A Naive Theory of Affixation and an Algorithm for Extraction

Harald Hammarström
Dept. of Computing Science
Chalmers University of Technology
412 96, Gothenburg Sweden
harald2@cs.chalmers.se

Abstract

We present a novel approach to the unsupervised detection of affixes, that is, to extract a set of salient prefixes and suffixes from an unlabeled corpus of a language. The underlying theory makes no assumptions on whether the language uses a lot of morphology or not, whether it is prefixing or suffixing, or whether affixes are long or short. It does however make the assumption that 1. salient affixes have to be frequent, i.e occur much more often than random segments of the same length, and that 2. words essentially are variable length sequences of random characters, e.g a character should not occur in far too many words than random without a reason, such as being part of a very frequent affix. The affix extraction algorithm uses only information from fluctuation of frequencies, runs in linear time, and is free from thresholds and untransparent iterations. We demonstrate the usefulness of the approach with example case studies on typologically distant languages.

1 Introduction

The problem at hand can be described as follows:

Input : An unlabeled corpus of an arbitrary natural language

Output : A (possibly ranked) set of prefixes and suffixes corresponding to true prefixes and suf-

fixes in the linguistic sense, i.e well-segmented and with grammatical meaning, for the language in question.

Restrictions : We consider only concatenative morphology and assume that the corpus comes already segmented on the word level.

The theory and practice of the problem is relevant or even essential in fields such as child language acquisition, information retrieval and, of course, the fuller scope of computational morphology and its further layers of application (e.g Machine Translation).

The reasons for attacking this problem in an unsupervised manner include advantages in elegance, economy of time and money (no annotated resources required), and the fact that the same technology may be used on new languages.

An outline of the paper is as follows: we start with some notation and basic definitions, with which we describe the theory that is intended to model the essential behaviour of affixation in natural languages. Then we describe in detail and with examples the thinking behind the affix extraction algorithm, which actually requires only a few lines to define mathematically. Next, we present and discuss some experimental results on typologically different languages. The paper then finishes with a brief but comprehensive characterization of related work and its differences to our work. At the very end we state the most important conclusions and ideas on future components of unsupervised morphological analysis.

2 A Naive Theory of Affixation

Notation and definitions:

- $w, s, b, x, y, \dots \in \Sigma^*$: lowercase-letter variables range over strings of some alphabet Σ and are variously called words, segments, strings, etc.
- $s \triangleleft w$: s is a terminal segment of the word w i.e there exists a (possibly empty) string x such that $w = xs$
- $W, S, \dots \subseteq \Sigma^*$: capital-letter variables range over sets of words/strings/segments
- $f_W(s) = |\{w \in W | s \triangleleft w\}|$: the number of words in W with terminal segment s
- $S_W = \{s | s \triangleleft w \in W\}$: all terminal segments of the words in W
- $|\cdot|$: is overloaded to denote both the length of a string and the cardinality of a set

Assume we have two sets of random strings over some alphabet Σ :

- Bases $B = \{b_1, b_2, \dots, b_m\}$
- Suffixes $S = \{s_1, s_2, \dots, s_n\}$

Such that:

Arbitrary Character Assumption (ACA) Each character $c \in \Sigma$ should be equally likely in any word-position for any member of B or S .

Note that B and S need not be of the same cardinality and that any string, including the empty string, could end up belonging to both B and S . They need neither to be sampled from the same distribution; pace the requirement, the distributions from which B and S are drawn may differ in how much probability mass is given to strings of different lengths. For instance, it would not be violation if B were drawn from a distribution favouring strings of length, say, 42 and S from a distribution with a strong bias for short strings.

Next, build a set of affixed words $W \subseteq \{bs | b \in B, s \in S\}$, that is, a large set whose members are concatenations of the form bs for $b \in B, s \in S$, such that:

Frequent Flyer Assumption (FFA) : The members of S are frequent. Formally: Given any $s \in S$: $f_W(s) \gg f_W(x)$ for all x such that 1. $|x| = |s|$; and 2. not $x \triangleleft s'$ for all $s' \in S$.

In other words, if we call $s \in S$ a *true suffix* and we call x an *arbitrary segment* if it neither a true suffix nor the terminal segment of a true suffix, then any true suffix should have much higher frequency than an arbitrary segment of the same length.

One may legitimately ask to what extent words of real natural languages fit the construction model of W , with the strong ACA and FFA assumptions, outlined above. For instance, even though natural languages often aren't written phonemically, it is not hard to come up with languages that have phonotactic constraints on what may appear at the beginning or end of a word, e.g, Spanish **st-* may not begin a word and yields *est-* instead. Another violation of ACA is that (presumably all (Ladefoged, 2005)) languages disallow or disprefer a consonant vs. a vowel conditioned by the vowel/consonant status of its predecessor. However, if a certain element occurs with *less* frequency than random (the best example would be click consonants which, in some languages e.g Eastern !Xõo (Traill, 1994), occur only initially), this is not a problem to the theory.

As for FFA, we may have breaches such as Biblical Aramaic (Rosenthal, 1995) where an old $-\bar{a}$ element appears on virtually everywhere on nouns, making it very frequent, but no longer has any synchronic meaning. Also, one can doubt the requirement that an affix should need to be frequent; for instance, the Classical Greek inflectional (lacking synchronic internal segmentation) alternative medial 3p. pl. aorist imperative ending $-\sigma\theta\omega\nu$ (Blomqvist and Jastrup, 1998), is not common at all.

Just how realistic the assumptions are is an empirical question, whose answer must be judged by experiments on the relevant languages. In the absense of fully annotated annotated test sets for diverse languages, and since the author does not have access to the Hutmegs/CELEX gold standard sets for Finnish and English (Creutz and Lindén, 2004), we can only give some guidelining experimental data.

ACA On a New Testament corpus of Basque (Leizarraga, 1571) we computed the probability of a character appearing in the initial, sec-

Positions	Distance
$ p_1 - p_2 $	0.47
$ p_1 - p_3 $	0.36
$ p_1 - p_4 $	0.37
$ p_2 - p_3 $	0.34
$ p_2 - p_4 $	0.23
$ p_3 - p_4 $	0.18

Table 1: Difference between character distributions according to word position.

ond, third or fourth position of the word. Since Basque is entirely suffixing, if it complied to ACA, we’d expect the distributions to be similar. However, if we look at the difference of the distributions in terms of variation distance between two probability distributions ($||p - q|| = \frac{1}{2} \sum_x |p(x) - q(x)|$), it shows that they differ considerably – especially the initial position proves more special (see table 1).

FFA As for the FFA, we checked a corpus of bible portions of Warlpiri (Yal, 1968 2001). This was chosen because it is one of the few languages known to the author where data was available and which has a decent amount of frequent suffixes which are also long, e.g case affixes are typically bisyllabic phonologically and five-ish characters long orthographically. Since the orthography used marked segmentation, it was easy to compute FFA statistics on the words as removed from segmentation marking. Comparing with the lists in (Nash, 1980, Ch. 2) it turns out that FFA is remarkably stable for all grammatical suffixes occurring in the outermost layer. There are however the expected kind of breaches; e.g a tense suffix *-ku* combined with a last vowel *-u* which is frequent in some frequent preceding affixes making the terminal segment *-uku* more frequent than some genuine three-letter suffixes.

The language known to the author which has shown the most systematic discord with the FFA is Haitian Creole (also in bible corpus experiments (Hai, 2003 1999)). Haitian creole has very little morphology of its own but owes the lion’s share of it’s words from French. French derivational morphemes abound in

these words, e.g *-syon*, which have been carefully shown by (Lefebvre, 2004) not to be productive in Haitian Creole. Thus, the little morphology there is in Haitian creole is very difficult to get at without also getting the French relics.

3 An Algorithm for Affix Extraction

The key question is, if words in natural languages are constructed as W explained above, can we recover the segmentation? That is, can we find B and S , given only W ? The answer is yes, we can partially decide this. To be more specific, we can compute a score Z such that $Z(x) > Z(y)$ if $x \in S_W$ and $y \notin S_W$. In general, the converse need not hold, i.e if both $x, y \in S_W$, or both $x, y \notin S_W$, then it may still be that $Z(x) > Z(y)$. This is equivalent to constructing a ranked list of all possible segments, where the true members of S_W appear at the top, and somewhere down the list the junk, i.e non-members of S_W , start appearing and fill up the rest of the list. Thus, it is not said *where* on the list the true-affixes/junk border begins, just that there is a consistent such border.

Now, how should this list be computed? Given the FFA, it’s tempting to look at frequencies alone, i.e just go through all words and make a list of all segments, ranking them by frequency? This won’t do it because 1. it doesn’t compensate between segments of different length; naturally, short segments will be more frequent than long ones, solely by virtue of their shortness 2. it overcounts ill-segmented true affixes, e.g *-ng* will invariably get a higher (or equal) count than *-ing*. What we will do is a modification of this strategy, because 1. can easily be amended by subtracting estimated prior frequencies (under ACA) and there is a clever way of tackling 2. Note that, to amend 2., when going through w and each $s \triangleleft w$, it would be nice if we could count s only when it is well-segmented in w . We are given only W so this information is not available to us, but, the FFA assumption let’s us make a local guess of it.

We shall illustrate the idea with an example of an evolving frequency curve of a word “playing” and its segmentations “playing”, “aying”, “ying”, “ing”, “ng”, “g” (W being the set of words from an English bible corpus (Eng, 1977)). Figure 1 shows a

Figure 1: The observed $f_W(s)$ and expected $e_W(s)$ frequency for $s \triangleleft w = \text{playing}$.

frequency curve $f_W(s)$ and its expected frequency curve $e_W(s)$. The expected frequency of a suffix s doesn't depend on the actual characters of s and is defined as:

$$e_W(s) = |W| \cdot \frac{1}{r^{|s|}}$$

Where r is the size of the alphabet under the assumption that its characters are uniformly distributed. We don't simply use 26 in the case of lowercase English since not all characters are equally frequent. Instead we estimate the size of a would-be uniform distribution from the entropy of the distribution of the characters in W . This gives $r \approx 18.98$ for English and other languages with a similar writing practice.

Next, define the adjusted frequency as the difference between the observed frequency and the expected frequency:

$$f'_W(s) = f_W(s) - e_W(s)$$

It is the slope of this curve that predicts the presence of a good split. Figure 2 shows the appearance of this curve again exemplified by "playing".

After these examples, we are ready to define the segmentation score of a suffix relative to a word $Z : S_W \times W \rightarrow \mathbf{Q}$:

$$Z_W(s, w) = \begin{cases} 0 & \text{if not } s \triangleleft w \\ \frac{f'(s_i) - f'(s_{i-1})}{|f'(s_{i-1})|} & \text{if } s = s_i(w) \text{ for some } i \end{cases}$$

Table 2 shows the evolution of exact values from the running example.

To move from a Z -score for a segment that is relative to a word we simply sum over all words to get

Figure 2: The slope of the $f'_W(s)$ curve for $s \triangleleft w = \text{playing}$.

Input: A text corpus C

Step 1. Extract the set of words W from C (thus all contextual and word-frequency information is discarded)

Step 2. Calculate $Z_W(s, w)$ for each $w \in W$ and $s \triangleleft w$

Step 3. Accumulate $Z_W(s) = \sum_{w \in W} Z(s, w)$

Table 3: Summary of affix-extraction algorithm.

the final score $Z : S_W \rightarrow \mathbf{Q}$:

$$Z_W(s) = \sum_{w \in W} Z(s, w) \quad (1)$$

To be extra clear, the FFA assumption is "exploited" in two ways. On the one hand, frequent affixes get many opportunities to get a score (which could, however, be negative) in the final sum over $w \in W$. On the other hand, the frequency is what make up the appearance of the slope that predicts the segmentation point.

The final Z -score in equation 1 is the one that purports to have the property that $Z(x) > Z(y)$ if $x \in S_W$ and $y \notin S_W$ – at least if purged (see below). A summary of the algorithm described in this section is displayed in table 3.

The time-complexity bounding factor is the number of suffixes, i.e the cardinality of S_W , which is linear (in the size of the input) if words are bounded in length by a constant and quadratic in the (really) worst case if not.

s	<i>playing</i>	<i>laying</i>	<i>aying</i>	<i>ying</i>	<i>ing</i>	<i>ng</i>	<i>g</i>
$f(s)$	1	4	12	40	706	729	756
$e_W(s)$	0.00	0.00	0.00	0.10	1.90	36.0	684
$f(s) - e_W(s)$	0.99	3.99	11.9	39.8	704	692	71.0
$Z(s, \text{"playing"})$	0.00	2.99	1.99	2.32	16.6	-0.0	-0.8

Table 2: Exact values of frequency curves and scores from the running “playing” example.

1028682.0	<i>ing</i>	111264.0	<i>ling</i>	1288402.4	<i>i</i>	33756.55	<i>ler</i>
594208.0	<i>ed</i>	111132.0	<i>ent</i>	151056.9	<i>er</i>	29816.53	<i>da</i>
371145.0	<i>s</i>	109725.0	<i>ating</i>	142552.6	<i>in</i>	29404.49	<i>di</i>
337464.0	<i>'s</i>	109125.0	<i>ate</i>	141603.3	<i>im</i>	28337.89	<i>le</i>
326250.0	<i>ation</i>	108228.0	<i>an</i>	134403.2	<i>en</i>	26580.41	<i>dan</i>
289536.0	<i>es</i>	97020.0	<i>ies</i>	130794.5	<i>e</i>	26373.54	<i>r</i>
238853.5	<i>e</i>	94560.0	<i>ts</i>	127352.0	<i>an</i>	24183.99	<i>ti</i>
222256.0	<i>er</i>	81648.0	<i>ically</i>	113482.6	<i>a</i>	22527.26	<i>un</i>
191889.0	<i>ers</i>	81504.0	<i>ment</i>	82581.95	<i>ya</i>	21388.71	<i>iniz</i>
172800.0	<i>ting</i>	78669.0	<i>led</i>	78447.74	<i>ar</i>	20993.87	<i>sin</i>
168288.0	<i>ly</i>	77900.0	<i>ering</i>	76353.77	<i>ak</i>	20117.60	<i>ik</i>
159408.0	<i>ations</i>	74976.0	<i>er's</i>	68730.00	<i>n</i>	18612.14	<i>li</i>
143775.0	<i>ted</i>	73988.0	<i>y</i>	64761.37	<i>ir</i>	18316.45	<i>ek</i>
130960.0	<i>able</i>	...	...	53021.67	<i>la</i>	...	...
116352.0	<i>ated</i>	-26137.9	<i>l</i>	47218.78	<i>ini</i>	-38091.8	<i>t</i>
113364.0	<i>al</i>	-38620.6	<i>m</i>	44858.18	<i>lar</i>	-240917.5	<i>l</i>
113280.0	<i>ness</i>	-78757.3	<i>a</i>	37229.14	<i>iz</i>	-284460.1	<i>s</i>

Table 4: Top 30 and bottom 3 extracted suffixes for English. 47178 unique words yielded a total of 154407 ranked suffixes.

4 Experimental Results

For a regular English 1 million token newspaper corpus we get the top 30 plus bottom 3 suffixes as shown in table 4.

English has little affixation compared to e.g Turkish which is at the opposite end of the typological scale (Dryer, 2005). The corresponding results for Turkish on a bible corpus (Tur, 1988) is shown in table 5.

The results largely speak for themselves but some comments are in order. As is easily seen from the lists, some suffixes are suffixes of each other so one could *purge* the list in some way to get only the most “competitive” suffixes. One purging strategy would be to remove x from the list if there is a z

Table 5: Top 30 and bottom 3 extracted suffixes for Turkish. 56881 unique words yielded a total of 175937 ranked suffixes.

such that $x = yx$ and $Z(z) > Z(x)$ (this would remove e.g *-ting* if *-ing* is above it on the list). A more sophisticated purging method is the following, which does slightly more. First, for a word $w \in W$ define its best segmentation as: $Segment(w) = \argmax_{s \in S_W} Z(s)$. Then purge by keeping only those suffixes which are the best parse for at least one word: $S'_W = \{s \in S_W | \exists w Segment(w) = s\}$.

Such purging kicks out the bulk of “junk” suffixes. Table 4 shows the numbers for English, Turkish and the virtually affixless Maori (Bauer et al., 1993). It should be noted that “junk” suffixes still remain after purging – typically common stem-final characters – and that there is no simple relation between the number of suffixes left after purging and the amount of morphology of the language in question. Otherwise we would have expected the morphology-less Maori to be left with no, or 28-ish,

Language	Corpus	Tokens	$ W $	$ S_W $	$ S'_W $
Maori	(Mao, 1996)	1101665	8354	23007	78
English	(Eng, 1977)	917634	12999	39845	63
Turkish	(Tur, 1988)	574592	56881	175937	122

Table 6: Figures for different languages on the effects on the size of the suffix list after purging.

suffixes or at least less than English.

A good sign is that the purged list and its order seems to be largely independent of corpus size (as long as the corpus is not very small) but we do get some significant differences between bible English and newspaper English.

We have chosen to illustrate using affixes but the method readily generalizes to prefixes as well and even prefixes and suffixes at the same time. As an example of this, we show top-10 purged prefix-suffix scores in the same table also for some typologically differing languages in table 7. Again, we use bible corpora for cross-language comparability (Swedish (Swe, 1917) and Swahili (Swa, 1953)). The scores have been normalized in each language to allow cross-language comparison – which, judging from the table, seems meaningful. Swahili is an exclusively prefixing language but verbs tend to end in *-a* (whose status as a morpheme is the linguistic sense can be doubted), whereas Swedish is suffixing, although some prefixes are or were productive in word-formation.

A full discussion of further aspects such as a more informed segmentation of words, peeling of multiple suffix layers and purging of unwanted affixes requires, is beyond the scope of this paper.

5 Related Work

For reasons of space we cannot cite and comment every relevant paper even in the narrow view of highly unsupervised extraction of affixes from raw corpus data, but we will cite enough to cover each line of research. The vast fields of word segmentation for speech recognition or for languages which do not mark word boundaries will not be covered. In our view, segmentation into lexical units is a different problem than that of affix extraction since the frequencies of lexical items are different, i.e occur

Swedish		English		Swahili	
<i>för-</i>	0.097	<i>-eth</i>	0.086	<i>-a</i>	0.100
<i>-en</i>	0.086	<i>-ing</i>	0.080	<i>wa-</i>	0.095
<i>-na</i>	0.036	<i>-ed</i>	0.063	<i>ali-</i>	0.065
<i>-ade</i>	0.035	<i>-est</i>	0.036	<i>nita-</i>	0.059
<i>-a</i>	0.034	<i>-th</i>	0.035	<i>aka-</i>	0.049
<i>-ar</i>	0.033	<i>-es</i>	0.034	<i>ni-</i>	0.046
<i>-er</i>	0.033	<i>-s</i>	0.033	<i>ku-</i>	0.044
<i>-as</i>	0.032	<i>-ah</i>	0.026	<i>ata-</i>	0.042
<i>-s</i>	0.031	<i>-er</i>	0.026	<i>ha-</i>	0.032
<i>-de</i>	0.031	<i>-ation</i>	0.019	<i>a-</i>	0.031
...	...	...	...	...	...

Table 7: Comparative figures for prefix vs. suffix detection. The high placement of English *-eth* and *-ah* are due to the fact that the bible version used has drinketh, sitteth etc and a lot of personal names in *-ah*.

much more sparsely. Results from this area which have been carried over or overlap with affix detection will however be taken into account. A lot of the papers cited have a wider scope and are still useful even though they are criticized here for having a non-optimal affix detection component.

Many authors trace their approaches back to two early papers by Zellig Harris (Harris, 1955; Harris, 1970) which count *letter successor varieties*. The basic procedure is to ask how many different phonemes occur (in various utterances e.g a corpus) after the first n phonemes of some test utterance and predict that segmentation(s) occur where the number of successors reaches a peak. For example, if we have *play*, *played*, *playing*, *player*, *players*, *playground* and we wish to test where to segment *plays*, the successor count for the prefix *pla* would be 1 because only *y* occurs after whereas the number of successors of *play* peak at three (i.e $\{e, i, g\}$). Although the heuristic has had some success it was shown (in various interpretations) as early as (Hafer and Weiss, 1974) that it is not really sound – even for English. A slightly better method is to compile a set of words into a *trie* and predict boundaries at nodes with high activity (e.g (Johnson and Martin, 2003; Schone and Jurafsky, 2001; Kazakov and Manandhar, 2001) and earlier papers by the same authors), but this not sound either as non-morphemic short common character sequences also show significant branching.

The algorithm in this paper differs significantly from the Harris-inspired varieties. First, we do not record the number of phonemes/character of a given prefix/suffix but the total number of continuations. In the example above, that would be the set $\{ed, ing, er, ers, ground\}$ rather than the three-member set of continuing phonemes/characters. Secondly, segmentation of a given word is not the immediate objective and what amounts to identification of the end of a lexical (thus generally low-frequency) item is not within the direct reach of the model. Thirdly, and most importantly, the algorithm in this paper looks at the *slope* of the frequency curve not at peaks in absolute frequency.

A different approach, sometimes used in complement of other sources of information, is to select *aligned pairs* (or sets) of strings that share a long character sequence (work includes (Jacquemin, 1997; Yarowsky and Wicentowski, 2000; Baroni et al., 2002; Clark, 2001)). A notable advantage is that one is not restricted to concatenative morphology.

Many publications (Ćavar et al., 2004; Brent et al., 1995; Goldsmith et al., 2001; Déjean, 1998; Snover et al., 2002; Argamon et al., 2004; Goldsmith, 2001; Creutz and Lagus, 2005; Neuvel and Fulop, 2002; Baroni, 2003; Gaussier, 1999; Sharma et al., 2002; Wicentowski, 2002; Oliver, 2004), and various other works by the same authors, describe strategies that use frequencies, probabilities, and optimization criteria, often Minimum Description Length (MDL), in various combinations. So far, all these are unsatisfactory on two main accounts; on the theoretical side, they still owe an explanation of why compression or MDL should give birth to segmentations coinciding with morphemes as linguistically defined. On the experimental side, thresholds, supervised/developed parameters and selective input still cloud the success of reported results, which, in any case, aren't wide enough to sustain some too rash language independence claims.

To be more specific, some MDL approaches aim to minimize the description of the set of words in the input corpus, some to describe all tokens in the corpus, but, none aims to minimize, what one would otherwise expect, the set of possible words in the language. More importantly, none of the reviewed works allow any variation in the description language ("model") during the minimization

search. Therefore they should be more properly labeled "weighting schemes" and it's an open question whether their yields correspond to linguistic analysis. Given an input corpus and a traditional linguistic analysis, it is trivial to show that it is possible to decrease description length (according to the given schemes) by stepping away from linguistic analysis. Moreover, various forms of codebook compression, such as Lempel-Ziv compression, yield shorter description but without any known linguistic relevance at all. What is clear, however, apart from whether it is theoretically motivated, is that MDL approaches are *useful*.

A systematic test of segmentation algorithms over many different types of languages has yet to be published. For three reasons, it will not be undertaken here either. First, as e.g. already Manning (1998) notes for sandhi phenomena, it is far from clear what the gold standard should be (even though we may agree or disagree on some familiar European languages). Secondly, segmentation algorithms may have different purposes and it might not make good sense to study segmentation in isolation from induction of paradigms. Lastly, and most importantly, all of the reviewed techniques (Wicentowski, 2004; Wicentowski, 2002; Snover et al., 2002; Baroni et al., 2002; Andreev, 1965; Ćavar et al., 2004; Snover and Brent, 2003; Snover and Brent, 2001; Snover, 2002; Schone and Jurafsky, 2001; Jacquemin, 1997; Goldsmith and Hu, 2004; Sharma et al., 2002; Clark, 2001; Kazakov and Manandhar, 1998; Déjean, 1998; Oliver, 2004; Creutz and Lagus, 2002; Creutz and Lagus, 2003; Creutz and Lagus, 2004; Hirsimäki et al., 2003; Creutz and Lagus, 2005; Argamon et al., 2004; Gaussier, 1999; Lehmann, 1973; Langer, 1991; Flenner, 1995; Klenk and Langer, 1989; Goldsmith, 2001; Goldsmith, 2000; Hu et al., 2005b; Hu et al., 2005a; Brent et al., 1995), as they are described, have threshold-parameters of some sort, explicitly claim **not** to work well for an open set of languages, or require noise-free all-form input (Albright, 2002; Manning, 1998; Borin, 1991). Therefore it is not possible to even design a fair test.

In any event, we wish to appeal to the merits of developing a theory in parallel with experimentation – as opposed to only ad hoc result chasing. If we have a theory and we don't get the results we want,

we may scrutinize the assumptions behind the theory in order to modify or reject it (understanding why we did so). Without a theory there's no telling what to do or how to interpret intermediate numbers in a long series of calculations.

6 Conclusion

We have presented a new theory of affixation and a parameter-less efficient algorithm for collecting affixes from raw corpus data of an arbitrary language. Depending on one's purposes with it, a cut-off point for the collected list is still missing, or at least, we do not consider that matter here. The results are very promising and competitive but at present we lack formal evaluation in this respect. Future directions also include a more specialized look into the relation between affix-segmentation and paradigmatic variation and further exploits into layered morphology.

7 Acknowledgements

The author has benefited much from discussions with Bengt Nordström.

References

- Adam C. Albright. 2002. *The Identification of Bases in Morphological Paradigms*. Ph.D. thesis, University of California at Los Angeles.
- Nikolai Dmitrievich Andreev, editor. 1965. *Statistiko-kombinatornoe modelirovanie iazykov*. Akademia Nauk SSSR, Moskva.
- Shlomo Argamon, Navot Akiva, Amihod Amit, and Oren Kapah. 2004. Efficient unsupervised recursive word segmentation using minimum description length. In *COLING-04, 22-29 August 2004, Geneva, Switzerland*.
- Marco Baroni, Johannes Matiassek, and Harald Trost. 2002. Unsupervised discovery of morphologically related words based on orthographic and semantic similarity. In *Proceedings of the Workshop on Morphological and Phonological Learning of ACL/SIGPHON-2002*, pages 48–57.
- Marco Baroni. 2003. Distribution-driven morpheme discovery: A computational/experimental study. *Yearbook of Morphology*, pages 213–248.
- Winifred Bauer, William Parker, and Te Kareongawai Evans. 1993. *Maori*. Descriptive Grammars. Routledge, London & New York.
- Jerker Blomqvist and Poul Ole Jastrup. 1998. *Grekisk Grammatik: Graesk grammatik*. Akademisk Forlag, København, 2 edition.
- Lars Borin. 1991. *The Automatic Induction of Morphological Regularities*. Ph.D. thesis, University of Uppsala.
- Michael R. Brent, S. Murthy, and A. Lundberg. 1995. Discovering morphemic suffixes: A case study in minimum description length induction. In *Fifth International Workshop on Artificial Intelligence and Statistics, Ft. Lauderdale, Florida*.
- Damir Čavar, Joshua Herring, Toshikazu Ikuta, Paul Rodrigues, and Giancarlo Schrementi. 2004. On induction of morphology grammars and its role in bootstrapping. In Gerhard Jäger, Paola Monachesi, Gerald Penn, and Shuly Wintner, editors, *Proceedings of Formal Grammar 2004*, pages 47–62.
- Alexander Clark. 2001. Learning morphology with pair hidden markov models. In *ACL (Companion Volume)*, pages 55–60.
- Mathias Creutz and Krista Lagus. 2002. Unsupervised discovery of morphemes. In *Proceedings of the 6th Workshop of the ACL Special Interest Group in Computational Phonology (SIGPHON), Philadelphia, July 2002*, pages 21–30. Association for Computational Linguistics.
- Mathias Creutz and Krista Lagus. 2003. Unsupervised discovery of morphemes. In *Proceedings of the 6th Workshop of the ACL Special Interest Group in Computational Phonology (SIGPHON), Philadelphia, July 2002*, pages 21–30. Association for Computational Linguistics.
- Mathias Creutz and Krista Lagus. 2004. Induction of a simple morphology for highly-inflecting languages. In *Proceedings of the 7th Meeting of the ACL Special Interest Group in Computational Phonology (SIGPHON)*, pages 43–51. Barcelona.
- Mathias Creutz and Krista Lagus. 2005. Unsupervised morpheme segmentation and morphology induction from text corpora using morfessor 1.0. Technical report, Publications in Computer and Information Science, Report A81, Helsinki University of Technology, March.
- Mathias Creutz and Krister Lindén. 2004. Morpheme segmentation gold standards for finnish and english. publications in computer and information science, report a77, helsinki university of technology. Technical report, Publications in Computer and Information Science, Report A77, Helsinki University of Technology, October.

- Hervé Déjean. 1998. *Concepts et algorithmes pour la découverte des structures formelles des langues*. Ph.D. thesis, Université de Caen Basse Normandie.
- Matthew S. Dryer. 2005. Prefixing versus suffixing in inflectional morphology. In Bernard Comrie, Matthew S. Dryer, David Gil, and Martin Haspelmath, editors, *World Atlas of Language Structures*, pages 110–113. Oxford University Press.
1977. The holy bible, containing the old and new testaments and the apocrypha in the authorized king james version. Thomas Nelson, Nashville, New York.
- Gudrun Flenner. 1995. Quantitative morphsegmentierung im spanischen auf phonologischer basis. *Sprache und Datenverarbeitung*, 19(2):63–78. Also cited as: *Computatio Linguae II*, 1994, pp. 1994 as well as *Sprache und Datenverarbeitung* 19(2):31–62, 1994.
- Éric Gaussier. 1999. Unsupervised learning of derivational morphology from inflectional lexicons. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics (ACL-1999)*. Association for Computational Linguistics, Philadelphia.
- John Goldsmith and Yu Hu. 2004. From signatures to finite state automata. Technical report TR-2005-05, Department of Computer Science, University of Chicago.
- John Goldsmith, Derrick Higgins, and Svetlana Soglasnova. 2001. Automatic language-specific stemming in information retrieval. In Carol Peters, editor, *Cross-Language Information Retrieval and Evaluation: Proceedings of the CLEF 2000 Workshop*, Lecture Notes in Computer Science, pages 273–283. Springer-Verlag, Berlin.
- John Goldsmith. 2000. Linguistica: An automatic morphological analyzer. In A. Okrent and J. Boyle, editors, *Proceedings from the Main Session of the Chicago Linguistic Society's thirty-sixth Meeting*.
- John Goldsmith. 2001. Unsupervised learning of the morphology of natural language. *Computational Linguistics*, 27(2):153–198.
- Margaret A. Hafer and Stephen F. Weiss. 1974. Word segmentation by letter successor varieties. *Information and Storage Retrieval*, 10:371–385.
- 2003 [1999]. Bib la. American Bible Society.
- Zellig S. Harris. 1955. From phoneme to morpheme. *Language*, 31(2):190–222.
- Zellig S. Harris. 1970. Morpheme boundaries within words: Report on a computer test. In Zellig S. Harris, editor, *Papers in Structural and Transformational Linguistics*, volume 1 of *Formal Linguistics Series*, pages 68–77. D. Reidel, Dordrecht.
- Teemu Hirsimäki, Mathias Creutz, Vesa Siivola, and Mikko Kurimo. 2003. Unlimited vocabulary speech recognition based on morphs discovered in an unsupervised manner. In *Proceedings of Eurospeech 2003*, Geneva, pages 2293–2996. Geneva, Switzerland.
- Yu Hu, Irina Matveeva, John Goldsmith, and Colin Sprague. 2005a. Refining the SED heuristic for morpheme discovery: Another look at Swahili. In *Proceedings of the Workshop on Psychocomputational Models of Human Language Acquisition*, pages 28–35, Ann Arbor, Michigan, June. Association for Computational Linguistics.
- Yu Hu, Irina Matveeva, John Goldsmith, and Colin Sprague. 2005b. Using morphology and syntax together in unsupervised learning. In *Proceedings of the Workshop on Psychocomputational Models of Human Language Acquisition*, pages 20–27, Ann Arbor, Michigan, June. Association for Computational Linguistics.
- Christian Jacquemin. 1997. Guessing morphology from terms and corpora. In *Proceedings, 20th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '97)*, Philadelphia, PA.
- Howard Johnson and Joel Martin. 2003. Unsupervised learning of morphology for english and inuktitut. In *HLT-NAACL 2003, Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, May 27 - June 1, Edmonton, Canada*, volume Companion Volume - Short papers.
- Dimitar Kazakov and Suresh Manandhar. 1998. A hybrid approach to word segmentation. In C. D. Page, editor, *Proceedings of the 8th International Workshop on Inductive Logic Programming (ILP-98) in Madison, Wisconsin, USA*, volume 1446 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin.
- Dimitar Kazakov and Suresh Manandhar. 2001. Unsupervised learning of word segmentation rules with genetic algorithms and inductive logic programming. *Machine Learning*, 43:121–162.
- Ursula Klenk and Hagen Langer. 1989. Morphological segmentation without a lexicon. *Literary and Linguistic Computing*, 4(4):247–253.
- Peter Ladefoged. 2005. *Vowels and Consonants*. Blackwell, Oxford, 2 edition.
- Hagen Langer. 1991. *Ein automatisches Morphsegmentierungsverfahren für deutsche Wortformen*. Ph.D. thesis, Georg-August-Universität zu Göttingen.
- Claire Lefebvre. 2004. *Issues in the study of Pidgin and Creole languages*, volume 70 of *Studies in Language Companion Series*. John Benjamins, Amsterdam.

- Hubert Lehmann. 1973. *Linguistische Modellbildung und Methodologie*. Max Niemeyer Verlag, Tübingen. Pp. 71-76 and 88-93.
- Joanes Leizarraga. 1571. Iesus krist gure iaunaren testamentu berria. Pierre Hautin, Inprimizale, Roxellan. [NT only].
- Christopher D. Manning. 1998. The segmentation problem in morphology learning. In Jill Burstein and Claudia Leacock, editors, *Proceedings of the Joint Conference on New Methods in Language Processing and Computational Language Learning*, pages 299–305. Association for Computational Linguistics, Somerset, New Jersey.
1996. Maori bible. The British & Foreign Bible Society, London, England.
- David G. Nash. 1980. *Topics in Warlpiri Grammar*. Ph.D. thesis, Massachusetts Institute of Technology.
- Sylvain Neuvel and Sean A. Fulop. 2002. Unsupervised learning of morphology without morphemes. In *Workshop on Morphological and Phonological Learning at Association for Computational Linguistics 40th Anniversary Meeting (ACL-02)*, July 6-12, pages 9–15. ACL Publications.
- A. Oliver. 2004. *Adquisició d'informació lèxica i morfosintàctica a partir de corpus sense anotar: aplicació al rus i al croat*. Ph.D. thesis, Universitat de Barcelona.
- Franz Rosenthal. 1995. *A grammar of biblical Aramaic*, volume 5 of *Porta linguarum Orientalium*. Harrassowitz, Wiesbaden, 6 edition.
- Patrick Schone and Daniel Jurafsky. 2001. Knowledge-free induction of inflectional morphologies. In *Proceedings of the North American Chapter of the Association for Computational Linguistics*, Pittsburgh, PA, 2001.
- Utpal Sharma, Jugal Kalita, and Rajib Das. 2002. Unsupervised learning of morphology for building lexicon for a highly inflectional language. In *Proceedings of the 6th Workshop of the ACL Special Interest Group in Computational Phonology (SIGPHON)*, Philadelphia, July 2002, pages 1–10. Association for Computational Linguistics.
- Matthew G. Snover and Michael R. Brent. 2001. A bayesian model for morpheme and paradigm identification. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics (ACL-2001)*, pages 482–490. Morgan Kaufmann Publishers.
- Matthew G. Snover and Michael R. Brent. 2003. A probabilistic model for learning concatenative morphology. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 1513–1520. MIT Press, Cambridge, MA.
- Matthew G. Snover, Gaja E. Jarosz, and Michael R. Brent. 2002. Unsupervised learning of morphology using a novel directed search algorithm: Taking the first step. In *Workshop on Morphological and Phonological Learning at Association for Computational Linguistics 40th Anniversary Meeting (ACL-02)*, July 6-12. ACL Publications.
- Matthew G. Snover. 2002. An unsupervised knowledge free algorithm for the learning of morphology in natural languages. Master's thesis, Department of Computer Science, Washington University.
1953. Maandiko matakatifu ya mungu yaitwaya biblia, yaani agano la kale na agano jipya, katika lugha ya kiswahili. British and Foreign Bible Society, London, England.
1917. Gamla och nya testamentet: de kanoniska böckerna. Norstedt, Stockholm.
- Anthony Traill. 1994. *A !Xóǀ Dictionary*, volume 9 of *Quellen zur Khoisan-Forschung/Research in Khoisan Studies*. Rüdiger Köppe Verlag, Köln.
1988. Turkish bible. American Bible Society, Tulsa, Oklahoma.
- Richard Wicentowski. 2002. *Modeling and Learning Multilingual Inflectional Morphology in a Minimally Supervised Framework*. Ph.D. thesis, Johns Hopkins University, Baltimore, MD.
- Richard Wicentowski. 2004. Multilingual noise-robust supervised morphological analysis using the word-frame model. In *Proceedings of the ACL Special Interest Group on Computational Phonology (SIGPHON)*, pages 70–77.
- 1968–2001. Bible: selections in warlpiri. Summer Institute of Linguistics. Document 0650 of the Aboriginal Studies Electronic Data Archive (ASEDA), AIATSIS (Australian Institute of Aboriginal and Torres Strait Islander Studies), Canberra.
- David Yarowsky and Richard Wicentowski. 2000. Minimally supervised morphological analysis by multimodal alignment. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics (ACL-2000)*, pages 207–216.