

Learning Phonotactics Using ILP

Stasinos Konstantopoulos
Alfa-Informatica, Rijksuniversiteit Groningen
konstant@let.rug.nl

3rd October 2001

Abstract

This paper describes experiments on learning Dutch phonotactic rules using Inductive Logic Programming, a machine learning approach based on the notion of inverting resolution. Different ways of approaching the problem are experimented with, and compared against each other as well as with related work on the task. Further research is outlined.

1 Introduction

The *Phonotactics* of a given language is the set of rules that identifies what sequences of phonemes constitute a possible word in that language. The problem can be broken down to the *syllable structure* (i.e. what sequences of phonemes constitute a possible syllable) and the processes that take place at the syllable boundaries (e.g. assimilation).

Previous work on the syllable structure of Dutch includes hand-crafted models [9, 1], but also machine-learning approaches such as abduction [8] and neural networks ([7] and [6, ch. 4]).

This paper describes experiments on the task of constructing from examples a model of Dutch monosyllabic words. The reason for restricting the domain is to avoid the added complexity of handling syllable boundary phonological processes. Furthermore by not using polysyllables no prior commitment is made to any one particular syllabification (and thus syllable structure) theory.

2 Inductive Logic Programming and Aleph

Inductive Logic Programming (ILP) is a machine learning discipline. It is logic programming in the sense that the target concept to be learned is a logic program, i.e. a set of Horn clauses. It is inductive because the core operator used is that of *induction*.

Induction can be seen as the inverse of deduction. For example from the clauses ‘All humans die’ and ‘Socrates is human’ the fact that ‘Socrates will die’ can be deduced. Inversely, induction uses background knowledge (e.g. ‘Socrates is human’) and a set of observations (training data) (e.g. ‘Socrates died’) to search for a hypothesis that, in conjunction with the background knowledge, can deduce the data. In more formal terms, given a logic program B modelling the background knowledge and a set of ground terms D representing the training data, ILP constructs a logic program H , such that $B \wedge H \models D$.

If the deductive operator used is resolution (as defined by Robinson [4]), then the inductive operator necessary to solve the equation above is the inverse resolution operator, as defined by Muggleton and De Raedt [2].

Aleph [5] is an ILP system implementing the Progol algorithm [3]. This algorithm allows for single-predicate learning only, without background theory revision or predicate invention. It incrementally constructs the clauses of a single-predicate hypothesis that describes the data, by iterating through the following basic algorithm:

- Saturation: pick a positive example from the training data and construct the most specific, non-ground clause that entails it. This is done by repeated application of the inverse resolution operator on the example, until all its ground terms have been replaced by variables which are appropriately bound by the body of the clause, so that the original ground positive and only that is covered by the bottom clause. This minimally generalised clause is called the *bottom clause*.
- Reduction: search between the maximally general, empty-bodied clause and the maximally specific bottom clause for a ‘good’ clause. The space between the empty clause and the bottom clause is partially ordered by θ -subsumption, and the search proceeds along the lattice defined by this partial ordering. The ‘goodness’ of each clause encountered along the search path is evaluated by an *evaluation function*.
- Cover removal: add the new clause to the theory and remove all examples covered by it from the dataset.
- Repeat until all positive examples are covered.

The *evaluation function* quantifies the usefulness of each clause constructed during the search and should be chosen so that it balances between overfitting (i.e. covering the data too tightly and making no generalisations that will yield coverage over unseen data) and overgeneralising (i.e. covering the data too loosely and accepting too many negatives.) It can be simple coverage (number of positive minus number of negative examples covered by

the clause) or the Bayes probability that the clause is correct given the data (for learning from positive data only) or any function of the numbers of examples covered and the length of the clause (for implementing bias towards shorter clauses.)

Syntactic bias is applied during reduction to either prune paths that are already known to not yield a winning clause, or to enforce restrictions on the constructed theory, for example, conformance to a theoretical framework.

3 Setting up the Experiments

As a starting point, a rough template matching all syllables is assumed. This template is $\mathcal{C}_3\mathcal{V}\mathcal{C}_5$, where \mathcal{C}_n represents any consonant cluster of length up to n and \mathcal{V} any vowel or diphthong. The problem can now be reformulated as a single-predicate learning task where the target theory is one of acceptable prefixes to a given vowel and partial consonant cluster. The rules for prevocalic and postvocalic affixing are induced in two separate learning sessions.

The training data is derived from 5095 monosyllabic words taken from the Dutch section of the CELEX Lexical Database, with an additional 597 reserved for evaluation.

The positive examples are constructed by breaking the phonetic transcriptions down to three parts: a prevocalic and a postvocalic consonant cluster (consisting of zero or more consonants) and a vowel or diphthong. The consonant clusters are treated as ‘affixes’ to the vowel, so that syllables are constructed by repeatedly affixing consonants, if the *context* (the vowel and the pre- or post-vocalic material that has been already affixed) allows it. So, for example, from the word /ma:kt/ the following positives would be generated:

```

prefix( m, [], [a,:] ).    suffix( k, [], [:,a] ).
prefix( ^, [m], [a,:] ).  suffix( t, [k], [:,a] ).
                           suffix( ^, [tk], [:,a] ).

```

Note that the context lists in suffix rules is reversed, so that the two processes are exactly symmetrical and can use the same background predicates.

The caret, ^, is used to mark the beginning and end of a word. The reason that the affix termination needs to be explicitly licensed is so that it is not assumed by the experiment’s setup that all partial sub-affixes of a valid affix are necessarily valid as well.

In Dutch, for example, a monosyllable with a short vowel has to be closed, which means that the null suffix is not valid. The end-of-word mark allows for this to be expressible as a theory that does not have the following clause: `suffix(^, [], [V])`.

The positives are, then, all the prefixes and suffixes that must be al-

lowed in context, so that all the monosyllables in the training data can be constructed: 11067 and 10969 instances of 1428 and 1653 unique examples, respectively.

The negative data is randomly generated words that match the $\mathcal{C}_3\mathcal{V}\mathcal{C}_5$ template and do not appear as positives. The random generator is designed so that the number of examples at each affix length are balanced, in order to avoid having the large numbers of long, uninteresting sequences overwhelm the shorter, more interesting ones.

The negative data is also split into evaluation and training data, and the negative examples are derived from the training negative data by the following deductive algorithm:

1. For each example, find the maximal substring that is provable by the positive `prefix/3` and `suffix/3` clauses in training data. So, for example, for `/mtratk/` it would be `trat` and for `/mlat/`, `lat`.
2. Choose the clause that should be a negative example, so that this word is not accepted by the target theory. Pick the inner-most one on each side, i.e. the one immediately applicable to the maximal substring computed above. For `/mlat/` that would be `suffix(m, [l], [a])`. `/mtratk/`, however, could be negative because either `prefix(m, [tr], [a])` or `suffix(k, [t], [a])` are unacceptable. In such cases, pick one at random. This is bound to introduce false negatives, but no alternative that does not presuppose at least part of the solution could be devised.
3. Iterate, until enough negative examples have been generated to disprove all the words in the negative training data.

Since the problem is, in effect, that of identifying the sets of consonants that may be prefixed or suffixed to a partially constructed monosyllable, the clauses of the target predicate must have a means of referring to various subsets of \mathcal{C} and \mathcal{V} in a meaningful and intuitive way. This is achieved by defining a (possibly hierarchical,) linguistically motivated partitioning of \mathcal{C} and \mathcal{V} . Each partition can then be referred to as a feature-value pair, for example `LAB+` to denote the set of the labials or `VOIC+` for the set of voiced consonants. Intersections of these basic sets can then be easily referred to by feature-value vectors; the intersection, for example, of the labials and the voiced consonants (i.e. the voiced labials) is the feature-value vector `[VOIC+,LAB+]`.

The *background knowledge* is, as seen in section 2, playing a decisive role in the quality of the constructed theory, by implementing the theoretic framework to which the search for a solution will be confined. In more concrete terms, the background predicates are the building blocks that will be used for the construction of the hypothesis' clauses and they must be defining all the relations necessary to discover an interesting hypothesis.

For the purposes of this task, they have been defined as relations between individual phones and feature values, e.g. `labial(m,+)` or `voiced(m,+)`. Feature-value vectors can then be expressed as conjunctions like, for example, `labial(C,+) ∧ voiced(C,+)` to mean the voiced labials.

Except for the linguistic features predicates, the background knowledge also contained the `head/2` and `rest/2` list access predicates. This approach was chosen over direct list access with the `nth/3` predicate, as bias towards rules with more local context dependencies.

The experiments described in sections 3.1, 3.2 and 3.3 below, were conducted with background knowledge that encodes increasingly more information about Dutch phonology as well as Dutch phonotactics: for the experiment in 3.1 the learner has access to the way the various symbols are arranged in the IPA table, whereas for the experiment in 3.2 a classification that is sensitive to Dutch phonological processes was chosen. And, finally, in section 3.3 the *sonority level* feature is implemented, which has been proposed with the explicit purpose of solving the problem of Dutch syllable structure.

The quantitative evaluation shown for three experiments was done using the 597 words and the part of the randomly generated negative data that have been reserved for this purpose.

3.1 The IPA segment space

In this experiment the background knowledge reflects the way that the IPA table is organised: the phonetic inventory of Dutch consists of two disjoint spaces, one of consonants and one of vowels, with three and four orthogonal dimensions of differentiation respectively.

The consonant space varies in *place* and *manner of articulation*, and *voicing*. The manner of articulation can be *plosive*, *nasal*, *lateral approximant*, *trill*, *fricative* or *approximant*. The place can be *bilabial*, *alveolar*, *velar*, *labiodental*, *postalveolar* or *palatal*. Voicing can be present or not.

Similarly for vowels, where there are four dimensions: place (front, centre, back) and manner of articulation (open, mid-open, mid-closed, closed), length and roundedness.

The end-of-word mark has no phonological features whatsoever and it does not belong to any of the partitions of either \mathcal{C} or \mathcal{V} .

This schema was implemented as one background predicate per dimension relating each phone with its value along that dimension, for example:

```
manner( plosive, p ).    place( bilabial, p ).
manner( nasal, m ).
```

etc.

The evaluation function used was the *Laplace function* $\frac{P+1}{P+N+2}$, where P and N is the coverage of positive and negative examples, respectively.

Since the randomly generated negatives must also contain false negatives, it cannot be expected that even a good theory will fit it perfectly. In order to avoid overfitting, the learning algorithm was set to only require an accuracy of 85% over the training data.

The resulting hypothesis consisted of 199 prefix and 147 suffix clauses and achieved a recall rate of 99.3% with 89.4% precision. All the false negatives were rejected because they couldn't get their onset licensed, typically because it only appears in a handful of loan words. The /ɕ/ onset necessary to accept 'jeep' and 'junk', for example, was not permitted and so these two words were rejected.

The most generic rules found were:

```
prevoc(A,B,C) :- A= '^'.
prevoc(A, [], C).
```

```
postvoc(A,B,C) :- A= '^'.
postvoc(A, [], C).
```

meaning that (a) the inner-most consonant can be anything, and (b) all sub-prefixes (-suffixes) of a valid prefix (suffix) are also valid.

Other interesting rules include pairs like these two:

```
prevoc(A,B,C) :-
  head(B,D), manner(trill,D), head(C,E), length(short,E),
  manner(closed,E), manner(plosive,A).
prevoc(A,B,C) :-
  head(B,D), manner(trill,D), head(C,E), length(short,E),
  manner(open_mid,E), manner(plosive,A).
```

and

```
prevoc(A,B,C) :-
  head(B,D), manner(approx,D), head(C,E), length(short,E),
  place(front,E), voiced(minus,A).
prevoc(A,B,C) :-
  head(B,D), manner(approx,D), head(C,E), length(short,E),
  place(front,E), manner(plosive,A), place(alveolar,A).
```

that could have been collapsed if a richer feature system would include features like 'closed or mid-open vowel' and 'devoiced consonant or plosive alveolar', respectively. These particular disjunctions might be unintuitive or even impossible to independently motivate, but they do suggest that a redundant feature set might allow for more interesting theories than the minimal, orthogonal one used for this experiment. This is particularly true for a system like Aleph, that performs no predicate invention or background theory revision.

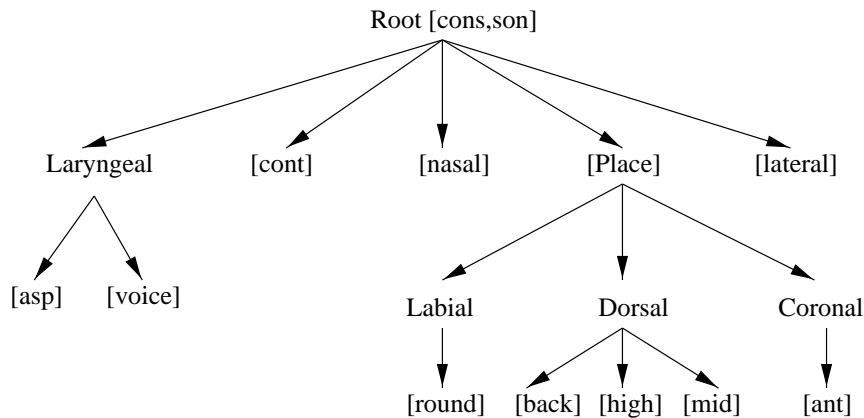


Figure 1: The feature geometry of Dutch

3.2 Feature Classes

For this experiment a richer (but more language-specific) background knowledge was made available to the inductive algorithm, by implementing the feature hierarchy suggested by Booij [1, ch. 2] and replicated in figure 1.

The most generic features are the *major class features* (CONSONANT and SONORANT) that are placed on the root node and divide the segment space in vowels [CONS-,SON+], obstruents [CONS+,SON-] and sonorant consonants [CONS+,SON+]. Since all vowels are sonorous, [CONS-,SON-] is an invalid combination.

The features specifying the continuants, nasals and the lateral /l/ are positioned directly under the root node, with the rest of the features bundled together under two *feature classes*, those of the *laryngeal* and the *place* features. These classes are chosen so that they collect together features that behave as a unit in phonological processes of Dutch. The class of laryngeal features is basically making the voiced-voiceless distinction, while the ASPIRATION feature is only separating /h/ from the rest. The *place* class bundles together three subclasses of place of articulation features, one for each articulator. Furthermore some derived or redundant features such as GLIDE, APPROXIMANT and LIQUID are defined. The vowels do not include the schwa, which is set apart and only specified as SCHWA+.

Using the Laplace evaluation function and this background, the constructed theory consisted of 13 prefix and 93 suffix rules, accepting 94.2% of the test positives and under 7.4% of the test negatives.

Among the rejected positives are loan words ('jeep' and 'junk' once again), but also all the words starting with perfectly Dutch /s/ - obstruent - liquid clusters.

The prefix rule with the widest coverage is:

`prefix(A,B,C) :-`
 `head(C,D), sonorant(D,plu), rest(B, []).`

or, in other words, ‘prefix anything before a single consonant before a nucleus other than the schwa’.

The suffix rules were less strict, with only 3 rejected positives, ‘branche’, ‘dumps’ and ‘krimpst’ (the first two of which are loan words) that failed to suffix /ʃ/, /s/ and /s/ respectively. Some achieve wide coverage (although never quite as wide as that of the prefix rules,) but some are making reference to individual phonemes and are of more restricted application. For example:

`suffix(A,B,C) :-`
 `rest(C,D), head(D,E), rest(B, []), A=t.`

or, ‘suffix a /t/ after exactly one consonant, if the nucleus is a long vowel or a diphthong’.

Of some interest are also the end-of-word marking rules (see in section 3 above about the $\hat{\quad}$ mark), because of the fact that open, short monosyllables are very rare in Dutch (there are four in CELEX: ‘schwa’, ‘ba’, ‘hè’, and ‘joh’). This would suggest that the best way to treat those is as exceptions, and have the general rule disallow open, short monosyllables. What was learned instead was a whole set of 29 rules for suffixing $\hat{\quad}$, the most general of which is:

`postvoc(A,B,C) :-`
 `head(B,t), larynx(t,E), rest(B,F),`
 `head(F,G), larynx(G,E), A= ' ^ '.`

or ‘suffix an end-of-word mark after at least two consonants, if the outer-most one is a /t/ and has the same values for all the features in the LARYNGEAL feature class as the consonant immediately preceding it’.

A final note that needs to be made regarding this experiment, is one regarding its computational complexity. Overlapping and redundant features might be offering the opportunity for more interesting hypotheses, but are also making the search space bigger. The reason is that overlapping features are diminishing the effectiveness of the inverse resolution operator at keeping uninteresting predicates out of the bottom clause: the more background predicates can be used to prove the positive example on which the bottom clause is seeded, the longer the latter will get.

3.3 Sonority Scale

This experiment implements and tests the syllabic structure model suggested by van der Hulst [9, ch. 3]. The Dutch syllable is there analysed as having three prevocalic and 5 postvocalic positions, (some of which may be empty) and constraints are placed on the set of consonants that can occupy each.

phoneme	obstruents	m	n	l	r	glides	vowels
sonority	1	2	2.25	2.5	2.75	3	4

Table 1: The Sonority Scale

The most prominent constraint is the one stipulating a high-to-low *sonority* progression from the nucleus outwards. Each phoneme is assigned a sonority value (table 1) based not only on language-independent features such as it being a SONORANT or an OBSTRUENT, but also because of syllable structure of Dutch itself. Especially the fine tuning done with respect to the sonority values of the nasals and the liquids is explicitly justified on filtering out impossible consonant clusters that would otherwise be predicted by the simpler model. It must, therefore, be noted that the background knowledge for this experiment is not only language-specific, but is also directly aimed at solving the very problem that is being investigated.

In addition to the high-to-low sonority level progression from the nucleus outwards, there are both *filters* and explicit licensing rules. The former are restrictions referring to sonority (e.g. ‘the sonority of the three left-most positions must be smaller than 4’) or other phonological features (e.g. the ‘no voiced obstruents in coda’ filter in p. 92) and are applicable in conjunction to the sonority rule. The latter are typically restricted in scope rules, that take precedence over the sonority-related constraints mentioned so far. The left-most position, for example, may be /s/ or empty, regardless of the contents of the rest of the onset.

Implementing the basic sonority progression rule as well as the most widely-applicable filters and rules¹ yielded impressive compression rates matched with results lying between those of the two previous experiments: 93.1% recall, 83.2% precision.

4 Conclusions and Further Work

The quantitative results from the machine learning experiments presented above are collected in table 2, together with those of Tjong Kim Sang and Nerbonne [8]² and the results from the sonority scale experiment. Those last ones in particular, are listed for comparison’s sake and as the logical end-point of the progression towards more language- and task-specific prior assumptions.

¹Some were left out because they were too lengthy when translated from their fixed-position framework to the affix licensing one used here, and were very specifically fine tuning the theory to individual onsets or codas.

²From experiments on phonetic data in the ‘Experiments without Linguistic Constraints’ section.

	(Tjong 2000)	IPA	Feat. Classes	Sonority
Recall	99.1%	99.3%	94.2%	93.1%
Precision	74.8%	79.8%	92.6%	83.2%
Num. Clauses	577 + 577	145 + 36	13 + 93	3+8

Table 2: Results

The first two columns are directly comparable, because they are both only referring to phonetic primitives with no linguistically sophisticated background knowledge. The fact that the $\mathcal{C}_3\mathcal{V}\mathcal{C}_5$ template assumed in this work is not taken for granted in [8], is compensated in terms of compactness as well as performance. Compactness because the numbers of rules quoted in the first column do not include the 41 extra rules (besides the 1154 prefix and suffix rules) that describe the "basic word" on which the affix rules operate. Performance because the precision given is measured on totally random strings, whereas in this work only strings matching the $\mathcal{C}_3\mathcal{V}\mathcal{C}_5$ template are used.

As can be seen, then, the ILP-constructed rules compare favourably (in both performance and hypothesis compactness) with those constructed by the deductive approach employed in [8].

What can be also seen by comparing the two ILP results with each other, is that the drop in recall between the the second and third column is compensated by higher precision and compression, suggesting a direct correspondence between the quality of the prior knowledge encoded in the background theory and that of the constructed hypothesis.

One interesting follow-up to these experiments would be attempting to expand their domain to that of syllables of multisyllabic words and, eventually, full word-forms. In the interest of keeping the problems of syllabic structure and syllable-boundary phonology apart, a way must be devised to derive from the positive data (i.e. a corpus of Dutch word-forms) examples for a distinct machine learning session for each task.

Furthermore, it would be interesting to carry out the same (or rather the equivalent) experiments on other parts of the CELEX corpus (e.g. English or German) and see to which extend the results-to-background relation follows the same patterns.

On the purely computational aspect of the problem, and as has already been mentioned in section 3.2, including overlapping and redundant features in the background might be interesting, but is also implying a very fast growth of the search space. It would, therefore, be useful to employ the search path pruning facilities of Aleph to avoid directions along the search path that would yield inconsistent or unlikely feature combinations.

References

- [1] Geert Booij. *The Phonology of Dutch*. The Phonology of the World's Languages. Clarendon Press, Oxford, 1995.
- [2] Stephen Muggleton and Luc De Raedt. Inductive Logic Programming: Theory and methods. *Journal of Logic Programming*, 19(20):629–679, 1994. Updated version of technical report CW 178, May 1993, Department of Computing Science, K.U. Leuven.
- [3] Stephen H. Muggleton. Inverse entailment and Progol. *New Generation Computing*, 13:245–286, 1995.
- [4] J. A. Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, 1965.
- [5] Ashwin Srinivasan. *Aleph*. <http://www.comlab.ox.ac.uk/oucl/research/areas/machlearn/Aleph/>, Last update: May 18, 2001.
- [6] Ivilin Stoianov. *Connectionist Lexical Processing*. PhD thesis, Rijksuniversiteit Groningen, 2001.
- [7] Ivilin Stoianov and John Nerbonne. Exploring phonotactics with Simple Recurrent Networks. In van Eynde, Schuurman, and Schelkens, editors, *Proceedings of Computational Linguistics in the Netherlands 1998*, 1999.
- [8] Erik Tjong Kim Sang and John Nerbonne. Learning the logic of simple phonotactics. In James Cussens and Saso Dzeroski, editors, *Learning Language in Logic*, volume 1925 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, 2000.
- [9] Harry van der Hulst. *Syllable Structure and Stress in Dutch*. PhD thesis, Rijksuniversiteit Leiden, 1984.