Phonology Assistant 3.0.1

from SIL International

Reviewed by MARK DINGEMANSE, *Max Planck Institute for Psycholinguistics*

1. OVERVIEW. Phonology Assistant (PA) is a free phonology tool by SIL that (as of version 3.0) works interactively with the data stored in Toolbox, Fieldworks Language Explorer, and Speech Analyzer. It automates many of the cumbersome and repetitive tasks associated with doing phonological analysis, and it does so in a most systematic and revealing way. The things it does more or less automatically include drawing up a phone[1] inventory; computing relative frequencies of phones; computing syllable structures; generating phonotactic charts for every conceivable combination of positions, phones, or features; and finding minimal pairs along various dimensions. A powerful search function allows the user to search for phonetic patterns within specified environments.

2. GETTING STARTED: DATA SOURCES AND DATA. To automate this type of meticulous phonological bookkeeping, the data that PA works with have to meet some requirements. First of all they have to be Unicode.[2] Secondly, the data have to reside either in one or more Toolbox databases (PA can work with both interlinearized and non-interlinearized data) or in a Fieldworks Language Explorer project. Third, the transcriptions need to be phonetic or at least phonemic (preferably both), and in either case close to IPA.[3] The kind of analytic results one can obtain will obviously depend on the nature of the input; with only a broad phonemic transcription, PA will not be able to extract many interesting distributional facts that might be available in a narrow phonetic transcription.

PA plays very well with the applications from which it reads data. It automatically recognizes when sources have been modified and will update its views accordingly. Designed strictly as a tool for analyzing rather than editing, PA simply loads the data directly from the source without even being able to modify it. However, if the need arises to fix a quirky record (and PA is an excellent tool to discover such records), there is a command "Edit source record" through which the record can be opened in either Toolbox or Fieldworks Language Explorer. PA is project-based, i.e. metadata such as data sources, relevant source

---

[1] "Phone" is a conveniently agnostic term used throughout PA to refer to any character found in the data — conveniently agnostic because PA cannot be sure whether the transcription is phonetic (in which case a phone would be a speech sound) or phonemic (in which case it would be a phoneme), or even a mixture of both.

[2] There are apparently older, less feature-rich versions of the program available that can work with non-Unicode data.

[3] Note that the grapheme /y/, often used for the palatal glide [j] in broad phonemic transcriptions of language data, by default ends up as a high rounded vowel [y] in the Vowel chart. This problem will not arise when there is a separate field for phonetic data in IPA

fields, and other settings are stored in a .pa project file. It is possible to work with only one project at a time.

The PA interface has a task-based structure consisting of six main parts. The first part is the Data Corpus view. As soon as PA is told about the data source and the fields where it can find phonetic and phonemic transcriptions, this view is dynamically populated with records. Also, a column containing CV patterns is automatically generated. Here the first question may arise: how does PA know that, say, the digraph /kp/ is a doubly-articulated C instead of a sequence CC? The answer is that it doesn't unless it is told otherwise. This is done in the Phone Inventory view, the best place to go after loading the data. In this view there is a pane called "Ambiguous Sequences," where all digraphs, trigraphs, and the like can be registered. Another important feature in that view is "Experimental Transcriptions." Here, custom diacritics can be mapped onto their IPA equivalents so that PA can link them with the relevant phonological features (on which see below).

3. PHONE CHARTS, THE PHONE INVENTORY, AND THE SEARCH VIEW. When the source is specified and PA knows about any idiosyncratic transcriptions, it is time to check what it makes of the data. PA will automatically display consonant and vowel charts together with very useful frequency histograms (see figure 1 above). If the chart needs rearranging or if there are empty rows or columns, this can be fixed by dragging around phones, removing unused rows and columns, and editing captions. Note that when a phone is moved to another place in the chart, the effect is merely visual — i.e., the underlying
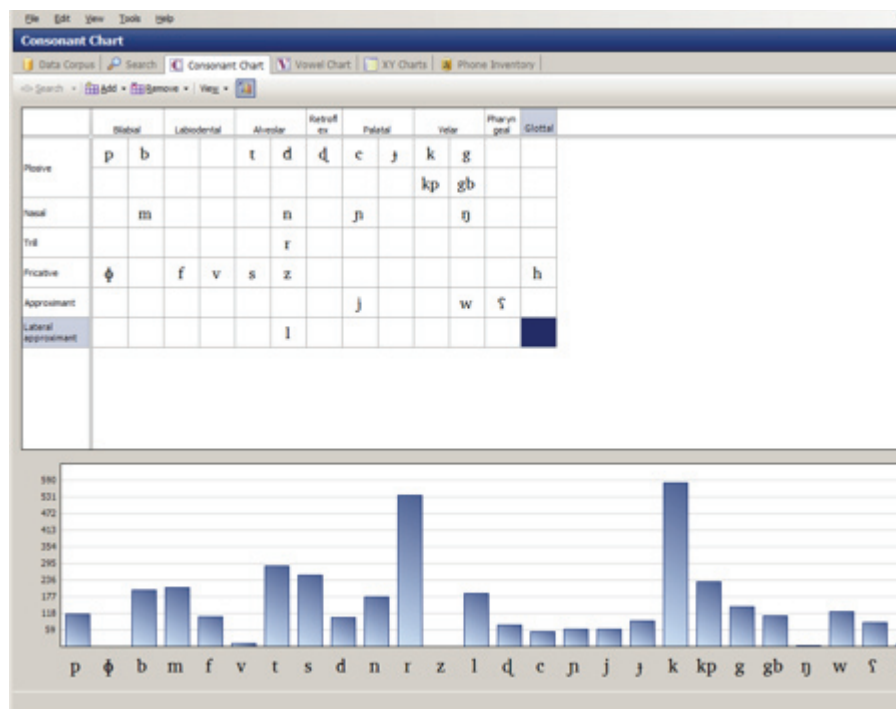

FIGURE 1: The consonant chart view

features are not changed. Phones cannot be removed from the chart as long as they occur in the data.

A brief word may be in order about these underlying features that PA gets from a built-in library of phonetic characters. All of these characters are assigned default articulatory features (e.g. bilabial, consonant, plosive) as well as binary features (e.g. -approx, +cons, -constr, -cont). The default library will be sufficient for most purposes, but PA does offer the option to edit the features in the Phone Inventory view. Moreover, characters can be added to the inventory with a separate editor.[4]

Double-clicking on a phone in the consonant or vowel chart will bring up a search for all occurrences of that phone in the data corpus. This is the Search view, where one can search for phones within specified environments. Searching for specific phones is just one possibility; much more complex searches can be constructed with reference to the articulatory and binary features mentioned above, or with reference to classes of phones or features defined by the user. This makes the Search view an excellent place to quickly test certain hypotheses. For example, one could do a search for all occurrences of low tones after voiced and voiceless plosives to see if depressor consonants play a role in the tone system. Search patterns are constructed in a simplified version of Regular Expressions (regex), and they can be saved for later use. The Search view comes with a copious set of example search patterns.

4. X/Y CHARTS: AUTOMATIC ACCOUNTING. To work out a good phonology, one needs to check in which environments each single phoneme can occur. This is a cumbersome task that can grow enormously complex if one wants to be comprehensive and look not just at simple positional distribution (initial, medial, final) but also at its occurrence in different environments (intervocalic, before voiced fricatives, after a nasal consonant, etc.). PA greatly simplifies this job through XY charts, probably its most powerful feature. An XY chart plots phonemes, articulatory features, or sequences of features against each other and produces for each combination the number of occurrences in the data. They can be customized and saved.

For example, the XY chart in figure 2 below plots the Siwu[5] consonants by manner of articulation against several possible environments in which they may occur. Each cell displays the number of occurrences found of that particular pattern. The selected cell shows that there is one instance of a voiceless consonant (search pattern [[C][Voiceless]]) occurring word-finally (context *_#). The results pane below the table shows the item in which this pattern occurs: the loanword *tsîs* [cîs] 'cheese'. Cells with zero results are yellow. Thus, plosives never occur word-finally in my Siwu data.

A convenient touch in the Search and XY Chart views is the "minimal pairs" button;

---

[4] This built-in library should remind one of the fact that PA cannot actually *do* the phonetics and phonology. It should always be kept in mind that PA is essentially no more than a meticulous accountant with a formidable memory and a handy library of phonetic and articulatory features. It is the user who is, in the end, responsible for the actual analysis of the data; PA just takes on part of the bookkeeping chores — no more, no less.

[5] Siwu is a na-Togo (GTM, Kwa, Niger-Congo) language spoken in eastern Ghana.

if this is activated, the results are ordered such that minimally different words end up next to each other. This feature can be configured to be very strict (requiring identical preceding and following environments) or more lenient (requiring only identical preceding or identical following environments). Here, as elsewhere in searches, one can choose to ignore diacritics and features like stress, tone, and length if needed.
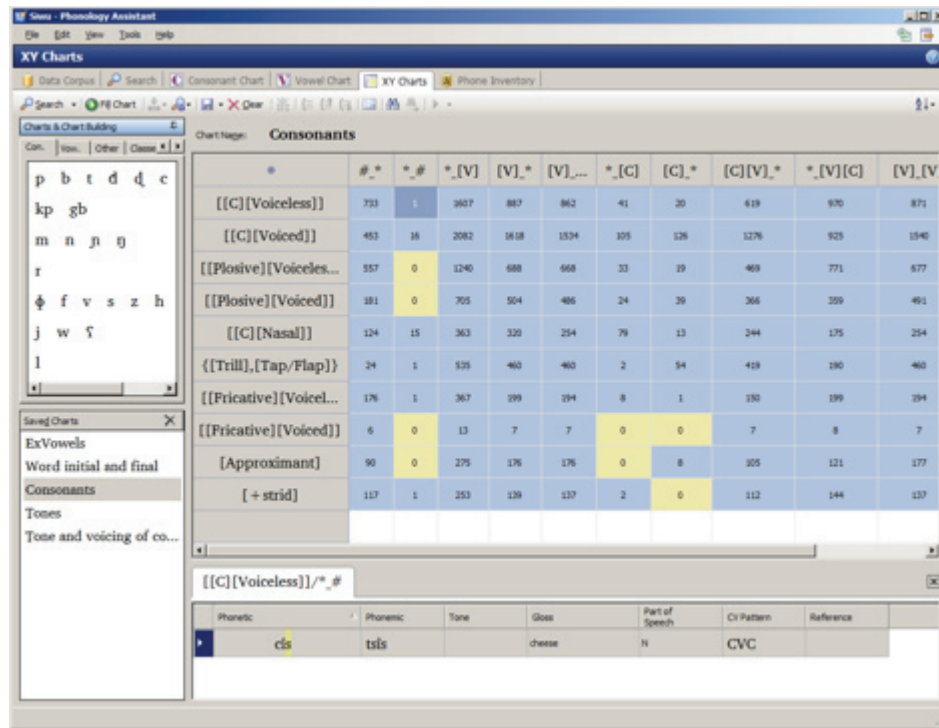
FIGURE 2: An XY chart showing the distribution of different types of consonants in various environments.

5. GREAT INTERFACE DESIGN. While working with PA, it is very easy to overlook the conceptual complexity of the tasks it carries out. This is because everything takes place in a beautifully laid out user interface. The interface has been designed with the workflow of the user in mind, and as a result it works intuitively. Not only are the main views laid out in a very sensible way, there is also a lot of attention to detail. Efficient use is made of the screen real estate: there is no unused space, yet all important features are easily findable; and panes can often be hidden, docked, or undocked (large XY charts for example can be maximized on their own, or even dragged to the secondary monitor if one has a twomonitor setup).

Care is also taken to give recurring panes in the different parts of the application a uniform, consistent look. An example is the record listing, which is the basis of the Data Corpus view, but which is also used to display search results in the Search view and examples in the XY Charts view. In all of these places, a soon familiar set of buttons enables the user

to sort records by place or manner of articulation (a convenient feature) or alphabetically by Unicode order; and furthermore to group records (for example by part of speech) and to collapse groups for increased clarity. PA is also very user-friendly in that difficult concepts like search patterns and XY charts come with helpful sets of built-in examples. Moreover, contextual help is available everywhere, and a handy Student Manual is included in the installation.

I feel that the interface design really deserves high praise. Perhaps some of us have grown too familiar with software, open-source or commercial, that comes with utterly un-intuitive application logic and an interface design consisting of little more, in some cases, than hardcoded clickable rectangles; this may have made us forget how nice it can be to work with a visually pleasing, well thought-out, and well laid-out interface. Phonology Assistant sets a standard here that one can only hope will be followed by other providers of software for linguistic analysis.

6. LIMITATIONS AND GLITCHES. With all the praise out of the way, it is time to point out some limitations of the program. The biggest missing feature, and this is quite an essential one, is the ability to filter the data by word class (part of speech). That is, it is not yet possible to display phone frequency histograms and phonotactic patterns by part of speech, or to search for certain sequences in one word class only. Let me give an example of why this would be useful. In Siwu, certain speech sounds are only found in ideophones. They are thus extra-systemic in the sense that they are not part of the phoneme inventory of non-ideophonic vocabulary. I suspect that it might also be true that the relative frequencies of phonemes and phonotactic patterns are significantly different in ideophones than in other word classes. Moreover, nouns, verbs, and ideophones may have different sets of possible tonal melodies. Finding out distributional facts of this type is usually a very tedious job, and it is precisely the sort of task that one would want a phonology assistant to take on. At present, however, PA cannot do it.[6]

Two further omissions become apparent after some playing around with the Search and XY Charts views. First, it is strange that there seems to be no way to reference the syllable patterns that are so kindly generated in the Data Corpus view. It should be easy, given these representations, to refer to V1, V2, or C2 in search patterns as well as in XY charts. This would considerably simplify certain searches. Second, and related to this, the watered-down version of regular expressions used in the search patterns is not always powerful enough. I have yet to find a way to search for things like *#(nothing OR C)_* (which should be a pattern that yields the first vowel of a word whether preceded by a consonant or not); also, it seems there is no NOT operator; and for unclear reasons, one can only use one * wildcard in a preceding or following environment. These limitations make it quite

---

[6] Except indirectly. For instance, if I search for the pattern CVN# and out of the 30 search results 27 happen to be ideophones, this is telling. But I cannot look up syllable structures in the class of ideophones alone to check how this pattern relates to other phonotactic structures. The only work-around I can think of is to duplicate the data source and modify it to contain only the words I want to look at with PA. But since this would involve creating a separate Fieldworks project or Toolbox database for each word class, and since only one PA project can be opened at a time, this would be far too cumbersome.

difficult to look into vowel harmony, to mention just one thing one might want to do.

Furthermore, as much as I love the interface, I have found some unexpected omissions. For example, right-click is not functional in the Data Corpus view. I would expect it to give me a context menu including, for example, the option to edit the record in the source application — an option which is now buried in the Edit menu or reachable via the shortcut Shift+F2 (why not just F2?). Also, copying to the clipboard doesn't work as expected: it always copies a whole row of data instead of just the selected cell. Talking about rows, it is not possible to move around columns and rows in XY charts; and it doesn't seem possible to hide columns in the Data Corpus view (and in results charts), even if they are empty.

I would like to see a few more sorting options, of which secondary sorts (allowing one to sort first by part of speech and within that by place of articulation) would be the most essential. At present, sorting by manner of articulation while the primary sort is part of speech unexpectedly makes the phonetic column the primary sort. Also, sorting by length or from end (as in Fieldworks Language Explorer) would be nice to have. In the Phone Inventory view, I find that changes to articulatory or binary features cannot easily be undone. Fiddling with the features can be useful, but it would be convenient if there were an option to revert to the default values so that one knows that there is a way back if things go awry.

Last but not least, although most of the tabular views can be nicely exported to RTF/CSV, XML, or HTML, exporting some other types of data from PA is difficult. For the neat phone charts, the "Export to RTF" in the File menu is grayed out, leaving only HTML and PA XML export options. Fortunately the HTML export works fine,[7] and it is possible to copy/paste it from there into a word processor. There is no way except a screenshot to get the frequency histograms out of PA. It would be nice to have SVG export for that (or at least to be able to export the underlying data to a spreadsheet application).

7. CONCLUSION. Phonology Assistant is a tremendously useful and beautifully executed application. With the exception of the missing filtering feature, the limitations mentioned above are largely negligible, and I am confident that Phonology Assistant 3.0 will be appreciated by countless colleagues. Many thanks to SIL International for making this tool freely available, and special praise to the developers for the interface design.

---

[7] I did note however that the title attribute is not set correctly.

Primary function:	"[Phonology Assistant] assists users in phonetics and phonology. It creates consonant and vowel inventory charts from a data corpus and allows a user to search that corpus for phonetic patterns within specified environments." (quotation from the readme file)
Pros:	Direct reading and dynamic updating of data from Toolbox, Fieldworks Language Explorer, and Speech Analyzer. Multiple data sources possible. Powerful search function that can find phones, articulatory features, binary features, and user-defined sets in all sorts of environments. XY charts plot phonemes, articulatory features, or sequences of features against each other. Easily find minimal pairs. Support for ambiguous sequences, experimental transcriptions, and uncertain transcriptions. All tabular data can be exported to RTF/CSV, HTML, or XML. Very complete documentation, including training guides and saved examples of search patterns and XY charts. Unicode support. The task-based interface is beautiful and well thought-out. Distributed as freeware.
Cons:	Filtering by part of speech is not possible. Only one project can be open at a time. Computed syllable structures cannot be referenced. Search patterns may not be powerful enough for some users. Phone charts cannot be exported to RTF; frequency histograms cannot be exported at all. PC/Windows only.
Platforms:	PC/Windows only.
Open Source:	No.
Proprietary:	Relies on .NET technology.
Available from:	SIL International (http://www.sil.org/computing/pa/)
Cost:	None; distributed as freeware.
Reviewed version:	3.0.1
Application size:	8 Mb (without the .NET framework)
Documentation:	http://www.sil.org/computing/pa/

Mark Dingemanse
Mark.Dingemanse@mpi.nl