



# CORAA ASR: a large corpus of spontaneous and prepared speech manually validated for speech recognition in Brazilian Portuguese

Arnaldo Candido Junior<sup>4</sup> · Edresson Casanova<sup>2</sup> ·  
Anderson Soares<sup>3</sup> · Frederico Santos de Oliveira<sup>3</sup> · Lucas Oliveira<sup>1</sup> ·  
Ricardo Corso Fernandes Junior<sup>1</sup> · Daniel Peixoto Pinto da Silva<sup>1</sup> ·  
Fernando Gorgulho Fayet<sup>2</sup> · Bruno Baldissera Carlotto<sup>2</sup> ·  
Lucas Rafael Stefanel Gris<sup>1</sup> · Sandra Maria Aluísio<sup>2</sup>

Accepted: 20 September 2022 / Published online: 21 November 2022  
© The Author(s) 2022

## Abstract

Automatic Speech recognition (ASR) is a complex and challenging task. In recent years, there have been significant advances in the area. In particular, for the Brazilian Portuguese (BP) language, there were around 376 h publicly available for the ASR task until the second half of 2020. With the release of new datasets in early 2021, this number increased to 574 h. The existing resources, however, are composed of audios containing only read and prepared speech. There is a lack of datasets including spontaneous speech, which are essential in several ASR applications. This paper presents CORAA (Corpus of Annotated Audios) ASR with 290 h, a publicly available dataset for ASR in BP containing validated pairs of audio-transcription. CORAA ASR also contains European Portuguese audios (4.6 h). We also present a public ASR model based on Wav2Vec 2.0 XLSR-53, fine-tuned over CORAA ASR. Our model achieved a Word Error Rate (WER) of 24.18% on CORAA ASR test set and 20.08% on Common Voice test set. When measuring the Character Error Rate (CER), we obtained 11.02% and 6.34% for CORAA ASR and Common Voice, respectively. CORAA ASR corpora were assembled to both improve ASR models in BP with phenomena from spontaneous speech and motivate young researchers to start their studies on ASR for Portuguese. All the corpora are publicly available at <https://github.com/nilc-nlp/CORAA> under the CC BY-NC-ND 4.0 license.

**Keywords** Automatic speech recognition · Spontaneous speech · Prepared speech · Brazilian Portuguese · Public datasets · Public speech corpora

---

✉ Arnaldo Candido Junior  
arnaldo.candido@unesp.br

Extended author information available on the last page of the article

# 1 Introduction

Automatic Speech Recognition (ASR) is complex and challenging. Significant progress in techniques and models for the task had occurred in recent years. The main reasons for this progress include (but are not limited to) the availability of large-scale datasets and advances in deep learning methods running over powerful computing platforms.

Despite significant advances in ASR benchmarking solutions, the main and large datasets available for training and evaluating ASR systems are English due to the predominance of the language in science and business, although there are some current efforts to build multilingual speech corpora (Ardila et al., 2020; Pratap et al., 2020; Wang et al., 2020a, 2020b). Another problem is the environment of the recording, mostly composed of clean speech. Regarding the style of speaking, they are read speech, such as (Ardila et al., 2020; Panayotov et al., 2015; Pratap et al., 2020; Wang et al., 2020a; Zanon Boito et al., 2020) or prepared speech like (Hernandez et al., 2018; Salesky et al., 2021).

In this paper, we focus on a specific language—the Brazilian Portuguese (BP)—, which was struggling with only a few dozen hours of public data available until the middle of 2020. The previous open dataset to train speech models in BP were much smaller than American English datasets, with only 10 h for speech synthesis (TTS)<sup>1</sup> and 60 h for ASR. The resource commonly used to train ASR models for BP is an ensemble of four small, non-conversational speech datasets: the Common Voice Corpus version 5.1 (Mozilla)<sup>2</sup>, Sid dataset<sup>3</sup>, VoxForge<sup>4</sup>, and LapsBM1.4<sup>5</sup>.

In the second half of 2020, three new datasets were made available: (i) the BRSD v2 which includes the CETUC dataset (Alencar & Alcaim, 2008) (with almost 145 h), plus 12 h and 30 min of non-conversational speech from 3 small open datasets<sup>6</sup> (Macedo Quintanilha et al., 2020), (ii) the Multilingual LibriSpeech (MLS), derived from reading LibriVox audiobooks in 8 languages, including BP (Pratap et al., 2020) with 169 h, and (iii) the dataset Common Voice version 6.1<sup>7</sup> (Ardila et al., 2020), with 50 validated hours, composed of recordings of read sentences which were displayed on the screen. These three datasets total 376 h. Given this recent public availability of large audio databases for BP language, the lack of resources has been gradually reduced, although it is still far from ideal when compared to resources for the English language.

In early 2021, a new dataset with prepared speech, called the Multilingual TEDx Corpus (Salesky et al., 2021), was made publicly available, providing 765 h to support speech recognition and speech translation research. The Multilingual TEDx

<sup>1</sup> <https://github.com/Edresson/TTS-Portuguese-Corpus>.

<sup>2</sup> <https://commonvoice.mozilla.org/pt/datasets>.

<sup>3</sup> <https://doi.org/10.17771/PUCRio.acad.8372>.

<sup>4</sup> <http://www.voxforge.org/pt/downloads>.

<sup>5</sup> <https://laps.ufpa.br/falabrasil/>.

<sup>6</sup> BRSD v2 also includes CSLU: Spoltech Brazilian Portuguese Version 1.0 — <https://catalog ldc.upenn.edu/LDC2006S16>.

<sup>7</sup> <https://commonvoice.mozilla.org/pt/datasets>.

Corpus is composed of a collection of audio recordings from TEDx talks in 8 source languages, including 164 h of Portuguese. Moreover, a new version of the dataset Common Voice (Common Voice Corpus 7.0) was launched with 84 validated hours, which is an increment of 34 h over the previous version. Therefore, currently, BP language is well represented with 574 h of speech data which can be used to train new ASR models.

Another interesting resource is the Spoken Wikipedia Corpus (Baumann et al., 2019). The official release describes audios for English, German and Dutch. However, some Portuguese language audios without text/audio alignments are available<sup>8</sup>.

However, there is still a lack of datasets with audio files that record spontaneous speech of various genres, from interviews to informal dialogues and conversations, i.e., conversational speech recorded in natural contexts and noisy environments to train robust ASR systems. Spontaneous speech presents several phenomena such as laughter, coughs, filled pauses, word fragments resulting from repetitions, restarts and revisions of the discourse. This gap makes difficult the development of both high-quality dialog systems and automatic speech recognition systems capable of handling spontaneous speech recorded in noisy environments. The latter ones are called rich transcription-style ASR (RT-ASR) when they explicitly convert those phenomena cited above into special tokens (Fujimura et al., 2018; Inaguma et al., 2017; Tanaka et al., 2021). Dialog systems, for example, must deal with several types of speech disfluencies, preserving them instead of removing filled pauses and word fragments (Baumann et al., 2016). In general, it is expected that ASR systems trained on read style and clean speech (or even on prepared speech used in lectures and stage talks) will face a drop of performance when dealing with informal conversations in contexts of free interactions and noisy environments.

The TaRSila project is an effort of the Center for Artificial Intelligence<sup>9</sup> (C4AI) to make available language resources to bring natural language processing of BP to the state-of-the-art. The project aims at growing speech datasets for BP language, to achieve state-of-the-art results for automatic speech recognition, multi-speaker synthesis, speaker identification, and voice cloning. In a joint effort of two research centers, the C4AI and the CEIA<sup>10</sup> (Center of Excellence in Artificial Intelligence), four speech corpora composed of prepared, guided interviews and spontaneous speech from academic projects were manually validated to serve as an ASR benchmark for BP. The projects are: (i) ALIP (Gonçalves, 2019); (ii) C-ORAL-BRASIL I (Raso & Mello, 2012); (iii) Nurc-Recife (Oliviera Jr., 2016); and (iv) SP2010 (Mendes & Oushiro, 2012). We also validated 76 h of prepared speech from a collection of TEDx Talks<sup>11</sup> in Brazilian Portuguese, including 4.69 h of European Portuguese, to allow experiments with Portuguese language variants.

<sup>8</sup> [https://pt.wikipedia.org/wiki/Wikip%C3%A9dia:Projetos/Wikip%C3%A9dia\\_aud%C3%ADvel](https://pt.wikipedia.org/wiki/Wikip%C3%A9dia:Projetos/Wikip%C3%A9dia_aud%C3%ADvel).

<sup>9</sup> <http://c4ai.inova.usp.br/pt/nlp2-pt/>.

<sup>10</sup> <http://centrodeia.org/>.

<sup>11</sup> <https://www.ted.com/>.

## 1.1 Goals

In this paper we present a new publicly available dataset called CORAA ASR version 1.1. CORAA ASR has 290 h of validated pairs of audio-transcription and is composed of five corpora: ALIP (Gonçalves, 2019), C-ORAL-BRASIL I (Raso & Mello, 2012), NURC-Recife (Oliviera Jr., 2016), SP2010 (Mendes & Oushiro, 2012), TEDx Portuguese talks. Information about each corpus is presented in Table 1. The original sizes of each dataset in hours are presented as reported in their respective original papers, when reported by the authors. Regarding SP2010, the total duration is estimated, since the authors report 60 recordings from 60 to 70 min each and the total hours of ALIP was computed after download.

All the corpora are publicly available at <https://github.com/nilc-nlp/CORAA> under the CC BY-NC-ND 4.0 license. Two of the academic projects (C-ORAL-BRASIL I and ALIP) have explicit academic licenses and we assured the TED Media Requests Team that the TEDx dataset in Brazilian Portuguese would be released under a CC BY-NC-ND 4.0 license. Therefore, while it would be great to release all CORAA ASR subcorpora under a less restrictive license, we decided to standardize all the licenses as CC BY-NC-ND 4.0, as SP2010 and NURC-Recife were also funded by Brazilian government agencies in the same way as C-ORAL-BRASIL I and ALIP.

These corpora were assembled with the purpose of improving ASR models in BP with phenomena from spontaneous speech and noise in order to motivate young researchers in this exciting research area.

As an example of the feasibility of speech recognition with CORAA ASR, we present a speech recognition experiment using the Wav2vec 2.0 XLSR-53 (Baevski et al., 2020; Conneau et al., 2020). Furthermore, we compared our model with the state of the art in automatic speech recognition in Brazilian Portuguese (Gris et al., 2021). These two models are evaluated according to three main scenarios: (a) testing audios with different characteristics from training; (b) focusing on model performance for each of the five corpora, considering noise level and accent; (c) analyzing spontaneous and prepared speech styles impacts on the trained models.

## 1.2 Highlights

The main contributions made in this work are summarised as follows.

1. A large BP corpus of validated audio-transcription pairs containing 290 h, composed of five corpora (ALIP, C-ORAL-BRASIL I, NURC Recife, SP2010, and TEDx Portuguese talks), adapted for the task of ASR in BP. We also included 4.69 h of European Portuguese (in TEDx Portuguese corpus).
2. The first corpus, according to our knowledge, tackling spontaneous speech for ASR in BP.
3. An ASR model, publicly available, based on the presented corpus.

**Table 1** Speech genres, accents, speaking styles and hours (in decimal) in each original CORAA ASR corpus

Corpus	ALIP	C-ORAL BRASIL I	NURC Recife	SP2010	TEDx Portuguese
Speech genres	Interviews, dialogues	Monologues, dialogues, conversations	Dialogues, interviews, conference and class talks	Conversations, interviews, reading	Stage talks
Speaking styles	Spontaneous speech	Spontaneous speech	Spont. and prepared speech	Spont. and read speech	Prepared speech
Accent	São Paulo state cities	Minas Gerais	Recife	São Paulo capital	Misc.
Original (hrs)	78	21.13	279	65	249

Section 2 details both related work on datasets available for ASR in BP and the five spoken corpora projects used in CORAA ASR. Section 3 describes the steps followed in preparing the CORAA ASR corpus. Section 4 presents the statistics of the five sub-corpora that make up CORAA ASR version 1.1, after the revision process described in Sect. 3.5.2. Section 5 presents the final numbers of train, development and test splits of CORAA ASR version 1.1 and the experiment on ASR for BP. Finally, Sect. 6 presents the final remarks of the work.

## 2 Related work on speech datasets and spoken corpora for BP

### 2.1 Open datasets for speech recognition in BP

Three new datasets were released for BP at the end of 2020. CETUC dataset (Alencar & Alcaim, 2008) contains 145 h of 100 speakers, half males, and half females. The sentence set is composed of 1,000 sentences (3,528 words). The sentences are phonetically balanced and extracted from CETEN-Folha<sup>12</sup> corpus. Each speaker uttered all sentences from the sentence set exactly once. CETUC was recorded in a controlled environment, using a sample rate of 16kHz. The audios are publicly available,<sup>13</sup> without an explicit license. Regarding the environment of recording and speaking style, CETUC delivers clean and read speech.

Common Voice Corpus 6.1, version pt\_63h\_2020-12-11, contains 63 h of audio, 50 of which were considered validated. The dataset comprises 1,120 BP speakers, 81% males and 3% females (some audios are not sex labeled). The audios were collected using the Common Voice website<sup>14</sup> or using a mobile app. The speakers read aloud sentences presented on the screen. A maximum of 3 contributors analyzed each audio-transcription pair, and simple voting is applied: two votes for acceptance validate the audio; two votes for rejection invalidate the audio. A given release may also contain samples that were analyzed but did not receive enough votes to be validated/invalidated — these samples have the status “OTHER” (Ardila et al., 2020). Releases are distributed under the CC-0<sup>15</sup> license and contain MP3 files, originally collected at 48kHz sampling rate but downsampled to 16kHz. The following meta-data is also available: ID\_speaker, path\_audio\_file, read\_sentence, up\_votes, down\_votes, age, sex, and accent. Where up\_votes and down\_votes refer to the voting result, and the last three fields are optional. Regarding the speaking style, Common Voice Corpus has read speech. As for recording environment, both noise level and sound clarity is very heterogeneous. The current version of the dataset (Common Voice Corpus 7.0) has 84 validated hours, 34 h more than version 6.1.

<sup>12</sup> [https://www.linguateca.pt/cetenfolha/index\\_info.html](https://www.linguateca.pt/cetenfolha/index_info.html).

<sup>13</sup> <https://figormq.github.io/datasets/>.

<sup>14</sup> <https://commonvoice.mozilla.org/pt/speak>.

<sup>15</sup> <https://commonvoice.mozilla.org/pt/datasets>.

The Multilingual LibriSpeech (MLS) dataset (Pratap et al., 2020) is composed of audios extracted from Librivox<sup>16</sup> audiobooks. The Librivox project releases audiobooks in the public domain. MLS dataset encompasses eight languages, including BP, and is released under the CC BY 4.0<sup>17</sup> license. MLS can be used for developing both ASR and TTS models. There are 160.96 h for training models, 3.64 h for tuning and 3.74 for testing for Portuguese. It provides 26 male and 16 female speakers in the training dataset; 5 female, and 5 male speakers for tuning; and the same for testing. The audios were downsampled from 48kHz to 16kHz for easy processing. Regarding the environment of the recording and speaking style, MLS is made of clean and read speech.

In early 2021, a new dataset was made publicly available — the Multilingual TEDx Corpus, licensed under the CC BY-NC-ND 4.0.<sup>18</sup> This dataset has recordings of TEDx talks in 8 languages, BP being one of them, represented with 164 h and 93K sentences. Each TEDx talk is stored as a 44 or 48kHz sampled wav file. Available metadata include source language, talk title, speaker name, audio length, keywords, and a short talk description. Multilingual TEDx Corpus was built to advance ASR and speech translation research, with multilingual models and baseline models being distributed for ASR and speech translation. Regarding the speaking style and the environment of the recording, Multilingual TEDx Corpus is composed of prepared and clean speech.

## 2.2 Spoken corpora projects used in CORAA ASR

### 2.2.1 ALIP

The project ALIP<sup>19</sup> (Amostra Linguística do Interior Paulista – Language Sample of the Interior of São Paulo, in English) (Gonçalves, 2019) was proposed in 2002 and was responsible for building the database called Iboruna (Gonçalves, 2007), composed of two types of speech samples:

- A sample of 151 interviews (each with about 20 minutes, being 76 male and 76 female voices) from the northwest region of the São Paulo state;
- Another sample consisting of 11 dialogues, involving from two to five informants. It was recorded in contexts of free social interactions. This sample has 28 informants (10 men and 18 women).

This corpus totals 78 h and it is characterized by the spontaneous speech of the linguistic variety of Brazilian Portuguese spoken in the interior of São Paulo. It was compiled between the years of 2004 and 2005. The informants, residents of 7

<sup>16</sup> <https://librivox.org/>.

<sup>17</sup> <http://www.openslr.org/94/>.

<sup>18</sup> <http://www.openslr.org/100>

<sup>19</sup> <https://www.alip.ibilce.unesp.br/>.

different cities, range in age from 7 to over 55 years, with a considerable variety of income and education.

The speech samples were recorded with GamaPower and PowerPack digital recorders. For interviews, the consent of the informants was obtained before recording, while, for the dialogues, the consent was obtained after recording. The interviewer conducted the interviews, and the dialogues were free, with topics defined by the participant interactions.

The corpus is available for academic use without a defined license, but with defined Terms of Use and Privacy Policy.<sup>20</sup> It is available via download from the project website. The two types of samples have a dedicated folder for each, in the following formats. Each folder contains .mp3 files (the audios are sampled in 8kHz), as well as .doc and .pdf files (transcriptions, informant's socio-demographic information, among others). It is important to note that audio files are not aligned with their transcriptions.

### 2.2.2 C-ORAL-BRASIL I

C-ORAL-BRASIL I is a corpus published in 2012, resulting from the project C-ORAL-BRASIL<sup>21</sup> (Raso & Mello, 2012; Raso et al., 2012, 2015). This synchronic corpus was recorded between 2008 and 2011 and is composed of informal and spontaneous speech, representative of the linguistic variation in Minas Gerais, especially in the city of Belo Horizonte.

It is composed of 139 texts, totaling 21.13 h and 208, 130 words, averaging 1,500 words per text. C-ORAL-BRASIL I has 362 informants. There is a balance regarding number of uttered words: 50.36% words are uttered by 159 males and 49.64% words by 203 females.

Its content is divided into private-family (about 3/4 of the corpus) and public (1/4) contexts. In addition, there is a separation of interaction types by number of participants: monologues (amounting to about 1/3 of recordings), dialogues and conversations, i.e. more than two active participants (about 2/3 of recordings).

The speech flow was segmented into tonal units and terminal units according to the prosodic criterion, based on the Language Into Act Theory (L-ActT) (Emanuela Cresti & Panunzi, 2018) which designates the utterance as the reference unit of speech. The boundary between tonal units results from a prosodic break with a non-conclusive value, while the boundary between terminal units corresponds to the perception of a prosodic break with a conclusive value.

In order to obtain a great diaphasic diversity, i.e., according to the communicative context, the project brought a remarkable variety of communicative contexts, compiling scenarios such as communication between players in a football match, the preparation of a drag queen for a presentation, a conversation between a realtor and a client, among others. In addition, a considerable balance was reached regarding the demographic criterion concerning the informants' education and sex. There are 362

<sup>20</sup> <https://www.alip.ibilce.unesp.br/termos-de-uso>.

<sup>21</sup> <http://www.c-oral-brasil.org/>.



informants in the corpus, 138 from the city of Belo Horizonte, 89 from other cities in Minas Gerais, and the rest from other states, countries, or of unknown origin.

There was an effort to use high-quality acoustic equipment at the time. The project used PMD660 Marantz digital recorders and Sennheiser Evolution EW100 G2 wireless kits. It also used non-invasive “clip-on” microphones to create a more natural environment, essential for recording high diaphasic variation in spontaneous speech.

C-ORAL-BRASIL I is available via download from the project website in raw format, morphosyntactically annotated by the Parser Palavras (Bick, 2000), in addition to metadata. The C-ORAL-BRASIL I corpus is licensed under CC BY-NC-SA 4.0. The following files are of special interest for this work: (i) audio in .wav format, with a sampling rate of 48kHz, transcription in .rtf and .txt formats, audio-transcription alignment in XML format generated by the software WinPitch.<sup>22</sup>

### 2.2.3 NURC-Recife

The NURC-Recife corpus has its origins in the 1969 NURC (*Norma Urbana Oral Culta*) project, which documents the spoken language in five Brazilian capitals: Recife, Salvador, Rio de Janeiro, São Paulo and Porto Alegre. NURC-Recife corresponds to the part referring to the linguistic variety spoken in the city of Recife. The corpus is available on the website of the NURC Digital project,<sup>23</sup> developed between 2012–2016. The project NURC Digital was responsible for processing, organizing and releasing the data of the NURC-Recife project in digital form (Oliviera Jr., 2016).

The project is comprised of 290 h spread over 346 recordings (called inquiry in the project) obtained between the years of 1974 and 1988. In fact, this value would be the total duration in hours if all audios and their transcriptions were available on the website. An analysis of all audio-transcription pairs raised one inquiry lacking audio and transcription and 11 inquiries lacking transcriptions, resulting in 279 h available.

The recordings follow NURC guidelines and are categorized as follows:

- Formal utterances (EF), consisting of 37 recordings of lectures and talks given by one speaker;
- Dialogues between two informants (D2) conducted by a mediator, with 71 recordings;
- Dialogues between an informant and an interviewer (DID), with 238 recordings.

The informant ages range from 25 to over 56 years, all of them with higher education and initially selected with equal division (originally 300–300) for male and female voices.

<sup>22</sup> <https://www.winpitch.com/>.

<sup>23</sup> <https://fale.ufal.br/projeto/nurcdigital/>.

The environment of the recordings varied, depending on the type of inquiry: specific rooms, classrooms, auditoriums or even the informants' homes. It also has very heterogeneous noise levels and sound clarity, whether from the equipment used, the recording environment or deterioration of the recording tapes.

The original recordings were captured with omnidirectional dynamic microphones with table support. The reel-to-reel tape recorders used were: AKAI 4000 DS Mk-II, SONY TC-366, and Philips N 4416, the first being the most frequent. The audios were recorded on professional reel magnetic tapes, 0.0018mm thick, 6.35mm wide, and 540m long (BASF TP 18 LH). However, within the scope of the NURC Digital project, they were digitized following the recommendations of the Open Archival Information System (OAIS), in the ISO standard (14721 : 2003), with a sampling rate of 96kHz and quantization of 24 bits. For this digitization, were used the software Audacity, Audiofile Specter, the AKAI 4000 DS Mk-II reel-to-reel recorder, a USB Audio Interface Sound Devices USBPre 2, and the RCA Diamond Cable JX-2055.

NURC Digital is available for academic use, without a defined license, via download from the project website, which allows a search by recording year (1974 to 1988), recording topic, and type of inquiry (D2, DID, and EF). There is also information about the age range of the informants, sex, and audio quality. Within each inquiry folder there are: (i) the digitized version of the specific recording (metadata), in .pdf format; (ii) a file in textgrid format, containing the audio timestamps with the transcriptions; (iii) the audio file of the recording in .wav format (48kHz); (iv) a copy of the audio file, also in .wav format, compressed at a frequency of 44kHz; and (v) the original transcription in .pdf format.

#### 2.2.4 SP2010

The SP2010 project (Mendes & Oushiro, 2012; Projeto SP2010, 2021) started in 2009 and ended in 2013 to document and study the Portuguese spoken in the city of São Paulo. The project was supported by the FAPESP agency between 2011 and 2013, generating a corpus publicly available for academic research.

The corpus contains 60 recordings of 60 to 70 minutes each, collected between 2012 and 2013,<sup>24</sup> with equal division for female and male voices. Each recording identifies an interview with an informant, comprising two parts:

- an informal and spontaneous conversation, with questions about the informant's neighborhood, family, childhood, work and leisure, seeking personal involvement;
- the continuation of the conversation, but exploring a more argumentative speech, with questions on more objective themes about the city of São Paulo, involving problems, solutions, characterizations of the city and its inhabitants. In addition, there are three reading recordings: a list of words, a news article and a statement.

<sup>24</sup> <http://projeto2010.fflch.usp.br/corpus>.

Finally, specific questions about the sociolinguistic varieties of the city are proposed.

The informants were selected to represent 12 sociolinguistic profiles characterized by distinct combinations of the following variations: age group, (with three age groups encompassing individuals from 19 to 89 years), education, (with two school stages represented — up to elementary school and with higher education), and sex (male and female). Each sociolinguistic profile has five informants as representatives, each with a recording. The informants' region of residence within the city was also considered, and a balance of informants was sought in this regard, considering the division of São Paulo into 3 regions: *Centro Velho*, *Centro Expandido* and *Periferia*.

For the recording, the authors used TASCAM DR100 MK2 digital recorders and Sennheiser HMD25-1 microphones, having varied recording conditions, with some interviews being more noisy than others, as they were not conducted in specialized and isolated environments.

The material collected in the SP2010 project is made available via download from the project website, free of charge to the academic community of researchers. Eight files are available for each interview: two audio files — in .wav stereo format, 44kHz, and also in .mp3; four transcription files (in .eaf, .doc, .txt and textGrid formats); the informant and the recording forms (in .xls format); and a .zip file that contains all of the interview materials except the .wav file.

## 2.2.5 TEDx Portuguese

TEDx Portuguese is a new corpus compiled specifically for CORAA ASR. It should not be confounded with the BP audios available in Multilingual TEDx Corpus (described in Sect. 2.1). TEDx Portuguese is based on the TEDx Talks,<sup>25</sup> which are events in which presentations on a wide range of topics take place, and in the same format as the TED Talks,<sup>26</sup> but in languages other than English.

Although they are independent meetings, they are licensed and guided by the TED organization, that is, they are short presentations, containing prepared speech, with a duration recommendation of less than 18 minutes, typically presented by a single presenter. The “x” at the end indicates that the event is carried out by autonomous entities worldwide. More than 3,000 new recordings are made annually.<sup>27</sup>

To create this dataset, we selected presentations spoken in Portuguese, both from Brazil and Portugal, with available preexisting subtitles. After selecting the presentations, they were downloaded, the audios were extracted and converted to .wav format, mono, with a sampling rate of 44kHz. BP presentations have accents from practically all regions of Brazil.

<sup>25</sup> <https://www.ted.com/watch/TEDx-talks>.

<sup>26</sup> <https://www.ted.com/>.

<sup>27</sup> <https://www.ted.com/about/programs-initiatives/TEDx-program>.

The subtitles were also downloaded, with the text extracted exclusively, that is, the timestamps were discarded. The dataset is composed of excerpts from 908 talks (671 of which are in BP), totaling at least 908 different speakers, since there are also talks with more than one speaker. The variant (PT-PT or PT-BR) is annotated in the dataset metadata. Considering both variants, there are 543 male and 375 female voices.

### 3 Data processing pipeline

In this section, we present the processing steps of the CORAA ASR corpus:

1. Normalization of transcriptions,
2. Segmentation and removal of silence and untranscribed parts of speech;
3. Forced alignment between audio and corpora transcription for two corpora<sup>28</sup>;
4. Specific processing in the ALIP and NURC-Recife corpora. For example, (i) to maintain the capitalization of letters indicative of names, to aid in the expansion of names, (ii) to preserve the slashing annotation indicative of truncation in the speaker's speech, to aid the identification of truncated audios, and (iii) to discard audios with duration less than 0.3 s in the NURC-Recife<sup>29</sup>;
5. Validation of audio-transcription pairs, via the web interface created in the project, so that the CORAA ASR corpus can be used for training ASR methods;
6. Evaluation of agreement between annotators and between annotators and the gold-standard annotation, performed by a trained annotator;
7. Error analysis and corpus revision, generating the version 1.1 of CORAA ASR.

All corpora described in Sect. 2.2 were obtained from their respective official websites. After downloading, all transcripts were converted to .csv format and the organization of audio files was standardized. Additionally, due to the differences between the transcription rules of each corpus, text normalization was performed, described in Sect. 3.1. Furthermore, as the ALIP corpus does not originally have alignment between the transcription and the audio file, we performed the forced alignment between the transcription and the audio. TEDx Portuguese has the alignment provided by the subtitles. However, this alignment is limited to 42 characters per line to optimize screen display, and may not correspond to sentence boundaries; for this reason we also performed forced alignment in TEDx Portuguese. We describe the forced alignment process in these two corpora in Sect. 3.2. The validation of the audio-transcription pairs is presented in Sect. 3.3 and the evaluation of agreement between annotators and between annotators and the gold standard annotation is presented in Sect. 3.4. In order to assess the corpus quality, we trained an ASR model using the initial version (version 1.0) of the CORAA ASR corpus. We analysed the

<sup>28</sup> ALIP audios were not originally aligned with their transcripts and TEDx Portuguese was available with segmentation to optimize on-screen presentation.

<sup>29</sup> The original duration of the corpus (279 h) dropped to 216 h.

errors of this ASR model using the test set; this error analysis informed the revision of the corpus. The error analysis and corpus revision process are presented in Sect. 3.5.

### 3.1 Text normalization

The four academic project corpora used their own transcription criteria. The oldest and most widely cited transcription standards are those of the NURC Project, which were used by NURC-Recife. NURC-Recife follows the orthographic transcription and its rules can be found in Preti (1999). During the NURC Digital project, NURC-Recife went through new processing steps, including: quality verification of digitized audio, manual alignment between audio and transcription, spelling revision using a spell checker, which are described by Oliviera Jr. (2016).

The corpus C-ORAL-BRASIL I follows the orthographic-based transcription criteria, but with the implementation of some non-orthographic criteria to capture grammaticalization or lexicalization phenomena (Raso & Mello, 2009). For example, there are aphereses (disappearance of a phoneme at the beginning of a word), reduced prepositions, absence of plural mark in noun phrases, cliticizations of pronouns and pre-verbal negation and articulations of preposition with article.

The SP2010 project uses semi-orthographic transcriptions, using the following criteria: (i) no change in the spelling of words, as phonetic transcription is not used; (ii) no grammatical corrections; (iii) use of parentheses to indicate the deletion of /r/ in syllabic coda, syllable /es/ of the verb “estar” (to be), in all tenses and verb modes, and syllable “vo” of “você(s)” (you). Other deletions were not indicated with marks. Filled pauses, interjections, and conversational markers such as “right?”, “okay?” were pervasively used.

The ALIP project follows the orthographic conventions of the written language, but uses capital initials only for proper names. The transcription annotates the following variable phenomena (Gonçalves, 2019): (i) vowel raising in contexts of medial postonic of nouns, as in “c[o]zinha ~ c[u]zinha” and of verbs, as in “d[e]via ~ d[i]via”; (ii) postonic lifting and syncope medial, as in “pes.s[e].go ~ pes.s[i].go ~ pes.go”; (iii) gerund reduction, as in “canta[ndo] ~ canta[no]”, a striking feature of São Paulo speech.

Results for variable phenomena of morphosyntactic order include, for example, the realization of prepositions with and without contraction, as in “com a ~ cu’a ~ c’a”, “para ~ pra ~ pa”. The corpus proposed a transcription system based on the NURC project and reports the transcription conventions grouped in the following criteria: (i) word spelling, which includes, for example, question and exclamation marks next to the markers discursive and interjections, use of “/” for word truncations; (ii) prosodic elements where it uses an ellipsis for pauses, double-typed colons for lengthening vowels, and interrogation for questions; (iii) interaction in which it identifies the participants of the interaction and use square brackets for voice overlappings; (iv) transcriber’s comments where parentheses are used for hypotheses of what is heard and double parentheses for descriptive comments for laughs, for example.

Considering these differences between the transcriptions and seeking to maintain standardization, we performed the following normalizations in the texts of all CORAA ASR corpora. Some normalizations were performed before validation (items (1), (2), (3)) and practically the entire list below was performed at the end of the entire process, since the ALIP and TEDx Portuguese corpora had their transcriptions revised:

1. Removal of extra annotations that do not belong to the alignment of transcripts and audios, such as annotations that indicate the speech of the interviewer and interviewee, truncations, laughter and extra information provided by the annotators of the projects that make up CORAA ASR corpus;
2. Normalization of texts to lower case;
3. Removal of duplicate spaces;
4. Expansion of acronyms for their forms of pronunciation (standardization applied after validation, to guarantee the expansion of all acronyms);
5. Standardization of some uses of filled pauses, using a reduced set of these: *ah*, *eh* and *uh*. Some variations of these representations have been replaced by the closest of the three above (e.g.: *hum*, *hm*, *uhm* was replaced by *uh*; *éh*, *ehm*, *ehh*, was replaced by *eh*; *huh*, *uh*, *ã* was replaced by *ah*);
6. Expansion of cardinal and ordinal numbers, using the num2words library<sup>30</sup>;
7. Percentage sign expansion (%) for its transcribed form (percentage);
8. Removal of characters such as punctuation and non-language symbols (such as parenthesis and hyphen).

It is important to note that the corpus also brings a great variety of filled pauses forms, so that the model can learn to vary its use, although this richness penalizes the evaluation of models trained with CORAA ASR version 1.0 corpus, as detailed in Sect. 3.5.1.

### 3.2 Automatic forced alignment

As mentioned before, in the ALIP and TEDx Portuguese corpora the alignment between the transcripts and audio was performed using an automatic forced alignment method. For this, we use the tool Aeneas.<sup>31</sup> This tool requires the text segmented into sentences or excerpts.

In the ALIP corpus, the text was segmented using the annotations of pauses or hesitations, indicated by ellipses (“...”) and turn-shifts between speakers, indicated by a line break followed by the next speaker identification abbreviation, present in the original annotated corpus.

In the TEDx Portuguese corpus, the segmentation of text into sentences was performed using the punctuations present in the subtitles, if any. For this, a maximum

<sup>30</sup> <https://github.com/savoirfairelinux/num2words>.

<sup>31</sup> Available at <http://www.readbeyond.it/aeneas>.

limit of 30 words was defined for each sentence and, when this limit was reached, the sentence was divided in the point before this limit. In the case of no punctuation, the sentences were divided in an arbitrary way, for example, in silent passages, or with music, or based on variations in speech rate.

### 3.3 Human validation via web-based platform

The validation of audio-transcription pairs was performed in a simple web interface<sup>32</sup> through two tasks: binary annotation (VALID - INVALID) and transcription to correct automatic alignment effects, as was the case with ALIP corpus, or to review manual transcripts, previously made, as was the case for the TEDx Portuguese corpus.

The binary annotation was carried out by: listening to an audio file that could be listened to as many times as necessary and the reading of the original transcription. The annotation was binary, that is, the pairs were classified as valid or invalid, and it was necessary to point out the reason for such choice, which provided a guide for the choice itself.

There are 3 main reasons an audio is considered invalid:

1. Voice overlapping;
2. Low volume of the main speaker's voice, making the audio incomprehensible;
3. Word truncation.

There are also 3 causes for considering a transcript as invalid, i.e. when it is not aligned with the audio, because there are:

1. Too many words in the transcript;
2. Too few words;
3. Words swapped.

The following options were given to validate an audio/transcript pair:

1. Valid without problems.
2. Valid with filled pause(s).
3. Valid with hesitation.
4. Valid with background noise/low voice but understandable.
5. Valid with little voice overlapping.

In cases where there is an audio with hesitation but the transcription does not correspond to the pauses made, the pair must be invalidated. After one pair has been annotated, another is provided and this process continues until the user wants to stop the annotation and/or disconnect.

---

<sup>32</sup> Public available at: <https://github.com/nile-nlp/BrazSpeechData>.

In the web interface for validation, the transcription task has a screen composed of the original transcription, a player for the audio file that can be repeated as many times as necessary, an editing window initially filled with the original transcription, which is used by the annotator to transcribe, and a button to send the transcription. To complete the task of transcribing an audio, the annotator must listen to the audio.

The annotator must also analyze if this audio fits into any of the types below: music, clapping, word truncation in the audio, loud noise or another language other than Portuguese, very low voice, incomprehensible voice, foul words, hate speech, and loud second voice. If so, the annotator should insert the symbols “####” (denoting invalid audio) in the edit window and send its response. As we focused on the BP, we decided to keep 4.69 h of European Portuguese, so during most of the project, annotators were instructed to discard European Portuguese audios.

The annotators were instructed to comply with the following eight guidelines:

1. Do not change to the grammar normative form the following signs of orality in the audio: “tá/tó, né, cê, cês, pro, pra, dum, duma, num, numa”.
2. Transcribe filled pauses, such as “hum, aham, uh” as heard.
3. Transcribe repetitive hesitations such as “da da”, or “do do” as heard.
4. Write numbers in full form.
5. Letters that appear alone should be spelled out.
6. Acronyms and abbreviations should be transcribed in full form, using the English alphabet for those in English and the Portuguese alphabet for those that appear in Portuguese.
7. Foreign words should be transcribed normally, in the language in which they appeared.
8. Punctuation and case sensitivity could be applied, as normalization is performed in post-processing phase.

### 3.4 Kappa evaluation: subjectivity of the human annotation

The validation of audio-transcription pairs of the CORAA ASR version 1.0 corpus, using the binary annotation and transcription tasks (see Sect. 3.3), was performed from October 2020 to July 2021, when the database export was generated.

The number of annotators varied during the project duration. In total, 63 different annotators performed the validation, which could be divided into 4 main annotation groups according to the start and end dates of each annotator on the project. Two groups validated the corpora for 3 months in 2020 (October to December), with some annotators in this group continuing the validation in 2021. There was a 1-month annotation task-force during December 2020. The final group started the CORAA ASR version 1.0 validation work in May 2021 and ended in July 2021.

Each group attended a lecture on the validation process, read the tutorials for the two tasks (annotation and transcription) and received instructions to ask elucidate doubts via the project email throughout the process.

At the beginning of the validation process, from October to December 2020, each audio-transcription pair was annotated by two or three annotators, so that we could



use the majority vote to export the data, discarding the divergent pairs, in this initial phase of learning how to validate. Agreement between annotators was calculated in two ways: between annotators who annotated the same pairs (Sect. 3.4.1) and based on a gold-standard annotation of samples from all datasets, performed by a project member (Sect. 3.4.2).

### 3.4.1 Kappa among annotators

Two Fleiss kappa values were calculated for the annotation from October to December 2020, to separate the groups of annotators. The project started with two groups in October, totaling 28 annotators, but with the entry of a new group on November 23, 2020 the number of annotators went to 63. Thus, it was decided to calculate a kappa value to evaluate each period of annotation — from October 1st to November 23rd and from November 24th to December 31st, 2020. The hypothesis was that the annotation would become easier and with high agreement as the practice increased. However, there is another variable that influenced the agreement: the different transcription rules for each corpus of the CORAA ASR corpus (see Sect. 3.1) also influenced the agreement. We calculated the agreement value via Fleiss' kappa twice, once considering only two annotators and the other considering only three annotators, according to the total number of annotators of a given audio. The values are shown in Table 2.

It is observed that there are absent values on the table, because the specific corpus was not being annotated in the referred period. The great disagreement between the annotators showed a more subjective task than previously imagined. By manually comparing audios in which annotators agreed with audios in which they disagreed, some points became clear: (i) the human ear naturally tends to complete truncated words, so that different annotators may disagree in defining whether an audio is in fact truncated or not, (ii) background noise level and voice pitch (low/high) are very subjective concepts, and different people are expected to consider different noise levels as tolerable, (iii) naturally, due to the ease of understanding different accents, annotators from different regions of the country tend to understand more or less of the audio according to their accent, which can also be a source of disagreement.

### 3.4.2 Kappa for the gold-standard annotation

The gold standard was built to maintain the representativeness of all validated corpora, and all participating annotators, according to the following process:

1. For each annotated corpus, we generated a list of all annotators in that corpus;
2. For each name present in the list, five pairs annotated by the annotator were randomly selected (annotators with less than 5 pairs annotated per corpus had their pairs discarded);
3. The selected pairs were duplicated and annotated by an experienced annotator of the project, creating a gold-standard annotation with the following distribution:
  - **Alip**: 15 annotators and 75 pairs

**Table 2** Kappa values for each dataset in two annotation periods, separated by number of annotators

	1/10 - 23/11		24/11 - 31/12	
	2 annotators	3 annotators	2 annotators	3 annotators
Number of pairs	6785	29,835	26,974	4224
Number of annotators	25	25	51	51
Kappa values				
C-ORAL-BRASIL I	0.394	0.353	—	—
SP-2010	0.420	0.394	—	—
NURC-Recife	—	—	0.317	0.314
Total	0.391	0.392	0.317	0.314

In the last three months of 2020, the order of validation of the corpora was C-ORAL-BRASIL I, SP-2010 and NURC-Recife

- **C-ORAL-BRASIL I**: 24 annotators and 120 pairs,
- **NURC-Recife**: 55 annotators and 275 pairs,
- **SP-2010**: 25 annotators and 125 pairs,
- **TEDx Portuguese**: 50 annotators and 250 pairs.
- **Total**: 845 pairs (520 from the binary annotation task and 325 from the transcription task)

The consensus pairs between the annotators were included in the exported dataset, that is, if the absolute majority chose to validate the pairs. Thus, we analyzed the degree of agreement of the annotators together (exported values) in comparison with the gold-standard corpus. The value obtained was **0.514**, showing a “moderate agreement”, according to Landis and Koch (1977). Even though the task is subjective, the final result obtained from the annotation of the exported pairs was satisfactory.

### 3.5 Error analysis and corpus revision

In order to assess the dataset quality, we trained a model based on the architecture Wav2Vec 2.0 XLSR-53 (Conneau et al., 2020; Baevski et al., 2020) using the dataset version 1.0. For training this preliminary ASR model, we used the same procedure described in the Sect. 5.1. Before model training, the dataset was divided into three subsets: train, development and test. Table 3 presents the approximate number of hours for these sets for each sub-dataset, as well as the number of speakers from each sex. Sub-dataset validation sets were adjusted to have approximately 1 h. Test sets were built in a similar way, but having approximately 2 h. This decision is supported by the work of Sheshadri et al. (2021), which recommends that test sets should have at least 2 h. NURC-Recife test set contains more than 3 h of audios, because this sub-dataset has more speech genres than the others. All the audios from European Portuguese were included in the training set.

**Table 3** Statistics of Train/Dev/Test partitions of each CORAA ASR subcorpus (version 1.0)

Subset	Duration (h)			Num. Speakers (M—F)		
	Train	Dev	Test	Train	Dev	Test
ALIP	36.62	1.08	1.96	80—87	2—2	4—5
C-ORAL Brasil I	6.69	1.16	2.15	138—181	9—9	12—13
NURC-Recife	138.68	1.32	3.12	295—296	2—1	3—3
SP2010	28.17	1.15	2.41	27—27	1—1	2—2
TEDx Portuguese	73.45	1.05	1.99	532—364	4—4	7—7
Total	283.61	5.76	11.63	1072—955	18—17	28—30

### 3.5.1 Error analysis

The test dataset used for error analysis is composed of 13,931 pairs of audio-transcription pairs, totaling 11.63 h, with parts from all CORAA ASR version 1.0 dataset. As this is the first time that a dataset composed of spontaneous speech samples was used to train an ASR model for BP, we performed a more detailed analysis of the errors from our model in a sample of the test dataset.

The 13,931 test pairs were ordered by the Character Error Rate (CER)<sup>33</sup> values of our model to illustrate the different types of errors and to analyze whether there is a relationship of error types with CER values. The automatic transcription was analyzed using the typology of da Mota et al. (2000), adapted for the task of evaluating ASR models.

The typology used here to illustrate the model errors is composed of 11 error types, grouped into 6 more general classes: Alphabetical, Lexical, Morphological, Language and Spontaneous Speech, Semantic, and Diacritic Placement Errors. Below, we present a description of the 11 types of errors with examples.

- **Alphabetical errors** are alphabetic writing application errors. (1) Alphabetical errors occur in 3 situations: by transcribing speech directly into writing, in complex syllables or even with ambiguous letters (“ce” versus “sse” or “sa” versus “za”, in Portuguese). An example of this type of error is related to the sound /k/ in Portuguese which is represented by the letter “c” before some vowels and by “qu” before other vowels. Thus, the use of “c” in place of “qu” is associated with the speaking/writing relationship.
- **Lexical errors** occur in an excerpt transcribed by the ASR where there is: (2) omission or addition of words; (3) exchange of words. An example from our dataset regarding addition of a word in the automatic transcription is “que legal” instead of “legal”. Also from our dataset, an example of word exchange is “e que mais que a gente vida” instead of “e que mais que a gente viu”.

<sup>33</sup> CER is very similar to word error rate, however, this metric considers the number of substitutions, insertions, deletions and the number of characters instead of words.

**Table 4** Examples of the 11 different error types

Error type	Original transcription	ASR transcription
1	- uma maneira de saber o que e como o indivíduo <b>identifica</b> algo	- uma maneira de saber o que e como o indivíduo <b>identifica</b> algo
2	- ou pra dar um apoio moral	- ou pra dar um apoio <b>im</b> moral
3	- o outro <b>foi</b> morar um pouco mais longe	- o outro <b>prai</b> morar um pouco mais longe
5,4	- <b>que lhe dão</b> ora <b>dor</b>	- <b>ciridão</b> ora <b>do</b>
5	- criança é mais <b>coca cola</b> * biscoito	- criança é mais <b>cocacola</b> biscoito
6	- que levaria a uma resposta <b>aquele</b> estímulo - ah legal faz <b>tempão</b> já	- que levaria a uma resposta <b>a quele</b> estímulo - ah legal faz <b>tem pão</b> já
7	- na teoria de <b>osgood</b> é que - de <b>jazz</b>	- na teoria de <b>osguot</b> é que - de <b>dez</b>
8	- e essa daí <b>eh - eh - eh - ham</b>	- e essa daí <b>é - é - ahn - uhn</b>
9	- <b>pra</b> área específica que é o curso diz que é um curso excelente	- <b>para</b> área específica que é o curso diz que é um curso excelente
10	- entendeu era eles <b>suavam</b> mais a camisa pelo clube entendeu e	- entendeu era eles <b>soavam</b> mais a camisa pelo clube entendeu e
11	- então <b>é</b> conhecer a população usuária do equipamento urbano	- então <b>e</b> conhecer a população usuária do equipamento urbano

\*The lack of a hyphen in the test set is only for the calculation of CER/WER

**Table 5** Intervals of CER and frequencies of the different error types

Intervals	CER	Analysed samples	Error types (occurrences)
1–4613	0*	—	—
4614–8397	$0 < \text{CER} < 0.1$	110	2 (3), 3 (3), 4 (13), <b>5 (64)</b> , 6 (29), 8 (2), 9 (1), 10 (1), 11 (2)
8398–10,723	$0.1 \leq \text{CER} < 0.2$	10	<b>2 (5)</b> , <b>3 (5)</b> , 7 (1)
10,724–11,991	$0.2 \leq \text{CER} < 0.3$	10	2 (2), <b>3 (7)</b> , 8 (1)
11,992–12,666	$0.3 \leq \text{CER} < 0.4$	10	2 (5), <b>3 (16)</b> , 9 (1)
12,667–13,049	$0.4 \leq \text{CER} < 0.5$	25	2 (7), <b>3 (27)</b> , 4 (5), 5 (3), 8 (2)
13,050–13,335	$0.5 \leq \text{CER} < 0.6$	10	2 (3), <b>3 (5)</b> , 4 (2), 7 (1), 9 (1)
13,336–13,507	$0.6 \leq \text{CER} < 0.7$	10	2 (3), <b>3 (4)</b> , 4 (1), 5 (1), 7 (1), 8 (1), 11 (1)
13,508–13,931	$0.7 \leq \text{CER} \leq 12$	76	2 (4), <b>3 (38)</b> , 7 (1), 8 (36), 9 (2), 11 (1)

\*CER and WER are equal to zero

The most frequent error type in each interval appears in bold

- **Morphological errors** are errors that occur due to the violation of writing rules that is linked to the morphological structure of words. These are errors from: (4) omitting morphemes (e.g. “come” written instead of “comer”); (5) concatenation of morphemes (e.g., “agente” instead of “a gente”, or “acasa” instead of “a casa” ); (6) separation of morphemes, as in the example: “de ele” written instead of the contraction “dele”).
- **Language and spontaneous speech errors** are errors of: (7) Words in English (or in a language other than BP) wrongly transcribed; (8) Filled pause errors (e.g., “á” versus “eh” ) where the transcription and model responses diverge; (9) Spontaneous speech errors (e.g., “tá” versus “está”; “té” versus “até”; “cê” versus “você”) in which transcription and model responses diverge.
- **Semantic errors** occur when two words are spelled similarly but have different meanings. (10) Semantic errors (e.g. “Ela complimentou o diretor assim que chegou.”, where the correct form would be “cumprimentou”).
- **Diacritic placement errors** occur due to missing accents or improperly adding them. They are problematic because the five training corpora were built at different times, in which there were different spelling rules for the Portuguese

**Table 6** Sub-corpus and frequencies of the different error types

Sub-corpus	Analysed Samples	Error types (occurrences)
ALIP	58	2 (10), <b>3 (38)</b> , 4 (2), 5 (10), 6 (2), 7 (10), 8 (7), 10 (1)
C-ORAL BRASIL I	50	2 (9), <b>3 (24)</b> , 4 (6), 5 (11), 6 (4), 7 (1), 8 (8), 9 (4), 11 (1)
NURC Recife	66	2 (5), <b>3 (24)</b> , 4 (5), 5 (21), 6 (16), 7 (2), 8 (1), 9 (10), 11 (1)
SP2010	70	2 (7), 3 (19), 4 (7), 5 (14), 6 (5), 7 (1), <b>8 (25)</b> , 11 (2)
TEDx Portuguese	17	2 (1), 4 (2), <b>5 (12)</b> , 6 (3), 8 (1)

The most frequent error type in each sub-corpus appears in bold

language. For example, the last orthographic agreement for the Portuguese language came into force in Brazil in 2016. (11) Accent-marks errors.

Table 4 shows examples of 11 errors presented above (column 1), in which the original transcription (column 2) and the model response (column 3) diverge. The word(s) in focus in the original transcription and the error(s) in the ASR transcription appear in bold.

A sample of 708 audio-transcription pairs was analyzed, of which 133 contained some errors in the audio transcription and thus they were not framed in the typology. Also, 314 pairs were annotated for deletion as their audios were compromised (because of truncation, very loud noise or overlapping voices). In the remaining 261 pairs, the error types were analysed both by the CER intervals shown in column 1 of Table 5 and by the sub-corpus shown in column 1 of Table 6.

Error types are based on the typology presented above. For some pairs more than one error occurs and for some excerpts with high CER values only one error was annotated (the most frequent) although the transcription had many more.

Table 5 shows, in the last column, the variety of error types in each range presented in column 1; its frequency is shown in parentheses. We present in bold the most frequent type. The lexical error of type 3 — exchange of words — is the most frequent one, which is expected given that the task is automatic transcription, and the training process of these models favors the recognition of frequent and well formed words. Moreover, omission and addition of words (error type 2) is pervasive as it appears in all the intervals. However, the second and third errors classified by frequency are: concatenation (error type 5) and filled pause swap error (error type 8). The latter is related to the fact that the CORAA ASR dataset has a large percentage of spontaneous speech samples in which both the number and variety of filled pauses are high.

Table 6 shows, in the last column, the variety of error types in each sub-corpus presented in column 1; its frequency is shown in parentheses. We present in bold the most frequent error type. TEDx Portuguese is a prepared speech corpus with a duration recommendation for the talks of less than 18 minutes; the talks are fast, and therefore it presents high frequency of concatenation of morphemes. ALIP, C-ORAL-BRASIL and NURC-Recife present the most common type of error for an ASR (lexical error of type 3). For SP2010, the most common type of error is filled pause swap error. This corpus has many filled pauses, so it is natural that it presents a high filled pause error rate.

### 3.5.2 Revision of the dataset

After the error analysis, it became clear the need for more normalization rules for filled pause representations so that the model accuracy increases. Moreover, this

initial analysis resulted in a decision to make a revision in all pairs of the test set, and a partial revision in the development and training set. All the audios in the test set where the model predictions differed from the data set labels were reviewed by the annotators. Regarding the training and development sets, the audios were sorted by CER in descending order and the annotators reviewed all the audios with  $CER > 0.3$ . Finally, a post-processing step included normalization rules for filled pause representations, generating the version 1.1 of the dataset.

## 4 Dataset statistics

Overall, CORAA ASR version 1.1 has 290 h of validated audios, containing at least 65% of its contents in the form of spontaneous speech. We refer to the processed version of corpora in CORAA ASR as sub-datasets. NURC-Recife sub-dataset includes conference and class talks, considered prepared speech (see Table 1). Currently, no other dataset for BP includes audios with this speaking style. Therefore, the task of ASR is more challenging than for other datasets. Another feature of CORAA ASR is the presence of noise in some of its sub-datasets, which is also more challenging for creating models for this task. Table 7 presents statistics for each validated sub-dataset in CORAA ASR. The resulting set encompasses almost 1,700 speakers.



Audio durations range, in average, for 2.4 to 7.6 s according to sub-dataset. Audios having more than 200 words or 40 s were automatically filtered from the dataset. Figure 1 presents estimated speaker distribution in each sub-dataset according to sex. Overall, the distribution is similar for males and females.<sup>34</sup> Figure 2 presents audio duration distributions by sub-dataset. The audios are ranked by duration and their relative position (percentile) is shown in the  $x$  axis. Audios duration are presented in the  $y$  axis. Percentiles are used to simplify sub-dataset comparisons. Figure 3 is similar, but presenting word distribution per dataset.

Regarding duration, the segmentation process play a role in the obtained durations. Only ALIP and TEDx Portuguese were automatically segmented. The other sub-datasets were manually segmented. For the automatic segmentation, the parameters were adjusted aiming at better segmentation of informational units. ALIP had a similar duration than the others dataset. However, TEDx Portuguese audios tended to be longer. Speech style and genre also play a role in the obtained results. When pronunciation is faster and with less pauses, there are less places in the audio that the segmentation software is confident to break the utterances. TEDx Portuguese is the main source of prepared speech in CORAA ASR and had the longest audios and the same applies to word distribution, which is natural since the audios are longer. The remaining sub-corpora presented similar distributions among them.


<sup>34</sup> In the corpus C-ORAL-BRASIL I there is a balance regarding number of uttered words — 50.36% words are uttered by 203 females and 49.64% words are uttered by 159 males.

**Table 7** Statistics for each processed version of the projects included in CORAA ASR (hours in decimal)

Corpus	ALIP	C-ORAL BRASIL I	NURC Recife	SP2010	TEDx Port.	Total
Original (h)	78	21.13	279	65	249	692.21
Validated (hrs)	35.96	9.64	141.31	31.14	72.74	290.79
BP Speakers	179	362	417	60	671	1,689
Audios (segmented)	45,006	13,668	261,906	46,482	35,404	402,466
Audio duration (s)	2.90	2.46	1.94	2.43	7.55	3.39
Avg Tokens	53.910	60.079	20.418	48.118	166.369	41.546
Avg Types	6.391	7.188	3.733	6.002	8.807	5.581
Total Tokens	335,664	99,954	1,378,558	339,890	610,639	2,764,705
Total Types	14,189	8,715	41,903	12,351	27,469	58,237
Type/Token Ratio	0.042	0.087	0.030	0.036	0.046	0.022

**Fig. 1** Estimated Speaker Distribution by Sex**Fig. 2** Duration distribution per sub-dataset. Audios are sorted by duration and their percentiles are presented in the  $x$  axis. Duration in seconds is presented in the  $y$  axis. For example, 36% of audios in TEDx Portuguese contain six or less seconds





**Fig. 3** Word distribution per sub-dataset. Audios are sorted by absolute number of words and their percentiles are presented in the  $x$  axis. Absolute number of words is presented the  $y$  axis. For example, 64% of audios in C-ORAL-BRASIL I contain ten or less words

## 5 Baseline model development

We performed an experiment over CORAA ASR version 1.1 dataset in order to assess the dataset quality, potentials and limitations. For this we used the final numbers of hours for Train/Dev/Test after revision of the dataset. Table 8 presents the number of hours for each sub-dataset, as well as the number of speakers from each sex for the CORAA ASR version 1.1 dataset.

### 5.1 Proposed experiment

Our proposed experiment is based on the work of Gris et al. (2021). These authors fine-tuned the model Wav2Vec 2.0 XLSR-53 (Baevski et al., 2020; Conneau et al., 2020) for ASR, using publicly available resources for BP. One of their experiments consisted of training on 437.2 h of Brazilian Portuguese. Wav2Vec 2.0 is a model that learns quantized latent space representation from audios by solving a contrastive task. First, the model is pre-trained using an unsupervised approach in large datasets. Then, it is fine-tuned for the ASR task using supervised learning. Wav2Vec XLSR-53 is pre-trained over 53 languages, including Portuguese.

In our approach, Wav2Vec XLSR-53 is fine-tuned for CORAA ASR version 1.1. We also evaluated the fine-tuned public model developed by Gris et al. (2021) against CORAA ASR version 1.1, using the sets presented in Table 3.

Using the proposed training, development and testing divisions for CORAA ASR version 1.1, we explored training Wav2Vec 2.0 XLSR-53 model using CORAA ASR version 1.1 during 40 epochs. Similarly to the work of Conneau et al. (2020) and Gris et al. (2021), we opted to freeze the model feature extractor.

**Table 8** Statistics of Train/Dev/Test partitions of each CORAA ASR subcorpus (version 1.1)

Subset	Duration (h)			Num. Speakers (M—F)		
	Train	Dev	Test	Train	Dev	Test
ALIP	33.40	0.99	1.57	80—87	2—2	4—5
C-ORAL-BRASIL I	6.54	1.13	1.97	138—181	9—9	12—13
NURC-Recife	137.08	1.29	2.94	295—296	2—1	3—3
SP2010	27.83	1.13	2.18	27—27	1—1	2—2
TEDx Portuguese	68.67	1.37	2.70	532—364	4—4	7—7
Total	273.51	5.91	11.35	1072—955	18—17	28—30

To train the model, we use the framework HuggingFace Transformers (Wolf et al., 2020). The model was trained with GPU NVIDIA TESLA V100 32GB using a batch size of 8 and gradient accumulation over 24 steps. We used the optimizer AdamW (Loshchilov & Hutter, 2019) with a linear learning rate warm-up from 0 to 3e-05 in the first two epochs and after using linear decay to zero. During training the best checkpoint was chosen, using the loss in the development set. The code used to perform the experiment as well as the checkpoint of the trained model are publicly available at: <https://github.com/Edresson/Wav2Vec-Wrapper>.

## 5.2 Results and discussions

Section 5.2.1 presents a comparison of our results with the work of Gris et al. (2021). The models are tested against the entire test subset of CORAA ASR version 1.1 and Common Voice version 7.0 (Portuguese audios). Therefore, our model is evaluated *in-domain* using CORAA ASR version 1.1 test set, a dataset in which it was fine-tuned for specific recording characteristics. At the same time, our model is also evaluated *out-of-domain* in Common Voice, a dataset completely new to our model.

Additionally, Sect. 5.2.2 focuses on evaluating the models in test sets of CORAA ASR sub-datasets. This enables a more detailed analysis on factors such as audio quality and accents. Finally, Sect. 5.2.3 investigates the two speech styles: prepared or spontaneous.

### 5.2.1 In/out of domain evaluation

Table 9 presents the comparison of our experiment with the work of Gris et al. (2021). First, we performed an in-domain analysis of our model using CORAA ASR version 1.1 test set. Then, our model is evaluated out-of-domain using Common Voice test set. It is important to observe that, for the compared work, the analysis is mirrored, there is, CORAA ASR version 1.1 is the out-of-domain evaluation and Common Voice is the in-domain analysis.

**Table 9** Results for the In/Out of Domain Analysis

Datasets	Common voice		CORAA ASR		Mean	
	CER	WER	CER	WER	CER	WER
Gris et al. (2021)	<b>4.15</b>	<b>13.85</b>	22.32	43.7	13.23	28.77
Our	6.34	20.08	<b>11.02</b>	<b>24.18</b>	<b>8.68</b>	<b>22.13</b>

The best values for CER and WER appear in bold

In the Common voice dataset, as expected, Gris et al. (2021) model performed better. Regarding Word Error Rate (WER), it can be noted that our model is less than 7% above their work. We also focus our analysis on the metric CER, because for smaller audios, with just a few words, this metric tends to be more reliable. In this scenario, our models are approximately 2% worse than the model from Gris et al. (2021). On the other hand, in the CORAA ASR dataset, our model presented a much superior performance (more than 19% in WER and 11% in CER). Furthermore, our experiment managed to generalize better for audio characteristics not seen during training, achieving an average higher than the performance of the Gris et al. (2021). This is very interesting especially because the Gris et al. (2021) model was trained with approximately 147 h of speech more than our model.

We believe that models trained with the CORAA ASR version 1.1 dataset generalize better than a model trained with existing publicly available datasets for BP due to the spontaneous speech phenomenon and the wide range of noise and different acoustic characteristics present in CORAA ASR. Furthermore, accent can be a factor since the datasets used in the training of the Gris et al. (2021) model may not cover in depth all accents present in the CORAA ASR version 1.1.

### 5.2.2 Sub-dataset analysis

There are important differences in the recording environment for each sub-dataset. Additionally, they also vary on accents. Table 10 presents the performance in the test set for each sub-dataset of CORAA ASR version 1.1.

Regarding datasets, ALIP presented the greatest challenge for the models, both for CER and WER metrics. We believe this occurred because audios from ALIP presented more noise than the other sub-datasets.

Regarding accents, we have different results. On one hand, our model presented similar performances in NURC-Recife and SP2010, which have two distinct accents (Recife and São Paulo city). On the other hand, C-ORAL-BRASIL presented higher WERs and CERs than the other two. Two factors may have influenced this result. First, audio quality and noise presence tend to play a major role in model performances. Second, C-ORAL-BRASIL accent (Minas Gerais) has two characteristics that are difficult for models: speech rate is faster and there is more word agglutinations. As a consequence, the analysis was inconclusive for this accent, since the results are influenced both by the accent and the speech rate.

**Table 10** Results in the CORAA ASR test set for all subsets

Datasets	Gris et al. (2021)		Our	
	CER	WER	CER	WER
ALIP	33.72	59.30	17.30	34.06
C-ORAL-BRASIL I	23.53	45.9	13.62	28.88
NURC-Recife	19.46	42.17	9.09	22.03
TEDx Portuguese	9.75	22.69	7.43	19.36
SP2010	23.11	42.44	9.57	20.00

**Table 11** Results for Spontaneous versus Prepared Speech

Speech style	Spontaneous speech		Prepared speech	
	CER	WER	CER	WER
Gris et al. (2021)	25.75	49.18	<b>5.30</b>	<b>15.89</b>
Our	<b>12.44</b>	<b>26.5</b>	6.07	18.7

The best values for CER and WER appear in bold

Regarding experiments, our model presented results varying from 19 to 34% in WER and from 7 to 17% in CER. On the other hand, Gris et al. (2021) presented higher error rates, which is expected considering the training of their model had no previous contact with CORAA ASR version 1.1 audios.

### 5.2.3 Spontaneous versus prepared speech analysis

Table 11 presents an analysis in which sub-datasets are merged according to speech style. The Spontaneous Speech column is obtained from the merging of ALIP, C-ORAL-BRASIL I, SP2010 and parts of NURC-Recife. The prepared speech column contains TEDx Portuguese and parts of NURC-Recife. As expected, the models perform better on prepared speech. However, for several ASR applications, spontaneous speech is more relevant (for example, ASR of phone call and meetings). This can also be observed in Sect. 5.2.2, as TEDx Portuguese presented the lowest error rates.

## 6 Conclusions and future work

In this paper we presented and made publicly available a new dataset called CORAA ASR version 1.1, with 290 h of validated pairs of audio-transcription, composed of public corpora in BP and TEDx Talks in European and Brazilian Portuguese.

Counting on the cooperation among research centers, universities, private companies and The São Paulo Research Foundation (FAPESP), we made publicly available this new and large dataset for training BP speech recognition models, closing the gap of previous datasets, i.e., the lack of spontaneous and informal speech used

in conversations, dialogues and interviews. Informed by the error analysis, we normalized filled pauses representations and performed a revision over the test, development and train datasets, in order to increase future ASR model accuracy. We also proposed an ASR Challenge including CORAA ASR version 1.1 to further develop research in ASR for the Portuguese language, in order to motivate young researchers in this exciting research area.

Our work has the following limitation. C-ORAL-BRASIL I and NURC-Recife had extra annotations on the morphosyntactic and syntactic levels. However, we could not keep these annotations in CORAA ASR for the following reasons. First, some audio fragments were removed, for example, due to voice overlapping. Second, some audio fragments were edited, for example, Arabic numerals were changed to the number in full format. Third, there are some corrections in transcriptions, because even in the original corpora, transcriptions errors may still occur.

As for future work, we plan to enlarge CORAA ASR with new corpora from Tarsila Project<sup>35</sup> such as Museu da Pessoa<sup>36</sup> and NURC-SP.<sup>37</sup> Moreover, with the current availability of new forced phonetic aligners for Brazilian Portuguese ((McAuliffe et al., 2017), (Dias et al., 2020), (Kruse & Barbosa, 2021), (Batista et al., 2022)) we intend to evaluate the performance of these new tools in order to choose the best forced aligner for a specific corpus, speech genre and accent.<sup>383940414243</sup>

**Acknowledgements** This research was funded by CEIA with support by the Goiás State Foundation (FAPEG grant #201910267000527)<sup>41</sup>, Department of Higher Education of the Ministry of Education (SESU/MEC), Copel Holding S.A.,<sup>42</sup> and Cyberlabs Group<sup>43</sup>. In addition, this research was carried out at the Center for Artificial Intelligence (C4AI-USP), with support by the São Paulo Research Foundation (FAPESP grant #2019/07665-4) and by the IBM Corporation. Also, this study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001. We also would like to thank Nvidia Corporation for the donation of Titan V GPU used in CORAA ASR related projects. The coauthor Anderson da Silva Soares thanks to CNPq for Productivity Scholarship in Technological Development and Innovative Extension - number 308808/2020-7. Finally, the authors would like to thank all the members of the TARSILA project that contributed with discussions and insights regarding the compilation of CORAA ASR v1 corpus.

**Author Contributions** ACJ: Conceptualisation, Investigation, Project administration, Methodology, Resources, Supervision, Validation, Software Development, Validation and Writing—original paper. EC: Conceptualisation, Investigation, Project administration, Methodology, Resources, Software Development and Writing—original paper. AS: Funding acquisition, Project administration and Writing—original paper. FSdO: Resources and Writing—original paper. LO, RCFJ, DPPdS and LRSG: Software Development and Writing—original paper. BBC and FGF: Data curation, Resources, Software Development and Writing—original paper. SMA: Conceptualisation, Data curation, Investigation, Project administration, Methodology, Resources, Supervision, Validation and Writing—original paper.

<sup>35</sup> <https://sites.google.com/view/tarsila-c4ai>.

<sup>36</sup> <https://museudapessoa.org/>.

<sup>37</sup> <https://nurc.flch.usp.br/>.

<sup>38</sup> <http://centrodeia.org/>.

<sup>39</sup> <https://www.copel.com>.

<sup>40</sup> <https://cyberlabs.ai/>.

<sup>41</sup> <http://centrodeia.org/>.

<sup>42</sup> <https://www.copel.com>.

<sup>43</sup> <https://cyberlabs.ai/>.

**Funding** This research was funded by CEIA with support by the Goiás State Foundation (FAPEG grant #201910267000527)<sup>38</sup>, Department of Higher Education of the Ministry of Education (SESU/MEC), Copel Holding S.A.,<sup>39</sup> and Cyberlabs Group<sup>40</sup>. The coauthor Anderson da Silva Soares thanks to CNPq for Productivity Scholarship in Technological Development and Innovative Extension - number 308808/2020-7. This research was carried out at the Center for Artificial Intelligence (C4AI-USP), with support by the São Paulo Research Foundation (FAPESP grant #2019/07665-4) and by the IBM Corporation.

**Data availability** CORAA ASR dataset is available, in csv format, in the file DATA at <https://github.com/nilc-nlp/CORAA>, under CC BY-NC-SA 4.0 license.

**Code availability** Source Code of the models are available at <https://github.com/Edresson/Wav2VecWrapper>.

## Declarations

**Conflict of interest** The authors have no conflicts of interest to declare.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Alencar, V. F. S., & Alcaim, A. (2008). Lsf and lpc—Derived features for large vocabulary distributed continuous speech recognition in brazilian portuguese. In *2008 42nd Asilomar Conference on Signals, Systems and Computers* (pp. 1237–1241). <https://doi.org/10.1109/ACSSC.2008.5074614>
- Ardila, R., Branson, M., Davis, K., Kohler, M., Meyer, J., Henretty, M., Morais, R., Saunders, L., Tyers, F., & Weber, G. (2020). Common voice: A massively-multilingual speech corpus. In *Proceedings of the 12th language resources and evaluation conference, European language resources association, Marseille, France* (pp. 4218–4222). <https://www.aclweb.org/anthology/2020.lrec-1.520>
- Baevski, A., Zhou, Y., Mohamed, A., & Auli, M. (2020). wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33, 12449–12460.
- Batista, C. T., Dias, A. L., & Neto, N. (2022). Free resources for forced phonetic alignment in Brazilian Portuguese based on kaldi toolkit. *EURASIP Journal on Advances in Signal Processing*, 1, 11. <https://doi.org/10.1186/s13634-022-00844-9>
- Baumann, T., Kennington, C., Hough, J., & Schlangen, D. (2016). Recognising conversational speech: What an incremental ASR should do for a dialogue system and how to get there. In Jokinen, K., & Wilcock, G. (Eds.), *Dialogues with social robots—Enablers, analyses, and evaluation, seventh international workshop on spoken dialogue systems, IWSDS 2016, Saariselkä, Finland, January 13–16, 2016*. Springer. Lecture Notes in Electrical Engineering (Vol. 427, pp. 421–432). [https://doi.org/10.1007/978-981-10-2585-3\\_35](https://doi.org/10.1007/978-981-10-2585-3_35).
- Baumann, T., Köhn, A., & Hennig, F. (2019). The spoken Wikipedia corpus collection: Harvesting, alignment and an application to hyperlistening. *Language Resources and Evaluation*, 53(2), 303–329.
- Bick, E. (2000). The parsing system “Palavras”. Automatic grammatical analysis of Portuguese in a constraint grammar framework. University of Aarhus, Århus






- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, É., Ott, M., Zettlemoyer, L., & Stoyanov, V. (2020). Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th annual meeting of the association for computational linguistics* (pp. 8440–8451).
- Cresti, E., Gregori, L., Moneglia, M., & Panunzi, A. (2018). The language into act theory: A pragmatic approach to speech in real-life. In Koiso, H., & Paggio, P. (Eds.), *Proceedings of the eleventh international conference on language resources and evaluation (LREC 2018)*, European Language Resources Association (ELRA), Paris, France
- da Mota, M., Moussatchè, A. H., de Castro, C. R., de Moura, M. L. S., & D'Angelis, T. (2000). Erros de escrita no contexto: uma análise na abordagem do processamento da informação. *Psicologia: Reflexão e Crítica* [online] 13(1). <https://doi.org/10.1590/S0102-79722000000100002>.
- Dias, A. L., Batista, C., Santana, D., & Neto, N. (2020). Towards a free, forced phonetic aligner for Brazilian Portuguese using kald tools. In R. Cerri & R. C. Prati (Eds.), *Intelligent systems* (pp. 621–635). Springer. [https://doi.org/10.1007/978-3-030-61377-8\\_44](https://doi.org/10.1007/978-3-030-61377-8_44)
- Fujimura, H., Nagao, M., & Masuko, T. (2018). Simultaneous speech recognition and acoustic event detection using an lstm-ctc acoustic model and a wfst decoder. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 5834–5838)
- Gonçalves, S. C. L. (2007). Banco de dados Iboruna: amostras eletrônicas do português falado no interior paulista. Retrieved January 1, 2021, from <https://www.alip.ibilce.unesp.br/>.
- Gonçalves, S. C. L. (2019). Projeto ALIP (amostra linguística do interior paulista) e banco de dados iboruna: 10 anos de contribuição com a descrição do português brasileiro. *Revista Estudos Linguísticos*, 48(1), 276–297.
- Gris, L. R. S., Casanova, E., de Oliveira, F. S., da Silva Soares, A., Junior, A. C. (2021). Brazilian Portuguese speech recognition using wav2vec 2.0. 2107.11414.
- Hernandez, F., Nguyen, V., Ghannay, S., Tomashenko, N. A., Estève, Y. (2018). TED-LIUM 3: Twice as much data and corpus repartition for experiments on speaker adaptation. In Karpov, A., Jokisch, O., & Potapova, R. (Eds.), *Speech and computer—20th international conference, SPECOM 2018, Leipzig, Germany, September 18–22, 2018, Proceedings*. Lecture Notes in Computer Science (vol. 11096, pp. 198–208). Springer. [https://doi.org/10.1007/978-3-319-99579-3\\_21](https://doi.org/10.1007/978-3-319-99579-3_21).
- Inaguma, H., Inoue, K., Mimura, M., Kawahara, T. (2017). Social signal detection in spontaneous dialogue using bidirectional LSTM-CTC. In *Proceedings of the Interspeech 2017* (pp. 1691–1695). <https://doi.org/10.21437/Interspeech.2017-457>
- Kruse, J., & Barbosa, P. (2021). Alinha-pb. *Journal of Communication and Information Systems*, 36(1), 192–199. <https://doi.org/10.14209/jcis.2021.21>
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174.
- Loshchilov, I., & Hutter, F. (2019). Decoupled weight decay regularization. In *7th international conference on learning representations, ICLR 2019*, New Orleans, LA, USA, May 6–9, 2019, OpenReview.net. <https://openreview.net/forum?id=Bkg6RiCqY7>
- Macedo Quintanilha, I., Lima Netto, S., & Pereira Biscainho, L. (2020). An open-source end-to-end ASR system for Brazilian Portuguese using DNNs built from newly assembled corpora. *Journal of Communication and Information Systems*, 35(1), 230–242. <https://doi.org/10.14209/jcis.2020.25>
- McAuliffe, M., Socolof, M., Mihuc, S., Wagner, M., & Sonderegger, M. (2017). Montreal forced aligner: Trainable text-speech alignment using kald. *Interspeech*, 2017, 498–502.
- Mendes, R. B., & Oushiro, L. (2012). Mapping paulistano Portuguese: The sp2010 project. In *Proceedings of the VIIIth GSCP international conference: speech and corpora* (pp. 459–463). Fizenze University Press, Firenze, Italy
- Oliviera, M., Jr. (2016). Nurc digital um protocolo para a digitalização, anotação, arquivamento e disseminação do material do projeto da norma urbana linguística culta (nurc). *CHIMERA: Revista de Corpus de Linguas Romances y Estudios Lingüísticos*, 3(2), 149–174.
- Panayotov, V., Chen, G., Povey, D., & Khudanpur, S. (2015). Librispeech: An ASR corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 5206–5210)
- Pratap, V., Xu, Q., Sriram, A., Synnaeve, G., Collobert, R. (2020) Mls: A large-scale multilingual dataset for speech research. Interspeech 2020 <https://doi.org/10.21437/interspeech.2020-2826>
- Preti, D. (1999). Normas para transcrição dos exemplos. In: Preti D (ed) Análise de Textos Oraís, Série Projetos Paralelos, vol 1, 4th edn, Humanitas Publicações - FFLCH/USP (pp. 11–12)



- Projeto SP2010. (2021). Projeto SP2010: Amostra da fala paulistana. Retrieved July 11, 2021, from <https://projetosp2010.flch.usp.br/corpus>.
- Raso, T., & Mello, H. (2009). Parâmetros de compilação de um corpus oral: o caso do c-oral-brasi. *Vere-das*, 13, 20–35.
- Raso, T., & Mello, H. (2012). *C-oral - Brasil I: Corpus de Referência do Português Brasileiro Falado Informal*. Editora UFMG.
- Raso, T., Mello, H., & Mittmann, M. M. (2012). The C-ORAL-BRASIL I: Reference corpus for spoken Brazilian Portuguese. In *Proceedings of the eighth international conference on language resources and evaluation (LREC'12)*, European Language Resources Association (ELRA) (pp. 106–113), Istanbul, Turkey. [http://www.lrec-conf.org/proceedings/lrec2012/pdf/624\\_Paper.pdf](http://www.lrec-conf.org/proceedings/lrec2012/pdf/624_Paper.pdf).
- Raso, T., Mello, H., & Mittmann, M. (2015). O projeto c-oral-brasil. *CHIMERA: Revista de Corpus de Linguas Romances y Estudios Lingüísticos*, 1, 31–67.
- Salesky, E., Wiesner, M., Bremerman, J., Cattoni, R., Negri, M., Turchi, M., Oard, DW., & Post, M. (2021). The multilingual tedx corpus for speech recognition and translation. CoRR abs/2102.01757. [arxiv:2102.01757](https://arxiv.org/abs/2102.01757)
- Sheshadri, A. K., Rao Vijjini, A., & Kharbanda, S. (2021). WER-BERT: Automatic WER estimation with BERT in a balanced ordinal classification paradigm. In *Proceedings of the 16th conference of the european chapter of the association for computational linguistics: Main volume, association for computational linguistics, online* (pp. 3661–3672). <https://aclanthology.org/2021.eacl-main.320>.
- Tanaka, T., Masumura, R., Ihori, M., Takashima, A., Orihashi, S., & Makishima, N. (2021). End-to-end rich transcription-style automatic speech recognition with semi-supervised learning. In *Proceedings of the Interspeech 2021* (pp. 4458–4462). <https://doi.org/10.21437/Interspeech.2021-1981>
- Wang, C., Pino, J., Wu, A., Gu, J. (2020a). CoVoST: A diverse multilingual speech-to-text translation corpus. In *Proceedings of the 12th language resources and evaluation conference, European language resources association* (pp. 4197–4203), Marseille, France. <https://www.aclweb.org/anthology/2020.lrec-1.517>
- Wang, C., Wu, A., Pino, J. (2020b). Covost 2: A massively multilingual speech-to-text translation corpus. 2007.10310.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Fun-towicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Drame, M., Lhoest, Q., Rush, A. (2020). Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: System demonstrations, association for computational linguistics, online* (pp. 38–45). <https://doi.org/10.18653/v1/2020.emnlp-demos.6>, <https://aclanthology.org/2020.emnlp-demos.6>
- Zanon Boito, M., Havar, W., Garnerin, M., Le Ferrand, E., & Besacier, L. (2020). Mass: A large and clean multilingual corpus of sentence-aligned spoken utterances extracted from the bible. In *Proceedings of The 12th language resources and evaluation conference, European Language Resources Association* (pp. 6486–6493), Marseille, France. <https://www.aclweb.org/anthology/2020.lrec-1.799>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## Authors and Affiliations

Arnaldo Candido Junior<sup>4</sup>  · Edresson Casanova<sup>2</sup>  ·  
Anderson Soares<sup>3</sup>  · Frederico Santos de Oliveira<sup>3</sup>  · Lucas Oliveira<sup>1</sup> ·  
Ricardo Corso Fernandes Junior<sup>1</sup> · Daniel Peixoto Pinto da Silva<sup>1</sup> ·  
Fernando Gorgulho Fayet<sup>2</sup> · Bruno Baldissera Carlotto<sup>2</sup> ·  
Lucas Rafael Stefanel Gris<sup>1</sup> · Sandra Maria Aluísio<sup>2</sup> 

Edresson Casanova  
edresson@usp.br



Anderson Soares  
andersonsoares@ufg.br

Frederico Santos de Oliveira  
fred.santos.oliveira@gmail.com

Lucas Oliveira  
lucasoliveira.2018@alunos.utfpr.edu.br

Ricardo Corso Fernandes Junior  
ricardocorsojr@gmail.com

Daniel Peixoto Pinto da Silva  
peixoto.daniel@hotmail.com

Fernando Gorgulho Fayet  
fer.fayet@usp.br

Bruno Baldissera Carlotto  
bruno.baldissera@usp.br

Lucas Rafael Stefanel Gris  
gris@alunos.utfpr.edu.br

Sandra Maria Aluísio  
sandra@icmc.usp.br

- <sup>1</sup> Federal University of Technology — Paraná (UTFPR), Medianeira, Brazil
- <sup>2</sup> Instituto de Ciências Matemáticas e de Computação - University of São Paulo, São Carlos, Brazil
- <sup>3</sup> Federal University of Goiás, Goiânia, Brazil
- <sup>4</sup> São Paulo State University, São José do Rio Preto, Brazil