

# Dialogs Re-enacted Across Languages

Nigel G. Ward, Jonathan E. Avila, Emilia Rivas

Department of Computer Science, University of Texas at El Paso

Draft of November 18, 2022

## Summary

To support machine learning of cross-language prosodic mappings and other ways to improve speech-to-speech translation, we present a protocol for collecting closely matched pairs of utterances across languages, a description of the resulting data collection, and some observations and musings. This report is intended for:

- people using this corpus
- people extending this corpus
- people designing similar collections of bilingual dialog data

## Contents

- 1 Motivation
- 2 Related Work and Aims
- 3 Strategy
- 4 Data Collection Procedure
  - 4.1 overview
  - 4.2 participants
  - 4.3 equipment and configuration
  - 4.4 step-by-step procedure
  - 4.5 post-processing
- 5 Data Released to Date
- 6 Observations
- 7 Towards Increased Variety
- 8 Cost Components
- 9 Other-Language Explorations
- 10 An Exploration in Remote Collection

## 1 Motivation

Future speech-to-speech translation systems should be able to effectively support dialog, but current systems are not specifically built for this. One reason is the lack of bilingual data on how dialog behaviors differ across languages. Thus there is a need for parallel corpora of conversational speech across languages. This report presents the Dialogs Re-enacted Across Languages (DRAL) protocol and corpus. The key idea is gathering unscripted conversations among bilingual speakers who later re-enact selected utterances from those conversations in their other language.

We hope that this data will be useful for many purposes, but it is worth noting some specific likely uses. First, we ourselves plan to use these utterance pairs to train models for mapping Spanish to English prosody, and conversely. For this we will try modeling phrases both in isolation and considering the local context. Second, this data could also support linguistic studies of the prosodic differences between the two languages. Third, it could support study of other language phenomena, such as the meanings and translation of dialog markers and discourse connectives [15]. Fourth, this data could be used to fine-tune translation models trained on large, non-dialog data. Finally, this data could be used to evaluate the performance of different speech-to-speech translation systems for dialog relevant scenarios.

We acknowledge that this data collection method is expensive and could never produce corpora large enough to train lexical translation models or language models. However small, well-designed corpora can be very informative regarding many aspects of conversational behavior [9, 6, 7].

## 2 Related work and corpus desiderata

As already noted, an important use case for future speech-to-speech translation systems is support for people in dialog, as they interact across languages. While speech-to-speech translation systems so far have been developed using mostly monologue speech data, participants in dialog seldom speak in monologue style; rather their utterances are often pragmatically rich, as the speakers take turns, shift topics, take stances, try to amuse or persuade, and so on. Systems which are blind to such intents will be of only limited value for supporting dialog. Corpus-based approaches could help us model these behaviors across languages, but corpora of equivalent spoken dialog data across languages have been lacking.

Classic corpora for speech translation research, as surveyed by [8], have been largely designed to support research on speech-to-text translation. However the development of true speech-to-speech systems requires corpora that include target-language audio.

Recent years have seen the development of several true multilingual speech corpora, often motivated by the need for data to train end-to-end speech-to-speech translation systems [11, 14, 25, 13]. Most such corpora are exclusively monologue and feature mostly read speech, such as short prompts, Wikipedia articles, books, and religious texts [1, 20, 17, 3]. As reading lacks interactive aspects, these corpora exhibit no dialog phenomena and little pragmatic richness. Even monologue corpora which are not just reading, for example political discussions and informative talks [21, 5, 18], exhibit only modest style variation. Parallel corpora derived from dubbed television shows [16, 2], exhibit more pragmatic variety, but these are scripted, and thus neither spontaneous nor truly interactive, and are moreover performed by professional actors and interpreters, and thus depart significantly from natural conversational behavior. Some interesting corpora have been described but are unfortunately not available [8]. The corpora which come closest to meeting the need are those which start with actual spoken conversations and include target-language equivalents produced by translating the dialog transcripts and then feeding the resulting target-language text to a speech synthesizer [25, 10]. While suitable for some purposes, the synthesized outputs lack meaningful prosodic variation, and in that respect, at least, they fail to faithfully match the intent of the original.

Thus, although useful in various ways, no existing resources support development of speech-to-speech translation systems for dialog scenarios. In particular, there is an unmet need for corpora that are natural, faithful, exhibit wide pragmatic diversity, and are available.

### 3 Strategy

For the sake of both pragmatic diversity and naturalness, we start by recording nonprofessional speakers in unstructured conversations. We then collect re-enactments in another language.

For the sake of fidelity, we use bilingual speakers. Thus each matched utterance pair in our corpus was the product of a single speaker. This avoids the need at modeling time to design mechanisms to disentangle differences due to speaker variation from differences due to the language spoken.

Having each speaker re-enact their own previous utterances has the further advantage of increasing the likelihood of faithfully recognizing or recalling the intent of the original utterance, with all of its nuances.

Achieving reasonable naturalness required some thought on how to obtain the re-enactments. In theory we could have had participants re-enact an entire conversation in one go, but this would have required them to transcribe and translate the words as a memory aid, and likely to go through multiple passes of rehearsal, making the result more like read speech and less spontaneous [19]. Therefore we chose to have them re-enact the conversation fragment-by-fragment. While we considered having participants re-enact utterances as they went along, turn by turn, this would have broken the flow of the conversation. Thus we had them do the re-enactments immediately after the original conversation finished.

While participants might, in theory, be able to manage everything themselves, we chose to have each session moderated by a “director,” variously below also called the “operator” and “producer,” as one person performed all these roles.

Choosing the length of the fragments involved a trade-off: too long, and it would be hard for the participants to retain the full intention in working memory; too short, and the context would be lacking. We did not set any fixed target, but let the operator and participants find what lengths worked well. In the data collected so far, the re-enactments average about 3.7 seconds.

## 4 Data collection procedure

### 4.1 Overview

To reiterate, the basic idea is that pairs of participants have a short conversation in one language and reenact fragments of the conversation in another language. An operator handles the recording and guides the participants.

To help participants get in the mood, at each stage the operator instructs them in the language they are about to be using.

Since overlap is common in spontaneous conversation, we chose to record the speakers in separate channels, with the best audio separation we could reasonably attain. For this reason, we have participants wear head-mounted headsets, and sit in separate rooms. The signals from the microphones are fed into a stereo digital recorder and also played back to the headphones of the other speaker so they can hear each other.

The original conversations last about 10 minutes. If they were longer, we would be gathering more data than we could ever re-enact. If they were shorter, participants would not have time to get comfortable with each other and get into interesting topics. Also if they were shorter, we would at times lack enough interesting utterances to re-enact. We also find that 10 minutes of free conversation puts them in a good mood and perhaps more willing to work hard on the difficult task of re-enacting utterances.

The original conversations are done with the participants in two rooms which have audio separation but a glass window so they can see each other. We close the doors to the main room not only to improve audio separation but to enable them to have a private conversation, without feeling monitored or feeling the pressure to put on a performance.

For the re-enactment stage, the doors to the small rooms are opened so that the participants can interact naturally with the operator; this makes coordination easier, although at some cost in the quality of the audio separation. After identifying any segments of the conversation that the participants want redacted, the operator scrubs through the recording, using ELAN, and together with the participants, chooses utterances or utterance pairs for re-enactment.

Selecting utterances for re-enactment involves many factors. Since we do not have a specific use case in mind, our primary consideration is “interestingness.” Uninteresting utterances are those of types of which we already have many in the corpus, such as initial formalities and utterances from the midst of telling stories. Thus the operator preferentially chooses fragments that are interesting in the sense of exhibiting behaviors or attitudes less commonly seen. These are also typically more interesting for the participants. As a rule of thumb, utterances in regions of active engagement, often marked by the presence of laughter, tend to be interesting. Also, for the sake of diversity of speakers, and to keep both participants engaged, when one person dominated the original conversation, the operator attempts to balance out the re-enactments across both speakers.

Another consideration in utterance selection is the turn-taking. On the one hand, we generally avoid sequences with rapid back-and-forth turn-taking, both because they are hard to replicate and because they are unlikely to occur in any use case of conversations mediated by a system. On the other hand, we often choose fragments that have speech from both participants. This helped keep both participants engaged throughout the process. In most such fragments, one participant’s contribution was only a back-channel, laughter, or short response. Such short utterances are very hard to produce realistically in isolation, so re-enacting them in context helped improve their fidelity and naturalness. Further, including these served to help encourage the main speaker to re-create the exact same feeling in order to elicit the exactly equivalent response.

The operator instructs the participants to make their re-enactments “feel the same way” as the original speech, allowing for different word choice. The consent form uses the wording “re-create the mood and feeling of each utterance as closely as possible,” and the operator reinforces this idea many times, in different words. The participants are also encouraged to re-enact overlaps, disfluencies and pauses as much as possible.

Each selected fragment is played repeatedly for the participants, until they indicate that they are ready to re-create it. The participant or participants translate and reenact the utterance or utterance pair. Re-enactments occasionally require multiple attempts to attain enough naturalness and fidelity to satisfy the participants and the operator. When five attempts are not enough, that fragment is abandoned. There is something of a trade-off here, and we might explore whether to limit play back of each fragment to just a couple of repetitions before re-enactment, to reduce the pressure to, and opportunity to, mimic it too faithfully, in order to increase the chances of obtaining equally natural but more distinct target-language forms.

To avoid fatiguing the participants, sessions last no more than one hour. For this reason also, not all utterances are re-enacted. At the end, the participants are asked again whether they consent to share the recordings. Finally, each is paid; this was 15 USD and is now 20 USD.

Later, the operator uses his or her handwritten notes on the start times of the re-enacted fragments to use ELAN to mark up the spans and identify them with the corresponding spans in the original conversation. In a second pass, the operator marks individual phrases within both the original and the re-enacted utterances, typically pause-separated, and cross-identifies them. These phrases average 2.3 seconds in duration. Each re-enactment comprises one or more such

phrases. While re-enactments may include speech from both participants, each phrase is from only one participant. Finally, we run a script to pull out the matching phrases from the recordings, described in Section 4.6.

## 4.2 Participants

We recruit participants from among students at the University of Texas at El Paso, which is located on the US-Mexico border. Many students have family members on both sides of the border, many students have received part of their education on both sides, and many commute across it daily. Even on the US side, where the dominant language is English, many neighborhoods are Spanish-dominant.

We advertised for Spanish-English bilinguals. Our advertisement is in Appendix C. Their self-described dialects were, for English, overwhelmingly El Paso and, for Spanish, mostly “El Paso / Juarez,” with a few listing Chihuahua City or Guanajuato.

We encouraged participants to bring a partner to talk with, but if they came alone, the operator served as their conversation partner. She grew up in El Paso and is native and fluent in both languages.

## 4.3 Hardware and room configuration

The goal of the hardware setup is to record a conversation such that interlocutors are recorded on separate audio tracks, while still being able to hear each other. Our hardware setup (Figure 1) includes two identical single-sided headsets (Shure BRH441M) and a digital recorder with microphone inputs (TASCAM DR-40). Audio is recorded in WAV format, sampling at 44.1 kHz. Working with the cables we had, we added cable adapters where needed. These adapters should not have a significant effect on quality.

The data collection takes place in a research lab with a main room connected to two separate rooms with a glass window between them, as suggested by Figure 2.

## 4.4 Data collection instructions

### 4.4.1 Setting up for the initial recording

1. Take inventory of the hardware:
  - one desktop computer with stereo speakers (C1)
  - one recorder (C2)
  - two headsets (C3, C4)
  - two microphone cables (shielded, thick, three-prong) (C6, C7)
  - two audio cables (thin) (C8, C9)
  - one USB upload cable (C13)
  - three adapters (C5, C11, C12)

The desktop with stereo speakers (C1) is next to the bookshelf. On the bookshelf, the box labeled “TASCAM DR-40” contains (C2) and the “Under Armour” shoebox contains (C3) through (C13). Refer to Figures 1 and 2 when connecting components.

2. Take the components out of the boxes. Remove the Velcro and twist ties and keep them to store the components later. Some components are labeled “L” (left) and “R” (right).

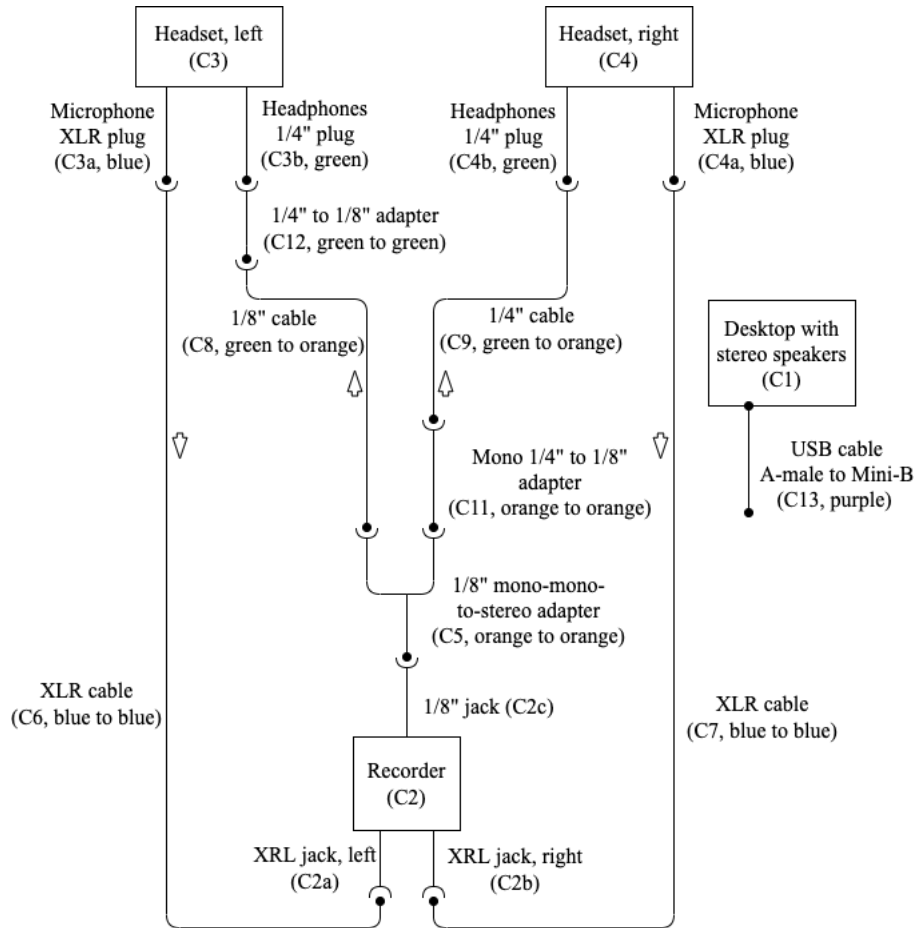


Figure 1: Hardware diagram. Component numbers and color labels are in parenthesis.

3. Place the recorder (C2) and cable (C13) in the main room, within reach. The USB upload cable (C13) will not be used until the next stage.
4. Place a chair in each of the two smaller rooms, near the glass window and facing each other.
5. Place the left headset (C3) on the chair in the left room.
6. Place the right headset (C4) on the chair in the right room.
7. Connect the left headset microphone to the recorder.
  - (a) Connect plug (C3a) to cable (C6).
  - (b) Connect cable (C6) to jack (C2a).
8. Connect the right headset microphone to the recorder.
  - (a) Connect plug (C4a) to cable (C7).
  - (b) Connect cable (C7) to jack (C2b).
9. Connect the left headset headphones to the mono-mono-to-stereo adapter.

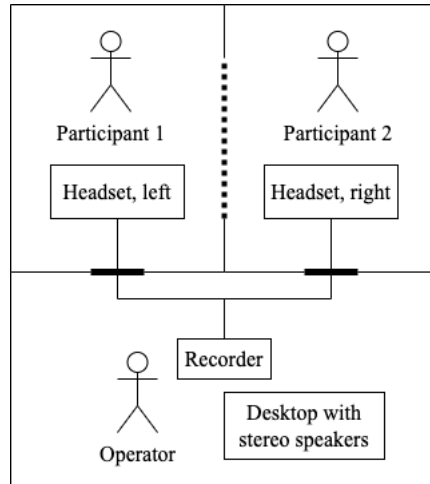


Figure 2: Configuration of rooms. The dotted line is the glass window and the thick lines are doors.

- (a) Connect plug (C3b) to adapter (C12).
  - (b) Connect adapter (C12) to cable (C8).
  - (c) Connect cable (C8) to adapter (C5).
10. Connect the right headset headphones to the mono-mono-to-stereo adapter.
  - (a) Connect plug (C4b) to cable (C9).
  - (b) Connect cable (C9) to adapter (C11).
  - (c) Connect adapter (C11) to adapter (C5).
11. Connect the mono-mono-to-stereo adapter (C5) to the recorder (C2c). Leave this step to last, because the weight of the cables can damage the recorder.
12. Lay all the cables flat and separated, so they do not catch when the doors are closed.
13. Turn on the recorder by pressing and holding the HOME button on the front of the recorder.
14. Sound-check the headsets. Press the RECORD button on the front of the recorder, then snap your fingers or speak close to each microphone in turn and check the recorder display. The bars on the display stretch to the right if the microphones are picking up sound. The top bar corresponds to the left microphone and the bottom bar corresponds to the right microphone.
15. Turn on the desktop to have it ready for the next stages. To unlock the desktop, press CTRL + ALT + Delete.
16. Sound-check the desktop speakers. Play a stereo audio, e.g. from a YouTube video, and check that each speaker is outputting sound.
17. From the green folder in the desk, take out a blank piece of paper to write notes in the next stage. Look up the most recent assigned subject numbers, and prepare to assign the sequential next numbers to the new participants.

#### 4.4.2 When participants arrive

1. Welcome participants into the main room and have them sit at the round table to go over consent forms and any questions. Let them know that, as they read on the flyer, there will be a lot of translating from one language to the other, and it is crucial to the study that they are fluent in both, for example, by saying:

Before we begin and sign anything, I'm instructed to inform you that there will be a lot of translating from one language to the other, so if you feel like if that is something you won't be able to do, then I wouldn't want to waste your time. Otherwise, we can proceed.

2. Casually speak in both languages to gauge how fluent they are in each. Ask them to practice translating with you.
  - If both participants are fluent, continue with the procedure.
  - If only one of the participants is fluent, add a note of this and continue with the procedure. We may still use the audio.
  - If both of the participants are not fluent, inform the participants that you have "gathered what we needed for the study" then give them the 20 USD and dismiss them.
3. Give each participant a consent form (Appendix D) and a language background form (Appendix B), found in the bookshelf.
4. When participants hand you their filled-out forms, enter the participants' assigned participant ID on their language background form: the same participant IDs you assigned before they arrived.
5. Go over things one more time, or as many times as they need in order to understand the purpose and procedure. For example, you might say:

The purpose of this data collection is to further speech-to-speech translation research by creating an open collection of translated conversations, something that has not been done before.

Today you will have a conversation with your partner in one language, then re-enact parts of it in another language. I will select some snippets of the audio and replay them for you to translate re-record in the other language. It is important that you try your best to make it sound natural while also keeping the same feeling as in the original. Try to recreate pauses, laughs, long breaths, or anything of that sort during the second recording if possible.

I can replay the audio as many times as you need, and give you as much time as you need to translate. If either of us feel like you can translate the words better or if the prosody was not as faithful in feeling as it could be, then we can redo it until we are satisfied. Please be vocal of any opinions that you have about the process and ask any questions that may arise.

6. Optionally suggest a topic or dialog activity (Section 7).
7. Ask the participants to move into their separate rooms and put on their headsets.



8. Press the RECORD button on the recorder. Make sure it is recording. Allow the participants to converse until they run out of things to say, or after 10 minutes of recording, whichever comes first.
9. While the participants are talking, on the blank note sheet, record the following information: the current date, the ID of the conversation (following the format from Table 4), the left-track (left room) participant ID, and the right track (right room) participant ID.
10. Let the participants know that the recording is over, and now they will begin translating and re-recording fragments of the conversation. Ask them to wait a few minutes while you set up.

#### 4.4.3 Setting up for the re-enactment recording

1. Disconnect the recorder (C2) from the other components (C5, C6, C7). To disconnect (C6, C7), press and hold the release tabs next to the jacks (C2a, C2b).
2. Turn on the desktop speakers by raising the volume with the volume knob.
3. Connect the recorder (C2) to the desktop (C1) using cable (C13).
4. Press the HOME button on the front of the recorder. From the options that appear on the recorder display, select STORAGE. If the menu does not appear, disconnect and reconnect cable (C13), then press the HOME button again.
5. When the recorder storage window appears on the desktop, open the MUSIC folder.
6. Copy the most recent WAV file to the `IsgAudioRecordings` folder on the desktop.
7. Rename the audio file with its conversation ID, following the format from Table 4.
8. Eject the recorder from the desktop, then disconnect cable (C13) from the recorder (C2) and desktop (C1).
9. Reconnect the recorder (C3) to the other components (C5, C6, C7), as they were before.
10. Open ELAN. In the blank window that appears, select “File” > “New”. Navigate to to the `IsgAudioRecordings` folder, select the WAV file you copied, and select “OK” to finish.
11. An ELAN window should open with the audio on the bottom half and an annotation tier named “default”.
12. Set up the annotation tiers for the next stage.
  - (a) Rename the “Default” to “Utterance”.
  - (b) Add a new tier named “LittleLeft”.
  - (c) Add a new tier named “LittleRight”.

#### 4.4.4 Recording re-enactments

1. Ask the participants if there are any parts of their conversation that they would like to have removed. Annotate these sections with “DELETE” (see Appendix E for keyboard shortcuts).
2. Emphasize that they should keep the same feeling in their re-enactments.
3. Turn on the recorder and press the RECORD button.
4. If at any point during recording the re-enactments you notice that one or both of the participants are not as fluent as expected, follow the instruction in Step 2 of Section 4.4.2.
5. Play the audio for the participants and you to hear, stopping at a fragment to re-enact, preferably one that is easy to understand and differs from the rest of the recording in terms of prosody or feeling. (What makes a “good” fragment to re-enact was discussed in Section 4.1.)
6. Mark this region the Utterance tier, starting with #1 and incrementing.
7. Replay this audio fragment at least two or three times, or more if the participants request.
8. When the participant(s) succeed in producing a good re-enactment, jot down the time on your notes sheet, to reference in the next stage.
9. Repeat this process for the remainder of the session, leaving some time to debrief and pay the participants.
10. Press the RECORD button to stop recording re-enactments.
11. Ask the participants if they have any questions. If everything is clear, ask them to sign the re-consent forms. After they sign the re-consent forms, thank them for their participation, give them each 20 USD, and dismiss them.

#### 4.4.5 When participants leave

If there will be another recording session soon, set up the hardware again like in Section 4.4.1. If there will not be other recording sessions soon, disconnect all the components and store them as they were found. Wrap cables carefully using the same Velcro and twist ties. Turn off the speakers completely (until the volume knob clicks).

#### 4.5 Entering metadata and adding annotations

1. On the desktop, open the metadata Excel workbook in the “IsgDatabaseTables” folder.
2. If this is your first time as producer, enter two new rows in the “producer” sheet. Table 6 describes each field in the “producer” sheet.
3. In the “participant” sheet, enter two new rows. Table 3 describes each field in the “participant” sheet.
4. In the “conversation” sheet, enter a new row. Table 4 describes each field in the “conversation” sheet.

Now you can start mapping fragments to their translations.

1. Create a new ELAN file for the re-enacted conversation and configure the three annotation tiers, as done in Section 4.4.3 Step 12.
2. Open the ELAN file for the original conversation, so that you have both the ELAN file for the original conversation and the re-enacted conversation open, side-by-side. The original will already have utterances labeled. Typically these will be integers starting with 1 and going up to 30 or so.
3. For each fragment in the ELAN file for the original conversation, annotate the re-enactment in the re-enactment file with the same number as the fragment it recreates. Use the times that you noted in Section 4.4.4 to find the re-enactments quickly. See Appendix E for keyboard shortcuts.
4. In addition, when possible or necessary, identify matching phrases within each original-re-enactment pair. This is done for two reasons: first, the prosody within a single phrase is expected to be simpler to model than the prosody of concatenated phrases, and second, we want these matched pairs to contain speech from only one participant. There are three cases. In the first case, there are re-enactments which already include speech from only one participant. If this is one inseparable unit, then we simply annotate the same range in the tier for that speaker, LittleLeft for the left speaker, and LittleRight for the right speaker. If however, there is a clear pause (or occasionally more than one), we mark the individual phrases as separate regions in the appropriate Little tier. The second case is for reenactments that include speaker changes. For these we mark the appropriate regions in the Little tiers. For example, if the left participant talks then the right participant, the first region will be marked in LittleLeft and the second in LittleRight. The third case are reenactments that involve overlap. Again, we mark all speech regions for the left participant in LittleLeft and for the right in LittleRight. In all three cases, we do this first the original audio and then in the same way for the re-enactment audio.
5. Save both files before closing the windows.

## 4.6 Post-processing procedure

The script for creating a DRAL release reads the conversation audio, markup, and metadata files, and then outputs: phrase clip pairs (also called “short fragment audios”), reenactment clip pairs (also called “long fragment audios”), short fragment audios concatenated per conversation and track, a copy of the complete original and re-enacted audio recordings, and updated metadata as CSV files. The script is written in Python and uses SoX for manipulating audio.

### 4.6.1 Set up

1. Create a root directory for the project somewhere, e.g. `~/Desktop/DRAL`.
2. Download the following files from the GitHub repository <https://github.com/joneavila/DRAL> into the project root directory.
  - `requirements.txt`
  - `make_DRAL_release.py`

3. Install SoX. See: <https://github.com/rabitt/pysox#install> for instructions for installing SoX on MacOS and Linux. If installing on Windows, add the directory containing the SoX executable to the PATH environment variable.
4. Install Python 3.10.4 or create a Python 3.10.4 virtual environment for the project.
5. Install the required Python modules and their dependencies with pip: `pip install -r requirements.txt`. The main Python modules are: *pandas* for data manipulation, *pympi* for processing ELAN files, and *pysox*, a Python SoX wrapper.
6. Copy the the directory `recordings/`, containing the conversation audio (`.wav`) and markup (`.eaf`) files, to the project root directory.
7. Copy the metadata worksheet `metadata.xlsx` to the project root directory.

#### 4.6.2 Creating a new release

The script `make_DRAL_release.py` assumes that the audio, markup, and metadata files are in the same parent directory, as described in Section 4.6.1. It writes the contents of the release to the output directory `release/`. To create a new release, simply run the script: `python3 make_DRAL_release.py`.

The script will print warnings for problem conversations and fragments. Problem conversations and fragments will be skipped and excluded from the release. Each warning will hint at a possible fix (Table 1).

#### 4.6.3 Quality control

We do not have a systematic quality control process. Periodically all three of us sit down for an hour to listen to the matching pairs from a few dialogs. We comment on interesting cases and boring cases, and note cases which lapsed somewhat in fidelity or, more rarely, naturalness. In the rare case that incorrect pairings are found, as may result from errors in the markup stage, we mark these for re-annotation and thus correction in the next release. More commonly, we discuss what went well or less well, and may discuss adjusting the balance of instructions to prioritize naturalness in the target language, for example, or ways to diversify the dialog acts present.

## 5 Data Released to Date

While our data collection is not complete, we are releasing what we have so far. Table 2 summarizes the status as of November 15, 2022.

We anticipate that the individual phrases will be the primary unit of analysis for most purposes, but we also release the recordings of the entire original conversations and the entire re-enactment sessions. We also release all mark-up and scripts, so that others can reprocess the data as convenient for their purposes. We also release the metadata, notably the ID numbers for the participants in each dialog, cross-referenced to their dialect information, etc. Tables 3, 4, and 5 describe the columns in the CSV files.

We make this corpus public as a resource for research and for training and evaluating speech-to-speech translation models and systems. Releases are available at [www.cs.utep.edu/nigel/dral/](http://www.cs.utep.edu/nigel/dral/). The releases web page will be updated with additions to the corpus as they become available. Each

additional download will contain additional conversation recordings and their phrases along with revised, up-to-date metadata.

## 6 Observations

The original conversations seemed very natural and the participants generally seemed to enjoy them; indeed sometimes it was hard to get them to stop. While not all participants were equally fluent in both languages, code switching was rare, mostly occurring when a participant had a lexical gap, for example not knowing how to say *3D printer* in Spanish. All succeeded in the task. The re-enactment process required much more effort, but our participants generally seemed to enjoy the challenges involved. They understood that the purpose of our data collection was better translation systems, and they seemed eager to support it, setting high standards for the fidelity and naturalness of their productions, encouraged in this by the operator and probably some peer pressure.

In terms of quality, many of the original-conversation utterances were disfluency-ridden, partly unintelligible, or strange from a grammatical, semantic, or pragmatic viewpoint. This is to be expected, as our participants were not professional speakers. However the utterances were generally successful in terms of interacting appropriately with the interlocutor and keeping the conversation going.

Participants appeared to vary in how they interpreted our instructions regarding re-enactment. Many appeared to strive for (or relax into) lexical fidelity, with almost every word in the original represented by a single word in the other, as in Appendix A.3 (but see also Appendix A.1 as a partial counterexample). Many also seemed to strive for prosodic fidelity, producing re-enactments where the overall prosodic contour was very similar to that of the original. We subsequently honed our instructions to stress our desire not for low-level fidelity but rather fidelity in terms of the feeling and function conveyed.

The topics and interaction styles varied. We noticed that some participants deliberately avoided possible contentious or embarrassing topics (Appendix A.3), and others seemed to naturally gravitate to “safe” topics and a “safe” speaking style, such as telling stories or asking questions and giving answers. These tendencies worked against our aim of obtaining diverse pragmatic functions, so we are experimenting with adding additional instructions to include rotating suggestions for topics and dialog activities that would promote diversity, as discussed in Section 7.

We wanted the re-enactments to be natural, and in this we largely succeeded. In informal experiments presenting original-reenactment pairs to bilinguals, we found they were mostly unable to identify which was which, indicating that most of the re-enactments were highly natural. One common exception was utterances involving laughter, which seems often very hard to re-create convincingly.

In terms of functional diversity, in the first dozen dialogs we found that we had at least 29 of the 42 dialog acts in Jurafsky *et al.*’s (1997) list. In addition, we observed many functions and forms not listed in such taxonomies, such as saliently leaving something unsaid (Appendix A.2), strongly asserting oneself (Appendix A.3), enacting reported speech, and expressing sympathy. In terms of Ward’s dimensions of interaction style [23], at various times in these dialogs there appeared to be interactions at each extreme of all 8 dimensions.

We observed that often much of the information in these utterances was conveyed largely by the prosody (Examples 2 and 3). In terms of prosodic diversity, we observed examples of at least the 14 most frequent prosodic constructions of English in Ward’s (2019) listing. As interesting prosody is largely lacking from existing speech-to-speech translation corpora, the presence of many samples with meaningful and varied prosody is a unique strength of this corpus. We also did a preliminary

scan over the data to judge the extent to which utterances with identical intent would differ in prosody between English and Spanish utterances. Somewhat surprisingly, given the literature on this topic [4, 24], most often the differences were minor, most saliently involving lexical stress and utterance-initial and utterance-final lengthening. This may reflect properties of our local dialects or perhaps the “playing it safe” topics that most participants favored.

## 7 Towards Increased Variety

Although our data collection protocol has given fairly good coverage of the dialog acts and activities common in casual conversation, these do not exhaust the interesting things that can happen in dialog, and there are many use cases involve dialog activities that we’re unlikely to observe by just allowing the participants to talk about whatever they like.

Accordingly, for the more recent recording sessions we have worked to increase the diversity in the data by occasionally making suggestions to the participants, notably those seen in Table 7. To comment briefly on the dialog activities that these tend to elicit: 1 involves making plans together, including making suggestions, filtering them, and reaching agreement. 2 is about information exchange. 3 is about establishing responsibility. 4 is about persuasion, requests, and negotiation. 5 is about self-disclosure, and probably being supportive. 6 is about discussing contentious topics and disagreeing or resolving disagreement. 7 is about broadening the topics through eliciting an uncommon emotion. It also turns out that ad hoc topics suggested by the operator to appeal specifically to the participants can work very well: for example “whether or not it’s a good idea to be in a relationship while in college.”

Other prompts that we have considered but not tried yet include: a) Work together to write out a plan for a brief tour of campus etc. for a first-time visitor from Durango (or Arizona) who is considering coming to UTEP for graduate school. b) Pick one of you to pretend to want to travel by bus from UTEP to somewhere in the city (his/her choice), and ask the other person to use their phone to find a way to get there, and write it down. c) Have one person describe how to get to his/her house from UTEP, while the other takes notes well enough to get there without using an app. Beyond such suggestions we could assign role-play, but that would be a little unnatural, or set up complex situations, such as games or puzzles, but that would be time-consuming.

The producer tends to make suggestions based on her knowledge of what dialog activities have been rare recently.

## 8 Cost Components

The major cost component is the time spent by the producer/operator/director/annotator. (The other two are the stipends for the participants, and various “overhead,” including post-processing, organization and quality control, and paperwork.)

Currently we are spending about 75 producer hours per hour of yield, where an hour of yield means 1 hours worth of Spanish utterances matched with about 1 hours worth of English utterances.

This is based on the facts that our producer labor time for a 1-hour recording session is almost 2 hours, and that over our first 28 sessions our yield was just over 45 minutes. The breakdown: Each session takes 5 minutes for the consent and demographic forms, 10 minutes for the initial conversation, and 45 minutes for the re-enactments. After that comes annotation of the data to identify the matched pairs for our program to extract them. This takes 30–45 minutes. Most of this is in the second phase, where the annotator has to carefully check that the phrases correspond across the languages and carefully delimit them. The time varies based on factors such as the

quality of the re-enactments: when this is lower overall, the annotator has to be more judgmental and careful.

As noted earlier, we don't have a separate quality control step; rather we rely on the producer to enforce quality as she goes. However in real production, with a diverse set of producers, quality control would almost certainly be necessary. This would probably require someone to mark each matched pair for quality on a scale from 0 to 3, or something like that, which might take 10 minutes, or more if the quality control person is not the director, and therefore has to listen to more context. We'd then also need to modify our workflow to copy over such ratings into the release.

The total time per session is thus 60 minutes for participant A, 60 minutes for participant B, and 120 minutes for the staff person. Some efficiencies may be possible, but we feel that most of this time is probably not hugely squeezable. In addition we have other staff time costs which are rather specific to our local procedure, including time for recruiting subjects, dealing with no-shows, setting up and breaking down the environment and equipment, and running the software to output the pairs. Perhaps some of these could be avoided or better designed with a different setup or environment. We note that these time-cost estimates exclude the start-up costs, which were huge, but should be much less for future protocol modifications.

The next two sections discuss explorations with protocol alternatives, motivated in part by the goal of reducing cost.

## 9 Other-Language Explorations

This section describes experimental data collections in Japanese and Bengali. While included in the release, these are experimental and intended to illustrate what we can get, rather than to be used for modeling or evaluation.

Our main reason for collecting these was to explore the possibility of scaling up to more languages without having to hire operators knowing each language. To put it another way, a possible way to reduce costs would be to reduce the qualifications required of the operator hired to handle the recording. Our current protocol assumes that the producer/director/operator/annotator is bilingual and thus fully able to monitor quality in both languages, but it is possible to increase the division of labor, such that the director/operator needs only technical skills and people skills, not language skills.

Our second reason for doing these was to find any other issues that might arise for language pairs beyond Spanish-English.

Accordingly we did two experimental data collections, one in English and Japanese, and one in English and Bengali. Our operator speaks neither Japanese nor Bengali. For Ja-En, one participant was the first author, hence highly familiar with the protocol. For Bn-En, one participant was a lab member, hence broadly familiar with our aims and methods. All four participants in these experiments were less truly bilingual than our average Spanish-English participant. We followed the protocol, with two exceptions: in both sessions we tried both directions, switching the source and target languages, and we didn't do the second annotation pass to get the phrase-level correspondences. For both experiments, after processing was complete, we had a meeting to review a sampling of the pairs collected, with each meeting including one of the participants of the respective experiment.

The next subsection presents observations relating to quality, the following subsection presents observations relating more to language issues, and the last subsection summarizes the findings regarding the possibility of recording with a language-unfamiliar operator.

## 9.1 Quality-Related Observations

Reduced fidelity: Fidelity was often weaker than in the En-Es re-enactments. This may have been in part due to language differences making fidelity harder. It may have also been due to reduced effortfulness. In the Ja-En, we observed a couple of cases where the participants failed to re-create all the nuances, for example, treating something as new information in the original utterance, but not in the re-enactment. The ability of the director to assess fidelity and request retakes was very limited, although in one case she was able to suggest this after noticing that the re-enactment was suspiciously shorter and more fluent than the original.

Regarding naturalness: There were cases where aspects of the prosody of Japanese leaked into the English re-enactment. There appeared to be more of these than in the Es-En. This may have been partly due to a language skill deficit.

Regarding throughput: 1) We were concerned that it would take more time to identify interesting and complete utterances to re-enact. However, the director was able to estimate when an utterance started and ended by using the sound wave display in Elan, and most candidate utterances that she selected on this basis were confirmed by the participants to be natural units to re-enact. 2) The rate of data collection was good overall, however there were two confounding factors: having a knowledgeable participant and having apparently lower quality aspirations.

Regarding annotation quality: Here the usual operator did the annotation, based on her notes during the re-enactment phase. This went well, except in one case where there was an error in her notes which, not knowing the language, she was unable to later correct for.

## 9.2 Language-Related Observations

As these were our first excursions beyond English and Spanish, we noted some other issues, which likely relate more to the language pairs than to the operator’s knowledge. These observations are impressionistic and need to be revisited through systematic analysis.

First, regarding both languages, we observed that word order differences between Bn-En and Ja-En seem more extreme than between Es-En. This would make it harder to identify sub-utterance phrases that match up across the two languages.

Regarding Japanese: 1) Japanese and English are so different in structure and that it was at times impossible to translate fully faithfully. 2) Japanese is much more interactive than English, so many behaviors, such as backchannelling, overlapped speech, and laughter, were often simply missing in the English re-enactments. 3) Japanese is much less explicit than English, so sometimes the English came out significantly longer.

Regarding Bengali: 1) The Bengali re-enactments were often significantly longer, which may reflect a language or cultural tendency to speak in full sentences. 2) Prosodically, we noticed that Bengali seemed to differ from English in the expressions of pausing for a moment to think of a word, of listing things, and of marking questions. 3) Grammatically, Bengali has different word order, with the verb frequently at the end, and with prepositional phrases sometimes fronted.

Other quality-related comments: 1) The Bengali data included several quite long utterances, which the participants sometimes seemed to struggle to re-enact. This may have been in part because the topic chosen was “my research.” Our confederate had suggested lighter topics such as movies, but the other participant had declined these. 2) There were times where the re-enactment included thinking of the Bengali word while speaking. This led to gratuitous mismatches, which the operator couldn’t catch. A possible mitigation to this specific problem could be to add an instruction to “completely formulate the utterance in your mind, before you start speaking,” perhaps augmented with “this can be hard, so feel free to practice once or twice until you’re ready.”



### 9.3 Conclusion and Possible Mitigations

Overall, having a language-unfamiliar operator clearly reduces the quality obtained. We haven’t quantified this effect, but, in comparison to Es-En, where probably 95–98% were fully acceptable, for Jp-En this was probably closer to 75%. The lower-quality pairs could still have some utility as part of a large collection, but are of course less valuable. For Bn-En the yield was probably even lower, with some of the low-quality pairs so poor as to be of probably no value to any currently conceivable machine learning scheme.

This quality deficit occurred despite a situation that was nearly ideal in some respects: most of the participants were motivated and highly cooperative, and the operator/director was experienced, motivated, engaged, and sensitive enough to prosody to notice likely lapses in fidelity or naturalness. Despite this, the quality was frequently poor. The major factor is probably whether or not the operator/director can judge quality, apply “peer pressure,” and request retakes. Without this, even for motivated participants, the quality will suffer.

If it is nevertheless necessary to record with a language-unfamiliar operator, there may be ways to partly mitigate the quality loss, including:

Mitigation via procedure modification: One might periodically nag the participants about the goal of fidelity and occasionally ask them to rate the quality of their reenactments; and similarly for the goal of naturalness.

Mitigation via participant selection: With a language-unfamiliar director, it becomes more important to have participants who are highly cooperative, sensitive to nuance, and diligent.

Mitigation via using experienced participants: Our confederate later participated in another Bn-En session, and reported that her experience in the first session helped her become a better guide. Thus, it could be good to ensure that at least one participant in each pair has experience with the protocol.

Mitigation via participant training: We may also wish to produce a more polished, formal description of what we want, perhaps in the form of a video. This could stress our goals and provide illustrations of what we consider good re-enactments and of common “error” types (too literal, slavishly similar prosody, too polished, too unpolished, etc.). We might even back this up with a follow-on quiz, and only consider people qualified to participate if they pass it. We could go further and require participants to be indoctrinated and trained in our aims and procedure, including perhaps having participated in quality review of a dozen utterance pairs, so that they truly understand what we want.

Mitigation in the annotation phase: While our operator was able to produce utterance-level annotations, to obtain phrase-level annotations would of course require hiring an annotator who knows both languages.

Mitigation via a quality-control phase: In our standard protocol the quality is ensured during the re-enactment phase. If the operator is language-unfamiliar, one mitigation would be to add a formal quality-control phase, to identify and probably discard low-quality utterances.

As all of these mitigations would increase the time cost, and likely recover only some of the lost quality, our bottom line conclusion is that using a language-familiar operator is likely still the best strategy overall.

## 10 An Exploration in Remote Collection

Currently our recordings are done in-lab, which requires participants to come to us. On campus this is not a major problem, but to scale up, it would be nice to develop a protocol for remote data

collection. This section discusses the issues expected and those observed, based in part on a small experimental data collection using Zoom, also described.

## 10.1 Concerns

A priori, based on our general knowledge of recording methods and of remote conversations, many things could go wrong in remote recording.

First, there are some basic issues of audio quality:

1. Participants' connectivity issues may hurt audio quality or introduce delay. (observed)
2. Participants' audio environments will vary, due to room reverberation, background noise levels, transient noise, etc. (observed)
3. The audio quality will be poorer, since we cannot provide high-quality microphones. (observed)
4. The recorded audio quality in the interaction may be poorer, due to the videochat software's compression and packet-loss concealment algorithms. (observed)
5. Depending on the videochat software, it may be impossible to record the participants and director in different channels. The issues this could cause include: preventing analysis of speech in overlapping region (which is arguably is not truly needed in the data we produce, for most use cases), and slowing down the post-processing. (not observed in the Zoom setup described below)
6. Transmission delay may make it impossible, depending on the videochat software, to easily synch up the different tracks. (intermittantly observed)

Second, there are considerations of impact on the conversations; remote conversations will be different from the face-to-face ones. Given that our experiment was mostly about getting the procedure to work we did not specifically confirm whether or not these were happening.

1. Altered dialog dynamics, due to transmission delay — both audio and video, with the former mattering more — that affect the conversation dynamics of the original conversation, and maybe also limiting what can be re-enacted. Previous research suggests that the most severe effect will be on the turn-taking, but there is a silver lining in that reduced overlap, laughter and backchanneling will make postprocessing and modeling easier. In this way, remote data collection may have the benefit of excluding data that would be hard to model, and which is not representative of any real use case.
2. Altered speech and interactive behaviors, due to limited video bandwidth, such as seeing the face only, without hand gestures etc.
3. Altered behavior due to lack of a shared environment and shared visual references.
4. Reduced ability for pre-dialog synching of mood, which is likely to make the dialogs more formal. This could be mitigated in various ways, but a more formal style may any better serve some use cases.
5. Reduced ability to reinforce expectations, which may reduce the extent to which the participants take the re-enactment task seriously, which may reduce both the naturalness and fidelity of the pairs collected.

Third, there may be implications for participant behavior that affect process efficiency:

1. Channel lag and limited bandwidth interaction may reduce the ability of the director to deftly and swiftly guide re-enactments, making it a clunky, frustrating, or slow process. (observed)
2. Remote effects may reduce the enjoyability of the process. (observed)
3. Remote collection may increase the cognitive load on the director. (observed)
4. Remote collection may increase the time required for post processing. (not observed)

So, there are numerous quality-cost trade-offs and time-cost trade-offs, that need to be considered.

Incidentally, while we here we have been assumed that the director and the participants will all be remote from each other. There is also the possibility of allowing, or encouraging, the participants to be co-located. This would hurt audio separation, and decrease the likelihood of getting participants who are strangers to each other, which is important for the main use case. But it could be easier to recruit subjects and perhaps more enjoyable for them.

The rest of this section describes the process that we developed, in part to offer one remote workflow, and in part to describe the context in which we made the above observations.

## 10.2 Remote Collection Preparation using Zoom

The subsection describes the procedure we used in our pilot experiment. There are probably better tools for doing remote collections.

### 10.2.1 Install and Configure Zoom

1. Install the Zoom desktop client, version 2.0 or higher.
2. Open the Zoom desktop client and sign in. Create an account if you do not already have one.
3. Click the user account icon on the top right, then click “Settings”.
4. In the “Settings” window, click the “General” tab.
5. Enable “Show my meeting duration”.
6. In the “Settings” window, click the “Recording” tab.
7. Note the location meeting recordings will be stored to, next to “Store my recordings at:”. This folder will open automatically after ending a meeting.
8. Enable “Record a separate audio file for each participant”.
9. Close the “Settings” window.
10. (macOS) Open System Settings, then click the “Privacy & Security” tab. Under “Screen Recording”, enable screen recording for the Zoom application.

### 10.2.2 Install and Configure ELAN

1. Install ELAN, version 6.4 or higher.
2. The default settings may result in strange-looking waveform visualizations. This may be specific to the audio coding format (MPEG-4) or may be specific to macOS. To fix the visualizations, click “Preferences...”, then click the “Platform/OS” tab, and under “Extraction for audio samples from video for audio visualizations (waveform, spectrogram)” select “Use AV Foundation” (Figure 3). Compare the waveform visualization before (Figure 4) and after (Figure 5) changing this setting.




Figure 3: ELAN “Platform/OS” settings after changing the “Extraction of audio...” setting to “Use AV Foundation”.

Also create a folder to store all conversation audio and markup files, e.g. “DRAL”.

### 10.2.3 Prepare to record the original conversation

1. Collect the participants’ email addresses.
2. Like in the in-person procedure, enter a new row for each participant in the “participant” metadata sheet.

### 10.2.4 Record the original conversation

1. Start a Zoom meeting as the host.
2. Invite the participants to the meeting. Click “Participants”, then click “Invite”, then click the “Email” tab, then send an email invitation to the participants.
3. Once the participants join the meeting, click “Record”.
4. After the participants end their conversation, let them know they will take a break and join a second meeting shortly.

5. End the meeting. Zoom will begin converting the recording.
6. Once the recording has been processed, the folder containing the recording files will open.
7. Within the recording folder, open the “Audio Record” folder. In the “Audio Record” folder, each participant’s recorded audio will be listed as its own audio file (.m4a), with a filename containing the name they used in the meeting (Figure 6).

### 10.2.5 Prepare to record the re-enactments

1. In the “DRAL” folder, create a folder for the conversation with the conversation ID, e.g. “EN\_633”.
2. Copy both participants’ audio files to the conversation folder.
3. Rename each participant’s audio file with their participant ID (Figure 7).
4. Delete the Zoom meeting recording folder with the remaining files: the audio of the operator, the combined meeting audio, the recording configuration file, and the meeting video.
5. Open ELAN.
6. Click “File”, then click “New...”.
7. Import the participants’ audio files. In the “New” window, click “Add Media File...”, then browse to the conversation folder, then select the audio of the first participant, then click “Open”. Repeat for the second audio.
8. After selecting both audio files (Figure 9), click “OK”.
9. By default, ELAN will set the volume of the first audio to 100% and the volume of all other audios to 0%. Click the “Controls” tab, then raise the volume of the second audio to 100%.
10. Configure the annotation tiers, like in the in-person procedure.
11. After adjusting the volume and configuring the tiers, the ELAN window should look like Figure 8.
12. Click “File”, then click “Save” and save the markup file (.eaf) in the conversation folder, with the same name as the conversation ID.

### 10.2.6 Record the re-enactments

1. Start a Zoom meeting as the host, invite the participants to the meeting, and start recording the meeting, as in Section 10.2.4. Note the time into the meeting when the recording starts.
2. Share the ELAN window with the original conversation audios so that you and the participants can hear its audio. Click “Share Screen”, then select “ELAN – ELAN [version number] – [ELAN filename]”, then click “Share”.
  - Note: The ELAN waveform display may appear glitchy after sharing your screen. We were not able to resolve this issue in our tests.

3. Like in the in-person procedure, select the fragments to re-enact and note down the times of the re-enactments, using the time into the meeting displayed by Zoom.
  - Note: The time into the meeting displayed by Zoom is not the same as the time into the recording. Use the time into the meeting when the recording started you noted in Step 1 above to adjust accordingly.
  - Note: When using multiple audio files, ELAN displays just one waveform at a time. To change the waveform display to a different audio, select the audio from the drop-down list to the left of the waveform display.
  - Note: The two audios may sound out of sync at some times. We were not able to resolve this issue in our tests.
  - Note: The audio recorded and processed by Zoom is exactly what was heard during the meeting. Any problems due to lag or interruptions during this phase should be fixed by re-re-enacting.
4. End the meeting and wait for the folder containing the recording files to open, as in Section 10.2.4.

### 10.2.7 Markup the conversation fragment pairs

1. In the “DRAL” folder, create a folder for the re-enactment conversation with the conversation ID, e.g. “ES\_633”.
2. Copy both participant’s audio files to the conversation folder and rename each participant’s audio file with their participant ID, as in Section 10.2.4.
3. Create a new ELAN file, import the two audios, adjust the volume and configure the annotation tiers, and save the markup file (.eaf) with the conversation ID, as in Section 10.2.4.
4. Like in the in-person procedure, use your notes to annotate the conversation fragment pairs.
5. Like in the in-person procedure, enter a new row for the original conversation in the “conversation” metadata sheet and enter another row for the re-enactment conversation.

### 10.2.8 Prepare the data for the person who runs the post-processing script

1. After completing the data collection, the file structure of the “DRAL” folder should look like Figure 10.
2. Archive the “DRAL” folder with `tar` or another archiving utility, e.g.
 

```
tar -vcf 'DRAL-2.0.tgz' [path to DRAL folder]
```

**Acknowledgments:** We thank Ann Lee and Justine Kao for comments and discussion.

## References

- [1] Rosana Ardila et al. “Common Voice: A massively-multilingual speech corpus”. In: *LREC (12th)*. 2019, pp. 4218–4222.

- [2] Massa Baali, Wassim El-Hajj, and Ahmed Ali. “Creating Speech-to-Speech Corpus from Dubbed Series”. In: *arXiv preprint arXiv:2203.03601* (2022).
- [3] Marcelly Zanon Boito et al. “MaSS: A large and clean multilingual corpus of sentence-aligned spoken utterances extracted from the Bible”. In: *LREC*. 2020.
- [4] J. Donald Bowen. “A Comparison of the Intonation Patterns of English and Spanish”. In: *Hispania* 39.1 (Mar. 1956), p. 30. ISSN: 00182133. DOI: 10.2307/335189. JSTOR: 335189.
- [5] Roldano Cattoni et al. “MuST-C: A multilingual corpus for end-to-end speech translation”. In: *Computer Speech & Language* 66 (2021), p. 101155.
- [6] Mark Dingemanse and Simeon Floyd. “Conversation across cultures”. In: *The Cambridge handbook of linguistic anthropology*. Cambridge University Press, 2014, pp. 447–480.
- [7] Mark Dingemanse and Andreas Liesenfeld. “From text to talk: Harnessing conversational corpora for humane and diversity-aware language technology”. In: *ACL*. 2022, to appear.
- [8] Kosuke Doi, Katsuhito Sudoh, and Satoshi Nakamura. “Large-Scale English-Japanese Simultaneous Interpretation Corpus: Construction and Analyses with Sentence-Aligned Data”. In: *Proceedings, 18th International Conference on Spoken Language Translation (IWSLT)*. 2021, pp. 226–235.
- [9] Simeon Floyd. “Conversation and Culture”. In: *Annual Review of Anthropology*, 50. Annual Reviews, 2021, pp. 219–240.
- [10] Ye Jia et al. “CVSS Corpus and Massively Multilingual Speech-to-Speech Translation”. In: *LREC*. arXiv preprint arXiv:2201.03713. 2022.
- [11] Ye Jia et al. “Direct speech-to-speech translation with a sequence-to-sequence model”. In: *arXiv preprint arXiv:1904.06037* (2019).
- [12] Dan Jurafsky, Liz Shriberg, and Debra Biasca. *Switchboard SWBD-DAMSL shallow-discourse-function annotation coders manual*. Tech. rep. Institute of Cognitive Science Technical Report, 1997.
- [13] Takatomo Kano, Sakriani Sakti, and Satoshi Nakamura. “Transformer-based direct speech-to-speech translation with transcoder”. In: *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE. 2021, pp. 958–965.
- [14] Ann Lee et al. “Direct speech-to-speech translation with discrete units”. In: *arXiv preprint arXiv:2107.05604* (2021).
- [15] Gabriel Luthier and Andrei Popescu-Belis. “Chat or Learn: a Data-Driven Robust Question-Answering System”. English. In: *Proceedings of the Twelfth Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, May 2020, pp. 5474–5480. ISBN: 979-10-95546-34-4. URL: <https://aclanthology.org/2020.lrec-1.672>.
- [16] Alp Öktem, Mireia Farrús, and Antonio Bonafonte. “Corpora compilation for prosody-informed speech processing”. In: *Language Resources and Evaluation* 55 (2021), pp. 925–946.
- [17] Vineel Pratap et al. “MLS: A Large-Scale Multilingual Dataset for Speech Research”. In: *Proc. Interspeech 2020*. 2020, pp. 2757–2761. DOI: 10.21437/Interspeech.2020-2826. URL: <https://arxiv.org/abs/2012.03411v2>.

- [18] Elizabeth Salesky et al. “The Multilingual TEDx Corpus for Speech Recognition and Translation”. In: *Interspeech 2021*. Interspeech 2021. ISCA, 2021, pp. 3655–3659. DOI: 10.21437/Interspeech.2021-11. URL: [https://www.isca-speech.org/archive/interspeech\\_2021/salesky21\\_interspeech.html](https://www.isca-speech.org/archive/interspeech_2021/salesky21_interspeech.html).
- [19] Petra Wagner and Andreas Windmann. “Acted and spontaneous conversational prosody—Same or Different”. In: *Proc. 8th International Conference of Speech Prosody, Boston, USA*. 2016, pp. 1–5.
- [20] Changhan Wang, Anne Wu, and Juan Pino. “CoVoST 2 and Massively Multilingual Speech-to-Text Translation”. In: *LREC*. 2020, pp. 4197–4203. DOI: 10.48550/ARXIV.2007.10310. URL: <https://arxiv.org/abs/2007.10310>.
- [21] Changhan Wang et al. “VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation”. In: *ACL*. 2021, pp. 993–1003.
- [22] Nigel G. Ward. *Prosodic Patterns in English Conversation*. Cambridge University Press, 2019.
- [23] Nigel G. Ward. “Using Interaction Style Dimensions to Characterize Spoken Dialog Corpora”. In: *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*. 2022, pp. 225–230.
- [24] Nigel G. Ward and Paola Gallardo. “Non-native differences in prosodic-construction use”. In: *Dialogue & Discourse* 8.1 (2017), pp. 1–30.
- [25] Chen Zhang et al. “UWspeech: Speech to speech translation for unwritten languages”. In: *AAAI 2021*. Vol. 59. 2020, p. 132.



Table 1: Post-processing script warnings and causes.

Warning	Cause
These conversations are missing markup...	The conversation does not have a markup file with the same name as its ID.
These conversations are missing audio...	The conversation does not have an audio file with the same name as its ID.
These conversations have an unexpected ID...	The conversation does not have an ID following the pattern: a valid language code, followed by an underscore, followed by three digits.
These conversations have an unexpected original or re-enacted code...	The conversation does not have an original or re-enacted code following the pattern: a value of “OG” or “RE”.
These conversations do not have a translation with a valid ID and original and re-enacted code or have multiple translations...	The conversation does not have exactly one translation with an ID that matches the pattern: a valid language code other than the language code in the ID of the current conversation, followed by an underscore, followed by the same three digits in the ID of the current conversation.
These markups have unexpected values...	The markup has a value that does not match the pattern: an optional pound symbol, followed by any number of digits. The markup value might contain a typo or contain a comment to be excluded from the release.
These markups are in an unexpected tier...	The markup is in a tier with a name that does not follow the pattern: a value of “LittleLeft” or “LittleRight” or “Utterance”.
These markups have duplicate values...	The markup shares a value with one or more markups in the same tier (the markup value is not unique).
These markups have zero or more than one translation...	The fragment does not have exactly one translation with a markup value equal to the current fragment and in a tier with the same name as the current fragment.

Table 2: Corpus statistics.

Conversations	32
Participants	60
Matched re-enactment pairs	815
Mean re-enactment duration	3.8 s
Matched phrase pairs	1173
Mean phrase duration	2.3 s

Table 3: Participant table fields. All fields are entered manually.

Field	Description	Value
<i>id</i>	a unique identifier for a participant	sequence of integers
<i>lang1</i>	the first language the participant will speak	ISO 639–1 two-letter code
<i>lang2</i>	the second language the participant will speak	ISO 639–1 two-letter code
<i>lang_strength</i>	the participant’s strength in language 1 and language 2, self-reported in a questionnaire	an integer value 1 through 5, representing a spectrum: 1 “language 1 stronger”, 2 “language 1 slightly stronger”, 3 “equal”, 4 “language 2 slightly stronger”, and 5 “language 2 stronger”
<i>dialect_note1</i>	the note on the participant’s language 1 dialect, self-reported in a questionnaire	string
<i>dialect_note2</i>	the note on the participant’s language 2 dialect, self-reported in a questionnaire	string
<i>is_producer</i>	indicates whether the participant is also the producer	boolean value, asterisk for True, blank for False
<i>notes</i>	additional notes, not used in post-processing	string

Table 4: Conversation table fields. The field *trans\_id* field is entered automatically by a post-processing script.

Field	Description	Value
<i>id</i>	a unique identifier for the conversation, matches the filename of the conversation audio	a string following the format: <language code> - <three digits>
<i>date</i>	the date of the recording session	string dd_mm_yyyy
<i>original_or_reenacted</i>	the type of conversation	“OG” for original or “RE” for reenacted
<i>participant_id_left</i>	the unique identifier of the participant recorded on the left track	same as <i>id</i> field from Table 3
<i>participant_id_right</i>	the unique identifier of the participant recorded on the right track	same as <i>id</i> field from Table 3
<i>producer_id</i>	the unique identifier for the producer (operator)	same as <i>id</i> field from Table 6
<i>trans_id</i>	the conversation ID of the translation conversation	same as <i>id</i> field from Table 4

Table 5: Short fragments table and long fragments table fields. All fields are entered automatically by a post-processing script.

Field	Description	Value
<i>id</i>	a unique identifier for the fragment	a string in the format <language code> _ <conv_id> _ <markup value>
<i>time_start</i>	the time into the source conversation the utterance starts	duration with format: mm:ss.ms
<i>time_end</i>	the time into the source conversation the utterance ends	duration with format: mm:ss.ms
<i>duration</i>	the duration of the fragment	duration with format: mm:ss.ms
<i>conv_id</i>	the unique identifier of the conversation the utterance	same as <i>id</i> field from Table 4
<i>trans_id</i>	the unique identifier of the matching, translation fragment	same as <i>id</i> field from Table 5

Table 6: Producer table fields. All fields are entered manually.

Field	Description	Value
<i>id</i>	a unique identifier for the producer	integer
<i>name</i>	the name of the producer	string

1. Agree on 3 bits of advice to give all new students joining the CS department (or coming to UTEP in general).
2. Agree on a strategy for appealing the grading decisions of a hopelessly unfair TA.
3. Think of a bad outcome that happened to you or someone you know, and together discuss who might share part of the blame, and who deserves most of the blame.
4. Ask the other person to show you everything in their purse or backpack. Be persistent.
5. Talk about something that's on your mind, for example an intimidating homework assignment.
6. Find something that the two of you disagree about.
7. Talk about something sad.

Table 7: Suggestions for participants

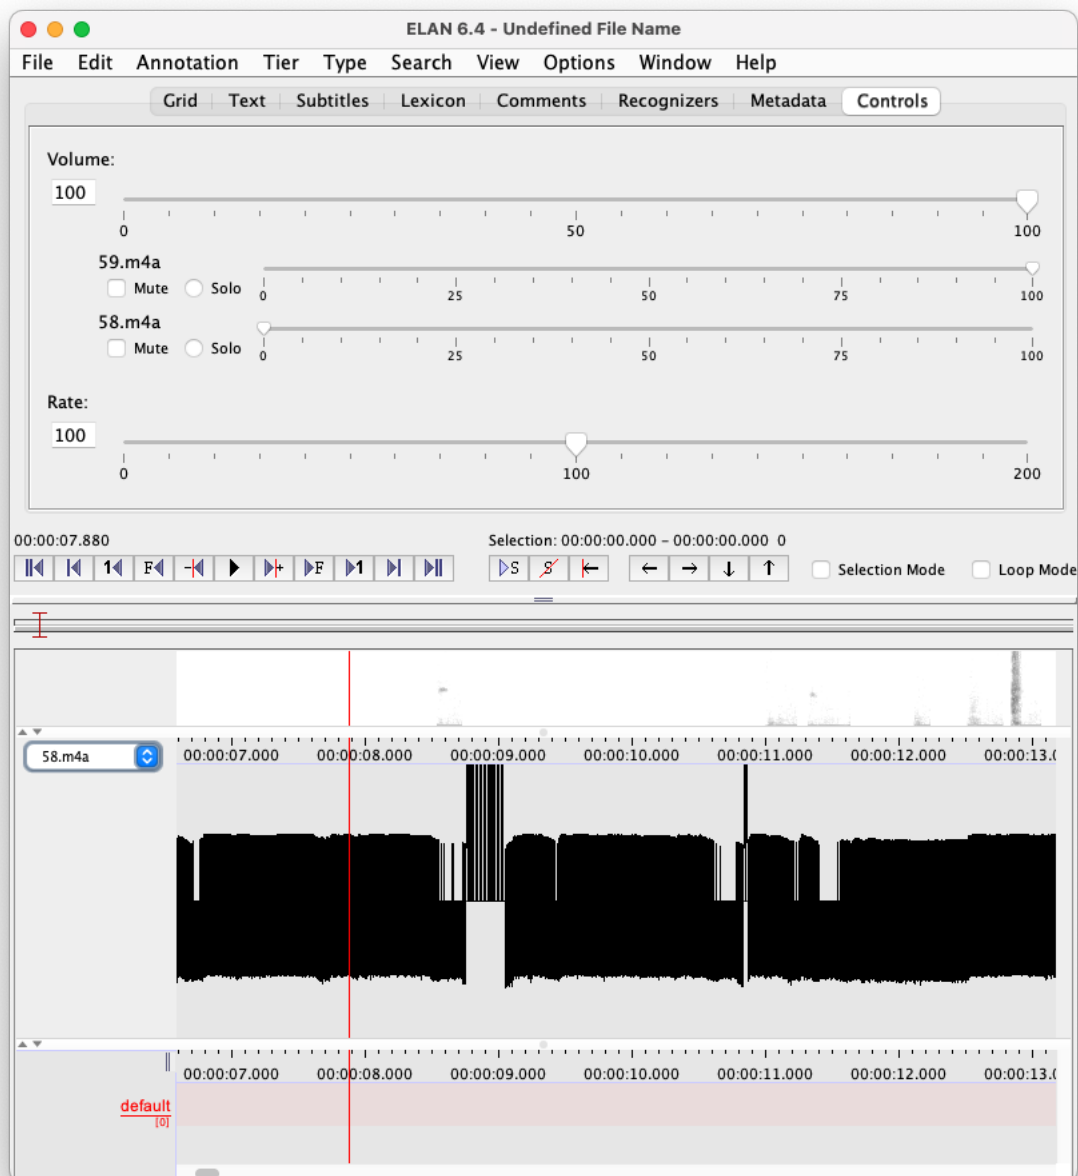


Figure 4: ELAN waveform visualization using the default “Use FFmpeg” setting.

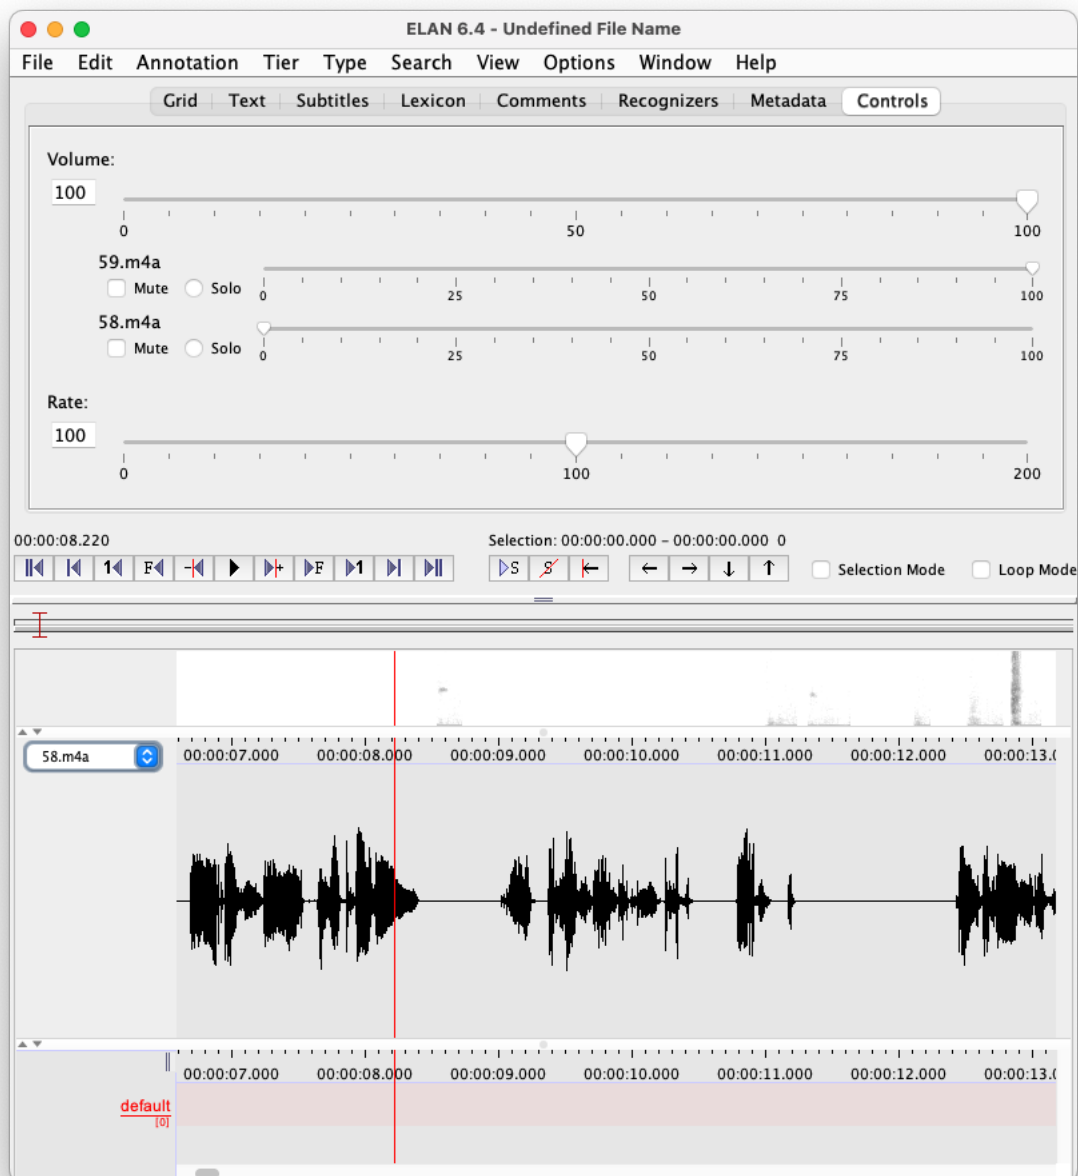


Figure 5: ELAN waveform visualization using the “Use AV Foundation” setting.

< > 2022-10-19 16.13.01 Jonathan Avila's Zoom Meeting			
Name	^	Size	Kind
Audio Record	--		Folder
audioCarlosOrtega21794516692.m4a		5.7 MB	Apple MPEG-4 audio
audioJonathanAvila31794516692.m4a		5.2 MB	Apple MPEG-4 audio
audioNigelWard11794516692.m4a		5.2 MB	Apple MPEG-4 audio
audio1794516692.m4a		5.2 MB	Apple MPEG-4 audio
recording.conf		126 bytes	Document
video1794516692.mp4		35.9 MB	MPEG-4 movie

Figure 6: Zoom recording files contents.

< > DRAL			
Name	^	Size	Kind
EN_633	↑ 10.4 MB		Folder
58.m4a	5.2 MB		Apple MPEG-4 audio
59.m4a	5.2 MB		Apple MPEG-4 audio

Figure 7: Conversation folder contents before opening ELAN.

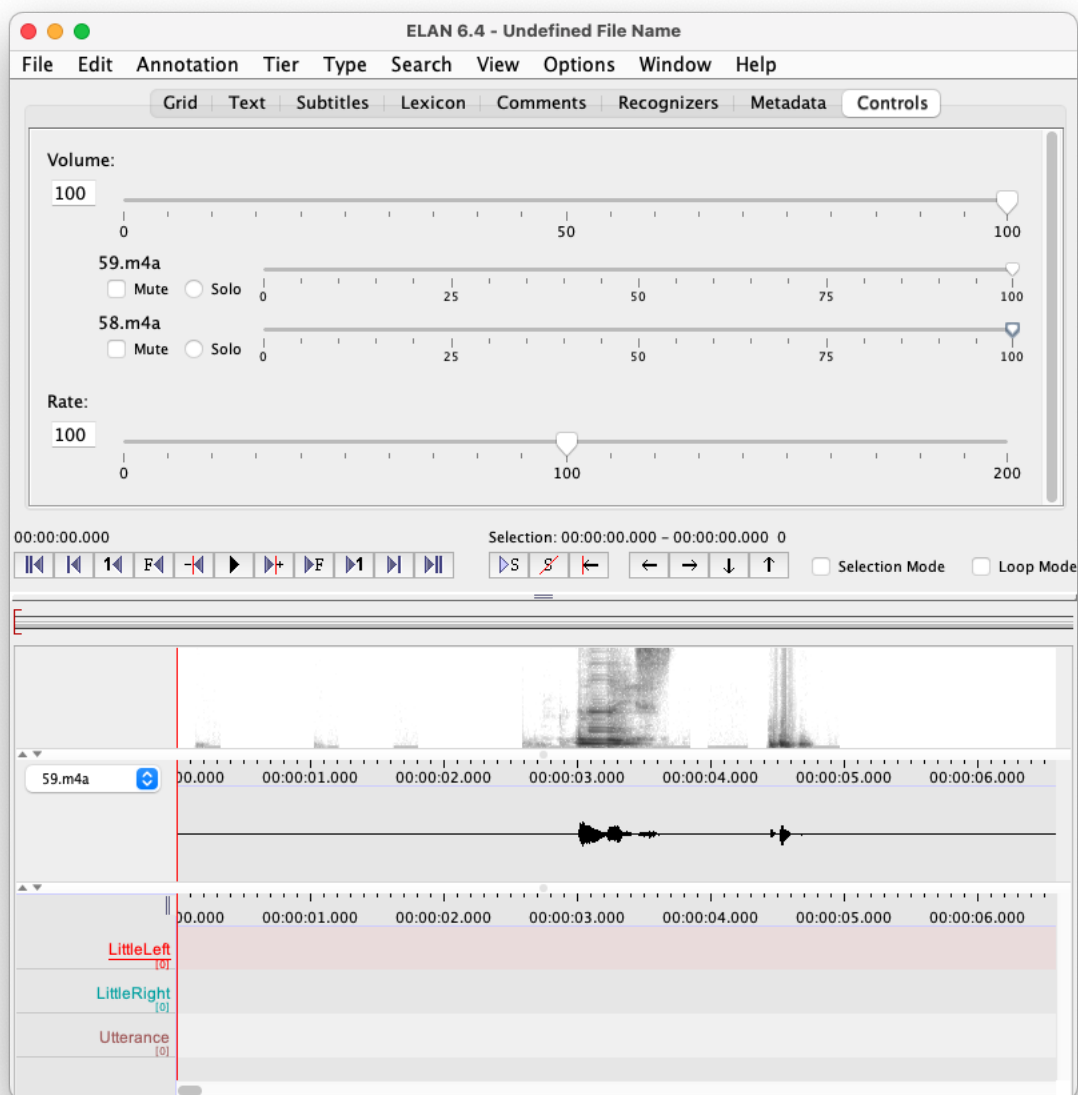


Figure 8: ELAN main window after importing audios, adjusting volume, and configuring annotation tiers.

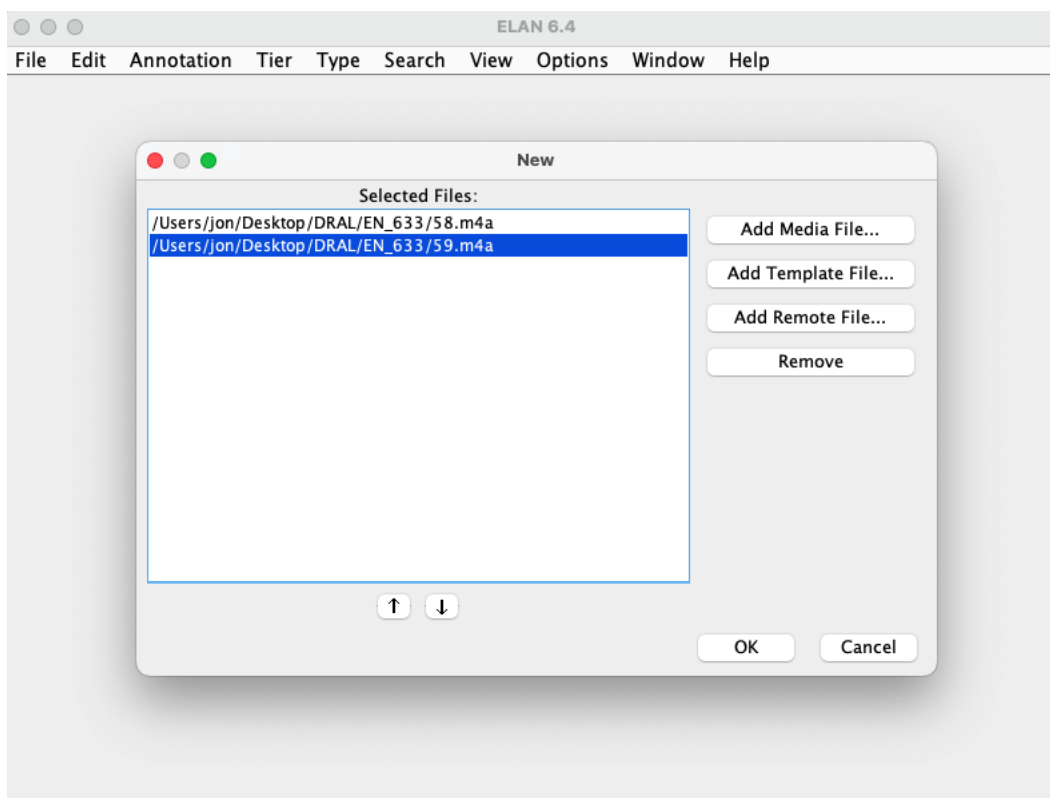


Figure 9: ELAN “New” window after importing participant audios.

< > DRAL		
Name	^	Kind
✓ EN_734		Folder
61.wav		Waveform audio
62.wav		Waveform audio
EN_734.eaf		EAF Document
✓ ES_734		Folder
61.wav		Waveform audio
62.wav		Waveform audio
ES_734.eaf		EAF Document

Figure 10: Data collection files, before archiving, ready for shipping, with one example original conversation and re-enactment conversation.



## A Examples from the corpus

Below are examples from the corpus. Inferable elided continuations are in brackets. Matched phrases are in bold.

### A.1 Example 1 (ES\_008\_18, EN\_008\_18)

Original: *¿Van a venir, venir tus papás para ...*

Gloss: will-they to come, come your parents to  
[help you move in]?

Re-enactment: *Are you parents gonna come, or  
[not]?*

Prosody Note: both original and re-enactment end with a saliently lengthened syllable, seemly inviting the listener to complete the phrase

In context:

X: *Vas a tener tu propio,*

Y: *Ai, si cierto.*

X: *departamento.*

Y: *Ya el jueves.*

X: *¿El jueves?*

Y: *El jueves me lo van a dar, el jueves a las tres de la tarde.*

X: *¿**Van a venir, venir tus papás para?***

Translation of context:

X: *You're going to have your own,*

Y: *Ah, that's right.*

X: *apartment.*

Y: *Already on Thursday.*

X: *On Thursday?*

Y: *On Thursday they're going to give it to me, on Thursday at three in the afternoon.*

X: ***Are you parents gonna come, or?***

### A.2 Example 2 (EN\_006\_29, ES\_006\_29)

Context: X spent some time working with inmates at the county jail

Original, in context:

X: *So I worked over there, it was pretty... **interesting.***

Y: (laughter)

X: *It was, it was a cool experience.*

Y: (overlapping) *Good choice of words.*

Prosody Note: The word *interesting* came after a 2.0 s pause.

Inferred Pragmatic Intent: indicate that there were things that he doesn't want to talk about

Re-enactment:

X: ***interesante.***

### A.3 Example 3 (ES\_002\_3.34, EN\_002\_4.34)

Original: *Lo estoy interpretando como yo quiero y como yo pienso.*

Gloss: It I-am construing as I wish and  
as I am-thinking.

Inferred Pragmatic Intent: assert his intention to tell his story without being censored

Prosody Note: both original and re-enactment have extra punch on the stressed syllables, giving an impression of annoyance and assertiveness

Reenactment: *I am interpreting things how I want and how I think.*

In context:

X: *Todo que es incómodo a trabajar en la tierra, a tumbar becerros,*

Y: (overlapping) *A sí.*

X: *cortarle los ...*

Y: (gasp)

X: *Eso viene siendo cosas que a muchas personas los incomoda, y yo ya estoy acostumbrado hacerlo.*

Y: *Sabes que me acabo de acordar que este audio lo vamos a, a distribuir, osea, entero, así que van a oír todo lo que dijimos desde el principio.*

X: *Pero pues yo estoy diciendo cosas. **Lo estoy interpretando como yo quiero y como yo pienso.***

Y: *A ok, bueno pues, perdón.*

Translation of context:

X: *Everything that is uncomfortable when working on the land, throwing down calves,*

Y: (overlapping) *Oh yes.*

X: *cutting off their ...*

Y: (gasp)

X: *Those are things that make many people uncomfortable, and I'm used to doing it.*

Y: *You know, I just remembered that this audio we're going to, to distribute, like, whole, so they're going to hear everything we said from the beginning.*

X: *Well, I'm saying things. **I am interpreting things how I want and how I think.***

Y: *Ah okay, well then, sorry.*

## B Language/dialect background form

## Language/Dialect Background Form

Subject number \_\_\_\_\_

**Language 1** \_\_\_\_\_

**Language 2** \_\_\_\_\_

### Language 1 Accent/Dialect

What variety of Language 1 do you speak? \_\_\_\_\_

### Language 2 Accent/Dialect

What variety of Language 2 do you speak? \_\_\_\_\_

### Overall Conversational Skills/Confidence

①

Language 1  
stronger

②

Language 1  
slightly stronger

③

Equal

④

Language 2  
slightly stronger

⑤

Language 2  
stronger

## C Advertisement

Parallel English-Spanish Conversational Speech Corpus Collection Project

## Help our research towards improving **SPEECH TRANSLATION TECHNOLOGY**

We are collecting the world's first closely matched collection of equivalent English and Spanish conversational utterances!

We need **Spanish-English Bilingual adults(18+)**  
to come make conversation recordings for our  
project!

### **What will you do?**

1. Come to our lab in the Chemistry-CS building
2. Have a short conversation with someone
3. Re-enact selected fragments of that conversation  
in the other language
4. Give permission for your recordings to be  
incorporated into our dataset for public release

Come with a conversation partner, or come alone and  
we'll provide a partner.

The audio recording session will last just one hour!

Before you leave, you'll receive \$15 USD in cash  
To arrange a time contact me at [erivas6@miners.utep.edu](mailto:erivas6@miners.utep.edu) or  
scan the QR code below!



UTEP Computer Science Department  
Interactive Systems Group  
Undergraduate Research Assistant Emilia Rivas  
PI Dr. Nigel Ward



## D Informed consent form

University of Texas at El Paso (UTEP) Institutional Review Board  
**Informed Consent Form for Research Involving Human Subjects**

---

**Protocol Title: Collection of a Parallel English-Spanish Conversational Speech Corpus**

**Principal Investigator: Dr. Nigel Ward**

**UTEP: Computer Science**

---

**Introduction**

---

You are being asked to take part voluntarily in the research project described below. You are encouraged to take your time in making your decision. It is important that you read the information that describes the study. Please ask the study researcher or the study staff to explain any words or information that you do not clearly understand.

**Why is this study being done?**

---

People who do not speak the same language may experience difficulty communicating with each other. This language barrier may be alleviated with a translating system or device. However, these mostly translate only the words, not reliably the intent or feeling. To support development of better systems UTEP is collecting the first-ever open-source speech corpus of parallel English and Spanish utterances from natural and reenacted conversations, thus we will be recording audio. We hope to support research that will eventually facilitate more accurate translations for real people in real-world situations.

We plan to enroll approximately 120 participants for this study at UTEP.

You are being asked to be in the study because you self-reported to be fluent in both English and Spanish, and being comfortable having a conversation and it being recorded. Age and gender are not exclusion factors, provided that you are age 18 or older .

If you decide to enroll in this study, your involvement will last one hour or less, in a single visit. However you can apply to participate again at a later time



### **What is involved in the study?**

If you agree to take part in this study:

You will fill out the Language/Dialect Background Form. This will take approximately thirty seconds.

In the case that you bring a conversation partner, the procedures described below will remain the same and both of you will be compensated \$15. In the case that you do not bring a partner and another participant signed up for the same time slot, you two will be partnered and follow the procedures below. In the case that you do not bring a partner and another participant does not sign up for the same time slot, the research assistant will take the place of the second participant and will also follow the procedures while still taking the part of the research assistant of handling the equipment and facilitating the process.

Before starting the recording process, there will be COVID precautions in place including the sanitizing of surfaces and equipment. After being briefed by the research assistant, and filling out a language background questionnaire, you will head into one of two rooms separated by a window from your conversation partner. You will then be asked to put on a headset with a microphone through which you will be able to hear each other.

You will then have a conversation in either Spanish or English about any topic you feel comfortable with. This will be recorded. After at least 5 but no more than 10 minutes you will end the first conversation. The research assistant will then ask you and the other participant if there are any parts of the conversation that should be omitted, such as anything you said that contained possibly sensitive or embarrassing information. Then you will wait a minute or two while the assistant uploads the files and processes the needed deletions.

Subsequently, the research assistant will select multiple snippets of audio, each about 5 seconds long or less, for you to mentally translate and re-record in the other language. Each snippet will be an utterance or utterance pair. The participant or participants who produced that utterance or pair will then mentally translate what they said, and then reenact it in the other language. The research assistant will play it aloud, as many times as you want, and when you are ready you can produce your re-enactment. We would like you to re-create the mood and feeling of each utterance as closely as possible. It is not important to use exactly appropriate words, or to produce perfectly grammatical Spanish or English. In fact, overlaps, pauses and other disfluencies should be re-enacted, when possible. This may be difficult, so we may ask you to do re-takes, and to accept suggestions from your conversation partner and/or the research assistant. This may require multiple attempts. Perfection will not be expected. If an utterance turns out to be too difficult to re-enact, the experimenter will move on to the next. The





session will end after the collection of 100 re-enacted snippets or after the hour is over. You will then be able to ask any questions of the research assistant. You will then have a final chance to withhold from inclusion any or all parts of either recording. You will then receive \$15 cash.

#### **What are the risks and discomforts of the study?**

---

For the recording, you will be instructed to have a natural conversation about any topics you choose. In the case that a you may say something that you or we feel might put them "at risk of criminal or civil liability or be damaging to your financial standing, employability, insurability, or reputation," we will remove that section of audio from the recording. You will be reminded of this before and after the recording session by the research assistant.

The research team will also take extensive COVID precautions before, during, and after the recording process. Before you enter the main room, we will sanitize surfaces and equipment that have been touched by other participants. Before you and your conversation partner put on the headsets, we will offer each participant a pair of latex gloves, hand sanitizer, a mask (if you are not already wearing one), and highly encourage you to follow current CDC guidelines. We will also provide sanitizing products like wipes for you to use at their discretion. We will follow the social distancing guidelines recommended by the CDC. The cash will be sprayed with Lysol or a similar disinfectant. After the recording process we will sanitize each surface and piece of equipment again, leaving it ready for the next set of participants. There will be no physical contact between you and the research assistant whatsoever.

#### **What will happen if I am injured in this study?**

---

The University of Texas at El Paso and its affiliates do not offer to pay for or cover the cost of medical treatment for research related illness or injury. No funds have been set aside to pay or reimburse you in the event of such injury or illness. You will not give up any of your legal rights by signing this consent form. You should report any such injury to Dr. Nigel Ward (915)747-6827 and to the UTEP Institutional Review Board (IRB) at (915-747-6590) or [irb.orsp@utep.edu](mailto:irb.orsp@utep.edu).

#### **Are there benefits to taking part in this study?**

---

You are not likely to benefit by taking part in this study other than being able to meet new people, have a conversation that you wouldn't have had otherwise, and enjoy the challenge of re-enacting a conversation across languages. This data collection will help us in research on



ways to reduce minimize language barriers in many situations, which you may ultimately benefit from, if this leads to better translation systems and products.

In particular, we will use this data to build models of the typical mappings between Spanish and English prosody. This is a research project, so we will try many things, and success is not guaranteed, but our aim is the open-source release of a general model that can then be used by other research and development teams to improve the quality of speech-to-speech translation. Your data will be used, in combination with the data from many other speakers, to learn the parameters of this model. It will thus be a general model of the relationship between Spanish and English prosody, and in no way a specific model of any single individual's behavior patterns..

---

**Who is paying for this study?**

Funding for this study is provided by UTEP Department of Computer Science internal funds.

---

**What are my costs?**

There are no direct costs.

---

**Will I be paid to participate in this study?**

You will be compensated for your participation in the form of \$15 in USD cash at the end of the recording process.

---

**What other options are there?**

You have the option not to take part in this study. There will be no penalties involved if you choose not to take part in this study.

Choosing to withdraw or not participate will not affect your grades, class, or university standing.



#### **What if I want to withdraw, or am asked to withdraw from this study?**

---

Taking part in this study is voluntary. You have the right to choose not to take part in this study. If you do not take part in the study, there will be no penalty or loss of benefit.

If you choose to take part, you have the right to skip re-enactment of any utterance or to stop at any time. However, if you do, we encourage you to talk to a member of the research group so that they know why you are leaving the study. If there are any new findings during the study that may affect whether you want to continue to take part, you will be told about them.

The researcher may decide to stop your participation without your permission, if he or she thinks that being in the study may cause you harm.

#### **Who do I call if I have questions or problems?**

---

You may ask any questions you have now. If you have questions later, you may call Dr. Nigel Ward at (915)747-6827 or via email at [nigel@utep.edu](mailto:nigel@utep.edu).

If you have questions or concerns about your participation as a research subject, please contact the UTEP Institutional Review Board (IRB) at (915-747-6590) or [irb.orsp@utep.edu](mailto:irb.orsp@utep.edu).

#### **What about confidentiality?**

---

Your participation in this study is confidential. The following procedures will be followed to keep their personal information confidential. 1) To protect your privacy, the research team will separate consent forms from the background forms and the recordings data and thus be unable to trace either back to you. The research team will delete any audio that you deem sensitive or otherwise wish not to be included. 2) To maintain the confidentiality of the pre-release data, the research team and all adjacent parties will handle the audio recordings and findings from this study with upmost care by not allowing access to the full recordings to anyone outside of the team as well as storing all information in locked rooms and password-protected laptops/ desktops.



The snippets of your recordings will be made publicly available, as part of our overall data release, projected to include 2000-10000 matched utterance pairs, through Github. However, your name will not be associated with this data. In addition, sample utterances from the recordings and the results of the research study may be presented at meetings or in publications; however, again, your name will not be disclosed in those presentations. Moreover, the entire recordings, excluding portions that you or we decide to delete, will be retained and possibly used for research in our lab or shared with other research teams without restriction. While your name will never be associated with the data, it is possible that someone who knows you may hear part of a recording and recognize your voice. For this reason we will proactively delete any utterances that we feel could embarrass you or anyone else if heard by anyone. We will also delete anything else you request. The rest of the recording will be kept indefinitely.

Every effort will be made to keep your information confidential. Your personal information may be disclosed if required by law.

Organizations that may inspect and/or copy your research records for quality assurance and data analysis include, but are not necessarily limited to:

- Office of Human Research Protections
- UTEP Institutional Review Board

Because of the potential need to release information to these parties, absolute confidentiality cannot be guaranteed.

All records will be kept in locked rooms and password-protected laptops/ desktops before and after release.

#### **Authorization Statement**

I have read each page of this paper about the study (or it was read to me). I will be given a copy of the form to keep. I know I can stop being in this study without penalty. I know that being in this study is voluntary.

UTEP IRB V.3 (01/2018)



Approved on: March 24, 2022  
Expires on: March 23, 2023  
Study Number: [1871718-1]

☐ Yes, I agree to participate in this research project. I have read the above informed consent form and I understand what the research entails and understand that I will be audio-recorded.

☐ No, I do not agree to participate in this research project.

\_\_\_\_\_  
Participant's Name (printed)

\_\_\_\_\_  
Participant's Signature

\_\_\_\_\_  
Date

\_\_\_\_\_  
Signature of Person Obtaining Consent (Witness)

\_\_\_\_\_  
Date

#### **Re-authorization Statement**

After participating in the recording and re-enactment sessions, and indicating all needed deletions, I agree that the remaining recordings may be released and used for research purposes.

\_\_\_\_\_  
Participant's Signature

\_\_\_\_\_  
Date

\_\_\_\_\_  
Signature of Person Obtaining Consent (Witness)

\_\_\_\_\_  
Date



## **E ELAN keyboard shortcuts**

- To highlight a span of the audio, left-click and drag the cursor across the span.
- CTRL + Space: Play / Pause
- Shift + Space: Play / Pause highlighted area
- CTRL + Alt + N: A new annotation is created for the selected and highlighted span.
- CTRL + Alt + < or >: Selects the previous or next annotated span.