma}}(\sigma)} \mathbb{1}\{\delta \in \mathcal{C}(\sigma)\} \quad (15)$$

Thus, we have provided an efficient solution to the prompt boundary problem. We also note that generating from p_{Σ} a character at a time is also a correct solution to the prompt boundary problem; however, it is slower, as it does not benefit from the fact that the generated string is shorter in token space. This is because once the minimally covering token string has been sampled, the method `sample_completion` will generate a complete sequence more efficiently than the character-at-a-time sample algorithm.

4. Experiments

This section investigates our algorithm’s running time and accuracy. We use the following setup:

- We use the following publicly available models: **Llama-3.2-1B**, **Meta-Llama-3.1-8B**, **DeepSeek-R1-Distill-Llama-8B**, and **phi-4 (14B)** from the 🤗 transformers library (Wolf et al., 2020). Each model was trained over

token strings created from byte-pair encoding (BPE; Sennrich et al. (2016); Gage (1994)).¹⁵

- We use the `wikitext-103-v1` corpus as a source of character strings; we used the version in the 🤗 datasets library. Specifically, we use the *test* portion.
- We use the **GenLM** library¹⁶ with the 🦙LLM (Kwon et al., 2023) backend to perform the efficient, batched evaluation of transformer language models on GPUs. We batch-evaluate all sequence extensions. Experiments were run on an L40S GPU with 40GB of memory.
- Our implementation utilizes a trie to efficiently represent all items in each bucket (see App. D), and the bucket-based pruning heuristic (§3).

To better understand the quality of the approximation our method provides, we perform the following experiments:¹⁷

- We measure the approximation error as the average **Jensen–Shannon distance (JSD)** to a reference model’s conditional distribution over the next byte (Fig. 1a). We use a large beam $K = 128$ as a reference model.
- We evaluate the average **surprisal** ($-\log_2$ probability) of our model’s estimated conditional distribution over the next byte in the corpus. As a baseline, we use the average surprisal (bits/bytes) of the canonical tokenization under the token-level language model (Fig. 1b).

¹⁵Note that in this section we use *bytes* as our set of characters to be compatible with the byte-pair encoding algorithm.

¹⁶<https://github.com/genlm/genlm-backend>

¹⁷Some additional baselines and experimental results are included in appendices (F and G).

Discussion. As expected, we observe that the speed (bytes/sec) decreases as K increases. We observe an inverse relationship between error (JSD/byte) and speed (bytes/sec): as the processing speed (bytes/sec) decreases, the error also decreases. Notably, this tradeoff is non-linear, with error increasing more sharply at higher processing speeds compared to lower speeds. This trend is evident in all models; in general, error appears to flatten out after $K \geq 8$. This indicates diminishing returns in reducing error as K gets large. We hypothesize that this occurs because the language model’s probability mass is concentrated on a limited set of tokenizations, which are adequately covered even with smaller beam sizes. We also observe (as expected) that larger models run more slowly; however, this appears to be due to their higher evaluation time rather than the need for larger covers.

We also observe that the average per-byte surprisal is significantly lower under all models than the canonically tokenized baseline. The reasons for this are twofold: (1) Each model assigns non-negligible probability to noncanonical tokenizations of the corpus, which are being thrown out in the baseline estimate, but that is accounted for in our estimate. (2) The most likely tokenization of the corpus is often noncanonical, which our method is better able to find, as our beam-summing method uses the probabilities assigned to tokenizations. In contrast, the baseline uses only the hard-coded canonical tokenization. Interestingly, increasing K does not appear to significantly decrease the surprisal, which we suspect is because the relatively greedy ($K = 2$) tokenizations adequately cover it.¹⁸

Conclusion

We have developed an effective method for ameliorating tensions between tokens and characters faced by engineers and users. We gave theory and algorithms that provide a character-level interface to tokenized language models. We characterized and resolved the prompt boundary problem. We investigated the empirical speed and error rates of our method on two modern language models. The primary limitation of our beam summing method is that it requires a very large beam size K if the language model does not favor a small number of tokenizations. The models that we explored in our experiments concentrate mass on a few tokenizations; thus, we did not require large K to estimate their character-level prefix probabilities accurately.

Acknowledgments

The authors would like to thank Andreas Opedal, Alex Lew, Ben Lipkin, Jacob Hoover Vigly, Luca Malagutti, Manuel de

¹⁸We found that $K = 1$ occasionally resulted in an empty beam, so we do not include it in this experiment.

Prada Corral, Vésteinn Snæbjarnarson, Samuel Kiegeand, and Yahya Emara for their helpful feedback and discussions. JT would like to thank Rycolab for its hospitality during a recent visit. The authors JLG and JT would like to thank Institut des Hautes Études Scientifiques (IHES) for their hospitality while revising this paper. MG was supported by an ETH Zürich Postdoctoral Fellowship. This research was enabled in part by compute resources provided by Mila (mila.quebec).

Impact Statement

The primary impact of this work is a more predictable, user-friendly interface for working with tokenized language models. By ensuring tools behave as expected, we mitigate certain unintended behaviors, such as the prompt-boundary problem highlighted in the introduction. More broadly, we do not anticipate any additional risks beyond those already inherent in the use of tokenized language models.

Limitations

Our method should work well with other tokenized language models, provided they were trained with a *deterministic* tokenizer. However, if they were trained with a *stochastic* tokenizer, such as UnigramLM (Kudo, 2018) or BPE-Dropout (Provilkov et al., 2020), we would expect that the probability mass over tokenizations in each covering would not be heavily concentrated on a small subset. Thus, these models may require a large beam size, making our approach expensive. Sampling-based methods may provide a better speed–accuracy tradeoff for this setting.

References

- Cao, K. and Rimell, L. You should evaluate your language model on marginal likelihood over tokenisations. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2021. URL <https://aclanthology.org/2021.emnlp-main.161>.
- Chirkova, N., Kruszewski, G., Rozen, J., and Dymetman, M. Should you marginalize over possible tokenizations? In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2023. URL <https://aclanthology.org/2023.acl-short.1>.
- Gage, P. A new algorithm for data compression. *C Users Journal*, 12(2), 1994. ISSN 0898-9788. URL https://web.archive.org/web/20230319172720/https://www.derczynski.com/papers/archive/BPE_Gage.pdf.
- Gastaldi, J. L., Terilla, J., Malagutti, L., DuSell, B., Vieira,

- T., and Cotterell, R. The foundations of tokenization: Statistical and computational concerns. In *Proceedings of the International Conference on Learning Representations*, 2025. URL <https://arxiv.org/abs/2407.11606>.
- Geng, S., Josifoski, M., Peyrard, M., and West, R. Grammar-constrained decoding for structured NLP tasks without finetuning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2023. URL <https://aclanthology.org/2023.emnlp-main.674.pdf>.
- Giulianelli, M., Malagutti, L., Gastaldi, J. L., DuSell, B., Vieira, T., and Cotterell, R. On the proper treatment of tokenization in psycholinguistics. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2024. URL <https://aclanthology.org/2024.emnlp-main.1032/>.
- Hale, J. A probabilistic Earley parser as a psycholinguistic model. In *Meeting of the North American Chapter of the Association for Computational Linguistics*, 2001. URL <https://aclanthology.org/N01-1021.pdf>.
- Koo, T., Liu, F., and He, L. Automata-based constraints for language model decoding. In *Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=BDBdb1myZY>.
- Kudo, T. Subword regularization: Improving neural network translation models with multiple subword candidates. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018. URL <https://aclanthology.org/P18-1007>.
- Kwon, W., Li, Z., Zhuang, S., Sheng, Y., Zheng, L., Yu, C. H., Gonzalez, J. E., Zhang, H., and Stoica, I. Efficient memory management for large language model serving with PagedAttention. In *Proceedings of the ACM SIGOPS Symposium on Operating Systems Principles*, 2023. URL <https://arxiv.org/abs/2309.06180>.
- Levy, R. Expectation-based syntactic comprehension. *Cognition*, 106(3), 2008. ISSN 0010-0277. URL <https://www.sciencedirect.com/science/article/pii/S0010027707001436>.
- Loula, J., LeBrun, B., Du, L., Lipkin, B., Pasti, C., Grand, G., Liu, T., Emara, Y., Freedman, M., Eisner, J., Cotterell, R., Mansinghka, V., Lew, A. K., Vieira, T., and O'Donnell, T. J. Syntactic and semantic control of large language models via sequential Monte Carlo. In *Proceedings of the International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=xoXn62FzD0>.
- Lundberg, S. and Ribeiro, M. T. The art of prompt design: Prompt boundaries and token healing. *Medium*, 2023. URL <https://towardsdatascience.com/the-art-of-prompt-design-prompt-boundaries-and-token-healing-3b2448b0be38>.
- Microsoft. Guidance. <https://github.com/microsoft/guidance>, 2023.
- Mikolov, T., Karafiát, M., Burget, L., Černocký, J., and Khudanpur, S. Recurrent neural network based language model. In *Proceedings of INTERSPEECH*, 2010. doi: 10.21437/Interspeech.2010-343. URL https://www.isca-archive.org/interspeech_2010/mikolov10_interspeech.html.
- Oh, B.-D. and Schuler, W. Leading whitespaces of language models' subword vocabulary poses a confound for calculating word probabilities, 2024. URL <https://arxiv.org/abs/2406.10851>.
- Phan, B., Havasi, M., Muckley, M. J., and Ullrich, K. Understanding and mitigating tokenization bias in language models. In *ICML Workshop on Theoretical Foundations of Foundation Models*, 2024. URL <https://openreview.net/forum?id=QqfdrBj1y1>.
- Pimentel, T. and Meister, C. How to compute the probability of a word. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2024. URL <https://aclanthology.org/2024.emnlp-main.1020/>.
- Poesia, G., Polozov, A., Le, V., Tiwari, A., Soares, G., Meek, C., and Gulwani, S. Synchromesh: Reliable code generation from pre-trained language models. In *Proceedings of the International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=KmtVD97J43e>.
- Provilkov, I., Emelianenko, D., and Voita, E. BPE-Dropout: Simple and effective subword regularization. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2020. URL <https://aclanthology.org/2020.acl-main.170/>.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8), 2019. URL <https://d4mucfpksyw.cloudfront.net/better-language-models/language-models.pdf>.
- Scholak, T., Schucher, N., and Bahdanau, D. PICARD: parsing incrementally for constrained auto-regressive decoding from language models. In *Proceedings of the Conference on Empirical Methods in Natural*, 2021. URL <https://doi.org/10.18653/v1/2021.emnlp-main.779>.

- Sennrich, R., Haddow, B., and Birch, A. Neural machine translation of rare words with subword units. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2016. URL <https://aclanthology.org/P16-1162>.
- Shannon, C. E. A mathematical theory of communication. *The Bell System Technical Journal*, 27(4), 1948. URL <https://doi.org/10.1002/j.1538-7305.1948.tb00917.x>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017. URL <https://arxiv.org/abs/1706.03762>.
- Willard, B. T. and Louf, R. Efficient guided generation for large language models. *CoRR*, abs/2307.09702, 2023. URL <https://doi.org/10.48550/arXiv.2307.09702>.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. Transformers: State-of-the-art natural language processing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2020. URL <https://aclanthology.org/2020.emnlp-demos.6>.

A. Supplementary Results

This section presents the data for Fig. 1 in tabular form.

A.1. Jensen–Shannon Distance

	Llama-3.2-1B		Meta-Llama-3.1-8B	
K	JSD / byte	bytes / sec	JSD / byte	bytes / sec
2	0.00025 (0.00010, 0.00048)	115.18 (110.54, 117.94)	0.00016 (0.00010, 0.00023)	36.03 (35.90, 36.16)
4	0.00006 (0.00005, 0.00007)	91.22 (88.36, 92.97)	0.00005 (0.00004, 0.00005)	32.06 (31.74, 32.27)
8	0.00004 (0.00004, 0.00005)	64.71 (63.50, 65.55)	0.00004 (0.00004, 0.00004)	26.72 (26.43, 26.95)
16	0.00004 (0.00003, 0.00004)	38.77 (37.77, 39.62)	0.00003 (0.00003, 0.00004)	19.03 (18.76, 19.28)
32	0.00004 (0.00004, 0.00004)	20.53 (19.93, 21.11)	0.00003 (0.00002, 0.00003)	11.42 (11.21, 11.62)
64	0.00003 (0.00003, 0.00004)	9.77 (9.52, 10.02)	0.00002 (0.00002, 0.00002)	5.91 (5.81, 6.02)
128	(not applicable)	4.91 (4.81, 5.02)	(not applicable)	2.97 (2.93, 3.02)

	DeepSeek-R1-Distill-Llama-8B		phi-4 (14B)	
K	JSD / byte	bytes / sec	JSD / byte	bytes / sec
2	0.00022 (0.00015, 0.00032)	36.64 (36.47, 36.83)	0.00065 (0.00034, 0.00116)	21.17 (21.03, 21.30)
4	0.00009 (0.00007, 0.00014)	32.60 (32.24, 32.87)	0.00012 (0.00008, 0.00017)	19.09 (18.98, 19.18)
8	0.00005 (0.00004, 0.00006)	26.81 (26.52, 27.04)	0.00004 (0.00003, 0.00006)	16.79 (16.69, 16.88)
16	0.00004 (0.00003, 0.00005)	19.26 (18.99, 19.52)	0.00008 (0.00004, 0.00012)	12.81 (12.68, 12.94)
32	0.00005 (0.00003, 0.00007)	11.62 (11.42, 11.83)	0.00003 (0.00003, 0.00004)	7.90 (7.78, 8.01)
64	0.00002 (0.00002, 0.00002)	5.96 (5.86, 6.07)	0.00002 (0.00002, 0.00002)	4.16 (4.10, 4.22)
128	(not applicable)	3.01 (2.97, 3.06)	(not applicable)	2.44 (2.40, 2.48)

A.2. Surprisal

Model	Canonically Tokenized Baseline bits / byte
Llama-3.2-1B	0.76 (0.69, 0.83)
Meta-Llama-3.1-8B	0.66 (0.59, 0.74)
DeepSeek-R1-Distill-Llama-8B	0.93 (0.83, 1.03)
phi-4 (14B)	0.68 (0.61, 0.76)

K	Llama-3.2-1B		Meta-Llama-3.1-8B	
	bits / byte	bytes / sec	bits / byte	bytes / sec
2	0.63539 (0.58988, 0.68280)	115.18 (110.54, 117.94)	0.53339 (0.49164, 0.57577)	36.03 (35.90, 36.16)
4	0.63593 (0.59076, 0.68238)	91.22 (88.36, 92.97)	0.53173 (0.49087, 0.57468)	32.06 (31.74, 32.27)
8	0.63540 (0.58932, 0.68020)	64.71 (63.50, 65.55)	0.53202 (0.49109, 0.57463)	26.72 (26.43, 26.95)
16	0.63518 (0.58985, 0.68186)	38.77 (37.77, 39.62)	0.53208 (0.49058, 0.57452)	19.03 (18.76, 19.28)
32	0.63524 (0.58988, 0.68187)	20.53 (19.93, 21.11)	0.53176 (0.49012, 0.57411)	11.42 (11.21, 11.62)
64	0.63525 (0.59041, 0.68259)	9.77 (9.52, 10.02)	0.53180 (0.48991, 0.57464)	5.91 (5.81, 6.02)
128	0.63550 (0.59101, 0.68157)	4.91 (4.81, 5.02)	0.53175 (0.49091, 0.57396)	2.97 (2.93, 3.02)

K	DeepSeek-R1-Distill-Llama-8B		phi-4 (14B)	
	bits / byte	bytes / sec	bits / byte	bytes / sec
2	0.76195 (0.70150, 0.82342)	36.64 (36.47, 36.83)	0.56296 (0.51968, 0.60938)	21.17 (21.03, 21.30)
4	0.76134 (0.70160, 0.82244)	32.60 (32.24, 32.87)	0.56345 (0.51961, 0.60944)	19.09 (18.98, 19.18)
8	0.76136 (0.70270, 0.82294)	26.81 (26.52, 27.04)	0.56194 (0.51601, 0.60854)	16.79 (16.69, 16.88)
16	0.76101 (0.70137, 0.82419)	19.26 (18.99, 19.52)	0.56228 (0.51854, 0.60802)	12.81 (12.68, 12.94)
32	0.76080 (0.70129, 0.82283)	11.62 (11.42, 11.83)	0.56192 (0.51807, 0.60742)	7.90 (7.78, 8.01)
64	0.76128 (0.70313, 0.82384)	5.96 (5.86, 6.07)	0.56206 (0.51752, 0.60760)	4.16 (4.10, 4.22)
128	0.76093 (0.70029, 0.82178)	3.01 (2.97, 3.06)	0.56185 (0.51753, 0.60815)	2.44 (2.40, 2.48)

B. Proof of Proposition 1

Proposition 1. Suppose $(\Sigma, \Delta, \tau, \kappa)$ is a tokenization model where κ is strict-prefix monotone and p_Δ is a token-level language model. Then, the prefix probability $\vec{p}_\Sigma(\sigma)$ for the character-level model Eq. (6) is given by

$$\vec{p}_\Sigma(\sigma) = \sum_{\delta \in \mathcal{C}(\sigma)} \vec{p}_\Delta(\delta), \quad \forall \sigma \in \Sigma^* \quad (13)$$

Proof. We prove the proposition through the following manipulations.

$$\vec{p}_\Sigma(\sigma) = \sum_{\delta' \in \Delta^*} \mathbb{1}\{\sigma \preceq \kappa(\delta')\} p_\Delta(\delta') \quad (16)$$

$$= \mathbb{1}\{\sigma = \varepsilon\} p_\Delta(\varepsilon) + \sum_{\delta' \in \Delta^+} \mathbb{1}\{\sigma \preceq \kappa(\delta')\} p_\Delta(\delta') \quad (17)$$

$$= \mathbb{1}\{\sigma = \varepsilon\} p_\Delta(\varepsilon) + \sum_{\delta \cdot \delta' \cdot \delta'' \in \Delta^+} \mathbb{1}\{\kappa(\delta) \prec \sigma \preceq \kappa(\delta \cdot \delta' \cdot \delta'')\} p_\Delta(\delta \cdot \delta' \cdot \delta'') \quad (18)$$

$$= \mathbb{1}\{\sigma = \varepsilon\} p_\Delta(\varepsilon) + \sum_{\delta \cdot \delta' \in \Delta^+} \mathbb{1}\{\kappa(\delta) \prec \sigma \preceq \kappa(\delta \cdot \delta')\} \sum_{\delta'' \in \Delta^*} p_\Delta(\delta \cdot \delta' \cdot \delta'') \quad (19)$$

$$= \mathbb{1}\{\sigma = \varepsilon\} p_\Delta(\varepsilon) + \sum_{\delta \cdot \delta' \in \Delta^+} \mathbb{1}\{\kappa(\delta) \prec \sigma \preceq \kappa(\delta \cdot \delta')\} \vec{p}_\Delta(\delta \cdot \delta') \quad (20)$$

$$= \sum_{\delta \in \mathcal{C}(\sigma)} \vec{p}_\Delta(\delta) \quad (21)$$

About the steps above: We start with the summation expression for the character-level prefix probability (i.e., Eq. (11)). We expand the summation into two cases (so that it will eventually match the two cases in the expression for the covering Eq. (12)). Next, for each summand, we consider its *unique* minimal prefix $\delta \cdot \delta'$ covering σ . We see why ε is handled separately, as it cannot be covered by a token sequence of that form. We exploit the key property of prefix monotone tokenizers (i.e., that once $\delta \cdot \delta'$ covers σ , each extension $\delta \cdot \delta' \cdot \delta''$ continues to cover it). This allows us to rearrange the summation to sum over the extension δ'' , which is conveniently equal to the prefix probability of $\delta \cdot \delta'$. The final step is to recognize that the summands can all be indexed by the covering $\mathcal{C}(\sigma)$. ■

C. The Size of the Covering

We now set about to bound the worst-case size of the covering function. To do so, we introduce additional definitions that characterize the different growth factors.

We define κ 's **fertility** as

$$F \stackrel{\text{def}}{=} \max_{\delta \in \Delta^*} \max_{\sigma \in \Sigma^*} |\{\delta' \in \Delta : \sigma = \kappa(\delta \cdot \delta')\}| \leq |\Delta| \quad (22)$$

Example 3. The BPE tokenizer has $F_{\text{BPE}} = 1$ because it is multiplicative, and its tokens represent distinct substrings. More formally,

$$F_{\text{BPE}} = \max_{\delta \in \Delta^*} \max_{\sigma \in \Sigma^*} |\{\delta' \in \Delta : \sigma = \kappa(\delta \cdot \delta')\}| \quad [\text{def of fertility}] \quad (23)$$

$$= \max_{\delta \in \Delta^*} \max_{\sigma \in \Sigma^*} |\{\delta' \in \Delta : \sigma = \kappa(\delta) \cdot \kappa(\delta')\}| \quad [\text{def multiplicativity}] \quad (24)$$

$$= |\{\delta' \in \Delta : \sigma' = \kappa(\delta')\}| \quad [\text{def function}] \quad (25)$$

$$= 1 \quad [\text{distinctness}] \quad (26)$$

Additionally, we define a κ 's **munch** as follows.

$$M \stackrel{\text{def}}{=} \max_{\delta \in \Delta^*} \max_{\delta \in \Delta} |\kappa(\delta \cdot \delta)| - |\kappa(\delta)| \quad (27)$$

In words, the munch measures the length of the largest number of characters that can be introduced by adding one more token to any given context.

Example 4. The munch of a multiplicative κ , such as BPE, is $\max_{\delta \in \Delta} |\kappa(\delta)|$. Put in words, it is the length of the longest detokenization. The munch for GPT-2 is surprisingly long (128).

Proposition 3. Let F and M be the fertility and munch of κ . Then, for all $\sigma \in \Sigma^*$,

$$|\mathcal{C}(\sigma)| \leq C(|\sigma|) \quad (28)$$

where

$$C(n) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } n < 0 \\ 1 & \text{if } n = 0 \\ F \sum_{j=n-M}^{n-1} C(j) & \text{otherwise} \end{cases} \quad (29)$$

Proof. The base cases $N \leq 0$ are straightforward. Consider the case of a string of length $N \geq 0$. *Inductive hypothesis:* Suppose for all strings σ' with $|\sigma'| < N$, $|\mathcal{C}(\sigma')| \leq C(|\sigma'|)$.

Let σ be an arbitrary string with length $N > 0$.

$$|\mathcal{C}(\sigma_1 \cdots \sigma_N)| \tag{30}$$

$$= |\{\delta \cdot \delta \in \Delta^+ : \kappa(\delta) \prec \sigma_1 \cdots \sigma_N \preceq \kappa(\delta \cdot \delta)\}| \tag{31}$$

$$= \left| \bigcup_{j=0}^N \underbrace{\{\delta \cdot \delta \in \Delta^+ : \kappa(\delta) = \sigma_1 \cdots \sigma_j, \sigma_1 \cdots \sigma_N \preceq \kappa(\delta \cdot \delta)\}}_{=\emptyset \text{ if } N-(j+1) > M \text{ or } N=j+1} \right| \tag{32}$$

$$\leq \sum_{j=N-M}^{N-1} \left| \underbrace{\{\delta \in \Delta^* : \kappa(\delta) = \sigma_1 \cdots \sigma_j\}}_{\subseteq \mathcal{C}(\sigma_1 \cdots \sigma_j)} \right| \cdot \underbrace{|\{\delta \in \Delta : \delta \in \Delta^*, \sigma_1 \cdots \sigma_N \preceq \kappa(\delta \cdot \delta)\}|}_{\leq F} \tag{33}$$

$$\leq F \cdot \sum_{j=N-M}^N \underbrace{|\mathcal{C}(\sigma_1 \cdots \sigma_j)|}_{\text{inductive hypothesis}} \tag{34}$$

$$\leq F \cdot \sum_{j=N-M}^N C(j) \tag{35}$$

$$= C(N) \tag{36}$$

Thus, the proposition holds true by the principle of induction. \blacksquare

Corollary 2. Let $N = |\sigma|$. Consider the following cases:

- When $M = N$ and $F = 1$, $C(N) = 2^N$.
- When $M = N$ and $F \geq 0$, $C(N) = F(1 + F)^N$.
- Otherwise, $C(N) = F^N \text{Fib}(N, M)$ where $\text{Fib}(N, M)$ is the N^{th} M^{th} -order Fibonacci number.¹⁹

In all cases, $C_M^F(N) < \infty$.

In the proposition below, we show that the covering can easily be exponential in size:

Proposition 4.

$$|\mathcal{C}(\sigma)| \in \Omega(2^{|\sigma|}) \tag{37}$$

Proof. We prove the proposition by constructing an example that achieves the lower bound.

- Let $\Sigma = \{a\}$, $\Delta = \left\{ \begin{smallmatrix} a & aa \\ 1 & 2 \end{smallmatrix} \right\}$.
- Let κ be multiplicative, and define $\kappa \begin{pmatrix} a \\ 1 \end{pmatrix} \stackrel{\text{def}}{=} a$, $\kappa \begin{pmatrix} aa \\ 2 \end{pmatrix} \stackrel{\text{def}}{=} aa$.
- Let σ be an arbitrary string from Σ^* . Let $N = |\sigma|$.

Then, $|\mathcal{E}(\sigma)| \geq$ the number of nonnegative integer solutions (n, m) to $1m + 2n = N$. This is because we can build the a^N using a sequence of $\begin{smallmatrix} a \\ 1 \end{smallmatrix}$ or $\begin{smallmatrix} aa \\ 2 \end{smallmatrix}$, but each $\begin{smallmatrix} a \\ 1 \end{smallmatrix}$ accounts for 1 a and each $\begin{smallmatrix} aa \\ 2 \end{smallmatrix}$ accounts for 2. So if n is the number of token 1 and m is the number of token 2, we must have that $1m + 2n = N$. The number of solutions grows like $\Omega(2^N)$. Lastly, because $\mathcal{C}(\sigma) \supseteq \mathcal{E}(\sigma)$, we have that $|\mathcal{C}(\sigma)| \in \Omega(2^{|\sigma|})$. Thus, the proposition holds. \blacksquare

¹⁹ M^{th} -order Fibonacci numbers are a variation of the well-known Fibonacci (i.e., $M = 2$) that sums the previous M numbers in the sequence instead of the previous two.

D. Bundled Beam Summing Implementation

Upon implementing this scheme, we observed that it is possible to efficiently reason about all the next tokens that extend a given token sequence in the cover in bulk. The key idea is to group the token sequences that fall into the same bucket in `prune_top_K_buckets` into a Bundle (see below) that represents them compactly. In particular, we can use a probability trie to efficiently filter out the next tokens that disagree with the next character. This improves the per-iteration cost of that filter as the data structure does the organizational work ahead of time (in bulk). We can regard the probability trie as a local language model that generates the next token character-by-character according to the probability assigned to it by $\vec{p}_\Delta(\cdot \mid \delta)$. Each bundle can be unbundled (see method `unbundle`) into the respective tuples that the `enum_cover` algorithm maintains. The algorithms are otherwise equivalent.

```

54 def beam( $\sigma_1 \dots \sigma_N$ ):
55     if  $N = 0$ : return [Bundle(1,  $\varepsilon$ ,  $\varepsilon$ , build_trie( $\varepsilon$ ))]
56     candidates  $\leftarrow$  []
57     for bundle in beam( $\sigma_1 \dots \sigma_{N-1}$ ):
58         filtered_bundle  $\leftarrow$  bundle.filter( $\sigma_N$ )
59         if filtered_bundle is not None:
60             candidates.append(filtered_bundle)
61         for extended_bundle in bundle.extend():
62             candidates.append(extended_bundle.filter( $\sigma_N$ ))
63     # Keep top- $K$  bundles according to their prefix probability
64     return top_K(candidates, key=lambda bundle: -bundle.p)
    
```

Each bundle is an instance of the following class with four member variables: the prefix probability p , token string δ , character string σ , and a reference to a local probability trie $trie$. The trie provides the character-level probabilities of the distribution over possible next tokens: $p_\Delta(\cdot \mid \delta)$. The trie is also augmented with a special symbol EOT to denote the end of this next token.²⁰

```

65 class Bundle( $p, \delta, \sigma, trie$ ):
66
67     def filter( $\sigma'$ ):
68         if  $trie.p(\sigma' \mid \sigma) = 0$ : return # no tokens give prefix  $\sigma \cdot \sigma'$  prob
69         return Bundle( $p \cdot trie(\sigma' \mid \sigma), \delta, \sigma \cdot \sigma', trie$ )
70
71     def extend():
72          $Z \leftarrow trie.p(\text{EOT} \mid \sigma)$ 
73         if  $Z > 0$ : # emit tokens that decode to  $\sigma$ 
74             for ( $\delta', p'$ ) in  $trie.tokens[\sigma]$ :
75                 yield Bundle( $\frac{p}{Z} \cdot p', \delta \cdot \delta', \varepsilon, build\_trie(\delta \cdot \delta')$ )
    
```

The code below builds the probability trie from the possible next tokens. It figures out the character strings associated with those next tokens and puts them into the trie with the respective probabilities.²¹

```

76 def build_trie( $\delta$ ):
77      $\sigma \leftarrow \kappa(\delta)$  # remember common prefix
78      $trie \leftarrow \text{ProbabilityTrie}()$  # uses EOT to mark the end of a token
79     for  $\delta' \in \Delta$ :
80          $\sigma \cdot \sigma' \leftarrow \kappa(\delta \cdot \delta')$  # use new characters, i.e., ignore prefix  $\sigma$ 
81          $trie.add(\sigma', \delta', p_\Delta(\delta' \mid \delta))$  # add  $\sigma'$  and  $\delta'$  to  $trie$  with this probability
82     return trie
    
```

The probability trie has the following functionality:

- $trie.tokens[\sigma']$ stores the set of tokens that decode to σ' along with their respective mass.

²⁰This is completely analogous to how EOS marks the end of a string.

²¹Note that EOS is handled as an indivisible symbol in the probability trie, whereas other tokens can be expanded into characters (or bytes).

- $trie.p(\sigma' | \sigma)$ returns the probability of the character σ' given σ , it is proportional to

$$trie.p(\sigma' | \sigma) \propto \sum_{\delta' \in \Delta} p_{\Delta}(\delta' | \delta) \begin{cases} \mathbb{1}\{\kappa(\delta \cdot \delta') = \sigma\} & \text{if } \sigma' = \text{EOT} \\ \mathbb{1}\{\kappa(\delta \cdot \delta') \succeq \sigma \cdot \sigma'\} & \text{otherwise} \end{cases} \quad (38)$$

We note that the trie implicitly depends on the token string δ used in its creation.

To aid in understanding how the bundled algorithm relates to the original algorithm, we give the following methods.

```

83 class Bundle(p,  $\delta$ ,  $\sigma$ , trie):
84     ...
85     def unbundle():
86         agenda  $\leftarrow$  [ $\sigma$ ]
87         while agenda:
88              $\sigma'$   $\leftarrow$  agenda.pop()
89             Z  $\leftarrow$  trie.p(EOT |  $\sigma'$ )
90             if Z > 0:
91                 for ( $\delta'$ , p')  $\in$  trie.tokens[ $\sigma'$ ]:
92                     yield (p · p',  $\kappa(\delta \cdot \delta')$ ,  $\delta \cdot \delta'$ )
93             for  $\sigma''$  in trie.p( $\cdot$  |  $\sigma'$ ):
94                 if  $\sigma'' \neq$  EOT:
95                     agenda.append( $\sigma' \cdot \sigma''$ )
96
97     def unbundle_beam(beam):
98         return [item for bundle in beam for item in bundle.unbundle()]

```

We note that `unbundle_beam(beam(σ))` gives precisely the same set of elements as the unbundled algorithm (i.e., `enum_cover(σ)`) run on the same string σ and the bucket-based pruning scheme with parameter K (up to reordering).

To compute the next-character probability, we use the following algorithm:

```

98 def next_character_probability( $\sigma$ , method=beam):
99      $\bar{p} \leftarrow \{\sigma': 0 \text{ for } \sigma' \in \Sigma \cup \{\text{EOS}\}\}$ 
100     for bundle in method( $\sigma$ ): # use beam approximation by default
101         for  $\sigma' \in \Sigma$ :
102              $\bar{p}(\sigma') +=$  bundle.p · bundle.trie.p( $\sigma'$  | bundle. $\sigma$ )
103         for ext_bundle in bundle.extend():
104             for  $\sigma' \in \Sigma$ :
105                  $\bar{p}(\sigma') +=$  ext_bundle.p · ext_bundle.trie.p( $\sigma'$  | ext_bundle. $\sigma$ )
106     Z  $\leftarrow$  sum( $\bar{p}$ .values())
107     return  $\{\sigma': \bar{p}(\sigma')/Z \text{ for } \sigma' \in \Sigma\}$ 

```

E. Proof of Proposition 2

Proposition 2. `conditional_token_generation`(σ) generates samples according to $p_{\Delta|\Sigma}(\cdot | \sigma)$ for all $\sigma \in \Sigma^*$.

Proof (Sketch). Choose an arbitrary $\delta \in \Delta^*$.

$$p_{\Delta|\Sigma}(\delta | \sigma) = \mathbb{P}_{Y \sim p_{\Delta}} [Y = \delta | \kappa(Y) \succeq \sigma] \quad (39)$$

$$= p_{\Delta}(\delta) \frac{\mathbb{1}\{\kappa(\delta) \succeq \sigma\}}{\mathbb{P}_{Y \sim p_{\Delta}} [\kappa(Y) \succeq \sigma]} \quad (40)$$

$$= p_{\Delta}(\delta) \frac{\mathbb{1}\{\kappa(\delta) \succeq \sigma\}}{\overrightarrow{p_{\Sigma}}(\sigma)} \quad (41)$$

Let $\delta' = \phi_{\sigma}(\delta)$ (i.e., the shortest prefix of δ such that $\kappa(\delta') \succeq \sigma$). Choose δ'' such that $\delta' \cdot \delta'' = \delta$.

$$= \underbrace{\overrightarrow{p_{\Delta}}(\text{EOS} | \delta' \cdot \delta'') \overrightarrow{p_{\Delta}}(\delta'' | \delta')}_{\text{sample completion}} \underbrace{\overrightarrow{p_{\Delta}}(\delta') \frac{\mathbb{1}\{\kappa(\delta') \succeq \sigma\}}{\overrightarrow{p_{\Sigma}}(\sigma)}}_{\text{sample from covering}} \quad (42)$$

We can see that the algorithm samples from this distribution because it samples a token string δ' from the covering in proportion to the right factor in Eq. (42) and then samples a completion δ'' of δ' in proportion to the left factor of the equation. Thus, the sample $\delta = \delta' \cdot \delta''$ has probability $p_{\Delta|\Sigma}(\delta | \sigma)$, and `conditional_token_generation`(σ) is a correct sampling procedure for it. ■

F. Token Healing Baseline

We implemented an additional baseline character-level model based on a token healing heuristic described in §1. While not competitive with our proposed methods, we include it here for completeness. Explain the idea, and provide pseudocode below. The experimental results are summarized in Tab. 1. We observe significantly higher surprisal (bits/byte) and error (JSD/bytes) compared to the methods in Fig. 1.

An approximate covering inspired by token healing. Our version of the token healing heuristic is based on an approximate covering. Given character string $\sigma \in \Sigma^*$ and its tokenization $\delta_1 \dots \delta_N = \tau(\sigma)$,²² we define its **token healing approximate covering** as follows:

$$\mathcal{C}_{\text{heal}}(\sigma) \stackrel{\text{def}}{=} \{\delta_1 \dots \delta_{N-1} \cdot \delta' \mid \delta' \in \Delta, \kappa(\delta_N) \preceq \kappa(\delta')\} \quad (43)$$

By construction, $\kappa(\delta_1 \dots \delta_{N-1}) \prec \sigma \preceq \kappa(\delta_1 \dots \delta_N)$, so it is straightforward to see that $\mathcal{C}_{\text{heal}}(\sigma) \subseteq \mathcal{C}(\sigma)$. We also note that the size of this approximate covering is bounded by the size of the token alphabet: $|\mathcal{C}_{\text{heal}}(\sigma)| \leq |\Delta|$.

Implementation. The code below implements a method for efficiently constructing $\mathcal{C}_{\text{heal}}(\sigma)$. It is based on the *bundled* implementation given in App. D.

```

108 def token_healing_cover( $\sigma$ ):
109     if  $\sigma = \varepsilon$ : # no tokens to heal
110         return [Bundle(1,  $\varepsilon$ ,  $\varepsilon$ , build_trie( $\varepsilon$ ))]
111     else
112          $\delta_1 \dots \delta_N \leftarrow \tau(\sigma)$ 
113         trie  $\leftarrow$  build_trie( $\delta_1 \dots \delta_{N-1}$ )
114          $\sigma' \leftarrow \kappa(\delta_N)$ 
115         return [Bundle( $\vec{p}_{\Delta}(\delta_1 \dots \delta_{N-1}) \cdot \text{trie.p}(\sigma' \mid \varepsilon)$ ,  $\delta_1 \dots \delta_{N-1}$ ,  $\sigma'$ , trie)]
    
```

We first tokenize the input character sequence σ into a sequence of tokens δ . We then construct a probability trie for $\delta_1 \dots \delta_{N-1}$ (i.e., all but the last token), allowing us to compute the probability of each possible next character σ' by querying the trie with the character suffix σ' corresponding to the final token.

Compared to the beam algorithm, token_healing_cover method represents an extreme form of pruning: it returns a beam of exactly one bundle. It also differs from the variant of token healing we described in §1, which produces token-level predictions rather than character-level predictions.

To estimate $\vec{p}_{\Sigma}(\cdot \mid \sigma)$ we use next_character_probability(σ , token_healing_cover). However, this approximation still fails to place an appropriately high probability on `d` when conditioned on `Hello, _world`). This happens for the same reasons as we described in §1.

Experimental results. The table below presents experimental results that may be compared to those of other methods in the paper. We found that the method performs poorly enough that we do not consider it a competitive approximation.

Model	Surprisal (bits/byte)	Error (JSD/byte)	Speed (bytes/sec)
DeepSeek-R1-Distill-Llama-8B	2.529 (2.391, 2.670)	0.1907 (0.1800, 0.2015)	36.70 (36.64, 36.76)
Llama-3.2-1B	2.126 (2.010, 2.246)	0.1790 (0.1691, 0.1891)	119.35 (118.73, 119.84)
Meta-Llama-3.1-8B	2.116 (1.992, 2.243)	0.1840 (0.1738, 0.1946)	36.61 (36.55, 36.67)
phi-4 (14B)	2.204 (2.075, 2.333)	0.1857 (0.1757, 0.1960)	21.24 (21.13, 21.30)

Table 1: Surprisal, JSD, and speed for the token-healing baseline across models using the same experimental settings as Fig. 1.

²²If $\sigma = \varepsilon$, $\mathcal{C}_{\text{heal}}(\varepsilon) \stackrel{\text{def}}{=} \Delta$.

G. Probability-Based Pruning Heuristic

This section presents some additional experimental results with a more extensive pruning method, based on the relative contribution of a token string to the current cover approximation.

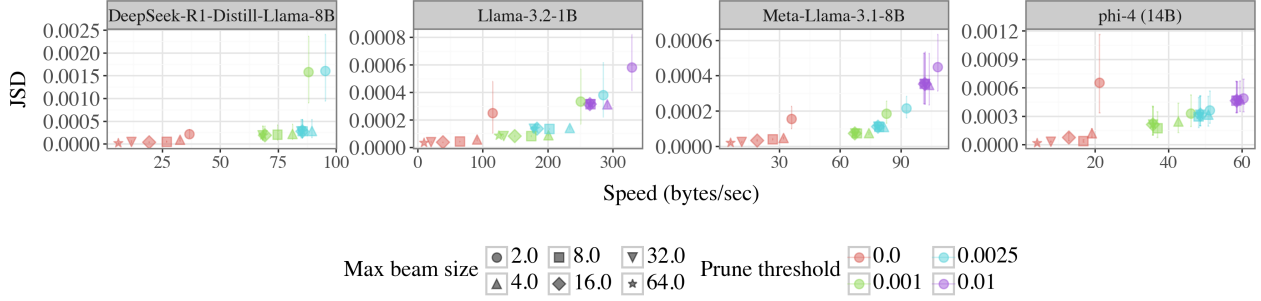
Probability-based pruning heuristic. The method works by augmenting the beam pseudocode with some additional pruning based on a threshold $\theta \in [0, 1]$. The benefit of this pruning is that it can avoid calls to `bundle.extend()`, which requires evaluating the language model (i.e., the performance bottleneck of our algorithm). It is often the case that for a given character string prefix, there is no reason to perform any **extend** operations because all tokenizations of that string are low probability; thus, filtering operations are needed.

```

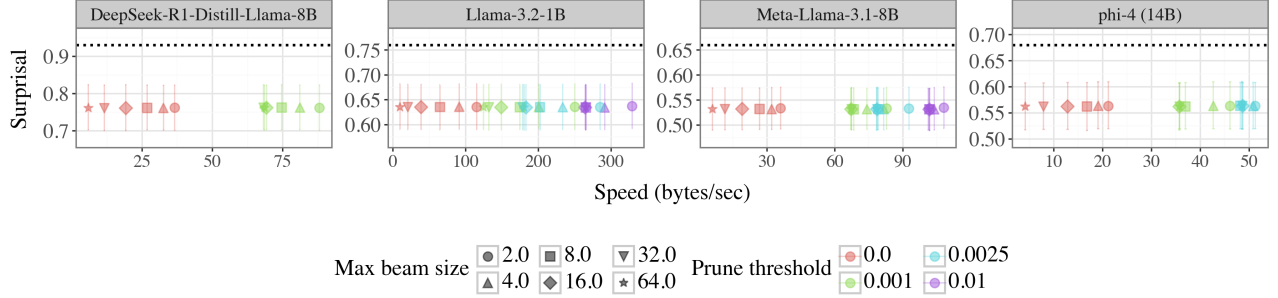
116 def beamK,θ( $\sigma_1 \dots \sigma_N$ ):
117     if  $N = 0$ : return [Bundle(1,  $\varepsilon$ ,  $\varepsilon$ , build_trie( $\varepsilon$ ))]
118     candidates ← []
119     B ← beam( $\sigma_1 \dots \sigma_{N-1}$ )
120     for bundle in B:
121         filtered_bundle ← bundle.filter( $\sigma_N$ )
122         if filtered_bundle is not None:
123             candidates.append(filtered_bundle)
124      $\tau \leftarrow \theta \cdot \text{sum}(\text{bundle.p for bundle in candidates})$ 
125     for bundle in B:
126         if bundle.p ≥  $\tau$ :
127             for extended_bundle in bundle.extend():
128                 candidates.append(extended_bundle.filter( $\sigma_N$ ))
129     # Keep top- $K$  bundles according to their prefix probability
130     return topK(candidates, key=lambda bundle: -bundle.p)

```

Experiments. Below, we extend the experiments of the main text with this additional pruning heuristic to better understand its speed and accuracy.



(a) Error (JSD/byte) vs. speed (bytes/sec). Here we compute the JSD between the byte-level conditional distributions computed with $\text{beam}_{K,\theta}$ for $(K, \theta) \in \{4, 8, 16, 32, 64\} \times \{0, 0.001, 0.0025, 0.01\}$ and the reference model ($K = 128, \theta = 0$). Missing values indicate “dead-ending” (i.e., empty beam) as a result of overly aggressive pruning.



(b) Average surprisal (bits/byte) vs. speed (bytes/sec) for $\text{beam}_{K,\theta}$ with $\{4, 8, 16, 32, 64\} \times \{0, 0.001, 0.0025, 0.01\}$. The dotted line represents surprisal of the canonically tokenized under the token-level language model. Missing values indicate infinite surprisal as a result of excessive pruning.

Figure 2: Additional experimental results showing the effects of the pruning thresholds on error vs. speed and surprisal vs. speed using the same settings as in §4.