

000

001

002

003

004

005

006

007

008

009

010

011

012

013

014

015

016

017

018

019

020

021

022

023

024

025

026

027

028

029

TELL ME HABIBI, IS IT REAL OR FAKE?

Anonymous authors

Paper under double-blind review

ABSTRACT

Deepfake generation methods are evolving fast, making fake media harder to detect and raising serious societal concerns. Most deepfake detection and dataset creation research focuses on monolingual content, often overlooking the challenges of multilingual and code-switched speech, where multiple languages are mixed within the same discourse. Code-switching, especially between Arabic and English, is common in the Arab world and is widely used in digital communication. This linguistic mixing poses extra challenges for deepfake detection, as it can confuse models trained mostly on monolingual data. To address this, we introduce **ArEnAV**, the first large-scale Arabic-English audio-visual deepfake dataset featuring intra-utterance code-switching, dialectal variation, and monolingual Arabic content. **It contains 387k videos and over 765 hours of real and fake videos.** Our dataset is generated using a novel pipeline integrating four Text-To-Speech and two lip-sync models, enabling comprehensive analysis of multilingual multimodal deepfake detection. We benchmark our dataset against existing monolingual and multilingual datasets, state-of-the-art deepfake detection models, and a human evaluation, highlighting its potential to advance deepfake research. The dataset and code will be made public.

1 INTRODUCTION

Deepfake technologies, involving the artificial generation and manipulation of audio-visual content, have rapidly advanced, significantly complicating the task of distinguishing real media from synthetic creations. The potential misuse of deepfakes for misinformation, defamation, or impersonation presents profound societal risks, driving substantial research into their detection. Although initial deepfake research primarily focused on manipulating individual modalities, audio-only (Todisco et al., 2019) or video-only (Jiang et al., 2020; Kwon et al., 2021; Li et al., 2020b), recent developments increasingly consider joint manipulation of audio and visual streams for more realistic synthesis.

A significant gap remains in existing deepfake datasets (Table 1), which largely overlook multilingual scenarios, particularly code-switching (CSW), despite its global prevalence among bilingual speakers. In the Arab world, CSW is a prominent feature of daily communication, serving not only as a linguistic tool but also as a marker of cultural identity and social context (Hamed et al., 2025). Arabic speakers frequently alternate between Arabic and English within the same sentence, such as:

مهم جداً *deepfake detection* (The topic of *deepfake detection* is very important).

This challenge is compounded by the diglossic nature of Arabic (Ferguson, 1959), comprising of two main varieties: Modern Standard Arabic (MSA) and Dialectal Arabic (DA). MSA functions as a lingua franca across the Arab world and is primarily used in formal settings. DA, belonging to each country, is used in everyday conversations and informal writing. Given that Arabic is one of the most widely spoken languages worldwide, ranked fifth by number of Standard Arabic speakers (Ethnologue, 2025), handling its diglossic variation and code-switching phenomena is essential for building deepfake detection systems that address the linguistic diversity in real-world media content.

Recent studies provide compelling evidence of how common CSW is among Arabic speakers. The ZAEBUC-Spoken corpus (Hamed et al., 2024) reveals that approximately 19% of spoken utterances exhibit CSW, having an average of 44% English words. The corpus also highlights the presence of diglossic CSW between Arabic variants. Similarly, the ArzEn corpus (Hamed et al., 2020)

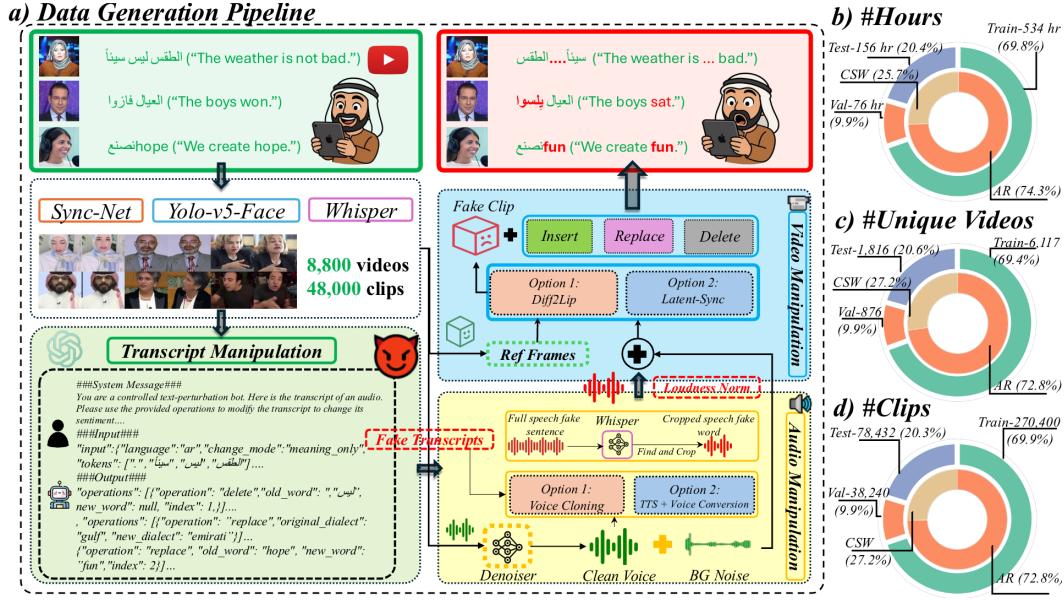


Figure 1: a) We show the data generation pipeline for ArEnAV dataset. In a) input videos are analysed for audio, face, and text extraction. Using few-shot prompts with GPT-4.1-mini, CSW-based spoken text manipulation is performed. This is followed by speech and face enactment generation. b-d) The plots show the data splits and CSW distribution. Here is an example of CSW input and manipulated text with translations in parentheses: $\text{اصنـح} \rightarrow \text{اصـنـع}$ hope (“We create hope.”) \rightarrow $\text{اصـنـع} \text{ fun}$ (“We create fun.”)

demonstrates high CSW frequency, where 63% of utterances involve CSW with approximately 19% of words being English. These findings highlight the extent to which Arabic-English CSW is not merely an incidental phenomenon, but a widespread communicative strategy. Despite its ubiquity, deepfake detection systems remain largely ill-equipped to handle such language alternation, focusing predominantly on monolingual data. Addressing this important oversight, our work seeks to bridge this critical gap by introducing the first Arabic-English CSW audio-visual deepfake dataset, thus advancing the field toward more relevant detection systems. Our core contributions are as follows:

- Introduction of ArEnAV, the first large-scale Arabic-English audio-visual deepfake dataset featuring intra-utterance code-switching and dialectal variation, including both bilingual and diglossic switching across Modern Standard Arabic, Egyptian, Levantine, and Gulf dialects, addressing a critical gap in multilingual deepfake research.
- A novel data generation pipeline tailored to English and Arabic (MSA and dialect-rich content), integrating four TTS (Text to Speech) and two lip-sync models.
- A comprehensive analysis contrasting our dataset against existing monolingual and multilingual datasets, existing state-of-the-art (SOTA) deepfake detection models, and a detailed User Study, underscoring its unique difficulty in detection by models and humans alike.

Table 1: **Details for publicly available deep-fake datasets in chronologically ascending order.**
 Cla: Binary classification, SL: Spatial localization, TL: Temporal localization, FS: Face swapping, RE: Face reenactment, TTS: Text-to-speech, VC: Voice conversion.

Dataset	Year	Tasks	Manip. Modality	Method	#Total	Multilingual	Code Switching
Google DFD (Nick & Andrew, 2019)	2019	Cla	V	FS	3,431	X	X
DFDC (Dolhansky et al., 2020b)	2020	Cla	AV	FS	128,154	X	X
DeeperForensics (Jiang et al., 2020)	2020	Cla	V	FS	60,000	X	X
Celeb-DF (Li et al., 2020b)	2020	Cla	V	FS	6,229	X	X
KoDF (Kwon et al., 2021)	2021	Cla	V	FS/RE	237,942	X	X
FakeAVCeleb (Khalid et al., 2022)	2021	Cla	AV	RE/FS	25,500+	X	X
ForgeryNet (He et al., 2021)	2021	SL/TL/Cla	V	Random FS/RE	221,247	X	X
ASVspoof2021DF (Liu et al., 2023)	2021	Cla	A	TTS/VC	593,253	X	X
LAV-DF (Cai et al., 2022)	2022	TL/Cla	AV	Content-driven RE/TTS	136,304	X	X
DF-Platter (Narayan et al., 2023)	2023	Cla	V	FS	265,756	X	X
AV-IM (Cai et al., 2023a)	2023	TL/Cla	AV	Content-driven RE/TTS	1,146,760	X	X
PolyGlotFake (Hou et al., 2024)	2024	Cla	AV	RE/TTS/VC	15,238	✓	X
Illusion (Thakral et al., 2025)	2025	Cla	AV	FS/RE/TTS	1,376,371	✓	X
ArEnAV (Ours)	2025	Cla/TL	AV	Content Driven RE/TTS/VC	387,072	✓	✓

108 2 RELATED WORK

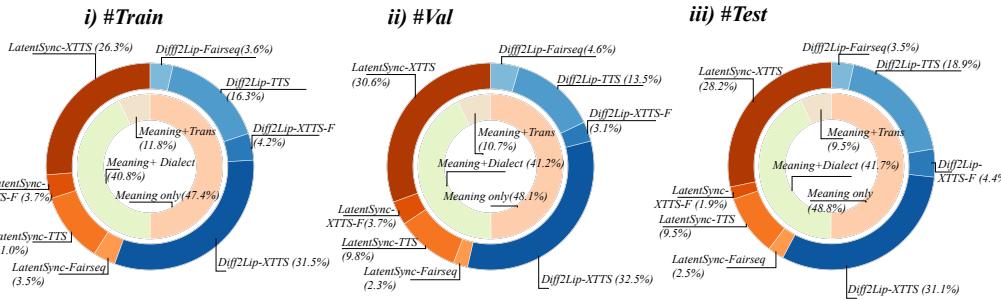
110 Early deepfake research was predominantly monolingual and modality-specific. Initial significant
 111 contributions included video manipulation techniques such as FaceSwap and Face2Face as introduced
 112 by Thies et al. (2020), which led to seminal datasets like FaceForensics++ (Rössler et al., 2019)
 113 and the DeepFake Detection Challenge (DFDC) (Dolhansky et al., 2020a). These datasets primarily
 114 provided facial manipulations within single-language contexts, focusing largely on visual realism.

115 Parallel to video deepfake advancements, audio deepfakes evolved rapidly, driven by progress in
 116 text-to-speech (TTS) synthesis, voice conversion, and generative audio models such as Tacotron
 117 (Wang et al., 2017). Datasets like ASVspoof (Wang et al., 2020) and WaveFake (Frank & Schönher, 2021) contributed significantly by providing benchmarks to evaluate audio manipulation detection
 118 methods, albeit still largely restricted to English.

120 In recent years, research has expanded towards joint audio-visual deepfake manipulations. Datasets
 121 such as FakeAVCeleb (Khalid et al., 2022) showcased realistic lip-synced speech synthesis in tandem
 122 with facial manipulations. AV-Deepfake1M (Cai et al., 2024a) further advanced this domain by
 123 automating transcript alterations to create nuanced, localized audio-visual deepfakes, highlighting the
 124 necessity of detecting temporally and contextually subtle manipulations.

125 Recently, there has been increased focus on multilingual audio deepfakes. These efforts have revealed
 126 key limitations in generalizing detection models across languages and proposed new resources to
 127 address these challenges. Marek et al. (2024) conducted a comprehensive study on cross-lingual
 128 deepfake detection, showing that models trained on one language often fail to generalize effectively
 129 to others, underscoring the role of language-specific phonetic and prosodic features in model per-
 130 formance. Multilingual audio-visual datasets emerged even more recently to address the global
 131 dimension of deepfake threats. The PolyGlotFake dataset (Hou et al., 2024) contains audio-visual
 132 deepfakes across seven languages. Although the dataset covers a wide range of language, the size of
 133 the real data is significantly small. Nonetheless, these multilingual datasets remain limited to either
 134 monolingual or monomodal scenarios within each single instance, ignoring the prevalent reality of
 135 intra-utterance language switching.

136 Our work directly responds to this critical gap. Unlike previous studies, we not only incorporate
 137 multilingual content but also explicitly generate intra-utterance code-switched audio-visual deepfakes.
 138 We leverage SOTA TTS and lip sync methodologies adapted for multilingual use, resulting in realistic,
 139 diverse, and challenging benchmarks.



150 Figure 2: Dataset distribution for i) *Train*, ii) *Val* and iii) *Test* split. The outer layer shows the
 151 split between various combinations of Text-to-Speech and Lip-Sync models used for audio-visual
 152 manipulation. The middle layer shows the distribution of language in the original transcript, which is
 153 either *Ar* (Arabic) or *CSW* (Code-Switched English-Arabic). The inner layer shows the distribution
 154 of different operations applied to the original transcripts, "meaning only", "dialect+meaning", and
 155 "meaning + translation" (For fine-grained detail about what they entail, refer to Table 16 in Appendix.)

158 3 ARENAV DATASET

160 ArEnAV is a large-scale audio-visual deepfake dataset specifically focused on Arabic–English CSW.
 161 Comprising approximately 765 hours of video data sourced from 8,809 unique YouTube videos,
 ArEnAV establishes itself as the first and most extensive benchmark for multilingual deepfake

162 detection (see Table 1 for dataset comparison). The dataset is constructed to preserve the original
 163 identity and environmental context of the source videos while systematically manipulating the
 164 semantic content to introduce Arabic-English CSW. Following the taxonomy proposed by Cai et al.
 165 (2024a), ArEnAV includes three manipulation strategies: **Fake Audio & Fake Video**: Both audio and
 166 visual content are synthetically generated, simulating complete audiovisual deepfakes. **Fake Audio**
 167 & **Real Video**: The audio track is manipulated to introduce anti-semantic and CSW content while
 168 maintaining the original visual content. **Real Audio & Fake Video**: The original audio is retained,
 169 while facial movements and lip synchronization are altered to create visually deceptive content.

170 3.1 DATA COLLECTION

171 We use the YouTube video links from VisPer’s Arabic Train subset (Narayan et al., 2024). We chose
 172 VisPer because it is the largest publicly available non-English audio-visual corpus, with over 1,200
 173 hours of Arabic alone. Its 200-keyword crawler (“interview,” “tutorial,” etc.) pulls videos spanning
 174 talk shows, vlogs, documentaries, and lectures, mirroring the broad-coverage strategy that is required
 175 for a fair and diverse representation of a culture-specific deepfake dataset. We first run a scene change
 176 detection model to split the video into clips, and then we use Yolo-v5 to obtain the faces in each
 177 frame as well as track them across frames. Since we did not have ground truth for transcripts for
 178 VisPER, we surveyed state-of-the-art Automatic Speech Recognition (ASR) models for Arabic, based
 179 on the Arabic ASR Leaderboard on HuggingFace. Following a qualitative comparison, we finalized
 180 the Whisper-v2 (finetuned on English-Arabic data) for our method, with the default output language
 181 set to Arabic. Following the transcripts, we apply Forced Alignment between the audio and text,
 182 using a multilingual wav2vec2 model (Baevski et al., 2020) supporting both Arabic and English. This
 183 provides us with word-level timestamps for code-switched Arabic and English data.

184 3.2 DATA GENERATION PIPELINE

185 The data generation pipeline roughly consists of three stages: transcript manipulation, audio
 186 generation, and video generation. First, we apply controlled modifications to the transcript. Secondly,
 187 we synthesize new audio for the altered transcript while preserving the speaker’s voice character-
 188 istics. Finally, we render a lip-synced video that matches the new audio, producing a realistically
 189 manipulated video clip. We detail each stage as follows.

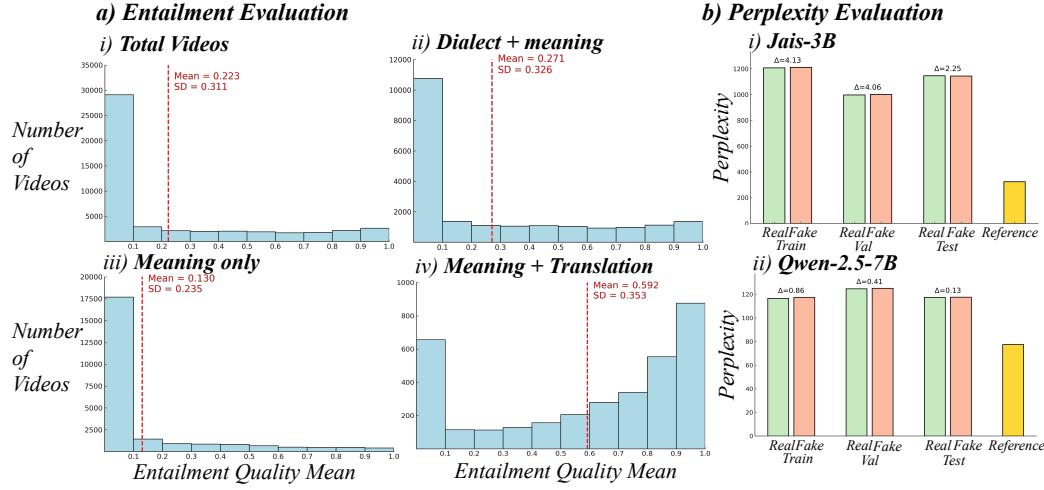
190 3.2.1 TRANSCRIPT MANIPULATION

191 We leverage GPT-4.1-mini (OpenAI, 2025) to perform content-driven modifications of our multilin-
 192 gual transcripts. We define eight distinct transcript change modes that span both code-switched and
 193 Arabic-only contexts, allowing fine-grained control over how the transcript is altered. These modes
 194 include three main operations: first, *meaning only*, which only involves changing the meaning of the
 195 word and keeping the language as it is, second, *meaning + dialect*, which involves changing the
 196 meaning of the word and changing its language to another Arabic variant (either MSA or any dialect),
 197 and lastly, *meaning + translation*, which asks the model to change the meaning of the word, and then
 198 translate it to English. Table 2 summarizes the eight modification modes with their intended effect.
 199 By categorizing edits in this way, we ensure a controlled and diverse set of manipulations ranging
 200 from subtle word substitutions to introducing or removing CSW instances. Due to the effectiveness
 201 of few-shot prompting, we prompt GPT-4.1-mini with 15 examples, explaining various kinds of
 202 transitions and possible changes. Examples of original and augmented transcripts achieved by these
 203 manipulation rules are shown in Appendix A.5. We provide the prompt in Appendix A.6. We report
 204

205 Table 2: Transcript manipulation rules in ArEnAV for Arabic (AR) and English (EN) words.

#	Original Transcript	Original Word	Inserted Word	Operation
1	CSW	EN	EN	Change meaning only (keep English)
2	CSW	AR	AR	Change meaning only (keep Arabic variant)
3	CSW	AR	AR	Change meaning + change Arabic variant
4	CSW	AR	AR/EN	multi-op; When 2-3 ops → edit 1 EN and 1-2 AR words
5	Arabic	AR	AR	Change meaning only (keep Arabic variant)
6	Arabic	AR	AR	Change meaning + change Arabic variant
7	Arabic	AR	EN	Change meaning + change language to English
8	Arabic	AR	AR/EN	multi-op; Apply all operations on Arabic words

216 text manipulations distributions as follows: replacement (94.6%), insertion (5.1%) and deletion
 217 (0.3%). These distributions reflect GPT’s manipulation choices, which were not manually enforced.
 218



236 Figure 3: a) Entailment distribution over i) All change modes, ii) Dialect + meaning, iii) Meaning only,
 237 and iv) Meaning + translation. b) Perplexity Evaluation distribution among dataset splits, showing
 238 perplexity calculated using i) Jais-3B, an Arabic-English LLM, and ii) Qwen-2.5-7B. b) Perplexity
 239 calculated using i) Jais-3B ii) Qwen-2.5-7B. Reference in both shows the perplexity calculated on an
 240 Arabic-English CSW text dataset (SDIANCAI, 2025).

Transcript Quality:

243 To quantify the impact of our LLM-based manipulations, we employ two complementary metrics:
 244 **Bidirectional Entailment Quality Mean**: the average of Real→Fake and Fake→Real NLI entailment
 245 scores (1.0 = full semantic entailment; 0.0 = direct contradiction) and **Perplexity**: how well a language
 246 model predicts a transcript (lower = more fluent/natural). Table 3a shows the distribution of entailment
 247 quality means over different types of perturbations. In every subset, a large fraction of samples lies
 248 below the 0.5 threshold, and many even in the contradiction zone, demonstrating that our pipeline
 249 reliably injects semantic change regardless of language or dialect.

250 Figure 3b reports average perplexities on real versus fake transcripts under two open-source LLMs;
 251 Jais-3B (Sengupta et al., 2023) and Qwen-2.5-7B (Qwen et al., 2025), across the data splits. The
 252 minimal difference in perplexity shows that our fake transcripts remain fluent and natural, despite
 253 major changes in meaning. This balance between altered content and surface-level fluency is essential
 254 for generating effective audio-visual deepfakes

3.2.2 AUDIO GENERATION

255 The next step involves generating a synthetic audio track that precisely follows the edited transcript
 256 while maintaining the voice characteristics of the original on-screen speaker. Initially, we segment the
 257 audio into clean speech and background noise using a Denoiser (Defossez et al., 2020). Conventional
 258 zero-shot voice cloning systems, such as YourTTS (Casanova et al., 2023), exhibit strong performance
 259 in English but struggle with Arabic phonetics and cross-lingual synthesis. To address this, we employ
 260 four targeted cloning strategies: a) **XTTS-v2 (Casanova et al., 2024)**: A multilingual, zero-shot
 261 TTS model natively supporting Arabic, English, and code-switching. b) **XTTS-v2 (Casanova et al.,
 262 2024) + OpenVoice-v2 (Qin et al., 2024)**: When a reference voice sample is available, we achieve
 263 higher fidelity by generating the utterance with XTTS-v2 and performing speaker conversion via
 264 OpenVoice-v2. c) **Fairseq Arabic TTS (Ott et al., 2019) + OpenVoice-v2 (Qin et al., 2024)**:
 265 For fully Arabic sentences, we generate audio with the Fairseq Arabic TTS, followed by speaker
 266 conversion using OpenVoice-v2. d) **GPT-TTS (OpenAI, 2023) + OpenVoice-v2 (Qin et al., 2024)**:
 267 We randomly select one voice out of 29, generate the sentence, and then convert the audio to the
 268 target speaker’s voice with OpenVoice-v2.

270 The audio-generation flow depends on the edit type. For *insert* or *replace* operations, we regenerate
 271 the complete sentence and validate the generated audio using Whisper-Turbo (Radford et al., 2022),
 272 retaining only samples that exactly match the intended transcript. This step ensures intelligibility
 273 and accurate timestamp alignment for splicing the segment into the original audio. If validation fails,
 274 we discard the sample. For a *delete* operation, we remove speech segments entirely, preserving only
 275 background noise. Finally, after each edit, we normalise the loudness of the manipulated segment
 276 relative to the original audio and recombine it with the extracted environment noise.

277 **278 *Audio Quality:***

279 Table 3 presents the comparison of audio quality for
 280 ArEnAV based on speaker similarity, signal quality
 281 and distribution realism with existing audio-visual
 282 deepfake datasets. We report speaker encoder co-
 283 sine similarity (SECS), Signal-to-Noise (SNR) and
 284 Fréchet audio distance (FAD) for recent Audio-Visual
 285 datasets. SECS measures the speaker’s voice simi-
 286 larity between a generated clip and the real reference
 287 (range $[-1, 1]$, higher is better), while FAD evaluates
 288 the distributional distance between the generated audio and real audio (lower is better). The metrics
 289 combined indicate that ArEnAV has high-quality audio samples.

290 **3.3 VISUAL MANIPULATION**

291 For video generation, after extensive experimentation, we chose two diffusion-based lip-sync ap-
 292 proaches: Diff2Lip (Mukhopadhyay et al., 2023) and LatentSync (Li et al., 2025). Both of these
 293 models perform high-quality zero-shot lip-sync and are open-sourced. Using the newly generated
 294 audio and the original video’s frames we generate the fake frames. For *replace* and *insert* word
 295 operations, we generate the fake frames for the new word, and for *delete* word operations, we generate
 296 a face with closed lips i.e., without audio.

297 **Visual Quality:** To evaluate visual quality, we use three stan-
 298 dard metrics: Peak Signal-to-Noise Ratio (PSNR), Structural
 299 Similarity Index (SSIM), and Fréchet Inception Distance (FID).
 300 Table 4 presents PSNR, SSIM, and FID results for the ArEnAV
 301 dataset. PSNR and SSIM measure pixel-level and structural simi-
 302 larity, respectively, between fake and original frames (higher is
 303 better) (ArEnAV lies nearby AV-1M). FID assesses realism by
 304 comparing the distributions of fake and real frames in a learned
 305 image feature space (lower is better) (ArEnAV slightly more
 306 than AV-1M). These scores highlight that ArEnAV attains higher
 307 / comparable visual quality compared to other deepfake datasets.

308 **Real Perturbations:** To mimic real-life video scenarios better, we add localized perturbations to
 309 both the real and the fake videos. We apply 15 different visual filters (eg, salt-pepper noise and
 310 camera shaking) and 10 different audio manipulations (eg, time-stretching, random loudness and
 311 pitch). For each video, we randomly sample one to three instances for visual perturbations and one to
 312 two instances for audio perturbations. Perturbation details are mentioned in Appendix A.4.

313 **3.4 USER STUDY**

316 To investigate whether humans can identify deepfakes in
 317 ArEnAV, we conducted a user study with 19 participants, out
 318 of which, 15 are native Arabic speakers, and 4 have basic
 319 knowledge of Arabic. We randomly sampled 20 videos, with
 320 either 0 or 1 manipulation. **Instructions for User Study:**
 321 Each participant was asked to 1) watch the video, and 2)
 322 answer 3 questions, i) Is the video real or fake, ii) If it is fake, localize where they think the fake
 323 region is, iii) Whether the given video contains Arabic-English code-switching or not, iv) Give a
 reason for labelling the video (if they have) as a deepfake.

Table 3: Audio quality comparison across different datasets.

Dataset	Language	SECS↑	SNR(dB)↑	FAD↓
FakeAVCeleb	English	0.543	2.16	6.598
LAV-DF	English	0.984	7.83	0.306
AV-Deepfake1M	English	0.991	9.39	0.088
ArEnAV	Arabic, English	0.990	7.65	0.140

Table 4: Visual quality comparison across different datasets.

Dataset	PSNR(dB)↑	SSIM↑	FID↓
FF++	24.40	0.812	1.06
DFDC	—	—	5.69
FakeAVCeleb	29.82	0.919	2.29
LAV-DF	33.06	0.898	1.92
AV-Deepfake1M	39.49	0.977	0.49
ArEnAV	37.70	0.971	0.68

Table 5: Detection and Localization results from our User Study.

Method	Acc.	AP@0.1	AP@0.5	AR@1
ArEnAV	60.00	8.35	0.79	1.38

The results in Table 5 reaffirm our hypothesis that identifying audiovisual deepfakes in multilingual (specially CSW) and multimodal settings is a non-trivial task, as even humans achieve only 60% accuracy, while it is even harder to localize the deepfakes, with AP@0.5 at 0.79. Further, Table 6 shows the primary reasons why people classified the videos as fake. We report that 85% of the users fail to identify deepfakes when the manipulation happens in the English word, in the CSW video, which can be attributed to a higher quality of voice cloning in English as well as the natural change in tone when a person code-switches, which makes it harder to detect. Further, localization is very tough due to the very high quality of lip-sync with diffusion models, as shown in Table 6, where the video being the reason for fake classification is only 8.7%.

3.5 DATASET STATISTICS

Table 7: Data distribution in ArEnAV and comparison with other multilingual datasets.

Subset	#Unique Videos	#Real Videos	#Fake Videos	#Non-English Clips	#CSW Videos	#Arabic Videos	Arabic Variants
PolyGlotFake (Hou et al., 2024)	766	766	14,472	11,941	0	1,403	NA
Illusion (Thakral et al., 2025)	–	141,440	1,234,931	4,385	0	–	NA
ArEnAV-Train	6,117	67,600	202,800	270,400	69,544	200,856	Egyptian, MSA,
ArEnAV-Validation	876	9,560	28,680	38,240	10,416	27,824	Levantine, Gulf
ArEnAV-Test	1,816	19,608	58,824	78,432	19,832	58,600	–
ArEnAV (total)	8,809	96,768	290,304	387,072	99,792	287,280	–

Table 7 compares ArEnAV with other multilingual deepfake detection datasets. Existing multilingual datasets like PolyGlotFake (Hou et al., 2024) and Illusion (Thakral et al., 2025) have significantly smaller multilingual content, containing limited Arabic data (1,400 Arabic videos in PolyGlotFake and minimal in Illusion across 26 languages). ArEnAV includes 387k videos sourced from 8,809 unique YouTube videos, totaling over 765 hours. Videos average approximately 7.7 sec each, with train, val, and test splits created via multilabel stratified sampling in a 7:1:2 ratio, ensuring no overlap.

Table 8: Fake segment duration comparison between ArEnAV and AV-1M.

Dataset	Mean (s)	Median (s)	Minimum (s)	Maximum (s)	Video-length (s)	Fake ratio (%)	Relative length (x)
ArEnAV	0.696	0.625	0.02	6.16	5.97	12.1	2.1
AV-1M	0.326	—	—	—	9.07	3.7	1.0

Fake Region Comparison: In Table 8, we summarize forged-segment duration statistics for ArEnAV and AV-1M. The Table highlights substantially longer and proportionally larger forged spans in ArEnAV (Fake segments are 2.1 times longer relative to AV-1M), confirming that performance drops stem from the intrinsic difficulty of detecting linguistically precise intra-utterance manipulations rather than from short spans.

Computational Cost: We spent around 50 GPU hours to generate the real transcript using Whisper-Large-V2 (Radford et al., 2022), 200 dollars worth of OpenAI credits, to generate fake transcripts and Text-to-Speech model, TTS-1 (OpenAI, 2023), and 650 GPU hours for video generation. Overall, we needed 800 GPU hours to generate AvEnAV with NVIDIA RTX-6000 GPUs.

4 BENCHMARK AND METRICS

We organize the data into *train*, *validation*, and *test* split. We use multilabel stratified sampling to divide the data in equal proportions based on the method type, the change mode, and the ground truth language. We also show evaluation on two subsets, *subset V*, which excludes videos with audio-only manipulation, and *subset A*, which excludes videos with visual-only manipulations. We evaluate models on two tasks, **temporal localization** and **detection** of audio-visual deepfakes. We use average precision (AP) and average recall (AR) metrics as prior works (He et al., 2021; Cai et al., 2022; 2023a) for temporal localization. For the task of deepfake detection, we use the standard evaluation protocol (Rossler et al., 2019; Dolhansky et al., 2020b; Cai et al., 2023a) to report video-level accuracy (Acc.) and area under the curve (AUC).

Table 6: Distribution of top reasons for predicting a video as Fake in our User Study.

Reason	Percentage (%)
Unintelligible speech (weird audio)	36.5
Video/audio mismatch (lip sync is off)	25.1
Audio sounds artificial	24.7
Video looks artificial	8.7
Code-switching is unnatural	3.0
Incoherent sentence	1.9

Table 9: Temporal localization results on the test set of ArEnAV.

Set	Method	Mod.	AP@0.5	AP@0.75	AP@0.9	AP@0.95	AR@50	AR@30	AR@20	AR@10	AR@5
Full dataset	Meso4	V	0.02	0.01	0.00	0.00	0.09	0.09	0.09	0.09	0.09
-	MesoInception	V	0.56	0.18	0.04	0.01	4.11	4.11	4.11	4.11	4.08
-	Xception	V	22.50	10.26	2.29	0.58	19.13	19.13	19.13	19.13	19.13
-	BA-TFD (ZS)	AV	0.17	0.01	0.00	0.00	9.72	5.20	3.07	1.46	0.73
-	BA-TFD+ (ZS)	AV	0.11	0.00	0.00	0.00	5.77	2.95	2.09	0.87	0.37
-	BA-TFD	AV	2.42	0.55	0.01	0.00	22.30	10.31	3.41	2.54	1.67
-	BA-TFD+	AV	3.74	1.10	0.06	0.01	30.75	9.42	4.55	3.05	1.83
Set V	Meso4	V	0.02	0.01	0.00	0.00	0.10	0.10	0.10	0.10	0.10
-	MesoInception	V	0.83	0.27	0.05	0.01	5.56	5.56	5.56	5.56	5.53
-	Xception	V	32.76	14.48	3.30	0.81	27.78	27.78	27.78	27.78	27.78
-	BA-TFD (ZS)	AV	0.12	0.00	0.00	0.00	8.44	4.34	2.44	1.13	0.49
-	BA-TFD+ (ZS)	AV	0.07	0.00	0.00	0.00	4.69	2.39	1.65	0.69	0.29
-	BA-TFD	AV	3.65	0.25	0.01	0.00	25.31	9.03	3.64	2.34	1.64
-	BA-TFD+	AV	5.65	1.89	0.08	0.02	31.09	13.21	5.91	3.05	2.05
Set A	Meso4	V	0.02	0.01	0.00	0.00	0.08	0.08	0.08	0.08	0.08
-	MesoInception	V	0.38	0.09	0.01	0.00	3.25	3.25	3.25	3.25	3.22
-	Xception	V	14.72	3.92	0.29	0.09	11.78	11.78	11.78	11.78	11.78
-	BA-TFD (ZS)	AV	0.23	0.01	0.00	0.00	12.14	6.46	3.85	1.83	0.95
-	BA-TFD+ (ZS)	AV	0.14	0.01	0.00	0.00	7.32	3.79	2.69	1.13	0.48
-	BA-TFD	AV	3.21	0.60	0.02	0.00	24.45	9.26	4.15	2.61	1.93
-	BA-TFD+	AV	4.35	1.10	0.10	0.00	28.35	11.23	4.85	3.11	2.00

Table 10: Deepfake detection results on the test set of ArEnAV.

Label Access For Training	Pretraining Data	Methods	Mod.	Fullset AUC	Fullset Acc.	Subset V AUC	Subset V Acc.	Subset A AUC	Subset A Acc.
Zero-Shot	ASVSpoof-19	XLSR-Mamba	A	39.19	52.77	52.73	40.68	52.50	42.59
-	Internet Scale	Video-LLaMA (7B)	V	51.48	26.29	51.47	34.21	51.43	34.18
-	Internet Scale	Video-LLaMA (7B)	AV	48.79	59.29	48.71	55.37	48.86	55.26
-	AV-1M	BA-TFD	AV	61.73	26.00	66.42	34.07	59.36	33.97
-	AV-1M	BA-TFD+	AV	60.96	25.84	64.49	34.28	59.44	33.80
Video Level	ArEnAV	XLSR-Mamba	A	73.00	61.00	57.47	66.16	86.33	78.00
-	ArEnAV	Meso4	V	49.30	75.00	49.15	66.67	49.30	66.67
-	ArEnAV	MesoInception4	V	50.34	46.23	50.28	47.48	50.35	47.67
-	ArEnAV	Xception	V	50.05	75.00	49.90	66.67	50.32	66.67
Frame level	ArEnAV	Meso4	V	49.55	26.60	49.60	34.40	49.53	34.36
-	ArEnAV	MesoInception4	V	51.14	41.25	50.77	51.84	45.28	44.09
-	ArEnAV	Xception	V	74.21	52.09	85.36	67.22	68.59	51.70
-	AV-1M & ArEnAV	BA-TFD	AV	75.91	44.31	77.64	58.29	72.21	45.21
-	AV-1M & ArEnAV	BA-TFD+	AV	79.97	27.44	84.20	36.47	72.89	34.56

Implementation Details: We benchmark **temporal detection** using SOTA models: Meso4, MesoInception4, Xception, BA-TFD, and BA-TFD+. BA-TFD and BA-TFD+ (Cai et al., 2023b) are evaluated in their original configurations, both in a zero-shot setting (pre-trained on AV-1M; (Cai et al., 2023a)) and after fine-tuning on our dataset. For image-based classifiers: Meso4, MesoInception4 (Afchar et al., 2018); and Xception (Chollet, 2017), we aggregate frame-level predictions to segments following Cai et al. (2023a). For benchmarking **deepfake detection**, image-based models (Meso4, MesoInception4, and Xception) are trained on video frames with corresponding labels, and predictions are aggregated to video-level using max voting, as suggested by Cai et al. (2023a). Additionally, we assess zero-shot performance of LLM-based models, VideoLLaMA2 and VideoLLaMA2.1-AV (Zhang et al., 2023), prompting them to produce a confidence score indicating the likelihood of a video being a deepfake. We include an audio-only baseline, XLSR-Mamba (Xiao & Das, 2025), the best open-source audio deepfake detection model on Speech DF Arena (Face, 2025), evaluating it both in zero-shot mode (pre-trained on ASVSpoof-2019 (Wang et al., 2020)) and after training with video-level labels from our dataset. BA-TFD and BA-TFD+ (Cai et al., 2022) are also evaluated using segmentation proposals treated as frame-level predictions and aggregated by max-voting, both pre-trained on AV-1M and fine-tuned on our dataset. For all finetuning runs, we subsample the frames so as to remove class imbalance.

432 Table 11: (a): Temporal localization comparison on ArEnAV, AV-1M and LAVDF. (b): Cross-Dataset
 433 comparison (% AUC) of recent SOTA models.

(a) Cross-dataset Deepfake Localization.					(b) Cross-dataset deepfake detection. P: DFDC-P set.						
Method	Dataset	AP@0.5	AP@0.95	AR@50	AR@10	Method	Venue	ArEnAV	DFDC	FF++	CelebDF
BA-TFD	LAV-DF	79.15	0.24	64.18	58.51	Capsule-v2	ICASSP-19	49.15	–	93.11	–
	AV-1M	37.37	0.02	45.55	30.66	Face-X-Ray	CVPR-20	55.56	80.92	98.52	80.58
	ArEnAV	2.42	0.01	22.30	2.54	LipForensics	CVPR-21	49.76	73.50	97.10	82.40
BA-TFD+	LAV-DF	96.30	4.44	80.48	78.75	M2TR	ICMR-22	50.12	–	99.92	–
	AV-1M	44.42	0.03	48.86	34.67	LAA-Net	CVPR-24	50.04	86.94(P)	99.96	–
	ArEnAV	3.74	0.04	30.75	3.05	ForensicsAdaptor	CVPR-25	50.58	88.70	–	94.00

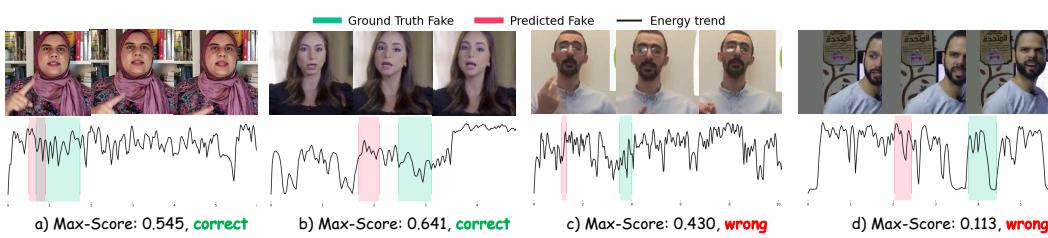
5 RESULTS AND ANALYSIS

432 **Audio-Visual Temporal Deepfake Localization.** The results for temporal localization are shown
 433 in Table 9. SOTA methods show significantly lower performance on ArEnAV as compared to other
 434 localization datasets (refer to Table 11a). BA-TFD and BA-TFD+, pretrained on AV-1M, show a
 435 drop in performance of more than 35% for AP@0.5 threshold, compared to evaluation on AV-1M.
 436 The image-based models, Meso4 and MesoInception4, also provide low performance, which can be
 437 attributed to the use of diffusion-based lip-sync models, which have been overlooked in previous
 438 data generation pipelines (Cai et al., 2023a;b). Through this benchmark, we claim that the highly
 439 realistic multimodal multilingual code-switched fake content in ArEnAV will open an avenue for
 440 further research on temporal multilingual deepfake localization methods.

441 **Audio-Visual Deepfake Detection.** The detection results are in Table 10. Image based models, that
 442 have access to video-level labels only, perform considerably worse, except XLSR-Mamba, which is
 443 designed to be trained on video-level labels for audio-deepfake detection. The best performing model
 444 is BA-TFD, pretrained on AV-1M and then further fine-tuned on our dataset, with AUC Score of 82%
 445 on the full subset. We also evaluate models on subsets V and A, as described in the implementation
 446 details. The audio-only model, XLSR-Mamba, performs better in the Audio-only *subset A*, while
 447 the image-only models perform better on *Subset V* for frame-level labels, compared to the *fullset*.
 448 XLSR-Mamba performs relatively worst when the audio is code-switched, compared to only Arabic.
 449

450 **Cross-Dataset Comparison for Deepfake Localization.** Table 11a shows the performance of
 451 BA-TFD and BA-TFD+ (Pretrained on AV-1M) on LAVDF, AV-1M and ArEnAV datasets. Both
 452 models perform significantly worse on ArEnAV, highlighting the poor generalizability (while it
 453 generalizes to LAV-DF) in multilingual and code-switching settings. BA-TFD and BA-TFD+ fail to
 454 generalize effectively, as the pretrained audio and video encoders struggle with out-of-distribution
 455 data encountered in both modalities of ArEnAV.

456 **Cross-Dataset Comparison for Deepfake Detection.** Table 11b shows the cross dataset performance
 457 of recent SOTA deepfake detection models, including Capsule-v2 (Nguyen et al., 2019), Face-X-Ray
 458 (Li et al., 2020a), LipForensics (Haliassos et al., 2021) and M2TR (Wang et al., 2022). All models
 459 were pretrained on FaceForensics++ (Rössler et al., 2019). While models show great cross-dataset
 460 performance on DFDC and CelebDF, they fail to perform better than guessing (50% AUC) on



461 Figure 4: Different output cases from evaluation of BA-TFD+ after finetuning on ArEnAV. Here,
 462 the ground truth of all samples is **FAKE**. Max-Score refers to the maximum score assigned to a
 463 candidate range during prediction. Correct means that the predicted class matches the ground truth
 464 class. The green region refers to the ground truth fake-segment, and the red region refers to the
 465 predicted fake-segment, based on the Max Score. a) Shows the model predicting the correct class,
 466 along with some overlap with the ground truth fake segment. b) Shows the model predicting the
 467 correct class, but with no overlap. c) and d) Show the model predicting the wrong class in the output.

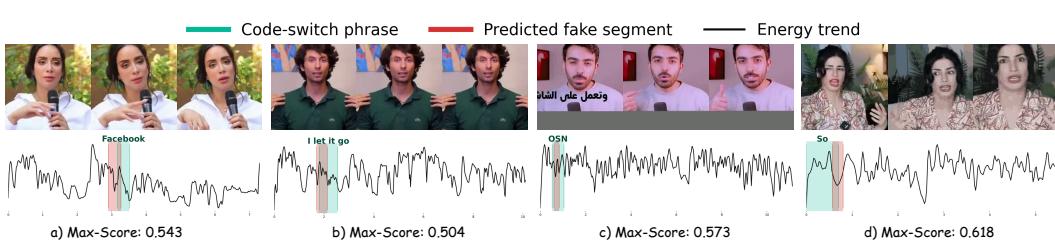


Figure 5: Different output cases from evaluation of BA-TFD+ after finetuning on ArEnAV. Here, the ground truth of all samples is **REAL**. Max-Score refers to the maximum score assigned to a candidate range during prediction. The green region refers to the code-switching region in the real video, and the red region refers to the predicted fake-segment. a),b),c),d) All samples show that the model misclassifies real videos as fakes, specifically at the code-switching regions.

ArEnAV. Even recent SOTA models, such as ForensicsAdaptor (CVPR-25) (Cui et al., 2025b) and LAA-Net (CVPR-24) (Nguyen et al., 2024), fail to generalize. The demographic and linguistic homogeneity of existing datasets (FF++, CelebDF, DFDC) limits model robustness. By incorporating multilingual audio and broader participant diversity, our dataset demonstrates why architectures must be designed to generalize beyond those biases.

6 QUALITATIVE ANALYSIS

Figures 4 and 5 show qualitative examples for different output cases of BA-TFD+, finetuned on ArEnAV. Figure 4 illustrates model predictions on fake samples, showing four representative cases comparing predicted fake segments with ground-truth regions. In some cases, the model correctly identifies the manipulated class with good temporal overlap, while in others, the predictions are either misaligned or incorrect. Importantly, these outputs demonstrate that the model’s behavior is not influenced by any boundary-level artifacts or splicing cues introduced during data generation. Further, spectrograms from BA-TFD+ outputs show no abrupt energy shifts or discontinuities near edit boundaries, indicating the absence of splice artifacts at the boundaries of manipulated content.

Figure 5 shows the model’s predictions on real code-switched videos. Here, the model frequently misclassifies real videos as fake, primarily due to the presence of code-switching between Arabic and English. The high predicted fake scores at these regions indicate that the model confuses natural linguistic transitions with synthetic inconsistencies. Together, these qualitative results confirm that the challenge in ArEnAV arises from the intrinsic complexity of code-switching rather than from generation artifacts.

7 CONCLUSION

This paper presents ArEnAV, a large multilingual and the first code-switching audio-visual dataset for temporal deepfake localization and detection. The comprehensive benchmark of the dataset utilizing SOTA deepfake detection and localization methods, indicates a significant drop in performance compared to previous monolingual datasets, indicating that the proposed dataset is an important asset for building the next generation of multilingual deepfake localization methods.

Limitations. Similar to other deepfake datasets, ArEnAV exhibits a misbalance in terms of the number of fake and real videos. Due to the limited performance of current SOTA Active-Voice-Detection (Whisper v2) models on Arabic (compared to English), the data generation pipeline can result in a few noisy transcripts. Due to limited instruction following in code-switching scenarios, LLMs might not produce the desired results, as visible in Figure 3 "Meaning + Translation Scenario". Compared to other subsets, Chat-GPT often fails to follow both instructions, making real and fake transcripts too similar and not always changing their meaning. Also, the dataset is currently limited to two languages only, where we hope to motivate further research in this direction.

Broader Impact. ArEnAV’s diverse and realistic English-Arabic fake videos will support the development of more robust audio-visual deepfake detection and localization models, better equipped to handle code-switched speech and real-world multilingual scenarios.

540 **8 ETHICS STATEMENT**
 541

542 Our work on ArEnAV raises important ethical considerations, especially given the sensitivity of deep-
 543 fake research. The dataset is built from publicly available YouTube content, in line with established
 544 practices in benchmarks. Use of such material for non-commercial research is covered under fair
 545 use, and access to ArEnAV is gated by a strict End-User License Agreement (EULA)(Section A.7).
 546 Below, we detail all the ethical considerations regarding our work:

547 **Use of YouTube Videos** : The ethical foundation of our data collection does not rely on VisPER but
 548 on established practices in prior peer-reviewed datasets such as LRS3-TED (Afouras et al., 2018b),
 549 VoxCeleb2 (Chung et al., 2018b), and AVSpeech (Ephrat et al., 2018a), which employ the same
 550 keyword-search and face-detection pipeline. We build our dataset from public YouTube videos under
 551 the research-focused “fair use” exception established in peer-reviewed work (e.g., Zhu et al. (2024)).
 552 Access is granted only after users agree to our EULA, which lists the following rules and regulations:

- 554 • Access will be granted only to researchers who supply their university IRB application ID,
 555 and every project member must use an individual account, safeguarding traceability and
 556 preventing misuse.
- 557 • Users are eligible to conduct independent research at their respective institutions and the
 558 Institution accepts responsibility for its Authorized Investigators’ actions related to the use
 559 of ArEnAV.
- 560 • Limits use to academic, non-commercial, not-for-profit research and education.
- 561 • Authorizes licensors to modify the data or license at any time and prohibits licensees from
 562 altering the database.
- 563 • Forbids any use that could cause subjects embarrassment or mental anguish.

564 This approach accords with established practice across the community, as evidenced by DF40 (Yan
 565 et al., 2024), which draws real videos and images from FaceForensics++ (Rössler et al., 2019),
 566 Celeb-DF (Li et al., 2020b), CelebA (Liu et al., 2015), FFHQ (Karras et al., 2019), and VFHQ (Xie
 567 et al., 2022); DeepfakeBench (Yan et al., 2023), which relies on FaceForensics++ and Celeb-DF;
 568 FaceForensics++; Celeb-DF; FakeAVCeleb (Khalid et al., 2021), which builds on VoxCeleb2 (Chung
 569 et al., 2018a); AVLIPS (Liu et al., 2024), which sources from LRS3 (Afouras et al., 2018a) and
 570 FaceForensics++; and AV-1M (Cai et al., 2024b), which is derived from VoxCeleb2. Together, these
 571 measures and precedents demonstrate that curating public YouTube content for non-commercial
 572 scientific inquiry is a responsible and widely adopted practice.

573 **Face Detection techniques applied on videos:** We acknowledge the risks of working with videos that
 574 contain faces, but face detection is used only as a preprocessing step and not for identification. In line
 575 with recent peer-reviewed works, using videos containing faces, for different research problems that
 576 involve face detection as a common prior step, is a standard practice, e.g.: a) LRS3-TED (Afouras
 577 et al., 2018b), VoxCeleb1 (Nagrani et al., 2017), VoxCeleb2 (Chung et al., 2018b) have been used
 578 for Speaker Identification, Verification, Recognition, and further, for Deepfake benchmark creation
 579 (AV-1M (Cai et al., 2024a) and FakeAVCeleb (Khalid et al., 2022)) b) MultiTalk (Sung-Bin et al.,
 580 2024) uses videos from YouTube for Talking Head generation c) AVSpeech (Ephrat et al., 2018b)
 581 used for Speech Separation d) Hallo3 (Cui et al., 2025a) used for Portrait Image Animation.

582 Since it is impractical to get individual consent for open-source content, we mitigate misuse by requir-
 583 ing institutional IRB approval, individual researcher accounts for access and a removal mechanism to
 584 request the removal of personal content.

585 **Human Study:** Our human study followed university IRB guidelines: participants were over 18
 586 years of age and were approached over email through connections in research groups within the
 587 affiliated universities of the authors. The participation was strictly voluntary and anonymous. All the
 588 details about the research project and conditions for participation in the study were clarified through
 589 an Explanatory Statement at the beginning of the user study form. Thus, the users consented to
 590 participate in the study by filling out and submitting the study form (Google form), and all material
 591 was screened to avoid disturbing content. No personal data were recorded and no compensation was
 592 provided.

593

594 **9 REPRODUCIBILITY STATEMENT**
595

596 Our data will be open-sourced. Data-generation code and evaluation scripts will be made public for
597 various open-sourced models evaluated.
598

599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

648 REFERENCES
649

650 Darius Afchar, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. MesoNet: a Compact Facial
651 Video Forgery Detection Network. In *2018 IEEE International Workshop on Information Forensics
652 and Security (WIFS)*, pp. 1–7, December 2018. ISSN: 2157-4774.

653 Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. Lrs3-ted: A large-scale dataset for
654 visual speech recognition. *arXiv preprint*, arXiv:1809.00496, 2018a.

655 Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. Lrs3-ted: a large-scale dataset for
656 visual speech recognition, 2018b. URL <https://arxiv.org/abs/1809.00496>.

658 Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework
659 for self-supervised learning of speech representations, 2020. URL <https://arxiv.org/abs/2006.11477>.

661 Zhixi Cai, Kalin Stefanov, Abhinav Dhall, and Munawar Hayat. Do You Really Mean That? Content
662 Driven Audio-Visual Deepfake Dataset and Multimodal Method for Temporal Forgery Localization.
663 In *2022 International Conference on Digital Image Computing: Techniques and Applications
664 (DICTA)*, pp. 1–10, Sydney, Australia, November 2022.

666 Zhixi Cai, Shreya Ghosh, Aman Pankaj Adatia, Munawar Hayat, Abhinav Dhall, and Kalin Stefanov.
667 AV-Deepfake1M: A Large-Scale LLM-Driven Audio-Visual Deepfake Dataset, November 2023a.
668 arXiv:2311.15308 [cs].

669 Zhixi Cai, Shreya Ghosh, Abhinav Dhall, Tom Gedeon, Kalin Stefanov, and Munawar Hayat. Glitch
670 in the matrix: A large scale benchmark for content driven audio-visual forgery detection and
671 localization. *Computer Vision and Image Understanding*, 236:103818, November 2023b. ISSN
672 1077-3142.

674 Zhixi Cai, Abhinav Dhall, Shreya Ghosh, Munawar Hayat, Dimitrios Kollias, Kalin Stefanov, and
675 Usman Tariq. 1m-deepfakes detection challenge, 2024a. URL <https://arxiv.org/abs/2409.06991>.

677 Zhixi Cai, Shreya Ghosh, Aman Pankaj Adatia, Munawar Hayat, Abhinav Dhall, Tom Gedeon,
678 and Kalin Stefanov. Av-deepfake1m: A large-scale llm-driven audio-visual deepfake dataset. In
679 *Proceedings of the 32nd ACM International Conference on Multimedia (MM '24)*, 2024b. doi:
680 10.1145/3664647.3680795.

682 Edresson Casanova, Julian Weber, Christopher Shulby, Arnaldo Candido Junior, Eren Gölge, and
683 Moacir Antonelli Ponti. Yourtts: Towards zero-shot multi-speaker tts and zero-shot voice conver-
684 sion for everyone, 2023. URL <https://arxiv.org/abs/2112.02418>.

685 Edresson Casanova, Kelly Davis, Eren Gölge, Görkem Göknar, Iulian Gulea, Logan Hart, Aya Aljafari,
686 Joshua Meyer, Reuben Morais, Samuel Olayemi, and Julian Weber. Xtts: a massively multilingual
687 zero-shot text-to-speech model, 2024. URL <https://arxiv.org/abs/2406.04904>.

689 Francois Chollet. Xception: Deep Learning With Depthwise Separable Convolutions. In *Proceedings
690 of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1251–1258, 2017.

694 Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. Voxceleb2: Deep speaker recognition. In
695 *Interspeech 2018*, pp. 1086–1090, 2018a. doi: 10.21437/Interspeech.2018-1929.

696 Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. VoxCeleb2: Deep Speaker Recognition. In
697 *Interspeech 2018*, pp. 1086–1090. ISCA, September 2018b.

698 Jiahao Cui, Hui Li, Yun Zhan, Hanlin Shang, Kaihui Cheng, Yuqi Ma, Shan Mu, Hang Zhou, Jingdong
699 Wang, and Siyu Zhu. Hallo3: Highly dynamic and realistic portrait image animation with video
diffusion transformer, 2025a. URL <https://arxiv.org/abs/2412.00733>.

700 Xinjie Cui, Yuezun Li, Delong Zhu, Jiaran Zhou, Junyu Dong, and Siwei Lyu. Forensics adapter:
701 Unleashing clip for generalizable face forgery detection, 2025b. URL <https://arxiv.org/abs/2411.19715>.

702 Alexandre Defossez, Gabriel Synnaeve, and Yossi Adi. Real time speech enhancement in the
 703 waveform domain. In *Interspeech*, 2020.

704

705 Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cris-
 706 tian Canton Ferrer. The deepfake detection challenge (dfdc) dataset, 2020a. URL <https://arxiv.org/abs/2006.07397>.

707

708 Brian Dolhansky, Joanna Bitton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cris-
 709 tian Canton Ferrer. The DeepFake Detection Challenge (DFDC) Dataset, October 2020b. arXiv:
 710 2006.07397 [cs].

711

712 Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T.
 713 Freeman, and Michael Rubinstein. Looking to listen at the cocktail party: a speaker-independent
 714 audio-visual model for speech separation. *ACM Transactions on Graphics*, 37(4):1–11, July 2018a.
 715 ISSN 1557-7368. doi: 10.1145/3197517.3201357. URL <http://dx.doi.org/10.1145/3197517.3201357>.

716

717 Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T.
 718 Freeman, and Michael Rubinstein. Looking to Listen at the Cocktail Party: A Speaker-Independent
 719 Audio-Visual Model for Speech Separation. *ACM Transactions on Graphics*, 37(4):1–11, August
 720 2018b. ISSN 0730-0301, 1557-7368.

721

722 Ethnologue. Ethnologue: Languages of the world. <http://www.ethnologue.com>, 2025.
 723 Accessed: 2025-09-25.

724

725 Hugging Face. Speech df arena - speech-area-2025. <https://huggingface.co/spaces/Speech-Arena-2025/Speech-DF-Arena>, 2025. Accessed: 2025-05-13.

726

727 Charles A. Ferguson. Diglossia. *Word*, 15(2):325–340, 1959.

728

729 Joel Frank and Lea Schönherr. Wavefake: A data set to facilitate audio deepfake detection, 2021.
 730 URL <https://arxiv.org/abs/2111.02813>.

731

732 Alexandros Haliassos, Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Lips don't lie:
 733 A generalisable and robust approach to face forgery detection, 2021. URL <https://arxiv.org/abs/2012.07657>.

734

735 Injy Hamed, Ngoc Thang Vu, and Slim Abdennadher. ArzEn: A speech corpus for code-switched
 736 Egyptian Arabic-English. In Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri,
 737 Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani,
 738 Hélène Mazo, Asuncion Moreno, Jan Odijk, and Stelios Piperidis (eds.), *Proceedings of the
 739 Twelfth Language Resources and Evaluation Conference*, pp. 4237–4246, Marseille, France, May
 740 2020. European Language Resources Association. ISBN 979-10-95546-34-4. URL <https://aclanthology.org/2020.lrec-1.523/>.

741

742 Injy Hamed, Fadhl Eryani, David Palfreyman, and Nizar Habash. ZAEBUC-Spoken: A multilingual
 743 multdialectal Arabic-English speech corpus. In *Proceedings of the 2024 Joint International
 744 Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING
 745 2024)*, pp. 17770–17782, 2024.

746

747 Injy Hamed, Caroline Sabty, Slim Abdennadher, Ngoc Thang Vu, Thamar Solorio, and Nizar Habash.
 748 A survey of code-switched Arabic NLP: Progress, challenges, and future directions. In *Proceedings
 749 of the 31st International Conference on Computational Linguistics*, pp. 4561–4585, 2025.

750

751 Yinan He, Bei Gan, Siyu Chen, Yichun Zhou, Guojun Yin, Luchuan Song, Lu Sheng, Jing Shao,
 752 and Ziwei Liu. ForgeryNet: A Versatile Benchmark for Comprehensive Forgery Analysis. In
 753 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
 754 4360–4369, 2021.

755

756 Yang Hou, Haitao Fu, Chuankai Chen, Zida Li, Haoyu Zhang, and Jianjun Zhao. Polyglotfake: A
 757 novel multilingual and multimodal deepfake dataset, 2024. URL <https://arxiv.org/abs/2405.08838>.

756 Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. DeeperForensics-1.0: A
 757 Large-Scale Dataset for Real-World Face Forgery Detection. In *Proceedings of the IEEE/CVF*
 758 *Conference on Computer Vision and Pattern Recognition*, pp. 2889–2898, 2020.

759

760 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative
 761 adversarial networks, 2019. URL <https://arxiv.org/abs/1812.04948>.

762

763 Hasam Khalid, Shahroz Tariq, Minha Kim, and Simon S. Woo. Fakeavceleb: A novel audio-video
 764 multimodal deepfake dataset. In *Proceedings of the Neural Information Processing Systems Track*
 765 *on Datasets and Benchmarks (NeurIPS Datasets and Benchmarks)*, 2021. arXiv:2108.05080.

766

767 Hasam Khalid, Shahroz Tariq, Minha Kim, and Simon S. Woo. Fakeavceleb: A novel audio-video
 768 multimodal deepfake dataset, 2022. URL <https://arxiv.org/abs/2108.05080>.

769

770 Patrick Kwon, Jaeseong You, Gyuhyeon Nam, Sungwoo Park, and Gyeongsu Chae. KoDF: A
 771 Large-Scale Korean DeepFake Detection Dataset. In *Proceedings of the IEEE/CVF International*
 772 *Conference on Computer Vision*, pp. 10744–10753, 2021.

773

774 Chunyu Li, Chao Zhang, Weikai Xu, Jingyu Lin, Jinghui Xie, Weiguo Feng, Bingyue Peng, Cunjian
 775 Chen, and Weiwei Xing. Latentsync: Taming audio-conditioned latent diffusion models for lip
 776 sync with syncnet supervision, 2025. URL <https://arxiv.org/abs/2412.09262>.

777

778 Lingzhi Li, Jianmin Bao, Ting Zhang, Hao Yang, Dong Chen, Fang Wen, and Baining Guo. Face
 779 x-ray for more general face forgery detection, 2020a. URL <https://arxiv.org/abs/1912.13458>.

780

781 Yuezun Li, Xin Yang, Pu Sun, Honggang Qi, and Siwei Lyu. Celeb-DF: A Large-Scale Challenging
 782 Dataset for DeepFake Forensics. In *Proceedings of the IEEE/CVF Conference on Computer Vision*
 783 and *Pattern Recognition*, pp. 3207–3216, 2020b.

784

785 Weifeng Liu, Tianyi She, Jiawei Liu, Boheng Li, Dongyu Yao, Ziyou Liang, and Run Wang. Lips are
 786 lying: Spotting the temporal inconsistency between audio and visual in lip-syncing deepfakes. In
 787 *Advances in Neural Information Processing Systems (NeurIPS), Poster / Datasets & Benchmarks*
 788 *Track*, 2024.

789

790 Xuechen Liu, Xin Wang, Md Sahidullah, Jose Patino, Héctor Delgado, Tomi Kinnunen, Massimiliano
 791 Todisco, Junichi Yamagishi, Nicholas Evans, Andreas Nautsch, and Kong Aik Lee. ASVspoof
 792 2021: Towards Spoofed and Deepfake Speech Detection in the Wild. *IEEE/ACM Transactions on*
 793 *Audio, Speech, and Language Processing*, 31:2507–2522, 2023. ISSN 2329-9304.

794

795 Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In
 796 *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 3730–3738,
 797 2015.

798

799 Bartłomiej Marek, Piotr Kawa, and Piotr Syga. Are audio deepfake detection models polyglots?,
 800 2024. URL <https://arxiv.org/abs/2412.17924>.

801

802 Soumik Mukhopadhyay, Saksham Suri, Ravi Teja Gadde, and Abhinav Shrivastava. Diff2lip: Audio
 803 conditioned diffusion models for lip-synchronization, 2023. URL <https://arxiv.org/abs/2308.09716>.

804

805 Arsha Nagrani, Joon Son Chung, and Andrew Zisserman. VoxCeleb: a large-scale speaker identifica-
 806 tion dataset. *Interspeech 2017*, pp. 2616–2620, August 2017.

807

808 Kartik Narayan, Harsh Agarwal, Kartik Thakral, Surbhi Mittal, Mayank Vatsa, and Richa Singh. DF-
 809 Platter: Multi-Face Heterogeneous Deepfake Dataset. In *Proceedings of the IEEE/CVF Conference*
 810 *on Computer Vision and Pattern Recognition*, pp. 9739–9748, 2023.

811

812 Sanath Narayan, Yasser Abdelaziz Dahou Djilali, Ankit Singh, Eustache Le Bihan, and Hakim
 813 Hacid. Visper: Multilingual audio-visual speech recognition, 2024. URL <https://arxiv.org/abs/2406.00038>.

810 Dat Nguyen, Nesryne Mejri, Inder Pal Singh, Polina Kuleshova, Marcella Astrid, Anis Kacem,
 811 Enjie Ghorbel, and Djamil Aouada. Laa-net: Localized artifact attention network for quality-
 812 agnostic and generalizable deepfake detection, 2024. URL <https://arxiv.org/abs/2401.13856>.

813

814 Huy H. Nguyen, Junichi Yamagishi, and Isao Echizen. Capsule-forensics: Using Capsule Networks
 815 to Detect Forged Images and Videos. In *IEEE International Conference on Acoustics, Speech and*
 816 *Signal Processing (ICASSP)*, pp. 2307–2311, May 2019. ISSN: 2379-190X.

817

818 Dufou Nick and Jigsaw Andrew. Contributing Data to Deepfake Detection Research, September
 819 2019.

820 OpenAI. Text-to-speech api: tts-1 model. <https://platform.openai.com/docs/guides/text-to-speech>, 2023. Accessed: 2025-05-13.

821

822 OpenAI. Introducing gpt-4.1 in the api. <https://openai.com/index/gpt-4-1/>, 2025.
 823 Accessed: 2025-05-13.

824

825 Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier,
 826 and Michael Auli. fairseq: A fast, extensible toolkit for sequence modeling. In Waleed Ammar,
 827 Annie Louis, and Nasrin Mostafazadeh (eds.), *Proceedings of the 2019 Conference of the North*
 828 *American Chapter of the Association for Computational Linguistics (Demonstrations)*, pp. 48–53,
 829 Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/
 830 N19-4009. URL <https://aclanthology.org/N19-4009/>.

831

832 Zengyi Qin, Wenliang Zhao, Xumin Yu, and Xin Sun. Openvoice: Versatile instant voice cloning,
 833 2024. URL <https://arxiv.org/abs/2312.01479>.

834

835 Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan
 836 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,
 837 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin
 838 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi
 839 Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan,
 Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL
 840 <https://arxiv.org/abs/2412.15115>.

841

842 Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever.
 843 Robust speech recognition via large-scale weak supervision, 2022. URL <https://arxiv.org/abs/2212.04356>.

844

845 Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias
 846 Niessner. FaceForensics++: Learning to Detect Manipulated Facial Images. In *Proceedings of the*
 847 *IEEE/CVF International Conference on Computer Vision*, pp. 1–11, 2019.

848

849 Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias
 850 Nißner. Faceforensics++: Learning to detect manipulated facial images, 2019. URL <https://arxiv.org/abs/1901.08971>.

851

852 SDAIANCAI. Ar-en code-switching textual dataset. <https://huggingface.co/datasets/SDAIANCAI/Ar-En-Code-Switching-Textual-Dataset>, 2025. Accessed: 2025-05-13.

853

854 Neha Sengupta, Sunil Kumar Sahu, Bokang Jia, Satheesh Katipomu, Haonan Li, Fajri Koto, William
 855 Marshall, Gurpreet Gosal, Cynthia Liu, Zhiming Chen, Osama Mohammed Afzal, Samta Kamboj,
 856 Onkar Pandit, Rahul Pal, Lalit Pradhan, Zain Muhammad Mujahid, Massa Baali, Xudong Han,
 857 Sondos Mahmoud Bsharat, Alham Fikri Aji, Zhiqiang Shen, Zhengzhong Liu, Natalia Vassilieva,
 858 Joel Hestness, Andy Hock, Andrew Feldman, Jonathan Lee, Andrew Jackson, Hector Xuguang
 859 Ren, Preslav Nakov, Timothy Baldwin, and Eric Xing. Jais and jais-chat: Arabic-centric foundation
 860 and instruction-tuned open generative large language models, 2023. URL <https://arxiv.org/abs/2308.16149>.

861

862 Kim Sung-Bin, Lee Chae-Yeon, Gihun Son, Oh Hyun-Bin, Janghoon Ju, Suekyeong Nam, and
 863 Tae-Hyun Oh. Multitalk: Enhancing 3d talking head generation across languages with multilingual
 864 video dataset, 2024. URL <https://arxiv.org/abs/2406.14272>.

864 Kartik Thakral, Rishabh Ranjan, Akanksha Singh, Akshat Jain, Richa Singh, and Mayank Vatsa.
 865 ILLUSION: Unveiling truth with a comprehensive multi-modal, multi-lingual deepfake dataset.
 866 In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=qnlG3zPQUy>.
 867

868 Justus Thies, Michael Zollhöfer, Marc Stamminger, Christian Theobalt, and Matthias Nießner.
 869 Face2face: Real-time face capture and reenactment of rgb videos, 2020. URL <https://arxiv.org/abs/2007.14808>.
 870

871 Massimiliano Todisco, Xin Wang, Ville Vestman, Md Sahidullah, Hector Delgado, Andreas Nautsch,
 872 Junichi Yamagishi, Nicholas Evans, Tomi Kinnunen, and Kong Aik Lee. ASVspoof 2019: Future
 873 Horizons in Spoofed and Fake Audio Detection, April 2019. arXiv:1904.05441 [cs, eess].
 874

875 Junke Wang, Zuxuan Wu, Wenhao Ouyang, Xintong Han, Jingjing Chen, Ser-Nam Lim, and Yu-
 876 Gang Jiang. M2tr: Multi-modal multi-scale transformers for deepfake detection, 2022. URL
 877 <https://arxiv.org/abs/2104.09770>.
 878

879 Xin Wang, Junichi Yamagishi, Massimiliano Todisco, Hector Delgado, Andreas Nautsch, Nicholas
 880 Evans, Md Sahidullah, Ville Vestman, Tomi Kinnunen, Kong Aik Lee, Lauri Juvela, Paavo Alku,
 881 Yu-Huai Peng, Hsin-Te Hwang, Yu Tsao, Hsin-Min Wang, Sébastien Le Maguer, Markus Becker,
 882 Fergus Henderson, Rob Clark, Yu Zhang, Quan Wang, Ye Jia, Kai Onuma, Koji Mushika, Takashi
 883 Kaneda, Yuan Jiang, Li-Juan Liu, Yi-Chiao Wu, Wen-Chin Huang, Tomoki Toda, Kou Tanaka,
 884 Hirokazu Kameoka, Ingmar Steiner, Driss Matrouf, Jean-François Bonastre, Avashna Govender,
 885 Srikanth Ronanki, Jing-Xuan Zhang, and Zhen-Hua Ling. Asvspoof 2019: A large-scale public
 886 database of synthesized, converted and replayed speech, 2020. URL <https://arxiv.org/abs/1911.01601>.
 887

888 Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J. Weiss, Navdeep Jaitly, Zongheng
 889 Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, Quoc Le, Yannis Agiomyrgiannakis, Rob Clark,
 890 and Rif A. Saurous. Tacotron: Towards end-to-end speech synthesis, 2017. URL <https://arxiv.org/abs/1703.10135>.
 891

892 Yang Xiao and Rohan Kumar Das. Xlsr-mamba: A dual-column bidirectional state space model for
 893 spoofing attack detection, 2025. URL <https://arxiv.org/abs/2411.10027>.
 894

895 Liangbin Xie, Xintao Wang, Honglun Zhang, Chao Dong, and Ying Shan. Vfhq: A high-quality
 896 dataset and benchmark for video face super-resolution. In *Proceedings of the IEEE/CVF Conference
 897 on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 657–666, 2022.
 898

899 Zhiyuan Yan, Yong Zhang, Xinhang Yuan, Siwei Lyu, and Baoyuan Wu. Deepfakebench: A
 900 comprehensive benchmark of deepfake detection, 2023. URL <https://arxiv.org/abs/2307.01426>.
 901

902 Zhiyuan Yan, Taiping Yao, Shen Chen, Yandan Zhao, Xinghe Fu, Junwei Zhu, Donghao Luo,
 903 Chengjie Wang, Shouhong Ding, Yunsheng Wu, and Li Yuan. Df40: Toward next-generation
 904 deepfake detection, 2024. URL <https://arxiv.org/abs/2406.13495>.
 905

906 Hang Zhang, Xin Li, and Lidong Bing. Video-LLaMA: An Instruction-tuned Audio-Visual Language
 907 Model for Video Understanding, June 2023. arXiv:2306.02858 [cs, eess].
 908

909 Liyun Zhu, Lei Wang, Arjun Raj, Tom Gedeon, and Chen Chen. Advancing video anomaly detection:
 910 A concise review and a new dataset, 2024. URL <https://arxiv.org/abs/2402.04857>.
 911

912

913

914

915

916

917

918 **A APPENDIX**
919920 **A.1 IDENTITY ANALYSIS IN ARENAV**
921922 We conducted a face analysis using InsightFace and DBSCAN to better understand potential identity
923 overlaps across dataset splits. While each split (train, validation, and test) already contains indepen-
924 dent YouTube video IDs, we found a 41.7% identity match between the combined train+val and test
925 sets at a similarity threshold of 0.7.926 To verify whether this overlap influences model performance, we evaluated BA-TFD and BA-TFD+
927 separately on overlapping and non-overlapping subsets. As shown in Table 12, the results remain
928 consistent across both groups, indicating that identity overlap does not affect the models' behavior. We
929 will release a separate identity-disjoint test set in the final version to further strengthen reproducibility
930 and fairness.931 Table 12: Performance of BA-TFD and BA-TFD+ models on ArEnAV test set. Each column shows
932 Accuracy@0.5 and AUC, with the number of evaluated videos in parentheses.
933

Model	Overall (78,400)		Non-Overlapping (45,512)		Overlapping (32,888)	
	Acc	AUC	Acc	AUC	Acc	AUC
BA-TFD+	27.44	79.97	27.58	79.96	27.31	79.99
BA-TFD	44.31	75.91	44.35	75.96	44.27	75.85

934 **A.2 AFFECT OF LIMITED INSTRUCTION FOLLOWING OF GPT-4O-MINI FOR "MEANING +**
935 **TRANSLATION" TASK:**
936937 In some cases, the “meaning + translation” mode produced sentences that were semantically similar
938 to the originals. This can be improved by reprocessing the LLM outputs and filtering based on
939 entailment or cosine distance to enforce greater semantic variation. However, this does not affect our
940 main conclusions. The detection task in ArEnAV focuses on identifying audio-visual inconsistencies,
941 not the extent of semantic change.942 Only a small fraction of the data (less than 8,000 out of 280,000 samples, and about 2,000 out of 78k
943 in the test set) (Figure 3a) has an entailment quality mean > 0.8 .944 To reaffirm this, we compared the best-performing models across these subsets. As shown in Table
945 13, the performance remains consistent, indicating that these few cases do not lower the dataset’s
946 overall difficulty.947 Table 13: Performance of BA-TFD and BA-TFD+ models on ArEnAV test set. Each column shows
948 AUC scores.
949

Model	Overall	entailment <0.8	entailment >0.8
BA-TFD+	79.97	79.96	79.98
BA-TFD	75.91	75.92	75.89

950 **A.3 AFFECT OF DIFFERENT LIP-SYNC AND TTS MODELS ON THE PERFORMANCE:**
951952 Table 14 reports model accuracy across four TTS and 2 lip-sync generation methods (finetuned
953 on ArEnAV). For audio-only detection, XLSR-Mamba achieves the highest accuracy with
954 Fairseq-OpenVoice (95.8%), indicating that lower-quality TTS outputs are more easily detected.
955 Higher-fidelity systems such as XTTS-OpenVoice and XTTS-v2 reduce detectability, suggesting
956 improved synthesis quality. For audio-visual models (BA-TFD and BA-TFD+), both Diff2Lip and
957 LatentSync yield similar accuracy levels, confirming stable and high-quality visual performance.
958 Overall, the results show that better generation quality leads to lower detectability, aligning with our
959 goal of creating challenging, realistic benchmarks.
960961
962
963
964
965
966
967
968
969
970
971

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

Table 14: Accuracy across different Audio and Video generation models used in ArEnAV.

Model Type	Audio Models (Subset A)			Video Models (Subset-V)		
	Fairseq-OpenVoice	GPT-TTS	XTTS-OpenVoice	XTTS-v2	Dif2clip	Latentsync
XLSR-Mamba	95.8	78.8	19.4	77.9	38.6	42.8
BA-TFD+	69.5	58.9	7.7	68.9	31.7	32.9
BF-TFD	33.7	32.2	29.3	31.9	58.1	56.2

A.4 REAL PERTURBATIONS

Table 15: List of video and audio perturbation types with descriptions.

Category	Perturbation Type	Description
Video Perturbations	Gaussian Blur	Applies Gaussian smoothing to simulate out-of-focus capture.
	Salt and Pepper Noise	Random white and black pixel noise, mimicking sensor errors.
	Low Bitrate Compression	Blocky, artifact-heavy images due to compression.
	Gaussian Noise	Electronic sensor noise typical in low-light conditions.
	Poisson Noise (Shot Noise)	Noise from photon-limited imaging environments.
	Speckle Noise	Multiplicative noise creating granular interference effects.
	Color Quantization	Banding effects from limited color palettes.
	Random Brightness	Simulates variations in exposure and lighting.
	Motion Blur	Imitates camera or object motion during capture.
	Rolling Shutter	Distortion effects due to CMOS sensor movements.
	Camera Shake	Minor frame shifts from handheld camera vibrations.
	Lens Distortion	Optical distortions like barrel or pincushion effects.
	Vignetting	Darkening of image edges typical of certain lenses.
	Exposure Variation	Adjusts brightness and contrast, simulating exposure issues.
	Chromatic Aberration	Color channel shifts causing fringing effects.
Audio Perturbations	Compression Artifacts	Quality loss from low bitrate compression.
	Pitch/Loudness Distortion	Gain or frequency alterations simulating recording issues.
	White Noise	Constant background electronic interference noise.
	Time Stretch	Audio speed adjustments without pitch change.
	Reverberation	Echo and reverb modeling room acoustics.
	Ambient Noise	Background environmental sounds added.
	Clipping	Distortion from exceeding audio amplitude limits.
	Frequency Filter	Filtering effects simulating transmission equipment variations.
	Doppler Effect	Pitch modulation due to relative motion.
	Interference	Static-like bursts mimicking external disturbances.
	Room Impulse Response	Complex echo patterns modeling specific environments.

1026 A.5 AUGMENTATION EXAMPLES
10271028 In Table 16, we provide examples of augmentations achieved through the manipulation rules pre-
1029 viously outlined in Section 3.2.1.1030
1031 Table 16: Examples of augmentations achieved through the different transcript manipulation rules,
1032 showing the original (orig) and augmented (aug) transcriptions.

1034 Original 1035 Transcription	1034 Original 1035 Word	1034 Inserted 1035 Word	1034 Operation	1034 Example
				Edit: Telephone → Radio [اِتَّشَغَلَتِ الْهَانِمُ فِي الْتَّلْفُونْ] (orig) (The lady got busy on the telephone)
1036 CSW	1037 EN	1038 EN	1039 Change meaning only (keep English)	1040 [اِتَّشَغَلَتِ الْهَانِمُ فِي الرَّادِيوْ] (aug) (The lady got busy with the radio)
1041	1042 CSW	1043 AR	1044 AR	1045 Edit: محدودة → منتشرة (MSA) [اَصَبَحَتِ اَدَاءُ مُنْتَشِرَةً جَدًّا] (orig) (Mirroring has become a popular tool) [اَصَبَحَتِ اَدَاءُ مُحَدَّدَةً جَدًّا] (aug) (Mirroring has become a limited tool)
1046	1047 CSW	1048 AR	1049 AR	1050 Edit: بَكَرَه → بشَكَرْ (MSA) [بَشَكَرْ كُلَّ الِّلَّيْ مُوجَدِينْ] (orig) (I hate all the present sponsors) [بَكَرَه كُلَّ الِّلَّيْ مُوجَدِينْ] (aug) (I hate all the present sponsors)
1051	1052 Arabic	1053 AR	1054 AR	1055 Edit: حَزِين → سَعِيد (MSA) [وَهِيَكُونُ هَذَا الشَّخْصُ رَاضِيٌّ وَسَعِيدٌ] (orig) (And this person will be content and happy) [وَهِيَكُونُ هَذَا الشَّخْصُ رَاضِيٌّ وَحَزِينٌ] (aug) (And this person will be content and sad)
1056	1057 Arabic	1058 AR	1059 AR	1060 Edit: تَاقَه → جَوَهْرِي (MSA) [كَانَتْ تَشْتَرِكُ بِعَمَلِ اَسَاسِيِّ جَوَهْرِي] (orig) (She was involved in a core and essential task) [كَانَتْ تَشْتَرِكُ بِعَمَلِ اَسَاسِيِّ تَاقَه] (aug) (She was involved in a core and non-essential task)
1061	1062 Arabic	1063 EN	1064 EN	1065 Edit: النَّاسُ → friends [اَنَا بِرُوحِ قَابِلِ النَّاسِ] (orig) (I go meet people) [اَنَا بِرُوحِ قَابِلِ اِنْدِرِنِي] (aug) (I go meet friends)

1080 A.6 PROMPT FOR TEXT PERTURBATION
10811082 Prompt for Fake Transcript Generation.
1083

```

1084     ###SYSTEM MESSAGE###
1085     You are a controlled text-perturbation bot.
1086     Here is the transcript of an audio.
1087     Please use the provided operations to modify
1088     the transcript to change its sentiment.
1089     The operation can be one of `delete`,
1090     `insert` and `replace`.
1091     Please priority modify adjectives and adverbs.
1092     -----CHANGE-MODES-----
1093     • meaning_only
1094         - Change the *meaning* of one word.
1095         - Keep the same language/script and dialect.
1096     • dialect_only
1097         - Swap a word for a dialectal equivalent of *identical meaning*.
1098         - Example: <syArT> → <rbyT> (Gulf dialect, same meaning).
1099     • dialect_plus_meaning
1100         - Change *both* dialect *and* meaning in a single word.
1101         - Example: <jmyl> (msa, 'nice') → <wH$> (Egyptian, 'awful').
1102     • meaning_plus_translation
1103         - In Arabic-only sentences, pick a word that
1104             is **commonly code-switched
1105             to English** in everyday speech (e.g., <mwbayl>, <syArT>, <Antrnt>).
1106             - Translate that word to English and change the
1107                 meaning simultaneously.
1108             Example: <syArT> ('car') → bike.
1109     -----CSW MULTI-OP LOGIC-----
1110     If language == 'csw':
1111         num = 1 → edit exactly one token matching target_token_script.
1112         num = 2 → edit 1 English + 1 Arabic token.
1113         num = 3 → edit 1 English + 2 Arabic tokens.
1114     -----OTHER RULES-----
1115     • Only modify tokens that are *commonly code-switched* in real speech
1116         (brand names, technology, everyday nouns, etc.).
1117     • Each operation targets ONE word (delete / insert / replace).
1118     • Number of operations for INSERT, DELETE and REPLACE
1119         should be equal across
1120             the data.
1121     • If sentiment can be changed with INSERT or DELETE,
1122         prefer it over REPLACE.
1123     • When dialect shifts, include original_dialect and new_dialect.
1124     • Never alter tense or add restricted content.
1125     • Return **only** a JSON object that matches the schema.
1126
1127
1128
1129
1130
1131
1132
1133

```

Figure 6: System prompt for text-perturbation bot

1134 A.7 END USER LICENSE AGREEMENT (EULA FORM)
1135

1136

1137 End User License Agreement.

1138

1139 End User License Agreement
(Academic, non-commercial, not-for-profit licence)

1140

Copyright (c) 2025[AUTHORS]

1141

All rights reserved.

1142

The goal of the ArEnAV database is to develop new techniques, technology, and algorithms
for multimodal, code-switched deepfake detection and localization, as most of the existing research
→ focuses on monolingual content, often overlooking the challenges of multilingual and
→ code-switched speech, where multiple languages are mixed within the same discourse. The licensors
→ are involved in an ongoing effort to strengthen detection algorithms against highly realistic
→ deepfakes. The dataset is meant to aid research efforts in the general area of developing,
→ testing and evaluating algorithms for multilingual code-switched deepfake detection and
→ localization.

1143

1144 To receive a copy of the dataset, the requester must agree to observe the conditions listed Below.

1145

The goal of the ArEnAV database is to develop new techniques, technology, and
algorithms for predicting and locating (with timestamps) where a video has been
manipulated, particularly when it has Arabic-English code-switching. Use is permitted of the
→ databases and annotations above in source and binary form, provided that the following
conditions are met:

1146

- The database is provided under the terms of this license strictly for academic, non-commercial, not-for-profit purposes.
- Requestor needs to supply their university IRB application ID, and every project member must use an individual account, safeguarding traceability and preventing misuse. Attach the IRB approval in the email along with the signed EULA form.
- Redistribution, republishing, or dissemination in any form, source or binary, is not permitted without prior written approval by the licensors. Linking to the webpage of the database [WEB LINK → HERE] is permitted.
- The names of the licensors may not be used to endorse or promote products derived from this software without specific prior written permission.
- The licensors reserve the right to modify the data/license at any point.
- Modification of the database by licensees is not permitted.
- In no case should the still frames or videos be used in any way that could cause the original subject embarrassment or mental anguish.
- You understand that the ArEnAV dataset is a deepfake dataset generated based on VisPer ([2406.00038] ViSpeR: Multilingual Audio-Visual Speech Recognition) dataset's Arabic Train subset. You also agree to all agreements of the VisPer dataset.
- The authors of the dataset make no representations or warranties regarding the dataset, including but not limited to warranties of non-infringement or fitness for a particular purpose.
- You accept full responsibility for your use of the dataset and shall defend and indemnify the Authors of ArEnEV, against any and all claims arising from your use of the dataset, including but not limited to your use of any copies of copyrighted images that you may create from the dataset.
- Any publications arising from the use of this software, including but not limited to academic journal and conference publications, technical reports and manuals, must cite the following works:
[CITATION]

1147

THE DATABASE IS PROVIDED BY THE AUTHORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS DATABASE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. THE PROVIDER OF THE DATABASE MAKES NO REPRESENTATIONS AND EXTENDS NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED. THERE ARE NO EXPRESS OR IMPLIED WARRANTIES THAT THE USE OF THE MATERIAL WILL NOT INFRINGE ANY PATENT, COPYRIGHT, TRADEMARK, OR OTHER PROPRIETARY RIGHTS. If you have read and understood the user agreement and will comply with it.

1148

Signed

1149

Print Name

1150

Institution Name

1151

Date

1152

Addition Researcher 1

1153

Addition Researcher 2

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

Figure 7: End User License Agreement for accessing ArEnAV.

1188 A.8 LLM USAGE
11891190 Along with the use of Large Language Models (LLMs) as described in our Data-Creation process, we
1191 made limited use of LLMs to enhance the clarity and readability of the text. They were not involved
1192 in the conception of ideas, the design of experiments, analysis, or the production of results.

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241