

You should evaluate your language model on marginal likelihood over tokenisations

Kris Cao and Laura Rimell

DeepMind, London, UK

{kriscao, laurarimell}@deepmind.com

Abstract

Neural language models typically tokenise input text into sub-word units to achieve an open vocabulary. The standard approach is to use a single canonical tokenisation at both train and test time. We suggest that this approach is unsatisfactory and may bottleneck our evaluation of language model performance. Using only the one-best tokenisation ignores tokeniser uncertainty over alternative tokenisations, which may hurt model out-of-domain performance.

In this paper, we argue that instead, language models should be evaluated on their marginal likelihood over tokenisations. We compare different estimators for the marginal likelihood based on sampling, and show that it is feasible to estimate the marginal likelihood with a manageable number of samples. We then evaluate pretrained English and German language models on both the one-best-tokenisation and marginal perplexities, and show that the marginal perplexity can be significantly better than the one best, especially on out-of-domain data. We link this difference in perplexity to the tokeniser uncertainty as measured by tokeniser entropy. We discuss some implications of our results for language model training and evaluation, particularly with regard to tokenisation robustness.

1 Introduction

Neural end-to-end language models have largely done away with traditional pipeline approaches towards building NLP systems. However, one component which stubbornly remains is the tokenisation step, used right at the start of preprocessing. At the time of writing, the most widely used tokenisers, such as BPE (Sennrich et al., 2016) and unigram (Kudo, 2018), break up the input text into subword units, potentially backing off to character-level segmentation if necessary. This allows for coverage of every possible input sequence; on the downside, a single input sequence may now have multiple possible tokenisations.

Typically, language models are trained and evaluated using a single canonical tokenisation out of the multitude of possible ones, but this tokenisation may be suboptimal (Bostrom and Durrett, 2020) for many reasons. For example, different tokenisations – that is, different surface segmentations – can reveal different morphological analyses of the word in question (think *un-ion-izeable* vs. *union-izable*), and committing to a particular analysis can discard useful information, particularly if the best analysis from the tokeniser is erroneous (Dyer, 2010).

Further, tokenisers themselves are trained using an objective which optimises the likelihood of the data. This can be explicit (the unigram tokeniser of Kudo (2018) optimises a unigram language modelling objective) or implicit (BPE aims to minimise the description length of the training data, which has close connections to probabilistic methods; MacKay 2003). In this sense they are also language models, albeit far less powerful than the neural language models we train on their outputs. This raises a difficult question: to what extent are our large language models bottlenecked by the tokenisers that we use to train them?

We argue that rather than evaluating language models using the one-best tokenisation from the tokeniser, one should evaluate language models using the marginal likelihood over all possible tokenisations of an input. This divorces language model performance from the performance of the tokenisation model, and we believe this gives a better indicator of the intrinsic quality of the language model.

In this paper, we take a language model pretrained using a single tokenisation, and estimate the marginal likelihood of the model on test data, taking multiple tokenisations of each input into account. While summing exactly over exponentially many tokenisations is intractable, we can estimate the marginal likelihood using importance sampling. One contribution of this paper is to show-

case low-variance estimators of the marginal likelihood based on sampling without replacement. We cast the tokeniser as the proposal distribution for our importance sampling estimator, which clearly delimits the role of the tokeniser. Indeed, as the number of samples we consider increases, the language model becomes less and less coupled to the tokeniser, and our evaluation becomes more intrinsic to the language model itself, rather than the language model + tokeniser combination.

We demonstrate that there can be a significant difference – which we call the *marginal gap* – in marginal likelihood compared to one-best tokenisation likelihood, especially on out-of-domain evaluation sets. This suggests that the tokeniser is failing to generalise well to out-of-domain data, and is therefore a significant bottleneck to the generalisation capability of the language model. Thus, taking the one-best tokenisation likelihood is a poor proxy for the true language model performance.

We next show that there is a correlation between the uncertainty of the tokeniser (as measured by the entropy of the segmentation lattice) and the marginal gap. We give an efficient dynamic program to calculate the entropy of the segmentation lattice, and show that this entropy is predictive of how poorly the tokeniser fails to generalise. This suggests that measuring tokeniser entropy can be a useful signal for adding additional samples to our estimate of the marginal likelihood. We also use our sampled tokenisations to demonstrate that language models are particularly sensitive to variations in tokenisation, a challenge that must be mitigated for marginal likelihood evaluation.

Finally, we investigate how many samples are necessary to obtain an accurate estimate of the marginal likelihood. We show that many samples are necessary, but only relatively few samples contribute significantly to this estimate. This shows that the tokeniser distribution over tokenisations differs significantly from the language model posterior distribution over tokenisations – indeed, taking only the best tokenisation from the samples can recover most of the performance increase obtained by marginalisation. This gives weight to our finding that tokenisers generalise poorly, and that the one-best tokenisation can often be suboptimal.

We conclude by discussing some implications of our results, particularly for languages with richer morphology than English. Finally, we sketch potential future directions to bridge this gap by using

sampled tokenisations at training time, and how this might improve language model robustness.

2 Taking multiple tokenisations into consideration

We denote by D (for document) a string of text whose score we would like to calculate. Given a vocabulary V of sub-word tokens (which is usually induced by the tokeniser), we denote by T_i potential tokenisations of D – i.e. sequences of tokens $t_1 t_2 \dots t_{n_i}$ such that each $t_i \in V$ and the sequence detokenises to D . An autoregressive neural language model (with parameters θ) is a model which decomposes the probability of the full sequence into a series of left-to-right predictions: $P_\theta(T, D) = \prod_{i=1}^n P_\theta(t_i | t_{<i})$. Crucially, neural language models P_θ do not score D directly, but rather token sequences $P_\theta(T, D)$. For any input document D , a tokeniser will define a canonical tokenisation T^* , and one usually approximates $P_\theta(D)$ with $P_\theta(T^*, D)$.

We believe, on the other hand, that it is more principled to marginalise over all possible tokenisations; that is, calculate $\sum_T P_\theta(T, D)$ directly. There could be significant tokeniser uncertainty over the correct tokenisation; we can view the uncertainty as either caused by ambiguity in local context imposed by the strong independence assumptions made by tokenisers, or because of inherent tokeniser uncertainty when confronted with out-of-domain input. In either case, incorporating additional analyses in the form of extra tokenisations can give the language model extra information compared to the one-best tokenisation. We believe that the marginal likelihood better represents the true capability of the language model, without the constraint of the tokeniser.

However, exactly calculating the marginal likelihood is infeasible, as the number of possible tokenisations is exponential in the length of the input text. Whenever calculating a marginal exactly is infeasible, the classical approach is to approximate it using samples. The best distribution to sample from would be the model posterior distribution over tokenisations given text, as this gives the lowest variance estimator; unfortunately, we are unaware of any methods that would let us sample directly from this distribution. Therefore, to estimate the marginal language model likelihood, we turn to importance sampling. Given some proposal distribution $Q(T|D)$ of possible tokenisations, we can

use the importance sampling estimator

$$P(D) = \sum_T P(T, D) = \mathbb{E}_{T \sim Q} \frac{P(T, D)}{Q(T|D)} \quad (1)$$

Now, it remains to find a suitable proposal distribution $Q(T|D)$. In this paper, we use the unigram tokeniser of [Kudo \(2018\)](#), as this is the only probabilistic tokeniser that we are aware of. This tokeniser first constructs a lattice of all possible tokenisations given an input and a lexicon of word pieces. Distinct tokenisations of the input correspond to paths through this lattice, and the score of a tokenisation is the sum of the scores of the tokens along the path. As the score decomposes along lattice segments, many interesting quantities, such as $Q(D)$ (the marginal likelihood of an input text under the tokeniser), are exactly calculable. This allows not only for sampling from the lattice of possible tokenisations, but also calculating the score of a given tokenisation (i.e. estimate $Q(T|D) = Q(T, D)/Q(D)$), which is necessary to estimate the importance weight.

Tokenising consistently There is prior evidence ([Lazaridou et al., 2021](#)) to suggest that Transformer language models are able to effectively leverage memory, and that perplexities of repeated words in a document can be much lower than the perplexity of the first occurrence of that word. We show in [Section 4.3](#) that this copying ability is tied to the exact tokenisation of that word: if a word reoccurs in a document with a different tokenisation, its perplexity is much higher than if it reappears with the same tokenisation.

Armed with this insight, we design an alternative proposal distribution which samples a single tokenisation for each unique whitespace-delimited type in a document, and then shares that tokenisation for each token of that type in the document. We note that it is possible to adapt a pre-trained unigram tokeniser to do this, by passing in only the unique whitespace types in a document to the tokeniser and reconstructing the document from the sampled tokenisations. This is possible because the unigram tokeniser does not consider context when tokenising, and whitespace tokens are tokenised independently. We note that this two-stage word generation process, where first we generate the vocabulary for a document, and then generate the document from that vocabulary, has close connections to the two-stage language models proposed

in [Goldwater et al. \(2011\)](#). The problem of tokenising consistently only arises when sampling from the tokeniser; the one-best tokenisation of an input from the unigram tokeniser will always tokenise each occurrence of a type identically.

2.1 Lowering the variance of the estimator

A naive approach to estimating the marginal likelihood using [Equation 1](#) would be to sample n tokenisations T_1, \dots, T_n at random from $Q(T|D)$, score the resulting tokenisations using the language model $P_\theta(T_i, D)$, and average the resulting importance weighted scores. However, due to Jensen’s inequality, this is only a lower bound of the true marginal likelihood. We can obtain a tighter bound with the same number of samples by taking the average in probability space rather than log space (as in [Burda et al. \(2016\)](#))

$$\log P_\theta(D) \approx \log \left(\frac{1}{n} \sum_i \frac{P_\theta(T_i, D)}{Q(T_i|D)} \right) \quad (2)$$

Changing the sampling procedure Taking n independent samples from Q can result in high-variance estimates if the entropy of Q is low and it assigns low probability to tokenisations with high posterior probability under the language model P_θ . In this case, one would expect to see multiple repeated samples, which do not sufficiently explore the sample space. One option to lower the variance of the estimate is to instead sample without replacement (WOR). By enforcing that all samples are distinct, we can explore the sample space better,

However, sampling without replacement without exactly enumerating all possible sample outcomes is tricky. [Kool et al. \(2019\)](#) show how to sample without replacement for sequence models using stochastic beam search (SBS). Unfortunately, the segmentation lattice used in the unigram tokeniser is not locally normalised, and we cannot naively use SBS. We therefore adapt the SBS algorithm by first running the forward algorithm on the segmentation lattice to calculate the normalising constant at each point of the lattice; we can then combine Viterbi backwards n -best search with the constrained Gumbel-max trick used in SBS to exactly sample n tokenisations WOR.

If we sample without replacement, the inclusion probability of a tokenisation T_i is no longer equal to $Q(T_i|D)$. [Kool et al. \(2019\)](#) show that, for the expectation of a function f under a distribution Q , an unbiased estimator using a set of k samples

without replacement is given by

$$\mathbb{E}_{T \sim Q} f(T) \approx \sum_{i=1}^k \frac{Q(T_i)}{q_\kappa(T_i)} f(T_i) \quad (3)$$

κ is the perturbed score of the $k + 1$ th item during search and $q_\kappa(T) = 1 - \exp(-\exp(\log Q(T) - \kappa))$ is the probability that a Gumbel variable with location $\log Q(T)$ takes a value greater than κ . In our case, $f(T) = P_\theta(T, D)/Q(T)$, and if we calculate this sum before taking the logarithm to obtain a tighter bound, then the $Q(T)$ terms cancel and we obtain the following estimator for the marginal likelihood of a document:

$$\log P_\theta(D) \geq \log \left(\sum_i \frac{P_\theta(T_i, D)}{q_\kappa(T_i)} \right) \quad (4)$$

Including the best tokenisation To lower the variance of the estimate further (at the cost of introducing some bias), we can always include the best tokenisation from the tokeniser in our set of samples (Botev et al., 2017). This method decomposes estimating $\sum_T P_\theta(T, D)$ as $P_\theta(T^*, D) + \sum_{T \neq T^*} P_\theta(T, D)$. We can then estimate the sum over all tokenisations using exactly the same methods as before, using the new distribution Q^* which places 0 mass on T^* and renormalises the resulting probabilities for other tokenisations. It remains to simulate samples from Q^* using samples from Q . We note that for sampling with replacement, a simple technique to sample from Q^* is simple rejection sampling, where we discard any sample from Q that equals T^* . However, if $Q(T)$ is particularly peaked around T^* , then this procedure may require many rejection steps. Therefore, we do not investigate this estimator further.

When sampling without replacement, we have to be a little more careful. We note that the following scheme samples k times exactly without replacement from Q^* :

1. Take $k + 1$ items T_1, \dots, T_{k+1} WOR from Q .
2. If any $T_i = T^*$, discard it from the sample.
3. Otherwise discard T_{k+1}

We also note (by conditioning on the event that T^* appears in the sample) that the inclusion probabilities are easily calculated (if T^* appears in the sample, take κ to be the perturbed score of the $k + 2$ th item; otherwise take it to be the perturbed score of the $k + 1$ th item).

Algorithm 1: Recursive algorithm for lattice entropy

Result: entropy H_n of segmentation lattice
init $H_0 = 0$, $\alpha[i]$ the forward marginals ;
for $i = 1$ to n **do**
 for w token terminating at position i **do**
 $j =$ start position of w ;
 // $\varphi(w)$ is the score of token w ;
 $p(w) = \exp(\alpha[j] + \varphi(w) - \alpha[i])$;
 $H_i += p(w)(H_j + \log p(w))$
 end
end
return H_n

2.2 Summing over the n -best tokenisations

An alternative approach to estimating $\sum_T P_\theta(T, D)$ is to restrict the sum to a smaller set of suitable candidates. As the unigram tokenisation objective decomposes over segments, one can use Viterbi search to find exactly the n highest scoring tokenisations from the tokeniser. We then score each tokenisation using the language model, and sum the contribution of each estimate to obtain a (lower bound) estimate of the marginal likelihood. This estimator is high-bias and low-variance compared to the sampling-based estimators; we show in Section 4.1 that, although the n -best estimator performs well, it is possible to tune the sample-based estimators to perform better by trading bias for variance.

3 Measuring segmentation lattice entropy

We believe that the entropy of the tokeniser segmentation lattice is an important quantity to measure. The entropy quantifies the uncertainty of the tokeniser, and has a nice interpretation as the (logarithm of the) size of the set of alternatives the tokeniser is choosing uniformly over. While the entropy over hidden states of other structured models like HMMs and CRFs have previously been published (Hernando et al., 2005; Mann and McCallum, 2007; Ilic, 2011), and a uniform treatment in terms of expectation semirings is given in Li and Eisner (2009), we are unaware of previous elementary derivations of the entropy of a segmentation lattice. We give the algorithm in Algorithm 1.

Note that the recursion has a particularly nice interpretation in terms of information theory. Recall that the entropy of a random variable can be

thought of as the necessary number of bits to transmit the random variable. The recursion states that, to transmit the lattice up to position i (which takes H_i bits), we can transmit a prefix of the lattice (using H_j bits), and then transmit the token w that goes from j to i (using $\log P(w)$ bits). The total number of bits necessary is then the weighted sum of all possible ways of doing this, where the weights are given by the probability of that particular decomposition.

4 Experiments

For our experiments, we first pretrain language models using one-best tokenisations from a tokeniser using WMT news shared task data (Barraut et al., 2020). We train models on both English and German data up to September 2017, reserving the rest of the 2017 data for validation and model selection. We use a Transformer-XL (Dai et al., 2019) model with 18 layers and a hidden size of 1024. During evaluation time, we do not use Transformer-XL memory, due to the interaction of batching and sampled tokenisation. While this may depress our results, we are not interested in absolute model performance *per se*, but rather in the relative performance of the marginal likelihood vs. the one-best likelihood.

The tokeniser we use at both training and evaluation time is a unigram tokeniser as implemented in the SentencePiece package (Kudo, 2018), with a vocabulary size of 50529. We train the tokeniser on the same training set, with a random sample of 100 million sentences for English, and 10 million documents for German.

4.1 Measuring the marginal likelihood

For both English and German, we use 500 documents sampled randomly from the WMT train and test data and 500 randomly sampled Wikipedia documents (WIKI). For English, we also use 500 documents from the CUSTOMNEWS and arXiv abstracts (ARXIV) datasets of Lazaridou et al. (2021), and for German, we additionally use 200 documents from the MC4 dataset in Xue et al. (2020).

For each method outlined in Section 2, we sample 128 different tokenisations of each document, and calculate $P_\theta(T_i, D)$ for each sample, before aggregating the sample scores into an estimate of the marginal likelihood. We parallelise evaluating all the samples for a document on a multi-host TPU setup; each dataset takes 15-30 minutes to evaluate.

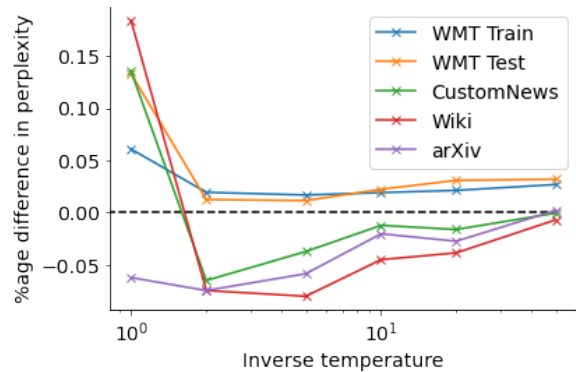


Figure 1: The effect of temperature scaling on the estimated perplexity on all English datasets, using WOR 1-best. The y -axis is the percentage difference in perplexity relative to the n -best baseline (lower is better). Note the x -axis is scaled as $1/\tau$, rather than τ .

Further, to ensure results are comparable across different tokenisations with potentially different numbers of tokens, we calculate perplexity by dividing the total likelihood across all documents by the total number of whitespace-delimited tokens. We present our results in Table 1.

Our results show that there can be a significant difference between the one-best tokenisation likelihood and the marginal likelihood, particularly as one moves further away from the training data domain. Indeed, the relative perplexity improvement reaches up to 1.9% on EN-ARXIV, and 0.9% on DE-MC4. Further, tokenising words consistently in a document has a large impact on the marginal likelihood estimation. We investigate this effect further in Section 4.3. While the n -best estimator appears to perform the best in this comparison, we show in the next section that by tuning the sampling temperature of the WOR 1-best estimator, it is possible to obtain even better estimates of the marginal likelihood.

The effect of sampling temperature We also investigate sharpening the tokeniser distribution before sampling by multiplying the log-probability of each tokenisation by a factor of $1/\tau$ before sampling. Using $\tau < 1$ has often shown to give improved results in various tasks (Kool et al., 2019; Melis et al., 2019; Adlam et al., 2020), and can be understood as a way of tuning the bias-variance tradeoff with the n -best estimator at the high-bias, low variance end, and independently sampling at the other. We compare the WOR with 1-best estimator at a various rate of temperatures on our English datasets, and show the results in Figure

		Consistent tokenization					Inconsistent tokenization				
		WR	WOR	WOR	1-best	<i>n</i> -best	WR	WOR	WOR	1-best	<i>n</i> -best
English	WMT train (16.49)	16.59	16.58		16.48	16.47	16.81	16.79		16.48	16.48
	WMT test (22.62)	22.73	22.72		22.59	22.56	23.07	23.01		22.60	22.58
	CUSTOMNEWS (37.09)	37.11	37.12		36.93	36.88	37.90	37.89		37.03	36.95
	WIKI (60.22)	61.09	61.02		59.82	59.71	63.37	63.33		60.06	59.92
	ARXIV (179.20)	176.38	176.11		175.87	175.98	179.76	179.74		177.52	176.90
German	WMT train (31.84)	32.51	32.58		31.80	31.77	33.04	33.12		31.80	31.78
	WMT test (37.16)	37.68	38.16		37.12	37.08	38.87	38.91		37.13	37.09
	WIKI (66.08)	69.44	69.30		65.86	65.63	72.37	72.41		66.01	65.78
	MC4 (194.02)	206.89	207.15		192.84	192.21	219.63	219.19		193.68	192.87

Table 1: Comparing the different estimators of model marginal perplexity on evaluation sets. The number in brackets represents the one-best tokenisation perplexity. Consistent vs. inconsistent tokenisation refers to whether we tokenise each appearance of a whitespace-delimited type consistently in a document or not.

1. One can see that it is possible to improve on the *n*-best estimator by trading some bias for variance, and this can result in a better estimate of the marginal, especially for out of domain datasets.

4.2 Tokeniser entropy and the marginal gap

Next, we investigate what causes the gap between marginal likelihood and one-best likelihood, and whether there are easily measurable factors that might predict this difference. We hypothesise that, the more uncertain the tokeniser is, the bigger this gap becomes. We pool together the documents in all our evaluation sets, and test whether there is a correlation between tokeniser entropy and marginal gap. Our results, shown in Figure 2, demonstrate that there is a correlation between entropy and the marginal gap (Spearman $r = 0.57$ for English, 0.49 for German); interestingly, it appears that high tokeniser entropy is predictive of a bigger marginal gap, but large marginal gaps are possible even if the tokeniser has low entropy.

4.3 Analysing the caching behaviour of language models

Our results show that tokenising word types consistently within a document leads to significantly tighter estimates of the marginal likelihood compared to independently tokenising input tokens. We analyse this phenomenon in this section, by investigating the loss language models assign to repeated tokens in a document, conditioned on whether the token appears in the same tokenised form or not.

Concretely, let w_1, \dots, w_m be the whitespace-delimited words in a document D , and let $\mathcal{T}_1, \dots, \mathcal{T}_n$ be the sampled tokenisations of the document. Each word w_i appears as a token sequence

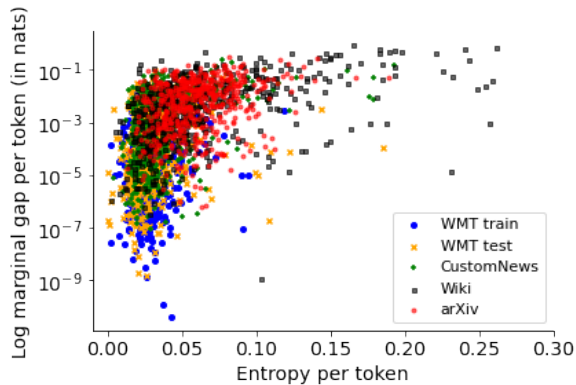
	All words			Multi-token words		
	First	(1)	(2)	First	(1)	(2)
WMT Tr	3.88	2.59	17.01	10.73	4.07	21.11
WMT Te	4.19	2.59	16.69	12.15	4.11	20.40
CNEWS	6.31	2.99	16.19	17.01	4.88	20.36
WIKI	7.84	3.62	16.54	17.80	5.63	19.81
ARXIV	9.94	3.97	14.93	17.56	5.41	18.03

Table 2: Investigating the caching ability of language models. For words which appear multiple times with different tokenisations, we show the average loss of the first occurrence of that word, of subsequent occurrences of that word with the same tokenisation (1), and subsequent occurrences of that word in a different tokenisation (2). WMT Tr and WMT Te are the WMT training and test evaluation sets respectively.

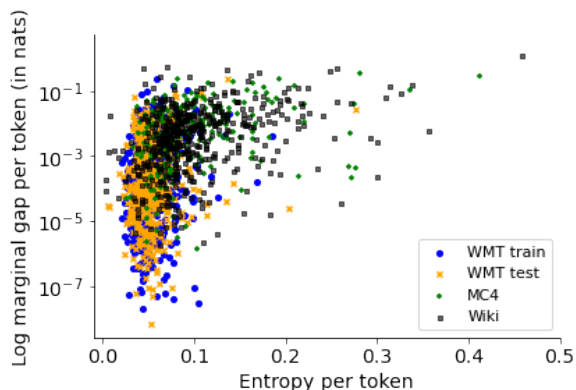
$T_{w_i} = \langle t_{w_i}^1 \dots t_{w_i}^{n_i} \rangle$, and each sampled tokenisation \mathcal{T}_i can have different token sequences $T_{w_i}^i$ for the same underlying word. We look for words $w_k \in (w_i, \dots, w_n)$ such that:

1. For some tokenisation T_l^i of w_i , for some $l < k$, $w_l = w_k$ and $T_{w_k}^i = T_{w_l}^i$ (the word has appeared before with the same tokenisation).
2. For some other tokenisation T_j^i , for all $l < k$ such that $w_l = w_k$, $T_{w_k}^j \neq T_{w_l}^j$ (all previous occurrences of this word in the document were tokenised differently).

We then calculate $P_\theta(w_k | w_{<k})$ for each tokenisation \mathcal{T}_i (by summing the scores of the tokens in w_k), and microaverage separately the loss for tokenisations which fulfill condition (1) and condition (2). The microaveraged loss for (1) represents the language model being able to copy the word as a sequence of tokens from its memory, while the microaveraged loss for (2) represents the model having to generate the word afresh as a new sequence



(a) English



(b) German

Figure 2: The correlation between entropy per token and the marginal gap per token in nats (not in perplexity), categorised by evaluation dataset. Some data points which extend beyond the right of the graph are truncated; they follow the same trend.

of tokens. By comparing the loss of words paired in this way, we can control for extra confounding factors (such as token unigram probability), and isolate the ability of the language model to recognise whether different token sequences correspond to the same underlying form.

We show our results for our various datasets, together with selected subsets of words, in Table 2. We see that, if the language model sees a word after already seeing it in the same tokenisation, its loss is significantly lower than the loss associated with the first time the word is seen (as was also reported in Lazaridou et al. (2021)). However, this ability is strongly tied to the exact tokenisation of the word: if it appears again, but in a different tokenisation, then its loss can in fact be even greater.

4.4 How many samples are necessary?

Next, we investigate how many samples are necessary to obtain an accurate estimate of the marginal likelihood. We experiment on the EN-ARXIV

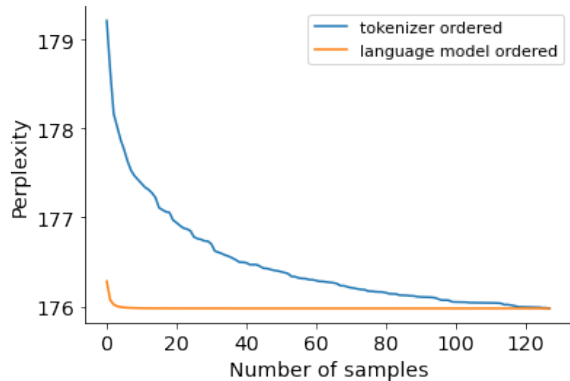


Figure 3: The performance of the n -best marginal likelihood estimator on the ARXIV evaluation set as we vary the number of samples, taken in order of $Q(T|D)$ in orange and $P_\theta(T, D)$ in blue.

dataset, as this showed the biggest relative improvement between the marginal likelihood and the one-best likelihood. We take the samples from our n -best estimator with $n = 128$, and incrementally sum the samples (which are given in decreasing order of likelihood under the tokeniser) to simulate having smaller n . As an oracle experiment to see how many samples contribute significantly to the marginal likelihood, we also order the samples by their language model scores (i.e. we order according to $P_\theta(T, D)$ rather than $Q(T|D)$) before taking the incremental sum. We show the results in Figure 3. Our results show that, although ostensibly many samples are necessary to estimate the marginal likelihood accurately, only very few samples (in the order of 5) actually contribute significantly.

In practical terms, our results suggest that one needs to take many samples with current tokenisers to accurately estimate the marginal likelihood, but that many of these samples are not effective. We therefore believe that a prerequisite for more widespread adoption of marginal likelihood as an evaluation metric is tokenisers that better fit the language model posterior over tokenisations. Current tokenisers make very strong independence assumptions to make learning and inference tractable, and we believe there is significant scope to design tokenisers which relax these assumptions.

5 Related Work

5.1 Tokenisation and segmentation

Unsupervised word segmentation has a long and illustrious history. The earliest motivations were in information retrieval, and the motivation was that collapsing a set of related query terms might help

smooth counts over each of those terms individually and result in better retrieval results. The earliest approaches, such as the Porter stemmer (Porter, 1997), were rule-based. However, the power of data-driven statistical methods quickly became apparent, and tools such as Morfessor (Virpioja et al., 2013) used likelihood-based objectives, typically with Bayesian smoothing methods (see also Goldwater et al. (2011)), to induce segmentations.

Sennrich et al. (2016) used a different algorithm to induce segmentations: byte-pair encoding (Gage, 1994). Originally designed as a data compression algorithm, BPE tokenisers are now some of the predominantly used tokenisation methods. Alternative approaches, such as WordPiece (Schuster and Nakajima, 2012) and SentencePiece (Kudo, 2018), explicitly use a language modelling objective to induce a token lexicon. Previous methods have used train-time tokenisation randomisation as a regularisation aid (Kudo, 2018; Provilkov et al., 2020), but still use the one-best tokenisation at test time.

Another strand of work has investigated whether tokenisers that capture linguistic morphology can improve language models. Bostrom and Durrett (2020) showed that unigram and BPE tokenisers for English and Japanese have low recall on recovering linguistic segments, since many morphologically complex words are treated as a single token. Linguistically aligned tokenisers have been shown to result in better language model perplexity (Schwartz et al., 2020; Park et al., 2021) and better downstream task performance (Alkaoud and Syed, 2020), especially for morphologically rich languages. These experiments also use one-best tokenisation at test time.

Rather than considering one-best or stochastic samples of tokenisations, one can use entire segmentation lattices as input to a model. This approach has been considered for morphological tagging (Seker and Tsarfaty, 2020), parsing (Goldberg and Tsarfaty, 2008), and spoken intent recognition (Ladhak et al., 2016), among others.

5.2 Tokenisation-free approaches

An alternative approach to inducing a tokenisation is to decompose input sequences into well-defined orthographic units, such as characters. These approaches circumvent the problem of inducing a lexicon, and have been used for text classification (Conneau et al., 2017), language modelling (Al-Rfou et al., 2019), machine translation (Lee

et al., 2017), and word representation (Cao and Rei, 2016). One downside is that dependency lengths become longer on the character-level, and lexical information has to be memorised by the compositional machinery of the model. For this reason, traditionally fully character-based approaches did not perform as well as their token-level counterparts, although recent progress suggests this may change soon (Choe et al., 2019; Clark et al., 2021). There also exist approaches which mix character-level and segment-level approaches (Buckman and Neubig, 2018; Kawakami et al., 2019; He et al., 2020), although these segmental language models require more complex inference procedures.

6 Conclusion

In this paper, we argue for using model marginal likelihood over tokenisations as an evaluation metric for language models, rather than one-best tokenisation likelihood. We introduce practical low-variance estimators for measuring the marginal likelihood, and demonstrate that there can be significant difference between the marginal and the one-best likelihoods, particularly on strongly out-of-domain evaluation sets. Evaluating with marginal likelihood thus goes some way toward loosening the bottleneck imposed by tokeniser quality in the currently dominant language modelling paradigm, and our results suggest that the field may be underestimating the generalisation capability of modern language models. We further demonstrate that tokeniser entropy is a good predictor of this “marginal gap”, suggesting that tokeniser entropy, especially when out-of-domain, can be a guide to the number of samples needed for evaluation.

More broadly, our experiments suggest that the field should continue seeking better ways to incorporate tokenisation into end-to-end language modelling. Sampling from the tokeniser during training is an obvious possibility; alternatively, one could incorporate the segmentation lattice into the model directly, which has been beneficial for parsing morphologically rich languages (Goldberg and Tsarfaty, 2008; Tsarfaty et al., 2020). Further, developing more contextual tokenisers which make fewer independence assumptions can also result in both better language models trained on their one-best tokenisation, and better evaluation estimates of the marginal likelihood with fewer samples.

We conduct experiments on German and English corpora in this paper. However, these two

languages are only a small sample in the full space of language typology. English is a morphologically impoverished language, and while German compounding and inflection offer some additional challenges, many languages have more complex patterns of word formation and inflection. We believe that estimating marginal likelihood will be important for morphologically richer languages, where tokenisation makes a bigger difference (Gerz et al., 2018; Mielke et al., 2019).

Finally, improved understanding of the interaction between tokenisation and language modelling has implications for evaluating language models on both downstream tasks and language generation tasks. Evidence has shown that gains in language modelling, as measured in perplexity, often lead to improvements in downstream task performance (Radford et al., 2019). It would be instructive to extend our marginal likelihood approach to downstream task evaluation. On generation tasks, since the tokeniser affects language model training but is only implicitly used when sampling (via the tokeniser vocabulary), the effect of tokenisation algorithms requires careful investigation.

Acknowledgements

The authors would like to thank Dani Yogatama and the rest of the Language group at DeepMind for comments and discussion, Gábor Melis and Phil Blunsom for comments on an earlier draft, and Mark Rowland for clarification remarks on sampling without replacement. We would also like to thank our anonymous reviewers.

References

- Ben Adlam, Jasper Snoek, and Samuel L. Smith. 2020. [Cold posteriors and aleatoric uncertainty](#).
- Rami Al-Rfou, Dokook Choe, Noah Constant, Mandy Guo, and Llion Jones. 2019. [Character-level language modeling with deeper self-attention](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):3159–3166.
- Mohamed Alkaoud and Mairaj Syed. 2020. On the importance of tokenization in Arabic embedding models. In *Proceedings of the Fifth Arabic Natural Language Processing Workshop (WANLP)*, pages 119–129.
- Loïc Barrault, Magdalena Biesialska, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Matthias Huck, Eric Joanis, Tom Kocmi, Philipp Koehn, Chi-kiu Lo, Nikola Ljubešić, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Santanu Pal, Matt Post, and Marcos Zampieri. 2020. [Findings of the 2020 conference on machine translation \(WMT20\)](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1–55, Online. Association for Computational Linguistics.
- Kaj Bostrom and Greg Durrett. 2020. [Byte pair encoding is suboptimal for language model pretraining](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4617–4624, Online. Association for Computational Linguistics.
- Aleksandar Botev, Bowen Zheng, and David Barber. 2017. [Complementary Sum Sampling for Likelihood Approximation in Large Scale Classification](#). In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1030–1038, Fort Lauderdale, FL, USA. PMLR.
- Jacob Buckman and Graham Neubig. 2018. Neural lattice language models. *Transactions of the Association for Computational Linguistics*.
- Yuri Burda, Roger B. Grosse, and Ruslan Salakhutdinov. 2016. [Importance weighted autoencoders](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Kris Cao and Marek Rei. 2016. [A joint model for word embedding and word morphology](#). In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pages 18–26, Berlin, Germany. Association for Computational Linguistics.
- Dokook Choe, Rami Al-Rfou, Mandy Guo, Heeyoung Lee, and Noah Constant. 2019. [Bridging the gap for tokenizer-free language models](#). *CoRR*, abs/1908.10322.
- Jonathan H. Clark, Dan Garrette, Iulia Turc, and John Wieting. 2021. [CANINE: pre-training an efficient tokenization-free encoder for language representation](#). *CoRR*, abs/2103.06874.
- Alexis Conneau, Holger Schwenk, Loïc Barrault, and Yann Lecun. 2017. [Very deep convolutional networks for text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 1107–1116, Valencia, Spain. Association for Computational Linguistics.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.

- Christopher Dyer. 2010. *A Formal Model of Ambiguity and its Applications in Machine Translation*. Ph.D. thesis, University of Maryland.
- Philip Gage. 1994. A new algorithm for data compression. *C Users J.*, 12(2):23–38.
- Daniela Gerz, Ivan Vulić, Edoardo Maria Ponti, Roi Reichart, and Anna Korhonen. 2018. [On the relation between linguistic typology and \(limitations of\) multilingual language modeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 316–327, Brussels, Belgium. Association for Computational Linguistics.
- Yoav Goldberg and Reut Tsarfaty. 2008. [A single generative model for joint morphological segmentation and syntactic parsing](#). In *Proceedings of ACL-08: HLT*, pages 371–379, Columbus, Ohio. Association for Computational Linguistics.
- Sharon Goldwater, Thomas L. Griffiths, and Mark Johnson. 2011. [Producing power-law distributions and damping word frequencies with two-stage language models](#). *Journal of Machine Learning Research*, 12(68):2335–2382.
- Xuanli He, Gholamreza Haffari, and Mohammad Norouzi. 2020. [Dynamic programming encoding for subword segmentation in neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3042–3051, Online. Association for Computational Linguistics.
- D. Hernando, V. Crespi, and G. Cybenko. 2005. [Efficient computation of the hidden markov model entropy for a given observation sequence](#). *IEEE Transactions on Information Theory*, 51(7):2681–2685.
- Velimir M. Ilic. 2011. [Entropy semiring forward-backward algorithm for HMM entropy computation](#). *CoRR*, abs/1108.0347.
- Kazuya Kawakami, Chris Dyer, and Phil Blunsom. 2019. [Learning to discover, ground and use words with segmental neural language models](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6429–6441, Florence, Italy. Association for Computational Linguistics.
- Wouter Kool, Herke Van Hoof, and Max Welling. 2019. [Stochastic beams and where to find them: The Gumbel-top-k trick for sampling sequences without replacement](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 3499–3508. PMLR.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. Association for Computational Linguistics.
- Faisal Ladhak, Ankur Gandhe, Markus Dreyer, Lambert Mathias, Ariya Rastrow, and Björn Hoffmeister. 2016. [Latticernn: Recurrent neural networks over lattices](#). In *Interspeech 2016*, pages 695–699.
- Angeliki Lazaridou, Adhiguna Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Sebastian Ruder, Dani Yogatama, Kris Cao, Tomas Kocisky, Susannah Young, and Phil Blunsom. 2021. [Pitfalls of static language modelling](#).
- Jason Lee, Kyunghyun Cho, and Thomas Hofmann. 2017. [Fully character-level neural machine translation without explicit segmentation](#). *Transactions of the Association for Computational Linguistics*, 5:365–378.
- Zhifei Li and Jason Eisner. 2009. [First- and second-order expectation semirings with applications to minimum-risk training on translation forests](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 40–51, Singapore. Association for Computational Linguistics.
- David J. C. MacKay. 2003. [Information theory, inference, and learning algorithms](#).
- Gideon Mann and Andrew McCallum. 2007. [Efficient computation of entropy gradient for semi-supervised conditional random fields](#). In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 109–112, Rochester, New York. Association for Computational Linguistics.
- Gábor Melis, Charles Blundell, Tomáš Kočiský, Karl Moritz Hermann, Chris Dyer, and Phil Blunsom. 2019. [Pushing the bounds of dropout](#).
- Sabrina J. Mielke, Ryan Cotterell, Kyle Gorman, Brian Roark, and Jason Eisner. 2019. [What kind of language is hard to language-model?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4975–4989, Florence, Italy. Association for Computational Linguistics.
- Hyunji Hayley Park, Katherine J. Zhang, Coleman Hayley, Kenneth Steimel, Han Liu, and Lane Schwartz. 2021. [Morphology matters: A multilingual language modeling analysis](#). *Transactions of the Association for Computational Linguistics*, 9:261–276.
- M. F. Porter. 1997. *An Algorithm for Suffix Stripping*, page 313–316. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Ivan Provilkov, Dmitrii Emelianenko, and Elena Voita. 2020. [BPE-dropout: Simple and effective subword regularization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1882–1892, Online. Association for Computational Linguistics.

- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, and Dario Amodei and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. OpenAI Technical Report.
- Mike Schuster and Kaisuke Nakajima. 2012. [Japanese and korean voice search](#). In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5149–5152.
- Lane Schwartz, Francis Tyers, Lori Levin, Christo Kirov, Patrick Littell, Chi-kiu Lo, Emily Prud’hommeaux, Hyunji Hayley Park, Kenneth Steimel, Rebecca Knowles, Jeffrey Micher, Lonny Strunk, Han Liu, Coleman Haley, Katherine J. Zhang, Robbie Jimmerson, Vasilisa Andriyanets, Aldrian Obaja Muis, Naoki Otani, Jong Hyuk Park, and Zhisong Zhang. 2020. [Neural polysynthetic language modelling](#). Final Report of the Neural Polysynthetic Language Modelling Team at the 2019 Frederick Jelinek Memorial Summer Workshop.
- Amit Seker and Reut Tsarfaty. 2020. [A pointer network architecture for joint morphological segmentation and tagging](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4368–4378, Online. Association 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 Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Reut Tsarfaty, Dan Bareket, Stav Klein, and Amit Seker. 2020. [From SPMRL to NMRL: What did we learn \(and unlearn\) in a decade of parsing morphologically-rich languages \(MRLs\)?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7396–7408, Online. Association for Computational Linguistics.
- Sami Virpioja, Peter Smit, Stig-Arne Grönroos, and Mikko Kurimo. 2013. [Morfessor 2.0: Python Implementation and Extensions for Morfessor Baseline](#). Technical report.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2020. [mt5: A massively multilingual pre-trained text-to-text transformer](#). *CoRR*, abs/2010.11934.