

RESEARCH ARTICLE

Warning: Humans cannot reliably detect speech deepfakes

Kimberly T. Mai^{1,2*}, Sergi Bray^{1,2}, Toby Davies¹, Lewis D. Griffin²¹ Department of Security and Crime Science, University College London, London, United Kingdom,² Department of Computer Science, University College London, London, United Kingdom* kimberly.mai@ucl.ac.uk

Abstract

Speech deepfakes are artificial voices generated by machine learning models. Previous literature has highlighted deepfakes as one of the biggest security threats arising from progress in artificial intelligence due to their potential for misuse. However, studies investigating human detection capabilities are limited. We presented genuine and deepfake audio to $n = 529$ individuals and asked them to identify the deepfakes. We ran our experiments in English and Mandarin to understand if language affects detection performance and decision-making rationale. We found that detection capability is unreliable. Listeners only correctly spotted the deepfakes 73% of the time, and there was no difference in detectability between the two languages. Increasing listener awareness by providing examples of speech deepfakes only improves results slightly. As speech synthesis algorithms improve and become more realistic, we can expect the detection task to become harder. The difficulty of detecting speech deepfakes confirms their potential for misuse and signals that defenses against this threat are needed.

OPEN ACCESS

Citation: Mai KT, Bray S, Davies T, Griffin LD (2023) Warning: Humans cannot reliably detect speech deepfakes. *PLoS ONE* 18(8): e0285333. <https://doi.org/10.1371/journal.pone.0285333>

Editor: Yogan Jaya Kumar, Universiti Teknikal Malaysia Melaka Fakulti Teknologi Maklumat dan Komunikasi, MALAYSIA

Received: February 9, 2023

Accepted: June 30, 2023

Published: August 2, 2023

Copyright: © 2023 Mai et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The data underlying this study are available at DOI: [10.17605/OSF.IO/AKM2P](https://doi.org/10.17605/OSF.IO/AKM2P).

Funding: KM and SB are supported by the Dawes Centre for Future Crime (<https://www.ucl.ac.uk/future-crime/>). KM is supported by EPSRC under grant EP/R513143/1 (<https://www.ukri.org/councils/epsrc/>). SB is supported by EPSRC under grant EP/S022503/1. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Introduction

Adversaries are already using speech deepfakes to commit fraud. In 2020, a bank manager in Hong Kong received a phone call from someone sounding like a company director he had spoken to before [1]. The purported director requested the bank manager to authorize transfers totaling \$35 million. Based on their existing relationship, the bank manager transferred \$400,000 until he realized something was wrong. The bank manager was a victim of an elaborate hoax: fraudsters had used deepfake technology to clone the director's voice. This incident is not isolated. In 2019, the CEO of a UK-based firm was swindled by a speech deepfake of his manager into transferring €220,000 to a Hungarian supplier [2].

Speech deepfakes are artificial voices generated by machine learning models. Due to rapid research progress, it is possible to produce a realistic-sounding clone using only a few audio samples [3]. This development raises the prospect of exploiting speech deepfakes for various criminal activities. Alongside impersonation, criminals may use deepfakes for spear phishing, propagating fake news, and bypassing biometric authentication systems [4–6].

Competing interests: The authors have declared that no competing interests exist.

Existing speech deepfake detection research focuses on developing machine learning systems in the context of voice authentication [7–9]. Comparisons beyond this biometric setting and studies which measure human detection capabilities are sparse [10].

The state of existing research raises questions. Firstly, machine learning systems require large amounts of data for training [11] and are hard to interpret [12]. When analyzing these systems, it is unclear which characteristics distinguish synthesized speech from bona fide. Therefore, knowing what humans use to identify deepfakes could provide a better understanding of how black-box machine learning systems work.

Secondly, focusing on automated biometric authentication does not quantify the threat of other potential criminal applications of speech deepfakes. Understanding the extent of these threats is critical. Multiple studies deem other uses of speech deepfakes as more concerning, such as misleading people through voice impersonations [5, 6]. Experts expect disinformation from deepfakes to erode trust on several levels: towards individuals, organizations, and even societies [13]. Moreover, it is estimated as much as 90% of online content will be synthetically generated by 2026 [14], meaning it will be challenging to moderate what gets produced. Therefore, understanding the risks of speech deepfakes will enable the development of better defenses and regulations to counteract hazards before they occur.

We seek to address these two questions by measuring how well humans distinguish bona fide speech from synthesized speech. We ran an online experiment where individuals listened to bona fide and fake audio clips and attempted to differentiate between them.

We randomly assigned the participants to two configurations. In the first configuration, we presented participants with one audio clip at a time and asked them to decide if the clip was fake. In the second configuration, we presented participants with audio clip pairs containing the same speech (one bona fide and one synthesized) and asked them to identify the synthesized audio.

We ran the experiment in English and Mandarin to understand if listeners used language-specific attributes to detect deepfakes and to observe if deepfake detection is more manageable in one language than another. Finally, we incorporated randomized interventions to evaluate whether familiarizing participants with examples of speech deepfakes boosts detection performance.

Our results suggest the listeners had limited detection capabilities, and performance is similar between languages. Additionally, familiarizing participants improved performance but only to a small extent.

Background

Deepfake media

Deepfakes are synthetic media produced in the likeness of a person. They fall under the field of generative artificial intelligence (AI). Generative AI is a subset of machine learning (ML) algorithms that learn the patterns and characteristics of a dataset [11]. The algorithms use this knowledge to generate synthetic content similar to the original data. Deepfakes specifically refer to the outputs of generative AI that resemble humans and their actions.

Deepfake media occur in different modalities:

1. **Images:** This modality contains static faces generated using varying techniques. These techniques include:
 - **Generation from scratch:** A generative adversarial network [15] or diffusion model [16] synthesizes a fictional identity.

- **Morphing:** Blending similar-looking faces to produce an identity containing the characteristics of the sources [17].
 - **Swaps:** A source face replaces the target in a different image [18].
2. **Video:** This modality features individuals performing actions. Currently, the techniques used to synthesize videos are similar to those used in images. Image synthesis techniques are applied at a frame level and stitched together to form a video.
 3. **Speech:** This modality conveys information in a manner that sounds like a genuine person's voice. Although audio can refer to general sound synthesis, the terms “audio”, “speech”, and “voice” deepfakes are used interchangeably in academic literature. We refer to them as “speech deepfakes” for consistency throughout the text.

In addition, deepfakes are either produced in the likeness of a known identity (**targeted**) or do not resemble a familiar identity (**untargeted**). For example, we can categorize video deepfakes of politicians as targeted. Conversely, a generic face created from scratch and not conditioned to resemble a specific individual is untargeted.

We refer the reader to Zhang (2022) [19] for further information on deepfake terminology. As fewer works focus on speech deepfakes, we concentrate on this modality.

Synthesizing speech

Generative models are often used to synthesize speech. Speech synthesizers which use generative models follow a common framework:

1. **Data collection:** Several audio recordings of the speaker are collected.
2. **Pre-processing:** The audio recordings are converted into alternative formats to make it easier for the generative model to work with them.
3. **Training:** Processed audio recordings are fed to the generative model to learn the patterns and characteristics of the data. The trained model is often called a vocoder.

The frameworks often include text-to-speech (TTS) modules to make it easier to generate speech. The generative model also sees text transcriptions corresponding to the audio recordings in this setting.

We depict a visualization of this framework in Fig 1.

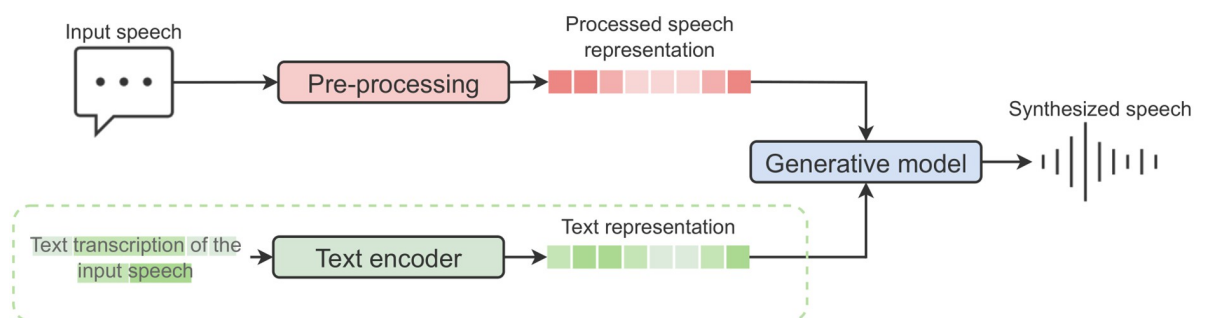


Fig 1. Diagram of a typical generative speech synthesis model.

<https://doi.org/10.1371/journal.pone.0285333.g001>

Related work on human deepfake detection capabilities

Most deepfake detection studies which examine human performance use visual media. When faced with deepfake content of politicians, participants rely on contextual knowledge in the form of political literacy to identify spoofs [20, 21].

Removing such background knowledge makes the detection task more difficult. In the context of images, multiple studies show humans do not perform much better than chance [22, 23]. There is no improvement when evaluating videos either [24–26]. Moreover, these studies suggest humans are overconfident in their deepfake detection abilities [25].

Several of the above studies examine if interventions can boost detection performance. However, the effectiveness of these interventions is debatable. Bray et al. [22] familiarized participants by showing examples of deepfakes before the main task. The authors also drew participants' attention to errors often present in bogus images. Although these interventions improved deepfake detection performance, they also increased overall skepticism as a higher proportion of bona fide images were falsely classified. One could also note that pointing out errors biases the participants and prevents them from independently identifying the tell-tale characteristics of deepfakes. Köbis et al. [25] presented interventions by informing participants about the impact of deepfakes and rewarding correct guesses. Neither intervention led to improved performance.

In contrast, other authors found interventions derived from ML model outputs improve detection. Tahir et al. [26] produced educational material containing indicators of bogus images with the assistance of ML interpretability tools. The authors found detection performance improved compared to the initial control group. However, a recent study [27] contests the reliability of these tools, as the authors show it is possible to manipulate the output visualizations. Groh et al. [24] allowed participants to amend their choices after viewing the predictions of an ML model. This form of cooperation improved results significantly.

There are fewer studies that examine how well humans can detect speech deepfakes. Watson et al. [28] presented eight clips to college students and asked them to decide whether the clips were real or fake. They found that shorter clips were easier to identify. However, the sample size of their study was small and skewed towards a younger, college-educated demographic. The ASVspoof challenge organizers ran an experiment with a larger sample size [29]. They asked 1,145 participants to imagine they worked in a call center and decide whether the incoming calls were spoken by humans or by an AI. However, the experiment is limited to the speaker verification setting.

Müller et al. [30] ran a game where 378 participants competed against an ML model to decide if an audio clip was fake. Similarly to Groh et al. [24], they found that feedback from the ML model improved human performance. In their experiment, Müller et al. [30] found that the difference between human and AI accuracy was about 10%. However, their study only used English-language clips, only presented one audio clip to participants at a time, and did not collect information about participant confidence.

We summarize the relevant literature in Table 1. We note that Barari et al. [20] mention fake speech stimuli in their analysis. However, they used actors to create the speech instead of generative AI. Therefore we excluded this from our analysis.

Materials and methods

Our research questions were as follows:

1. How well can humans detect speech deepfakes?
2. Are there differences in detection capabilities depending on the language?

Table 1. Summary of related literature measuring human capabilities to detect deepfakes.

Modality	Year	Author	Deepfake stimuli
Image	2021	Nightingale & Farid [23]	Faces generated using StyleGAN2 [31]
	2022	Bray et al. [22]	Faces generated using StyleGAN2
Video	2021	Barari et al. [20]	Face-swap videos of politicians
	2021	Groh et al. [24]	Face-swap videos from the Deepfake Detection Challenge dataset [32]
	2021	Köbis et al. [25]	Face-swap videos from the Deepfake Detection Challenge dataset
	2021	Tahir et al. [26]	Face-swap videos from Celeb-DF [33], FaceForensics++ [34] and DeepFaceLab [35]
	2022	Appel & Prietzel [21]	Face-swap videos of politicians
Speech	2020	Wang et al. [29]	Spoofed utterances generated from TTS and voice conversion systems used in ASVspoof2019
	2021	Watson et al. [28]	Audio clips generated using MelGAN [36]
	2022	Müller et al. [30]	Spoofed utterances from the ASVspoof2019 dataset [29]

<https://doi.org/10.1371/journal.pone.0285333.t001>

3. Do interventions in the form of examples and added context improve detection performance?

Our experiments focused on human performance rather than the performance of automated detectors. Through this setup, we could quantify the threat of speech deepfakes when humans interact with them.

Stimuli

Bona fide stimuli. We collected bona fide stimuli from two publicly available datasets. Both datasets consist of one female speaker reading generic sentences. The datasets also include text transcriptions of the audio. We chose such datasets to prevent participants from using external cues for the detection task.

We used LJSpeech [37] as the English dataset. The dataset consists of a speaker reading passages from seven non-fiction books, varying between one and ten seconds in length.

We used the Chinese Standard Mandarin Speech Corpus (CSMSC) [38] as the Mandarin dataset. The corpus used in the dataset aims to cover Mandarin tones and prosody as comprehensively as possible.

Deepfake stimuli. To create the deepfake stimuli, we used publicly available TTS models trained on the two datasets [39]. In particular, we chose pre-trained VITS models [40]. VITS is an end-to-end TTS model which combines the data pre-processing and vocoder into a single framework.

We randomly selected 50 sentences from the validation split of the two datasets to create the deepfakes. We used the same sentences for our bona fide stimuli. Therefore, we had 100 clips in total.

The validation split consists of samples not used for training the ML models. It is good practice to use unseen data because it indicates how well a trained model generalizes. Consequently, the resultant generated audio should contain artifacts that we would expect to hear from ML models. These artifacts might serve as informative features for distinguishing deepfakes. If we used samples previously seen during training, the model could potentially mimic the samples perfectly and would not contain representative artifacts.

Procedure

The setup for the English and Mandarin experiments was identical. We randomly assigned participants to two configurations: unary and binary. In both configurations, we asked

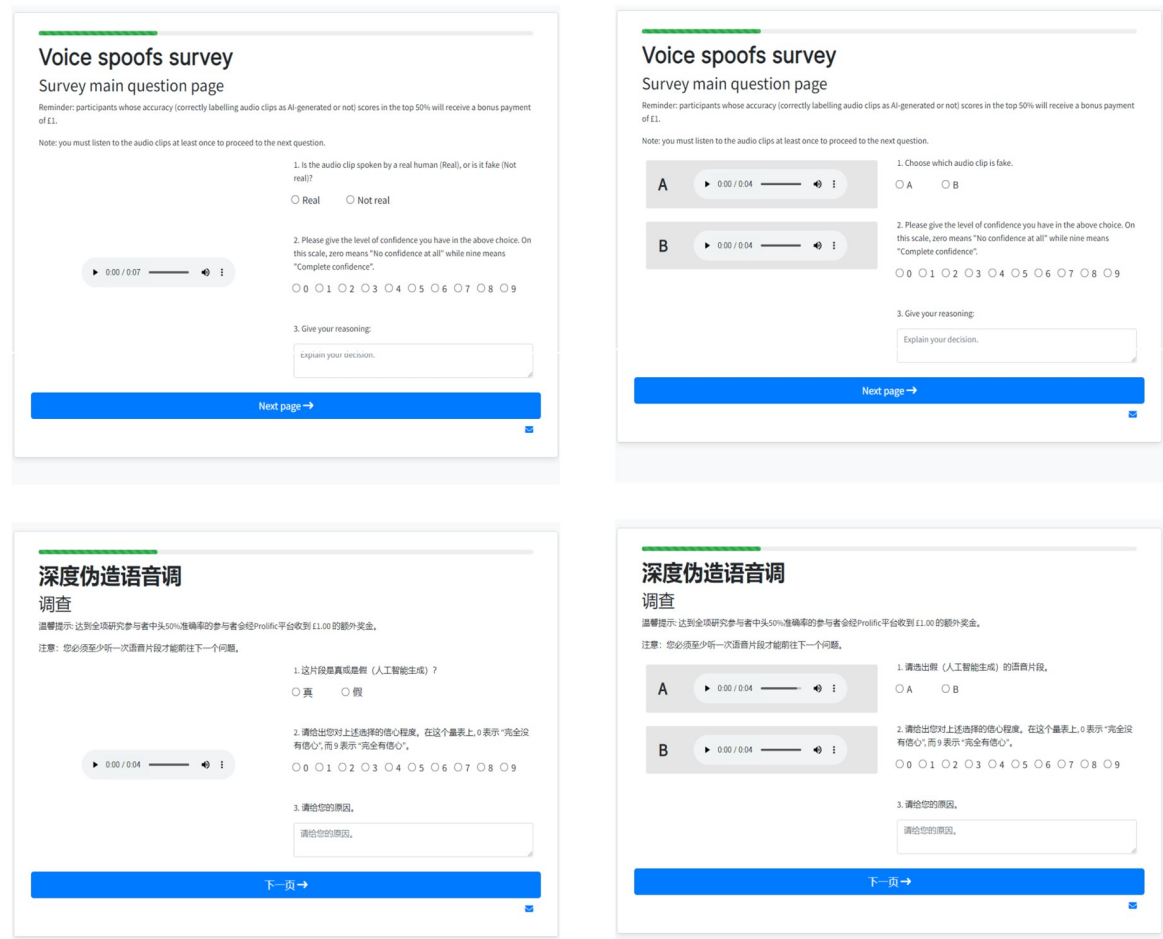


Fig 2. Screenshots of the task interface.

<https://doi.org/10.1371/journal.pone.0285333.g002>

participants to rate the confidence of their choice on a ten-point Likert scale and provide free-form text justifications. Participants were allowed to listen to the clips as often as they liked. We did not give feedback to the participants to inform them if their choices were correct. Compared to the setups described in Müller et al. and Groh et al. [24, 30], the lack of feedback creates a more realistic scenario. When encountering speech deepfakes in the wild (for example, through fraudulent calls), humans do not know that the voices are fake. We include screenshots of the two configurations in Fig 2.

Unary. We presented 20 randomly chosen distinct clips to each participant, each on separate pages. Participants listened to approximately an equal number of bona fide and synthesized clips, but we did not inform them about the proportion. We tasked the participants with deciding whether the clip they heard was real or fake.

Binary. We presented 20 randomly chosen clip pairs (labeled ‘A’ and ‘B’) comprising the same spoken sentence. Each pair contained a clip uttered by the human speaker and a clip produced by VITS. We randomized the order of the fake and real clips and asked the participants to decide which clip was fake. We included this scenario to see if contextual information helped detection.

Familiarization treatment. In addition to the two configurations, we randomly assigned half of the participants to a familiarization treatment group. We included the treatment to verify the existing literature and understand if humans could be trained to detect deepfakes like an ML model. We showed participants in the treatment group five deepfake utterances before commencing the main detection task. We informed the participants that these examples were synthesized and allowed them to listen to the clips multiple times. These clips were distinct from the stimuli used in the main task.

For participants in the control group, we gave them a filler task. In this task, we asked participants to list potential applications of synthesized speech and to provide their opinion about whether synthesized audio will positively or negatively impact society.

Participants

We recruited participants via the Prolific platform. We filtered for participants fluent in English and Mandarin, as fluency affects detection performance [30]. We paid participants at a rate of £7.25 per hour. To encourage more thoughtful responses, we informed participants they could receive a £1.00 bonus if their detection scores were in the top 50%. Overall, we recruited 529 participants. The mean age was 28.9 years old, and 50.6% identified as male. [Table 2](#) contains a more detailed breakdown of the demographics by treatment group.

Ethics statement

The study was reviewed and exempted by the Department of Security and Crime Science's ethics board at University College London. All participants were notified about the purpose of the study and were over the age of 18. Prior to participating, the participants were asked to tick a series of checkboxes to provide informed written consent.

Benchmarking against automated deepfake detectors

To compare the performance of the human participants to automated methods, we trained two artificial neural networks which specialized in detecting speech deepfakes. The two networks used an LFCC-LCNN architecture [41]. LFCC-LCNNs convert raw audio waveforms into two-dimensional representations. They learn by seeing bona fide and deepfake samples and are rewarded for correctly classifying a sample's authenticity. The ASVspoof 2021 challenge used LFCC-LCNNs as baseline models for spoof detection [9]. Hence, they are a reasonable benchmark for our experiments. For more detail about the top-performing speech deepfake detection architectures, we refer the reader to the article summarizing ASVspoof 2021 [9].

We used two versions for each language:

Table 2. Number of participants by group.

Group	English			Mandarin		
	<i>n</i>	<i>Age (SD)</i>	<i>Male (%)</i>	<i>n</i>	<i>Age (SD)</i>	<i>Male (%)</i>
Unary no familiarization	76	26.8 (8.1)	55.2	65	31.0 (10.4)	44.6
Unary familiarization	65	26.7 (7.3)	56.9	54	31.4 (8.7)	44.4
Binary no familiarization	60	27.5 (7.2)	53.3	70	31.8 (9.0)	48.8
Binary familiarization	80	27.4 (7.3)	57.5	59	29.1 (8.5)	39.6
Overall	281	27.1 (7.5)	55.8	248	30.9 (9.2)	44.5

<https://doi.org/10.1371/journal.pone.0285333.t002>

1. In-domain: We trained the networks using the training split of LJSpeech and CSMSC as bona fide samples and created deepfakes by passing the sentences of the training splits through VITS.
2. Out-of-domain: We trained the Mandarin network with FAD [42], another Mandarin-language dataset. We used the pre-trained ASVspoof network [43] for English-language evaluation.

Artificial neural networks are known to perform well when evaluating against samples similar to those seen during training. However, their performance often drops when encountering different examples [44], even if they are in the same language. These differences can be subtle to a human listener and include changes in the speaker's identity or environment. Therefore, we introduce the out-of-domain version for a fairer comparison with human performance, especially as it is unlikely that the participants in our study recognize the LJSpeech and CSMSC identities.

Results

Overall performance

Fig 3 summarizes human performance across all of the different groups. We provide breakdowns of the classification choices in Tables 3 and 4, which aggregate the English and Mandarin results. We completed the analysis using the SciPy [45] and statsmodels [46] Python packages. For further details, the Supporting information contains results per stimulus.

Participants made the correct classifications 70.35% of the time in the unary scenario. They were better at identifying deepfakes (73% accuracy). In comparison, participants correctly identified bona fide examples 67.78% of the time. We speculate the high number of misclassified bona fide samples is partly due to increased skepticism, as participants were aware of the presence of deepfakes through the task briefing. This behavior aligns with observations in Bray et al. [22].

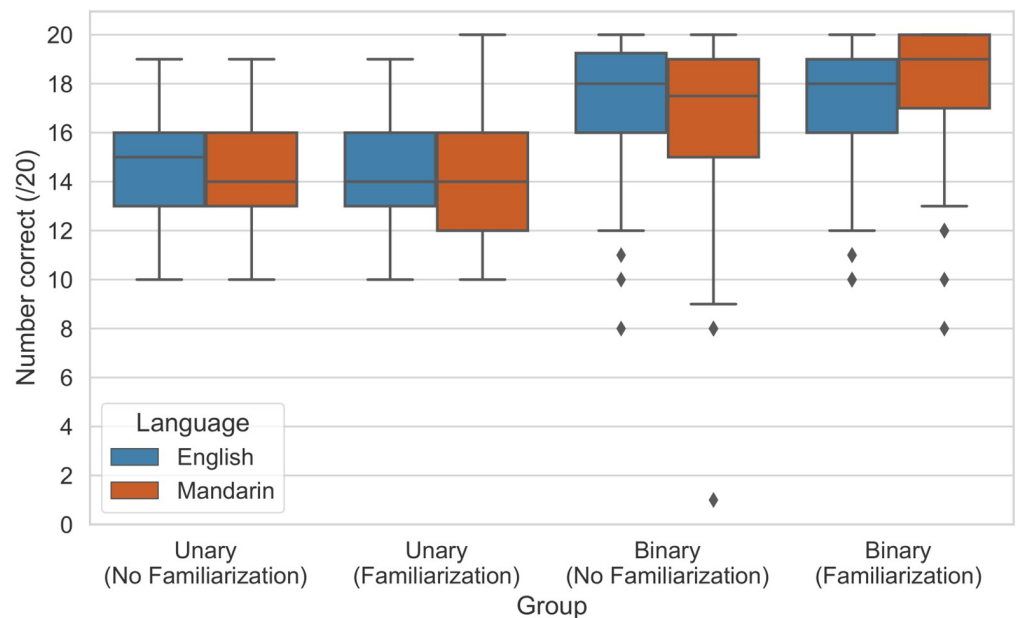


Fig 3. Box plot summarizing human performance across the different groups.

<https://doi.org/10.1371/journal.pone.0285333.g003>

Table 3. Confusion matrix for the unary group responses.

		Predicted class	
		<i>Real (2,442)</i>	<i>Fake (2,678)</i>
True class	<i>Real (2,598)</i>	1,761	837
	<i>Fake (2,522)</i>	681	1,841

$n = 5,120$.

Overall accuracy = 70.35%.

Reals correctly identified = 67.78%.

Fakes correctly identified = 73.0%.

<https://doi.org/10.1371/journal.pone.0285333.t003>

Table 4. Confusion matrix for the binary group responses.

		Predicted class	
		<i>Real</i>	<i>Fake</i>
True class	<i>Real</i>	-	-
	<i>Fake (5,380)</i>	775	4,605

True *real* class labels are not defined in this scenario as participants were asked to choose the fake clip every time.

Overall accuracy is equivalent to fakes correctly identified = 85.59%.

<https://doi.org/10.1371/journal.pone.0285333.t004>

Performance improved under the binary scenario. Participants correctly recognized the deepfake audio in 85.59% of trials. However, the binary setup represents an unrealistic scenario. Even if the speaker's identity is known, reference utterances containing the same speech as the test clip we would like to evaluate are unlikely to be available.

Measuring the effects of interventions

We follow a similar approach to Groh et al. [24] to disentangle the effects of each intervention on performance. We transformed the correct/incorrect results into continuous values by weighting each participant's decision with their provided confidence scores.

The ten-point confidence scale participants completed serves as the mapping function. The lowest score of 0 signals that the participant's choice is a guess, so their confidence in making the right decision corresponds to 50%. In contrast, the highest score of 9 corresponds to 100% belief.

The resulting transformed scores depended on whether the participants made the correct classification. For example, if the participant rated their confidence as 7, this maps to a belief of 88%. If they make the right decision, the adjusted score is 0.88. Conversely, if they make the wrong decision, we subtract the value from 1, resulting in an adjusted score of 0.12.

The revised scores also enable fairer comparisons with the automated deepfake detectors, which output scores between 0 and 1 when evaluating examples. We refer to the revised scores as accuracy scores for the remainder of the text. We also rescale the scores to percentages.

After transforming the results, we analyzed the effects of different interventions on the accuracy scores of participants on each audio clip using linear regression. In addition to language, familiarization and binary intervention, we analyzed the impact of the clip duration. Table 5 outlines the results at the overall, unary and binary levels.

Reference audio helps with deepfake detection. The linear regression results indicate the improvement gained from the binary scenario is statistically significant ($p < 0.001$).

Table 5. Linear regression results of interventions on confidence-scaled accuracy.

Independent variable	Dependent variable: Confidence-scaled accuracy		
	All (SD)	Unary (SD)	Binary (SD)
Constant	43.742*** (1.897)	46.394*** (2.791)	71.217*** (2.530)
Mandarin ^a	1.790 (1.404)	1.477 (1.882)	2.152 (2.118)
Familiarization	3.840*** (1.191)	3.758** (1.571)	3.854** (1.802)
Clip length	0.797*** (0.209)	0.375 (0.358)	1.168*** (0.230)
Binary intervention	29.830*** (1.186)	-	-
Observations	10,500	5,120	5,380
R ²	0.165	0.003	0.010
Adjusted R ²	0.165	0.002	0.009
F-Statistic	171.102***	2.455*	11.794***

^aDummy variable indicating which language was used in the task. 1 = Mandarin, 0 = English.

* $p < 0.1$.

** $p < 0.05$.

*** $p < 0.01$.

<https://doi.org/10.1371/journal.pone.0285333.t005>

Consequently, the results suggest contextual information via reference audio is beneficial for uncovering quirks in synthesized speech.

Training humans to detect deepfakes only helps slightly. The familiarization treatment increases detection accuracy by 3.84% on average ($p = 0.001$). This effect is also present in the unary and binary regression results, improving accuracy by 3.76% ($p = 0.017$) and 3.85% ($p = 0.032$), respectively. However, incorporating familiarizations equates only to an accuracy slightly above chance (52.31%) in the unary setting for the mean clip length (5.76 seconds), *ceteris paribus*.

It is equally challenging to detect deepfakes in Mandarin and English. Fig 3 shows that performance in English and Mandarin is comparable across the different treatment groups. This observation is supported by Table 5, which shows Mandarin-speaking participants only outperform their English counterparts by 1.79%, and this effect is not statistically significant ($p = 0.202$).

Shorter speech deepfakes are not easier to identify. As our stimuli varied from 2 to 11 seconds, we included clip length in the regression to verify whether it is easier to discriminate shorter clips. Our results suggest clip length has a negligible impact on accuracy, improving performance by only 0.80% for each additional second. Our scatter plot (Fig 4) supports this and shows no relationship between the two variables. These findings conflict with Watson et al. [28], who suggest it is easier to identify shorter deepfakes.

Analyzing performance against time

In addition to analyzing the treatment effects, we examine whether the hypothesis of spending more time on the task improves performance.

Listening to the clips more frequently does not aid detection. We recorded the number of times participants clicked on each audio clip and compared the values to accuracy. As shown in Fig 5, there is no relationship between the two variables ($\rho = -0.05$, $p < 0.001$).

Spending more time on the task also does not affect performance. Similar to the above analysis, we compared the time taken to complete the entire task to the total number of clips correctly identified. Fig 6 does not indicate a relationship between the two variables ($\rho = 0.10$,

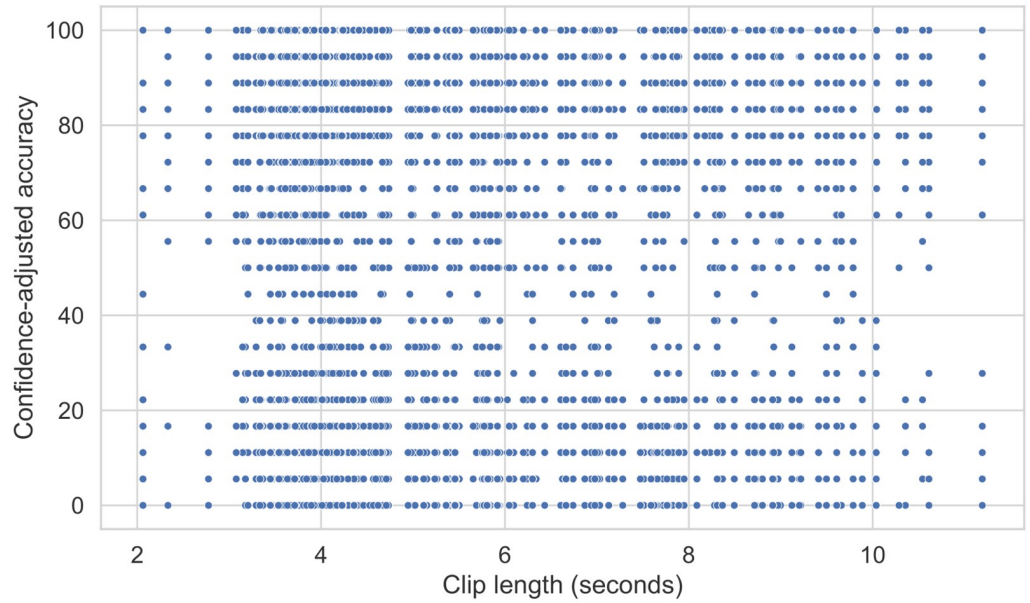


Fig 4. Scatter plot showing the relationship between clip length and confidence-scaled accuracy.

<https://doi.org/10.1371/journal.pone.0285333.g004>

$p = 0.018$), suggesting investing more time to complete the task does not improve performance.

Participants do not get better throughout the task without explicit feedback. To understand whether participants improved as they saw more examples and progressed further in the task, we calculated the number of correct responses per question number. If so, we would expect more correct answers in question 20 compared to question 1. Fig 7 illustrates the

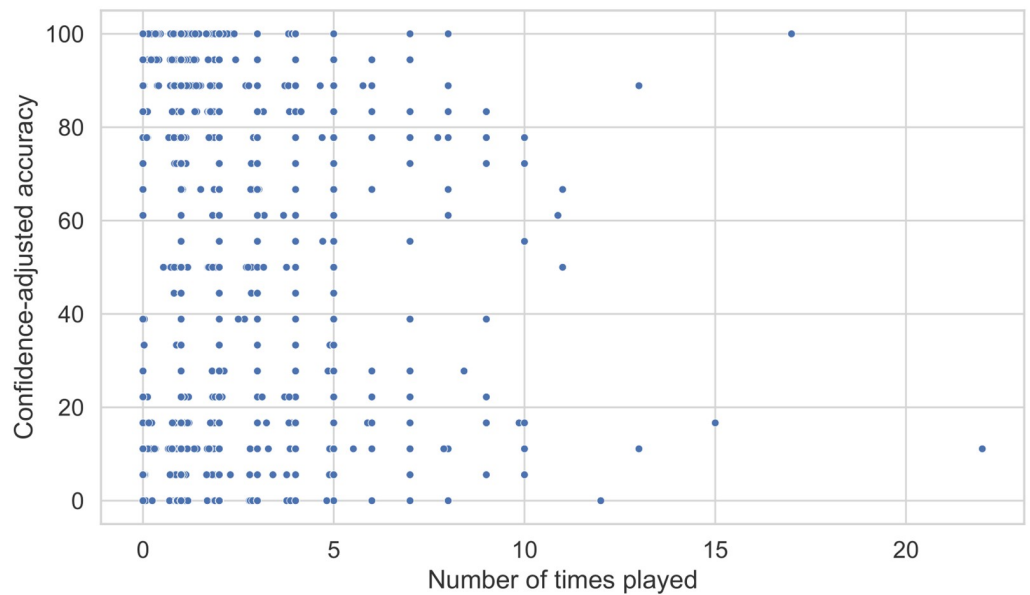


Fig 5. Scatter plot showing the relationship between the number of times played and confidence-scaled accuracy.

<https://doi.org/10.1371/journal.pone.0285333.g005>

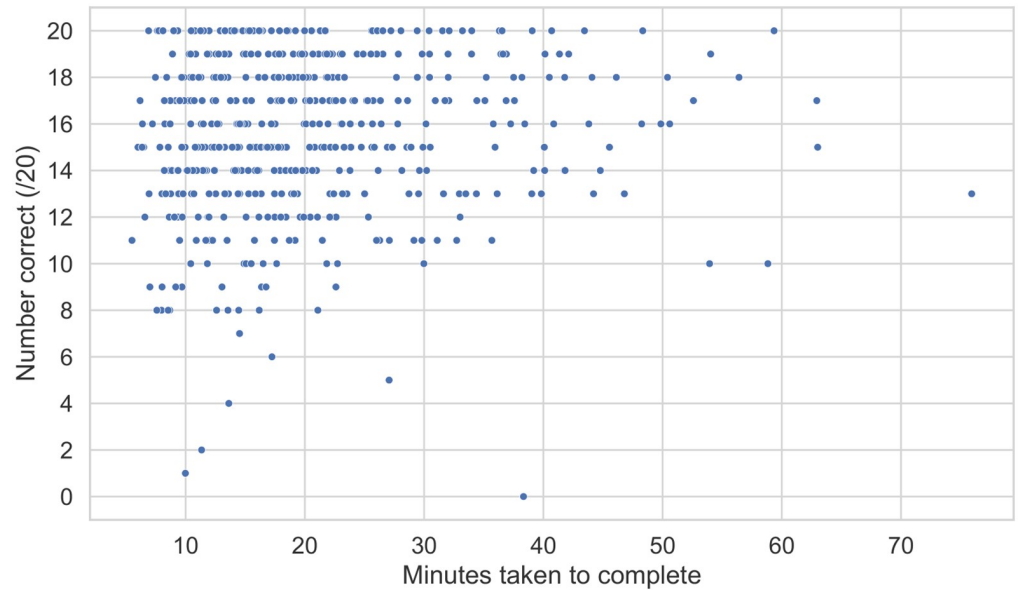


Fig 6. Scatter plot showing the relationship between minutes taken to complete and correctness scores.

<https://doi.org/10.1371/journal.pone.0285333.g006>

resulting histogram. The histogram shows performance is relatively stable across the questions. This observation indicates participants do not improve throughout the task unless they have explicit feedback, as examined by Groh et al. [24] and Müller et al. [30]. We quantitatively verified the result by conducting a one-way chi-squared hypothesis test against the uniform distribution, which was not statistically significant ($\chi^2 = 6.19, p = 0.997$).

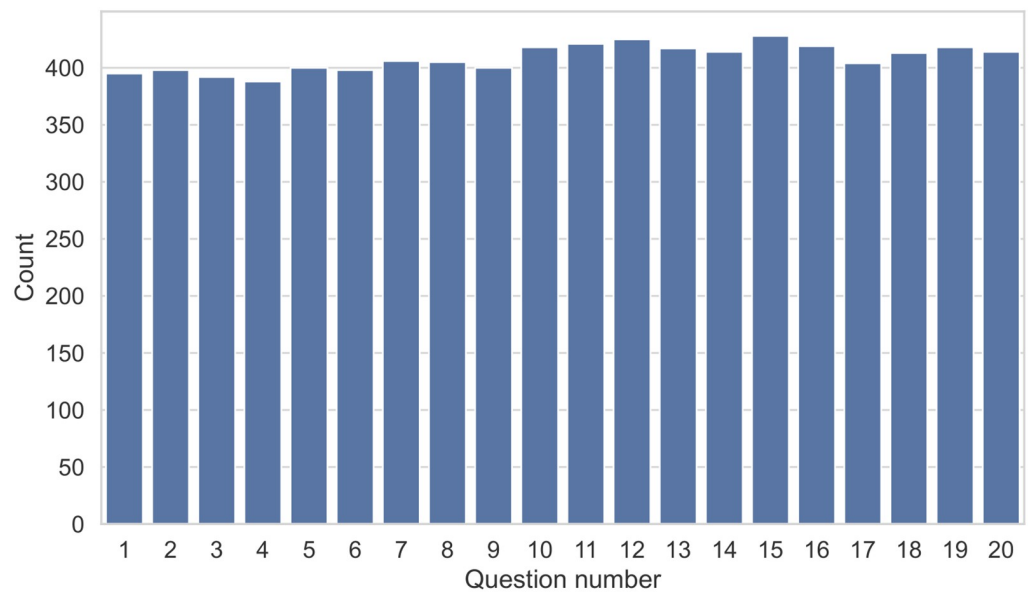


Fig 7. Histogram of correct responses across question number.

<https://doi.org/10.1371/journal.pone.0285333.g007>

Comparing human performance to automated detectors

The following section compares human performance to automated deepfake detectors. For comparability, we use commonly-reported performance metrics found in ML literature.

- **Receiver operating characteristic (ROC):** These plots represent discriminatory ability. They compare true positive rates against false positive rates at different thresholds.
- **The area under the receiver operating characteristic (AUROC):** This score summarizes ROCs into a single value. 50% AUROC indicates all predictions are guesses, whereas 100% AUROC means perfect discrimination between bona fides and deepfakes in all trials.
- **Equal error rate (EER):** This describes the point on ROCs where the true positive and false positive rates are equal.

Fig 8 displays the AUROC and EER scores. We include only the unary scenario in this analysis as the inference setup between humans and automated detectors is more comparable. Both evaluate one clip at a time. We aggregated the English and Mandarin results as we observed similar results.

Human performance is less sensitive to unknown conditions compared to automated detectors. The no familiarization (AUROC = 73.83%) and familiarization curves (AUROC = 75.54%) confirm humans performed better than chance. The curves also support the linear regression result. Showing participants examples of deepfakes only had a minute impact on performance. However, performance was quite unreliable: on average, humans incorrectly classified clips a quarter of the time. Humans underperformed the in-domain automated detectors, which had perfect discrimination ability (AUROC = 100% for both languages). However, out-of-domain detectors often incorrectly classified bona fides as deepfakes (AUROC = 25.31%). Based on this behavior, humans are more robust to unknown factors, such as speaker identity.

Crowd speech deepfake detection is comparable to the top-performing automated detectors. Per Groh et al. [24], we averaged participants' accuracy scores per clip to calculate the crowd-sourced responses. Like the results observed with video stimuli [24], crowd performance is on par with the in-domain detector. However, the benefit of familiarizing participants dissipates when averaging responses. The crowd no familiarization and crowd familiarization AUROCs are similar at 95.51% and 94.04%, respectively.

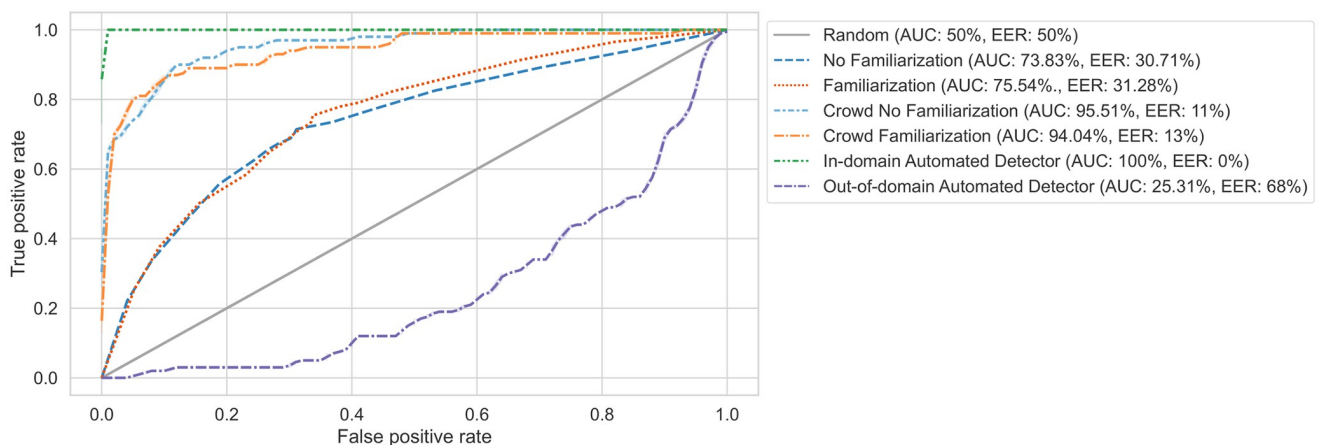


Fig 8. Receiver operator curves under the unary scenario.

<https://doi.org/10.1371/journal.pone.0285333.g008>

Freeform text analysis

To understand how participants assessed the genuineness of audio clips, we analyzed their freeform text responses. We grouped responses by language, clip authenticity, and whether participants made the correct choice. We then created word clouds using tf-idf weightings. Tf-idf measures the importance of a word within a document compared to a collection of documents to account for frequently appearing words [47]. Figs 9 and 10 show the English and Mandarin word clouds.

Participants referred to the same characteristics regardless of whether they made the correct decisions. For example, in Fig 9, participants who correctly classified bona fide utterances as legitimate (in the top left of Fig 9) mentioned pauses, tone and intonation. However, participants who incorrectly categorized bona fide utterances as fake (top right of Fig 9) also referred

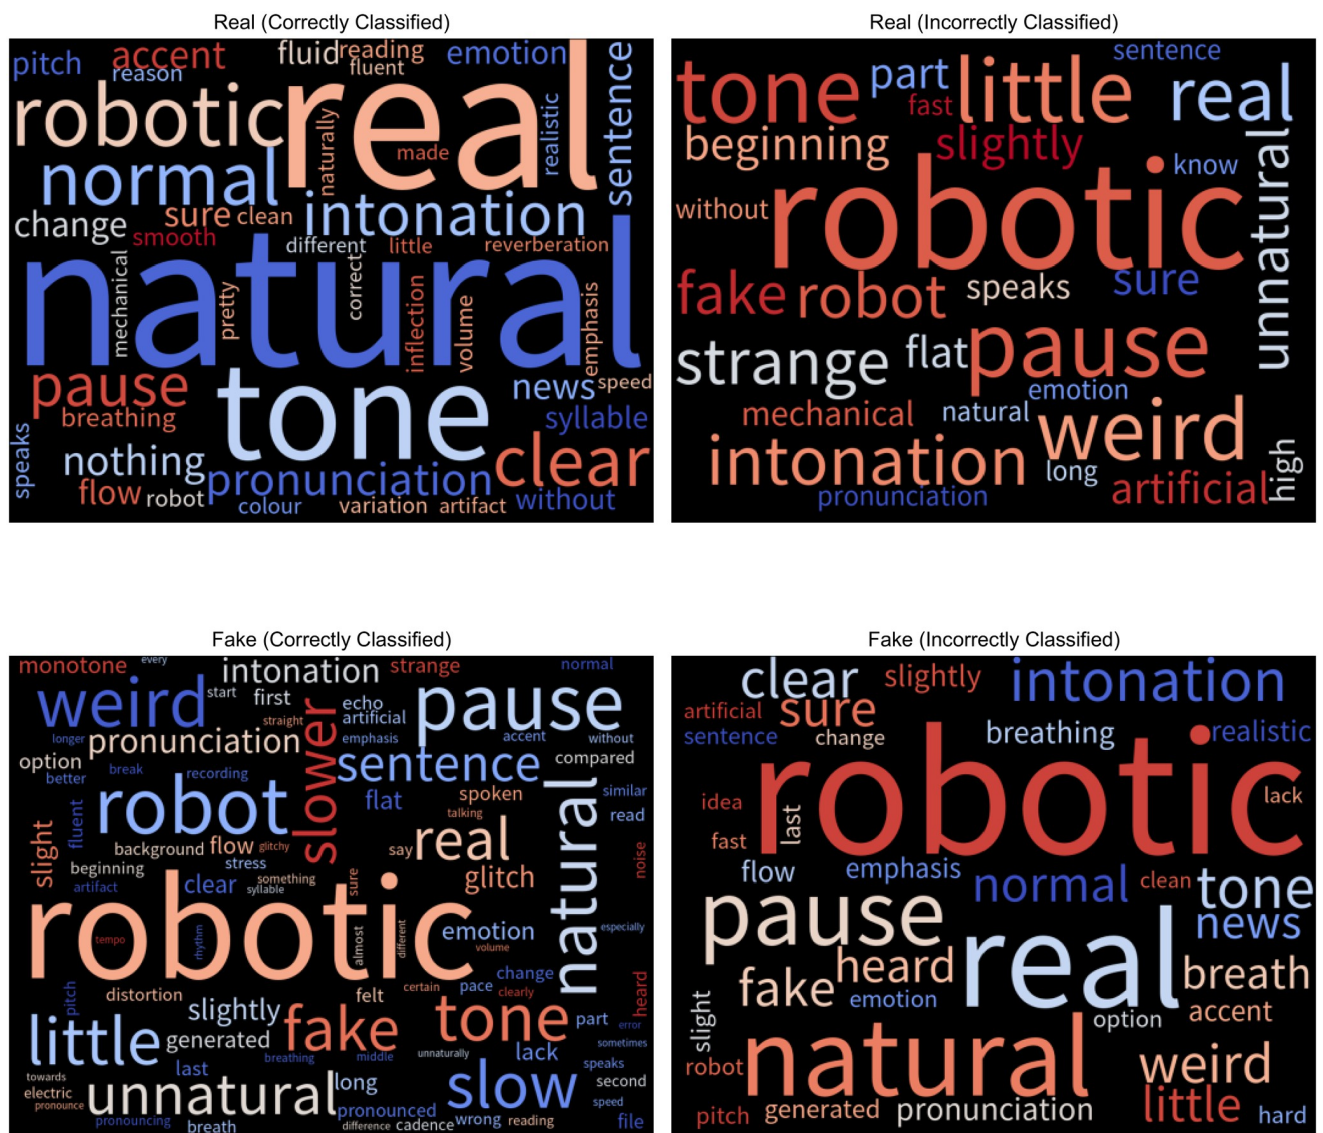


Fig 9. Word clouds containing justifications for the English-language clips.

<https://doi.org/10.1371/journal.pone.0285333.g009>

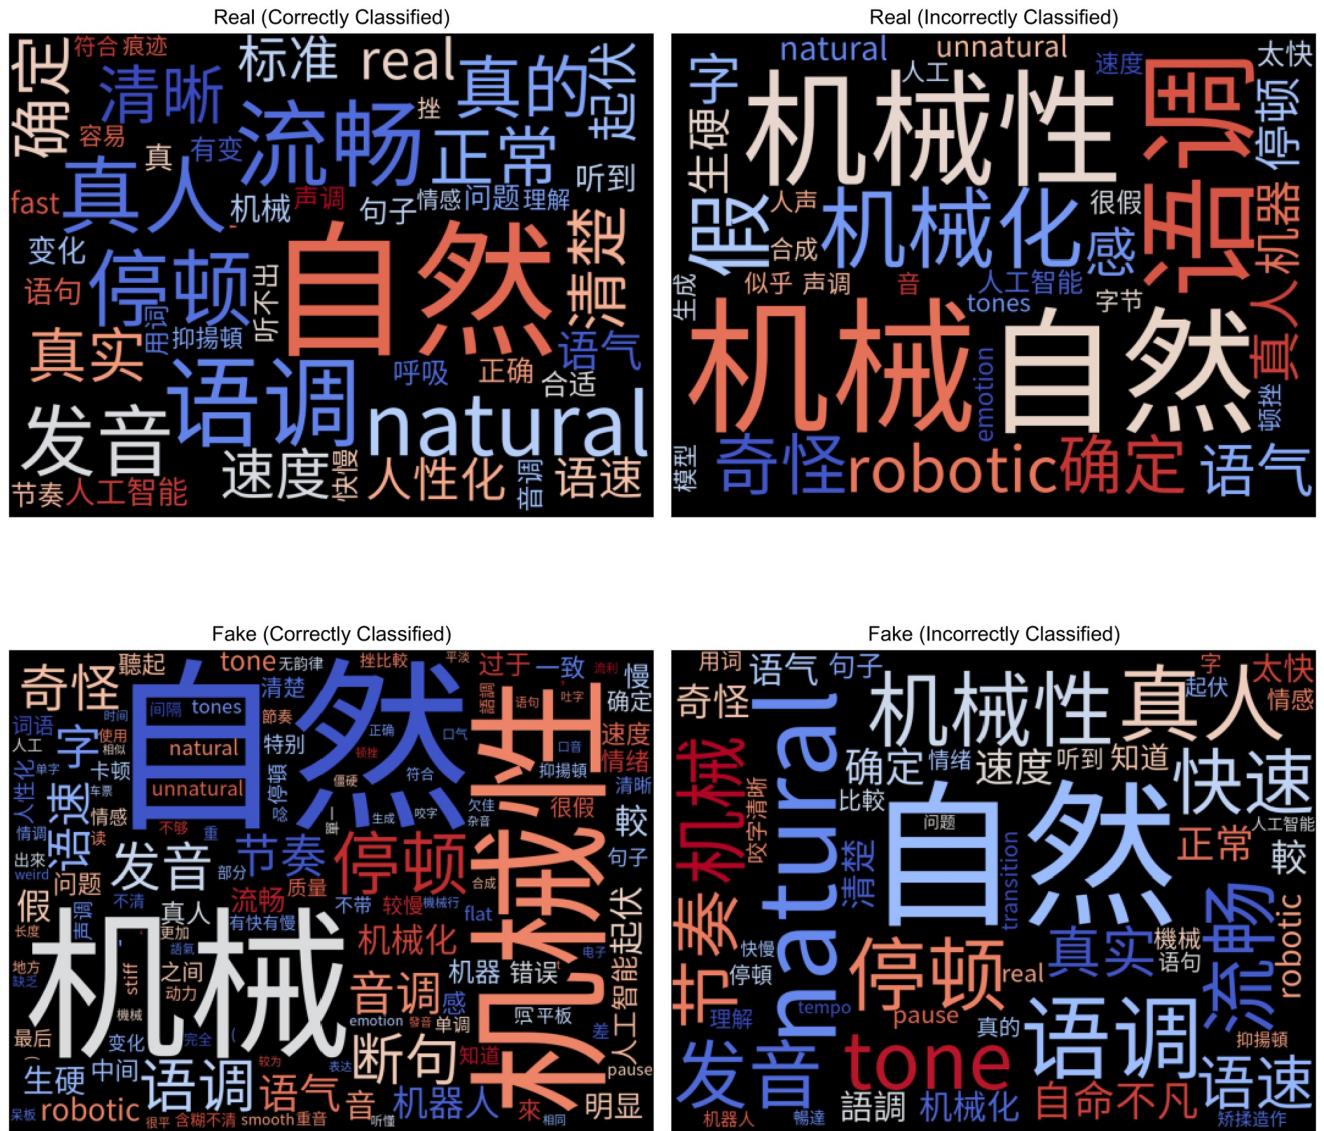


Fig 10. Word clouds containing justifications for the Mandarin-language clips. Note participants for the Mandarin tasks provided justifications in both Mandarin and English.

<https://doi.org/10.1371/journal.pone.0285333.g010>

to these exact attributes. We compared responses by the actual label of the clips and whether participants made the correct response. We did not find substantial differences between these segments. Therefore, automated detectors that incorporate these human characteristics would produce limited improvements. We observed this activity in both English and Mandarin. They tended to rely on intuition to make classifications, referring to naturalness (自然) and robotic (机械) sounds. Beyond intuition, English and Mandarin participants also commonly referenced pauses (停顿), intonation (语调), pronunciation (发音), and speed (速度).

Regarding differences between languages, there were more references to breathing among the English-speaking participants. In contrast, Mandarin-speaking participants mentioned the speaker’s cadence (节奏), pacing between words (断句), and fluency (流畅). This result may be due to differences in timing properties between the two languages. English is stress-timed, while Mandarin is syllable-timed [48].

Limitations

Although our setup enabled comparison with automated detectors, it does not necessarily reflect more realistic scenarios where a listener may encounter speech deepfakes.

Firstly, the balance of deepfakes we presented in our experiment does not reflect the proportion that occurs in the wild. Participants were equally likely to encounter deepfakes as bona fides in the task. However, AI-generated content (including the use of deepfakes for nefarious purposes) is still rare for now. In addition, we could expect participants to be much more attentive to the occurrence of deepfakes as we informed them about the nature of the task.

Moreover, we minimized contextual information in our stimuli. For example, we do not examine situations where the listeners' contextual knowledge (such as awareness of the speaker's identity, emotional status, the number of parties in a conversation, or political affiliations) may have informed their decisions. These aspects may be relevant to typical use cases where speech deepfakes may arise, such as false news propagation [5]. Future work could look at exploring how these characteristics influence detection.

Additionally, we asked participants in both languages to listen to utterances purporting to originate from a single female speaker. Given that age and gender of speakers influence speech perception [49, 50], future work could consider how varying speaker identity affects deepfake detection performance.

To generate our deepfake stimuli, we used an older approach which is not necessarily illustrative of the state-of-the-art speech synthesis algorithms. Although our results indicate how well humans can detect speech deepfakes generated with limited-computational resources, they may not faithfully reflect performance under the most current conditions.

Discussion

Humans can detect speech deepfakes, but not consistently. They tend to rely on naturalness to identify deepfakes regardless of language. As speech synthesis algorithms improve and become more natural, it will become more difficult for humans to catch speech deepfakes.

Although there are some differences in the features that English and Mandarin speakers use to detect deepfakes, the two groups share many similarities. Therefore, the threat potential of speech deepfakes is consistent despite the language involved.

It will be easier for adversaries to generate more deepfakes as the computational barrier for synthesizing data lowers. More deepfakes in the wild will have a knock-on effect. Adversaries will have more opportunities to scale their operations, particularly for disinformation such as impersonations and spear phishing [6].

Ultimately, the battle between deepfake creation and detection is an arms race [51]. How can we defend against falling prey to deepfake trickery? Our binary scenario shows that comparing against reference audio is helpful if we know the speaker's identity. However, we do not always have this information.

Increasing awareness by showing people examples of deepfake audio has a limited effect, as demonstrated by our familiarization results. Spending more time evaluating the clips does not seem to help either.

To summarize, attempting to improve human detection capabilities is unrealistic. We show that even in a controlled environment where the task is easier (participants are aware of the presence of speech deepfakes and the deepfakes are not created using state-of-the-art speech synthesizers), deepfake detection is not high. Our results suggest the need for automated detectors to mitigate a human listener's weaknesses. Automated detectors' performance on in-domain data indicates they can pick up on subtleties that humans cannot. However, we show they are brittle and fail to work when there are changes in the test audio's

environmental conditions. Given the extent of human limitations and the increasing availability of computational resources for deploying detectors, research should focus on improving these detectors. In the meantime, crowd-sourcing is a reasonable mitigation. We confirm crowd performance is on par with the top-performing automated detectors and is not as brittle. Extending fact-checking tools to include audio evaluations is one way to protect against deepfake threats.

Supporting information

S1 Fig. Confidence-adjusted accuracy scores per clip (English, unary, no familiarization).
(TIF)

S2 Fig. Confidence-adjusted accuracy scores per clip (English, unary, familiarization).
(TIF)

S3 Fig. Confidence-adjusted accuracy scores per clip (English, binary, no familiarization).
(TIF)

S4 Fig. Confidence-adjusted accuracy scores per clip (English, binary, familiarization).
(TIF)

S5 Fig. Confidence-adjusted accuracy scores per clip (Mandarin, unary, no familiarization).
(TIF)

S6 Fig. Confidence-adjusted accuracy scores per clip (Mandarin, unary, familiarization).
(TIF)

S7 Fig. Confidence-adjusted accuracy scores per clip (Mandarin, binary, no familiarization).
(TIF)

S8 Fig. Confidence-adjusted accuracy scores per clip (Mandarin, binary, familiarization).
(TIF)

Acknowledgments

Kimberly T. Mai thanks Kelvin Ma for assistance with translating the survey materials.

Author Contributions

Conceptualization: Kimberly T. Mai, Sergi Bray, Toby Davies, Lewis D. Griffin.

Data curation: Kimberly T. Mai.

Formal analysis: Kimberly T. Mai.

Investigation: Kimberly T. Mai.

Methodology: Kimberly T. Mai, Sergi Bray, Lewis D. Griffin.

Software: Kimberly T. Mai, Sergi Bray.

Supervision: Toby Davies, Lewis D. Griffin.

Visualization: Kimberly T. Mai.

Writing – original draft: Kimberly T. Mai.

Writing – review & editing: Toby Davies, Lewis D. Griffin.

References

1. Brewster T. Fraudsters Cloned Company Director's Voice In \$35 Million Bank Heist, Police Find. 2021 Oct 14 [Cited 2023 Jan 19]. Available from: <https://www.forbes.com/sites/thomasbrewster/2021/10/14/huge-bank-fraud-uses-deep-fake-voice-tech-to-steal-millions/?sh=7dfbccf67559>.
2. Stupp C. Fraudsters Used AI to Mimic CEO's Voice in Unusual Cybercrime Case. 2019 Aug 30 [Cited 2023 Jan 19]. Available from: <https://www.wsj.com/articles/fraudsters-use-ai-to-mimic-ceos-voice-in-unusual-cybercrime-case-11567157402>.
3. Choi S, Han S, Kim D, Ha S. Attention: Few-Shot Text-to-Speech Utilizing Attention-Based Variable-Length Embedding. In: Meng H, Xu B, Zheng TF, editors. Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020. ISCA; 2020. p. 2007–2011.
4. Alspach K. Does your boss sound a little funny? it might be an audio deepfake; 2022 Aug 18 [Cited 2023 Jan 19]. Available from: <https://www.protocol.com/enterprise/deepfake-voice-cyberattack-ai-audio>.
5. Caldwell M, Andrews J, Tanay T, Griffin L. AI-enabled future crime. *Crime Science*. 2020; 9(1):1–13. <https://doi.org/10.1186/s40163-020-00123-8>
6. Mirsky Y, Demontis A, Kotak J, Shankar R, Gelei D, Yang L, et al. The threat of offensive ai to organizations. *Computers & Security*. 2022; p. 103006.
7. Wu Z, Yamagishi J, Kinnunen T, Hanilçi C, Sahidullah M, Sizov A, et al. ASVspooF: The Automatic Speaker Verification Spoofing and Countermeasures Challenge. *IEEE Journal of Selected Topics in Signal Processing*. 2017; 11(4):588–604. <https://doi.org/10.1109/JSTSP.2017.2671435>
8. Nautsch A, Wang X, Evans NWD, Kinnunen TH, Vestman V, Todisco M, et al. ASVspooF 2019: Spoofing Countermeasures for the Detection of Synthesized, Converted and Replayed Speech. *IEEE Transactions on Biometrics, Behavior, and Identity Science*. 2021; 3(2):252–265. <https://doi.org/10.1109/TBIOM.2021.3059479>
9. Yamagishi J, Wang X, Todisco M, Sahidullah M, Patino J, Nautsch A, et al. ASVspooF 2021: accelerating progress in spoofed and deepfake speech detection. arXiv preprint arXiv:210900537. 2021 Sep 1 [Cited 2023 Jan 19].
10. Gamage D, Chen J, Ghasiya P, Sasahara K. Deepfakes and Society: What lies ahead? In: *Frontiers in Fake Media Generation and Detection*. Springer; 2022. p. 3–43.
11. Goodfellow I, Bengio Y, Courville A. *Deep Learning*. MIT Press; 2016.
12. Zhang Y, Tiño P, Leonardi A, Tang K. A survey on neural network interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence*. 2021. <https://doi.org/10.1109/TETCI.2021.3100641>
13. van Huijstee M, van Boheemen P, Das D, Nierling L, Jahnke J, Karaboga M, et al. Tackling Deepfakes in European Policy. European Parliament; 2021. Available from: [https://www.europarl.europa.eu/thinktank/en/document/EPRS_STU\(2021\)690039](https://www.europarl.europa.eu/thinktank/en/document/EPRS_STU(2021)690039).
14. Schick N. *Deep Fakes and the Infocalypse: What You Urgently Need To Know*. Hachette; 2020.
15. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, et al. Generative Adversarial Nets. In: *Advances in Neural Information Processing Systems*; 2014. Available from: https://proceedings.neurips.cc/paper_files/paper/2014/file/5ca3e9b122f61f8f06494c97b1afcc3-Paper.pdf
16. Sohl-Dickstein J, Weiss E, Maheswaranathan N, Ganguli S. Deep Unsupervised Learning using Non-equilibrium Thermodynamics. In: *Proceedings of the 32nd International Conference on Machine Learning*. Available from: <https://proceedings.mlr.press/v37/sohl-dickstein15.html>
17. Damer N, Saladié AM, Braun A, Kuijper A. MorGAN: Recognition Vulnerability and Attack Detectability of Face Morphing Attacks Created by Generative Adversarial Network. In: 2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS).
18. Bitouk D, Kumar N, Dhillon S, Belhumeur P, Nayar SK. Face Swapping: Automatically Replacing Faces in Photographs *ACM Trans Graph*. 2008; 27(3):1–8.
19. Zhang T. Deepfake generation and detection, a survey. *Multimedia Tools and Applications*. 2022; 81(5):6259–6276. <https://doi.org/10.1007/s11042-021-11733-y>
20. Barari S, Lucas C, Munger K. Political Deepfakes Are As Credible As Other Fake Media And (Sometimes) Real Media [Preprint]; 2021 [Cited 2023 Jan 19]. Available from: osf.io/cdfh3.
21. Appel M, Prielzel F. The detection of political deepfakes. *Journal of Computer-Mediated Communication*. 2022; 27(4). <https://doi.org/10.1093/jcmc/zmac008>
22. Bray SD, Johnson SD, Kleinberg B. Testing Human Ability To Detect Deepfake Images of Human Faces. arXiv preprint arXiv:221205056. 2022 Dec 7 [Cited 2023 Jan 19].

23. Nightingale SJ, Farid H. AI-synthesized faces are indistinguishable from real faces and more trustworthy. *Proceedings of the National Academy of Sciences*. 2022; 119(8):e2120481119. <https://doi.org/10.1073/pnas.2120481119> PMID: 35165187
24. Groh M, Epstein Z, Firestone C, Picard R. Deepfake detection by human crowds, machines, and machine-informed crowds. *Proceedings of the National Academy of Sciences*. 2021; 119(1):e2110013119. <https://doi.org/10.1073/pnas.2110013119>
25. Köbis NC, Doležalová B, Soraperra I. Fooled twice: People cannot detect deepfakes but think they can. *iScience*. 2021; 24(11):103364. <https://doi.org/10.1016/j.isci.2021.103364> PMID: 34820608
26. Tahir R, Batool B, Jamshed H, Jameel M, Anwar M, Ahmed F, et al. Seeing is believing: Exploring perceptual differences in deepfake videos. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*; 2021. p. 1–16.
27. Geirhos T, Zimmerman RS, Bilodeau BL, Brendal W, Kim, B. Don't trust your eyes: on the (un)reliability of feature visualizations. *arXiv preprint arXiv: 230604719*. 2023 Jun 13 [Cited 2023 Jun 13].
28. Watson G, Khanjani Z, Janeja VP. Audio Deepfake Perceptions in College Going Populations. *arXiv preprint arXiv:211203351*. 2021 Dec 6 [Cited 2023 Jan 19].
29. Wang X, Yamagishi J, Todisco M, Delgado H, Nautsch A, Evans N, et al. ASVspooF 2019: A large-scale public database of synthesized, converted and replayed speech. *Computer Speech & Language*. 2020; 64:101114. <https://doi.org/10.1016/j.csl.2020.101114>
30. Müller NM, Pizzi K, Williams J. Human perception of audio deepfakes. In: *Proceedings of the 1st International Workshop on Deepfake Detection for Audio Multimedia*; 2022. p. 85–91.
31. Karras T, Laine S, Aittala M, Hellsten J, Lehtinen J, Aila T. Analyzing and improving the image quality of StyleGAN. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; 2020. p. 8110–8119.
32. Dolhansky B, Bitton J, Pflaum B, Lu J, Howes R, Wang M, et al. The Deepfake Detection Challenge Dataset. *arXiv preprint arXiv:200607397*. [Cited 2023 Mar 27]. Dataset.
33. Li Y, Yang X, Sun P, Qi H, Lyu S. Celeb-df: A large-scale challenging dataset for deepfake forensics. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; 2020. p. 3207–3216.
34. Rossler A, Cozzolino D, Verdoliva L, Riess C, Thies J, Nießner M Faceforensics++: Learning to detect manipulated facial images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; 2019. p. 1–11.
35. Perov I, Gao D, Chervonyi N, Liu K, Marangonda S, Umé C, et al. DeepFaceLab: Integrated, flexible and extensible face-swapping framework. *arXiv preprint arXiv:2005.05535*. [Cited 2023 Mar 27]. Dataset.
36. Kumar K, Kumar R, De Boissiere T, Gestin L, Teoh WZ, Sotelo J, et al. MelGAN: Generative adversarial networks for conditional waveform synthesis. In: *Proceedings of Advances in Neural Information Processing Systems*; 2019; 32.
37. Ito K, Johnson L. The LJ Speech Dataset; 2017 [Cited 2023 Jan 19]. Dataset. <https://keithito.com/LJ-Speech-Dataset/>.
38. Databaker. Chinese Standard Mandarin Speech Corpus; 2019 [Cited 2023 Jan 19]. Dataset. https://www.data-baker.com/open_source.html.
39. Watanabe S, Hori T, Karita S, Hayashi T, Nishitoba J, Unno Y, et al. ESPnet: End-to-End Speech Processing Toolkit. In: *Proceedings of Interspeech*; 2018. p. 2207–2211.
40. Kim J, Kong J, Son J. Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech. In: *International Conference on Machine Learning*. PMLR; 2021. p. 5530–5540. Available from: <https://proceedings.mlr.press/v139/kim21f.html>.
41. Wang X, Yamagishi J. A Comparative Study on Recent Neural Spoofing Countermeasures for Synthetic Speech Detection. In: *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association, Brno, Czechia, 30 August—3 September 2021*. ISCA; 2021. p. 4259–4263.
42. Ma H, Yi J, Wang C, Yan X, Tao J, Wang T, et al. FAD: A Chinese Dataset for Fake Audio Detection. *arXiv preprint arXiv:220712308*. 2022 Jul 12. [Cited 2023 Jan 19].
43. Delgado H, Evans N, Kinnunen T, Lee KA, Liu X, Nautsch A, et al. ASVspooF 2021: Automatic speaker verification spoofing and countermeasures challenge evaluation plan. *arXiv preprint arXiv:210900535*. 2021 Sep 1 [Cited 2023 Jan 19].
44. Shen Z, Liu J, He Y, Zhang X, Xu R, Yu H, et al. Towards out-of-distribution generalization: A survey. *arXiv preprint arXiv:210813624*. 2021 Aug 31. [Cited 2023 Jan 19].
45. Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*. 2020; 17:261–272. <https://doi.org/10.1038/s41592-019-0686-2> PMID: 32015543

46. Seabold S, Perktold J. statsmodels: Econometric and statistical modeling with python. In: 9th Python in Science Conference; 2010.
47. Leskovec J, Rajaraman A, Ullman JD. Mining of Massive Datasets. 2nd ed. USA: Cambridge University Press; 2014.
48. Mok PPK, Dellwo V. Comparing native and non-native speech rhythm using acoustic rhythmic measures: Cantonese, Beijing Mandarin and English. In: Proc. Speech Prosody 2008; 2008. p. 423–426.
49. Hummert ML, Shaner JL, Garstka TA, Henry C. Communication with older adults: The influence of age stereotypes, context, and communicator age. *Human communication research*. 1998; 25(1):124–151. <https://doi.org/10.1111/j.1468-2958.1998.tb00439.x>
50. Strand EA. Uncovering the role of gender stereotypes in speech perception. *Journal of language and social psychology*. 1999; 18(1):86–100. <https://doi.org/10.1177/0261927X99018001006>
51. Chesney R, Citron D. Deepfakes and the new disinformation war: The coming age of post-truth geopolitics. *Foreign Affairs*. 2019 Jan/Feb;98:147.