Article:

Montgomery, C. orcid.org/0000-0002-3017-2909, Walker, G. and Woods, H. (2025) Salient language in context (SLIC): a web app for collecting real-time attention data in response to audio samples. Linguistics Vanguard. ISSN 2199-174X

https://doi.org/10.1515/lingvan-2025-0028

Chris Montgomery*, Gareth Walker and Harry Woods

Salient Language in Context (SLIC): a web app for collecting real-time attention data in response to audio samples

https://doi.org/10.1515/lingvan-2025-0028
Received February 22, 2025; accepted May 29, 2025; published online August 12, 2025

Abstract: This article introduces Salient Language in Context (SLIC), a web application designed for collecting real-time attention data from listeners in response to audio samples. SLIC enables researchers to capture timestamped attention points as listeners engage with spoken language, allowing for an innovative approach to studying real-time perception of speech. The application facilitates the collection of both quantitative and qualitative data, providing researchers with not only precise reaction timing but also listener justifications for their responses. SLIC addresses long-standing challenges in sociolinguistics and perceptual dialectology by enabling researchers to examine which linguistic features capture attention, when they do so, and how listeners interpret them in context. The technical aspects of SLIC are discussed in the article and include its secure online infrastructure and integration with scripts written for Praat, which allow for efficient data visualization and interpretation. A case study using SLIC demonstrates its effectiveness in examining listener perceptions of segmental features in Liverpool English, revealing distinct patterns of attention and commentary strategies. By providing a robust method for investigating real-time language perception, SLIC has a broad range of potential applications.

Keywords: real-time attention data; sociolinguistics; speech perception; web-based research; linguistic salience

1 Introduction

This article describes the functionality of the Salient Language in Context web app (SLIC). SLIC was conceived and commissioned by the first author of this paper, and written by a team including the third author. In this article we also introduce companion scripts for processing and visualizing results of experiments using SLIC that have been written by the second author. SLIC is currently free to use for non-commercial research purposes. The companion scripts are made available for use under the GNU General Public License (http://www.gnu.org/licenses/). SLIC has the capacity to allow a step-change in understanding listeners' attention to spoken language, as well as in the ways in which data can be extracted and analysed. A brief use case is presented at the end of the article in order to illustrate how SLIC can be used to isolate features of import to listeners, to understand when these features are attended to.

The specific features that drive listener reactions to speakers have been an object of research in sociolinguistics for many years. Campbell-Kibler (2006: 64) notes that in attitudes studies "it can be difficult to establish which aspects of the speech trigger which aspects of the evaluation". Attempts to identify features that are important for listener evaluation or classification have been varied and numerous. Research has generally taken two directions: some research preselects features based on the results of production studies (e.g., Campbell-Kibler 2008; Lawrence 2015; Pharao et al. 2014; Plichta and Preston 2005), while other research seeks to understand levels

***Corresponding author: Chris Montgomery**, School of English, The University of Sheffield, Jessop West, Upper Hannover Street, Sheffield, South Yorkshire, S3 7RA, UK, E-mail: c.montgomery@sheffield.ac.uk
Gareth Walker and Harry Woods, School of English, The University of Sheffield, Jessop West, Upper Hannover Street, Sheffield, South Yorkshire, S3 7RA, UK

of attention to features after listeners have heard a speaker (e.g., Braber et al. 2022; Mackay et al. 2024) or as they hear a speaker (e.g., Levon et al. 2022; Soukup 2011; Watson and Clark 2013, 2015).

Research using preselection generally examines features that have been demonstrated to index a stable social meaning in production studies and manipulates these in some way to examine the effect on listeners. Plichta and Preston (2005), for example, synthesized increasingly monophthongal tokens of /ay/ and played them to listeners who were asked to locate them along a north–south continuum in the United States. Other studies have taken a less abstract approach, embedding features of interest in longer stretches of talk. Labov et al. (2011; replicated in a UK context by Levon and Fox 2014) examined the effect of <ing> realization and the extent to which the frequency of the apical variant [n] would affect "professionalism" ratings for a speaker assumed to be auditioning for a job as a newsreader amongst listeners in the USA. Pharao et al. (2014) examined the role of fronted /s/ in Copenhagen, finding that contextual information (i.e., the features that co-occurred with the target of the research) were important in determining the evaluative reactions to speakers. Pharao and Maegaard (2017) built on these findings, demonstrating the importance of both the order and co-occurrence of fronted /s/ and palatalized variants of /t/ for listener perceptions.

Questions remain about the extent to which listeners are aware of the features focused on in the research above. There have been some attempts to address this. Braber et al. (2022) asked listeners to write down features of note after they heard a speaker, although they found that when asked to recall features after listening to a whole a sample, listeners "did not have the meta-linguistic ability to describe linguistic features of the voices they are hearing" (Braber et al. 2022: 287). Mackay et al. (2024) added a post hoc attention task to a version of Labov et al.'s (2011) experiment, finding variable listener attention to the manipulated features in their task.

The use of post hoc tasks means that the point at which listeners arrive at an impression of the speaker is unclear, and can also introduce problems associated with recall of features. Some research has therefore attempted to examine attention in real time and investigate in-the-moment listener perceptions. An early attempt to examine this can be seen in Soukup's (2011) work in Austria, in which listeners were encouraged to annotate a transcript as they heard speakers. Other approaches have been developed using audience response methods from other fields, as seen in Watson and Clark (2013, 2015) and Levon et al. (2022). These approaches use a slider presented on screen as listeners hear voice samples. Listeners are instructed to move this slider left or right according to the extent to which a percept was cued as they heard a speaker (e.g., listeners are asked to move the slider along a "poshness" scale). Reviewing this method, Austen and Campbell-Kibler (2022) found expected results, but "poor time granularity" (Austen and Campbell-Kibler 2022: e128), meaning that features of import for listeners are often inaccessible to researchers.

The literature reviewed above demonstrates that when it comes to responding to Campbell-Kibler's (2006: 64) initial problem of finding the important features for listeners, there remains a gap. Research is still some way from:

1. Isolating the features of import to listeners
2. Understanding when these features are attended to
3. Understanding the impact that these features have when heard in context

SLIC is designed to address this gap. SLIC provides a platform in which experiments can be designed that allow researchers to capture real-time attention points collected as listeners hear an audio sample, and for listeners to subsequently review these attention points in order to justify why they added them. This provides granular time detail for reactions, and adds certainty about these reactions. This responds to the problems associated with recall in post hoc tests, and addresses the time lag problem inherent to slider-based audience response methods. SLIC is the only approach we are aware of that is able to collect this type of data.

A pilot version of SLIC was used in Montgomery and Moore (2018) and Moore and Montgomery (2018). That research demonstrated that the approach is able to collect meaningful data, allowing listeners to be more specific about what they have attended to as they heard a sample. SLIC is being used to collect data on listener attention to speech in research in England, an excerpt of which forms the basis of the use case below, and in Denmark (Boas et al. 2024). As well as important improvements to the pilot tool used in Montgomery and Moore (2018), SLIC also has companion tools for data visualization which we present below after describing SLIC.

2 Overview of SLIC

SLIC provides an environment in which researchers can conduct experiments that gather listeners' perceptions of speakers. SLIC affords researchers the ability to host sound files and transcripts, and to design online questionnaires in an easy-to-use graphical interface. A variety of question types are available to researchers. These different question types mean that SLIC is well suited to research in the language attitudes/language regard paradigms, and has the potential to be used in other applied fields (e.g., speech and language therapy training).

Once real-time experiments are set up, SLIC provides an interface through which listeners can play an audio sample and use a "click" button each time they attend to a feature of the sample they are listening to (Figure 1). This allows researchers to collect real-time attention data (RTAD) from listeners as a number of discrete timestamps. SLIC has a review and comment function that uses the timestamped data points to display each attention point back to the listener, accompanied with a fragment of the audio file that they heard leading up to their click, and a transcript fragment of the same. Listeners are invited to use a text box to type an explanation for why they clicked when they did, or to use tick boxes to indicate an error or that they are unsure of why they clicked when they did (Figure 2). Listeners can also add a comment after listening to the whole sample if they did not click. The resulting dataset comprises a set of time-aligned attention points accompanied by commentary data for each of these points.

3 Technical aspects

SLIC is hosted online and runs in standard web browsers on computers, tablets, and smartphones. It can therefore be deployed remotely over the web or in a lab-based environment. SLIC is written in the Ruby programming language using the rails framework (Ruby on Rails). It is hosted on a virtual Ubuntu server hosted by the University of Sheffield. Server security is kept up to date. The platform does not collect any personally identifiable information (PII) from research participants except their email address, and this is only used to collect and send a user their information should they request it. This is required for compliance with the UK General Data Protection Regulation (GDPR). All website traffic is sent over SSL (Secure Sockets Layer) so it cannot be seen by third parties.

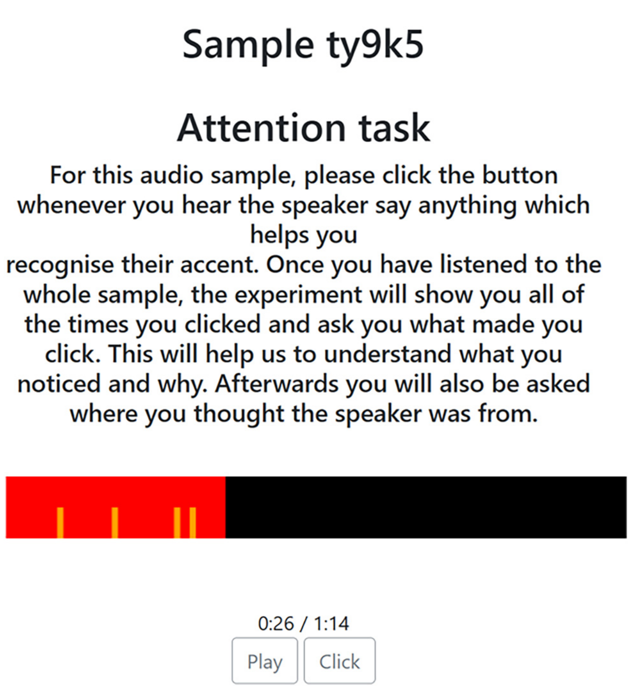

Figure 1: The SLIC real-time interface.

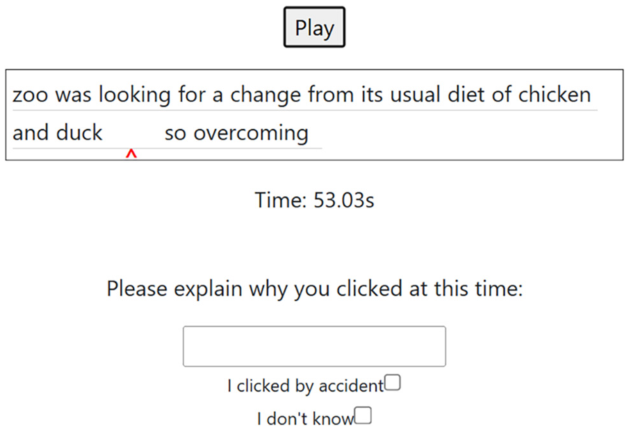

Figure 2: The SLIC review interface.

All PII in the database, most of which concerns the researcher running the experiment, is encrypted so that it is illegible should the server be compromised. Since researchers cannot know in advance if participants' responses will contain PII, these are also encrypted and hidden from the server logs. Participants' responses are only made available to the researcher who collected them. Researchers using SLIC therefore need to be able to justify their reasons for collecting that data when they create the experiment, to follow local ethics procedures, and to ensure that data remains GDPR compliant if they remove that data from the platform. Researcher PII is only available to the researcher themself and to the SLIC admin (currently the first author of this article).

4 Set-up

SLIC is available at https://slic.shef.ac.uk, which provides full user documentation and links to the companion scripts we describe below. SLIC is currently free to use for non-commercial research purposes, on the condition that this paper is cited in any publication, presentation, report, or other form of research output when it has been used to collect data. Researchers must create an account to begin using SLIC. Account creation at present requires an invitation code, which can be requested by contacting the first author of this paper via email. Once the invitation code is provided, researchers are required to enter some brief details, including country, institution (if any), and academic profile, as well as to create login details (email and password). Once their account is created, researchers are able to access the full functionality of SLIC.

New experiments are created in SLIC by selecting the "New experiment" option, and entering an experiment title and description. From there, researchers are able to add consent terms in order to comply with ethics procedures, and to generate the different elements of their experiment. Researchers can add two types of element: "question pages" or "audio samples". Researchers can then configure the elements of the experiment to suit their needs.

For "question pages" a variety of question types are available to researchers, including text questions, option questions (using drop-down boxes, radio buttons, or checkboxes), and ratings scales. This means that questions can be added to collect a range of different types of information from listeners.

The "audio samples" options provide researchers with the ability to upload audio files and transcripts. In order for the real-time capabilities to function, audio samples and transcript files must be uploaded. Audio samples must be WAV files and corresponding transcripts must be prepared at the word level in ELAN (2025), and saved as EAF files. For each audio sample, researchers can specify if they wish it to be presented as an "audio hearing" task with no real-time click data requested, or as a "reaction task" with the real-time click and review functionality included. The listening prompt can be specified by the user for the reactions tasks, allowing them to request button clicks for whatever they wish listeners to focus on. In the use case below, listeners were provided with the prompt that can be seen in Figure 1. Researchers may also wish to provide a preview of the interface for listeners before they encounter samples that are the object of the research: to do this, a radio button can be clicked during set-up to indicate that an audio sample is a “calibration sample”.

Once question pages and audio samples have been created, researchers can review the whole experiment and can choose to randomize specific elements within it (Figure 3). This permits researchers to maintain the order of certain elements (e.g., biographical details and a calibration sample) whilst randomizing others (e.g., the presentation of audio samples).

Once deployed, experiments are available in the researcher’s dashboard. Clicking on a deployed experiment provides details of the total responses, completed responses, a percentage completion rate, and the timestamp of the last response. There is also a button that will download the results of the experiment. A dummy experiment is available at https://slic.shef.ac.uk/868a1107-eda5-45e3-baeb-9d726f9113f8/home in order to demonstrate SLIC from the point of view of listeners.

5 Processing results

5.1 Format of the results files

SLIC outputs a Microsoft Excel XLSX file. For illustrative purposes, a real sample results file, with changes to ages and the removal of some data to preserve the anonymity of the listeners, is available at https://docs.google.com/spreadsheets/d/1rQgTTbMe3Ad3flJXlRFSUmTFS4xiWqhW. The file contains responses to audio recordings of speakers from five English cities as well as the calibration sample. Tabs in results files correspond to questions asked of listeners. In the sample file, Tab 1 contains responses given to demographic questions. Tabs 2, 4, 6, 9, and 11 contain information about each listener’s clicks and comments, arranged by listener ID, in response to five audio samples. These tabs contain the following information: the time at which the click occurred (_*time*), the listener ID of the listener who clicked (_*uid*), the comment that they subsequently made on reviewing their clicks (_*comment*), columns to show that, on review, the listener decided they did not know why they clicked (_*dk*) or felt that they had clicked by accident (_*accident*), and a column to show that a comment had been given after hearing the whole sample despite no click being made over its course (_*no_click*). Tabs 3, 5, 7, 10, and 12 contain listeners’

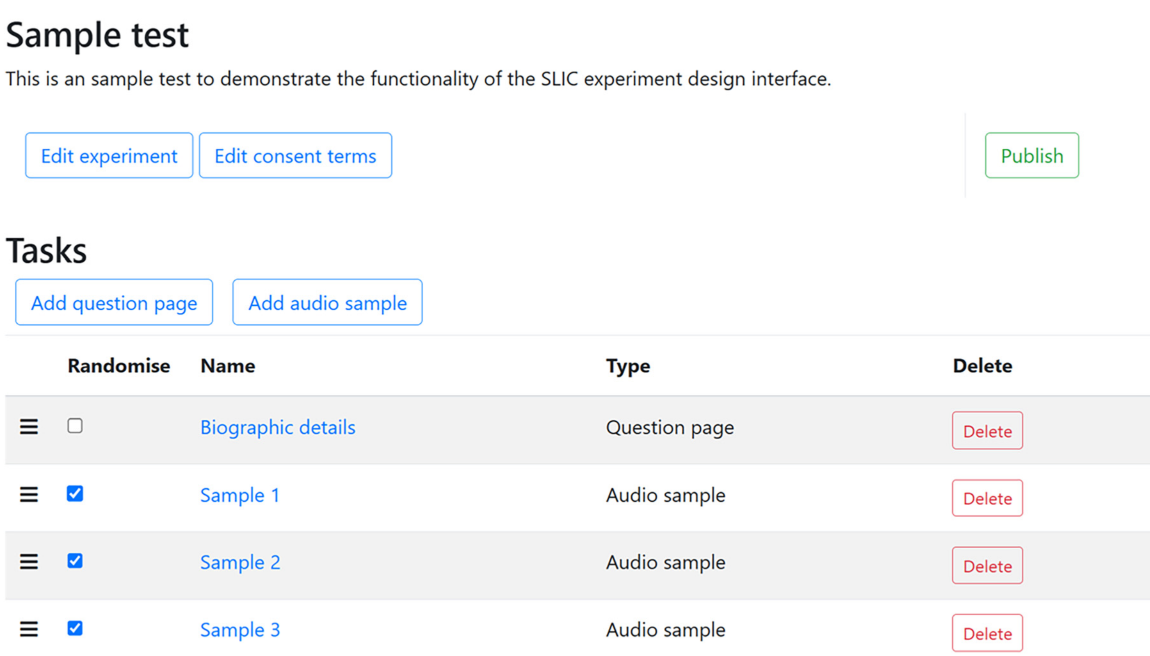

Figure 3: Experiment builder in SLIC.

answers to questions about where they thought the speaker was from. Tab 8 contains responses to a question asking listeners for any other information they wanted to share. Tab 13 provides timestamped click data for the calibration sample, and Tab 14 contains listeners' guesses for the calibration sample provenance.

5.2 Viewing the results

We provide resources (scripts and documentation) to create visual representations of data from SLIC in Praat (Boersma and Weenick 2025). The scripts and documentation are available via the SLIC website. Currently we provide resources to allow RTAD to be inspected in dynamic editor windows alongside, for example, waveforms, spectrograms and pitch traces (Section 5.2.1), to prepare static publication-ready graphics showing RTAD (Section 5.2.2), and to prepare TextGrids containing click and comment data (makeReactionTextGrids.praat: see the SLIC website).

To read RTAD into Praat, the relevant spreadsheet tab needs to be saved as either a comma-separated or tab-separated plain text file (CSV or TSV file, respectively). This can be done via the save options of spreadsheet software (e.g., Microsoft Excel, LibreOffice). The comma-separated file of reactions to one sample, from a Liverpool speaker (Tab 9 of the sample results file), is provided at https://drive.google.com/file/d/167cLWZ1AMi7cHjPERhiKLI9iUyaxLXrA. For the purposes of this discussion, all click data from that tab has been retained and will be plotted. This includes cases where listeners have clicked but not offered a comment, where they have reported that they do not know why they clicked, where they have reported that they clicked by accident, and where they have made a comment but not clicked.

5.2.1 Viewing real-time attention data in dynamic editor windows

The script makeReactionRealTiers.praat allows RTAD to be viewed in dynamic editor windows in Praat. The script runs on a Sound object (e.g., containing the sample that listeners reacted to) and a Table object in Praat, such as that created by reading in the CSV file referred to above, via the Open menu in Praat, then "Read table from comma-separated file…". With both the Sound and the Table objects selected in the List of Objects, the script can be run. On running the script, a dialog box appears which allows various settings to be changed; these settings are described in the script documentation. Once the script has finished running successfully, two new objects will be created: a Table containing counts of the reactions, and a RealTier based on those counts. The RealTier can be viewed in a Praat editor window (the RealTierEditor) to scroll, zoom in and out, make selections, and get information about the numbers of reactions at any point. (Researchers may find it helpful to set the range of the RealTierEditor window: from the RealTier window, RealTier > Set range….) The RealTierEditor window can be moved synchronously with other editor windows by using Praat's Group function. Figure 4 shows a SoundEditor window containing the waveform and spectrogram of the audio sample in the top panel; the bottom panel shows listeners' reactions in a dynamic RealTierEditor.

Clicking on the top of the highest peak in the RealTierEditor window allows the user to see that there were 43 reactions within ±0.25 s of the cursor (a bin width of 0.5 s is the default value in the script; it can be changed via the dialog box which appears when running the script, or changed in the script itself). Considering what is happening at the same point in the SoundEditor window allows the user to see that this peak has occurred after the speaker has produced "chicken and duck" with noticeably fricated productions of /k/.

5.2.2 Presenting real-time attention data as graphics

The script drawReactionData.praat allows reaction data to be drawn to the Praat picture window in order to create publication-quality graphics runs on a Table object with the same format as required to view the data in dynamic windows, as described above. Running the script on the same Table object as before, adjusting the options in a dialog box which pops up on running the script in order to select draw just a portion of the data, and

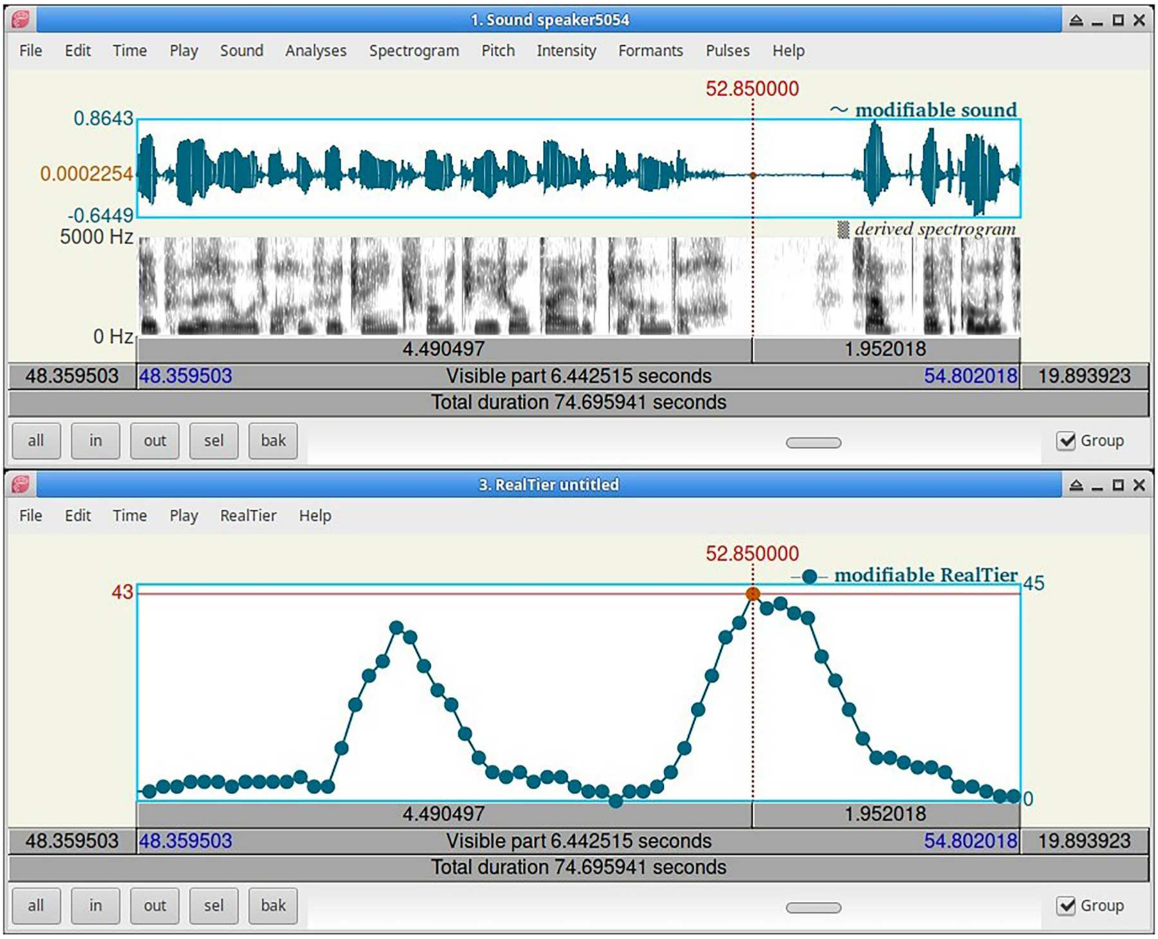

Figure 4: Screenshot showing Praat's SoundEditor window (top) containing a waveform and spectrogram of a speech sample, and RealTierEditor window (bottom) containing listeners' reactions to the speech sample.

adjusting the drawing of the *y*-axis produces the image Figure 5, which is equivalent to what was shown in a dynamic editor window in Figure 4.

By selecting a Sound object (the audio sample to which listeners were reacting) and/or a Spectrogram object, the script can be used to draw these acoustic records alongside the plot of the RTAD, as shown in Figure 6 below – which is in effect a publication-ready version of the screenshot in Figure 4.

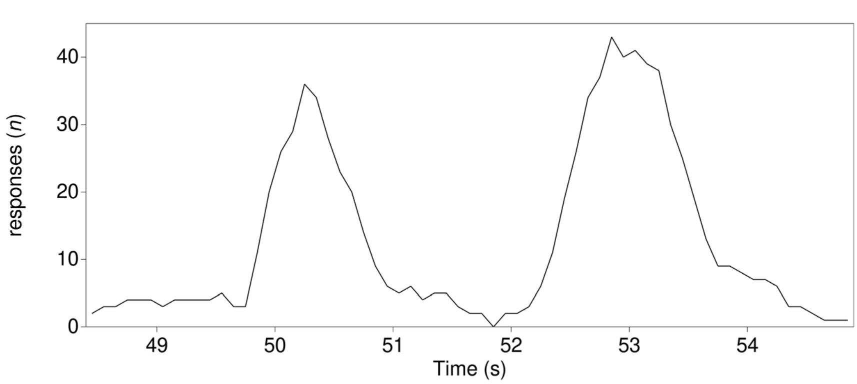

Figure 5: Static image of listeners' reactions to a speech sample.

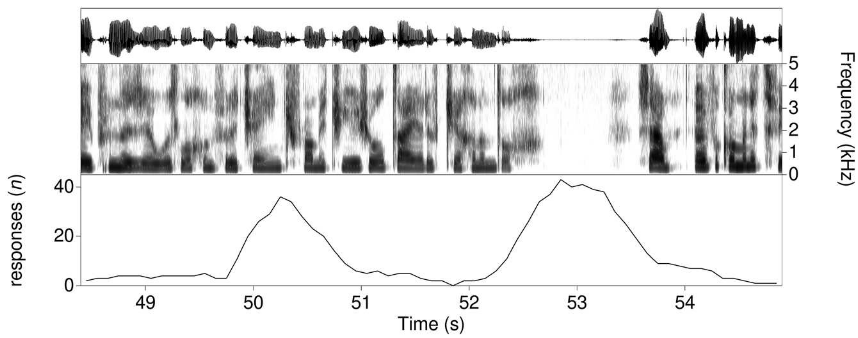

Figure 6: Static image of listeners' reactions to a speech sample (*bottom panel*), accompanied by a waveform (*top panel*) and spectrogram (*middle panel*).

On running the script, a dialog box appears allowing the user to change further options, including the style and colour of the lines for the number of reactions. These options allow for multiple sets of results to be drawn, such as lines for different groups of listeners. The options are described more fully in the script documentation linked to above.

6 A use case

We now present a brief use case using SLIC from data collected in England that focuses on listener attention to a feature of Liverpool speech, as part of a larger real-time attention study of speech from different cities. The experiment was hosted on Prolific and used five voice samples from the English Dialects App Corpus (Leemann et al. 2018), comprising speakers reading the short story "The Boy Who Cried Wolf" (Deterding 2006: 193). One hundred and one listeners were recruited, split into two groups: those from Yorkshire and the Humber (49) and those from the South East of England (52). The listener sample had an equal gender split and had a mean age of 41.9 years old. Education level was skewed towards higher education or equivalent: 55 listeners had a degree or higher qualification while 46 did not.

Segmental coding was performed independently by two researchers,[1] with differences flagged and discussed. A high bar was set for segmental coding, which resulted in 21.17 % of the clicks made by listeners being usable (167 of 789 clicks). Only those clicks that could be definitively identified by the two coders as relating to a particular segmental feature are reported on here. Thematic coding was performed by a single researcher, and was able to include a greater number of clicks as the criteria for inclusion was less strict (i.e., a comment such as "pronunciation of mountain" was unusable in relation to the segmental data but could be coded as "general phonetic commentary" for the thematic coding). The thematic coding resulted in 549 usable clicks (69.58 %). The thematic coding revealed five strategies used by listeners in comment data, which we return to below.

Here we concentrate only on attention to /k/, and in particular the strategies used by listeners in their review comments for their clicks. This feature was the most likely to be attended to for the Liverpool speaker, attracting 56 clicks. This is perhaps unsurprising given the social meanings associated with the specific (lenited) realization of the feature in Liverpool accents (see Honeybone 2001; Watson 2007). Figure 7 shows the distribution of those clicks. It reveals variable levels of attention over the time course of the sample, with the largest peak occurring at the point in the sample at which there were two instances of heavily affricated /k/ in quick succession, as noted in Section 5.2.1.

The comments on the clicks for /k/ revealed a range of strategies used by listeners for this feature for the Liverpool speaker, as follows (note that these are the original comments, with no modifications made in their reproduction here):

1. Explicit phonetic descriptions: e.g., "the ck in duck" (Listener 100, responding to token d in Figure 7)
2. General phonetic descriptions: e.g., "change and duck the vowels seem elongated" (Listener 26)

1 The two researchers were Dr. Chris Montgomery and Dr. Hielke Vriesendorp.

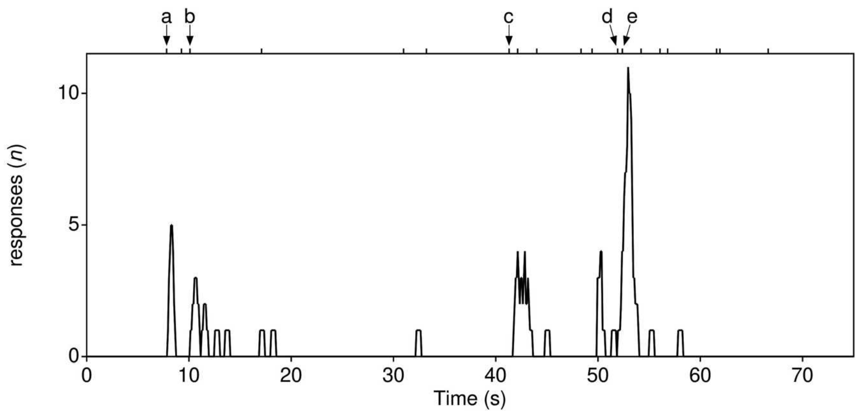

Figure 7: Attention to coded instances of /k/ in the Liverpool sample (n = 56) for all listeners. A secondary x-axis running along the top of the chart shows the time of each occurrence of /k/ in the sample, with selected tokens indicated with a letter.

3. Metalinguistic commentary: e.g., "Flocks (not pronouncing the 'ck' clearly" (Listener 60, responding to token a in Figure 7)
4. Phonetic respellings: e.g., "exachly the same" (Listener 33, responding to token c in Figure 7); "Chicken & duck' sounds like chihen & duhk" (Listener 60, responding to the tokens d and e in Figure 7)
5. Relational comparisons: e.g., "the gutteral rk sound or dark indicates a liverpool accent" (Listener 23, responding to token b in Figure 7)

Although we do not have space to document them here, there were important differences in the number of comments provided by listeners from the two different regions, as well as differences in the types of strategies used. Taken together, though, the coded data, Figure 7, and the thematic analysis of comment data demonstrate the ability of SLIC to address the research gaps we identified above.

7 Conclusions

Over the course of this paper we have described SLIC and companion scripts for data interpretation and visualization created by the authors of this article. SLIC is the only method of its type and we have shown how it is configured and used to create experiments that can access RTAD for audio samples. We have introduced a number of companion tools that permit analysis in Praat in order to rapidly investigate the distribution of click data. The use case demonstrates the RTAD that SLIC is able to collect and presents two approaches to analysing comment data, either to focus on explicit attention to segmental features, or to understand the metalinguistic strategies used by listeners when adding comments explaining their clicks. We have illustrated the use of SLIC in a study of responses to segmental features of a single urban accent of English, though we foresee the application of SLIC in many research contexts where researchers wish to collect RTAD in response to audio samples, and especially where they want to supplement reaction data with comments from listeners.

Acknowledgments: We would like to thank Dr. Hielke Vriesendorp (Utrecht University) for assisting with data coding for the use case presented in this article. We are also grateful to the University of Sheffield's Software Hut initiative and to Prof. Emma Norling (University of Sheffield) for their support in the development of the SLIC web app. We would also like to thank the two anonymous reviewers for their insightful comments and suggestions. Any remaining errors are our own.

References

Austen, Martha & Kathryn Campbell-Kibler. 2022. Real-time speaker evaluation: How useful is it, and what does it measure? *Language* 98(2). e108–e130.

Boas, Kirstine, Inger Schoonderbeek Hansen & Ea Overgaard Lindhardt. 2024. Testing for salience and stereotypes in real-time Funen speech. In *Paper presented at ICLaVE, Vienna, Austria, 8–11 July*.

Boersma, Paul & David Weenick. 2025. Praat: Doing phonetics by computer [computer software]. http://www.praat.org/ (accessed 12 April 2025).

Braber, Natalie, Harriet Smith, David Wright, Alexander Hardy & Jeremy Robson. 2022. Assessing the specificity and accuracy of accent judgments by lay listeners. *Language and Speech* 66(2). 267–290.

Campbell-Kibler, Kathryn. 2006. *Listener perceptions of sociolinguistic variables: The case of (ING)*. Stanford, CA: Stanford University PhD thesis.

Campbell-Kibler, Kathryn. 2008. I'll be the judge of that: Diversity in social perceptions of (ING). *Language in Society* 37(5). 637–659.

Deterding, David. 2006. The North Wind versus a Wolf: Short texts for the description and measurement of English pronunciation. *Journal of the International Phonetic Association* 36(2). 187–196.

ELAN. 2025. *version 6.9 [computer program]*. Nijmegen: Max Planck Institute for Psycholinguistics, The Language Archive. https://archive.mpi.nl/tla/elan.

Honeybone, Patrick. 2001. Lenition inhibition in Liverpool English. *English Language and Linguistics* 5(2). 213–249.

Labov, William, Sharon Ash, Maya Ravindranath, Tracey Weldon, Maciej Baranowski & Naomi Nagy. 2011. Properties of the sociolinguistic monitor. *Journal of Sociolinguistics* 15(4). 431–463.

Lawrence, Daniel. 2015. Phonetic variation as a cue to regional identity: An experimental approach. *York Papers in Linguistics* 2. 42–63.

Leemann, Adrian, Marie-José Kolly & David Britain. 2018. The English Dialects App: The creation of a crowdsourced dialect corpus. *Ampersand* 5. 1–17.

Levon, Erez & Sue Fox. 2014. Social salience and the sociolinguistic monitor: A case study of ING and TH-fronting in Britain. *Journal of English Linguistics* 42(3). 185–217.

Levon, Erez, Devyani Sharma & Yang Ye. 2022. Dynamic sociolinguistic processing: Real-time changes in judgments of speaker competence. *Language* 98(4). 749–774.

Mackay, Brad, Erik Schleef & Jana Pflaeging. 2024. Awareness and sociolinguistic monitoring: Comparing (ing), (t)-deletion, *like, you know*, filled and unfilled pauses. In *Paper presented at ICLaVE, Vienna, Austria, 8–11 July*.

Montgomery, Chris & Emma Moore. 2018. Evaluating S(c)illy voices: The effects of salience, stereotypes, and co-present language variables on real-time reactions to regional speech. *Language* 94(3). 629–661.

Moore, Emma & Chris Montgomery. 2018. The dialect of the Isles of Scilly: Exploring the relationship between language production and language perception in a Southern insular variety. In Laura Wright (ed.), *Southern English varieties then and now*, 39–73. Berlin: Mouton de Gruyter.

Pharao, Nicolai & Marie Maegaard. 2017. On the influence of coronal sibilants and stops on the perception of social meanings in Copenhagen Danish. *Linguistics* 55(5). 1141–1167.

Pharao, Nicolai, Marie Maegaard, Janus Spindler Møller & Tore Kristiansen. 2014. Indexical meanings of [s+] among Copenhagen youth: Social perception of a phonetic variant in different prosodic contexts. *Language in Society* 43(1). 1–31.

Plichta, Bartlomiej & Dennis R. Preston. 2005. The /ay/s have it: The perception of /ay/ as a north-south stereotype in United States English. *Acta Linguistica Hafniensia* 37(1). 107–130.

Soukup, Barbara. 2011. Austrian listeners' perceptions of standard-dialect style-shifting: An empirical approach. *Journal of Sociolinguistics* 15(3). 347–365.

Watson, Kevin. 2007. Liverpool English. *Journal of the International Phonetic Association* 37(3). 351–360.

Watson, Kevin & Lynn Clark. 2013. How salient is the NURSE~SQUARE merger? *English Language and Linguistics* 17(02). 297–323.

Watson, Kevin & Lynn Clark. 2015. Exploring listeners' real-time reactions to regional accents. *Language Awareness* 24(1). 38–59.