



PUC

ISSN 0103-9741

Monografias em Ciência da Computação

nº 05/2021

**Building a Noisy Audio Dataset to Evaluate
Machine Learning Approaches for Automatic
Speech Recognition Systems**

Julio Cesar Duarte

Sérgio Colcher

Departamento de Informática

PONTIFÍCIA UNIVERSIDADE CATÓLICA DO RIO DE JANEIRO

RUA MARQUÊS DE SÃO VICENTE, 225 - CEP 22451-900

RIO DE JANEIRO - BRASIL

Building a Noisy Audio Dataset to Evaluate Machine Learning Approaches for Automatic Speech Recognition Systems

Julio Cesar Duarte and Sérgio Colcher

jduarte@inf.puc-rio.br, colcher@inf.puc-rio.br

Abstract. Automatic speech recognition systems are part of people's daily lives, embedded in personal assistants and mobile phones, helping as a facilitator for human-machine interaction while allowing access to information in a practically intuitive way. Such systems are usually implemented using machine learning techniques, especially with deep neural networks. Even with its high performance in the task of transcribing text from speech, few works address the issue of its recognition in noisy environments and, usually, the datasets used do not contain noisy audio examples, while only mitigating this issue using data augmentation techniques. This work aims to present the process of building a dataset of noisy audios, in a specific case of degenerated audios due to interference, commonly present in radio transmissions. Additionally, we present initial results of a classifier that uses such data for evaluation, indicating the benefits of using this dataset in the recognizer's training process. Such recognizer achieves an average result of 0.4116 in terms of character error rate in the noisy set (SNR = 30).

Keywords: Machine Learning, Datasets, Automatic Speech Recognition Systems, Noisy Audio.

Resumo. Sistemas de reconhecimento automático de fala fazem parte da vida cotidiana das pessoas, embutidos em assistentes pessoais e telefones celulares, auxiliando como um facilitador para a interação homem-máquina ao mesmo tempo que permitem o acesso à informação de forma praticamente intuitiva. Tais sistemas são normalmente implementados utilizando técnicas de aprendizado de máquina, especialmente com redes neurais profundas. Mesmo com o seu alto desempenho na tarefa de transcrever o texto da fala, poucos trabalhos levam em consideração a questão do seu reconhecimento em ambientes ruidosos e, normalmente, as bases utilizadas não apresentam exemplos de áudios com ruídos, apenas mitigando tal questão, por intermédio de técnicas de aumento de dados. Este trabalho tem por objetivo apresentar o processo de construção de uma base de dados de áudios ruidosos, em um caso específico de áudios degenerados em função de interferências, comumente presentes em transmissões rádio. Adicionalmente, apresentamos resultados iniciais de um classificador que utiliza tais dados para avaliação, indicando os benefícios de se utilizar este conjunto já no processo de treinamento de um reconhecedor. Tal reconhecedor alcança um resultado médio de 0.4116 em termos de taxa de erro de caracteres no conjunto ruidoso (SNR = 30).

Palavras-chave: Aprendizado de Máquina, Conjuntos de Dados, Sistemas de Reconhecimento Automático de voz, Áudios Ruídosos. bigskip

In charge of publications:

PUC-Rio Departamento de Informática - Publicações

Rua Marquês de São Vicente, 225 - Gávea

22451-900 Rio de Janeiro RJ Brasil

Tel. +55 21 3527-1516 Fax: +55 21 3527-1530

E-mail: publicar@inf.puc-rio.br

Web site: <http://bib-di.inf.puc-rio.br/techreports/>

Table of Contents

1	Introduction	1
2	Related Work	2
2.1	Portuguese datasets for speech recognition applications	2
2.2	Speech recognition with noisy audio	3
3	Building the Dataset	4
3.1	1st Scenario - White Noise	5
3.2	2nd Scenario - Noise Collected from Live Antennas	6
3.3	3rd Scenario - Software Simulation Noise	8
3.4	4th Scenario - Hardware Simulation Noise	10
3.5	Final considerations in building the dataset	11
4	Dataset Validation	11
5	Conclusions and Future Work	14
	References	15

1 Introduction

Automatic Speech Recognition (ASR) systems are composed of tools and techniques that allow machines to understand the human voice through speech, allowing them to mimic a common human interaction called conversation. It is a field of computer science that has been researched for decades and consists mainly of deriving a speech transcript of an audio waveform from an usually digitized signal [14].

Figure 1 shows the outline of the application of an ASR. Usually a microphone, or an array of them, is used to capture speech, digitizing its contents in a wave file, that can be fed to the software that will process it. Finally, this software's output consists of the speech's transcription, that may be stored in a text file.

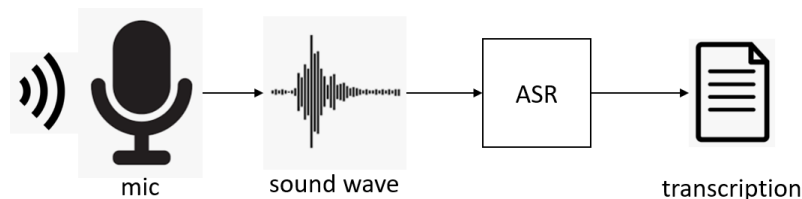


Figure 1: ASR application pipeline

An ASR module, in any system, is very convenient for data entry, access to remote information, and interactive services such as *chatbots*. ASR modules also help handicapped people, allowing or helping them interact with systems in a way that could be hard or nearly impossible for them [16].

Even with such practical and noble objectives, and having achieved high performance in their task in normal controlled environments, reliably transcription systems in realistic environments are still a challenge and their research and product development have been progressing for more than 30 years [14].

As an important part within the concept of Natural Language Processing [2], this task is now embedded in several equipment used daily, such as mobile phones and personal assistants and even with the excellent performance on the task obtained by the machine learning derived classifiers usually being used, a problem persists, speech recognition in noisy environments. Normally, such classifiers use datasets with noise-free audio that are built with good quality microphones in controlled ambiances. This reality may not be found in real-world noisy environments, such as, for example, when engines are present or in radio transmissions.

Current ASR systems often rely on online services such as Google Cloud, Microsoft Azure, and IBM Watson Speech To Text APIs [32]. They typically perform well in controlled environments and can be constantly updated. However, as they are online services, they depend on a stable Internet connection, in addition to having an operating cost.

There are also other offline solutions, such as CMU Sphinx [17], which are limitedly updated and with less support for languages other than English. Among the *generic* solutions that allow the development of an ASR for any language, DeepSpeech [13] stands out, by allowing, through the training of deep neural networks implemented using the Tensorflow framework, the development of an ASR for any language by just feeding the tool with a dataset in the desired language, in addition to adjusting its hyper-parameters.

Deepspeech has even proven performance close to the state of the art for both English and Mandarin [4].

Although ASRs prove to be a very important task in NLP, most of the developments are based on the English language, although some recent works are investigating the performance of such techniques for other languages, such as Portuguese [10, 11, 27].

Likewise, ASRs can be used in a military environment, incorporated in a solution for audio transcription on a military radio, which usually suffers from a lot of interference. In this context, there is the Software-Defined Radio Project [7], which today is Brazilian Army’s priority 1.1 [9], where radio families with such technology could automatically identify the audio that is received (or retransmitted), presenting the corresponding text on its display, in addition to storing the information for future utilization. A speech recognition technology that enables its use in noisy environments, such as a high frequency (HF) communication channel, commonly used in the Brazilian theater of operations, would be, then, extremely useful.

The main objective of this work is to present the methodological process of building a noisy audio dataset that allows the evaluation of ASRs derived from machine learning algorithms. As a case study, radio transmission paths on HF channels will be used, especially, but not limited, for use in military operations, with the Portuguese language. As an additional contribution, we will also present initial experiments with the generated dataset that illustrate the importance of adapting the ASRs to the real environment in which they are used.

In the remainder of this article, we will present additional steps toward the construction of this dataset. First, in Section 2, we show related work, focusing on Portuguese audio datasets, as well as works that deal with ASRs in noisy environments. Section 3 details the construction process showing the assumptions we took, as well as the simulation tools used to represent the chosen environment. We, then, conclude this article in Section 4 with our experimental results that validates the dataset that is followed by our final remarks and possible future work in Section 5.

2 Related Work

ASRs are a hot topic for research in computer science and many papers can be reviewed in order to help build custom systems, but only a small amount deals with noise in audio. Conversely, few works are related to the Portuguese language, mainly because, but not limited to, the limited availability of voice speech data with their validated transcriptions [10]. If we consider both restrictions, the Portuguese idiom and noisy data, this limits even more the comparison with similar articles. Nevertheless, this section summarizes the most related works found in our research.

2.1 Portuguese datasets for speech recognition applications

The commonly called dataset CETUC [3] was the first initiative found for a Portuguese-speaking dataset derived from the CETEN-folha text corpus. It consists of 145 hours of noise-free speech text from 100 speakers repeating once the same 1,000 sentences, or 3,528 words. These phrases were chosen from CETEN-folha, having phonetic balance as their main characteristic [8]. All recordings were made in a studio where the audio was captured

at a sampling rate of 16 kHz and 16 bits per sample.

LibriSpeech [24] is a multilingual dataset benchmark for speech related tasks. It comprises of about 50.5K hours of speech spread in eight languages. The Portuguese subset of Librispeech contains approximately 170 hours of voice from 62 different speakers. The dataset was derived from several audiobooks downloaded from the Internet.

The Common Voice Dataset [5] is a multilingual open-source project to build voice datasets for basically any language. Anyone, through the project’s site¹, can record or listen to and evaluate a sentence, in order to contribute to the project. Its current version 7.0 supports more than 75 different languages. The Portuguese subset comprises 3GB of over 84 hours of validated speech and approximately 2,000 different speakers. This work uses Common Voice as its base dataset because of two major factors. First, this is an ongoing project that everyone can contribute to, and as a result, the project is expected to have good longevity. Second, this is the project officially supported by the DeepSpeech tool (Section 4), also used for our ASR experiments.

Table 1 summarizes the main characteristics of the datasets depicted here. While these datasets can be used to build well-performing ASRs, they are nowhere near the size when compared to similar datasets in English. Furthermore, they are built with audio files that were recorded in a controlled environment and do not represent possible noisy environments in which the systems can be employed.

Table 1: Portuguese voice datasets main characteristics

Dataset	Total Hours	Speakers	Audio Source
CETUC [3]	145	100	studio recording
LibriSpeech [24]	170	62	audiobooks
Common Voice [5]	84	2,000	on-line recording

2.2 Speech recognition with noisy audio

Oliveira Santos[22] presents a comparison between three techniques for speech recognition in noisy environments applied to a limited set of Portuguese words corresponding to the digits of the alphabet. Three types of noise were applied: white noise, chitchat and background industry noise. Of the techniques that were used, the one that applies the Maximum Likelihood Linear Regression (MLLR) obtained the best results in conjunction with Hidden Markov Models (HMM).

Menêses Santos[19] analyzes the use of a hybrid model of Convolutional Neural Networks (CNN) and HMM in conjunction with noisy audios for speaker-dependent speech recognition. The experiments were performed on audio datasets that contain utterances and digits. The results presented with noises that represent chitchat, engine and industry noises, using a signal-to-noise ratio (SNR) of 6db, showed that the use of this hybrid model improves the quality of the system even without the explicit incorporation of a noise suppression technique.

Prodeus et al.[25, 26] show that the use of noisy audio samples during the training phase of an ASR can improve its performances. In their work, several types of noise were used, such as those coming from grinders, computers, trucks, among others, in a range of

¹<https://commonvoice.mozilla.org/>

SNR from 0 to 40 db. Different techniques were used where the signal-to-noise ratios and the types of noise were combined in the training and test (recognition) sets.

Wang et al.[31] proposes the use of Deep Neural Networks (DNNs) as a way to improve the performance of an ASR that was already trained using only noiseless audios. The system was evaluated using a dataset of 21,800 English utterances with added noise ranging from a SNR of 5 to 20 db. By applying the technique in conjunction with an ASR trained using a 6-layer context-dependent (CD)-DNN-HMM hybrid model with 2048 nodes per layer, a word error rate of 6.56% was obtained from a system that already performed at 17.89%.

Ribeiro[29] investigates the use of neural networks in the development of ASRs using a throat microphone, which is a more suitable device for capturing audio in noisy environments. The system uses audio samples of simple commands (“OK” and “Cancel”) and digits for the Portuguese Language, which were noise filtered prior to being applied to a classification system composed by an ensemble of three classifiers: Multilayer Perceptron, Binary Multilayer Perceptron and Self-Organizing Maps. The performance of 96.6% in the hit rate indicates that the use of the device can help with voice recognition in noisy environments.

Maruf et al.[18] analyzes the effects of noise in ASRs using two feature extraction techniques called Mel Frequency Cepstral Coefficients (MFCC) and Relative Spectral Transform-Perceptual Linear Prediction (RASTA-PLP). The proposed CNN-derived ASR was trained using digits and 11 utterances of the Bangla language from 120 different speakers. Urban noises were added from different environments, like libraries, street and rail stations. The best results of 93.18% in terms of accuracy were found using the RASTA-PLP model.

Finally, Pervaiz et al.[23] investigates the use of data augmentation to improve noise robustness of ASRs. It uses several public utterance datasets to assess the quality of different machine learning approaches. The approaches tested were traditional deep neural networks such as Deep Feed-Forward Networks, Long Term Memory Networks (LSTMs) and CNNs that were evaluated on a noise dataset based on an Asian accent.

In summary, the effects of noise on ASRs are a common research topic. However, few works deal with all aspects of ASRs, such as any type of transcription, and none propose a methodology for building datasets to be used for the purpose of evaluating machine learning techniques. Furthermore, only a few focus on the Portuguese language. In this sense, this work intends to fill this gap by proposing a process of building noisy audio datasets to be used in training ASRs. In addition, a specific use case for the Portuguese language is implemented, where we present some results of proposals inspired by the related works reported here.

3 Building the Dataset

In this section, we present the steps taken in our proposal to build a noisy audio dataset. This process is generic and can be customized depending on the type of noise you want to model in the ASR. First, it is necessary to start with the relationship between noise and the original signal that is used, since it plays an important role in the dataset construction process.

In our methodology, we can vary the signal-to-noise ratio (SNR) when combining or

simulating the chosen environment with a noiseless audio dataset. SNR is a metric for comparing signal power to noise and provides an important indication of the degree to which the the signal was contaminated with additive noise [6]. if we want to mathematically formalize, SNR is the quotient (ratio) between the mean power of a signal and the power of the background noise. Equation 1 illustrates the mathematical formula to evaluate the SNR of any signal.

$$SNR = 20 \cdot \log \frac{RMS(s)}{RMS(n)} \quad (1)$$

where RMS is a function that evaluates the square root of the averages of data being analyzed, the signal (s) or noise (n).

When noise has an approximately Gaussian distribution, it is called an Additive White Gaussian Noise (AWGN) and it is the most commonly assumed model [6].

In our dataset generation approaches, we use the following SNRs as input: {-30, -20, -10, -5, 0, 5, 10, 20, 30}. The positive sub-interval of this range was chosen because it is a set of values widely used in similar works. However, we made sure to also include the same negative sub-interval, in which the noise is stronger than the signal itself. Such range can be interesting in works that perform audio pre-processing for noise extraction, although they are not very useful in machine learning approaches, since, when listening to an audio from these subsets, even a person is not able to perform its transcription.

In order to build a speech corpus with noisy audio, we first need a noise-free corpus. In our process, we use the Common Voice 6.1 dataset [5] (Section 2.1). The main reason for this choice is the fact that it is a dataset with active and collaborative development, keeping the release of new versions every 6 months, approximately. Furthermore, this set is officially supported by the DeepSpeech platform, used in our experiments. More details of DeepSpeech will be explained in Section 4.

Our case study of interest is voice transmission over a military analog radio transmission channel. In Brazil, one of the most used transmission bands is HF, due to its wide use in areas such as the Amazon. Thus, scenarios are created that allow the simulation of this channel in an audio (voice) transmission. We created a total of four scenarios to run this simulation, each with its own degree of reliability. These four scenarios, which can be replicated for different case studies, are explained in the next subsections.

3.1 1st Scenario - White Noise

Our first scenario is a simple AWGN model where we have a randomly generated signal with equal strength at different frequencies with a flat spectral density. This is the simplest model, since it does not emulate a particular noise, but can be used as a generic approach for modeling noise in any channel.

Adding this kind of noise to a signal is simple and can be done with the Python snippet presented in Listing 3.1.

Listing 3.1: An excerpt of a python code that adds white noise to a signal

```
def addWhiteNoise(signal, SNR):
    RMSs = sqrt(mean(signal**2))
    RMSn = sqrt(RMSs**2/(pow(10, SNR/10)))
    noise = random.normal(0, RMSn, signal.shape[0])
    return signal + noise
```

Note that the SNR parameter influences noise degradation, since the lower the value (higher in terms of negative value), the more degraded the signal will be. This can be better observed in Figure 3 that depicts the application of different SNR values in the sample wav file shown in Figure 2.




Figure 2: Sample .wav file without noise

Observe that the higher the SNR value (Figure 3i), the more similar are the obtained signal and the original one.

Finally, we show in Figure 4 all generated files, in order to compare them within the same scale.

3.2 2nd Scenario - Noise Collected from Live Antennas

Although AWGN is the most used scenario to model generic noise, it does not replace the use of noise found in a communication channel. Thus, models trained with white noise may not perform well when applied in the real world. Using a more reliable representation of the noise in the environment, in this case, leads to ASRs with a better performance.

An easy way to generate audio with more representative noises is to collect this noise directly on the communication channel that the system is used. Such captured noise can then be merged with the original audio files, generating a dataset closer to the reality of applying the recognizer.

In order to accomplish this idea, we used a remote HD radio service provided by Halász [12]. Through this site, we control a multi-user SDR receiver software with a web interface, called OpenWebRX² that uses two Kiwi SDRs connected to dipoles (40m and 80m) and vertical (20m) antennas. The site of these antennas is located in Ilha dos Ratos, Paraty, Rio de Janeiro. Figure 5 illustrates both the radio control interface (5a) and the antenna's site (5b).

While using this setup, we sampled receptions on the following frequencies / bands: 500 KHz / AM, 1 MHz / AM, 1,836 MHz / USB, 8,060 MH / CW and 7,794 MHz / USB.

²<https://github.com/ha7ilm/openwebrx>

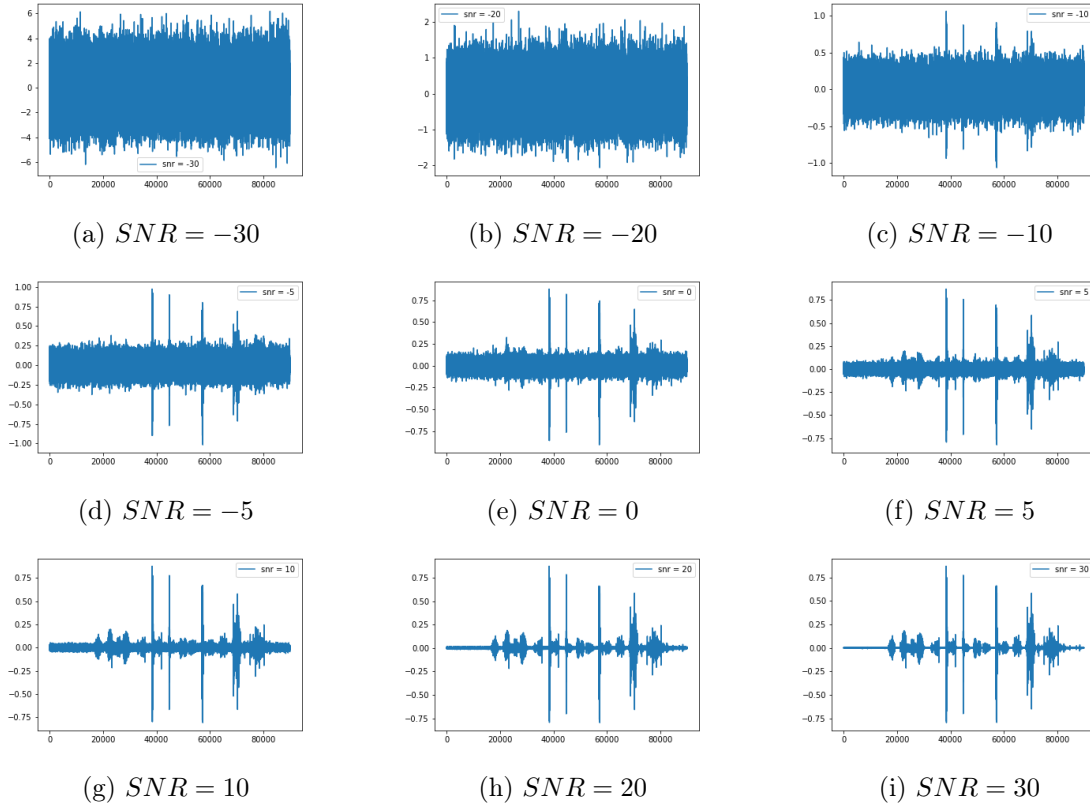


Figure 3: AWGN applied to a sample .wav file with different SNR values

These frequencies were chosen since they were active during the collection period, although they were recorded when no transmission was detected.

Those recording samples form a pool of noise sources which can then be applied to any audio file by, randomly, choosing an excerpt of this noise pool. A Python snippet of these method is presented in Listing 3.2.

Listing 3.2: An excerpt of a python code that adds a random cut of noise stored in a file to a signal

```
def addNoise(signal, noiseSource, SNR):
    signal = interp(signal, (signal.min(), signal.max()),
                    (-1, 1))
    startCut = randrange(0, len(noiseSource) - len(signal))
    noise = noiseSource[startCut:startCut + len(signal)]
    RMSs = sqrt(np.mean(signal**2))
    RMSn = sqrt(RMSs**2/(pow(10, SNR/10)))
    noise = noise*(RMSn/sqrt(mean(noise**2)))
    signal_n = signal + noise
```

For instance, we show the application of this method in Figure 6, with a SNR value of 10. The same sample .wav file depicted in Figure 2 was also used as input.


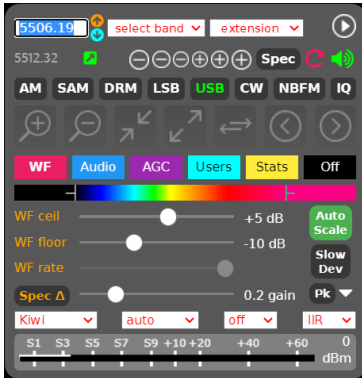


Figure 4: AWGN applied to a sample .wav file with all SNR values



(a) Radio Web Interface



(b) Antenna's site. Source: [12]

Figure 5: Equipment used to sample noise files




Figure 6: Sample .wav with added noise from samples (SNR = 10)

3.3 3rd Scenario - Software Simulation Noise

Using collected files is a good strategy to represent a communication channel when you do not have the implementation of a mathematical model for the communication channel, however, such files can limit the possibilities of noise used in the training of a speech recognizer. Thus, when using a software that implements such models, the variability of the nuances of the noise being analyzed is increased.

PathSim is a software that simulates HF radio propagation paths using an audio file or the computer's sound card [20]. Three HF paths are available with different spreading frequencies parameters and frequency offset functions. The spreading uses a complex

AWGN signal to achieve the desired SNR which is filtered by a 3KHz low pass filter to limit the noise. Figure 7 shows PathSim’s interface, while the sound card input was injected with a 1 KHz sine wave test tone.

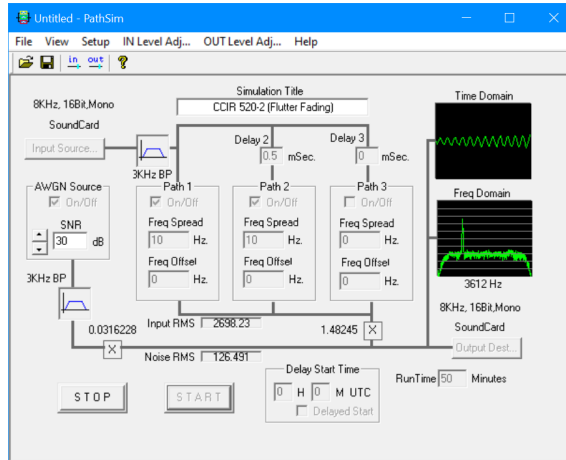


Figure 7: PathSim Interface

Initially, the intention was to process all audio files directly in PathSim, however some limitations prevented this approach. First, the software does not provide a programmable API, so it was necessary to manually feed the files into its interface. To get around the large amount of files needed to be processed, they were merged into files of maximum 500 minutes, PathSim’s limit. However, sometimes the software simply crashed when processing these big .wav files. It is important to note that this is older software from the 2000s, without active development, running on a new operating system. As a result, smaller files were generated and manually fed into the software. With these files, initial tests were performed on the recognizer, which detected problems in the generated .wav files. Some of these files were corrupt and some of which produced incompatible results. Finally, the software has the limitation of working only with 8 KHz .wav files, which generates files with worse quality and that need to be resampled after processing.

In order to circumvent these problems, a different approach was adopted, similar to that of Scenario 2 (Section 3.2). Several noise files were generated using the software and, then, were merged to the noise-free dataset using the same code in Listing 3.2.

In our study case, we use the *CCIR-Flutter* preset, which presented the most degenerative noises, following the International Radio Consultative Committee - CCIR 520-2 [15]. In this configuration, two paths are used with the same attenuation, differential time delay of 0.5 ms and a Doppler spread of 10.0 Hz with zero Doppler shift. Other setups can also be used.

With this approach, we maintain the advantage of using a custom mathematical model for the communication channel, even with its software limitations. Figure 8 shows the application of this approach, with a SNR value of 10, when merged to the sample .wav file represented in Figure 2.

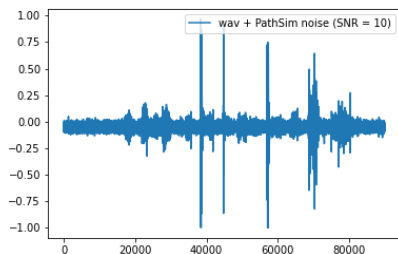
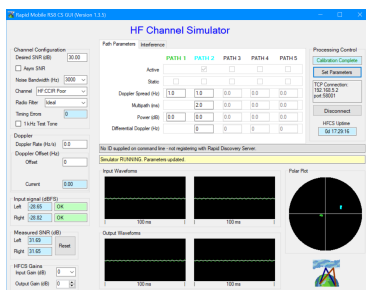


Figure 8: Sample .wav with added noise from PathSim (SNR = 10)

3.4 4th Scenario - Hardware Simulation Noise

Although the strategy used in Section 3.3 allowed the generation of audio files with customized noise through a mathematical model of the communication channel, we also choose to use a more robust solution, embedded in a dedicated hardware. A dedicated hardware unit is usually built for a single purpose and, therefore, excels at the task for which it was manufactured.

With that idea in mind, RapidM RS8 was used. The RS8 is a device that allows the simulation of HF and V/UHF channels in order to evaluate modems and waveforms at baseband which is in compliance with military standards like MIL-STD-188-110B [28]. Figure 9 presents the software and hardware interface of the equipment.



(a) Application Interface



(b) RS8 device. Source: [28]

Figure 9: RS8 Hardware and Software Interfaces

Using a device such as RS8 allows us to achieve the state-of-the-art in terms of simulation of communication channels, thus representing it more reliably. However, a disadvantage here is present. Unlike the software solutions presented in Sections 3.1, 3.2 and 3.3, all processing is performed on an external equipment, requiring the audio to be sent to the computer's sound card output, which is connected to the device through a customized cable, processed by the equipment, and, finally, captured in one of the inputs of the same sound card. All this takes place in real time, that is, for every second of audio in the dataset, one second of processing is required and the process is repeated for all the signal-to-noise ratios used.

Similar to Section 3.3, and to give some noise variance, we configured the device with the *CCIR-Poor* preset, also recommended in CCIR 520-2 [15], which has two independently fading paths, but with a differential time delay of 2 ms and a Doppler frequency spread of

1.0. As with PathSim, other setups may be used.

Using this equipment, we obtain .wav files like the one depicted in Figure 10, using a SNR value of 10. It is noteworthy that no file merge is done in this case, as all processing is done in real time.

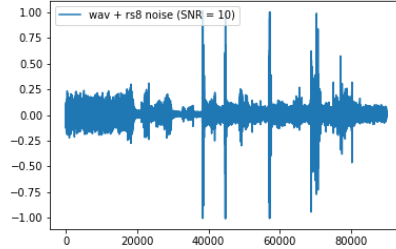


Figure 10: Sample .wav with added noise from RS8 (SNR = 10)

3.5 Final considerations in building the dataset

In this section, we present four ways to generate a noisy audio dataset. This 4-step methodology is generic and can be applied to any kind of noise. In our case study, we specialize it for noises generated in HF communication channels. These semi-synthetic datasets can be used in the training phases of machine learning algorithms, either alone, if they well reflect the noise being modelled, or together, feeding the algorithms with more examples, as in a specialized data augmentation scheme.

All script and noise files used are available in the project’s GitHub repository³ allowing the replication of the experiments, with the exception of the RS8 output files, which can be provided on request.

4 Dataset Validation

In order to evaluate the dataset in terms of its application potential for training with machine learning algorithms, we conducted some experiments on the construction of an ASR for Portuguese using the four approaches presented in Section 3.

In our experiments, we use Mozilla DeepSpeech [13, 21], which is an open source speech-to-text framework that can use models trained by machine learning algorithms implemented within Google’s TensorFlow [1]. DeepSpeech has a flexible architecture where you can try different hyperparameters to build ASRs for any specific language. It also supports the Common Voice Dataset, which makes our noisy datasets easier to use.

Deespech uses a Recurrent Neural Network architecture like the one depicted in Figure 11. It consists of 5 hidden layers, where the first three and the fifth are composed of regular neurons and use a clipped rectified-linear (ReLU) activation function. The fourth layer is a recurring LSTM layer that includes a set of hidden units with forward recurrence. Finally, the output layer uses the softmax function which outputs the probabilities of the characters of the alphabet.

³<https://github.com/duartejulio/noisy-asr>

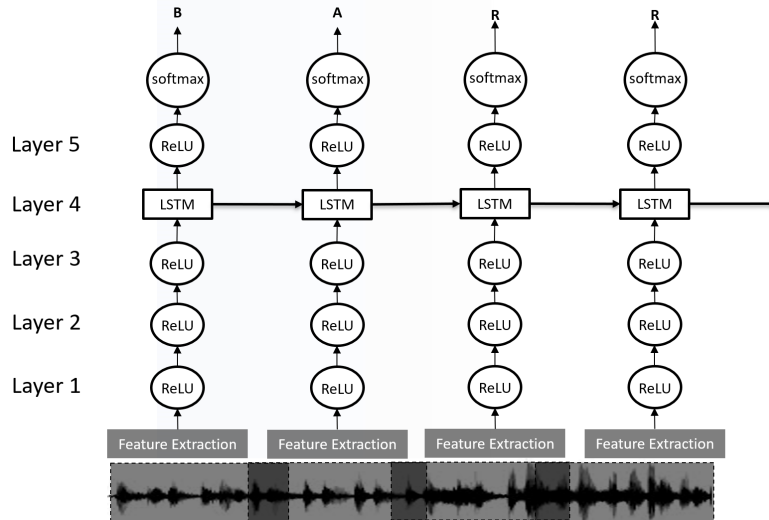


Figure 11: DeepSpeech basic architecture. Adapted from [21]

For our experiments, we use the following hyperparameters in order to train the neural networks: $epochs = 100$, $batch_size = 16$, $number_of_neurons_per_layer = 2048$, $feature_extraction_audio_window_length = 32ms$, $learning_rate = 0.0001$, and $dropout_rate = 0.40$. All other hyperparameters were set to the default values.

With this hyperparameter setup, we conducted two experiments. The first one uses only the train and dev subsets of Common Voice combined for training the neural networks. The dev subset is usually used for validation purposes, but in order to have a bigger training set, we opted to use it in conjunction to the train subset. Although Common Voice has about 50 hours of validated voice, its subsets are built in a way to ensure that the same speakers do not appear in different sets, nor that the same sentences appear repeated in the subsets. This considerably reduces the availability of data for training so, by combining the train and dev subsets, we get a total of 11,086 training examples.

Figure 12 shows the results of this model in the same test set but applied to all four approaches presented in Section 3 and all considered SNR values. In this graph, we use character error rate (CER) as our main metric for evaluation.

First, we can see that, for the negative SNR values, it is very hard for the models to learn anything. The ability to generalize of the models is very low, with erratic results. Normally, such results indicate a massive classification only of the most frequent letter of the alphabet ('a'). Next, it is possible to see the difference in terms of CER among the different approaches when comparing with the same SNR. Finally, it seems that the approach with hardware simulation noise is the hardest to recognize.

We show here, only the results in terms of CER, but the ones considering WER have similar behavior. We do not consider this comparison because evaluations with WER usually contemplate the use of a language model, in the output of the ASR, after the classification process. A language model assigns probabilities to translation hypotheses in the target language [30], and, in this context, it is used to correct transcriptions both at character and word level. As we would like to observe the behavior of different forms of noise in the final transcripts of the text, we chose to disable the external scorer, in the output of DeepSpeech, that comprises the language model being used, as such model could




Figure 12: CER results for the model trained with the train+dev sets

mask these results.

In our second experiment, we used an expanded version of the train+dev combined subset. As previously reported, Common Voice has over 50 hours of validated text, but only a small portion is used on those sets. We propose to use all versions of the same sentences spoken by different speakers in the same subset. As a result, we add all repeated sentences from the validated set, as long as they were spoken by a speaker from the same subset. In this way, it is possible to expand the combined set without such subsets having examples with sentences or speakers from the test subset. The final expanded train+dev combined set has a total of 28,068 training examples, more than doubling the amount of examples available for training.

Figure 13 shows the results with this new model. It is important to note that the test subset has not changed from previous results.




Figure 13: CER results for the model trained with the expanded train+dev sets

Analyzing the results of Figure 13, while comparing with the same results of Figure 12, we observe that the new model has a better overall performance. Note, nonetheless, that the results for the 4th approach, the one that uses hardware simulation noise, are still the worst, indicating that using a noisy audio dataset with those characteristics can improve such results.

Finally, it is not the intention of this work to generate an automatic speech recognizer for Portuguese, but rather to promote the discussion of the use of noisy audio data in its construction. However, it is interesting to show the possibilities of applying noisy audio datasets in a state-of-the-art speech recognizer.

Quintanilha et al.[27] presents an ASR derived from DeepSpeech with a performance of 10.49% and 25.45%, in terms of CER and WER, respectively. Also, Gris et al.[11] presents an ASR with a WER of 12.79%. The models used, however, are much more advanced, trained on larger datasets and use other strategies, such as language models, differently from this work. Nonetheless, such works can be improved using the noise generation strategy proposed, especially if applied in environments with frequent noise. Furthermore, as the proposal presented here is language-independent, it is possible to apply it in other languages under the same conditions.

5 Conclusions and Future Work

Automatic speech recognition systems are currently part of our lives embedded in several hardware and software. Such systems are generally derived from machine learning algorithms that require a large amount of data representative of the environment in which they are applied, in this case, speech samples with their respective transcriptions. Often, such systems can be used in environments with constant and characteristic noise, such as noise from urban areas, industries, or even radio communication channels.

The main objective of this work is to propose a methodology for the generation of noisy audio datasets to be used in conjunction with machine learning algorithms in order to make them more robust for application in noisy environments. In this sense, four approaches are proposed, each with distinct characteristics and representing differently the noise being added to the data. Depending on the type of noise to be modeled and the availability of the needed tools, one or more approaches can be used together in order to provide more examples for the algorithm responsible for training the ASR, working as a specialized data augmentation strategy. The evaluation of such sets in straightforward ASRs implemented points out that their use in the training phase can improve their performance, especially when applying the ASR in an noisy environment.

The main contribution of this work, in addition to the proposed methodology, is the generated dataset, as well as all the codes available for reproducing the construction of the dataset and applying it in a new context. The codes can be downloaded from <https://github.com/duartejulio/noisy-asr>, as well as the noise samples collected during this work.

Even with the fulfillment of the work objectives, future works are always foreseen in order to improve or expand the results presented here. With this purpose in mind, We also intend to use the noisy dataset in the ASR training phase, thus allowing a better evaluation of its behavior in the presence or not of the characterized noise. In addition, we also intend to carry out tests with other types of noise comparing the performance of the proposed methodology with traditional and generic data augmentation schemes. Finally, studies of how ASRs can be embedded in more compact hardware are also planned as future works, since such systems usually require a lot of computational power and the devices in which they can be installed do not always support such processing and memory requirements.

Acknowledgments

This material is based upon work supported by Air Force Office Scientific Research under award number FA9550-19-1-0020.

We would also like to thank *Centro Tecnológico do Exército* (CTEx) and *Indústria de Material Bélico do Brasil* (IMBEL) for providing the means for us to use the RS8 equipment.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] Steven P. Abney. Parsing by chunks. In Robert C. Berwick, Steven P. Abney, and Carol Tenny, editors, *Principle-Based Parsing: Computation and Psycholinguistics*, pages 257–278. Kluwer, Dordrecht, 1991.
- [3] Vladimir Alencar and Abraham Alcaim. LSF and LPC - Derived features for large vocabulary distributed continuous speech recognition in brazilian portuguese. In *Proceedings of the 42nd Asilomar Conference on Signals, Systems and Computers*, pages 1237–1241, Pacific Grove, California, USA, October 2008. IEEE.
- [4] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, et al. Deep speech 2 : End-to-end speech recognition in english and mandarin. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 173–182, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [5] Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis Tyers, and Gregor Weber. Common voice: A massively-multilingual speech corpus. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4218–4222, Marseille, France, May 2020. European Language Resources Association.
- [6] A. Bruce Carlson, Paul B. Crilly, and Janet C. Rutledge. *Communication systems: An introduction to signals and noise in electrical communication*. Boston: McGraw-Hill, 4th edition, 2002.
- [7] Centro Tecnológico do Exército. Rádio definido por software de defesa (RDS-DEFESA). <http://www.ctex.eb.mil.br/projetos-em-andamento/84-radio-definido-por-software-rds>, 2020. Accessed: 2021-09-02.

- [8] Ricardo J. R. Cirigliano, Clarisse Monteiro, Filipe Leandro Barbosa, Fernando Gil Vianna Resende Junior, Letícia Rebolo Couto, and João A. de Moraes. Um conjunto de 1000 frases foneticamente balanceadas para o português brasileiro obtido utilizando a abordagem de algoritmos genéticos. In *Proceedings of the 22nd Simpósio Brasileiro de Telecomunicações – SBrT*, 2005.
- [9] Exército Brasileiro. Plano estratégico do exército brasileiro 2020-2023. http://www.ceadex.eb.mil.br/images/legislacao/XI/plano_estrategico_do_exercito_2020-2023.pdf, 2019. Accessed: 2021-09-02.
- [10] Lucas Rafael Gris, Edresson Casanova, Frederico Oliveira, Anderson Soares, and Arnaldo Candido-Junior. Desenvolvimento de um modelo de reconhecimento de voz para o português brasileiro com poucos dados utilizando o wav2vec 2.0. In *Anais do XV Brazilian e-Science Workshop*, pages 129–136, Porto Alegre, RS, Brasil, 2021. SBC.
- [11] Lucas Rafael Stefanel Gris, Edresson Casanova, Frederico Santos de Oliveira, Anderson da Silva Soares, and Arnaldo Candido Junior. Brazilian portuguese speech recognition using wav2vec 2.0, 2021.
- [12] Bruce Halász. KiwiSDR by PY2BS and PY2PE. <http://py1eme.homeip.net:8073/>, 2021. Accessed: 2021-07-14.
- [13] Awni Y. Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng. Deep speech: Scaling up end-to-end speech recognition. *CoRR*, abs/1412.5567, 2014.
- [14] Martin G. Helander, Thomas K. Landauer, and Prasad V. Prabh. Chapter 14 - systems design for automated speech recognition. In *Handbook of Human-Computer Interaction*. Elsevier Science Inc., USA, 2nd edition, 1997.
- [15] International Telecommunication Union (ITU). *Recommendation F.520-2. Use of High Frequency Ionospheric Channel Simulators*, volume III. Recommendations and Reports of the CCIR, Genova, 1978-1982-1992.
- [16] Bing-hwang Juang and Sadaoki Furui. Automatic recognition and understanding of spoken language - a first step toward natural human-machine communication. *Proceedings of the IEEE*, 88:1142–1165, 2000.
- [17] Paul Lamere, Philip Kwok, Evandro Gouvêa, Bhiksha Raj, Rita Singh, William Walker, Manfred Warmuth, and Peter Wolf. The CMU SPHINX-4 speech recognition system, 2003.
- [18] Md. Raffael Maruf, Md. Omar Faruque, Salman Mahmood, Nazmun Nahar Nelima, Md. Golam Muhtasim, and Md. Jahedul Alam Pervez. Effects of noise on RASTA-PLP and MFCC based bangla ASR using CNN. In *2020 IEEE Region 10 Symposium (TENSYMP)*, pages 1564–1567, 2020.
- [19] Rafael Meneses Santos. *Uma abordagem híbrida CNN-HMM para reconhecimento de fala tolerante a ruídos de ambiente*. Dissertation, Universidade Federal de Sergipe, 2016.

- [20] Moetronix. PathSim – HF Path Simulator Project. <https://www.moetronix.com/ae4jy/pathsim.htm>, 2000. Accessed: 2021-07-19.
- [21] Mozilla Corporation. Welcome to DeepSpeech’s documentation! <https://deepspeech.readthedocs.io/en/r0.9/>, 2020. Accessed: 2021-07-19.
- [22] Debora Andrea de Oliveira Santos. *Reconhecimento de Voz em Presença de Ruído*. PhD thesis, Pontifícia Universidade Católica do Rio de Janeiro, 2001.
- [23] Ayesha Pervaiz, Fawad Hussain, Huma Israr, Muhammad Ali Tahir, Fawad Riasat Raja, Naveed Khan Baloch, Farruh Ishmanov, and Yousaf Bin Zikria. Incorporating noise robustness in speech command recognition by noise augmentation of training data. *Sensors*, 20(8), 2020.
- [24] Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. MLS: A large-scale multilingual dataset for speech research. *Interspeech 2020*, Oct 2020.
- [25] Arkadiy Prodeus and Kateryna Kukharicheva. Training of automatic speech recognition system on noised speech. In *2016 4th International Conference on Methods and Systems of Navigation and Motion Control (MSNMC)*, pages 221–223, 2016.
- [26] Arkadiy Prodeus and Kateryna Kukharicheva. Automatic speech recognition performance for training on noised speech. In *2017 2nd International Conference on Advanced Information and Communication Technologies (AICT)*, pages 71–74, 2017.
- [27] Igor Macedo Quintanilha, Sergio Lima Netto, and Luiz Pereira Biscainho. 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, Sep. 2020.
- [28] Rapid Mobile. RS8 HF & V/UHF channel simulator. <https://www.rapidm.com/product/rs8-wideband-hf-vuhf-channel-simulator>, 2019. Accessed: 2021-07-19.
- [29] Fábio Cisne Ribeiro. *Reconhecimento de comandos de voz em português brasileiro em ambientes ruidosos usando laringofone*. PhD thesis, Universidade Federal do Ceará, 2019.
- [30] David Talbot and Miles Osborne. Randomised language modelling for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 512–519, Prague, Czech Republic, June 2007. Association for Computational Linguistics.
- [31] Quandong Wang, Sicheng Wang, Fengpei Ge, Chang Woo Han, Jaewon Lee, Lianghao Guo, and Chin-Hui Lee. Two-stage enhancement of noisy and reverberant microphone array speech for automatic speech recognition systems trained with only clean speech. In *2018 11th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, pages 21–25, 2018.
- [32] Binbin Xu, Chongyang Tao, Zidu Feng, Youssef Raqui, and Sylvie Ranwez. A benchmarking on cloud based speech-to-text services for french speech and background noise effect. *CoRR*, abs/2105.03409, 2021.