

# A LEXICAL DATABASE TOOL FOR QUANTITATIVE PHONOLOGICAL RESEARCH

Steven Bird

The University of Edinburgh  
Centre for Cognitive Science  
Edinburgh EH8 9LW, UK

SIL Cameroon  
B.P. 1299  
Yaoundé, Cameroon

`Steven.Bird@ed.ac.uk`

## Abstract

A lexical database tool tailored for phonological research is described. Database fields include transcriptions, glosses and hyperlinks to speech files. Database queries are expressed using HTML forms, and these permit regular expression search on any combination of fields. Regular expressions are passed directly to a Perl CGI program, enabling the full flexibility of Perl extended regular expressions. The regular expression notation is extended to better support phonological searches, such as search for minimal pairs. Search results are presented in the form of HTML or L<sup>A</sup>T<sub>E</sub>X tables, where each cell is either a number (representing frequency) or a designated subset of the fields. Tables have up to four dimensions, with an elegant system for specifying which fragments of which fields should be used for the row/column labels. The tool offers several advantages over traditional methods of analysis: (i) it supports a quantitative method of doing phonological research; (ii) it gives universal access to the same set of informants; (iii) it enables other researchers to hear the original speech data without having to rely on published transcriptions; (iv) it makes the full power of regular expression search available, and search results are full multimedia documents; and (v) it enables the early refutation of false hypotheses, shortening the analysis-hypothesis-test loop. A life-size application to an African tone language (Dschang) is used for exemplification throughout the paper. The database contains 2200 records, each with approximately 15 fields. Running on a PC laptop with a stand-alone web server, the 'Dschang HyperLexicon' has already been used extensively in phonological fieldwork and analysis in Cameroon.

## INTRODUCTION

Initial stages of phonological analysis typically focus on words in isolation, as the phonemic inventory and syllable canon are established. Data is stored as a lexicon, where each word is entered as a transcription accompanied by at least a gloss (so the word can be elicited again) and the major syntactic category. In managing a lexicon, the working phonologist has a variety of computational needs: storage and retrieval; searching and sorting; tabular reports on distributions and contrasts; updates to database and to reports as distinctions are discovered or discarded. In the past the analyst had to do all this computation by hand using index cards kept in shoeboxes. But now many of these particular tasks are automated by software such as the SIL programs Shoebox (Buseman et al., 1996) and Findphone (Bevan, 1995),<sup>1</sup> or using commercial database packages.

Of course, many tasks other than those listed above have already benefitted from (partial) automation.<sup>2</sup> Additionally, it has been shown how a computational inheritance model can be used for structuring lexical information relevant for phonology (Reinhard & Gibbon, 1991). And there is a body of work on the use of finite state devices – closely related to regular expressions – for modelling phonological phenomena (Kaplan & Kay, 1994) and for speech processing (cf. Kornai's

<sup>1</sup>Unlike regular database management systems, these include international and phonetic character sets and user-defined keystrokes for entering them, and a utility to dump a database into an RTF file in a user-defined lexicon format for use in desktop publishing.

<sup>2</sup>For example, see (Ellison, 1992; Lowe & Mazaudon, 1994; Coleman, Dirksen, Hussain & Waals, 1996).

<code>\id</code>	1612	<i>identifier (used for hyperlinks)</i>
<code>\v</code>		<i>validation status</i>
<code>\w</code>	mbhū	<i>orthographic form</i>
<code>\as</code>	#m.bhU#	<i>ascii transcription</i>
<code>\rt</code>	#bhU#	<i>transcription of word root</i>
<code>\t</code>	LDH	<i>tone transcription</i>
<code>\sd</code>	mbhū	<i>southern dialect form</i>
<code>\pg</code>	*bú+à	<i>Proto-Grassfields form</i>
<code>\p</code>	n	<i>part of speech</i>
<code>\pl</code>	me-	<i>plural prefix</i>
<code>\cl</code>	9/6	<i>noun class (singular/plural)</i>
<code>\en</code>	dog	<i>english gloss</i>
<code>\fr</code>	chien	<i>french gloss (used with informants)</i>

Figure 1: Format of Database Records

work with HMMs (Kornai, 1995)). However, computational phonology is yet to provide tools for manipulating lexical and speech data using the full expressive power of the regular expression notation in a way that supports pure phonological research.

This paper describes a lexical database system tailored to the needs of phonological research and exemplified for Dschang, a language of Cameroon. An online lexicon (originally published as Bird & Tadadjeu, 1997), contains records with the format in Figure 1. Only the most important fields are shown.

The user interface is provided by a Web browser. A suite of Perl programs (Wall & Schwartz, 1991) generates the search form in

HTML and processes the query. Regular expressions in the query are passed directly to Perl, enabling the full flexibility of Perl extended regular expressions. A further extension to the notation allows searches for minimal sets, groups of words which are minimally different according to some criterion. Hits are structured into a tabular display and returned as an HTML or L<sup>A</sup>T<sub>E</sub>X document.

In the next section, a sequence of example queries is given to illustrate the format of queries and results, and to demonstrate how a user might interact with the system. A range of more powerful queries are then demonstrated, along with an explanation of the notations for minimal pairs and projections. Next, some implementation details are given, and the component

modules are described in detail. The last two sections describe planned future work and present the conclusions.

## EXAMPLE

This section shows how the system can be used to support phonological analysis. The language data comes from Dschang, a Grassfields Bantu language of Cameroon, and is structured into a lexicon consisting of 2200 records. Suppose we wished to learn about phonotactic constraints in the syllable rhyme. The following sequence of queries were not artificially constructed, but were issued in an actual session with the system in the field, running the Web server in a stand-alone mode. The first query is displayed below.<sup>3</sup>

### Search Attributes:

```

display: count
root: .*([$V])([$C])#
loanwords: exclude
suffixed: include
phrases: exclude
time-limit: 2 minutes
vars: $B = "\.#-"; # boundaries
      $S = "pbtckgcj'"; # stops
      $F = "zsvfZS"; # fricatives
      $O = $S.$F; # obstruents
      $N = "mnN"; # nasals
      $G = "wy"; # glides
      $C = $O.$N.$G."hl"; # cons
      $V = "ieaouEOU@"; # vowels

```

The main attribute of interest is the `root` attribute.<sup>4</sup> The `.*` expression stands for a sequence of zero or more segments. The expressions `$V` and `$C` are variables defined in the `vars` section of the query form. These are strings, but when surrounded with brackets, as in `[$V]` and `[$C]`, they function as wild cards which match a single element from the string. The `#` character is a boundary symbol marking the end of the root. Observe that the `root` attribute contains two parenthesised subexpressions. These will be called *parameters* and have a special role in structuring the search output. This is best demonstrated by way of an example. Consider the table below, which is the result

<sup>3</sup>The display is only a crude approximation to the HTML form. Note that the query form comes with the variables already filled in so that it is not necessary for the user to supply them, although they can be edited. The transcription symbols used in the system have the following interpretation: U=u, @=ə, E=ɛ, O=ɔ, N=ŋ, '=?.

<sup>4</sup>In the following discussion, 'attribute' refers to a line in the query form while 'field' refers to part of a database record.

of the above query. In this table, the row labels are all the segments which matched the variable \$V, while the column labels are just the segments that matched \$C.

**Search Results:**

	p	t	k	'	m	N
i	5		10	24	9	32
U	9			38	1	9
u	14			60	10	39
@			15		41	75
o			31	12		
E	51	14				
a	30	1	46	61		76
O	15	1	12	36		49

There are sufficient gaps in the table to make us wonder if all the segments are actually phonemes. For example, consider o and u, given that they are phonetically very similar ([ɔ] and [u] respectively). We can easily set up o as an allophone of u before k. Only the case of glottal stop needs to be considered. So we revise the form, replacing \$V with just the vowels in question, and replacing the \$C of the coda with apostrophe (for glottal stop). We add a term for the syllable onset and resubmit the query. See Figure 2. This time, several attributes are omitted from the display for brevity.






We can now conclude that o and u are in complementary distribution, except for the five words corresponding to pf and v onsets. But what are these words? We revise the form again, further restricting the search string as follows:

**Search Attributes:**

```
display: speech word gloss
root: .* (pf|v) [ou] ' #
```

The display parameter is set to speech word gloss allowing us to see (and hear) the individual lexical items. The results are shown below.

**Search Results:**

pf	 lepfo' mortar
	 mpfu' blood pact
v	 mvo' space in front of bed
	 avu' remainder
	 levu'té kitchen woodpile

The cells of the output table now contain fragments of the lexical entries. The first part is an icon which, when

clicked, plays the speech file. The second part is a gif of the orthographic form of the word. The third part is the English gloss. Note that the above nouns have different prefixes (e.g. le-, m-, a-). These are noun class prefixes and are not part of the root field. If we had wanted to take prefixes into consideration then the as attribute, containing a transcription of the whole word, could have been used instead.

Listening to the speech files it was found that the syllables pfo' and pfu' sounded exactly the same, as did vo' and vu'. The whole process up to this point had taken less than five minutes. After some quick informant work to recheck the data and hear the native-speaker intuitions, it was clear that the distinction between o and u in closed syllables was subphonemic.

## MORE POWERFUL QUERIES

### Constraining one field and displaying another

In some situations we are not interested in seeing the field which was constrained, but another one instead. The next query displays the tone field for monosyllabic roots, classed into open and closed syllables. Although the root attribute is used in the query, the root field is not actually displayed. (This query makes use of a projection function which maps all consonants onto C and all vowels onto V, as will be explained later.)

**Search Attributes:**

```
display: tone
root: #C+V(C?)# ($CV-proj)
```

The C+ expression denotes a sequence of one or more consonants, while C? denotes an optional coda consonant. By making C? into a parameter (using parentheses) the search results will be presented in a two column table, one column for open syllables (with a null label) and one for closed syllables (labelled C). A minor change to the root attribute, enlarging the scope of the parameter (\#C+(VC?)\#), will produce the more satisfactory column labels V and VC.

### Searching for near-minimal sets

Finding good minimal sets is a heuristic process. No attempt has been made to encode heuristics into the system. Rather, the aim has been to permit flexible interaction between user and system as a collection of minimal sets is refined. To facilitate this process, the regular expression notation is extended slightly.

**Search Attributes:**

```
display: count
root: .*([\$C]+)([ou])'##
axes: flip
```

**Search Results:**

	w	p	pf	b	t	ts	d	c	j	k	g	f	v	s	z	m	n	ŋ	l
u	6	8	1	1	6		1	6	4	5	3	5	2			4	1	1	5
o			1			6							1	1	3				

Figure 2: Query to Probe the Phonemic Status of the O/U Contrast

Recall the way that parameters (parenthesised sub-expressions) allowed output to be structured. One of the parameters will be said to be *in focus*. Syntactically, this is expressed using braces instead of parentheses. Semantically, such a parameter becomes the focus of a search for minimal sets.

Typically, this parameter will contain a list of segments, such as { [ou] }, or an optional segment whose presence is to be contrasted with its absence, such as {h?}.

In order for a minimal set to be found, the parameter in focus must have more than one possible instantiation, while the other parameters remain unchanged. To see how this works, consider the following example. Suppose we wish to identify the minimal pairs for o/u discussed above, but without having to specify glottal stop in the query, as shown in Figure 3. Note this example of a 3D table.

If this was not enough minimal pairs, we could relax the restrictions on the context. For example, if we do not wish to insist on the following consonant being identical across minimal pairs, we can remove the second set of parentheses thus: .\*([\\$C]+)([ou])[\\$C]#. This now gives minimal pairs like legōk *work* and ŋgu' *year*. Observe that the consonant preceding the o/u vowel is fixed across the minimal pair, since this was still parenthesised in the query string.

Usually, it is best for minimal pairs to have similar syntactic distribution. We can add a restriction that all minimal pairs must be drawn from the same syntactic category by making the whole part attribute into a parameter as follows.

**Search Attributes:**

```
display: word gloss
root: .*([\$C]+){[ou]}([\$C])#
```

**Search Results:**

	,
pf	<div style="border: 1px solid black; padding: 5px;"> <p>lepfo' <i>mortar</i></p> <hr/> <p>mpfu' <i>blood pact</i></p> </div>
v	<div style="border: 1px solid black; padding: 5px;"> <p>mvo' <i>space in front of bed</i></p> <hr/> <p>avu' <i>remainder</i></p> <hr/> <p>levu'tí <i>kitchen woodpile</i></p> </div>

Figure 3: Minimal Sets for O/U

**Search Attributes:**

```
display: tone
root: .*([\$C]+){[ou]}[\$C]#
part: (.*)
```

Making the part attribute into a parameter adds an extra dimension to the table of results. We now only see an o/u minimal pair if the other parameters agree. In other words, all minimal pairs that are reported will contain the same consonant cluster before the o/u vowel and will be from the same syntactic category.

**Variables across attributes**

There are occasions where we need to have the same variable appearing in different attributes. For example,

suppose we wanted to check where the southern dialect and the principal dialect have identical vowels:<sup>5</sup>

#### Search Attributes:

```
display: root s_dialect
root: .* (3[$V]+) .*
s_dialect: .*$3.*
```

This query makes use of another syntactic extension to regular expressions. An arbitrary one-digit number which appears immediately inside a parameter allows the parameter to be referred to elsewhere. This means that whichever sequence of vowels matches `[$V]+` in the `root` field must also appear somewhere in the `s_dialect` field.

#### Negative restrictions

The simplest kind of negative restriction is built using the set complement operator (the caret). However this only works for single character complements. A much more powerful negation is available with the `?!` zero-width negative lookahead assertion, available in Perl 5, which I will now discuss.

The next example uses the tone attribute. Dschang is a tone language, and the records in the lexicon include a field containing a tone melody. Tone melodies consist of the characters H (high), L (low), D (downstep) and F (fall). A single tone has the form `D?[HL]F?`, i.e. an optional downstep, followed by H or L, followed by an optional fall. The next example finds all entries starting with a sequence of unlike tones.

#### Search Attributes:

```
root: .* (1[$T]) (?!$1) [$T] .*
vars: $T = D?[HL]F?
```

The `(1[$T])` expression matches any tone and sets the `$1` variable to the tone which was matched. The `(?!$1)` expression requires that whatever follows the first tone is different, and the final `[$T]` insists that this same following material is a tone (rather than being empty, for example).<sup>6</sup>

<sup>5</sup>Roots are virtually all monosyllabic, so there will usually be a unique vowel sequence for the `[$V]+` in the regular expression to match with.

<sup>6</sup>Care must be taken to ensure that the alphabetic encodings of distinct tones are sufficiently different from each other, so that one is not an initial substring of another.

## Projections

I earlier introduced the notion of projections. In fact, the system allows the user to apply an arbitrary manipulation to any attribute before the matching is carried out. Here is the query again, this time with the `$cv-proj` variable filled out.

#### Search Attributes:

```
display: tone
root: #C+V(C?)# ($CV-proj)
vars: $CV-proj = {tr/$C/C/; tr/$V/V/;}
```

This causes the Perl `tr` (transliterate) function to be applied to the `root` attribute before the `#C+V(C?)#` regular expression is matched on this field.

Projections can also be used to simulate second order variables, such as required for place of articulation. Suppose that the language has three places of articulation: L (labial), A (alveolar) and V (velar). We are interested in finding any unassimilated sequences in the data (i.e. adjacent consonants with different places of articulation). The following query does just this. Prior to matching, the segments which have a place of articulation value are projected to that value, again using `tr`. The query expression looks for a sequence of any pair `$P$P`, where `$P` is a second order variable ranging over places of articulation.

#### Search Attributes:

```
display: word
root: .* (5$P) (?!$5) ($P) .* ($P-proj)
vars: $P-proj=tr/pbmtdnkgN/LLLAADV/;
$P = {LAV};
```

Observe that the second `$P` must be different from the first, because of the zero-width negative lookahead assertion `(?!$5)`. This states that immediately to the right of this position one does not find an instance of `$5`, where this variable is the place of articulation found in the first position. The output of the query is a  $3 \times 3$  table showing all words that contain unassimilated consonant sequences.

## SYSTEM OVERVIEW

### Lexicon compiler

The base lexicon is in Shoebox format, in which the fields are not required to be in a fixed order. To save on runtime processing, a preprocessing step is applied to each field. For example, the contents of the `\w`

field, comprising characters from the Cameroon character set, are replaced by a pointer to a graphics file for the word (i.e. a URL referencing a gif).<sup>7</sup> Each record is processed into a single line, where fields occur in a canonical order and a field separator is inserted, and the compiled lexicon is stored as a DBM file for rapid loading.

### The query string

The search attributes in the query form can contain arbitrary Perl V5 regular expressions, along with some extensions introduced above. A CGI program constructs a query string based on the submitted form data. The query string is padded with wild cards for those fields which were not restricted in the query form.

The dimensionality of the output and the axis labels are determined by the appearance of 'parameters' in the search attributes. These parenthesised subexpressions are copied directly into the query string. So, for example, the first query above contained the search expression `.*([V])([C])#` applied to the `root` field. This field occupies fifth position in the compiled version of a record, and so the search string is as follows. The variable `$e` matches any sequence of characters not containing the field separator.

```
$search = /^$e;$e;$e;$e;$e;.*([V])([C])#;
          $e;$e;$e;$e;$e;$e;$e;$e;$e$;
```

### The search loop

Search involves a linear pass over the whole lexicon `%LEX`.<sup>8</sup> The parameters contained in `$search` are tied to the variables `$1 - $4`. These are stored in four associative arrays `$dim1 - $dim4` to be used later as axis labels.

```
foreach $entry (keys %LEX) {
  if ($LEX{$entry} =~ /$search/) {
    $dim1{$1}++;
    $dim2{$2}++;
    $dim3{$3}++;
    $dim4{$4}++;
    $hits{"$1;$2;$3;$4"} .= ";$entry;
  }
}
```

<sup>7</sup>These gifs were generated using  $\LaTeX$  along with the utilities `pstogif` and `giftool`.

<sup>8</sup>Inverting on individual fields was avoided because of the runtime overheads and the fact that this prevents variable instantiation across fields.

Finally, a pointer to the entry is stored in the 4D array `$hits` (appended to any existing hits in that cell.) Here we see that the structuring of the output table using parameters is virtually transparent, with Perl itself doing the necessary housekeeping.

As an example, suppose that the following lexical entry is being considered at the top of the above loop:

```
$entry = 0107
$LEX{$entry} =
  0107; ;;
  #a.kup#;#kup#;LL; ;*k`ub`;n;7/6,8;
  skin, bark;peau,\`ecorce;
```

By matching this against the query string given in our first example we end up matching `.*([V])([C])#` with `#kup#`. This results in `$1=u` and `$2=p`. The entries `$dim1{u}` and `$dim2{p}` are incremented, recording these values for later use in the `$V` and `$C` axes respectively. Finally `$hits("u;p;;")` is updated with the index 0107.

### The display loop

This module cycles through the axis labels that were stored in `$dim1 - $dim4` and combines them to access the `$hits` array. At each level of nesting, code is generated for the HTML or  $\LaTeX$  table output. At the innermost level, the fields selected by the user in the `display` attribute are used to build the current cell.

## FUTURE WORK

A number of extensions to the system are planned. Since `Dschang` is a tone language, it would be particularly valuable to have access to the pitch contours of each word. These will eventually be displayed as small gifs, attached to the lexical entries.

Another extension would be to permit updates to the lexicon through a forms interface. A special instance of the search form could be used to validate existing and new entries, alerting the user to any data which contradicts current hypotheses.

The regular expression notation is sometimes cumbersome and opaque. It would be useful to have a higher level language as well. One possibility is the notation of autosegmental phonology, which can be compiled into finite-state automata (Bird & Ellison, 1994). The graphics capabilities for this could be provided on the client side by a Java program.

A final extension, dependent on developments with HTML itself, would be to provide better support for special characters and user-definable keystrokes for accessing them.

## CONCLUSION

This paper has presented a hypertext lexicon tailored to the practical needs of the phonologist working on large scale data problems. The user accesses the lexicon via a forms interface provided by HTML and a browser. A CGI program processes the query. The user can refine a query during the course of several interactions with the system, finally switching the output to L<sup>A</sup>T<sub>E</sub>X format for direct inclusion of the results in a research paper. An extension to the regular expression notation was used for searching for minimal pairs. Parenthesised subexpressions are interpreted as parameters which control the structuring of search results. These extensions, though intuitively simple, make a lot of expressive power available to the user. The current prototype system has been used heavily for substantive phonological fieldwork and analysis on the field, documented in (Bird, 1997). There are a number of ensuing benefits of this approach for phonological research: (i) it supports a quantitative method of doing phonological research; (ii) it gives universal access to the same set of informants; (iii) it enables other researchers to hear the original speech data without having to rely on published transcriptions; (iv) it makes the full power of regular expression search available, and search results are full multimedia documents; and (v) it enables the early refutation of false hypotheses, shortening the analysis-hypothesis-test loop.

## ACKNOWLEDGEMENTS

This research is funded by the the UK Economic and Social Research Council, under grant R00023 5540 *A Computational Model of Tone and its Relationship to Speech*. My activities in Cameroon were covered by a research permit with the Ministry of Scientific and Technical Research of the Cameroon government, number 047/MINREST/D00/D20. I am grateful to Dafydd Gibbon for helpful comments on an earlier version of this paper.

## References

- Bevan, D. (1995). *FindPhone User's Guide: Phonological Analysis for the Field Linguist, Version 6.0*. Waxhaw NC: SIL.
- Bird, S. (1997). Dschang Syllable Structure. In H. van der Hulst & N. Ritter (Eds.), *The Syllable: Views and Facts*. Oxford University Press. To appear.
- Bird, S. & Ellison, T. M. (1994). One level phonology: autosegmental representations and rules as finite automata. *Computational Linguistics*, 20, 55–90.
- Bird, S. & Tadadjeu, M. (1997). *Petit Dictionnaire Yémba-Français (Dschang-French Dictionary)*. Cameroon: ANACLAC.
- Buseman, A., Buseman, K., & Early, R. (1996). *The Linguist's Shoebox: Integrated Data Management and Analysis for the Field Linguist*. Waxhaw NC: SIL.
- Coleman, J., Dirksen, A., Hussain, S., & Waals, J. (1996). Multilingual phonological analysis and speech synthesis. In *Computational Phonology in Speech Technology: Proceedings of the Second Meeting of the ACL Special Interest Group in Computational Phonology*, (pp. 67–72). Association for Computational Linguistics.
- Ellison, T. M. (1992). *Machine Learning of Phonological Structure*. PhD thesis, University of Western Australia.
- Kaplan, R. M. & Kay, M. (1994). Regular models of phonological rule systems. *Computational Linguistics*, 20, 331–78.
- Kornai, A. (1995). *Formal Phonology*. New York: Garland Publishing.
- Lowe, J. B. & Mazaudon, M. (1994). The Reconstruction Engine: a computer implementation of the comparative method. *Computational Linguistics*, 20, 381–417.
- Reinhard, S. & Gibbon, D. (1991). Prosodic inheritance and morphological generalizations. In *Proceedings of the Fifth Conference of the European Chapter of the Association for Computational Linguistics*, (pp. 131–6). Association for Computational Linguistics.
- Wall, L. & Schwartz, R. L. (1991). *Programming Perl*. O'Reilly and Associates.

