





# AequeVox: Automated Fairness Testing of Speech Recognition Systems\*

Sai Sathiesh Rajan (✉) , Sakshi Udeshi , and Sudipta Chattopadhyay 

Singapore University of Technology and Design, Singapore 487372, Singapore  
{sai\_rajan, sakshi\_udeshi}@mymail.sutd.edu.sg  
sudipta\_chattopadhyay@sutd.edu.sg

**Abstract.** Automatic Speech Recognition (ASR) systems have become ubiquitous. They can be found in a variety of form factors and are increasingly important in our daily lives. As such, ensuring that these systems are equitable to different subgroups of the population is crucial. In this paper, we introduce, AEQUEVOX, an automated testing framework for evaluating the fairness of ASR systems. AEQUEVOX simulates different environments to assess the effectiveness of ASR systems for different populations. In addition, we investigate whether the chosen simulations are comprehensible to humans. We further propose a fault localization technique capable of identifying words that are not robust to these varying environments. Both components of AEQUEVOX are able to operate in the absence of ground truth data.

We evaluate AEQUEVOX on speech from four different datasets using three different commercial ASRs. Our experiments reveal that non-native English, female and Nigerian English speakers generate 109%, 528.5% and 156.9% more errors, on average than native English, male and UK Midlands speakers, respectively. Our user study also reveals that 82.9% of the simulations (employed through speech transformations) had a comprehensibility rating above seven (out of ten), with the lowest rating being 6.78. This further validates the fairness violations discovered by AEQUEVOX. Finally, we show that the non-robust words, as predicted by the fault localization technique embodied in AEQUEVOX, show 223.8% more errors than the predicted robust words across all ASRs.

## 1 Introduction

Automated speech recognition (ASR) systems have made great strides in a variety of application areas e.g. smart home devices, robotics and handheld devices, among others. The wide variety of applications have made ASR systems serve increasingly diverse groups of people. Consequently, it is crucial that such systems behave in a non-discriminatory fashion. This is particularly important because assistive technologies powered by ASR systems are often the primary mode of

---

\* This work is partially supported by Singapore Ministry of Education (MOE) grant number MOE2018-T2-1-098 and OneConnect Financial grant number RGOCFT2001.

$$ASR_{Err}(\text{male} + \text{noise}) \approx ASR_{Err}(\text{female} + \text{noise})$$

$$ASR_{Err}(\text{male} + \text{noise}) \ll ASR_{Err}(\text{female} + \text{noise})$$

Fig. 1: Fairness Testing in AEQUEVOX

interaction for users with certain disabilities [20]. Consequently, it is critical that an ASR system employed in such systems is effective in diverse environments and across a wide variety of speakers (e.g. male, female, native English speakers, non-native English speakers) since they are often deployed in safety-critical scenarios [18].

In this paper, we are broadly concerned with the *fairness properties in ASR systems*. Specifically, *we investigate whether speech from one group is more robustly recognised as compared to another group*. For instance, consider the example shown in Figure 1 for a system  $ASR$ . The metric  $ASR_{Err}$  captures the error rate induced by  $ASR$ . Consider speech from two groups of speakers i.e. *male* and *female*. We assume that the  $ASR$  has similar error rates for both the groups of speakers, as illustrated in the upper half of Figure 1. We now apply a small, constant perturbation on the speech provided by the two groups. Such a perturbation can be, for instance, addition of small noise, exemplifying the natural conditions that the ASR systems may need to work in (e.g. a noisy environment). If we observe that the  $ASR_{Err}$  increases disproportionately for one of the speaker groups, as compared to the other, then we consider such a behaviour a violation of fairness (see the second half of Figure 1). Intuitively, Figure 1 exemplifies the violations of *Equality of Outcomes* [36] in the context of ASR systems, where the male group is provided with a higher quality of service in a noisy environment as compared to the female group. Automatically discovering such scenarios of unfairness via simulating the ASR service in diverse environments is the main contribution of our AEQUEVOX framework.

AEQUEVOX facilitates fairness testing without having any access to ground truth transcription data. Although, text-to-speech (TTS) can be used for generating speech, we argue that it is not suitable for accurately identifying the bias towards speech coming from a certain group. Specifically, speakers may intentionally use enunciation, intonation, different degrees of loudness or other aspects of vocalization to articulate their message. Additionally, speakers unintentionally communicate their social characteristics such as their place of origin (through their accent), gender, age and education. This is unique to human speech and TTS systems cannot faithfully capture all the complexities inherent to human speech. Therefore, we believe that fairness testing of ASR systems should involve speech data from human speakers.

We note that human speech (and the ASRs) may be subject to adverse environments (e.g. noise) and it is critical that the fairness evaluation considers such adverse environments. To facilitate the testing of ASR systems in adverse environments, we model the speech signal as a sinusoidal wave and subject it to eight different metamorphic transformations (e.g. noise, drop, low/high pass filter) that are highly relevant in real life. Furthermore, in the absence of man-

ually transcribed speech, we use a differential testing methodology to expose fairness violations. In particular, AEQUEVOX identifies the bias in ASR systems via a two step approach: Firstly, AEQUEVOX registers the increase in error rates for speech from two groups when subjected to a metamorphic transformation. Subsequently, if the increase in the error rate of one group exceeds the other by a given threshold, AEQUEVOX classifies this as a violation of fairness. To the best of our knowledge, we are unaware of any such differential testing methodology. As a by product of our AEQUEVOX framework, we highlight words that contribute to errors by comparing the word counts from the original speech. This information can be further used to improve the ASR system.

Existing works [17,49] isolate certain sensitive attributes (e.g. gender) and use such attributes to test for fairness. Isolating these attributes is difficult in speech data, making it challenging to apply existing techniques to evaluate the fairness of ASR systems. AEQUEVOX tackles this by formalizing a unique fairness criteria targeted at ASR systems. Despite some existing efforts in testing ASR systems [5,13], these are not directly applicable for fairness testing. Additionally, some of these works require manually labelled speech transcription data [13]. Finally, differential testing via TTS [5] is not appropriate to determine the bias towards certain speakers, as they might use different vocalization that might be impossible (and perhaps irrational) to generate via a TTS. In contrast, AEQUEVOX works on speech signals directly and defines transformations directly on these signals. AEQUEVOX also does not require any access to manually labelled speech data for discovering fairness violations. In summary, we make the following contributions in the paper:

1. We formalize a notion of fairness for ASR systems. This formalization draws parallels between the Equality of Outcomes [36] and the quality of service provided by ASR systems in varying environments.
2. We present AEQUEVOX, which systematically combines metamorphic transformations and differential testing to highlight whether speech from a certain group (e.g. female) is subject to fairness violations by ASR systems. AEQUEVOX neither requires access to ground truth transcription data nor does it require access to the ASR model structures.
3. We propose a fault localization method to identify the different words contributing to fairness errors.
4. We evaluate AEQUEVOX with three different ASR systems namely Google Cloud, Microsoft Azure and IBM Watson. We use speech from the Speech Accent Archive [54], the Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS) [30], Multi speaker Corpora of the English Accents in the British Isles (Midlands) [11], and a Nigerian English speech dataset [2]. Our evaluation reveals that speech from non-native English speakers and female speakers exhibit higher fairness violations as compared to native English speakers and male speakers, respectively.
5. We validate the fault localization of AEQUEVOX by showing that the identified *faulty* words generally introduce more errors to ASR systems even when used within speech generated via TTS systems. The inputs to the TTS system are randomly generated sentences that conform to a valid grammar.

Table 1: Notations used

Notation	Description
$GR_B$	Base group
$GR_k$	$k \in (1, n)$ . Various comparison groups
$MT$	Metamorphic transformations
$ASR$	Automatic Speech Recognition system under test
$\tau$	A user specified threshold beyond which the difference in word error rate for the base and comparison groups is considered a violation of individual fairness

- We evaluate (via the user study) the human comprehensibility score of the transformations employed by AEQUEVOX on the speech signal. The lowest comprehensibility score was 6.78 and 82.9% of the transformations had a comprehensibility score of more than seven.

## 2 Background

In this section, we introduce the necessary background information.

**Fairness in ASR Systems:** A recent work, FairSpeech [26], uses conversational speech from black and white speakers to find that the word error rate for individuals who speak African American Vernacular English (AAVE) is nearly twice as large in all cases.

**Testing ASR Systems:** The major testing focus, till date has been on image recognition systems and large language models. Few papers have probed ASR systems. One such work, Deep-Cruiser [13] applies metamorphic transformations to audio samples to perform coverage-guided testing on ASR systems. Iwama et al. [23] also perform automated testing on the basic recognition capabilities of ASR systems to detect functional defects. CrossASR [5] is another recent paper that applies differential testing to ASR systems.

**The Gap in Testing ASR Systems:** There is little work on automated methods to formalise and test fairness in ASR systems. In this work, we present AEQUEVOX to test the fairness of ASR systems with respect to different population groups. It accomplishes this with the aid of differential testing of speech samples that have gone through metamorphic transformations of varying intensity. Our experimentation suggests that speech from different groups of speakers receives significantly different quality of service across ASR systems. In the subsequent sections, we describe the design and evaluation of our AEQUEVOX system.

## 3 Methodology

In this section, we discuss AEQUEVOX in detail. In particular, we motivate and formalize the notion of *fairness* in ASR systems. Then, we discuss our methodology to systematically find the violation of fairness in ASR systems. The notations used are described in Table 1.

**Motivation:** Equality of outcomes [36] describes a state in which all people have approximately the same material wealth and income, or in which the general

economic conditions of everyone’s lives are alike. For a software system, equality of outcomes can be thought of as everyone getting the same quality of service from the software they are using. For a lot of software services, providing the same quality of service is baked into the system by design. For example, the results of a search engine only depend on the query. The quality of the result generally does not depend on any sensitive attributes such as race, age, gender and nationality. In the context of an ASR, the quality of service does depend on these sensitive attributes. This inferior quality of service may be especially detrimental in safety-critical settings such as emergency medicine [18] or air traffic management [27,21].

In our work, we show that the quality of service provided by ASR systems is vastly different depending on one’s gender/nationality/accent. Suppose there are two groups of people using an ASR system, males and females. They have approximately the same level of service when using this service at their homes. However, once they step into a different environment such as a noisy street, the quality of service drops notably for the female users, but does not drop noticeably for the male users. This is a violation of the principle of equality of outcomes (as seen for software systems) and more specifically, group fairness [14]. Such a scenario is unfair (violation of group fairness) because some groups enjoy a higher quality of service than others.

In our work, we aim to automate the discovery of this unfairness. We do this by simulating the environment where the behaviour of ASR systems are likely to vary. The simulated environment is then enforced in speech from different groups. Finally, we measure how different groups are served in different environments.

**Formalising Fairness in ASRs:** In this section, we formalise the notion of fairness in the context of automated speech recognition systems (ASRs). The fairness definition in ASRs is as follows:

$$|ASR_{Err}(GR_i) - ASR_{Err}(GR_j)| \leq \tau \quad (1)$$

Here,  $GR_i$  and  $GR_j$  capture speech from distinct groups of people. If the error rates induced by  $ASR$  for group  $GR_i$  ( $ASR_{Err}(GR_i)$ ) and for group  $GR_j$  ( $ASR_{Err}(GR_j)$ ) differ beyond a certain threshold, we consider this scenario to be unfair. Such a notion of unfairness was studied in a recent work [26].

In this work, we want to explore whether different groups are fairly treated under varying conditions. Intuitively, we subject speech from different groups to a variety of simulated environments. We then measure the word error rates of the speech in such simulated environments and check if certain groups fare better than others. Formally, we capture the notion of fairness targeted by AEQUEVOX as follows:

$$\begin{aligned} D_i &\leftarrow ASR_{Err}(GR_i) - ASR_{Err}(GR_i + \delta) \\ D_j &\leftarrow ASR_{Err}(GR_j) - ASR_{Err}(GR_j + \delta) \end{aligned} \quad (2)$$

$$|D_i - D_j| \leq \tau$$

Here we perturb the speech of the two groups ( $GR_i$  and  $GR_j$ ) by adding some  $\delta$  to the speech. We compare the degradation in the speech ( $D_i$  and  $D_j$ ). If the

**Algorithm 1** AEQUEVOX Fairness Testing

---

```

1: procedure FAIRNESS_TESTING( $GR_B, MT, GR_1, \dots, GR_n, \tau, ASR_1, ASR_2$ )
2:    $Error\_Set \leftarrow \emptyset$ 
3:   for  $T \in MT$  do
4:      $GR_B^T \leftarrow T(GR_B)$ 
5:      $\triangleright \mathcal{L}$  computes the average word level levenshtein distance
6:      $\triangleright$  between the outputs of  $ASR_1$  and  $ASR_2$ 
7:      $d_B \leftarrow \mathcal{L}(ASR_1(GR_B), ASR_2(GR_B))$ 
8:      $d_B^T \leftarrow \mathcal{L}(ASR_1(GR_B^T), ASR_2(GR_B^T))$ 
9:      $D_B \leftarrow d_B^T - d_B$ 
10:    for  $k \in (1, n)$  do
11:       $GR_k^T \leftarrow T(GR_k)$ 
12:       $d_k \leftarrow \mathcal{L}(ASR_1(GR_k), ASR_2(GR_k))$ 
13:       $d_k^T \leftarrow \mathcal{L}(ASR_1(GR_k^T), ASR_2(GR_k^T))$ 
14:       $D_k \leftarrow d_k^T - d_k$ 
15:      if  $D_B - D_k > \tau$  then
16:         $Error\_Set \leftarrow Error\_Set \cup (GR_B, GR_k, T)$ 
17:      end if
18:    end for
19:  end for
20:  return  $Error\_Set$ 
21: end procedure

```

---

degradation faced by one group is far greater than the one faced by the other, we have a fairness violation. This is because speech from both groups ought to face similar degradation when subject to similar environments (simulated by  $\delta$  perturbation) when equality of outcomes [36] holds. More specifically, this is a group fairness violation because the quality of service (outcome) depends on the group [14,51].

**Example:** To motivate our system, let us sketch out an example. Consider texts of approximately the same length spoken by two sets of speakers whose native languages are  $L_1$  and  $L_2$  respectively. Let us assume that both sets of speakers read out a text in English. AEQUEVOX uses two ASR systems and obtains the transcript of this speech. AEQUEVOX then employs differential testing to find the word-level levenshtein distance [29] between these two sets of transcripts. Let us also assume that the average word-level levenshtein distance is two and four for  $L_1$  and  $L_2$  native speakers, respectively.

AEQUEVOX then simulates a noisy environment by adding noise to the speech and obtains the transcript of this transformed speech. Let us assume now that the average levenshtein distance for this transformed speech is 4 and 25 for  $L_1$  and  $L_2$  native speakers, respectively. It is clear that the degradation for the speech of native  $L_2$  speakers is much more severe. In this case, the quality of service that  $L_2$  native speakers receive in noisy environments is worse than  $L_1$  native speakers. This is a violation of fairness which AEQUEVOX aims to detect.

The working principle behind AEQUEVOX holds even if the spoken text is different. This is because AEQUEVOX just measures the relative degradation in ASR performance for a set of speakers. For large datasets, we are able to measure the average degradation in ASR performance with respect to different groups of speakers (e.g. *male, female, native, non-native English* speakers).

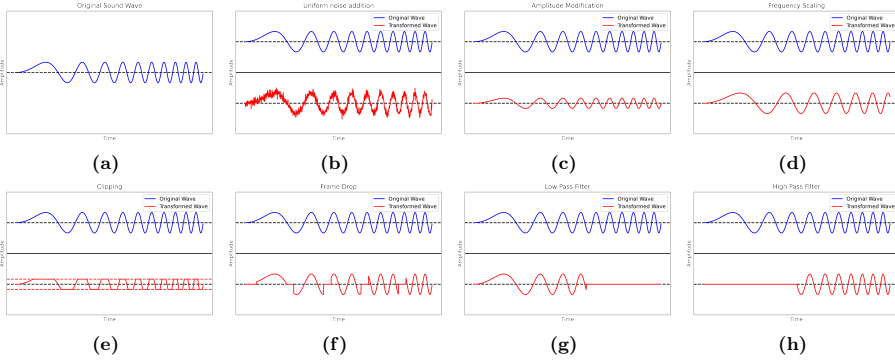


Fig. 2: Sound wave transformations




Fig. 3: AEQUEVOX System Overview

**Metamorphic Transformations of Sound:** The ability to operate in a wide range of environments is crucial in ASR systems as they are deployed in safety-critical settings such as medical emergency services [18] and air traffic management [21], [27], which are known to have interference and noise. Metamorphic speech transformations serve to simulate such scenarios. The key insight for our metamorphic transformations comes from how waves are represented and what can happen to these waves when they’re transmitted in different mediums. We realise this insight in the fairness testing system for ASR systems. To the best of our knowledge AEQUEVOX is the first work that combines this insight from acoustics, software testing and software fairness to evaluate the fairness of ASR systems. AEQUEVOX uses the addition of noise (Figure 2 (b)), amplitude modification (Figure 2 (c)), frequency modification (Figure 2 (d)), amplitude clipping (Figure 2 (e)), frame drops (Figure 2 (f)), low-pass filters (Figure 2 (g)), and high-pass filters (Figure 2 (h)) as metamorphic speech transformations. We choose these transformations because they are the most common distortions for sound in various environments [1].

**System Overview:** Algorithm 1 provides an outline of our overall test generation process. We realise the notion of fairness described in Equation (2) using differential testing. The error rates ( $ASR_{Err}$ ) for a particular speech clip are

found by finding the difference between the outputs of two ASR systems,  $ASR_1$  and  $ASR_2$ . It is important to note that we make a design choice to use *differential testing* to find the error rate ( $ASR_{Err}$ ). This helps us eliminate the need for ground truth transcription data which is both labor intensive and expensive to obtain. Furthermore, AEQUEVOX realises the  $\delta$  seen in Equation (2) by using metamorphic transformations for speech (see Figure 3). These speech metamorphic transformations represent the various simulated environments for which AEQUEVOX wants to measure the quality of service for different groups. Additionally, the user can customise this  $\delta$  per their requirements. In our implementation we use eight distinct metamorphic transformations as  $\delta$  (see Figure 2). Specifically, we investigate how fairly do two ASR systems ( $ASR_1$  and  $ASR_2$ ) treat groups ( $GR_k^T \mid k \in \{1, 2, \dots, n\}$ ) with respect to a base group ( $GR_B$ ). AEQUEVOX achieves this by taking a dataset of speech which contains data from two or more different groups (e.g. male and female speakers, Native English and Non-native English speakers) and modifies these speech snippets through a set of transformations ( $MT$ ). These are then divided into base group transformed speech ( $GR_B^T$ ) and the transformed speech for other groups ( $GR_k^T \mid k \in \{1, 2, \dots, n\}$ ). As seen in Algorithm 1, the average word-level levenshtein distance (word-level levenshtein distance divided by the number of words in the longer transcript) between the outputs of the two ASR systems is captured by  $d_B$  and  $d_B^T$  for the original and transformed speech respectively. Similarly, for the comparison groups  $GR_k^T (k \in \{1, 2, \dots, n\})$  the word-level levenshtein distance is captured by  $d_k$  and  $d_k^T$ . The higher the levenshtein distance the larger the error in terms of differential testing. In other words, larger error in differential testing would mean that the ASR systems disagree on a higher number of words.

To capture the degradation in the quality of service for the speech subjected to simulated environments ( $MT$ ), we compute the difference between the word-level levenshtein distance for the original and transformed speech. Specifically, we compute  $D_B$  as  $d_B^T - d_B$  and  $D_k$  as  $d_k^T - d_k (k \in \{1, 2, \dots, n\})$  for the base and comparison groups, respectively. The higher this metric ( $D_B$  and  $D_k$ ), the more severe the degradation in ASR quality of service because of the transformation  $T$ .

We compare these metrics and if  $D_B$  exceeds  $D_k$  by some threshold  $\tau$ , we classify this as an error for the *base group* ( $GR_B$ ) and more specifically a violation of fairness (see Figure 3). In our experiments we set each of the groups in our dataset as the base group ( $GR_B$ ) and run the AEQUEVOX technique to find errors with respect to that base group. The lower the errors (as computed via the violation of the assertion  $D_B - D_k \leq \tau$ ), the fairer the ASR systems are with respect to groups  $GR_B$ . As an example, let us say Russian speakers are the base group ( $GR_B$ ), English speakers are the comparison group ( $GR_k$ ) and the value of  $\tau$  is 0.1. If  $D_B$  is strictly greater than  $D_k$  by 0.1, then fairness violation is counted for the Russian speakers. Otherwise, no fairness errors are recorded.

**Fault Localisation:** AEQUEVOX introduces a word-level fault localisation technique, which does not require any access to ground truth data. We first illustrate a use case of this fault localisation technique.



**Algorithm 2** AEQUEVOX Fault Localizer

---

```

1: procedure FAULT_LOCALIZER( $WC, WC^{T_\theta}, \omega, param^T$ )
2:   Drop_Count  $\leftarrow \emptyset$ 
3:   Non_Robust_Words  $\leftarrow \emptyset$ 
4:   for  $word \in WC.keys()$  do
5:      $init\_count \leftarrow WC[word]$ 
6:      $\triangleright$  Returns the minimum count of  $word$  across all the parameter
7:      $\triangleright$  of transformation  $T$ 
8:      $min\_count \leftarrow get\_min(WC^{T_\theta}[word], param^T)$ 
9:      $count\_diff \leftarrow \max((init\_count - min\_count), 0)$ 
10:    if  $count\_diff > \omega$  then
11:      Non_Robust_Words  $\leftarrow$  Non_Robust_Words  $\cup \{word\}$ 
12:    end if
13:    Drop_Count  $\leftarrow$  Drop_Count  $\cup \{count\_diff\}$ 
14:  end for
15:  return Non_Robust_Words, Drop_Count
16: end procedure

```

---




Fig. 4: AEQUEVOX Fault Localization Overview

**Example:** Let us consider a corpus of English sentences by a group of speakers (say  $GR$ ) who speak language  $L_1$  natively. AEQUEVOX builds a dictionary for all the words in the transcript obtained from  $ASR_1$ . An excerpt from such a dictionary appears as follows:  $\{brother : 16, nice : 25, is : 33, \dots\}$ . This means the words *brother*, *nice* and *is* were seen 16, 25 and 33 times in the transcript respectively. Now, assume AEQUEVOX simulates a noisy environment by adding noise with various signal to noise (SNR) ratios as follows:  $\{10, 8, 6, 4, 2\}$ . This is the parameter for the transformation ( $param^T$ ).

Once AEQUEVOX obtains the transcript of these transformed inputs, it creates dictionaries similar to the ones seen in the preceding paragraph. Let the relevant subset of the dictionary for SNR *two* (2) be  $\{brother : 1, nice : 23, is : 32, \dots\}$ . We use this to determine that the utterance of the word *brother* is not robust for noise addition for the group  $GR$ . This is because, the word *brother* appears significantly less in the transcript for the modified speech, as compared to the transcript for the original speech.

**AEQUEVOX fault localisation overview:** Algorithm 2 provides an overview of the fault localization technique implemented in AEQUEVOX. The goal of the AEQUEVOX fault localisation is to find words for a group ( $GR$ ) that are not robust to the simulated environments. Specifically, AEQUEVOX finds words which are not recognised by the ASR when subjected to the appropriate speech transformations.

The transformation is represented by  $T_\theta$ . Here,  $T \in MT$  is the transformation and  $\theta \in param^T$  is the parameter of the transformation, which controls the severity of the transformation.

As seen in Algorithm 2, AEQUEVOX builds a word count dictionary for each word in  $WC$  and  $WC^{T_\theta}$  for the original speech and for each  $\theta \in param^T$  re-

spectively. For each word, AEQUEVOX finds the difference in the number of appearances for a word in  $WC$  and in  $WC^{T_\theta}$  for  $\theta \in param^T$ . To compute the difference, we locate the minimum number of appearances across all the transformation parameters  $\theta \in param^T$  (i.e. *min\_count* in Algorithm 2). This is to locate the worst-case degradation across all transformation parameters. The difference is then calculated between *min\_count* and the number of appearances of the word in the original speech (i.e. *init\_count*). If the difference exceeds some user-defined threshold  $\omega$ , then AEQUEVOX classifies the respective words as non robust w.r.t the group  $GR$  and transformation  $T$ .

We envision that practitioners can then review the data generated by fault localization (i.e. Algorithm 2) and target the non-robust words to further improve their ASR systems for speech from underrepresented groups [24] and accommodate for speech variability [22]. In **RQ3**, we validate our fault localization method empirically and in **RQ4**, we show how the proposed fault localization method can be used to highlight fairness violations.

## 4 Datasets and Experimental Setup

**ASR Systems under Test:** We evaluate AEQUEVOX on three commercial ASR systems from Google Cloud Platform (GCP), IBM Cloud, and Microsoft Azure. We use the standard models for GCP and Azure, and the *BroadbandModel* for IBM. In all three cases, the audio samples were identically encoded as .wav files using Linear 16 encoding.

In each of the following transformations, we vary a parameter,  $\theta$ . We call this the transformation parameter. Some of the transformations have abbreviations within parentheses. Such abbreviations are used in later sections to refer to the respective transformations.

**Amplitude Scaling (Amp):** For amplitude scaling, we scale the audio sequence by a constant by multiplying each individual audio sample by  $\theta$ .

**Clipping:** The audio samples are scaled such that their amplitude values are bound by  $[-1, 1]$ . AEQUEVOX then clips these samples such that the amplitude range is  $[-\theta, \theta]$ . These clipped samples are then rescaled and encoded.

**Drop/Frame:** For Drop, AEQUEVOX divides the audio into 20ms chunks.  $\theta\%$  of these chunks are then randomly discarded (amplitude set to zero) from the audio. For Frame, AEQUEVOX divides the audio into  $\theta$ ms chunks and 10% of these chunks are then randomly discarded. No two adjacent chunks are discarded.

**High Pass (HP)/ Low Pass (LP) Filter:** Here we apply a butterworth [7] filter of order two to the entire audio file with  $\theta$  determining the cut-off frequency.

**Noise Addition (Noise):**  $\theta$  represents signal to noise (SNR) ratio [25] of the transformed audio signal. A lower  $\theta$  means higher noise in the transformed audio.

**Frequency Scaling (Scale):** In this case,  $\theta$  is the sampling frequency. The lower the value of  $\theta$ , the slower the audio. In this transformation, the audio is slowed down  $\theta$  times.

Table 2: Transformations Used

Transformation Type	$\theta$ Used				
	Least Destructive $\rightarrow$ Most Destructive				
Amplitude	0.5	0.4	0.3	0.2	0.1
Clipping	0.05	0.04	0.03	0.02	0.01
Drop	5	10	15	20	25
Frame	10	20	30	40	50
HP	500	600	700	800	900
LP	900	800	700	600	500
Noise	10	8	6	4	2
Scale	0.9	0.8	0.7	0.6	0.5

Table 3: Datasets Used

Dataset	Duration(s)	#Clips	#Distinct Speakers
Accents	25-35	28	28
RAVDESS	3	32	8
Midlands	3-5	4	4
Nigerian English	4-6	4	4

Table 2 lists all the different values used for  $\theta$ . An additional parameter ( $\theta = 2.0$ ) is used for *Amp*.

**Datasets:** We use the Speech Accent Archive (Accents) [54], the Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS) [30], Multi speaker Corpora of the English Accents in the British Isles (Midlands) [11], and a Nigerian English speech dataset [2] to evaluate AEQUEVOX taking care to ensure male and female speakers are equally represented. Table 3 provides additional details about the setup.

## 5 Results

In this section, we discuss our evaluation of AEQUEVOX in detail. In particular, we structure our evaluation in the form of four research questions (**RQ1** to **RQ4**). The analysis of these research questions appears in the following sections.

### **RQ1: What is AEQUEVOX’s efficacy?**

We structure the analysis of this research question into three sections, each corresponding to a dataset we have used in our analysis. All of the relevant data is presented in Table 4 with the lowest errors for each dataset bolded. We first analyse the number of errors (used interchangeably with fairness violations) for each case. Subsequently, we analyse the sensitivity of the errors with respect to the values of  $\tau$  ( $\tau \in \{0.01, 0.05, 0.1, 0.15\}$ ). Detecting violations of fairness is regulated by parameter  $\tau$ . Lower values of  $\tau$  imply that the degradation of word error rates between two groups should be similar, and conversely higher values of  $\tau$  allow for the difference in degradation of word error rates to be more severe between two groups. Next, we analyse the sensitivity of the pairs of the ASR systems under test. Concretely, we analyse the errors found in the Microsoft Azure and IBM Watson (*MS\_IBM*), Google Cloud and IBM Watson (*IBM\_GCP*), and Microsoft Azure and Google Cloud (*MS\_GCP*) pairs. Finally, we analyse the sensitivity of the AEQUEVOX test generation with respect to the eight different types of transformations implemented (see Figure 2).

Table 4: Errors Discovered by AEQUEVOX

	Accents								RAVDESS		Nigerian/Midlands English	
	English	Ganda	French	Gujarati	Indonesian	Korean	Russian	Male	Female	Midlands	Nigerian	
$\tau$ Sensitivity												
	<i>0.01</i>	<b>168</b>	381	267	232	178	499	354	<b>12</b>	92	<b>36</b>	75
	<i>0.05</i>	<b>75</b>	245	99	101	85	340	227	<b>8</b>	53	<b>26</b>	65
	<i>0.10</i>	43	145	39	49	<b>34</b>	172	161	<b>5</b>	21	<b>17</b>	55
	<i>0.15</i>	26	73	<b>8</b>	24	14	75	111	<b>3</b>	10	<b>14</b>	44
ASR Sensitivity												
	<i>MS IBM</i>	<b>36</b>	369	128	126	64	388	303	<b>10</b>	57	<b>30</b>	86
	<i>GCP IBM</i>	131	325	123	147	<b>98</b>	342	361	<b>9</b>	64	<b>31</b>	96
	<i>MS GCP</i>	145	150	162	<b>133</b>	149	356	189	<b>9</b>	55	<b>32</b>	57
Transition Sensitivity												
	<i>Clipping</i>	<b>4</b>	81	38	159	72	182	237	<b>0</b>	24	50	<b>3</b>
	<i>Drop</i>	<b>8</b>	113	33	29	40	184	45	<b>0</b>	21	<b>4</b>	33
	<i>Frame</i>	<b>14</b>	106	61	25	36	170	26	<b>1</b>	13	<b>13</b>	19
	<i>Noise</i>	<b>5</b>	128	54	86	22	217	213	<b>0</b>	24	<b>5</b>	43
	<i>LP</i>	39	158	108	57	<b>14</b>	110	208	<b>0</b>	45	<b>4</b>	34
	<i>Amplitude</i>	81	19	44	33	<b>14</b>	40	26	<b>0</b>	27	<b>8</b>	40
	<i>HP</i>	114	168	29	<b>9</b>	61	87	57	<b>9</b>	20	<b>1</b>	51
	<i>Scale</i>	47	71	46	<b>8</b>	52	96	41	<b>18</b>	<b>2</b>	<b>8</b>	16
<b>Total Errors</b>		312	844	413	406	<b>311</b>	1086	853	<b>28</b>	176	<b>93</b>	239

It is important to note that we excluded the two most destructive Scale transformations. This is because the word error rate for these transformations is 0.89 on average out of 1. This degradation may be attributed to the transformation itself rather than the ASR. To avoid such cases, we exclude these transformations from this research question.

**Accents Dataset:** Native English speakers and Indonesian speakers have the lowest number of errors. On average, speech from non-native English speakers generates *109% more errors* in comparison to speech from native English speakers. For the two smallest values of  $\tau$ , speech from the native English speakers shows the *least number* of fairness violations. Speech from native English speakers has the lowest, second lowest and third lowest errors for the pairs of ASRs, (*MS\_IBM*), (*MS\_GCP*) and (*IBM\_GCP*) respectively. Speech from native English speakers has the lowest errors for the clipping, two types of frame drops and noise transformations and the second lowest errors for the low-pass filter transformation. The high-pass filter and scaling induce a comparable number of errors from native and a majority of the non-native English speakers. However, speech from native English speakers has the highest number of errors when subject to the amplitude transformation.

*Speech from non-native English speakers generally exhibits more fairness violations in comparison to speech from native English speakers.*

**RAVDESS Dataset:** Speech from male speakers has significantly lower errors than speech from female speakers. On average, speech from female speakers generates *528.57% more errors* in comparison to speech from male speakers. Speech from male speakers shows significantly fewer fairness violations for all values of  $\tau$ , and for all ASR pairs tested. Clipping, both types of frame drops, noise, low-pass, amplitude and high-pass transformations induce significantly

fewer errors on speech from male speakers. However, speech from male speakers has more errors when subject to scale transformations.

*Speech from female speakers has significantly more fairness violations in comparison to speech from male speakers.*

**Midlands/Nigeria Dataset:** Speech from UK Midlands English (ME) speakers has significantly fewer errors than speech from Nigerian English (NE) speakers. On average, speech from NE speakers generates *156.9% more errors* in comparison to speech from ME speakers. Speech from ME speakers has significantly fewer fairness errors for all values of  $\tau$ , and for all ASR pairs tested. For the transformations scale, drop, noise, amplitude, low pass and high pass filters, the speech from ME speakers has significantly fewer errors than speech from NE speakers. Clipping induces more errors in speech from ME speakers, while the frame transformation induces comparable number of errors in speech from both groups.

*Speech from Nigerian English speakers has significantly more fairness errors in comparison to speech from UK Midlands speakers.*

## RQ2: What are the effects of transformations on comprehensibility?

To better understand the effects of the transformations (*see Figure 2*) on the comprehensibility of the speech we conducted a user study. Speech of one randomly chosen *female native English* speaker from the Accents [54] dataset was used since the audio contains nearly all the sounds present in the English language [54]. Survey participants were presented with the original audio file along with a set of transformed speech files in order of increasing intensity. All the transformations (*see Figure 2*) and transformation parameters (*see Table 2*) were used. We asked 200 survey participants (sourced through Amazon mTurk) the following question:

*How comprehensible is (transformed) Speech with respect to the Original speech?*




Fig. 5: Average Transformation Comprehensibility Ratings

The rating of one (1) is *Not Comprehensible at all* and the rating of ten (10) is *Just as Comprehensible as the Original*.

Unsurprisingly, as seen in Figure 5, increasing the intensities of the transformation had a generally detrimental effect on the comprehensibility of the speech. But none of the transformations majorly affect the comprehensibility of the speech. All of the transformations had an average comprehensibility rating above 6.75 and 82.9% of the transformations had a comprehensibility rating above 7.

Table 5: Fairness errors where the transformations have a comprehensibility rating of at least 7.2

	Accents							RAVDESS		Nigerian/Midlands English	
	English	Ganda	French	Gujarati	Indonesian	Korean	Russian	Male	Female	Midlands	Nigerian
<b>Total Errors</b>	246	509	240	<b>166</b>	225	687	329	<b>28</b>	88	<b>55</b>	161

Table 6: Grammar-generated sentence examples

ASR	Microsoft	Google Cloud	IBM Watson
<i>Robust</i>	Ashley likes <b>fresh</b> smoothies Paul adores <b>spoons</b> of cinnamon	Karen loves <b>plastic</b> straws Donald hates <b>big</b> decisions	William detests <b>plastic</b> cups Steven detests <b>big</b> flags
<i>Non-robust</i>	Ashley detests <b>thick</b> smoothies Ryan likes <b>slabs</b> of cake	John loves <b>spoons</b> of cinnamon Robert loves <b>bags</b> of concrete	Betty likes <b>scoops</b> of ice cream Amanda is fond of <b>things</b> like groceries

The average degradation in comprehensibility for the least destructive parameter across all transformations was 24.36%. Noise was the most destructive at 27.75% and drop was the least destructive (20.96%).

The average degradation in comprehensibility for the most destructive parameter across all transformations was 29.18%. In this case, scaling was the most destructive at 32.23% whereas drop was the least destructive with 25.88%.

Additionally, for each transformation, we analyse the percentage drop of comprehensibility between the least and the most destructive transformation parameters. The average drop is 4.82% across all transformations. The scaling and drop transformations show high relative percentage drops of 10.05% and 8.32% respectively. Amplitude, clipping, noise, high-pass and low-pass filters show closer to average drops between 3.1% and 4.5%. Frame, on the other hand, shows very low relative drops at 0.76%.

*All the transformations, though destructive, are comprehensible by humans.*

For safety critical applications, we recommend that future work test the whole gamut of transformations. For other use cases, practitioners may choose the transformations that satisfy their needs. To aid this, AEQUEVOX allows the users to choose the comprehensibility threshold of the transformations. As seen in Table 5, our conclusion holds even if we choose the transformations with higher comprehensibility threshold (7.2). We highlight the group with the least errors in each dataset to aid in readability. In particular, we observe that speech from male and UK Midlands speakers generally exhibit fewer errors. Setting aside speech from native Gujarati speakers, speech from native English speakers exhibits comparable or better performance than speech from other groups.

### **RQ3: Are the outputs produced by AEQUEVOX fault localiser valid?**

To study the validity of the outputs of the fault localiser, we study the number of errors for the predicted robust and non-robust words. We do this by generating speech containing the predicted robust and non-robust words for each ASR tested. We choose an  $\omega$  of three, three and two for GCP, MS Azure and IBM respectively to choose the non-robust words (*see Algorithm 2*). We choose the robust words from the set of words that do not show any errors

in the presence of noise ( $count\_diff = 0$  in *Algorithm 2*) for these specific ASR systems. Specifically, we test whether the robust and non-robust words identified by the fault localiser in the Accents dataset are robust in the presence of noise. Our goal is to show that if noise is added to speech containing these non-robust words, the ASR will be less likely to recognise them. Vice-versa, if noise is added to the predicted robust-words they are less likely to be affected.

To generate the speech from the output we generate sentences containing the robust and non-robust words predicted by the fault localiser for each ASR using a grammar and then use a text-to-speech (TTS) service to generate speech. The actual randomly selected robust and non-robust words (in bold) and the examples of the sentences generated by the grammar can be seen in Table 6. We use the Google TTS for MS Azure and we use the Microsoft Azure TTS for GCP and IBM to generate the speech.

To evaluate the generality of outputs of the fault localisation technique, we use the speech produced by the TTS and then add noise to that speech. This speech is used to generate a transcript from the ASR and the transcript is used to evaluate how many of the predicted robust and non-robust words are incorrect in the transcript. We add the most noise possible to the TTS speech in our AEQUEVOX framework. Specifically, the signal to noise (SNR) ratio is 2. We use the TTS generated speech for 50 sentences for each of the robust and non-robust cases. Each sentence has either a robust or a non-robust word.

The results of the experiments are seen in Table 7. In the transcript of the speech with noise added at SNR 2, robust words show *zero* errors for the predicted robust words for Microsoft and Google Cloud and 21 errors for IBM. The non-robust words, on the other hand, had 23, 15 and 30 errors.

*The predicted non-robust words have a higher propensity for errors than the robust words.*

Table 7: Transcript Errors

ASR	No. of Erroneous Sentences
Microsoft (MS)	
<i>Robust</i>	0
<i>Non-Robust</i>	23
Google Cloud (GCP)	
<i>Robust</i>	0
<i>Non-Robust</i>	15
IBM Watson (IBM)	
<i>Robust</i>	21
<i>Non-Robust</i>	30

Table 8: Grammarly Scores

ASR	Overall Score	Correctness	Clarity
Microsoft (MS)			
<i>Robust</i>	99		
<i>Non-Robust</i>	99		
Google Cloud (GCP)		Looking	Very
<i>Robust</i>	100	Good	Clear
<i>Non-Robust</i>	99		
IBM Watson (IBM)			
<i>Robust</i>	100		
<i>Non-Robust</i>	96		

**Note on grammar validity:** Since the grammars used by us to validate the explanations of AEQUEVOX are handcrafted, they may be prone to errors. To verify these hand crafted grammars, we use 100 sentences produced by each grammar and use the online tool Grammarly [3] to investigate the semantic and syntactic correctness of the sentences and the clarity. The sentences generated by the grammars have a high overall average score of 98.33 out of 100, with the

Table 9: Average word mispredictions in the Accents dataset using the AE-QUEVOX localisation techniques

		Accents						
		English	Ganda	French	Gujarati	Indonesian	Korean	Russian
<b>ASR Sensitivity</b>								
	<i>GCP</i>	1.21	1.51	1.21	1.17	<b>1.07</b>	1.55	1.64
	<i>IBM</i>	<b>1.03</b>	1.94	1.38	1.35	1.48	1.92	1.70
	<i>MS Azure</i>	0.47	0.66	0.40	0.48	<b>0.36</b>	0.87	0.63
<b>Transition Sensitivity</b>								
	<i>Clipping</i>	<b>2.00</b>	2.53	2.12	2.60	2.29	2.81	3.13
	<i>Drop</i>	<b>0.30</b>	1.02	0.52	0.54	0.57	1.15	0.74
	<i>Frame</i>	<b>0.38</b>	0.89	0.68	0.56	0.51	1.19	0.65
	<i>Noise</i>	<b>0.57</b>	1.60	0.85	1.27	0.71	1.74	1.54
	<i>LP</i>	1.72	2.22	1.90	1.79	<b>1.58</b>	1.98	2.13
	<i>Amplitude</i>	0.17	0.15	0.11	0.12	<b>0.06</b>	0.20	0.16
	<i>HP</i>	0.74	0.75	0.38	<b>0.22</b>	0.49	0.64	0.76
	<i>Scale</i>	1.38	1.79	1.42	<b>0.90</b>	1.54	1.89	1.45

lowest being 96 (see Table 8). On the correctness and clarity measure, all the sentences generated by the grammars score *Looking Good* and *Very Clear*.

#### RQ4: Can the fault localiser be used to highlight unfairness?

The goal of this RQ is to investigate if the output of Algorithm 2 can call attention to bias between different groups. Specifically, we evaluate if some groups show fewer faults, on average than others. To this end, we use the fault localisation algorithm (Algorithm 2) on the accents dataset and record the number of words incorrect in the transcript, on average for each group of the accents dataset. This is done for each ASR under test. It is also important to note that this technique uses no ground truth data and requires no manual input. This technique is designed to work with just the speech data and metadata (groups).

Table 9 shows the average word drops across all transformations for the accents dataset for each ASR under test. We highlight the best performing groups by bolding the values. Speech from native-English speakers shows the lowest average word drops for the IBM Watson ASR and the third lowest for GCP and MS Azure ASRs. We also investigate the average word drops for each transformation in AEQUEVOX averaged across all ASRs. Speech from native English speakers has the lowest average word drops for the Clipping, two types of frame drops and noise transformations and the second lowest errors for the low-pass filter transformation. (see Table 9). For the rest of the transformations, namely amplitude, high-pass filter and scaling, we find that both speech from non-native English speakers and speech from native English speakers have comparable average word drops in the majority of cases (see Table 9). This result is consistent with results seen in RQ1.

*The technique seen in Algorithm 2 can be used to highlight bias in speech and the results are consistent with RQ1.*



## 6 Threats to Validity

**User Study:** In conducting the study, two assumptions were made. Firstly, we assume that the degree to which comprehensibility changes when subject to transformations is independent of the characteristics of the speaker’s voice. Secondly, we assume that the speech is reflective of the broader English language. In future work, a larger scale user study could be performed to verify the results.

**ASR Baseline Accuracy:** AEQUEVOX measures the degradation of the speech to characterise the unfairness amongst groups and ASR systems. If the baseline error rate is very high, then the room for further degradation is very low. As a result, AEQUEVOX expects ASR services to have a high baseline accuracy. To mitigate this threat, we use state-of-the-art commercial ASR systems which have high baseline accuracies.

**Completeness and Speech Data:** AEQUEVOX is incomplete, by design, in the discovery of fairness violations. AEQUEVOX is limited by the speech data and the groups of this speech data used to test these ASR systems. With new data and new groups, it is possible to discover more fairness violations. The practitioners need to provide data to discover these. In our view, this is a valid assumption because the developers of these systems have a large (and growing) corpus of such speech data. It is also important to note that AEQUEVOX does not need the ground truth transcripts for this speech data and such speech data is easier to obtain.

**Fault Localisation:** To test AEQUEVOX’s fault localisation, we identify the robust and non-robust words in the speech and subsequently construct sentences (with the aid of a grammar). These sentences are then converted to speech using a text-to-speech (TTS) software and the performance of the robust and non-robust words is measured. In the future, we would like to repeat the same experiment with a fixed set of speakers, which allows us capture the peculiarities of speech in contrast to the usage of TTS software.

## 7 Related Work

In the past few years, there has been significant attention in testing ML systems [35,48,32,47,55,34,50,40,56,16,52,8,41,19]. Some of these works target coverage-based testing [48,55,34,32] or leverage property driven testing [41], while others focus on effective testing in targeted domains e.g. text [50,40]. None of these works, however, are directly applicable for testing ASR systems. In contrast, the goal of AEQUEVOX is to automatically discover violations of fairness in ASR systems without access to ground truth data.

DeepCruiser [13] uses metamorphic transformations and performs coverage-guided fuzzing to discover transcription errors in ASR systems. Concurrently, CrossASR [5] uses text to generate speech from a TTS engine and subsequently employs differential testing to find bugs in the ASR system. In contrast to these systems, the goal of AEQUEVOX is to automatically find violations of fairness

by measuring the degradation of transcription quality from the ASR when the speech is transformed. AEQUEVOX compares this degradation across various groups of speakers and if the difference is substantial, AEQUEVOX characterises this as a fairness violation. Moreover, AEQUEVOX neither requires access to manually labelled speech data nor does it require any white/grey box access to the ASR model. Works on audio adversarial testing [23], [10], [9], [37], [28] aims to find an imperceptible perturbation that are specially crafted for an audio file. In contrast, AEQUEVOX aims to find fairness violations. Additionally, AEQUEVOX also proposes automatic fault localisation for ASR systems without using a ground truth transcript.

Unlike AEQUEVOX, recent works on fairness testing have focused on credit rating [17,49,4,57,42,44,43,41], computer vision [12,6] or NLP systems [33,45]. In the systems that deal with such data, it is possible to isolate certain sensitive attributes (gender, age, nationality) and test for fairness based on these attributes. It is challenging to isolate such sensitive attributes in speech data, necessitating the need for a separate fairness testing framework specifically for speech data.

Frameworks such as LIME [38], SHAP [31], Anchor [39] and DeepCover [46] attempt to reason why a model generates a specific output for a specific input. In contrast to this, AEQUEVOX's fault localisation algorithm identifies utterances spoken by a group which are likely to be not recognised by ASR systems in the presence of a destructive interference (such as noise). Recent fault localization approaches either aim to highlight the neurons [15] or training code [53] that are responsible for a fault during inference. In contrast, AEQUEVOX highlights words that are likely to be transcribed wrongly without having any access to the ground truth transcription and with only blackbox access to the ASR system.

## 8 Conclusion

In this work, we introduce AEQUEVOX, an automated fairness testing technique for ASR systems. To the best of our knowledge, we are the first work that explores considerations beyond error rates for discovering fairness violations. We also show that the speech transformations used by AEQUEVOX are largely comprehensible through a user study. Additionally, AEQUEVOX highlights words where a given ASR system exhibits faults, and we show the validity of these explanations. These faults can also be used to identify unfairness in ASR systems.

AEQUEVOX is evaluated on three ASR systems and we use four distinct datasets. Our experiments reveal that speech from non-native English, female and Nigerian English speakers exhibit more errors, on average than speech from native English, male and UK Midlands speakers, respectively. We also validate the fault localization embodied in AEQUEVOX by showing that the predicted non-robust words exhibit 223.8% more errors than the predicted robust words across all ASRs.

We hope that AEQUEVOX drives further work on systematic fairness testing of ASR systems. To aid future work, we make all our code and data publicly available at [github.com/sparkssss/AequeVox](https://github.com/sparkssss/AequeVox) and [10.5281/zenodo.5897347](https://zenodo.org/record/5897347)

## References

1. Audio data augmentation (2021), <https://www.kaggle.com/CVxTz/audio-data-augmentation>
2. Crowdsourced high-quality nigerian english speech data set (2021), <http://openslr.org/70/>
3. Grammarly (2021), <https://app.grammarly.com/>
4. Aggarwal, A., Lohia, P., Nagar, S., Dey, K., Saha, D.: Black box fairness testing of machine learning models. In: Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering. pp. 625–635 (2019)
5. Asyrofi, M.H., Thung, F., Lo, D., Jiang, L.: Crossasr: Efficient differential testing of automatic speech recognition via text-to-speech. In: 2020 IEEE International Conference on Software Maintenance and Evolution (ICSME). pp. 640–650 (2020). <https://doi.org/10.1109/ICSME46990.2020.00066>
6. Buolamwini, J., Gebru, T.: Gender shades: Intersectional accuracy disparities in commercial gender classification. In: Conference on fairness, accountability and transparency. pp. 77–91. PMLR (2018)
7. Butterworth, S., et al.: On the theory of filter amplifiers. *Wireless Engineer* **7**(6), 536–541 (1930)
8. Calò, A., Arcaini, P., Ali, S., Hauer, F., Ishikawa, F.: Simultaneously searching and solving multiple avoidable collisions for testing autonomous driving systems. In: Proceedings of the 2020 Genetic and Evolutionary Computation Conference. pp. 1055–1063 (2020)
9. Carlini, N., Wagner, D.: Audio adversarial examples: Targeted attacks on speech-to-text. In: 2018 IEEE Security and Privacy Workshops (SPW). pp. 1–7. IEEE (2018)
10. Chen, G., Chen, S., Fan, L., Du, X., Zhao, Z., Song, F., Liu, Y.: Who is real bob? adversarial attacks on speaker recognition systems. In: IEEE Symposium on Security and Privacy (2021)
11. Demirsahin, I., Kjartansson, O., Gutkin, A., Rivera, C.: Open-source Multi-speaker Corpora of the English Accents in the British Isles. In: Proceedings of The 12th Language Resources and Evaluation Conference (LREC). pp. 6532–6541. European Language Resources Association (ELRA), Marseille, France (May 2020), <https://www.aclweb.org/anthology/2020.lrec-1.804>
12. Denton, E., Hutchinson, B., Mitchell, M., Gebru, T., Zaldivar, A.: Image counterfactual sensitivity analysis for detecting unintended bias. In: CVPR 2019 Workshop on Fairness Accountability Transparency and Ethics in Computer Vision (2019)
13. Du, X., Xie, X., Li, Y., Ma, L., Zhao, J., Liu, Y.: Deepcruiser: Automated guided testing for stateful deep learning systems. CoRR **abs/1812.05339** (2018), <http://arxiv.org/abs/1812.05339>
14. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.: Fairness through awareness. In: Proceedings of the 3rd innovations in theoretical computer science conference. pp. 214–226 (2012)
15. Eniser, H.F., Gerasimou, S., Sen, A.: Deepfault: Fault localization for deep neural networks. In: Hähnle, R., van der Aalst, W.M.P. (eds.) Fundamental Approaches to Software Engineering - 22nd International Conference, FASE 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6–11, 2019, Proceedings. Lecture Notes in Computer Science, vol. 11424, pp. 171–191. Springer (2019)

16. Feng, Y., Shi, Q., Gao, X., Wan, J., Fang, C., Chen, Z.: Deepgini: prioritizing massive tests to enhance the robustness of deep neural networks. In: Proceedings of the 29th ACM SIGSOFT International Symposium on Software Testing and Analysis. pp. 177–188 (2020)
17. Galhotra, S., Brun, Y., Meliou, A.: Fairness testing: testing software for discrimination. In: Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2017, Paderborn, Germany, September 4–8, 2017. pp. 498–510 (2017). <https://doi.org/10.1145/3106237.3106277>, <http://doi.acm.org/10.1145/3106237.3106277>
18. Goss, F.R., Zhou, L., Weiner, S.G.: Incidence of speech recognition errors in the emergency department. *International journal of medical informatics* **93**, 70–73 (2016)
19. Guo, Q., Xie, X., Li, Y., Zhang, X., Liu, Y., Li, X., Shen, C.: Audee: Automated testing for deep learning frameworks. In: Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering (ASE). pp. 486–498. ACM (Dec 2020)
20. Hawley, M.S.: Speech recognition as an input to electronic assistive technology. *British Journal of Occupational Therapy* **65**(1), 15–20 (2002)
21. Helmke, H., Ohneiser, O., Mühlhausen, T., Wies, M.: Reducing controller workload with automatic speech recognition. In: 2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC). pp. 1–10. IEEE (2016)
22. Huang, C., Chen, T., Li, S.Z., Chang, E., Zhou, J.L.: Analysis of speaker variability. In: INTERSPEECH. pp. 1377–1380 (2001)
23. Iwama, F., Fukuda, T.: Automated testing of basic recognition capability for speech recognition systems. In: 2019 12th IEEE Conference on Software Testing, Validation and Verification (ICST). pp. 13–24. IEEE (2019)
24. Jain, A., Upreti, M., Jyothi, P.: Improved accented speech recognition using accent embeddings and multi-task learning. In: Interspeech. pp. 2454–2458 (2018)
25. Johnson, D.H.: Signal-to-noise ratio. *Scholarpedia* **1**(12), 2088 (2006)
26. Koenecke, A., Nam, A., Lake, E., Nudell, J., Quartey, M., Mengesha, Z., Toups, C., Rickford, J.R., Jurafsky, D., Goel, S.: Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences* **117**(14), 7684–7689 (2020)
27. Kopald, H.D., Chanen, A., Chen, S., Smith, E.C., Tarakan, R.M.: Applying automatic speech recognition technology to air traffic management. In: 2013 IEEE/AIAA 32nd Digital Avionics Systems Conference (DASC). pp. 6C3–1. IEEE (2013)
28. Kreuk, F., Adi, Y., Cisse, M., Keshet, J.: Fooling end-to-end speaker verification with adversarial examples. In: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 1962–1966. IEEE (2018)
29. Levenshtein, V.I.: Binary codes capable of correcting deletions, insertions, and reversals. In: *Soviet physics doklady*. vol. 10, pp. 707–710. Soviet Union (1966)
30. Livingstone, S.R., Russo, F.A.: The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english. *PloS one* **13**(5), e0196391 (2018)
31. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 30, pp. 4765–4774. Curran Associates, Inc. (2017), <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>

32. Ma, L., Juefei-Xu, F., Zhang, F., Sun, J., Xue, M., Li, B., Chen, C., Su, T., Li, L., Liu, Y., Zhao, J., Wang, Y.: Deepgauge: multi-granularity testing criteria for deep learning systems. In: Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018, Montpellier, France, September 3-7, 2018. pp. 120–131 (2018)
33. Ma, P., Wang, S., Liu, J.: Metamorphic testing and certified mitigation of fairness violations in NLP models. In: Bessiere, C. (ed.) Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020. pp. 458–465
34. Odena, A., Olsson, C., Andersen, D., Goodfellow, I.: Tensorfuzz: Debugging neural networks with coverage-guided fuzzing. In: International Conference on Machine Learning. pp. 4901–4911. PMLR (2019)
35. Pei, K., Cao, Y., Yang, J., Jana, S.: Deepxplore: Automated whitebox testing of deep learning systems. In: Proceedings of the 26th Symposium on Operating Systems Principles, Shanghai, China, October 28-31, 2017. pp. 1–18 (2017)
36. Phillips, A.: Defending equality of outcome. *Journal of political philosophy* **12**(1), 1–19 (2004)
37. Qin, Y., Carlini, N., Cottrell, G., Goodfellow, I., Raffel, C.: Imperceptible, robust, and targeted adversarial examples for automatic speech recognition. In: International conference on machine learning. pp. 5231–5240. PMLR (2019)
38. Ribeiro, M.T., Singh, S., Guestrin, C.: "why should I trust you?": Explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016. pp. 1135–1144 (2016)
39. Ribeiro, M.T., Singh, S., Guestrin, C.: Anchors: High-precision model-agnostic explanations. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018)
40. Ribeiro, M.T., Wu, T., Guestrin, C., Singh, S.: Beyond accuracy: Behavioral testing of NLP models with checklist. In: Jurafsky, D., Chai, J., Schluter, N., Tetreault, J.R. (eds.) Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020. pp. 4902–4912. Association for Computational Linguistics (2020)
41. Sharma, A., Demir, C., Ngomo, A.C.N., Wehrheim, H.: Mlcheck-property-driven testing of machine learning models. arXiv preprint arXiv:2105.00741 (2021)
42. Sharma, A., Wehrheim, H.: Testing machine learning algorithms for balanced data usage. In: 2019 12th IEEE Conference on Software Testing, Validation and Verification (ICST). pp. 125–135. IEEE (2019)
43. Sharma, A., Wehrheim, H.: Automatic fairness testing of machine learning models. In: IFIP International Conference on Testing Software and Systems. pp. 255–271. Springer (2020)
44. Sharma, A., Wehrheim, H.: Higher income, larger loan? monotonicity testing of machine learning models. In: Proceedings of the 29th ACM SIGSOFT International Symposium on Software Testing and Analysis. pp. 200–210 (2020)
45. Soremekun, E., Udeshi, S., Chattopadhyay, S.: Astraea: Grammar-based fairness testing. arXiv preprint arXiv:2010.02542 (2020)
46. Sun, Y., Chockler, H., Huang, X., Kroening, D.: Explaining image classifiers using statistical fault localization. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J. (eds.) Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXVIII. Lecture Notes in Computer Science, vol. 12373, pp. 391–406. Springer (2020)

47. Sun, Y., Wu, M., Ruan, W., Huang, X., Kwiatkowska, M., Kroening, D.: Concolic testing for deep neural networks. In: Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018, Montpellier, France, September 3-7, 2018. pp. 109–119 (2018)
48. Tian, Y., Pei, K., Jana, S., Ray, B.: Deeptest: automated testing of deep-neural-network-driven autonomous cars. In: Proceedings of the 40th International Conference on Software Engineering, ICSE 2018, Gothenburg, Sweden, May 27 - June 03, 2018. pp. 303–314 (2018)
49. Udeshi, S., Arora, P., Chattopadhyay, S.: Automated directed fairness testing. In: Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018, Montpellier, France, September 3-7, 2018. pp. 98–108 (2018)
50. Udeshi, S.S., Chattopadhyay, S.: Grammar based directed testing of machine learning systems. *IEEE Transactions on Software Engineering* (2019)
51. Verma, S., Rubin, J.: Fairness definitions explained. In: 2018 IEEE/ACM International Workshop on Software Fairness (Fairware). pp. 1–7. IEEE (2018)
52. Wang, J., Chen, J., Sun, Y., Ma, X., Wang, D., Sun, J., Cheng, P.: Robot: Robustness-oriented testing for deep learning systems. In: ICSE '21: 43rd International Conference on Software Engineering (2021)
53. Wardat, M., Le, W., Rajan, H.: Deeplocalize: Fault localization for deep neural networks. In: 43rd IEEE/ACM International Conference on Software Engineering, ICSE 2021, Madrid, Spain, 22-30 May 2021. pp. 251–262. IEEE (2021)
54. Weinberger, S.H., Kunath, S.A.: The speech accent archive: towards a typology of english accents. In: *Corpus-based Studies in Language Use, Language Learning, and Language Documentation*, pp. 265–281. Brill Rodopi (2011)
55. Xie, X., Ma, L., Juefei-Xu, F., Xue, M., Chen, H., Liu, Y., Zhao, J., Li, B., Yin, J., See, S.: Deephunter: a coverage-guided fuzz testing framework for deep neural networks. In: Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis. pp. 146–157 (2019)
56. Xie, X., Zhang, Z., Chen, T.Y., Liu, Y., Poon, P.L., Xu, B.: Mettle: a metamorphic testing approach to assessing and validating unsupervised machine learning systems. *IEEE Transactions on Reliability* **69**(4), 1293–1322 (2020)
57. Zhang, J., Harman, M.: "ignorance and prejudice" in software fairness. In: International Conference on Software Engineering. vol. 43. IEEE (2021)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), 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 license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license 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.

