Introducing Phon to the Sociophonetics Community*

Angelica Hernandez and Yvan Rose
Western University and Memorial University

1. Sociolinguistics and phonetics

Decades of research in variationist sociolinguistics have shown that phonetic and phonological variation may be driven by combinations of linguistic and social factors. For example, recent research by Tennant et al. (2015) shows that the aspiration of /s/ in coda position in Cuban Spanish is conditioned by linguistic factors such as the position of the segment within a word, and the consonant following the segment of interest, as well as by social factors such as the socioeconomic level, and the rural or urban residence of the speakers.

The field of sociophonetics, which focuses on the social factors that condition and influence phonetic variants, has expanded significantly since Labov's early studies in New York City and in Martha's Vineyard (1966, 1972). Many phoneticians now recognize the importance of investigating social factors in conjunction with phonetic aspects of speech (Müller and Ball 2012, Labov 2006, Hay and Drager 2007, Hay and Foulkes 2016). Along with the growth in interest from linguists on the interaction between social factors and phonetic variants, there has been and continues to be a need for the development of computer-assisted methods to study these phenomena which would otherwise be difficult to assess. For instance, more traditional, manual impressionistic judgements are now being supplemented, if not supplanted altogether, for computer-assisted measurements of speech samples which can provide precise data descriptions regarding various relevant acoustic parameters such as vowel formants, voice, or pitch, among others (Thomas 2013). These advances have significantly contributed to the field of sociophonetics by increasing analytical scope, precision and replicability.

Of particularly significant importance during the last two decades is Praat, an open-source, free software program available to students and researchers worldwide since the mid-1990s (http://www.fon.hum.uva.nl/praat/; Boersma and Weenink 2019). This program, which provides researchers with the ability to carry out acoustic analyses of speech data (among many other functions), offers powerful solutions to phoneticians and laboratory phonologists. The recent expansion of these fields of research, and advances within them, can be attributed at least in part to the unrestricted availability of Praat. In addition to Praat, which stands as the first and foremost example of open technology for acoustic analysis, other programs such as EMU-SDMS (Winkelmann et al. 2017) and

* We are thankful to Drs. Jeff Tennant, David Heap and Jorge Emilio Rosés Labrada for allowing us to use the data from the Holguin corpus for our practical demonstration. The early development of Phon was funded by grants from the Social Sciences and Humanities Research Council of Canada, the Canada Fund for Innovation, as well as by a Petro-Canada Award for Young Innovators. Since 2006, the development of Phon and PhonBank has been funded primarily through grants from the National Institutes of Health (USA; R01 HD051698, R01 HD051698-06A1, and R01 HD051698-11).

Actes du congrès annuel de l'Association canadienne de linguistique 2019.
Proceedings of the 2019 annual conference of the Canadian Linguistic Association.

LaBB-CAT (Fromont and Hay 2012), to name a few, have also contributed to research in similar ways.

However, for the field of sociophonetics in particular, analyses incorporating both the social and linguistic aspects pertinent to current research questions remain cumbersome. While software for acoustic analysis provides the tools to carry out precise acoustic analyses, they do not provide facilities for the coding of social or demographic variables; they also do not offer straightforward methods for the incorporation of these variables within data queries. These systems also typically offer very little by way of phonological annotations pertaining to syllable structure, stress, or segmental features, or to other potentially relevant linguistic variables such as position within the word, phrase or utterance. This results in the need to manually input large amounts of arbitrary annotations to the datafiles, a task particularly challenging in the context of multifactorial analyses.

More generally, the programs listed above are better suited for analyses of shorter utterances than for the long recordings typically involved in sociolinguistic interviews. In this regard, programs such as ELAN (Brugman and Russel, 2004) or Transcriber (Barras et al. 2001) provide convenient solutions for the transcription and annotation of audio and video files. These programs, however, do not have the capacity to carry out acoustic analyses, and do not provide much support for analyses based on linguistic factors; much manual coding is also required, and their query functions do not lend themselves easily to multifactorial analysis either. This leaves sociophonetic studies involving complex methodologies largely unsupported; most studies in this area require the creative combination of often labour-intensive annotation and data compilation methods, where the potential for human error is constantly present, in addition to challenges in database management, data and method sharing, scientific replicability, and so on. To the best of our knowledge, of the systems listed above, LaBB-CAT provides some of the most advanced solutions toward these challenges, for example through its functions integrating linguistic and social/demographic variables as part of its query system. LaBB-CAT is however relatively difficult to implement and maintain, as it involves a server-based computer architecture.

In this paper, we discuss some of these issues in light of another open-source (free) software program, Phon (https://www.phon.ca; Rose and MacWhinney 2014), which offers several ready solutions to the variationist community. Among other advantages, Phon incorporates tools that serve the social and linguistic aspects of sociophonetic research within a stand-alone application that supports all essential aspects of database creation and data queries combining social/demographic, linguistic and acoustic data within a uniform framework. In a nutshell, Phon greatly simplifies variationist sociolinguistic research, as we illustrate through the following sections.

2. Phon

Phon was initially introduced to the linguistics research community over a decade ago (Rose et al. 2006). It was initially developed as a tool to facilitate the phonological analysis of large speech corpora for phonetics and phonological research, more specifically within the context of the PhonBank database, which aims to support corpus-based research on first-language, second-language, or bilingual phonological development, as well as on developmental and acquired speech disorders (https://phonbank.talkbank.org; Rose and MacWhinney 2014).

Phon has since significantly evolved to accommodate more general methods supporting corpus-based research on phonetics and phonology, also in line with the trend noted in the introduction, to meet the now-common need for integration of social/demographic factors as part of phonological or acoustic analysis. Phon consists of a series of inter-connected modules assembled within a uniform graphical user interface which together offer functionality to assist the researcher in important tasks related to corpus transcription, coding and analysis (Rose et al. 2006). Furthermore, because Phon also functions a database system, it offers ready solutions in the area of corpus management as well as several functions to facilitate the sharing of both data and query methods among students and researchers, both within individual research projects and beyond. In this respect, and because of its user-friendly features as a standalone, integrated program, Phon readily supports open research, delivering powerful solutions directly in the hands of end users.

3. Traditional methods and challenges in sociophonetics research

Sociophonetic studies investigate the effects of social factors such as age, gender, and socioeconomic level, among others, on phonetic and phonological variants. These studies range anywhere from investigations of language variation and change regarding the use of different phonemes by different generations of speakers in a community, to clinical studies of language and communicative disorders documenting patients with language impairments. In all of these cases, the studies rely on the use of natural speech data, often collected by means of sociolinguistic interviews. Natural speech data is essential to be able to analyse real-world speech production and how the variation observed in the speech data is conditioned as a function of specific social characteristics of speakers.

Following the collection of speech data, researchers must transcribe and code the interviews for the linguistic, social and demographic variables of interest. For the transcription and annotation (or coding) of the data there exist a variety of software solutions, including those mentioned in the introduction. In the case of sociophonetic studies, it is common practice to use Praat to transcribe and annotate, a tradition that goes back to the advent of Praat as free, open-source software incorporating the tools essential to acoustic analysis. Prior to this, acoustic analysis was limited by the high cost and scarce availability of the solutions available at the time.

However, as already noted, the process of transcribing and annotating the data in Praat can be very time consuming. After opening the audio file of interest in the program and creating a TextGrid file, researchers must then manually segment the file in a tier (i.e. set the time intervals relevant to that specific tier) and then transcribe the data, whereby every time interval and annotation relevant to every tier must be set and/or adjusted manually by the researcher. In **Figure 1** below, we see an example of a sample of speech data from a Cuban Spanish corpus (Tennant et al. 2015) which has been transcribed orthographically (Tier 1). In this study, Tennant et el. (2015) investigated the social and linguistic factors that influence the aspiration [h] and deletion [Ø] of /s/ in coda position in the Spanish of Holguín, Cuba. For this purpose, the data were coded to investigate the effects of several morphological and phonetic factors influencing the speakers' speech productions, including position of coda /s/, the preceding and following segments, and the actual allophone used within the utterance (Tiers 3-9).

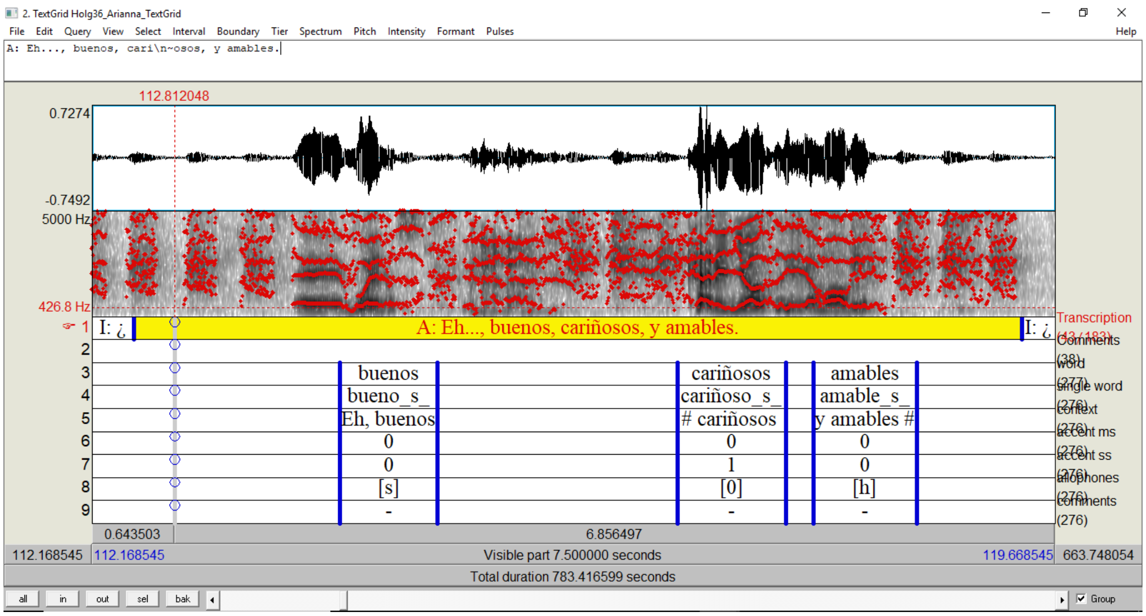

Figure 1. Praat TextGrid file with coding of linguistic factors for /s/ aspiration and deletion in coda position in Cuban Spanish (Tennant et al. 2015)

Once the data has been fully segmented and transcribed, the researchers must then manually look for and, at times, further annotate the transcription for all tokens of interest in the speech sample, which in the case of sociolinguistic interviews may be as long as an hour in duration or at times more. This process is not only time consuming but also susceptible to human error, and thus requires additional checks to ensure minimal errors, such as misidentified or miscoded tokens, before final acoustic analyses can be performed.

The results of these acoustic analyses must then be exported into spreadsheets or statistical analyses packages where the social/demographic factors relevant to each result are incorporated, again manually. This is illustrated in **Figure 2,** where we see an example of the result of this process, in which the spreadsheet includes the linguistic factors annotations which were extracted from Praat, along with annotations for social factors such as speaker ID, sex, age, and socioeconomic status coded together. Similar to the process of annotating the TextGrids, this manual incorporation of the social and linguistic factors into a unique document is both time consuming and prone to human error.

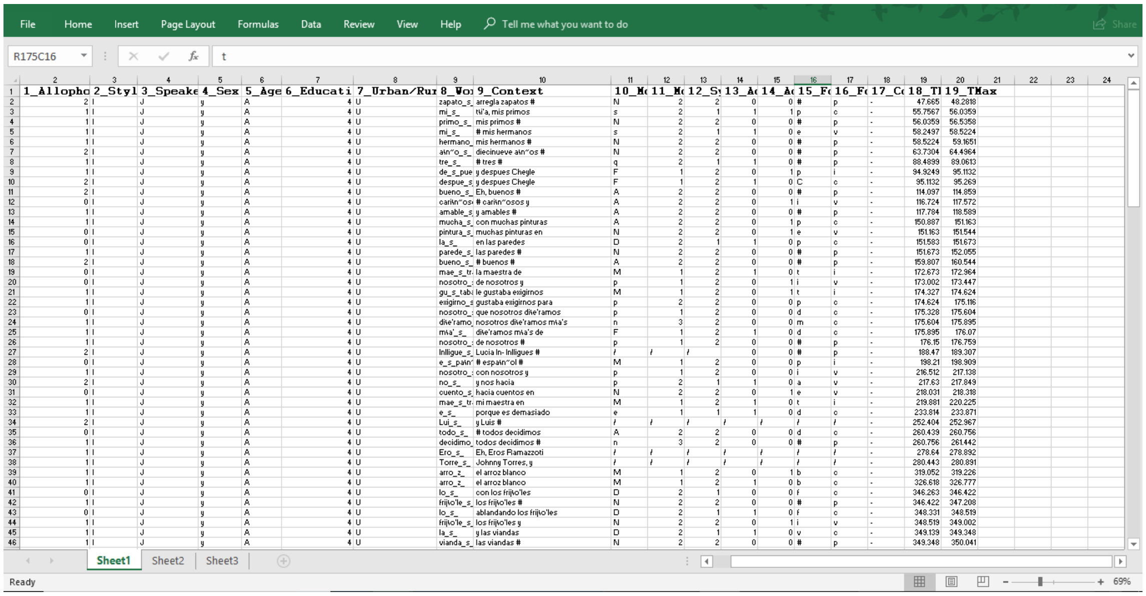

Figure 2. Spreadsheet showing annotations for linguistic factors and corresponding social factors for tokens of /s/ aspiration and deletion in coda position in Cuban Spanish (Tennant et al. 2015)

By the time this work is completed, researchers generally end up with multiple files, including the audio recordings of the interviews, their corresponding transcripts (e.g. Praat TextGrids), the results exported from the analyses, and the documents ultimately collating all of the data relevant to analysis. This renders the organization and subsequent sharing of the corpus complicated and impractical. More importantly, any change to any of the original transcripts can also have snowballing effects making the end-result documents extremely hard to update at times or, as it is too often the case, useless; the work has to then be performed anew. With this general context as a background, we highlight some of the advantages that Phon offers, in the next section.

4. Phon for corpus creation and sociophonetics research

It is generally agreed that when creating corpora for the linguistic analysis of speech production, a larger sample size is ideally preferred since it is more likely to be representative of the population of speakers under investigation. While previous conceptions of an ideal corpus size for sociolinguistic research was stipulated at around a million words, nowadays this notion has been replaced by the idea that a larger corpus is always better (Voorman and Gut 2008). Larger corpora make conclusions that are more generalizable to the real world — yet the larger the corpus is, the longer the time required to transcribe and annotate the data. Phon provides welcome facilities in this respect, combining power and flexibility, and makes these functions accessible through an intuitive graphical user interface. Data segmentation can be performed in real time (with plans to automate this task in the near future, through the incorporation of diarization software libraries; Rose 2019), and transcription is supported through a built-in map to easily access for all symbols and diacritics of the International Phonetic Alphabet (IPA;

https://www.internationalphoneticassociation.org/) as well as a large number of dictionaries of pronounced forms (19 languages and dialects are supported, all in standard IPA), and various functions to easily split, merge or other modify data records (e.g. find/replace functions, record list generation, etc.).

4.1 Automated functions and Praat integration

Once the orthographic transcription of a corpus is completed (a task that remains essentially manual in the absence (still today) of reliable and/or easily accessible speech recognition systems), Phon can automatically generate from its built-in dictionaries IPA-transcribed forms which greatly speed up this otherwise time-consuming task. Even the transcription of non-standard pronunciations becomes easier, as it only involves the modification of the standard forms provided by the built-in dictionaries as opposed to the full transcription of all the phones and diacritics otherwise required to obtain a phonological representation of the speech productions.

Based on the (automatically- or manually-input) phonetic transcriptions, Phon automatically labels utterances into syllables and syllable positions (e.g. syllable onsets vs. codas), and each time a researcher wants to compare produced forms against standard versions of these forms (e.g. *goin'* /goɪŋ/ produced as [goɪn]), Phon automatically performs a phone-by-phone alignment through which we can obtain descriptions of the non-standard forms based on systematic comparisons with their corresponding standard forms (here, the alveolar production of a standard velar nasal in the gerund). In all cases (IPA generation, syllabification labelling, and phone alignment), the forms and annotations generated by Phon are fully modifiable by the user. Phon also offers different syllabification algorithms to accommodate different theoretical views of syllabification. Finally, each phone and diacritic is assigned descriptive phonological features, such that [ɛ̃] is described through the features set {vowel, mid, lax, front, unrounded, nasal}. This feature-level labelling helps research on natural classes of phones (e.g. voiceless obstruents, tense vowels, etc.).

Finally, Phon integrates seamlessly with Praat functions for acoustic data analysis; TextGrids can be imported from Praat and serve as a basis for the generation of corresponding data records in Phon, or be generated directly from within Phon. At the moment, Phon-generated TextGrids must be aligned manually or using third-party forced alignment software (libraries for forced alignment are also planned for future updates to Phon); however, no additional annotation is required before acoustic analysis. In all cases where a Phon corpus includes TextGrids, the record data and associated TextGrid annotations can be used in tandem to incorporate as many linguistic criteria as needed into the queries (e.g. utterance-, phrase-, word-, syllable-, and phone-level information, including phonological features), as required by the question at hand. This can be seen in Figure 3 below, where the segmentation tool, the spectrogram and all tiers are visible and accessible within the same window.

In order to take advantage of this data structure, the researcher can use one of three simple query languages (simple text, regular expressions, phonological expressions), each of which offers a variety of methods to tap into textual, phonological or acoustic dimensions of the data.

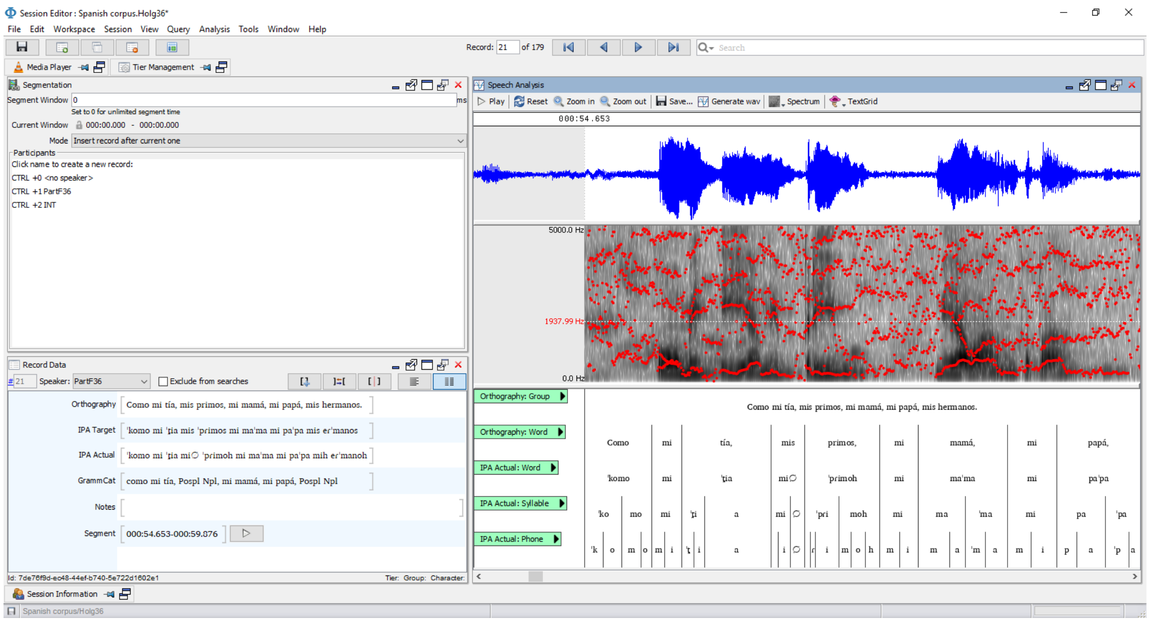

Figure 3. Phon interface used for segmentation and transcription. Figure shows all annotations for lexical, morpho-syntactic, acoustic, and social data (minimized)

For example, a phonological expression (also called a "phonex") can be used to identify the locus of acoustic data measurement. Figure 4 below illustrates the interface for a simple query, based on the "s:c" phonex, which translates into "look for /s/ in syllable codas", as produced by the adult speaker being interviewed.

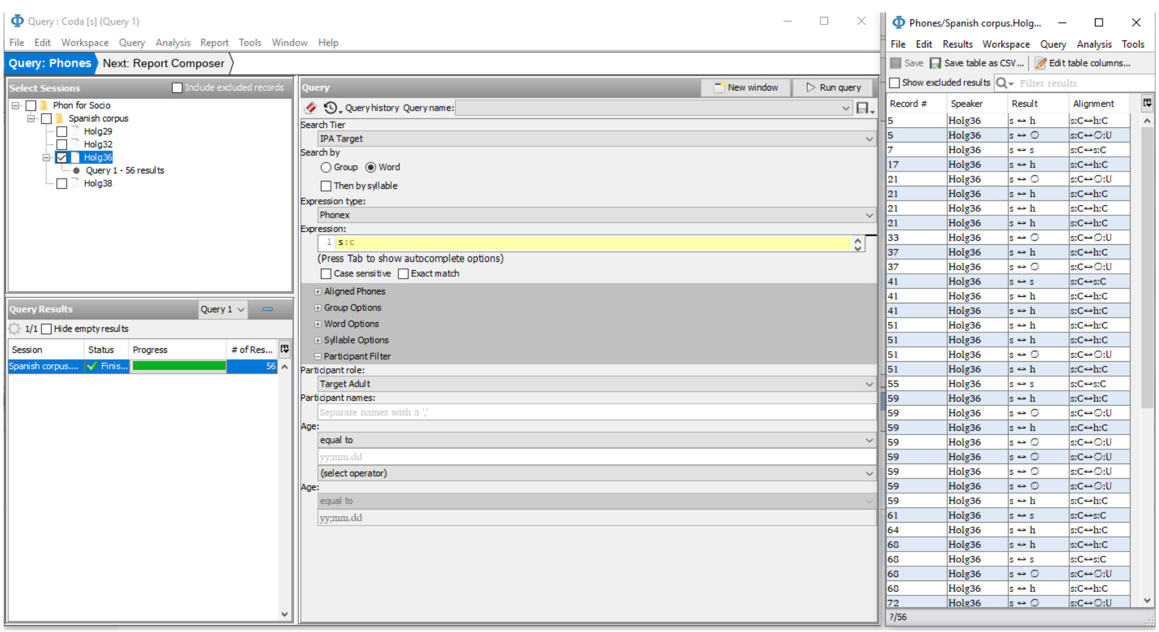

Figure 4. Display of Phon's query tool. In this case, the search was directed to find all instances of coda /s/ in one interview from the corpus of the Spanish of Holguín, Cuba

In sum, researchers can obtain a maximum number of query results and associated acoustic measurements based on a minimum of textual or phonological annotations, the latter of which can often be generated automatically. For instance, returning to the example of the Cuban Spanish corpus discussed earlier (Tennant et al. 2015), once orthographic transcription of the dataset was completed, Phon was used to automatically generate a set of IPA Target (or 'standard') forms and corresponding IPA Actual (speaker-produced) forms using the Latin American Spanish dictionary built into Phon. From this point, all we had to do was to search the corpus for the instances where /s/ is found in coda position (something easy given the phonex "s:c" mentioned above) and modify the Phon-generated forms within the IPA Actual tier to reflect the actual pronunciations used by speakers (listed in the right data panel in Figure 3). Because of the phone-by-phone alignment between IPA Target and Actual forms supported by Phon, we could then perform queries on the standard forms, their variants as produced by the interviewed speaker, or systematically compare the two forms (for example to derive the rate of coda /s/ aspiration or deletion through the program's easy-to-use query and reporting systems).

For added convenience, all individual query results generated by Phon are 'hyperlinked' to the corpus data itself, which can then be accessible at a click, which greatly facilitates further verifications of the data transcripts. Query reports follow a similar logic, such that even sophisticated analyses combining social/demographic, phonological and acoustic criteria can be repeated as needed, without any need for massive integration between different datafiles or for any scripting of the Praat functions. Finally, all queries composed by a researcher can be applied to either single transcripts or sets of transcripts (depending on the question at hand), and be saved for later use; saved queries and associated report layouts can also be shared between researchers (in the form of XML files that can be sent over email), making it easier to reproduce methods across different computers or research sites.

4.2 Phon and sociolinguistic data

As implied above, Phon can greatly facilitates research for sociophoneticians since it allows for the inclusion of social/demographic metadata alongside the linguistic annotations and corresponding audio and TextGrid files. Social/demographic metadata for sociolinguistic studies can be directly entered in Phon data transcripts (primarily through the Session Information panel, which incorporates the relevant information for each participant) so that it remains accessible next to the transcriptions, annotations and corresponding media files. Figure 5 below illustrates this interface.

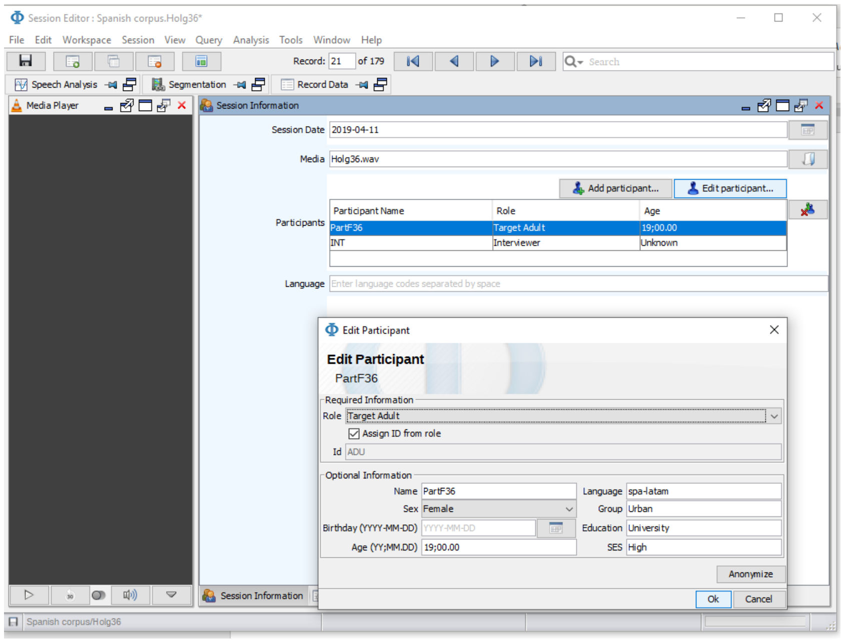

Figure 5. Display of Phon interface for entering social metadata

Factors such as the participant's name or ID, age, language, education level and/or socioeconomic level, as well as any other groups that may be relevant for specific research interests can be included under participant information. Further, it is possible to add multiple participants per interview while specifying the role of each participant in the interview (i.e., interviewer, adult or child participant, etc.). All of these characteristics are then available for specifying data queries later on (e.g. queries that focus uniquely on the interviewees, ignoring all data pertaining to the interviewer).

4.3 Queries and results in Phon

Once transcription and annotations for both linguistic and social factors are completed, researchers can conduct queries. Phon's query interface was already shown in Figure 4 above. Recall that using this window, continuing with the example of the Cuban Spanish corpus of Holguín, we searched for all instances of /s/ in coda position (phonex "s:c") within one interview. After queries such as this are completed, we can generate a report toward either data interpretation or for post-hoc processing within spreadsheet or statistical analysis programs.

To generate a report, the researcher needs to just click on the "Report Composer" tab at the top of the interface. This action triggers another window, shown in Figure 7, where researchers can choose to include different types of data reports and/or extract acoustic measurements relevant to the tokens of interest (returned by the query). In each case, both the linguistic data and the social/demographic information associated to it can be included in the report.

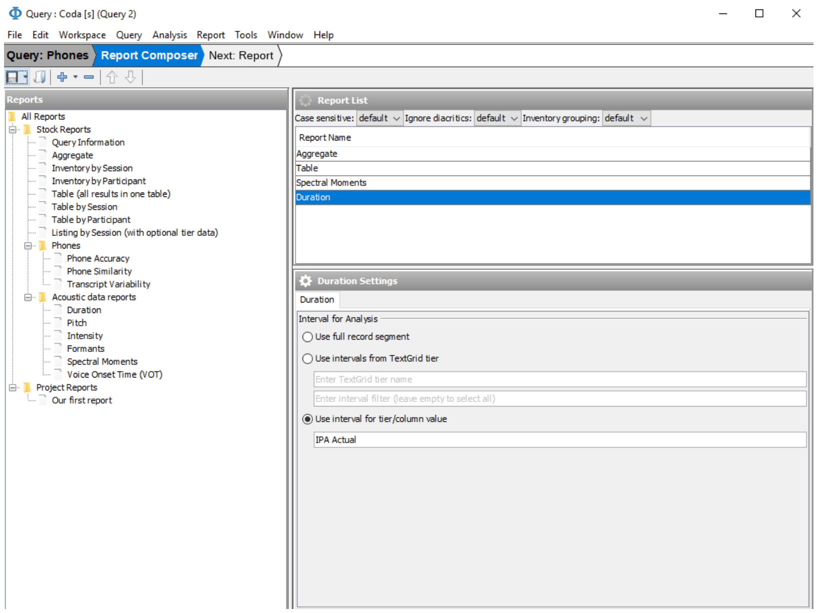

Figure 6. Phon's Report Composer interface showing different report types and formats in the left column, alongside the parameters for each report in the right column

Once all social and linguistic parameters are selected, then by clicking on the following tab "Report" researchers can generate structured reports, in different table or list formats, as shown in Figure 7 below.

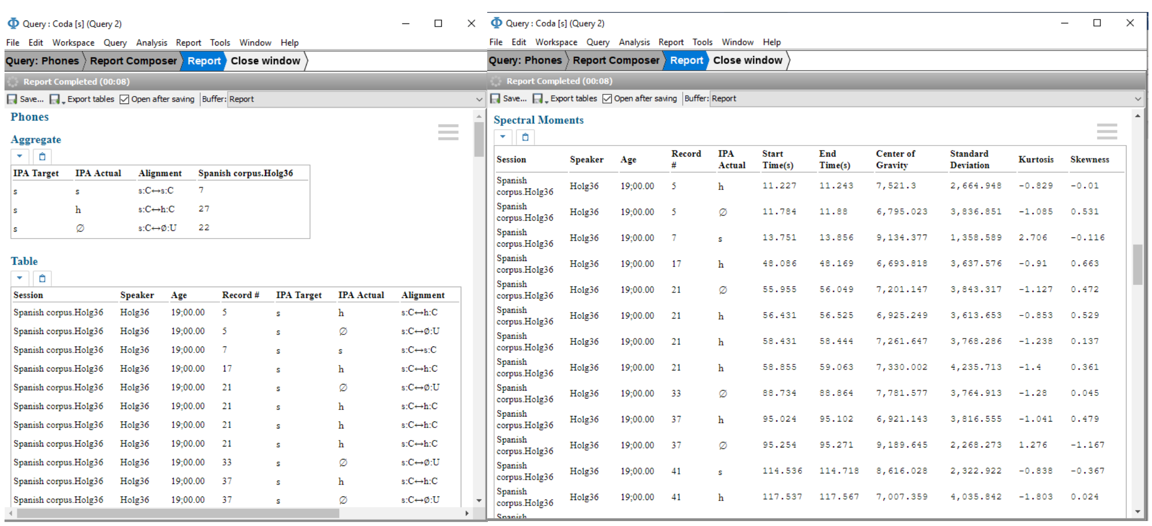

Figure 7. Example Phon report, in this case on the phonetic realization and acoustic measurement of the spectral moments of /s/ as they are produced in syllable codas

Researchers can, using a single, unified report, obtain a general summary of the query results in relation to all of their social/demographic metadata, location of the token within the interview (through the inclusion of beginning/end time values corresponding to the intervals where we can retrieve each token in the recorded sample), and all selected phonetic and acoustic results, all organized in a user-friendly format. In Figure 7 above, we can see two separate sections from a unique report, namely a summary of the frequency of the production of the three possible allophones for coda /s/ in Cuban Spanish and the acoustic measurements data concerning the centre of gravity of the tokens returned by the query, each listed by speaker (here, speaker ID and age are listed in the acoustic report).

The tables automatically generated in the report are also able to be exported in CSV (comma-separated value) format, or directly as an Excel workbook (Figure 8). Both the data transcribed and annotated within Phon and the results obtained from specific queries can thus be further analyzed in spreadsheets or statistical software packages.

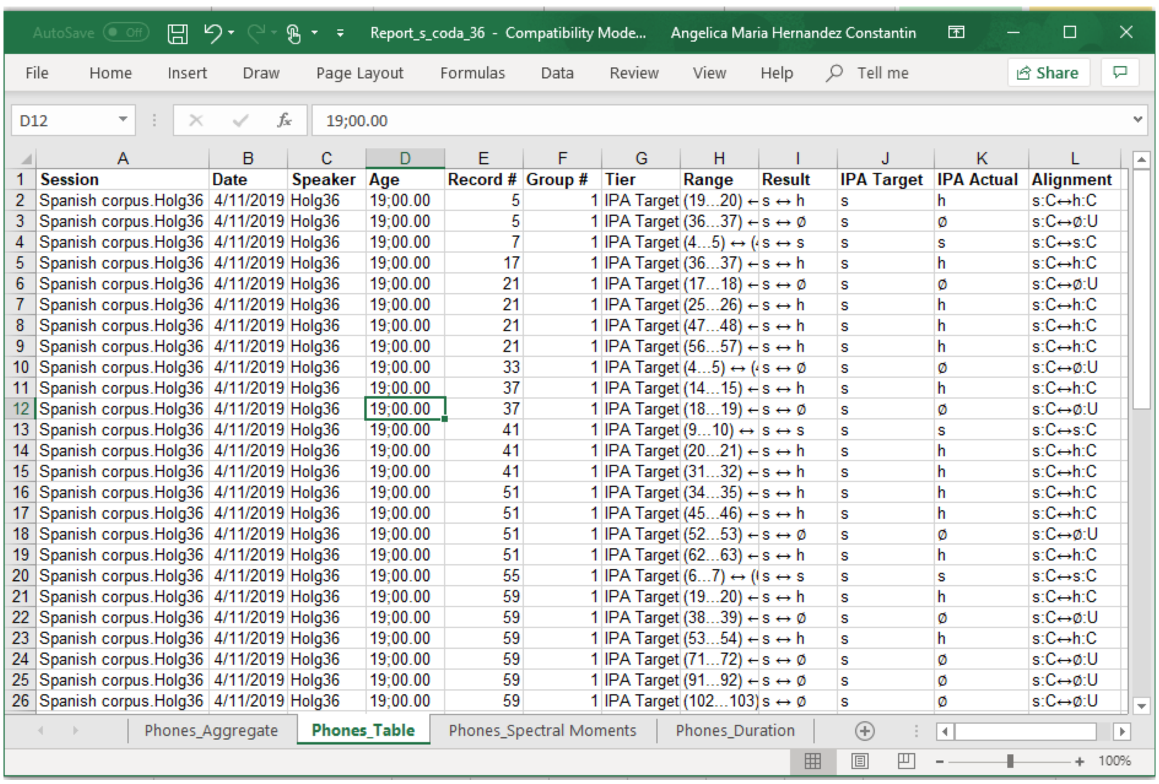

	Session	Date	Speaker	Age	Record #	Group #	Tier	Range	Result	IPA Target	IPA Actual	Alignment
2	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	5	1	IPA Target	(19…20) ←	s ↔ h	s	h	s:C↔h:C
3	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	5	1	IPA Target	(36…37) ←	s ↔ ø	s	ø	s:C↔ø:U
4	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	7	1	IPA Target	(4…5) ↔ (	s ↔ s	s	s	s:C↔s:C
5	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	17	1	IPA Target	(36…37) ←	s ↔ h	s	h	s:C↔h:C
6	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	21	1	IPA Target	(17…18) ←	s ↔ ø	s	ø	s:C↔ø:U
7	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	21	1	IPA Target	(25…26) ←	s ↔ h	s	h	s:C↔h:C
8	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	21	1	IPA Target	(47…48) ←	s ↔ h	s	h	s:C↔h:C
9	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	21	1	IPA Target	(56…57) ←	s ↔ h	s	h	s:C↔h:C
10	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	33	1	IPA Target	(4…5) ↔ (	s ↔ ø	s	ø	s:C↔ø:U
11	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	37	1	IPA Target	(14…15) ←	s ↔ h	s	h	s:C↔h:C
12	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	37	1	IPA Target	(18…19) ←	s ↔ ø	s	ø	s:C↔ø:U
13	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	41	1	IPA Target	(9…10) ↔	s ↔ s	s	s	s:C↔s:C
14	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	41	1	IPA Target	(20…21) ←	s ↔ h	s	h	s:C↔h:C
15	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	41	1	IPA Target	(31…32) ←	s ↔ h	s	h	s:C↔h:C
16	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	51	1	IPA Target	(34…35) ←	s ↔ h	s	h	s:C↔h:C
17	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	51	1	IPA Target	(45…46) ←	s ↔ h	s	h	s:C↔h:C
18	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	51	1	IPA Target	(52…53) ←	s ↔ ø	s	ø	s:C↔ø:U
19	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	51	1	IPA Target	(62…63) ←	s ↔ h	s	h	s:C↔h:C
20	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	55	1	IPA Target	(6…7) ↔ (	s ↔ s	s	s	s:C↔s:C
21	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	59	1	IPA Target	(19…20) ←	s ↔ h	s	h	s:C↔h:C
22	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	59	1	IPA Target	(38…39) ←	s ↔ ø	s	ø	s:C↔ø:U
23	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	59	1	IPA Target	(53…54) ←	s ↔ h	s	h	s:C↔h:C
24	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	59	1	IPA Target	(71…72) ←	s ↔ ø	s	ø	s:C↔ø:U
25	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	59	1	IPA Target	(91…92) ←	s ↔ ø	s	ø	s:C↔ø:U
26	Spanish corpus.Holg36	4/11/2019	Holg36	19;00.00	59	1	IPA Target	(102…103)	s ↔ ø	s	ø	s:C↔ø:U

Figure 8. Excel workbook directly exported from Phon query report

Considering the alternative, more traditional process of analysis of the same data through Praat, where each linguistic annotation needs to be performed manually on each tier, and where exporting each tier takes several minutes and requires storage of multiple files before importing them and then organizing them into a spreadsheet where we must also manually add in all social factors into further data columns, the types of report we can obtain from Phon from a handful of simple steps is enough to illustrate the extent of time savings, as well as the much lower likelihood of human error.

5. Conclusion

As we saw in this brief paper, technological progress has been propelling the field of (socio)linguistics in significant ways over the last few decades, including both the democratization of the tools at the disposal of scholars and the development of innovative methods in support for research that traditionally involved cumbersome methodologies. The Phon software program is particularly relevant to the latter challenge, as it provides the database infrastructure, the corpus-building and annotation functions as well as the query and reporting functions needed for research in sociophonetics (and beyond, across virtually all linguistic disciplines interested in phonetic or phonological evidence) within a unified framework. Not only does Phon make it easier to conduct research in our field, its functionality also greatly reduces the risk that human error might undermine research outcomes.

Beyond these technological improvements, perhaps the next significant barrier to sociolinguistic research lies in the limited availability of corpus data. In this area, the PhonBank project, itself inspired largely by the CHILDES project as well as all similar initiatives which, together, consist of the TalkBank database (https://talkbank.org), offers a model worth considering toward similar initiatives in the area of sociophonetics and sociolinguistics more generally. With technological barriers becoming a lesser challenge, it is indeed within the realm of open science, where data and tools for data analysis are freely available to researchers and their students, that our field could make its most significant progress in the years to come.

References

Barras, Claude, Edouard Geoffrois, Zhibiao Wu, and Mark Liberman. 2001. Transcriber: Development and Use of a Tool for Assisting Speech Corpora Production. *Speech Communication* 33 (1–2): 5–22. https://doi.org/10.1016/S0167-6393(00)00067-4.

Brugman, Hennie, and Albert Russel. 2004. Annotating Multi-Media / Multi-Modal Resources with ELAN. *Proceedings of the Fourth International Conference on Language Resources and Evaluation*, 2065-2068. https://tla.mpi.nl/tools/tla-tools/elan/.

Boersma, Paul and Weenink, David. 2019. Praat: doing phonetics by computer [Computer program]. Version 6.0.46. http://www.praat.org.

Fromont, Robert, and Jennifer Hay. 2012. LaBB-CAT: An Annotation Store. In *Proceedings of the Australasian Language Technology Association Workshop*, 113–117. https://www.aclweb.org/anthology/U12-1015.

Hay, Jennifer, and Paul Foulkes. 2016. The Evolution of Medial /t/ over Real and Remembered Time. *Language* 92(2). 298–330. doi:10.1353/lan.2016.0036.

Hay, Jennifer, and Katie Drager. 2007. Sociophonetics. *Annual Review of Anthropology* 36 (1): 89–103. https://doi.org/10.1146/annurev.anthro.34.081804.120633.

Labov, William. 2006. A sociolinguistic perspective on sociophonetic research. *Journal of Phonetics* 34 (4): 500–515. https://doi.org/10.1016/j.wocn.2006.05.002.

Labov, William. 1972. Sociolinguistic patterns. Philadelphia: University of Pensylvannia press.

Labov, William. 1966. The social stratification of English in New York City. Washington, DC: Center for applied linguistics.

Müller, Nicole, and Martin J. Ball. 2012. Research Methods in Clinical Linguistics and Phonetics. Chicester, UK. John Wiley and Sons Inc.

Rose, Yvan. 2019. PhonBank Data and Analyses: State of the Art and Potential for Future Expansions. Paper presented at the CHILDES, TalkBank, Competition, Emergentism: Honoring the Impact of Brian MacWhinney on Language Research, Pittsburgh, PA.

Rose, Yvan, and Brian MacWhinney. 2014. *The PhonBank Project*. https://doi.org/10.1093/oxfordhb/9780199571932.013.023.

Rose, Yvan, Brian MacWhinney, Rodrigue Byrne, Gregory Hedlund, Keith Maddocks, Philip O'Brien and Todd Wareham. 2006. Introducing Phon: A software solution for the study of phonological acquisition. In *Proceedings of the 30th Annual Boston university conference on language development,* ed. David Bamman, Tatiana Magnitskaia, and Colleen Zaller, 489-500. Somerville, MA: Cascadilla Press.

Thomas, Erik R. 2013. Phonetic Analysis in Sociolinguistics. In *Research Methods in Sociolinguistics: A Practical Guide*, ed. Janet Holmes and Kirk Hazen, 119-135. John Wiley and Sons, Inc.

Tennant, Jeff, David Heap, Angelica Hernandez and Jorge Emilio Rosés Labrada. 2015. Aspiration vs. Deletion of /s/ in Contemporary Eastern Cuban Spanish: Differing Constraints. *Poster presented at annual conference of New Ways of Analyzing Variation 44,* Toronto, ON.

Voormann, Holger, and Ulrike Gut. 2008. Agile Corpus Creation. *Corpus Linguistics and Linguistic Theory* 4 (2): 1-16. https://doi.org/10.1515/CLLT.2008.010.

Winkelmann, Raphael, Jonathan Harrington, and Klaus Jänsch. 2017. EMU-SDMS: Advanced Speech Database Management and Analysis in R. *Computer Speech and Language* 45: 392–410. https://doi.org/10.1016/j.csl.2017.01.002.