

Morfessor EM+Prune: Improved Subword Segmentation with Expectation Maximization and Pruning

Stig-Arne Grönroos¹, Sami Virpioja², Mikko Kurimo¹

¹Department of Signal Processing and Acoustics, Aalto University, Finland

²Department of Digital Humanities, University of Helsinki, Finland
{stig-arne.gronroos,mikko.kurimo}@aalto.fi, sami.virpioja@helsinki.fi

Abstract

Data-driven segmentation of words into subword units has been used in various natural language processing applications such as automatic speech recognition and statistical machine translation for almost 20 years. Recently it has become more widely adopted, as models based on deep neural networks often benefit from subword units even for morphologically simpler languages. In this paper, we discuss and compare training algorithms for a unigram subword model, based on the Expectation Maximization algorithm and lexicon pruning. Using English, Finnish, North Sami, and Turkish data sets, we show that this approach is able to find better solutions to the optimization problem defined by the Morfessor Baseline model than its original recursive training algorithm. The improved optimization also leads to higher morphological segmentation accuracy when compared to a linguistic gold standard. We publish implementations of the new algorithms in the widely-used Morfessor software package.

Keywords: Morphology, Statistical and Machine Learning Methods, Language Modelling, Unsupervised learning, Tools, Less-Resourced/Endangered Languages

1. Introduction

Subword segmentation has become a standard preprocessing step in many neural approaches to natural language processing (NLP) tasks, e.g Neural Machine Translation (NMT) (Sennrich et al., 2015) and Automatic Speech Recognition (ASR) (Smit et al., 2017). Word level modeling suffers from sparse statistics, issues with Out-of-Vocabulary (OOV) words, and heavy computational cost due to a large vocabulary. Word level modeling is particularly unsuitable for morphologically rich languages, but subwords are commonly used for other languages as well. Subword segmentation is best suited for languages with agglutinative morphology.

While rule-based morphological segmentation systems can achieve high quality, the large amount of human effort needed makes the approach problematic, particularly for low-resource languages. The systems are language dependent, necessitating use of multiple tools in multilingual setups. As a fast, cheap and effective alternative, data-driven segmentation can be learned in a completely unsupervised manner from raw corpora. Unsupervised morphological segmentation saw much research interest until the early 2010’s; for a survey on the methods, see Hammarström and Borin (2011). Semi-supervised segmentation with already small amounts of annotated training data was found to improve the accuracy significantly when compared to a linguistic segmentation; see Ruokolainen et al. (2016) for a survey. While this line of research has been continued in supervised and more grammatically oriented tasks (Cotterell et al., 2017), the more recent work on unsupervised segmentation is less focused on approximating a linguistically motivated segmentation. Instead, the aim has been to tune subword segmenta-

tions for particular applications. For example, the simple substitution dictionary based Byte Pair Encoding segmentation algorithm (Gage, 1994), first proposed for NMT by Sennrich et al. (2015), has become a standard in the field. Especially in the case of multilingual models, training a single language-independent subword segmentation method is preferable to linguistic segmentation (Arivazhagan et al., 2019).

In this study, we compare three existing and one novel subword segmentation method, all sharing the use of a unigram language model in a generative modeling framework. The previously published methods are Morfessor Baseline (Creutz and Lagus, 2002), Greedy Unigram Likelihood (Varjokallio et al., 2013), and SentencePiece (Kudo, 2018). The new Morfessor variant proposed in this work is called Morfessor EM+Prune. The contributions of this article are

- (i) a better training algorithm for Morfessor Baseline, with reduction of search error during training, and improved segmentation quality for English, Finnish and Turkish;
- (ii) comparing four similar segmentation methods, including a close look at the SentencePiece reference implementation, highlighting details omitted from the original article (Kudo, 2018);
- (iii) and showing that the proposed Morfessor EM+Prune with particular hyper-parameters yields SentencePiece.

1.1. Morphological Segmentation with Unigram Language Models

Morphological surface segmentation is the task of splitting words into morphs, the surface forms of meaning-bearing sub-word units, morphemes. The concatena-

	Morfessor BL	Greedy Unigram	SentencePiece	Morfessor EM+Prune
Model	Unigram LM	Unigram LM	Unigram LM	Unigram LM
Cost function	MAP	ML	MAP	MAP
Prior	MDL	–	DP	MDL+DP
Training algorithm	Local search	EM+Prune	EM+Prune	EM+Prune
Initialization	Words	Seed lexicon	Seed lexicon	Seed lexicon
EM variant	–	Lateen-EM once	EM	EM / Lateen-EM
Stopping criterion				
Cost change threshold	✓	✓	–	✓
Target lexicon size	Approximate	✓	✓	✓
N-best decoding	✓	–	✓	✓
Sampling decoding	–	–	✓	✓
Count dampening	✓	–	–	✓
Semi-supervised	✓	–	–	✓
Requires pretokenization	✓	✓	–	✓
Reference implementation	Python	C++	C++	Python

Table 1: Comparison of subword segmentation methods applying a unigram language model.

tion of the morphs is the word, e.g.

$$\textit{reliability} \mapsto \textit{reli} + \textit{abil} + \textit{ity}$$

Probabilistic generative methods for morphological segmentation model the probability $P(\mathbf{s})$ of generating a sequence of morphs (a word, sentence or corpus) $\mathbf{s} = [m_0, \dots, m_N]$, as opposed to discriminative methods that model the conditional probability of the segmentation boundaries given the unsegmented data.

This study focuses on segmentation methods applying a *unigram language model*. In the unigram language model, an assumption is made that the morphs in a word occur independently of each other. Alternatively stated, it is a zero-order (memoryless) Markov model, generalized so that one observation can cover multiple characters. The probability of a sequence of morphs decomposes into the product of the probabilities of the morphs of which it consists.

$$P_{\theta}(\mathbf{s}) = \prod_{i=1}^N P_{\theta}(m_i) \quad (1)$$

The Expectation Maximization (EM) algorithm (Dempster et al., 1977) is an iterative algorithm for finding Maximum Likelihood (ML) or Maximum a Posteriori (MAP) estimates for parameters in models with latent variables. The EM algorithm consists of two steps. In the E-step (2), the expected value of the complete data likelihood including the latent variable is taken, and in the M-step (3), the parameters are updated to maximize the expected value of the E-step:

$$Q(\theta, \theta^{(i-1)}) = \int_{\mathbf{y}} \log P(\mathbf{D}, \mathbf{y} | \theta) P(\mathbf{y} | \mathbf{D}, \theta^{(i-1)}) d\mathbf{y} \quad (2)$$

$$\theta^i = \arg \max_{\theta} Q(\theta, \theta^{(i-1)}). \quad (3)$$

When applied to a (hidden) Markov model, EM is called the forward-backward algorithm. Using instead

the related Viterbi algorithm (Viterbi, 1967) is sometimes referred to as *hard-EM*.¹ Spitzkovsky et al. (2011) present lateen-EM, a hybrid variant in which EM and Viterbi optimization are alternated.

Virpioja (2012, Section 6.4.1.3) discusses the challenges of applying EM to learning of generative morphology. Jointly optimizing both the morph lexicon and the parameters for the morphs is intractable. If, like in Morfessor Baseline, the cost function is discontinuous when morphs are added or removed from the lexicon, there is no closed form solution to the M-step. With ML estimates for morph probabilities, EM can neither add nor remove morphs from the lexicon, because it can neither change a zero probability to nonzero nor vice versa.

One solution to this challenge is to apply local search. Starting from the current best estimate for the parameters, small search steps are tried to explore near-lying parameter configurations. The choice that yields the lowest cost is selected as the new parameters. Greedy local search often gets stuck in local minima. Even if there are parameters yielding a better cost, the search may not find them, causing search error. The error remaining at the parameters with globally optimal cost is the model error.

Another solution is to combine EM with pruning (EM+Prune). The methods based on pruning begin with a seed lexicon, which is then iteratively pruned until a stopping condition is reached. Subwords cannot be added to the lexicon after initialization. As a consequence, proper initialization is important, and the methods should not prune too aggressively without reestimating parameters, as pruning decisions cannot be backtracked. For this reason, EM+Prune methods

¹An analogy can be drawn to clustering using k -means, which yields a hard assignment of data points to clusters, and using EM for clustering with a Gaussian Mixture Model (GMM), where the assignment is soft.

proceed iteratively, only pruning subwords up to a pre-defined iteration pruning quota, e.g. removing at most 20% of the remaining lexicon at a time.

2. Related Work

In this section we review three previously published segmentation methods that apply a unigram language model. Table 1 summarizes the differences between these methods.

2.1. Morfessor Baseline

Morfessor is a family of generative models for unsupervised morphology induction (Creutz and Lagus, 2007). Here, consider the Morfessor 2.0 implementation (Virpioja et al., 2013) of Morfessor Baseline method (Creutz and Lagus, 2002).

A point estimate for the model parameters $\hat{\theta}$ is found using MAP estimation with a Minimum Description Length (MDL) (Rissanen, 1989) inspired prior that favors lexicons containing fewer, shorter morphs. The MAP estimate yields a two-part cost function, consisting of a prior (the lexicon cost) and likelihood (the corpus cost). The model can be tuned using the hyper-parameter α , which is a weight applied to the likelihood (Kohonen et al., 2010):

$$\hat{\theta} = \arg \min_{\theta} \left\{ \overbrace{-\log P(\theta)}^{\text{prior}} - \alpha \overbrace{\log P(D|\theta)}^{\text{likelihood}} \right\} \quad (4)$$

The α parameter controls the overall amount of segmentation, with higher values increasing the weight of each emitted morph in the corpus (leading to less segmentation), and lower values giving a relatively larger weight to a small lexicon (more segmentation).

The prior can be further divided into two parts: the prior for the morph form properties and the usage properties. The form properties encode the string representation of the morphs, while the usage properties encode their frequencies. Morfessor Baseline applies a non-informative prior for the distribution of the morph frequencies. It is derived using combinatorics from the number of ways that the total token count ν can be divided among the μ lexicon items:

$$P(\tau_1, \dots, \tau_\mu | \mu, \nu) = 1 / \binom{\nu - 1}{\mu - 1}. \quad (5)$$

Morfessor Baseline is initialized with a seed lexicon of whole words. The Morfessor Baseline training algorithm is a greedy local search. During training, in addition to storing the model parameters, the current best segmentation for the corpus is stored in a graph structure. The segmentation is iteratively refined, by looping over all the words in the corpus in a random order and resegmenting them. The resegmentation is applied by recursive binary splitting, leading to changes in other words that share intermediary units with the word currently being resegmented. The search converges to a local optimum, and is known to be sensitive to the initialization (Virpioja et al., 2013).

In the Morfessor 2.0 implementation, the likelihood weight hyper-parameter α is set either with a grid search using the best evaluation score on a held-out development set, or by applying an approximate automatic tuning procedure based on a heuristic guess of which direction the α parameter should be adjusted.

2.2. Greedy Unigram Likelihood

Varjokallio et al. (2013) presents a subword segmentation method, particularly designed for use in ASR. It applies greedy pruning based on unigram likelihood. The seed lexicon is constructed by enumerating all substrings from a list of common words, up to a specified maximum length. Pruning proceeds in two phases, which the authors call *initialization* and *pruning*.

In the first phase, a character-level language model is trained. The initial probabilities of the subwords are computed using the language model. The probabilities are refined by EM, followed by hard-EM. During the hard-EM, frequency based pruning of subwords begins. In the second phase, hard-EM is used for parameter estimation. At the end of each iteration, the least frequent subwords are selected as candidates for pruning. For each candidate subword, the change in likelihood when removing the subword is estimated by resegmenting all words in which the subword occurs. After each pruned subword, the parameters of the model are updated. Pruning ends when the goal lexicon size is reached or the change in likelihood no longer exceeds a given threshold.

2.3. SentencePiece

SentencePiece (Kudo and Richardson, 2018; Kudo, 2018) is a subword segmentation method aimed for use in any NLP system, particularly NMT. One of its design goals is use in multilingual systems.

Although (Kudo, 2018) implies a use of maximum likelihood estimation, the reference implementation² uses the implicit Dirichlet Process prior called Bayesian EM (Liang and Klein, 2007). In the M-step, the count normalization is modified to

$$P(z) = \frac{\exp(\Psi(C_z))}{\exp(\Psi(\sum_{z'} C_{z'}))} \quad (6)$$

where Ψ is the digamma function.

The seed lexicon is simply the e.g. one million most frequent substrings. SentencePiece uses an EM+Prune training algorithm. Each iteration consists of two sub-iterations of EM, after which the lexicon is pruned. Pruning is based on Viterbi counts (EM+Viterbi-prune). First, subwords that do not occur in the Viterbi segmentation are pre-pruned. The cost function is the estimated change in likelihood when the subword is removed, estimated using the assumption that all probability mass of the removed subword goes to its Viterbi segmentation. Subwords are sorted according to the cost, and a fixed proportion of remaining subwords are pruned each iteration. Single character

²<https://github.com/google/sentencepiece>




Figure 1: Unweighted Morfessor cost function components (prior and likelihood). Log scale.

subwords are never pruned. A predetermined lexicon size is used as the stopping condition.

3. Morfessor EM+Prune

Morfessor EM+Prune³ uses the unigram language model and priors similar to Morfessor Baseline, but combines them with EM+Prune training.

3.1. Prior

The prior must be slightly modified for use with the EM+Prune algorithm. The prior for the frequency distribution (5) is derived using combinatorics. When using real-valued expected counts, there are infinite assignments of counts to parameters. Despite not being theoretically motivated, it can still be desirable to compute an approximation of the Baseline frequency distribution prior, in order to use EM+Prune as an improved search to find more optimal parameters for the original cost. To do this, the real valued token count ν is rounded to the nearest integer⁴. Alternatively, the prior for the frequency distribution can be omitted, or a new prior with suitable properties could be formulated. We do not propose a completely new prior in this work, instead opting to remain as close as possible to Morfessor Baseline.

³Software available at <https://github.com/Waino/morfessor-emprune>.

⁴An alternative would be to replace the factorial with the gamma function. This added precision serves no practical purpose, particularly as we already use Stirling's approximation of the factorial.

In Morfessor EM+Prune, morphs are explicitly stored in the lexicon, and morphs are removed from the lexicon only during pruning. This differs from Morfessor Baseline, in which a morph is implicitly considered to be stored in the lexicon if it has non-zero count.

The prior for the morph form properties does not need to be modified. During the EM parameter estimation, the prior for the morph form properties is omitted as the morph lexicon remains constant. During pruning, the standard form prior is applicable.

Additionally we apply the Bayesian EM implicit Dirichlet Process prior (Liang and Klein, 2007). We experiment with four variations of the prior:

1. the full EM+Prune prior,
2. omitting the Bayesian EM (noexp Ψ),
3. omitting the approximate frequency distribution prior (nofreqdistr),
4. and omitting the prior entirely (noprior).

3.2. Seed Lexicon

The seed lexicon consists of the one million most frequent substrings, with two restrictions on which substrings to include: pre-pruning of redundant subwords, and forcesplit. Truncating to the chosen size is performed after pre-pruning, which means that pre-pruning can make space for substrings that would otherwise have been below the threshold.

Pre-pruning of *redundant subwords* is based on occurrence counts. If a string x occurs n times, then any substring of x will occur at least n times. Therefore, if

	FS	Prior ↓	Likelihood ↓	W-sum ↓
Words		1.79×10^7	1.32×10^7	2.98×10^7
Characters		2.35×10^2	2.90×10^7	2.61×10^7
EM+P MDL	✓	4.69×10^5	2.09×10^7	1.92×10^7
Morfessor Baseline	✓	7.55×10^5	2.05×10^7	1.92×10^7
Morfessor Baseline		8.84×10^5	1.99×10^7	1.88×10^7
EM+P MDL		5.80×10^5	2.02×10^7	1.88×10^7
EM+P MDL lateen		6.35×10^5	2.01×10^7	1.88×10^7

Table 2: Morfessor cost results for English. $\alpha = 0.9$. FS is short for forcesplit, W-sum for weighted sum of prior and likelihood. ↓ means that lower values are better. The bolded method is our primary configuration.

	FS	Prior ↓	Likelihood ↓	W-sum ↓
Words		8.64×10^7	4.77×10^7	8.74×10^7
Characters		2.46×10^2	1.27×10^8	2.54×10^6
Morfessor Baseline	✓	8.31×10^4	8.60×10^7	1.80×10^6
Morfessor Baseline		8.36×10^4	8.59×10^7	1.80×10^6
EM+P MDL	✓	1.29×10^5	8.28×10^7	1.79×10^6
EM+P MDL lateen		1.41×10^5	8.22×10^7	1.79×10^6
EM+P MDL		1.31×10^5	8.26×10^7	1.78×10^6

Table 3: Morfessor cost results for Finnish. $\alpha = 0.02$.

the substring has a count of exactly n , we know that it is not needed in any other context except as a part of x . Such unproductive substrings are likely to be poor candidate subwords, and can be removed to make space in the seed lexicon for more useful substrings. This pre-pruning is not a neutral optimization, but does affect segmentation results. We check all initial and final substrings for redundancy, but do not pre-prune internal substrings.

To achieve *forced splitting* before or after certain characters, e.g. hyphens, apostrophes and colons, substrings which include a forced split point can be removed from the seed lexicon. As EM+Prune is unable to introduce new subwords, this pre-pruning is sufficient to guarantee the forced splits. While Morfessor 2.0 only implements force splitting certain characters to single-character morphs, i.e. force splitting on both sides, we implement more fine-grained force splitting separately before and after the specified character.

3.3. Training Algorithm

We experiment with three variants of the EM+Prune iteration structure:

1. EM,
2. Lateen-EM,
3. EM+Viterbi-prune

EM+Viterbi-prune is an intermediary mode between EM and lateen-EM in the context of pruning. The pruning decisions are made based on counts from a single iteration of Viterbi training, but these Viterbi counts are not otherwise used to update the parameters. In effect, this allows for the more aggressive pruning using the Viterbi counts, while retaining the uncertainty of the soft parameters.

Each iteration begins with 3 sub-iterations of EM. In

	Prior ↓	Likelihood ↓	W-sum ↓
Words	1.31×10^7	9.09×10^6	1.68×10^7
Characters	1.19×10^2	2.08×10^7	8.30×10^6
Morfessor Baseline	2.54×10^5	1.39×10^7	5.82×10^6
EM+P MDL lateen	2.79×10^5	1.37×10^7	5.78×10^6
EM+P MDL	2.71×10^5	1.37×10^7	5.77×10^6
EM+P MDL keep-redundant	2.97×10^5	1.36×10^7	5.73×10^6

Table 4: Morfessor cost results for Turkish. $\alpha = 0.4$

	FS	Prior ↓	Likelihood ↓	W-sum ↓
Words		2.12×10^7	1.03×10^7	3.15×10^7
Characters		1.38×10^3	2.98×10^7	2.98×10^7
Morfessor Baseline	✓	1.76×10^6	1.62×10^7	1.80×10^7
Morfessor Baseline		1.87×10^6	1.61×10^7	1.80×10^7
EM+P MDL	✓	9.52×10^5	1.70×10^7	1.79×10^7
EM+P MDL lateen		9.83×10^5	1.69×10^7	1.79×10^7
EM+P MDL		9.56×10^5	1.69×10^7	1.79×10^7

Table 5: Morfessor cost results for North Sámi. $\alpha = 1.0$

the pruning phase of each iteration, the subwords in the current lexicon are sorted in ascending order according to the estimated change in the cost function if the subword is removed from the lexicon. Subwords consisting of a single character are always kept, to retain the ability to represent an open vocabulary without OOV issues. The list is then pruned according to one of three available pruning criteria:⁵

1. (α -weighted) MDL pruning,
2. MDL with automatic tuning of α for lexicon size,
3. lexicon size with omitted prior or pretuned α .

In (α -weighted) Minimum Description Length (MDL) pruning, subwords are pruned until the estimated cost starts rising, or until the pruning quota for the iteration is reached, whichever comes first.

A subword lexicon of a predetermined size can be used as pruning criterion in two different ways. If the desired α is known in advance, or if the prior is omitted, subwords are pruned until the desired lexicon size is reached, or until the pruning quota for the iteration is reached, whichever comes first.

To reach a subword lexicon of a predetermined size while using the Morfessor prior, the new *automatic tuning* procedure can be applied. For each subword, the estimated change in prior and likelihood are computed separately. These allow computing the value of α that would cause the removal of each subword to be cost neutral, i.e. the value that would cause MDL pruning to terminate at that subword. For subwords with the same sign for both the change in prior and likelihood, no such threshold α can be computed: if the removal decreases both costs the subword will always be removed, and if it increases both costs it will always be kept. Sorting the list of subwords according to the estimated threshold α including the always kept subwords allows automatically tuning α so that a sub-

⁵MDL with or without automatic tuning is not compatible with omitting the prior.

	α	FS	Pre \uparrow	Rec \uparrow	F \uparrow	
EM+P MDL noexp Ψ	0.8	✓	82.9	71.8	77.0	
EM+P MDL nofreqdistr	0.8	✓	83.3	71.4	76.9	
EM+P MDL	0.9	✓	81.9	72.1	76.7	–
Morfessor Baseline	0.8	✓	85.0	68.5	75.9	–
Morfessor Baseline	0.7		83.8	69.4	75.9	\sim B
EM+P MDL	0.6		79.0	72.8	75.8	
SentencePiece 50k	–		75.9	61.9	68.2	

Table 6: Boundary Precision (Pre), Recall (Rec) and F_1 -score (F) results for English. \sim **E** indicates **not** significantly different (two-sided Wilcoxon signed-rank test, $p < 0.05$, zero splitting) from the bolded EM+Prune method, and \sim **B** from the bolded Baseline.

	α	FS	Pre \uparrow	Rec \uparrow	F \uparrow	
EM+P MDL	0.035	✓	72.0	55.8	62.9	–
EM+P MDL nofreqdistr	0.02	✓	68.7	58.0	62.9	\sim E
EM+P MDL noexp Ψ	0.02	✓	68.4	57.9	62.7	\sim E
EM+P MDL	0.015		66.7	58.5	62.3	\sim E
Morfessor Baseline	0.02	✓	62.3	58.2	60.2	–
SentencePiece 50k	–		75.7	49.3	59.7	\sim B
Morfessor Baseline	0.02		62.0	57.6	59.7	

Table 7: Boundary Precision (Pre), Recall (Rec) and F_1 -score (F) results for Finnish.

word lexicon of exactly the desired size is retained after MDL pruning. The automatic tuning is repeated before the pruning phase of each iteration, as retraining the parameters affects the estimates.

3.4. Sampling of Segmentations

Morfessor EM+Prune can be used in subword regularization (Kudo, 2018), a denoising-based regularization method for neural NLP systems. Alternative segmentations can be sampled from the full data distribution using Forward-filtering backward-sampling algorithm (Scott, 2002) or approximatively but more efficiently from an n -best list.

3.5. SentencePiece as a Special Case of Morfessor EM+Prune

Table 1 contains a comparison between all four methods discussed in this work. To recover SentencePiece, Morfessor EM+Prune should be run with the following settings: The prior should be omitted entirely, leaving only the likelihood

$$\hat{\theta} = \arg \min_{\theta} \{-\log P(\mathcal{D} | \theta)\} \quad (7)$$

As the tuning parameter α is no longer needed when the prior is omitted, the pruning criterion can be set to a predetermined lexicon size, without automatic tuning of α . Morfessor by default uses type-based training; to use frequency information, count dampening should be turned off. The seed lexicon should be constructed without using forced splitting. The EM+Viterbi-prune training scheme should be used, with Bayesian EM turned on.

	α	Pre \uparrow	Rec \uparrow	F \uparrow	
EM+P MDL keep-redundant	0.3	87.8	58.7	70.4	
EM+P MDL noexp Ψ	0.4	87.6	58.1	69.9	
EM+P MDL nofreqdistr	0.3	86.4	58.2	69.6	\sim E
EM+P MDL	0.2	84.8	58.7	69.4	–
Morfessor Baseline	0.2	78.2	58.4	66.9	–
SentencePiece 12k	–	75.2	60.0	66.8	\sim B

Table 8: Boundary Precision (Pre), Recall (Rec) and F_1 -score (F) results for Turkish.

	α	FS	Pre \uparrow	Rec \uparrow	F \uparrow	
Morfessor Baseline	1.4		75.7	60.7	67.4	\sim E –
EM+P MDL nofreqdistr	1.0	✓	73.7	61.8	67.2	\sim B
Morfessor Baseline	1.2	✓	75.7	60.4	67.2	\sim E \sim B
EM+P MDL	1.3	✓	73.0	62.1	67.1	– \sim B
EM+P MDL	1.2		72.8	62.0	66.9	
EM+P MDL noexp Ψ	0.4	✓	66.5	65.9	66.2	
SentencePiece 64k	–		65.3	61.3	63.3	

Table 9: Boundary Precision (Pre), Recall (Rec) and F_1 -score (F) results for North Sámi.

4. Experimental Setup

English, Finnish and Turkish data are from the *Morpho Challenge 2010* data set (Kurimo et al., 2010a; Kurimo et al., 2010b). The training sets contain ca 878k, 2.9M and 617k word types, respectively. As test sets we use the union of the 10 official test set samples. For North Sámi, we use a list of ca 691k word types extracted from *Den samiske tekstbanken* corpus (Sametinget, 2004). and the 796 word type test set from version 2 of the data set collected by (Grönroos et al., 2015; Grönroos et al., 2016).

In most experiments we use a grid search with a development set to find a suitable value for α . The exception is experiments using autotuning or lexicon size criterion, and experiments using SentencePiece. We use type-based training (dampening counts to 1) with all Morfessor methods.

For English, we force splits before and after hyphens, and before apostrophes, e.g. “*women’s-rights*” is force split into “*women +’s +- +rights*”. For Finnish, we force splits before and after hyphens, and after colons. For North Sámi, we force splits before and after colons. For Turkish, the Morpho Challenge data is preprocessed in a way that makes force splitting ineffectual.

4.1. Evaluation

The ability of the training algorithm to find parameters minimizing the Morfessor cost is evaluated by using the trained model to segment the training data, and loading the resulting segmentation as if it was a Morfessor Baseline model. We observe both unweighted prior and likelihood, and their α -weighted sum.

The closeness to linguistic segmentation is evaluated by comparison with annotated morph boundaries using *boundary precision*, *boundary recall*, and *boundary F_1 -score* (Virpioja et al., 2011). The boundary F_1 -score (F-score for short) equals the harmonic mean of precision (the percentage of correctly assigned boundaries with respect to all assigned boundaries) and recall

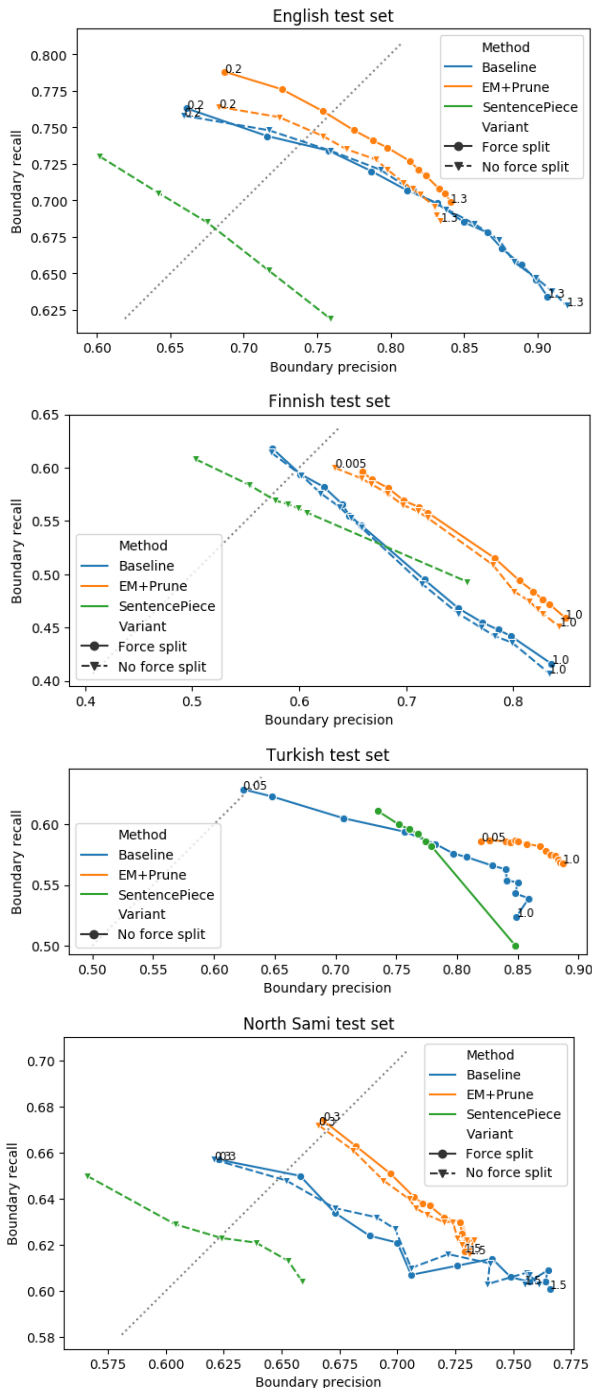


Figure 2: Boundary Precision–Recall curve at different tuning points, The smallest and largest α -values are labeled.

(the percentage of correctly assigned boundaries with respect to the reference boundaries). Precision and recall are calculated using macro-averages over the word types in the test set. In the case that a word has more than one annotated segmentation, we take the one that gives the highest score.

4.2. Error Analysis

We perform an error analysis, with the purpose of gaining more insight into the ability of the methods to model particular aspects of morphology. We follow

the procedure used by Ruokolainen et al. (2016). It is based on a categorization of morphs into the categories PREFIX, STEM, and SUFFIX. The category labels are derived from the original morphological analysis labels in the English and Finnish gold standards, and directly correspond to the annotation scheme used in the North Sámi test set.

We first divide errors into two kinds, *over-segmentation* and *under-segmentation*. Over-segmentation occurs when a boundary is incorrectly assigned within a morph segment. In under-segmentation, the a correct morph boundary is omitted from the generated segmentation. We further divide the errors by the morph category in which the over-segmentation occurs, and the two morph categories surrounding the omitted boundary in under-segmentation.

5. Results

Figure 1 compares the cost components of the Morfessor model across different α parameters. The lowest costs for the mid-range settings are obtained for the EM+Prune algorithm, but for larger lexicons, the Baseline algorithm copes better. As expected, using forced splits at certain characters increase the costs, and the increase is larger than between the training algorithms. As Turkish preprocessing causes the results to be unaffected by the forced splits, we only report results without them.

Tables 2 to 5 show the Morfessor cost of the segmented training data for particular α values. Again, the proposed Morfessor EM+Prune reaches a lower Morfessor cost than Morfessor Baseline. Using the lateen-EM has only minimal effect to the costs, decreasing the total cost slightly for English and increasing for the other languages. Turkish results include the “keep-redundant” setting discussed below in more detail.

Figure 2 shows the Precision–Recall curves for the primary systems, for all four languages. While increasing the Morfessor cost, forced splitting improves BPR. Tables 6 to 9 show test set Boundary Precision, Recall and F_1 -score (BPR) results at the optimal tuning point (selected using a development set) for each model, for English, Finnish, Turkish and North Sámi, respectively⁶. The default Morfessor EM+Prune configuration (“soft” EM, full prior, forcesplit) significantly outperforms Morfessor Baseline w.r.t. the F-score for all languages except North Sámi, for which there is no significant difference between the methods.

Morfessor EM+Prune is less responsive to tuning than Morfessor Baseline. This is visible in the shorter lines in Figures 1 and 2, although the tuning parameter takes values from the same range. In particular, EM+Prune can not easily be tuned to produce very large lexicons.

⁶Note that SentencePiece is not designed for aiming towards a linguistic morpheme segmentation. Neither does it attempt to minimize the Morfessor cost. Therefore, SentencePiece is included in the evaluations for context, not as a baseline method.

Language	Model	Over-segmentation					Under-segmentation						
		STM ↓	SUF ↓	PRE ↓	UNKNOWN ↓	PRE/TOT ↑	STM-SUF ↓	STM-STM ↓	SUF-SUF ↓	SUF-STM ↓	PRE-STM ↓	UNKNOWN ↓	REC/TOT ↑
eng	Characters	71.05	11.82	1.66	0.33	15.13	0.00	0.00	0.00	0.00	0.00		100.00
eng	Words	0.00	0.00	0.00	0.00	100.00	55.07	5.90	8.56	0.14	4.38		23.06
eng	SentencePiece 38k	17.60	10.25	0.18	0.24	71.74	26.40	2.48	2.74	0.05	2.78		65.26
eng	Morfessor Baseline	10.17	2.32	0.03	0.07	87.42	22.46	2.10	4.75	0.04	1.65		67.37
eng	EM+Prune MDL	15.46	2.75	0.05	0.13	81.61	19.93	1.82	4.32	0.04	1.46		70.84
fin	Characters	65.23	13.80	0.67	0.57	19.73	0.00	0.00	0.00	0.00	0.00	0.00	100.00
fin	Words	0.00	0.00	0.00	0.00	100.00	49.19	17.16	21.76	4.84	0.96	0.58	4.09
fin	SentencePiece 13k	35.11	3.71	0.08	0.41	60.69	25.96	1.45	16.18	0.35	0.08	0.16	55.81
fin	Morfessor Baseline	34.75	2.82	0.03	0.38	62.02	24.57	0.86	16.31	0.15	0.04	0.20	57.63
fin	EM+Prune MDL	29.34	2.20	0.03	0.26	68.18	24.68	0.90	15.95	0.29	0.05	0.19	57.60
sme	Characters	81.44	6.80			11.76	0.00	0.00	0.00	0.00			100.00
sme	Words	0.00	0.00			100.00	52.92	13.15	4.43	0.61			28.64
sme	SentencePiece 64k	30.10	4.52			65.38	31.35	3.96	3.09	0.20			61.40
sme	Morfessor Baseline	23.27	3.02			73.71	33.16	2.22	3.40	0.10			60.99
sme	EM+Prune MDL	23.35	4.41			72.25	30.48	3.10	3.23	0.17			62.84

Table 10: Error analysis for English (eng, $\alpha = 0.9$), Finnish (fin, $\alpha = 0.02$), and North Sámi (sme, $\alpha = 1.0$). All results without forcesplit. Over-segmentation and under-segmentation errors reduce precision and recall, respectively.

Pre-pruning of redundant substrings gives mixed results. For Turkish, both Morfessor cost and BPR are degraded by the pre-pruning, but for the other three languages the pre-pruning is beneficial or neutral. When tuning α to very high values (less segmentation), pre-pruning of redundant substrings improves the sensitivity to tuning. The same effect may also be achievable by using a larger seed lexicon. We perform most of our experiments with pre-pruning turned on.

To see the effect of pre-pruning on the seed lexicon, we count the number of subwords that are used in the gold standard segmentations, but not included in seed lexicons of various sizes. Taking Finnish as an example, we see 203 subword types missing from a 1 million substring seed lexicon without pre-pruning. Turning on pre-pruning decreases the number of missing types to 120. To reach the same number without using pre-pruning, a much larger seed lexicon of 1.7M substrings must be used.

Omitting the frequency distribution appears to have little effect on Morfessor cost and BPR. Turning off Bayesian EM (noexp Ψ) results in a less compact lexicon resulting in higher prior cost, but improves BPR for two languages: English and Turkish.

Table 10 contains the error analysis for English, Finnish and North Sámi. For English and North Sámi, EM+Prune results in less under-segmentation but worse over-segmentation. For Finnish these results are reversed. However, the suffixes are often better modeled, as shown by lower under-segmentation on SUF-SUF (all languages) and STM-SUF (English and North Sámi).

6. Conclusion

We propose Morfessor EM+Prune, a new training algorithm for Morfessor Baseline. EM+Prune reduces search error during training, resulting in models with lower Morfessor costs. Lower costs also lead to improved accuracy when segmentation output is compared to linguistic morphological segmentation.

We compare Morfessor EM+Prune to three previously published segmentation methods applying unigram language models. We find that using the Morfessor prior is beneficial when the reference is linguistic morphological segmentation.

In this work we focused on model cost and linguistic segmentation. In future work the performance of Morfessor EM+Prune in applications will be evaluated. Also, a new frequency distribution prior, which is theoretically better motivated or has desirable properties, could be formulated.

7. Acknowledgements

This study has been supported by the MeMAD project, funded by the European Union’s Horizon 2020 research and innovation programme (grant agreement № 780069), and the FoTran project, funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement № 771113) Computer resources within the Aalto University School of Science “Science-IT” project were used.

8. Bibliographical References

- Arivazhagan, N., Bapna, A., Firat, O., Lepikhin, D., Johnson, M., Krikun, M., Chen, M. X., Cao, Y., Foster, G., Cherry, C., Macherey, W., Chen, Z., and Wu, Y. (2019). Massively multilingual neural machine translation in the wild: Findings and challenges. *CoRR*, abs/1907.05019.
- Cotterell, R., Kirov, C., Sylak-Glassman, J., Walther, G., Vylomova, E., Xia, P., Faruqui, M., Kübler, S., Yarowsky, D., Eisner, J., and Hulden, M. (2017). CoNLL-SIGMORPHON 2017 shared task: Universal morphological reinflection in 52 languages. In *Proceedings of the CoNLL SIGMORPHON 2017 Shared Task: Universal Morphological Reinflection*, pages 1–30, Vancouver, August. Association for Computational Linguistics.
- Creutz, M. and Lagus, K. (2002). Unsupervised discovery of morphemes. In *ACL-02 Workshop on Morphological and Phonological Learning*, volume 6 of *MPL '02*, pages 21–30, Philadelphia, Pennsylvania, USA, July. Association for Computational Linguistics.
- Creutz, M. and Lagus, K. (2007). Unsupervised models for morpheme segmentation and morphology learning. *ACM Transactions on Speech and Language Processing*, 4(1), January.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B (Methodological)*, 39(1):1–38.
- Gage, P. (1994). A new algorithm for data compression. *C Users Journal*, 12(2):23–38, February.
- Grönroos, S.-A., Hiovain, K., Smit, P., Rauhala, I. E., Jokinen, P. K., Kurimo, M., and Virpioja, S. (2016). Low-resource active learning of morphological segmentation. *Northern European Journal of Language Technology*.
- Grönroos, S.-A., Jokinen, K., Hiovain, K., Kurimo, M., and Virpioja, S. (2015). Low-resource active learning of North Sámi morphological segmentation. In *Proceedings of 1st International Workshop in Computational Linguistics for Uralic Languages*, pages 20–33.
- Hammarström, H. and Borin, L. (2011). Unsupervised learning of morphology. *Computational Linguistics*, 37(2):309–350, June.
- Kohonen, O., Virpioja, S., and Lagus, K. (2010). Semi-supervised learning of concatenative morphology. In *Proceedings of the 11th Meeting of the ACL Special Interest Group on Computational Morphology and Phonology*, pages 78–86, Uppsala, Sweden, July. Association for Computational Linguistics.
- Kudo, T. and Richardson, J. (2018). SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium, November. Association for Computational Linguistics.
- Kudo, T. (2018). Subword regularization: Improving neural network translation models with multiple subword candidates. *arXiv:1804.10959 [cs]*, April. arXiv: 1804.10959.
- Kurimo, M., Virpioja, S., Turunen, V., and Lagus, K. (2010a). Morpho challenge 2005-2010: Evaluations and results. In Jeffrey Heinz, et al., editors, *Proceedings of the 11th Meeting of the ACL Special Interest Group on Computational Morphology and Phonology*, pages 87–95, Uppsala, Sweden, July. Association for Computational Linguistics.
- Kurimo, M., Virpioja, S., and Turunen, V. T. (2010b). Overview and results of Morpho Challenge 2010. In *Proceedings of the Morpho Challenge 2010 Workshop*, pages 7–24, Espoo, Finland, September. Aalto University School of Science and Technology, Department of Information and Computer Science.
- Liang, P. and Klein, D. (2007). Structured Bayesian nonparametric models with variational inference (tutorial). In *Association for Computational Linguistics (ACL)*.
- Rissanen, J. (1989). *Stochastic Complexity in Statistical Inquiry*, volume 15. World Scientific Series in Computer Science, Singapore.
- Ruokolainen, T., Kohonen, O., Sirts, K., Grönroos, S.-A., Kurimo, M., and Virpioja, S. (2016). A comparative study on minimally supervised morphological segmentation. *Computational Linguistics*.
- Scott, S. L. (2002). Bayesian methods for hidden markov models: Recursive computing in the 21st century. *Journal of the American Statistical Association*, 97(457):337–351.
- Sennrich, R., Haddow, B., and Birch, A. (2015). Neural machine translation of rare words with subword units. In *ACL16*, August. arXiv: 1508.07909.
- Smit, P., Virpioja, S., Kurimo, M., et al. (2017). Improved subword modeling for WFST-based speech recognition. In *INTERSPEECH 2017–18th Annual Conference of the International Speech Communication Association*.
- Spitkovsky, V. I., Alshawi, H., and Jurafsky, D. (2011). Lateen EM: Unsupervised training with multiple objectives, applied to dependency grammar induction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1269–1280. Association for Computational Linguistics.
- Varjokallio, M., Kurimo, M., and Virpioja, S. (2013). Learning a subword vocabulary based on unigram likelihood. In *Proc. 2013 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pages 7–12, Olomouc, Czech Republic, December. IEEE.
- Virpioja, S., Turunen, V. T., Spiegler, S., Kohonen, O., and Kurimo, M. (2011). Empirical comparison of evaluation methods for unsupervised learning of morphology. *Traitement Automatique des Langues*, 52(2):45–90.
- Virpioja, S., Smit, P., Grönroos, S.-A., and Kurimo, M. (2013). Morfessor 2.0: Python implementation and

- extensions for morfessor baseline. Report 25/2013 in Aalto University publication series SCIENCE + TECHNOLOGY, Department of Signal Processing and Acoustics, Aalto University, Helsinki, Finland.
- Virpioja, S. (2012). *Learning Constructions of Natural Language: Statistical Models and Evaluations*. PhD Thesis, Aalto University, Espoo, Finland, December.
- Viterbi, A. J. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2):260–269.

9. Language Resource References

- Grönroos, Stig-Arne and Hiovain, Katri and Smit, Peter and Rauhala, Ilona Erika and Jokinen, Päivi Kristiina and Kurimo, Mikko and Virpioja, Sami. (2015). *North Sámi active learning morphological segmentation annotations*. Aalto University, 2.0.
- Kudo, Taku and Richardson, John. (2018). *Sentence-Piece*. Taku Kudo.
- Kurimo, Mikko and Virpioja, Sami and Turunen, Ville T. (2010). *Morpho Challenge 2010 dataset*. Aalto University.
- Sametinget. (2004). *Den samiske tekstbanken*. UiT Norgga árkálaš universitehta.
- Virpioja, Sami and Smit, Peter and Grönroos, Stig-Arne and Kurimo, Mikko. (2013). *Morfessor 2.0*. Aalto University, 2.0.6.