



Developing Speech Processing Pipelines for Police Accountability

Anjalie Field, Prateek Verma, Nay San, Jennifer L. Eberhardt, Dan Jurafsky

Stanford University, USA

{anjalief, prateekv, nay.san, jleberhardt, jurafsky}@stanford.edu

Abstract

Police body-worn cameras have the potential to improve accountability and transparency in policing. Yet in practice, they result in millions of hours of footage that is never reviewed. We investigate the potential of large pre-trained speech models for facilitating reviews, focusing on ASR and officer speech detection in footage from traffic stops. Our proposed pipeline includes training data alignment and filtering, fine-tuning with resource constraints, and combining officer speech detection with ASR for a fully automated approach. We find that (1) fine-tuning strongly improves ASR performance on officer speech (WER=12-13%), (2) ASR on officer speech is much more accurate than on community member speech (WER=43.55-49.07%), (3) domain-specific tasks like officer speech detection and diarization remain challenging. Our work offers practical applications for reviewing body camera footage and general guidance for adapting pre-trained speech models to noisy multi-speaker domains.

Index Terms: speech recognition, accountability, policing, social applications, noisy domains

1. Introduction

Over the last decade, police departments across the United States have rapidly adopted body-worn cameras (BWCs) [1]. This rapid adoption has been spurred on by widespread protests demanding improved accountability and transparency following high-profile deaths of civilians involving officers' use of force [2, 3]. In some ways, BWCs have resulted in improvements: the footage is valuable evidence in instances such as litigation of excessive force cases [4, 5], and analysis of hand-transcribed footage can identify racial disparities in policing and failures to practice procedural justice [6, 7, 8]. However, in the absence of a lawsuit or high-profile incident, most footage is never reviewed. Further, reliance on manual transcriptions limits the scalability of existing automated analyses [6, 9, 8].

At the same time, large pre-trained speech models have achieved remarkable performance over standardized datasets [10, 11, 12, 13, 14]. Models like Whisper and Wav2Vec2 also have demonstrated potential in social good applications, e.g., in monitoring audio(visual) materials related to long-term elderly care [15] or child exploitation [16]. However, in applications involving multi-speaker conversations in noisy environments, models require application-specific adaptation and evaluation [17, 18, 19]. Little work has investigated the speech processing of police BWC footage specifically.

Here, we develop and evaluate automatic speech recognition (ASR) and police officer speech detection (diarization) for police BWC footage. Automatic transcription of officer speech would allow extending existing text analyses of racial bias in

hand-transcriptions to new data without requiring expensive transcription efforts [6, 8]. It would also allow departments to determine adherence to a procedure by using text classifiers [7] or keyword searches. Although most reviews are likely to be internal, some departments publicly release BWC footage or are mandated to provide access upon request [20, 1]. Thus, speech-processing technology could support independent audits.

Our primary data is footage from 1,040 vehicle stops conducted by one department in one month, where utterances spoken by officers and community members were previously hand-transcribed. We use the data to construct training and test data sets for ASR and officer speech detection. We evaluate ASR models, with and without in-domain fine-tuning, over the entire test set, dividing by role (officer or community member), race, and gender, and we examine the performance of officer speech detection in combination with ASR.

Our findings provide insight into the best practices and limitations of developing technology in this domain. For example, our training data processing pipeline is robust enough that fine-tuning improves ASR performance by 3-11 points. We also show evidence that Whisper models learn to mimic transcribers' representations of transcription confidence by marking difficult segments as unintelligible. Differences by gender and race are not significant; however, ASR over officer speech (WER=12-13% for officers unseen in training) is much more accurate than over community member speech (WER=43.55-49.07%), which suggests that models have a high potential for addressing accountability with less risk of compromising community member privacy [20]. Finally, we identify diarization, specifically officer speech detection, as a continued challenge.

2. Data

Video recordings of the 1,040 vehicle stops and hand-transcriptions were provided to us under a data use agreement for the management of such high-risk data and under IRB supervision. The data is generally noisy. Prior transcripts were intended for language analysis, rather than the development of speech processing tools, so not all speech was transcribed and diarized.¹ Stops contain background noise like wind and traffic. They contain multiple speakers, and secondary officers, as well as drivers and passengers, can be situated far from the recording device. Dispatch speech from officers' radios can often be heard, sometimes directly overlapping with utterances from the primary interaction. There is high variance in the clarity of speech and quality of footage across stops.

Test and Validation Sets. To create reliable test and valida-

¹The transcribers were instructed to transcribe only speech by officers and community members, not police dispatch; they inconsistently included officer speech to dispatch (vs. to the community member).

Alignment	Robust	Whisper	Prop. of Final data
	W2V2 WER	Large WER	
Unaligned	65.51	56.78	13.86
MFA	68.42	54.84	12.01
MFA Chunked	61.11	42.04	32.32
W2V2	60.25	43.27	12.72
W2V2 Chunked	68.0	52.27	29.10

Table 1: WER over the full training set (78K utterances) under each alignment method and what percentage of training data were ultimately aligned with each method.

tion sets, we hand-align existing transcribed utterances to timestamps and correct observed transcription errors. To facilitate analysis by race, we chose the test data to consist of 50%/50% stops of white and black drivers. We also choose each test file to be a stop by a distinct officer and withhold any other stops made by the same officers (whether as primary or secondary officers) from the training and validation sets. Thus, we also selected officers who made a small number of stops to minimize unusable data. Hand-aligning data is extremely time consuming, so we restrict test set stops to contain < 60 utterances. We similarly ensure there is no overlap in primary officers between the validation and training set, withholding data as needed, though we less strictly enforce the separation of secondary officers, who speak less frequently. We conduct evaluations over these aligned utterances, discarding un-transcribed speech.

Training Set Alignment. We build a training set by applying automated alignment tools and filtering poor-quality transcriptions. We determine the start and end time for each transcribed utterance using the best of 5 alignment methods:

- Unaligned: 1sec granularity timestamps hand-written by transcribers with heuristics to correct for obvious typos and extending the start and end by .25sec
- MFA: Montreal Forced Aligner [21] with unaligned timestamps as starting points
- MFA chunked: Many utterances are too short for the aligner to process correctly. Thus, using the unaligned timestamps, we chunk consecutive utterances up to a total of 20sec. We run MFA to obtain word-level timestamps and then divide chunks back into separate utterances, with start and end times determined by the word-level timestamps
- W2V2: Robust Wav2Vec2 [13] for forced alignment [22]
- W2V2 chunked: Same as MFA chunked, but using Robust Wav2Vec2 for forced alignment instead of MFA.

For each utterance, we use off-the-shelf Whisper Large [14] and Robust Wav2Vec2 (W2V2) [13] to transcribe the audio segment identified by each alignment method and compare the output with the hand-written transcript. We choose as the final alignment the one for which $\min(WER_{Whisper}, WER_{W2V2})$ is lowest. Table 1 reports training WER for each alignment method and the percent of the final training data aligned using each method.

Training Set Filtering. Even after alignment, the training data is noisy, containing, for example, transcription errors, overlapping speech, and unfixd alignment errors. We again use $\min_{WER} = \min(WER_{Whisper}, WER_{W2V2})$ over the best alignment to filter out training instances that are likely incorrect. We experiment with four filtering criteria, indicating filtered training data size in brackets:

1. Remove instances < 0.5 sec and > 10 sec [54,600]
2. #1, and remove instances where $\min_{WER} > 50\%$ [40,361]

	# Stops	# Utterances	Speech Time
Train	795	78,082	61.85hr
Train (filtered)	787	26,121	17.61hr
Validation	8	373	21.24min
Test	20	634	32.41min

Table 2: Final data set sizes. Across the full data set, there are an average of 91.73 utterances and 3.2 speakers per stop.

3. We define $WER_{[nosubs.]}$ as WER where we do not count substitutions as errors. This metric is designed to retain instances where there may be errors in the Whisper/Wav2Vec2 outputs (e.g., WER is high) but likely not alignment errors (e.g., WER is driven by substitutions rather than insertions or deletions). We then filter according to #1, and keep only instances where $(\min_{WER_{[nosubs.]}} < 10\% \text{ AND } \min_{WER} < 50\%)$. [26,121]
4. #1, and remove instances where $\min_{WER} > 10\%$ [19,759]

We compare each criteria by using the filtered training data to fine-tune Robust Wav2Vec2 and examining performance over the validation set. Criteria #3 (WER=45.23) and #4 (WER=44.92) perform similarly and both outperform #1 (WER=49.34) and #2 (WER=48.75). We use #3 when training subsequent models, favoring the criteria that keeps more training data. Table 2 reports the final sizes for each data split.

3. ASR

We compare the performance of ASR models off-the-shelf and fine-tuned on the training data set constructed in Section 2. We use two of the current best-performing and most popular architectures: Wav2Vec2 [10] and Whisper [14]. For Wav2Vec2, we use the Robust model [13], which was pre-trained using a self-supervised objective on Libri-Light, CommonVoice, Switchboard, Fisher and fine-tuned for ASR on Switchboard. For Whisper, which was trained on 680,000 hours of multilingual and multitask data, we compare *small*, *medium*, and *large* [14]. Thus, both models are intended to perform well in a variety of domains and over noisy data. We describe the model training parameters in detail, including the use of decoder-only training for Whisper large due to compute constraints.

3.1. Experimental Setup

To fine-tune Wav2Vec2, we use model default parameters with learning rate=1e-5, weight decay=0.005, warmup steps=500, batch size=32. We report performance with and without a 4-gram language model trained over the training data transcripts, implemented with KenLM and integrated with beam size=1500, lm weight=1.31 and word score=1.31.²

For Whisper models without fine-tuning, we hard-code the task as transcription and the language as English. For fine-tuning, we use model default parameters with learning rate=1e-5, and warmup steps=500. Our experiments are conducted in a resource-constrained environment. Data protocols mandate that the footage be stored on a secure restricted-access server, which does not have sufficient GPU memory to fine-tune Whisper large, even with reduced batch size and precision. Thus, we experiment with freezing the encoder and just training the decoder as well as the inverse. We use a batch size of 32 for Whisper small and 16 for medium and large. Finally, as Whisper is

²lm weight and word score were tuned following the Bayesian optimization procedure in [10]. We do no other hyperparameter tuning.

prone to outputting repeated words and phrases, we remove any words from the model output if they occur > 10 times.

As transcription norms vary between corpora and the body-camera gold transcripts contain bracketed terms like *[unintelligible]* and *[laughter]*, we remove all terms in brackets and use the Whisper text normalizer on both the reference and model output before computing WER for all models (including Wav2Vec2 models). For all models, we choose the checkpoint with the lower validation WER after 5 epochs and train using 1-2 A40 GPUs. Wav2Vec2 and Whisper small models trained in < 5 hrs; Whisper medium and large models trained in < 16 hrs.

3.2. Results

3.2.1. Overall ASR

Tuned Params	WER _S	CER _S	WER _L	CER _L
None	34.75	24.86	33.07	28.47
Encoder	34.30	23.53	22.82	17.58
Decoder	28.12	20.07	22.26	16.86
Enc+Dec	26.07	18.76	-	-

Table 3: *Whisper Small (S) and Large (L) validation performance with encoder-only or decoder-only training*

Table 3 reports validation results (reserving the test set for final configurations) of freezing either the encoder or decoder when fine-tuning Whisper large and small. For Whisper small, decoder-only tuning performs almost comparably to tuning the entire model (28.12 vs., 26.07), whereas tuning only the encoder performs less well (34.30). For Whisper large, freezing the encoder or decoder provides advantages over no fine-tuning, though decoder-only tuning converged faster (2 vs. 5 epochs). Subsequently, we use decoder-only training for the fine-tuned the Whisper large model.

Wav2Vec2	WER	CER	Whisper	CER	WER
[None]	45.01	31.57	Small	32.13	22.83
+LM	38.91	31.27	Small+Tune	22.09	16.30
+Tune	42.20	26.05	Med.	26.21	18.36
+Tune+LM	32.29	25.97	Med.+Tune	23.47	17.78
			Large	29.60	22.35
			Large+Tune	18.33	13.61

Table 4: *ASR Results over police test set.*

Table 4 reports the overall WER and CER for each model. Whisper large with fine-tuning performs the best overall. Fine-tuning gives improves performance by 3-11pts across models.

Reference	Whisper	Whisper(tuned)
Yeah. I know, I'm trying to-	I'll turn it.	Yeah. [unintelligible].
Yeah. [unintelligible] expired like la- December.	The fire started in December.	Yeah. [unintelligible] expired like December.
[unintelligible].	It's going to be a bad traffic.	It's going to be a bad traffic.

Table 5: *Test outputs of fine-tuned Whisper large*

As Whisper is a new model with yet-limited work on understanding model performance and fine-tuning effects, we highlight a few examples from the data in Table 5. In the original transcripts, transcribers mark segments they are unable to

	W2V2	W2V2 tuned+LM	Whisp.	Whisp. tuned
Role [Officer]	-.440*	-.435*	-.791*	-.383*
Race [Black]	-.028	-.026	-.350	-.034
Gender [F]	.106	.088	.215	.091
CM Black [120]	83.67	66.53	66.53	43.55
CM White [130]	88.45	74.02	75.05	49.07
Off. Black [175]	42.14	27.26	19.43	13.11
Off. White [166]	32.80	21.95	22.70	12.50

Table 6: *ASR by role/race/gender for Robust W2V2 and Whisper Large (not including 3 Hispanic officers). Top: ASR Mixed Effects Regression. A negative (starred if significant) coefficient indicates lower WER (better performance). Bottom: WER for each subgroup. Brackets indicate number of test utterances*

decipher as *[unintelligible]*. While we removed all bracketed text when computing WER rate for fair comparison of off-the-shelf and fine-tuned models, examining Whisper outputs reveals that the fine-tuned model sometimes outputs *[unintelligible]*. In some instances, the predicted *[unintelligible]* exactly aligns with hand-transcription. However, we also find examples where Whisper hallucinates transcriptions for difficult content, whereas Wav2Vec2 more often does not produce output. After fine-tuning, Whisper hallucinations are particularly difficult to identify without referring back to the audio, as they often appear to be plausible statements in an interaction.

3.2.2. Performance by officer/driver, gender, and race

We examine model performance over sub-populations of the test data, specifically distinguishing between officers and community members, black and white people, and men and women. As there is high variance in model performance depending on the quality of footage from each stop, we use a mixed effects linear regression model. Each data point in the regression is a single utterance. The dependent variable is model WER for the utterance. Role (officer or community member), race, gender are fixed effects, and the specific stop is a random effect.

Table 6 reports the learned regression coefficients and WER by sub-population for the best performing Wav2Vec2 and Whisper models, off-the-shelf and fine-tuned. ASR performance for officers is significantly better than performance for community members by a wide margin. Even the best-performing models perform poorly at transcribing community member speech. Community members are situated further from the camera and typically speak very few short utterances. Even hand-transcribers often mark their speech as unintelligible, and training a high-performing model on this type of data may be infeasible. This result suggests that ASR could be an extremely useful tool for police accountability with small potential privacy-reducing impact on community members.

In contrast to prior work, we do not find significant differences by race or by gender [23]. Subdividing the test data leads to small data set sizes, which could be skewed by a single outlying stop. This potential effect is greater when looking at race and gender than looking at role, since a low-quality video would decrease ASR performance for both the officer and the community member, whereas in examining race and gender, we are comparing across footage of different stops. Table 6 does show WER is lower for white than black officers for most models.

4. Officer Speech Detection

In Section 3, we use hand-aligned evaluation data, but in practice, we do not know segmentation or speaker identities in new footage. As our goal is police accountability, we develop two models to identify segments of speech by primary officers (e.g., officers wearing the camera) and evaluate them using the best-performing ASR model over the detected speech.

4.1. Methodology

Training Data Processing We adapt the training set introduced in Section 2. We remove any instances that do not contain active speech using an off-the-shelf acoustic scene understanding Mobile-Net [24] architecture trained on AudioSet [25] (AudioSet category $0 < 0.3$). We divide remaining samples into 250ms chunks with a 100ms hop and represent each 250ms segment as a mel-spectrogram with 64 mel-filters, computed with a hop of 10ms, and a window of 25ms. We create a balanced training corpus by randomly sampling 150K chunks each of officer/non-officer speech. Since officers are closer to body-camera microphones (near-field) than community members (far-field), we use volume-based data augmentation.

As the raw training data contains non-officer speech that was not transcribed (e.g., dispatch speech), we also *augment* the training set. We divide training files into 250ms chunks with a 100ms hop, keep chunks with a speech score (from the Mobile-Net model) ≥ 0.5 , and merge consecutive chunks that occur within 1sec of each other. We add all new segments (ones that were not transcribed) to the training data as instances of not-officer-speech and then filter and sample the data as described above. We use these data to train models to classify 250ms chunks as officer or not-officer speech (with cross-entropy loss).

In-domain classifier We train a custom model from scratch, which contains 7 convolutional layers with 128 3×3 filters in every layer and Relu activation followed by max-pooling of 2. The output of the last layer is passed onto a linear head of 1024 neurons, followed by softmax activation, and the posterior probability is taken as officer score for that instance.

Universal d-vectors We extract d-vectors as features from an off-the-shelf model trained over the VoxCeleb dataset for speaker recognition [26] and train an officer speech classifier, with the same linear-head architecture as the in-domain model.

Inference We predict voice activity detection (using the same Mobile-Net model) and officer scores for 250ms chunks with 100ms hops. We consider a chunk to be officer speech if its voice activity score is $> t_{\text{VAD}}$ and its officer score is $> t_{\text{officer}}$, and we merge positive chunks if they occur within t_{smooth} sec of each other.³ For evaluation, we concatenate the ASR model output for all identified segments and compute WER against similarly concatenated hand-aligned officer segments.

4.2. Results

Table 7 reports results for the best performing ASR model over the automatically detected officer speech segments. There is a substantial performance decrease between the hand-aligned segments and the detected segments. The d-vector model performance particularly poorly, likely due to the high difference in domain between VoxCeleb and police traffic stops. Augment-

³ $\{t_{\text{VAD}}, t_{\text{smooth}}, t_{\text{officer}}\}$ are hyperparameters chosen via 20-iteration Bayesian optimization over the validation set with range $[0,1]$ for $t_{\text{VAD}}/t_{\text{officer}}$ and $[0.25,2]$ for t_{smooth} . They are $\{0.93, 1.76, 0.16\}$ for d-vector, $\{0.4, 0.67, 1.2\}$ for in-domain, and $\{0.52, 0.51, 1.1\}$ for in-domain [aug.]

Model	WER	CER	%S	%D	%I
d-vector [26]	61.85	53.14	15.50	25.07	21.27
In-domain	49.47	39.83	15.15	21.10	13.22
In-domain [aug.]	31.52	25.29	7.94	11.64	11.93
Hand-aligned	12.80	8.68	5.98	3.97	2.85

Table 7: ASR results over officer detected speech using tuned Whisper Large. S:substitutions, D:deletions, and I:insertions

ing the training data does substantially improve performance (49.47 to 31.52 WER), though performance still may not be sufficient for applications. In reviewing model outputs, we identify that models often misidentify other near-field speech as officer speech, and the presence of multiple officers complicates the task, as speech by secondary officers is sometimes scored closer to community member speech. We also identified several annotation errors, such as segments attributed to the wrong person and inconsistencies in which speech was transcribed, suggesting these metrics may under-estimate performance. These errors could be removed in hand-aligned test data, but their presence in training data is still likely to degrade model performance, and manually re-cleaning training data (as opposed to automatic augmentation) would involve a substantial undertaking that may not generalize to other settings.

5. Discussion

We find pre-trained ASR models achieve low WER over police officer speech, particularly when fine-tuned on automatically cleaned training data. Whisper specifically achieves low WER and even learns to mimic transcribers in marking segments as unintelligible, but can still fail more dramatically over difficult segments than Wav2Vec2. While prior work has identified ASR as a limitation in noisy speech domains [17, 18], we instead find that officer speech detection is a significant challenge in this setting. There are potential avenues for improvement, such as explicitly modeling dispatch and secondary officer speech or using text-based classifiers over ASR outputs [19]. Further, although WER is worse over detected than hand-aligned officer speech, WER is an imperfect proxy metric for tasks actually of interest, such as determining officers’ adherence to procedure. As many errors are driven by misidentified or short utterances, performance may still be sufficient for tasks like dialog act classification [7]. While we focus on policing, our work has the potential to inform adapting ASR models to other noisy multi-speaker domains as well.

Limitations and Ethical Considerations Our data consists of traffic stops from one police department. We cannot predict if results generalize to data from other departments, time periods or types of police-community interactions. Also, although all work abides by IRB and data sharing protocols, it has high misuse potential. Models could be used for purposes other than police accountability, such as community surveillance. Because models were trained on private data and pending mitigation of potential misuse, we are not releasing them at this time.

Acknowledgements

This research was supported by the John D. and Catherine T. MacArthur Foundation (G-1512-150464 and G-1805-153038). We thank Stanford Data Science for fellowship funding as well as the city and the police department whose provision of camera footage enabled this research. This work was administered and supported by Stanford SPARQ, a center that builds research-driven partnerships to combat bias, reduce disparities, and drive culture change. We thank Martijn Bartelds, Dora Demszky, Rebecca Hetey, and Tolulope Ogunremi for helpful feedback.

6. References

- [1] S. Zansberg, "Public access to police body-worn camera recordings (status report 2020)," *American Bar Association*, January, vol. 22, 2021.
- [2] C. Lum, M. Stoltz, C. S. Koper, and J. A. Scherer, "Research on body-worn cameras: What we know, what we need to know," *Criminology & public policy*, vol. 18, no. 1, pp. 93–118, 2019.
- [3] N. Elkins-Brown, S. James, and L. James, "Validating a novel tool for coding body worn camera footage of police-community member interactions," *Criminology & Public Policy*, 2023.
- [4] M. Zamoff, "Assessing the impact of police body camera evidence on the litigation of excessive force cases," *Ga. L. Rev.*, vol. 54, p. 1, 2019.
- [5] S. Çubukçu, N. M. Sahin, E. Tekin, and V. Topalli, "Body-worn cameras and adjudication of citizen complaints of police misconduct," National Bureau of Economic Research, Tech. Rep., 2021.
- [6] R. Voigt, N. P. Camp, V. Prabhakaran, W. L. Hamilton, R. C. Hetey, C. M. Griffiths, D. Jurgens, D. Jurafsky, and J. L. Eberhardt, "Language from police body camera footage shows racial disparities in officer respect," *Proceedings of the National Academy of Sciences*, vol. 114, no. 25, pp. 6521–6526, 2017.
- [7] V. Prabhakaran, C. Griffiths, H. Su, P. Verma, N. Morgan, J. L. Eberhardt, and D. Jurafsky, "Detecting institutional dialog acts in police traffic stops," *Transactions of the Association for Computational Linguistics*, vol. 6, pp. 467–481, 2018. [Online]. Available: <https://aclanthology.org/Q18-1033>
- [8] E. H. Rho, M. Harrington, Y. Zhong, R. Pryzant, N. P. Camp, D. Jurafsky, and J. L. Eberhardt, "Escalated police stops of black men are linguistically and psychologically distinct in their earliest moments," *Proceedings of the National Academy of Sciences*, vol. 120, no. 23, 2023.
- [9] N. P. Camp, R. Voigt, D. Jurafsky, and J. L. Eberhardt, "The thin blue waveform: Racial disparities in officer prosody undermine institutional trust in the police," *Journal of personality and social psychology*, vol. 121, no. 6, p. 1157, 2021.
- [10] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [11] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," *Proc. of Interspeech*, pp. 3465–3469, 2019.
- [12] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [13] W.-N. Hsu, A. Sriram, A. Baevski, T. Likhomanenko, Q. Xu, V. Pratap, J. Kahn, A. Lee, R. Collobert, G. Synnaeve *et al.*, "Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training," *arXiv preprint arXiv:2104.01027*, 2021.
- [14] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," *arXiv preprint arXiv:2212.04356*, 2022.
- [15] C. Hacking, H. Verbeek, J. P. H. Hamers, and S. Aarts, "The development of an automatic speech recognition model using interview data from long-term care for older adults," *Journal of the American Medical Informatics Association*, vol. 30, no. 3, pp. 411–417, 12 2022. [Online]. Available: <https://doi.org/10.1093/jamia/ocac241>
- [16] J. C. Vásquez-Correa and A. Álvarez Muniain, "Novel speech recognition systems applied to forensics within child exploitation: wav2vec2.0 vs. Whisper," *Sensors*, vol. 23, no. 4, 2023. [Online]. Available: <https://www.mdpi.com/1424-8220/23/4/1843>
- [17] S. Dutta, S. A. Tao, J. Reyna, R. E. Hacker, D. Irvin, J. Buzhardt, and J. H. Hansen, "Challenges remain in building ASR for spontaneous preschool children speech in naturalistic educational environments," *Proc. of Interspeech*, pp. 4322–4326, 2022.
- [18] A. Abulimiti, J. Weiner, and T. Schultz, "Automatic speech recognition for ILSE-interviews: Longitudinal conversational speech recordings covering aging and cognitive decline." in *Proc. of Interspeech*, 2020, pp. 3795–3799.
- [19] J. Zuluaga-Gomez, S. S. Sarfjoo, A. Prasad, I. Nigmatulina, P. Motlicek, K. Ondrej, O. Ohneiser, and H. Helmke, "BERTTraffic: BERT-based joint speaker role and speaker change detection for air traffic control communications," *arXiv preprint arXiv:2110.05781*, 2021.
- [20] R. Lin, "Police body worn cameras and privacy: Retaining benefits while reducing public concerns," *Duke L. & Tech. Rev.*, vol. 14, p. 346, 2015.
- [21] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using Kaldi." in *Proc. of Interspeech*, vol. 2017, 2017, pp. 498–502.
- [22] L. Kürzinger, D. Winkelbauer, L. Li, T. Watzel, and G. Rigoll, "CTC-segmentation of large corpora for German end-to-end speech recognition," in *Speech and Computer: 22nd International Conference, SPECOM 2020, St. Petersburg, Russia, October 7–9, 2020, Proceedings 22*. Springer, 2020, pp. 267–278.
- [23] A. Koenecke, A. Nam, E. Lake, J. Nudell, M. Quartey, Z. Mengesha, C. Toups, J. R. Rickford, D. Jurafsky, and S. Goel, "Racial disparities in automated speech recognition," *Proceedings of the National Academy of Sciences*, vol. 117, no. 14, pp. 7684–7689, 2020. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.1915768117>
- [24] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.
- [25] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [26] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4879–4883.