

Which Pieces Does Unigram Tokenization Really Need?

Anonymous ACL submission

Abstract

The Unigram tokenization algorithm offers a probabilistic alternative to the greedy heuristics of Byte-Pair Encoding. Despite its theoretical elegance, its implementation in practice is complex, limiting its adoption to the SentencePiece package and adapters thereof. We bridge this gap between theory and practice by providing a clear guide to implementation and parameter choices. We also identify a simpler algorithm that accepts slightly higher training loss in exchange for improved compression.

 github.com/code/after/review

1 Introduction

The Unigram language model for tokenization (Kudo, 2018) was introduced as a principled, probabilistic alternative to the greedy heuristics of Byte-Pair Encoding (BPE; Sennrich et al., 2016). Despite its theoretical elegance, its implementation in practice is daunting: the original paper presents the training algorithm at a very high level of abstraction, online tutorials are often oversimplified to the point of being incorrect, and the only authoritative reference is a complex C++ implementation in SentencePiece (Kudo and Richardson, 2018). This has practical consequences: researchers often default to BPE or treat SentencePiece as a black box, limiting experimentation and exploration of tokenizer algorithms beyond small variations in BPE (Provilkov et al., 2020; Chizhov et al., 2024, inter alia).

This paper explains the reference Unigram training algorithm, suggests alternatives to key steps, and explores the effect of configurable and hard-coded parameters on tokenizer performance. We find that costly defaults such as multiple expectation-maximization steps have little effect on the output vocabulary. We also present a better heuristic for constructing seed vocabularies and a simplified pruning algorithm that accepts

slightly higher loss while compressing almost as well as BPE.

2 The Unigram Model in Theory

The Unigram model¹ assumes that the probability $P(\mathbf{x})$ of a tokenized sequence $\mathbf{x} = (x_1, \dots, x_k)$ is the product of the token probabilities $p(x_i)$:

$$P(\mathbf{x}) = \prod_{i=1}^k p(x_i). \quad (1)$$

When tokenizing text with a trained model, the Viterbi algorithm is typically used to determine the sequence of tokens which maximizes $P(\mathbf{x})$. During LLM training, this can be extended to sample lower-probability token sequences for regularization. In tokenizer training, both the vocabulary \mathcal{V} and the token probabilities $p(x)$ for all $x \in \mathcal{V}$ are jointly optimized to maximize the marginal log-likelihood over a training corpus C . For a given text X , let $\mathcal{S}(X)$ be the set of all possible segmentations using the vocabulary. The objective function \mathcal{L} to minimize is:

$$\mathcal{L} = \frac{-1}{\sum_{X \in C} |X|} \sum_{X \in C} \log \left(\sum_{\mathbf{x} \in \mathcal{S}(X)} P(\mathbf{x}) \right), \quad (2)$$

where $|X|$ is the length of X measured in atomic tokens (e.g. bytes). We normalize by the number of atomic tokens rather than the number of texts to ensure the metric remains independent of how the corpus is subdivided. For a more detailed mathematical treatment, see Meister (2025).

3 The Unigram Model in Practice

Since the vocabulary \mathcal{V} is not known in advance, the training process must simultaneously discover a good vocabulary and estimate token probabilities.

¹Sometimes referred to as *UnigramLM*. We use the shorter term as it is unambiguous in this context.

The algorithm described by Kudo (2018) is an iterative procedure that starts with a large number of seed tokens and gradually reduces them to the target vocabulary size. Both the original description and our summary of the reference implementation are listed in Algorithm 1.²

3.1 Seed Vocabulary Initialization

The SentencePiece implementation uses a combination of a suffix array and a longest prefix algorithm to extract frequent substrings. In short, the suffix array-based approach encodes a compact representation of the corpus’ sorted suffixes, which then allows for efficient frequency counting with a stack-based emission algorithm. Algorithm 2 describes these steps in more detail.

Unlike many other tokenizers, SentencePiece does not pretokenize (i.e., split text into *pretokens*) before training. Instead, it operates on the entire corpus, rejecting candidate tokens during initialization based on validity criteria such as not containing internal whitespace. As a result, the stack-based emission algorithm in the reference implementation (cf. Algorithm 2, lines 12–18) may exclude certain frequent tokens from the set of seed tokens. This happens when the shared prefix is not a valid token, often due to trailing spaces. For example, given the corpus “the_old_man_the_boat”, the substring “the_” will be emitted, and then rejected due to its trailing space, while “the” is never emitted as it is not a maximal prefix. We propose and test a ‘Maximal Valid Prefix Recovery’ variation of the algorithm where the longest valid prefixes of rejected seed tokens are included (see Algorithm 3). Note that in both variants, a single corpus occurrence of a word can introduce new high-frequency tokens in the seed vocabulary, by making a prefix ‘maximal’ in at least one case.

3.2 The EM Algorithm

Given a fixed vocabulary, token probabilities $p(x)$ are estimated using the expectation-maximization algorithm, which alternates between expectation and maximization steps. In the Expectation step, the forward-backward algorithm (Baum, 1972) is used to efficiently determine the expected count of each token, averaged over all segmentations of the corpus under the current probabilities $p(x_i)$. The Maximization step updates the token probabilities

²We defer detailed expansions such as pseudocode algorithms to the appendix, but the descriptions in the main text should suffice for understanding the experiments.

$p(x_i)$ based on these expected counts. In SentencePiece, at the start of the M-step, tokens with an expected count below $\tau_{\text{mp}} = 0.5$ are removed, which we refer to as the ‘early prune step’. This hyperparameter is sensitive to the absolute corpus size, i.e. duplicating a corpus would change thresholding behavior despite identical token distributions.

3.3 Iterative Vocabulary Pruning

After a fixed number N_{em} of EM steps,³ the main pruning step is performed. For each token, the pruning algorithm determines whether the token represents its own optimal tokenization and removes it otherwise. For the remaining tokens, the algorithm determines each token’s second-best tokenization and calculates the effect on total log-likelihood if all occurrences of the token were replaced with this second-best tokenization. Tokens that are least critical to maintain high overall likelihood are discarded. During this step, the probabilities $p(x_i)$ are only used to inform the optimal tokenization paths.⁴ The fraction of vocabulary preserved at each step is controlled by a hyperparameter α_{prune} (typically 0.75–0.9), ensuring that pruning happens gradually over several iterations.

In the reference implementation, the second-best tokenization is determined with a generic algorithm for determining the top k tokenizations. We note that a simple alternative is to apply the Viterbi algorithm over all paths over the token text except for the singleton.

3.4 Final Pruning

Once the iterative EM and pruning process reduces the vocabulary to at most α_{inter} times the target size (typically 1.1, i.e. 10% over), the SentencePiece implementation performs a final vocabulary selection step. At this stage, tokens with the lowest estimated probabilities $p(x_i)$ are pruned regardless of token interactions.

4 Evaluating Unigram Tokenizers

In this section, we evaluate the effect of various algorithm and parameter choices on Unigram tokenizer performance.

Tokenizers are generally difficult to evaluate, as intrinsic metrics such as compression rate do not

³The SentencePiece implementation uses $N_{\text{em}} = 2$.

⁴This is a local heuristic: it cannot account for cases where, for example, `thou` is removed from the vocabulary and now a segmentation like `court|house` becomes optimal.

Method	Loss	#Tokens	Overlap
<i>Mean over Monolingual 30 MB Corpora</i>			
Pretokens (Ours)	1.325	5.43	—
Pretok., Recovery	1.326 ^{+0.02}	5.43 ^{↓0.01}	90.2
Full-text (SP)	1.333 ^{+0.54}	5.47 ^{+0.75}	79.1
Full-text, Recov.	1.326 ^{+0.03}	5.43 ^{↓0.01}	88.6
<i>Mean over Monolingual 300 MB Corpora</i>			
Pretokens (Ours)	1.337	54.02	—
Pretok., Recovery	1.338 ^{+0.02}	54.02 ^{+0.00}	91.6
Full-text (SP)	1.338 ^{+0.05}	54.05 ^{+0.05}	95.7
Full-text, Recov.	1.338 ^{+0.03}	54.02 ^{↓0.01}	91.6

Table 1: Seed vocabulary initialization algorithms. Comparing pretoken-based (Ours) and full-text (SentencePiece-style) initialization, with and without Maximal Valid Prefix Recovery (Algorithm 3). Superscripts show % change vs baseline; ‘Overlap’ is vocabulary overlap with baseline. Pretoken-based initialization achieves better loss and compression on both corpus sizes. The recovery variation has minimal effect, suggesting that missing prefixes are recovered naturally in sufficiently large corpora.

161 always correlate well with downstream model per-
162 formance (Schmidt et al., 2024), but are neverthe-
163 less important for inference speed. We evaluate
164 Unigram tokenizer training algorithms based on
165 (1) the training objective (marginal log-likelihood,
166 Equation 2) and (2) compression as measured by
167 corpus token count. For both, lower is better.

168 4.1 Experimental Setup

169 We test the effect of training hyperparameters us-
170 ing 300 MB monolingual corpora in English, Ger-
171 man, Korean, Chinese, Arabic, and Hindi. We use
172 SCRIPT encoding and pretokenization with char-
173 acter boundaries (Land and Arnett, 2025) and a
174 trained vocabulary size of 32,768 for all experi-
175 ments.

176 4.2 Seed Vocabulary Initialization

177 We implement a pretoken-based initialization algo-
178 rithm using a pretoken-to-count mapping, follow-
179 ing the common approach in efficient BPE imple-
180 mentations for storing corpora. This approach uses
181 suffix arrays to track both text and indices, largely
182 sidestepping the issue of incorrectly-rejected pre-
183 fixes (§3.1) while improving efficiency. We also
184 test SentencePiece-style initialization with the Max-
185 imal Valid Prefix Recovery variant described in Al-
186 gorithm 3. Results in Table 1 show that both the
187 pretoken-count approach and the recovery varia-
188 tion improve both compression and objective op-

Method	n	Train Tok.	FW Tok.	Morph.
Baseline	16K	58.7	65.0	0.620
	32K	54.0	59.6	0.611
	64K	50.6	55.7	0.644
FSP	16K	58.2 ^{↓0.8}	64.1 ^{↓1.4}	0.497
	32K	53.7 ^{↓0.6}	58.9 ^{↓1.2}	0.514
	64K	50.3 ^{↓0.5}	55.1 ^{↓1.1}	0.559
BPE	16K	58.5 ^{↓0.4}	64.0 ^{↓1.5}	0.221
	32K	53.8 ^{↓0.3}	58.8 ^{↓1.3}	0.186
	64K	50.4 ^{↓0.4}	54.9 ^{↓1.4}	0.165

Table 2: Compression (M tokens, lower is better, super-
scripts show % change vs baseline at same n) and mor-
phological alignment across vocabulary sizes n . Base-
line is Unigram with our default settings, while FSP
uses Final-Style Pruning. *Train Tok.*: trained and eval-
uated on same 300 MB corpora. *FW Tok.*: trained on
FineWiki 1 GB, evaluated on 300 MB corpora. Note that
the gap with Baseline widens, while BPE gains a small
edge over FSP. *Morph.*: MorphScore boundary recall on
English (higher is better), where Unigram dominates.

189 timization over the SentencePiece-style approach.
190 However, the recovery variation has only a minor
191 effect, ranging from mildly beneficial to mildly
192 detrimental. We hypothesize that in larger corpora,
193 the few remaining important seed tokens that are
194 not caught by the corpus-based approach are likely
195 to appear at least once, mitigating the original is-
196 sue, while adding too many similar seed tokens
197 may crowd out other candidates. Therefore, in the
198 rest of our experiments we use the more efficient
199 pretoken-count approach without recovery.

200 The only other parameter in initialization is the
201 seed vocabulary size, which we implement as a
202 relative factor $n_{\text{seed}} = \beta_{\text{seed}} \cdot n$, where n is the
203 target vocabulary size. While our default of $\beta_{\text{seed}} =$
204 10 is sufficient, performance drops significantly at
205 lower values. This is particularly relevant given that
206 SentencePiece’s default value (10^6) corresponds to
207 $\beta_{\text{seed}} \approx 4$ for large multilingual tokenizers with
208 250k tokens.

209 4.3 Iterative Vocabulary Pruning

210 The main pruning loop involves several hyper-
211 parameters, many of which have minimal effect.
212 Varying the number of EM iterations within the
213 1–5 range shows no consistent differences. Sen-
214 tencePiece applies a digamma transformation to
215 expected counts in the M-step (cf. Algorithm 1)
216 which suppresses low-count tokens; omitting this
217 transformation yields essentially no difference.




Figure 1: Effect of training hyperparameters. The origin is our baseline settings, as defined in Appendix C. Hyperparameters with minimal impact on both metrics omitted. The parameters with larger effects reveal a fundamental tradeoff: improvements in compression come at the cost of Unigram objective. Final-Style Pruning achieves the best compression. No configuration in our search significantly improves both metrics simultaneously.

Similarly, the early prune threshold τ_{mp} (cf. subsection 3.2) can be varied between 0 and 10 with no apparent effect.

Among the more impactful parameters, a slower shrinking factor α_{prune} yields better objective values but worse compression, though both effects are small. A similar tradeoff appears when halting pruning earlier via a larger α_{inter} . Motivated by these results, we investigate Final-Style Pruning (FSP), which replaces *all* pruning steps with the simpler probability-based algorithm. Figure 1 visualizes all parameter variations, revealing that (1) parameters that significantly affect performance exhibit a clear tradeoff with no configuration improving both metrics, and (2) FSP is the only variation which achieves compression exceeding BPE.

Table 2 validates these findings across vocabulary sizes (16K–64K) and training on FineWiki (Penedo, 2025, 1 GB subsets across the same 6 languages). The compression–objective tradeoff persists across all settings. Notably, when training on different datasets, BPE appears to generalize better. We also report MorphScore (Arnett et al., 2025, higher is better) for morphological alignment: Unigram methods substantially outperform BPE, consistent with Unigram’s original motivation for subword regularization. Notably, FSP’s compression gains come at the cost of reduced morphological alignment, positioning it between standard Unigram and BPE on this metric. Appendix D shows some examples of segmentations.

5 Conclusion

We provide a detailed guide to Unigram tokenization, clarifying implementation details that bridge the gap between the original high-level description and the SentencePiece reference implementation.

Our analysis reveals that Unigram is remarkably robust to hyperparameter choices. Several computationally expensive settings, including the number of EM iterations, digamma transformation, and early pruning thresholds, have minimal impact and can be simplified to improve training efficiency. We also demonstrate that pretoken-based initialization consistently outperforms the full-text approach in both objective and compression.

We identify a fundamental tradeoff between training objective and compression: no parameter configuration improves both metrics simultaneously. Our proposed Final-Style Pruning achieves compression closer to BPE by simplifying the pruning procedure, though at the cost of a marginally higher loss. Given the uncertain relationship between intrinsic tokenization metrics and downstream performance, this simpler alternative warrants consideration when compression is prioritized. BPE appears to achieve slightly better compression in out-of-domain settings, which warrants further investigation.

Overall, Unigram compression closely matches BPE, suggesting the choice between them may depend more on the need for probabilistic segmentation than on compression. We hope this work lowers the barrier to broader research and experimentation in Unigram tokenization.

282 Limitations

283 Our evaluation focuses on intrinsic metrics (train-
284 ing objective and compression) rather than down-
285 stream language model performance, which re-
286 mains the ultimate measure of tokenizer effective-
287 ness.

288 References

289 Catherine Arnett, Marisa Hudspeth, and Brendan
290 O’Connor. 2025. [Evaluating morphological align-
291 ment of tokenizers in 70 languages](#). *Preprint*,
292 arXiv:2507.06378.

293 Leonard E Baum. 1972. An inequality and associated
294 maximization technique in statistical estimation for
295 probabilistic functions of markov processes. *Inequal-
296 ities*, 3(1):1–8.

297 Pavel Chizhov, Catherine Arnett, Elizaveta Korotkova,
298 and Ivan P. Yamshchikov. 2024. [Bpe gets picky: Effi-
299 cient vocabulary refinement during tokenizer training](#).
300 *Preprint*, arXiv:2409.04599.

301 Taku Kudo. 2018. [Subword regularization: Improv-
302 ing neural network translation models with multiple
303 subword candidates](#). In *Proceedings of the 56th An-
304 nual Meeting of the Association for Computational
305 Linguistics (Volume 1: Long Papers)*, pages 66–75,
306 Melbourne, Australia. Association for Computational
307 Linguistics.

308 Taku Kudo and John Richardson. 2018. [SentencePiece:
309 A simple and language independent subword tok-
310 enizer and detokenizer for neural text processing](#). In
311 *Proceedings of the 2018 Conference on Empirical
312 Methods in Natural Language Processing: System
313 Demonstrations*, pages 66–71, Brussels, Belgium.
314 Association for Computational Linguistics.

315 Sander Land and Catherine Arnett. 2025. [BPE stays on
316 SCRIPT: Structured encoding for robust multilingual
317 pretokenization](#). *Preprint*, arXiv:2505.24689.

318 Clara Meister. 2025. [UnigramLM: An attempt at writ-
319 ing the missing manual](#). Blog post.

320 Guilherme Penedo. 2025. [Finewiki](#). Source:
321 Wikimedia Enterprise Snapshot API
322 (<https://api.enterprise.wikimedia.com/v2/snapshots>).
323 Text licensed under CC BY-SA 4.0 with attribution
324 to Wikipedia contributors.

325 Ivan Provilkov, Dmitrii Emelianenko, and Elena Voita.
326 2020. [Bpe-dropout: Simple and effective subword
327 regularization](#). *Preprint*, arXiv:1910.13267.

328 Craig W Schmidt, Varshini Reddy, Haoran Zhang, Alec
329 Alameddine, Omri Uzan, Yuval Pinter, and Chris
330 Tanner. 2024. [Tokenization is more than compres-
331 sion](#). In *Proceedings of the 2024 Conference on
332 Empirical Methods in Natural Language Processing*,
333 pages 678–702, Miami, Florida, USA. Association
334 for Computational Linguistics.

Rico Sennrich, Barry Haddow, and Alexandra Birch.
2016. [Neural machine translation of rare words with
subword units](#). In *Proceedings of the 54th Annual
Meeting of the Association for Computational Lin-
guistics (Volume 1: Long Papers)*, pages 1715–1725,
Berlin, Germany. Association for Computational Lin-
guistics.

335
336
337
338
339
340
341

A The Unigram Tokenizer Training Algorithm

Algorithm 1 High-Level Unigram Model Training Algorithms, with **hardcoded** and **configurable** parameters highlighted.

Inputs and Parameters:

C ▷ Mandatory input: Training corpus
 n ▷ Mandatory parameter: Target vocabulary size
 $n_{\text{seed}} \leftarrow 10^6$ ▷ Initial seed vocabulary size
 $N_{\text{em}} \leftarrow 2$ ▷ Number of EM sub-iterations per pruning iteration
 $\alpha_{\text{prune}} \leftarrow 0.75$ ▷ Reduction factor for vocabulary size during pruning
 $\alpha_{\text{inter}} \leftarrow 1.1$ ▷ Pre-final vocab size as multiple of n
 $\tau_{\text{mp}} \leftarrow 0.5$ ▷ Minimum expected count threshold in M -step pruning

Output: Final vocabulary \mathcal{V} and probabilities $p(x_i)$.

Conceptual algorithm from Kudo (2018)

Initialize: Heuristically make a reasonably big seed vocabulary from the training corpus

while $|\mathcal{V}| > n$ **do** ▷ Pruning Loop
 (a) Fixing the set of vocabulary, optimize $p(x)$ with the EM algorithm.
 (b) Compute the loss for each subword x_i , which represents how likely the likelihood \mathcal{L} is reduced when the subword x_i is removed from the current vocabulary.
 (c) Sort the symbols by loss, and keep the top 80% of subwords.

SentencePiece Implementation

Initialize: $\mathcal{V} \leftarrow V_{\text{seed}}(C, n_{\text{seed}})$ using suffix arrays (Algorithm 2).

$n_{\text{inter}} \leftarrow \alpha_{\text{inter}} \cdot n$ ▷ Set an intermediate target size slightly larger than final.
while True **do** ▷ Pruning loop
for $i = 1$ **to** N_{em} **do** ▷ EM sub-iterations
E-step: Compute expected counts c_j for each token $x_j \in \mathcal{V}$ using forward-backward algorithm
Fast-prune step: Remove tokens x_j with expected count $c_j < \tau_{\text{mp}}$
M-step: Update token probabilities $p(x_j) \leftarrow F(c_j) / \sum_k F(c_k)$ ▷ digamma function F
if $|\mathcal{V}| \leq n_{\text{inter}}$ **then break**
Pruning step:
 $n_{\text{new}} \leftarrow \max(n_{\text{inter}}, \alpha_{\text{prune}} \cdot |\mathcal{V}|)$
 For each token x_j , calculate the effect on \mathcal{L} if it were replaced by its optimal sub-tokenization
 Prune \mathcal{V} to size n_{new} by removing tokens with the lowest effect on the objective \mathcal{L} if removed
Finalize: Prune vocabulary V to size n by selecting those with highest $p(x_i)$

B Seed Vocabulary Creation Algorithms

343
344
345
346
347
348
349

The algorithm below uses a suffix array to find frequent substrings efficiently. A *suffix array* sorts all suffixes of the corpus lexicographically, placing suffixes with shared prefixes adjacent to each other. The *longest common prefix* (LCP) between adjacent suffixes tells us how many characters they share. The algorithm maintains a stack to track ranges of suffixes sharing a common prefix, known as “open intervals”. When the LCP drops below a stack entry’s height, that prefix’s interval has closed, and we emit the pattern with its frequency.

Algorithm 2 SentencePiece Seed Vocabulary Creation

```

1: Input: Corpus  $C$ , maximum token length  $\ell_{\max}$ , Target seed vocabulary size  $n_{\text{seed}}$ 
2: Output: Seed vocabulary  $V_{\text{seed}}$  with initial scores
3: Initialization
4:  $C_{\text{concat}} \leftarrow$  Concatenation of  $C$ , separated by boundary markers.
5:  $\text{SA} \leftarrow \text{SuffixArray}(C_{\text{concat}})$   $\triangleright$   $\text{SA}[r]$  is the index of  $r^{\text{th}}$  suffix in lexicographical order
6:  $V_{\text{seed}} \leftarrow$  empty queue  $\triangleright$  Bounded priority queue, size  $n_{\text{seed}}$ , sorted by score
7:  $\text{stack} \leftarrow$  empty stack
8: Stack-based prefix emission loop
9: for  $r_{\text{curr}} = 1$  to  $|\text{SA}| + 1$  do  $\triangleright$  Process all adjacent suffix pairs in lexicographical order
10:    $h \leftarrow \text{LCP}(\text{SA}[r_{\text{curr}} - 1], \text{SA}[r_{\text{curr}}])$   $\triangleright$  Length of common prefix.  $h = 0$  in last iteration
11:    $r_{\text{lb}} \leftarrow r_{\text{curr}} - 1$ 
12:   while not  $\text{IsEmpty}(\text{stack})$  and  $\text{Top}(\text{stack}).\text{height} > h$  do
13:      $(\ell, r_{\text{start}}) \leftarrow \text{Pop}(\text{stack})$   $\triangleright$  pattern length  $\ell$ , interval start
14:      $\text{freq} \leftarrow r_{\text{curr}} - r_{\text{start}}$   $\triangleright$  Number of suffixes sharing prefix of length  $\geq \ell$ 
15:      $p \leftarrow$  prefix of length  $\ell$  of suffix at  $\text{SA}[r_{\text{start}}]$ 
16:     if  $1 < \ell \leq \ell_{\max}$  and  $\text{freq} \geq 2$  and  $p$  is a valid token then
17:        $\text{Add}(p, \text{freq} \times \ell)$  to  $V_{\text{seed}}$   $\triangleright$  Character coverage score: frequency  $\times$  length
18:        $r_{\text{lb}} \leftarrow r_{\text{start}}$   $\triangleright$  Propagate interval start for merging
19:     if  $\text{IsEmpty}(\text{stack})$  or  $\text{Top}(\text{stack}).\text{height} < h$  then
20:        $\text{Push}(\text{height} = h, \text{rank} = r_{\text{lb}})$  onto stack
21: return  $V_{\text{seed}}$ 

```

Algorithm 3 Fallback variant: including longest valid prefix of rejected tokens when creating seed vocabulary. (replacing lines 13–17 in Algorithm 2)

```

1:  $(\ell, r_{\text{start}}) \leftarrow \text{Pop}(\text{stack})$   $\triangleright$  Pattern length  $\ell$ , interval start rank
2:  $\text{freq} \leftarrow r_{\text{curr}} - r_{\text{start}}$ 
3:  $\ell_s \leftarrow \text{Top}(\text{stack}).\text{height} + 1$   $\triangleright$  if stack is empty, use  $h + 1$ 
4: for  $\ell_i = \ell$  down to  $\ell_s$  do  $\triangleright$  Change: Iterate over all intermediate lengths
5:    $p \leftarrow$  prefix of length  $\ell_i$  from suffix at  $\text{SA}[r_{\text{start}}]$ 
6:   if  $1 < \ell_i \leq \ell_{\max}$  and  $\text{freq} \geq 2$  and  $p$  is a valid token then
7:      $\text{score} \leftarrow \text{freq} \times \ell_i$ 
8:      $\text{Add}(p, \text{score})$  to  $V_{\text{seed}}$ 
9:     break  $\ell_i$  loop  $\triangleright$  Only include the longest valid token

```

C Detailed results of Unigram hyperparameter variation

Table 3: Seed Vocabulary Algorithm and Final Style Prune Results: Relative Change (%) vs Baseline, lower is better. Each corpus shows loss (relative change), token count (relative change), and vocabulary overlap (absolute %, 100% = identical to baseline), with **good** and **bad** changes above 0.5% highlighted along with the largest absolute change per parameter.

(a) Seed Vocabulary Algorithm Comparison: Relative performance compared to pretoken-based without prefix recovery.

Corpus	Metric	30 MB			300 MB		
		Full text	Full text recovery	Pretokens recovery	Full text	Full text recovery	Pretokens recovery
English	<i>Loss</i>	+0.27	+0.04	-0.01	-0.00	+0.02	+0.01
	<i>Tokens</i>	+0.26	+0.01	-0.00	-0.01	+0.01	+0.01
	<i>Vocabulary Overlap</i>	82.3	89.6	93.9	96.5	91.5	91.5
German	<i>Loss</i>	+0.14	+0.17	+0.21	+0.01	+0.17	<u>+0.19</u>
	<i>Tokens</i>	+0.32	+0.02	+0.04	-0.01	-0.03	-0.02
	<i>Vocabulary Overlap</i>	89.4	80.3	79.9	98.6	82.3	<u>81.3</u>
Arabic	<i>Loss</i>	+0.39	+0.02	+0.00	+0.02	+0.00	-0.00
	<i>Tokens</i>	+1.24	-0.05	-0.09	+0.03	-0.01	+0.01
	<i>Vocabulary Overlap</i>	74.4	83.5	84.8	95.6	89.7	89.6
Hindi	<i>Loss</i>	<u>+2.10</u>	-0.02	-0.03	+0.09	-0.00	-0.00
	<i>Tokens</i>	<u>+2.29</u>	-0.01	-0.01	<u>+0.23</u>	-0.00	+0.00
	<i>Vocabulary Overlap</i>	<u>57.9</u>	88.1	89.0	89.5	90.9	91.2
Korean	<i>Loss</i>	+0.33	+0.01	+0.01	+0.02	+0.02	+0.01
	<i>Tokens</i>	+1.24	-0.06	-0.03	+0.09	-0.01	-0.00
	<i>Vocabulary Overlap</i>	82.1	93.0	95.8	96.1	96.4	97.0
Chinese	<i>Loss</i>	+0.39	+0.02	+0.01	+0.09	+0.01	+0.00
	<i>Tokens</i>	+0.29	+0.02	+0.01	+0.07	-0.00	+0.00
	<i>Vocabulary Overlap</i>	88.6	97.4	98.0	97.7	98.9	99.3
Mean	<i>Loss</i>	+0.60	+0.04	+0.03	+0.04	+0.04	+0.04
	<i>Tokens</i>	+0.94	-0.01	-0.02	+0.07	-0.01	+0.00
	<i>Vocabulary Overlap</i>	79.1	88.6	90.2	95.7	91.6	91.6

(b) Final Style Prune: Relative performance compared to default settings, using varying pruning shrinking factor α_{prune} . We use $\alpha_{\text{inter}} = 1$ for these experiments, disabling the final pruning step, as it would be identical to normal pruning steps.

Corpus	Metric	0.0	0.25	0.5	0.75	0.9	0.95
English	<i>Loss</i>	+0.72	+0.64	+0.59	+0.55	+0.53	+0.53
	<i>Tokens</i>	+0.34	-0.06	-0.21	-0.24	-0.25	-0.25
	<i>Vocabulary Overlap</i>	73.4	75.8	75.9	75.5	74.5	74.4
German	<i>Loss</i>	+1.41	+1.06	+0.96	+0.85	+0.79	+0.77
	<i>Tokens</i>	+0.19	-0.47	-0.60	-0.70	-0.72	-0.73
	<i>Vocabulary Overlap</i>	69.0	70.4	70.9	70.5	69.9	69.7
Arabic	<i>Loss</i>	+1.19	+0.82	+0.74	+0.57	+0.49	+0.47
	<i>Tokens</i>	+0.19	-0.97	-1.02	-0.93	-0.85	-0.84
	<i>Vocabulary Overlap</i>	65.1	71.3	72.9	74.7	74.9	74.8
Hindi	<i>Loss</i>	+0.66	+0.49	+0.47	+0.45	+0.44	+0.43
	<i>Tokens</i>	+0.53	-0.00	-0.11	-0.13	-0.13	-0.14
	<i>Vocabulary Overlap</i>	74.5	77.2	77.6	77.0	76.3	76.2
Korean	<i>Loss</i>	<u>+1.77</u>	+1.08	+0.95	+0.84	+0.79	+0.77
	<i>Tokens</i>	-0.10	-1.14	-1.23	-1.30	-1.31	-1.32
	<i>Vocabulary Overlap</i>	<u>65.0</u>	70.1	70.8	71.4	71.5	71.5
Chinese	<i>Loss</i>	+1.05	+1.00	+0.97	+0.94	+0.93	+0.93
	<i>Tokens</i>	-0.09	-0.25	-0.26	-0.31	-0.33	-0.33
	<i>Vocabulary Overlap</i>	74.6	74.2	74.4	74.1	74.0	73.8
Mean	<i>Loss</i>	+1.13	+0.85	+0.78	+0.70	+0.66	+0.65
	<i>Tokens</i>	+0.18	-0.48	-0.57	-0.60	-0.60	-0.60
	<i>Vocabulary Overlap</i>	70.3	73.2	73.8	73.9	73.5	73.4

Table 4: Detailed Unigram hyperparameter results. For each corpus, we report relative change (%) in Loss and Token Count (lower is better) and absolute Vocabulary Overlap (100% = identical to baseline). Significant deviations ($> 0.5\%$) are highlighted (**good** and **bad**), with the largest absolute change per metric marked. Results are shown for relative seed size (default $\beta_{\text{seed}} = 10$), pruning overshoot (default $\alpha_{\text{inter}} = 1.1$), M-step digamma transformation (default on), pruning threshold (default $\tau_{\text{mp}} = 0.5$), EM sub-iterations (default $N_{\text{em}} = 2$), and pruning shrinking factor (default $\alpha_{\text{prune}} = 0.75$).

		Seed Vocabulary Parameters							
Corpus	Metric	$\beta_{\text{seed}} (10)$				$\alpha_{\text{inter}} (1.1)$			
		3	25	50	100	1.0	1.25	1.5	2.0
Arabic	<i>Loss</i>	+0.60	+0.03	+0.02	+0.02	-0.30	+0.44	+1.00	+2.00
	<i>Tokens</i>	+0.46	-0.17	-0.20	-0.20	+0.26	-0.21	-0.37	-0.50
	<i>Vocabulary Overlap</i>	67.1	86.7	87.3	87.3	88.0	89.0	83.0	78.2
German	<i>Loss</i>	+2.90	-0.09	-0.10	-0.11	-0.23	+0.45	+1.10	+2.09
	<i>Tokens</i>	+2.51	-0.06	-0.09	-0.09	+0.21	-0.14	-0.33	-0.35
	<i>Vocabulary Overlap</i>	<u>45.7</u>	87.1	86.4	86.6	87.1	88.2	78.8	<u>73.6</u>
English	<i>Loss</i>	+0.92	-0.02	-0.03	-0.03	-0.21	+0.32	+0.71	+1.21
	<i>Tokens</i>	+0.59	-0.01	-0.02	-0.02	+0.07	-0.04	-0.10	-0.11
	<i>Vocabulary Overlap</i>	62.1	92.2	92.4	92.4	87.4	89.3	83.0	79.2
Hindi	<i>Loss</i>	+0.79	-0.01	-0.01	-0.01	-0.20	+0.31	+0.63	+1.07
	<i>Tokens</i>	+0.57	-0.01	-0.01	-0.01	+0.05	-0.00	-0.05	-0.07
	<i>Vocabulary Overlap</i>	65.2	91.4	91.4	91.4	87.5	89.9	84.4	80.9
Korean	<i>Loss</i>	+0.63	-0.01	+0.01	+0.03	-0.38	+0.52	+1.14	+2.14
	<i>Tokens</i>	-0.59	+0.09	+0.14	+0.14	+0.26	-0.31	-0.60	-0.74
	<i>Vocabulary Overlap</i>	77.1	91.8	91.5	91.8	87.0	87.6	79.2	74.0
Chinese	<i>Loss</i>	+0.80	-0.04	-0.04	-0.03	-0.47	+0.66	+1.29	+2.06
	<i>Tokens</i>	-0.13	+0.03	+0.04	+0.04	+0.02	+0.10	+0.03	-0.08
	<i>Vocabulary Overlap</i>	75.7	96.4	96.5	96.4	87.0	88.0	81.0	76.7
Mean	<i>Loss</i>	+1.11	-0.02	-0.03	-0.02	-0.30	+0.45	+0.98	+1.76
	<i>Tokens</i>	+0.57	-0.02	-0.02	-0.02	+0.15	-0.10	-0.24	-0.31
	<i>Vocabulary Overlap</i>	65.5	90.9	90.9	91.0	87.4	88.7	81.6	77.1

Pruning Loop Parameters

Per-language results for pruning loop parameters show similarly negligible effects across all corpora, with no single variation exceeding 0.5% change in loss or token count. The mean results below summarize these findings:

Corpus	Metric	digamma (on)	$\tau_{\text{mp}} (0.5)$			$N_{\text{em}} (2)$			$\alpha_{\text{prune}} (0.75)$		
		Off	0.0	2.0	10.0	1	3	5	0.5	0.9	0.95
Mean	<i>Loss</i>	-0.01	+0.00	+0.02	-0.00	+0.02	-0.00	-0.00	+0.15	-0.07	-0.09
	<i>Tokens</i>	-0.00	+0.00	+0.00	-0.03	+0.02	-0.00	+0.00	-0.04	+0.03	+0.03
	<i>Vocabulary Overlap</i>	98.0	94.0	93.8	93.7	93.3	95.7	94.1	88.8	93.2	92.4

351
352
353

D MorphScore Evaluation

Table 5 shows MorphScore (version 2, Arnett et al., 2025) evaluation results comparing our Baseline unigram tokenizer, the FSP (Final-Style Prune) variant, and BPE on English morphological segmentation.

Word	Freq	Gold	Baseline	FSP	BPE
All Methods Match Gold (258 words covering 9.5% of corpus)					
been	2.2%	be en	be en	be en	be en
months	0.4%	month s	month s	month s	month s
following	0.3%	follow ing	follow ing	follow ing	follow ing
Only Baseline Matches (443 words covering 6.8% of corpus)					
attached	0.4%	attach ed	attach ed	atta ched	att ached
Minutes	0.2%	Minute s	Min ute s	Min utes	Min utes
animals	0.2%	animal s	animal s	anim als	anim als
Only FSP Matches (111 words covering 1.3% of corpus)					
Palestinians	0.1%	Palestinian s	P ale stin ians	Pal est inian s	Pal est in ians
instructions	0.1%	instruction s	in struct ions	in struction s	inst ru ctions
quoted	0.1%	quote d	quot ed	quote d	qu oted
Only BPE Matches (68 words covering 0.8% of corpus)					
intended	0.1%	intend ed	int ended	int ended	intend ed
bigger	0.1%	big ger	b igger	b igger	big ger
seemed	0.1%	seem ed	see med	see med	se em ed
No Method Matches (1000 words covering 12.6% of corpus)					
comments	0.2%	comment s	com ments	com ments	com ments
weapons	0.2%	weapon s	we ap ons	we ap ons	we ap ons
decided	0.1%	decide d	de cid ed	de cid ed	dec ided

Table 5: MorphScore examples for English, grouped by which methods match gold morpheme boundaries. Each category shows the three most frequent words meeting that criterion. Freq is the word’s share of total evaluation frequency.