

SpeechCompass: Enhancing Mobile Captioning with Diarization and Directional Guidance via Multi-Microphone Localization

Artem Dementyev*
Google Research
Mountain View, CA
artemd@google.com

Mathieu Parvaix
Google Research
Mountain View, CA
parvaix@google.com

Dimitri Kanevsky
Google Research
Mountain View, CA
dkanevsky@google.com

Chiong Lai
Google Research
Mountain View, CA
chionglai@google.com

Samuel J. Yang
Google Research
Mountain View, CA
samuely@google.com

Alex Olwal*
Google Research
Mountain View, CA
olwal@acm.org

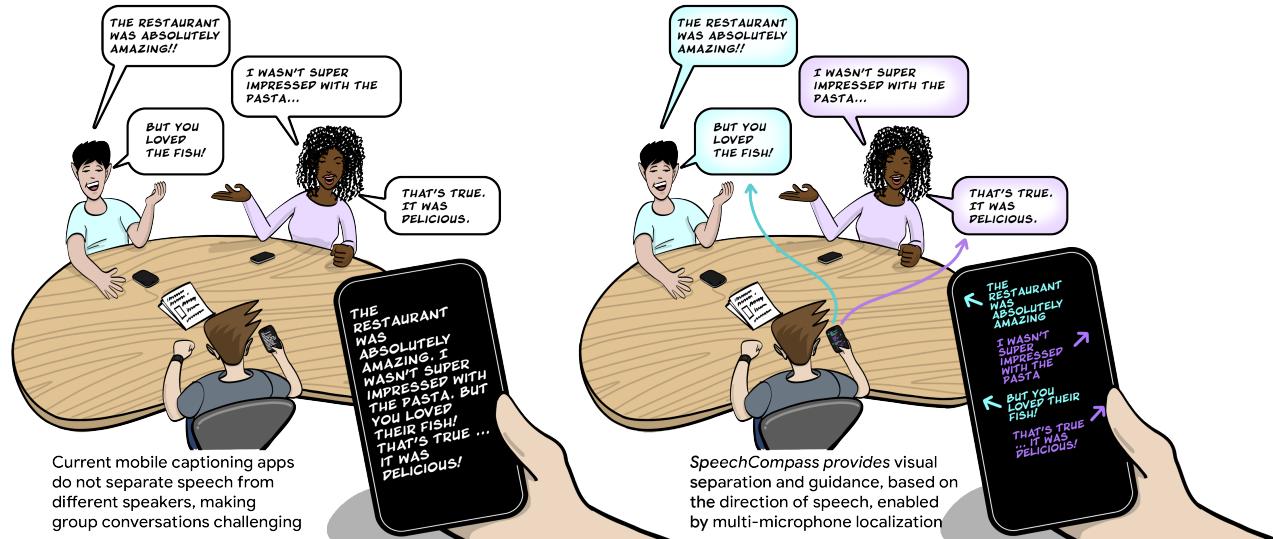


Figure 1: *SpeechCompass* creates user-friendly speech transcripts for group conversations with multiple speakers. **Left:** Current solutions concatenate and mix the transcribed speech when multiple people participate in a conversation, which makes it challenging to read and understand the transcript. **Right:** *SpeechCompass* addresses this limitation through real-time, multi-microphone speech localization, where the direction of speech allows diarization, visual separation, and guidance (e.g., arrows) in the user interface.

Abstract

Speech-to-text capabilities on mobile devices have proven helpful for hearing and speech accessibility, language translation, note-taking, and meeting transcripts. However, our foundational large-scale survey (n=263) shows that the inability to distinguish and indicate speaker direction makes them challenging in group conversations. *SpeechCompass* addresses this limitation through real-time, multi-microphone speech localization, where the direction of speech allows visual separation and guidance (e.g., arrows) in the

user interface. We introduce efficient real-time audio localization algorithms and custom sound perception hardware, running on a low-power microcontroller with four integrated microphones, which we characterize in technical evaluations. Informed by a large-scale survey (n=494), we conducted an in-person study of group conversations with eight frequent users of mobile speech-to-text, who provided feedback on five visualization styles. The value of diarization and visualizing localization was consistent across participants, with everyone agreeing on the value and potential of directional guidance for group conversations.

*First and last author contributed equally to this work.



This work is licensed under a Creative Commons Attribution 4.0 International License.
CHI '25, Yokohama, Japan

© 2025 Copyright held by the owner/author(s).
ACM ISBN 978-x-xxxx-xxxx-x/2025/04
https://doi.org/10.1145/3706598.3713631

CCS Concepts

• Human-centered computing → Accessibility technologies;
Human computer interaction (HCI); Mobile devices.

Keywords

Assistive technology, hearing accessibility, localization, diarization, microphone array, captioning

ACM Reference Format:

Artem Dementyev, Dimitri Kanevsky, Samuel J. Yang, Mathieu Parvaix, Chiong Lai, and Alex Olwal. 2025. SpeechCompass: Enhancing Mobile Captioning with Diarization and Directional Guidance via Multi-Microphone Localization. In *CHI Conference on Human Factors in Computing Systems (CHI '25), April 26-May 1, 2025, Yokohama, Japan*. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3706598.3713631>

1 Introduction

Recent advances have enabled real-time automatic speech recognition (ASR) in mobile and embedded hardware, supporting a range of conversational scenarios [2, 46]. Real-time captioning can enhance human communication in various ways, e.g., for real-time translation between languages, transcriptions during an interview, automatic subtitle generation, hearing accessibility [38, 55], and note-taking. However, there are still unsolved limitations with real-time speech-to-text – specifically, the ability to distinguish multiple speakers (*speaker diarization*) and tracking the direction of speech (*localization*).

Consider the following scenario: *Throughout the day, a person relies on a mobile phone with real-time captioning to understand speech. However, at a dinner table with multiple people, the conversation is difficult to follow since the app cannot distinguish between speakers and concatenates all speech into a single paragraph. Additionally, since the person needs to look at their phone regularly, they struggle with following the turn-taking across speakers and don't know where to look if there is a speaker change. Also, irrelevant nearby conversations get transcribed as well, which can cause confusion and privacy implications.*

Many of the challenges in the scenario stem from the spatial complexities of audio, which are challenging to capture with a single-microphone ASR system. The benefits of multi-microphone topologies for localization have been demonstrated in numerous applications, such as public safety [61], virtual reality [39], robot navigation [43], mobile computing [36, 49] and audio accessibility [27, 31]. In this work, we leverage arrays of multiple microphones and apply techniques for microphone array signal processing to demonstrate how this technology could improve ASR performance and usability in such scenarios. Specifically, we identified opportunities for improvements in three areas:

- (1) *Speaker diarization*. The transcript can visually separate different speakers based on the direction of the speech.
- (2) *Localization*. For spatial sound visualization, the screen can display the direction of the incoming sound.
- (3) *Selective attention*. The interface can allow the user to select speech of interest and filter out self-speech.

In this work, we developed *SpeechCompass*, a solution to add diarization and speech localization to mobile captioning. It includes three main parts. First, low-latency localization algorithms that can run on generic microcontrollers or mobile phones. Second, a compact 4-microphone phone case that allows 360-degree localization on a low-power microcontroller. Third, a mobile captioning app that shows how sound localization can be visualized in different

ways and used to support multi-speaker conversations through diarized transcripts. We also run our algorithm on an off-the-shelf mobile phone with only two microphones to demonstrate that limited 180-degree localization is possible in the app without additional hardware.

While machine learning approaches to single-source speaker diarization have been improving [52], our multi-mic approach has the advantage of lower computational cost, latency, and privacy, and thus is inherently suitable for real-time ASR applications on low-power, low-cost microcontrollers. It is also language agnostic and can work for sounds other than speech. Our approach is tied to the position of the phone and the speaker, which can be advantageous as diarization can be immediately reconfigured by moving the phone. This paper shows how traditional microphone array processing can significantly benefit diarization and localization for mobile ASR. While diarization with microphone arrays has been well studied [4], it has yet to be applied to mobile ASR and related interfaces.

1.1 Contributions

The contributions of this work are:

- **Real-time sound localization algorithms and optimized, embedded multi-microphone hardware**, implemented both as phone-case prototypes with a low-power microcontroller (<20 ms processing time) and on a mobile phone with constrained microphone hardware. The code and design files are available via Github.¹
- **Mobile captioning UIs for group conversations, enabled through speech localization and diarization**, and implemented as mobile speech-to-text applications with different visualization and interaction techniques.
- **Technical evaluation of localization and diarization using the optimized algorithms and hardware**. We characterize localization accuracy and estimation time under various signal-to-noise conditions and speaker configurations, and evaluate diarization accuracy.
- **Foundational large-scale survey (n=263) with frequent users of captioning technology**, showing that noise and speaker separation are important and frequent challenges with existing solutions.
- **Lab study (n=8) and large-scale survey (n=494) of mobile interfaces and visualizations with frequent users of captioning technology**. Informed by a large-scale survey, we designed a lab study with eight frequent users of captioning technology. We contribute insights into interface preferences, customization, and benefits since diarization and localization have not been previously studied for mobile ASR.

2 Related work

This section first discusses previous works in relevant audio applications like diarization, real-time audio transcription, and visualization interfaces. Then, we discuss the relevant literature on audio multi-microphone processing.

¹github.com/google/multi_mic_audio

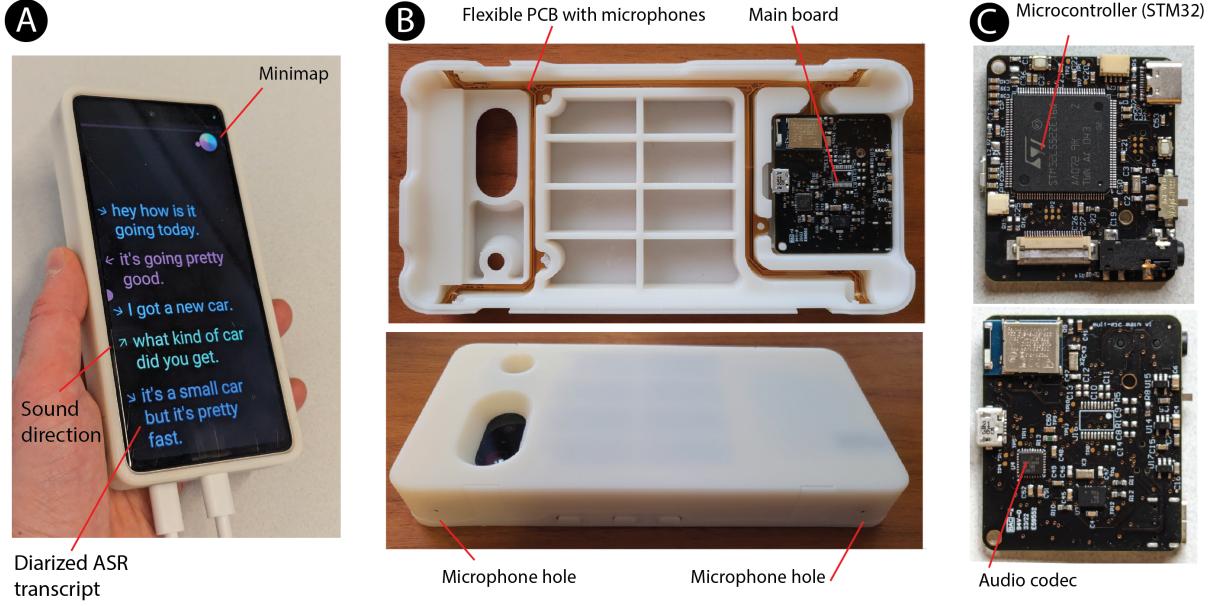


Figure 2: Overview of the SpeechCompass phone case prototype. A) A mobile phone application interface with a mounted multi-microphone phone case. B) Inside and outside view of the prototype with a flexible PCB microphone mount and a compact main printed circuit board (PCB). C) Pictures of the main PCB with a top and bottom view.

2.1 Real-time mobile speech recognition

CART (Communication Access Realtime Translation) is a well-established method for providing real-time captions. A trained professional, typically using specialized software and a stenography machine, transcribes speech into text as it is spoken. CART is frequently used in broadcasting scenarios such as lectures and presentations. Recent advances in machine learning enabled real-time automatic mobile speech captioning, such as Live Transcribe [2] and Microsoft Translator [46], and also on head-worn displays [50]. One of the main motivating uses for audio-to-text translation was audio accessibility, and research has shown the benefits of real-time ASR for hearing accessibility in scenarios like classrooms [13, 38]. Research has shown various improvements to real-time ASR, such as ways to display transcription confidence [9] and customizing fonts appearance to improve readability [8].

2.2 Speaker separation and diarization: Transcripts that distinguish speakers

Studies suggest that there is still room for improvements in, e.g., accuracy and usability [22]. We particularly observe that existing approaches have yet to leverage microphone arrays in mobile devices to augment ASR. We outline the existing techniques in Table 1.

Acoustic beamforming relies on classical signal processing techniques such as beamforming and localization from multiple microphones [4] to separate and diarize speakers. The main challenge for speaker separation is estimating a beamformer for each speaker using localization and other cues. Recently, neural networks have been employed to successfully aid in beamforming [28, 65]. However, acoustic beamforming has been mainly applied to meeting

room scenarios with a static microphone array, while SpeechCompass uses a similar multimicrophone technology for localization, which is applicable to mobile ASR throughout a user's everyday life.

Blind source separation. This approach separates the speakers using a single microphone without additional cues. The technique is challenging for classical signal processing, but various machine-learning techniques have been successful [19, 30]. Blind speaker separation is effective when done offline on the entire audio file, providing the model with access to both future and past content. This is, however, not possible in a real-time causal system, as in this paper, since the future is not accessible, and the model only processes a limited amount of past information. These constraints for real-time causal systems make blind separation inapplicable for real-time transcription of conversations.

Recently, multi-microphone approaches have been combined with speaker separation and diarization. In [59] and [60] Taherian et al. tackle the challenge of speaker separation in multi-speaker scenarios, focusing on conversational or meeting environments. They leverage multi-channel audio and deep learning models to improve separation accuracy to enhance the performance of downstream speech applications like ASR. In one project [59], the authors leverage an end-to-end diarization system to identify individual speaker activity and then use this information to guide a multi-talker separation model. In another approach [60] a multi-input multi-output (MIMO) complex spectral mapping model allows for robust speaker localization, and is used to reduce speaker splitting errors. The complexity of the processing and non-causal components make these solutions unsuitable for real-time processing.

Table 1: Comparison of speaker diarization and separation technologies. The comparison shows that only SpeechCompass can support diarization and visualize sound direction on mobile devices.

	SpeechCompass [this work]	Acoustic beamforming [5, 25]	Audio-visual [15, 18, 31]	Multiple mic blind source separation [55]	Single mic blind source separation [16, 27]	Audio embeddings as voice fingerprint [54, 58]
Sound direction (360°)	Yes	Yes	No	Yes	No	No
Mobile / on-the-go	Yes	No	No	No	Yes	Yes
Computation cost	Low	Low	High	High	High	Medium
Real-time	Yes	Yes	No	No	No	Yes
Calibration free	Yes	Yes	Yes	Yes	Yes	No
Privacy preserving	Yes	Yes	No	Yes	Yes	No
Number of mics	3-4	3+	1	2+	1	1

Voice fingerprint audio embeddings. This approach extracts unique speaker embeddings from a single microphone [57, 62] and uses them for diarization. Principal Component Analysis (PCA) or other unsupervised methods are typically done on the embeddings. The speaker embedding approach has been the primary go-to for real-time diarization since it can run causally. Adding multi-microphone data to speaker embeddings has improved diarization accuracy [57]. A key disadvantage of speaker embeddings is its reliance on implicit or explicit speaker enrollment, as the initial number of speakers is unknown. Requiring every conversation partner or nearby speaker to explicitly register through a voice sample is particularly impractical in dynamic mobile scenarios. Furthermore, there are privacy concerns as speaker embeddings can be considered biometric information, and asking people to enroll would be in conflict with the discreetness that is often desired for accessibility aids.

Audio-visual. Another approach has been to process audio and video using a multimodal model to separate speakers [18, 21]. The camera feed can help infer the active speaker from facial and lip motion when correlated with audio. Researchers have proposed deep learning models for audiovisual speaker separation that operate in the time-frequency domain and use cross-attention for audiovisual fusion [34]. The audio-visual model, while outperforming an audio-only model, however does not leverage spatial information from multiple microphones. The non-causal nature of the separation model (bi-LSTM) makes it ill-suited for real-time applications. This approach works best for meeting rooms and post-processing of recordings, since for mobile ASR, users typically do not point a camera at their conversation partners. There is potential for such applications for head-worn cameras (e.g., in smart glasses), although such approaches would still be dependent on a suitable field of view, sufficient bandwidth, and computing to process video streams. Such approaches will also have power and thermal implications for embedded devices with limited battery size.

Several commercial off-the-shelf solutions exist for mobile speaker diarization. The Ava mobile application [6], for example, allows diarization by connecting each speaker’s smartphone to a network. However, this solution requires every speaker to set up and use their phone, which adds setup overhead. Speaksee [58] is another solution that utilizes clip-on microphones for each speaker, with each microphone exclusively picking up speech from its designated wearer. These microphones connect to a central hub that provides

diarized transcripts. This solution, while effective, requires dedicated hardware, making it more appropriate for formal meeting room scenarios. While our approach may be less accurate than these solutions, because we are not using a dedicated microphone for each speaker, it does not depend on instrumentation or preparation by conversation partners, which is a crucial advantage in real-life situations whether with friends or strangers. It may even be prohibitive in certain scenarios where conversations are very brief (e.g., watercooler conversation) or the person wants to be discreet with their hearing accessibility needs.

2.3 Visualizing non-speech sounds for accessibility use cases

Speech is not the only aspect of sound that has been transcribed. For example, sound event recognition and visual alerts are helpful in hearing accessibility [11, 20, 23, 33]. Expressive captions [15], for example, visually convey the speech’s emotion by altering the text’s rendering based on the detected emotion. Closest to this work, research has looked into displaying sound localization cues from a microphone array using head-worn displays, such as HoloLens [27] and Head-Mounted displays (HMD) [32]. Another work integrates localization into a watch [35] and displays sound location with LEDs. These proof-of-concept research projects use off-the-shelf devices, which do not meet power consumption, ergonomics, or form factor requirements needed for all-day use in everyday lives. In contrast, we leverage mobile phones that users already carry and use daily. While we demonstrate a prototype in a phone case form factor, the embedded hardware could also be implemented in wearable devices, such as head-worn displays.

Due to captioning’s focus on speech, fine audio qualities and structures, such as rhythmicity, are lost. To communicate such qualities, researchers have proposed real-time visualizations such as audio spectrograms [25] or Stabilized Auditory Images (SAI) [44], a visualization grounded in models of the brain, or tactile cues [16]. However, these approaches not only require extensive user training but also do not convey sound direction or leverage spatial information. Our approach can, however, be combined with those existing methods to extend their expressiveness.

2.4 Sound localization strategies

Multiple research projects have investigated sound localization, which can be divided into active and passive localization. Active localization sends out a signal and listens to the response. For example, Cricket [53] showed an indoor localization system using ultrasound beacons and detectors around the room. Active localization does not apply to speech, which is the main target of this paper. As used in this work, a passive localization system listens with an array of microphones. Example of passive localization includes gunshot localization [61], rendering of a sound source in virtual reality [39], robot navigation [43], and people tracking [14]. The Time Difference of Arrival (TDOA) is a popular technique for localization, where measuring the arrival delay between synchronized microphone pairs, allows the use of simple geometry to estimate the angle of arrival.

Multiple localization techniques have been proposed in the literature based on the popular and effective TDOA estimation method. Most are variations [17, 40] on generalized cross-correlation [37], performed in the frequency domain. Other approaches have been demonstrated, such as localization with one microphone [54] by adding structures to the microphone that change the sound based on incoming angle, or using two thin wires that change resistance based on the direction of sound [1].

Recent high-end phones use two or three microphones, which allows limited localization on some phones [36, 49]. However, the microphone placement on these phones is typically optimized for phone calls and video capture, with placements that are less well suited for localization, as microphones on the front and back are blocked when the phone is held or placed on a table.

Several development platforms offer localization. For example, MiniDSP [47] makes eight and sixteen microphone devices for localization and beamforming. ODAS [26] is an open-source platform for microphone array localization. Those platforms are, however, not designed for portable, battery-powered devices, given that their digital signal processing (DSP) chips have high power consumption (~500 mA).

Although mobile phone ASR has many users, previous work shows that diarization and sound localization still need to be addressed, especially to make it practical for conversations with multiple speakers. While machine learning has improved diarization, the approaches are still challenging to apply to mobile ASR due to their numerous limitations (See Table 1). Traditional multi-microphone processing techniques provide promising strategies for mobile ASR, which are practical today as shown through our technical implementation and evaluations.

3 Challenges with mobile captioning: Large-Scale survey with 263 frequent users

As we were interested in exploring the potential for more advanced mobile speech perception, we conducted a brief large-scale survey to learn about the challenges of using captioning for speech understanding in in-person meetings and conversations.

3.1 Participants

We used Google Surveys [24] to deploy a survey to the general population in the US of all ages and genders ("Android users of the

Google Opinion Rewards app"), screening for individuals that use technology to understand speech in meetings and conversations, and are frequent users of captioning technology. Our goal was to recruit deaf or hard-of-hearing participants, as we believed that they would have the most relevant experience and insights around mobile captioning technology and interfaces. To mitigate spam, the survey system analyzes question response times. By considering the distribution of response times across questions, it adapts to different question types and response patterns, rejecting sessions with unusually fast responses.

We acknowledge that our survey's focus on frequent users of captioning technology and the growing user base for mobile captioning limits its relevance for other populations, such as individuals who identify with Deaf culture and might be less likely to rely on ASR technology [10]. Unfortunately, we cannot quantify the representation in the survey since restrictions from our institution do not allow us to collect participant hearing levels or use of sign language.

Of the 1502 respondents that met our criteria, we focus on the 263 participants (18%), who reported that they used captioning technology to understand people (not TV/video) multiple times per week or more frequently, and for 2 hours or more on the days that they used it. For these 263 participants, the platform reported that 49.8% were women, 48.7% men, and 1.5% unknown, across all age ranges (27%: 18–24, 33%: 25–34, 15%: 35–44, 12%: 45–54, 6%: 55–64, 7%: 65+).

The participants were prompted to select challenges among the choices from the list in Figure 3. The choices were synthesized by user feedback from our previous experience with mobile captioning and informed by previous work in mobile captioning [31, 50]

3.2 Survey results: The use of captioning to understand people in conversations

64% of the participants reported daily use of captions in meetings or conversations to understand people, whereas 36% used it multiple times per week. Half of the participants (49%) use technology to understand people face-to-face for 2-3 hours on the day of use. Almost a quarter (23%) of participants use captions for 4-5 hours on the days of use and another quarter (28%) for 6 or more hours. Real-time captions, such as CART (69%), and the *Android Live Transcribe and Sound Notifications* app (55%) were the top two technologies that were used daily to understand people.

The top two issues with current transcription technology, as reported by our participants, were background noise (60%) and the combining of text from different speakers (46%), without the ability to separate them. Participants selected all that apply from the choices shown in Figure 3.

Finally, we asked participants about scenarios that are known to be challenging with today's transcription technology but have the potential to be addressed with more advanced microphone arrays and speech perception algorithms. Scenarios of interest included conversations where ignoring music, noise, or adjacent speech would be critical. We were also interested in group conversations and situations where separating speech from two people is critical. 68-70% of participants experienced these scenarios multiple times per week or more frequently, whereas only 11-12% rarely or never experienced them, as shown in Figure 4.

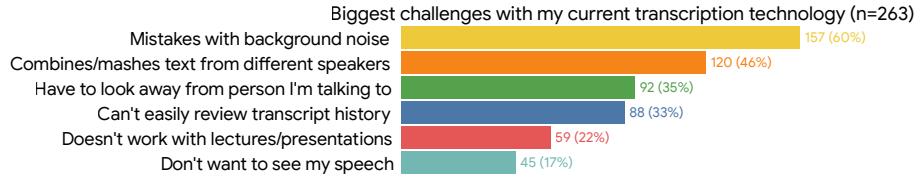


Figure 3: Participant responses to the question *What are the biggest challenges with your current captioning or transcription device/technology? (select all that apply)?*

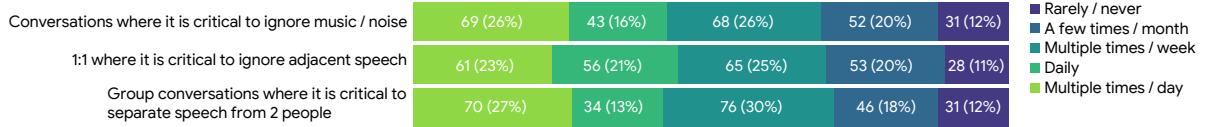


Figure 4: Survey results of how often participants encountered challenging scenarios with today's transcription technology. The number of participants and percentage is shown for each choice.

3.3 Discussion

Our large-scale survey enabled us to identify essential challenges with current transcription technology for frequent users of captioning and shows that 68-70% of the participants are frequently in situations that today's technology cannot adequately support. The findings suggest that more advanced speech technology for suppressing noise from adjacent speakers, music, or noise could help address those issues and that speech separation technology can potentially improve group conversations through more readable transcripts.

4 SpeechCompass system

This section details the SpeechCompass hardware, algorithms, and phone application. The SpeechCompass system diagram is shown in Figure 5.

4.1 Design considerations

Based on the foundational large-scale survey with 263 frequent users of mobile captioning (Section 3), previous work, and envisioned user journey, we outline the following design goals for the SpeechCompass prototype:

- **All-day use.** The prototype should be low-power, ensuring that the battery could last a day or more, as 28% of survey participants use captioning technology for six or more hours daily. Current DSP solutions [26, 47] draw about 500 mAh, which would require impractically large batteries, or significant impact on the phone battery.
- **Accurate 360° localization.** The survey participants indicated that a major challenge is how transcription apps combine text from different speakers, which makes it difficult to use in group conversations. To address this challenge, localization accuracy should be under 20° [45] to match human localization abilities and accurately distinguish different speakers. Also, 360° azimuth localization is needed as it avoids front-back confusion, as illustrated in Figure 6A.

Also, since group conversations are dynamic, speakers and the phone may be located at any angle around the user.

- **Integration with existing phones and applications.** 55% of the survey participants use mobile assistive apps, which is a strong motivator for us to target mobile phones, as they are already always with the user. The device should have multiple input and output options for audio and data, be designed to interface with external applications, and physically integrate with mobile phones.
- **Low-latency, real-time processing.** To integrate with existing applications, the localization needs to be faster than real-time transcription latency, which is expected to be in the 30–300 ms range [66]. Minimizing latency requires hardware-specific algorithm development and bypassing typical operating system latency.

For someone who uses mobile captions daily, adding SpeechCompass for group conversations should be a low effort. The person would use the SpeechCompass phone case instead of a regular one. Their favorite captioning app would automatically use sound localization data to diarize the captions and separate the speakers. Considering the scenario where the person struggles to use mobile captioning around a dinner table, the text from different speakers would be uniquely presented (e.g., using colors or arrows) based on their position around the table. As phones add more microphones, SpeechCompass could become a pure software solution, utilizing multiple microphones on the phone.

4.2 Multi-mic system for embedded perception

We chose a 110 MHz M33 ARM Core (STM32L55, STMicroelectronics) as the main microcontroller. This processor provides the low-power and high computing capabilities needed for localization.

Four digital microphones (MP34DT01-M, STMicroelectronics), which use the pulse density modulation (PDM) protocol, were arranged to have the largest distance between the microphones for finer resolution and support of 360° of azimuth angles. Although

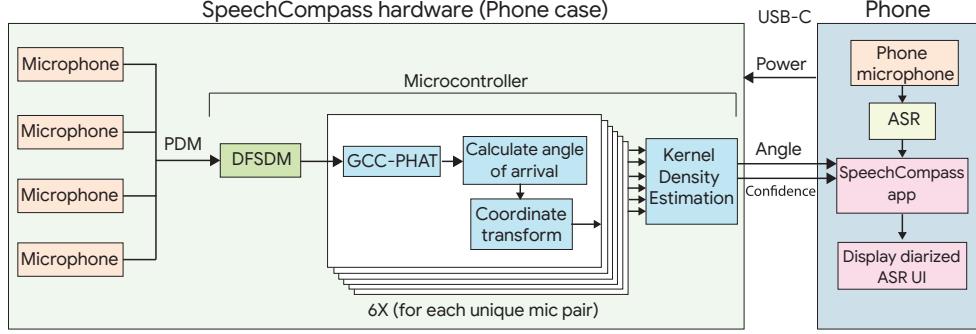


Figure 5: SpeechCompass system diagram. The phone case contains four microphones connected to a microcontroller. The audio localization algorithms run on the microcontroller, and the angle estimation is sent over USB to the phone. The SpeechCompass app combines ASR input and angle estimations to provide diarization and directional guidance for the mobile captioning UI.

three microphones could resolve similar angles, an additional microphone improves localization accuracy. Furthermore, when handheld, the four-microphone design is more robust to occlusions from the hand or fingers. The microphone data was collected at 16-bit resolution and 44.1kHz. The PDM to PCM (pulse-code modulation) conversion was done by the Digital Filter for Sigma-Delta Modulator (DFSDM) peripheral on the microcontroller. The same clock signal drove the four microphones, so they remained synchronized. We used flexible PCB-mounted microphones and a small main PCB, as shown in Figure 2B, C.

4.3 Localization

The localization is computationally intensive, as it requires estimating the delay between all unique pairs of microphones. The delay estimation is usually done using cross-correlation.

We implemented a variant of Generalized Cross-Correlation with Phase Transform (GCC-PHAT) [37]. This approach is more robust to noise than standard cross-correlation and takes advantage of ARM's CMSIS libraries [5], which enable more efficient computations. Then, the time delay between the two microphones is extracted from the cross-correlation. Since the microphone geometry is known, the time delay is converted into the angle of sound arrival. The localization equations are provided in Appendix A.1.

With just two microphones, there is front-back confusion. Different sources positioned at the inverse angle about the microphone pair axis will have exactly the same TDOA. (See Figure 6A).

To accurately determine the 360° azimuth angle, we need at least three microphones, positioned such that they span the largest plane possible (they cannot all be in-line).

Considering the uncertainties of sound propagation, a statistical approach has to be used to determine the actual source location. We determine the TDOA for each microphone pair (six unique pairs for four microphones) and add a second potential TDOA. Then, we perform a coordinate transform so that each angle of arrival is aligned with global azimuth angles. Finally, we do Kernel Density Estimation (KDE) with a Gaussian kernel (bandwidth = 25) with the 600 latest samples. The highest peak will correspond to the angle of arrival, as demonstrated in Figure 6. The KDE was wrapped around

between 0 and 360° to address the discontinuity around 0 and 360°. The localization is evaluated in Section 5.

4.4 SpeechCompass application: Mobile captioning with speaker separation

To investigate the SpeechCompass technique's potential, we implemented a mobile ASR application (Android) with localization features, as shown in Figure 7. We used the USB Serial for Android library [48] to establish a data connection between the phone and the SpeechCompass microcontroller, which transmits localization data over USB-C. The application improves over existing single-source ASR applications in three ways, which we describe below.

Real-time sound source visualization. We use a semicircle overlay around the edges of the screen to indicate the current sound direction. The semicircle moves according to the azimuth angle of the audio, as shown in Figure 7B. (Alternatively, a vertical line could be used if only 180° localization is available.) The radius of the semicircle is scaled according to localization confidence, corresponding to the KDE peak, with a larger semicircle corresponding to higher confidence.

Speaker diarization in transcription. APIs for mobile and web-based ASR have advanced significantly in recent years, but those APIs do not distinguish between multiple speakers or provide direction of the incoming speech. By combining SpeechCompass with mobile ASR, we enable legible transcripts by visually separating speech from different directions. We use the Android Speech Recognition API [3] to obtain real-time transcription and display the results on the mobile phone screen as vertically scrolling captions. The transcript is colored by directly mapping the sound arrival angle to the 360° color wheel, as shown in Figure 7C. In this mapping, the text is colored green if sound arrives from the top of the phone (around 90°) and red for bottom arrival (270°). Other direction indication options include showing a colored arrow next to the text or a directional glyph, as shown in Figure 7A, B. To accurately determine text color, we synchronize speech recognition and azimuth measurements by buffering azimuth angles after the *onBeginningOfSpeech()* callback and stopping after the *onEndOfSpeech()* callback. This way, only the angles detected during speech

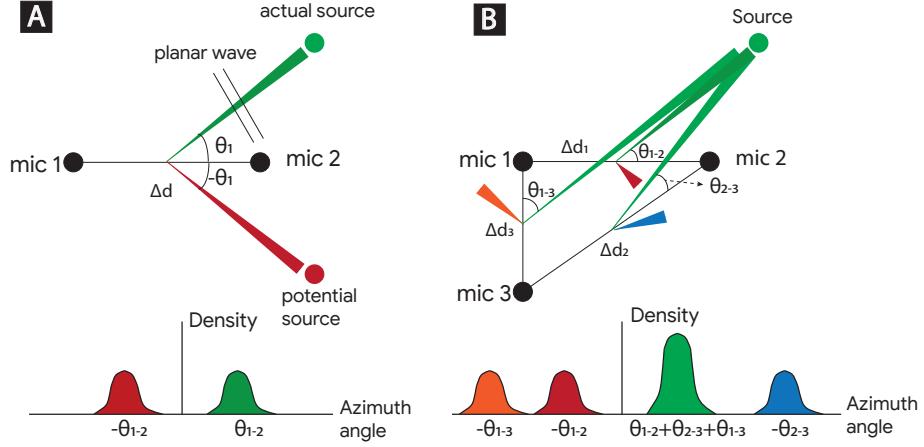


Figure 6: Visualization of localization methods with 2 and 3 microphone configurations. A) Localization with two microphones. The sound will arrive at microphone two before microphone one. This time difference could be used to estimate the angle of arrival. However, with two microphones, ambiguity exists, as the source could be at the inverse angle, shown as a "potential source." The graph on the bottom shows the kernel density estimation (KDE) with actual and potential sources. B) With three or more microphones, this angle ambiguity can be avoided. In our implementation, we use four microphones. The KDE from multiple microphone pairs will have the highest peak at the correct source.

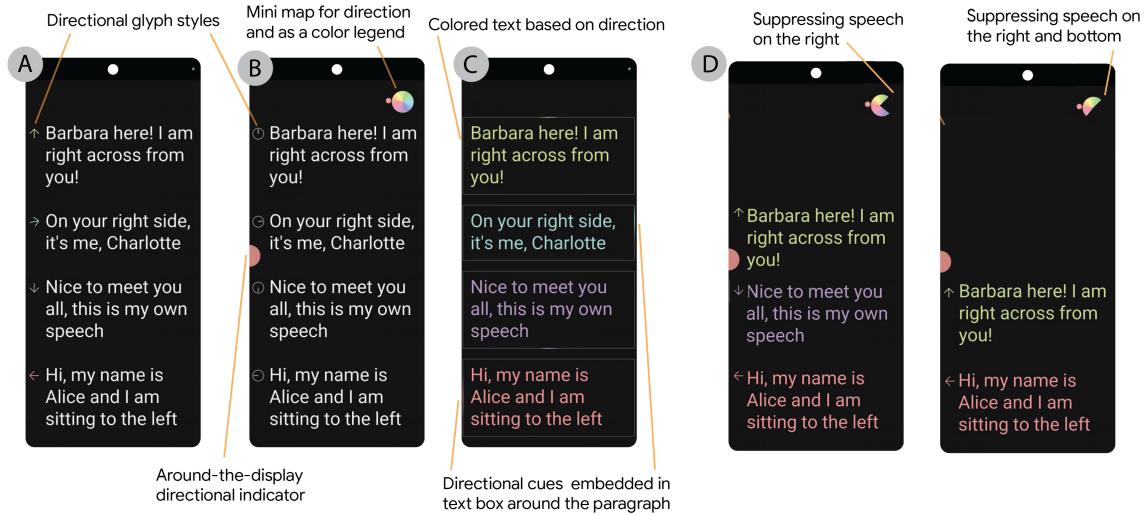


Figure 7: The mobile phone application with different direction visualization options. A) Directional glyphs are arrows next to the transcript, indicating the direction of speech. B) Minimap and directional glyph as radius inside the circle. C) Directional cues embedded in the color of the text and boxes around the text. D) Enabling speech suppression for right and bottom speech directions, as shown on the minimap.

moments are analyzed, ignoring changes during silence or background noise. We determine the resulting angle by computing the mode of the angles in the azimuth buffer.

Speaker suppression. As 60% of the participants found mistakes with background noise to be an issue, we added unwanted speech suppression. With the speaker diarization, we can simulate speaker suppression, by letting the user hide speech from certain

directions, or their own speech. Our implementation allows tapping on different edges of the screen. For example, as shown in Figure 7C, tapping on the right side of the screen will hide/show speech coming from the right.

5 Technical Evaluation

In this section, we report on our technical evaluations of the main features of SpeechCompass. First, we characterize the underlying

localization technology. We evaluate localization using an off-the-shelf mobile phone (Pixel 6, Google) and 360° localization with the SpeechCompass hardware. Then, we determine the compute time for reliable localization, as it is critical for real-time diarization and visualization. Finally, for the real-world applicability of SpeechCompass, we characterize diarization accuracy in a realistic scenario where we also measure power consumption.

5.1 Localization experiments setup

By its nature, localization can not be done precisely due to time resolution. For example, with 80 mm microphone spacing and 44.1 kHz audio sampling rate, the maximum delay is ± 10 samples (0.23 ms). This configuration provides a resolution of 9°. However, the actual accuracy is more nuanced, as it is obtained from multiple microphone pairs with different spacing and depends on many factors, such as the sound's frequency, environment, and loudness. To better characterize localization accuracy, we conducted a series of experiments with the phone and SpeechCompass hardware. In the setup, the device under test was mounted on a rotating platform. A servo motor (XM430-W350-T, Dynamixel) was used to rotate the device around its axis (azimuth angle) and was synchronized with the data collection. A stationary speaker (S120, Logitech) was mounted 1.5 meters away to provide a direct sound path. Using the servo mount, the device was rotated by 10 degrees from 0–350° to obtain azimuth angle error. At each position, data was collected for 15 seconds. The loudness was calibrated with a reference sound meter (VL6708, VLIKE). The speech audio came from the audiobook “Alice’s Adventure in Wonderland” [42]. Rain [56] was used as a realistic environmental sound with white noise properties. Each experimental condition was measured three times, and we report the mean value.

5.2 Localization: Off-the-shelf smartphone

Many smartphones today have two or three microphones, which has the potential to enable some localization. We investigate the feasibility and accuracy of localization with an unmodified, off-the-shelf phone (Pixel 6, Google). This phone has three microphones placed at the bottom, top and near the rear camera, as shown in Figure 8. The top and bottom microphones are on the side edge of the phone, while the rear camera microphone is on the backside.

To characterize the opportunity and limitations, we ported our localization algorithm to Java Native Interface (JNI), so it could run on the Android platform. The two-mic-localization runs on the phone in real time.

With only a 15 mm separation between microphones 1 and 2, the delay between those two microphones was too small for useful localization. However, the other two pairs – {mic 0, mic 2} and {mic 0, mic 1} – provided useful localization data. Given this geometry, only 180-degree azimuth localization is possible for this device. If the camera microphone (mic 1) had been placed on the left side of the phone (with a larger baseline distance from mic 2), 360-degree localization would have been feasible.

The mean error of environmental sound was 23.0 degrees, and speech was 16.0 degrees. In comparison, humans also have up to 20-degree azimuth error when localizing sound without visual cues [45]. This demonstrates that localization on the phone can

still be useful in limited use cases, for example, if the user does not expect any sounds from behind. We also observed that with such a limited amount of data, a simple histogram performed better than the Gaussian distribution approach, as there are not enough statistics for a reliable distribution.

5.3 Localization: SpeechCompass 4-mic system

We varied two variables; type of sound (speech and environmental), and loudness in sound pressure level (SPL). The loudness varied from 50–65 dB SPL in 5 dB increments, where 0 dB corresponds to the human hearing threshold. This range covered the span of loudness typically experienced in daily conditions [29], from a quiet office (50 dB) to a normal conversation (65 dB). The ambient loudness in the experiment’s room was 45 dB, so experimenting at this loudness or below was not practical.

The experimental results show that sound could be localized reliably for both speech and environmental sound, with different loudness, and with 360° coverage. The overall error is reported in Figure 10A. The error was 11.1–22.1° for normal conversational loudness (60–65 dB). In comparison, humans also have up to 20-degree errors when localizing sound [45]. The error increased significantly as the sound got quieter. The GCC-PHAT picks up the loudest sound, thus making it difficult to estimate TDOA if the sound source is under or at the environmental noise level.

Directional effect. The localization error has some dependence on the source angle, with an overall sinusoidal trend, as can be seen in Figure 10B. The lowest error was around 0, 90, 180, and 270°, whereas the highest error was observed at 45 and 135°. This sensitivity can be explained by the rectangular microphone geometry. With the current coordinate system, an angle such as 90° results in 4 out of 6 microphone pairs in the end-fire and/or broadside in relation to the sound source. In such a configuration, the maximum and minimum delays are obtained, thus easier to measure.

Type of sound. The error was lower for the environmental sound (11.3° at 65 dB) when compared to speech (18.3° at 65 dB). There are a few reasons why speech appears to be harder to localize reliably. First, speech has a very complex acoustic pattern in comparison to noise. There are many frequencies, sudden stops, and loudness variations [63]. This results in more reverberations and reflections, which can increase the error. Also, speech contains cyclic signals, which have ambiguity when subjected to cross-correlation, especially for high frequency, and speech contains signals up to 4 kHz.

Effects of elevation angle. In the above experiments, we test the azimuth angle error and keep the elevation angle at zero, so the audio source and microphone array are leveled. In practical use cases, the audio source might be located below or above the microphone array. For example, with the microphone array on the table, the speaker’s head is usually above the array. To characterize the effect of elevation, we move the source elevation from -40 (source below mics) to +40° (source above mics) in 20-degree increments, while keeping the azimuth angle at 90°. The distance is kept at 1.5 meters as in previous experiments and white noise is played at 65 dB. The results in Figure 9 show minor azimuth error (under 2 degrees) at -20 and +20° elevation. The error becomes more pronounced (up to 9.2°) at -40 to +40° elevation, as the source moves outside of the microphone 2-D plane. The 20° elevation would result in a height of

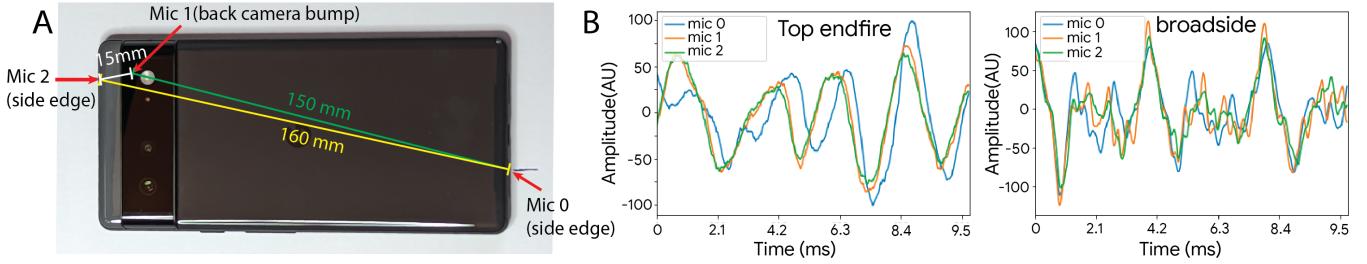


Figure 8: Localization on a mobile phone. A) Microphone positioning and distance. B) Examples of the raw microphone data from three microphones. The top graph shows the end-fire, with the bottom microphone facing the sound source, thus providing maximum delay between the microphones. The delay between mic 0 and mic 2 and 1 can be easily seen. The bottom graph shows broadside, where the microphones are the same distance from the sound, and thus have the same delay.

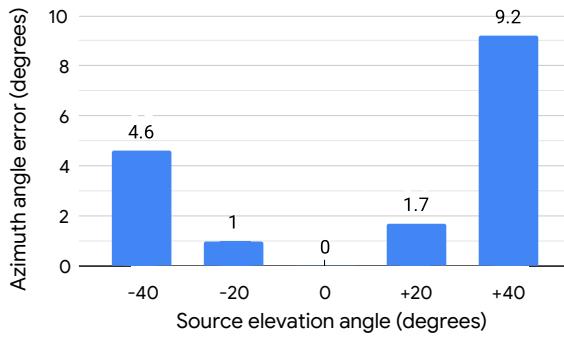


Figure 9: Effect of source elevation angle on the azimuth angle accuracy.

0.55m, which would correspond to speakers sitting around a table with the phone 1.5 meter away. The 40° elevation would result in a height of 1.26m, approximately corresponding to speakers standing around a table. The results therefore suggest that SpeechCompass could work well for seated conversation partners, but would need additional calibration to account for situations where speakers are further above the microphone array.

5.4 Compute time to estimate direction

It takes 2.9 ms to compute the delay between two microphones on the 110 MHz microcontroller. Therefore the time to collect one data frame for six unique microphone pairs is 17.4 ms. This provides a lower limit on time to estimate direction.

However, in practice, the estimation time could be much higher due to the uncertain nature of sounds like speech, as multiple measurements are needed for estimation. Therefore, we experimented to understand how quickly the device can react to the onset of sound in making an accurate direction estimate.

In a similar setup as the localization accuracy experiment, we positioned the device 1.5 meters from the sound source. As in the localization accuracy experiment, we used speech and noise, and measured them at different angles. We collected data for 10 seconds without sound, then turned on the sound and collected the data for

50 seconds at 44.1 kHz. This measurement was conducted at angles from 0–360 in 60° increments.

5.4.1 Results. We evaluated the number of samples it took for the sound source angle to converge around the actual angle by examining the KDE maximum after each additional sample.

With the noise audio, localization required a mean of 12.7 samples (max: 19, min: 7). This number of samples was typically collected within one frame. So the latency was just the computational latency of 17.4 ms.

The speech audio needed on average 5.7X more samples to stabilize, with a mean of 72.7 samples (max: 149, min: 24). The speech contains more sparse localization information, and most measurements (80%) did not have measurable microphone delays. As a result, a mean of 15 frames is required; thus, the mean estimation time is about 263 ms. However, estimation time can range approximately from 70–500 ms, depending on the speech characteristics.

5.5 Power consumption

As the device is designed to be portable and powered by its own power source or the phone's battery, it is essential to understand its power consumption. The power consumption of the whole system was 28 mA. The four microphones consumed 2.4 mA, while the microcontroller consumed 11.7 mA. The audio codec and other peripherals consumed the rest. Therefore, the device could run speaker localization continuously for 18 hours with a 500 mAh battery. When powered by a smartphone, considering a representative smartphone's power consumption of 154 mA (3700 mAh battery for 24 hours), this represents an additional 18 percent of the power consumption.

5.6 Diarization accuracy with localization

In this section, we evaluate diarization with the SpeechCompass hardware. A schematic view of the setup is shown in Figure 12A. We set up four speech sources at 0, 90, 180 and 270 degrees to mimic a turn-taking conversation with four talkers. The speech material is from LibriSpeech [51]. A diffuse noise field is simulated using four loudspeakers simultaneously playing uncorrelated noise content. The noise types are babble noise (pub and cafeteria) as well as traffic noise. One synthesized conversation lasts an average of 40 seconds with each speaker speaking in turn (i.e., no overlap in the speech

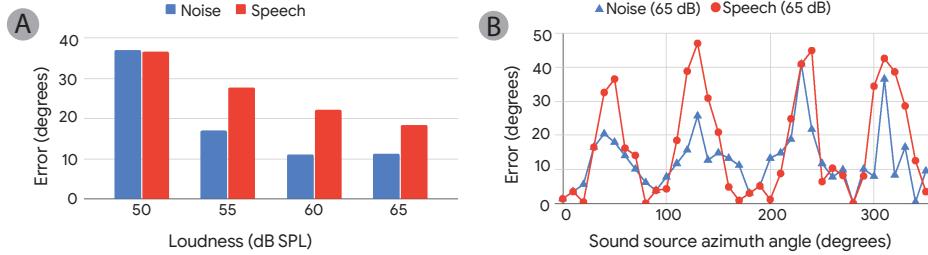


Figure 10: A) SpeechCompass azimuth angle measurement errors for speech and noise sound at different loudness levels. A mean of 0–350° azimuth measurements is reported. B) Azimuth localization error at different angles. Data from one noise and speech trial are shown. The multi-mic device was rotated around its axis from 0–350° while the source was stationary. A 15-second window was used for each estimate.

content). The noise level is set to create three signal-to-noise ratio (SNR) levels of 6, 12, and 18 dB respectively. A fourth scenario is generated in clean condition (i.e., no noise). For each scenario, 100 conversations are created, corresponding to about 1.2 hours of active speech.

An example diarized segment overlaying TDOA values is shown in Figure 12B, indicating the TDOA follows the speaker but still has some noise.

In addition to the SNR sweep, we compared two microphone configurations using respectively three and four microphones, resulting in a total of eight experiments. We computed the diarization error rate (DER) for each scenario using the PyAnnotate toolkit [12]. The DER is computed by summing the durations of false alarms, missed detections, and speaker confusions, and then dividing it by the total ground truth speech duration. The DER for the four-microphone configuration performed consistently better than the three-microphone one, across all four SNR conditions, with a relative DER improvement varying from 23 to 35% for an average of 32% as can be seen in Table 12. A four-microphone configuration allows for more unique microphone pairs (6) than a three-microphone configuration (3). This provides extra TDOA information that can be used by the diarization algorithm to segment and tag the speech content between the four talkers more accurately.

6 User evaluation with frequent users of mobile ASR: Lab study and online survey

To evaluate the usability of our approach, we decided to conduct an in-person lab evaluation of the SpeechCompass phone case and the speech-to-text application (described in Section 4.4), with frequent users of mobile transcription technology. We first conducted a large-scale online pilot study to inform the design of the in-person lab evaluation, which we conducted with eight deaf or hard-of-hearing participants, set up to mimic a realistic conversation scenario.

6.1 Large-scale, online survey (n=494)

In this survey, we use screenshots of our interactive UI prototypes to solicit initial user feedback on the potential for our proposed approach, to guide the design of a more realistic in-person lab study.

The study was conducted using the same Google Surveys deployment and screening methodology as for the foundational study,

detailed in Section 3. The participants were shown different UI renderings and were asked to rate them. The large-scale online survey could only show static images of the interfaces, due to limitations of the survey tool. Out of 985 respondents we focus our analysis on the 494 participants who use captioning technology multiple times per week or more frequently.

As shown in Figure 13A, the colored text was found to be valuable by 60% of participants. Glyph indicators for speech direction, which included arrow and circle+line indicators, were found valuable by 70%. The Edge indicator and the mini map had a less positive reception.

To better understand which glyph indicators were favored, we also asked targeted questions about them, as shown in Figure 13B. *Circle + line* was preferred by 13.1% more respondents than the *highlight box* (45.1% vs 32.0%), and the *arrow* was preferred by 21.9% more respondents than the *circle + line* (51.2% vs 29.3%).

6.2 Lab study (n=8)

We recruited 8 participants from our institution who were frequent users of captioning technology. Five were female, three were male, and all were deaf or hard of hearing. One participant was 25–34 years old, four were 34–44, one was 45–53, and two were 65+ (we are only allowed to collect age ranges at our institution).

The study took place in a quiet lab over approximately 60 minutes and used the phone-case prototype (Figure 2) with our mobile ASR application (Figure 7). First, the participant was introduced to the technology, prototype, and the purpose of the study. Then, the participant was asked to fill out a background survey, which included demographic questions and their current use and experienced challenges with transcription technology. Afterward, the participant was introduced to different visualization scenarios with the SpeechCompass application. The participant used the SpeechCompass transcription while sitting between the two experimenters, as they all sat around a small table with the SpeechCompass phone case in the center. In each of the seven conditions, which ran for 5 minutes, the experimenters sat across from each other and had short conversations about different topics. The participants were instructed to turn off hearing aid devices if they used any, and were asked to use the SoundCompass UI and transcript to follow the conversation. The experimenters' casual conversations included topics like weekend plans, hobbies, and the weather. The seven conditions, which

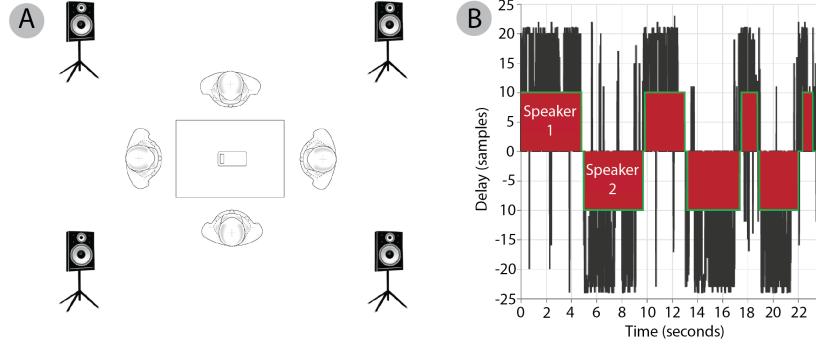


Figure 11: Diarization experiment. A) Data collection setup for diarization testing. The SpeechCompass phone case is placed on a small table in the center of the room. Head and torso simulators are used to play back speech material. Four loudspeakers are placed at the corners of the room to simulate a diffuse noise field. B) Example of TDOA data overlayed with diarization from two speakers for a conversation snippet. The extracted speaker labels from TDOA precisely follow the conversation turns. For TDOA, a frame size of 512 (11.6 ms) with no overlap at 44.1 KHz was used. Then, a running histogram of 522 ms was applied to get speaker labels.

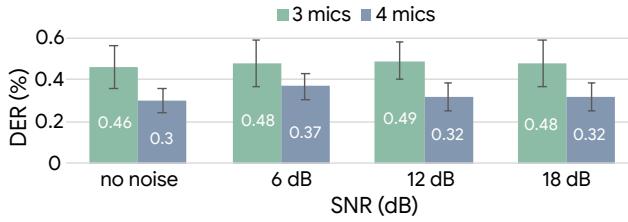


Figure 12: Diarization error rate (DER) for three and four microphone configurations, across a sweep of SNR levels.

used the ASR, diarization, and localization functionality for different visualization techniques, are shown in Figure 14 and presented with more UI context in Figure 7. The conditions were:

- (1) **Transcription only.** The transcribed text is shown in white on a black background.
- (2) **Edge indicator.** A circle (“dot”) that moves around the edge of the screen to point to the currently active speaker. The color of the dot changes based on the direction.
- (3) **Arrow indicator.** A glyph using a colored arrow next to a white text block. The glyph points in the direction of the associated speech.
- (4) **Circle + line indicator.** A glyph using a circle with a directional line next to a white text block. The glyph points in the direction of the speech associated with the text.
- (5) **Mini map.** A colored circle with a smaller circle (“dot”) moves around its edge to point to the currently active speaker. The color of the dot changes based on the direction.
- (6) **Colored text.** The text is colored based on the direction that the associated speech was coming from.
- (7) **Everything on.** All indicators are turned on (except the Circle + line, as it couldn’t be used simultaneously with the arrow).

After participants had completed all conditions, they filled out a form that asked them to rate how desirable each of the five visual

indicator styles (*Edge indicator*, *Arrow*, *Circle + line*, *Colored map*, and *Colored text*) were on a 7-point Likert scale, from -3: *Strongly dislike* to +3: *Strongly like*. Finally, they were asked to rate the overall value of directional feedback to the transcription experience, how strongly they would recommend these features to users of mobile captioning, and whether they had any general free-form feedback about SpeechCompass.

6.3 Results

Mobile transcription apps (e.g., Android Live Transcribe) were the most used communication technology for the participants. Specifically, three used them multiple times per day, one used them daily, three used them multiple times per week, and one used them rarely.

75% of participants frequently experienced the scenario where multiple people would get mixed up in the transcript (two multiple times per day, two daily, two multiple times per week). All participants agreed that it was challenging to participate in conversations when speech was combined from multiple people.

A Kruskal-Wallis (KW) test found a significant effect on participant preferences for visualization techniques ($P=.014$). The post-hoc pair-wise analyses using the Wilcoxon test with Bonferroni correction did, however, not show statistical significance between any pairs. Of the five visualization techniques, *Colored text* was the most well-received (mean (\bar{x}) = 2.625), as it was rated positively by all the participants. The *Arrow* indicator was also well-received (\bar{x} = 1.125), with six positive, one negative, and one neutral participant. Several participants noted that *Arrow* and *Colored text* worked well together: “*Arrows + color seem to be most easier way to indicate the direction.*” (P2) and “*The combination of the colored text with the arrow was the most effective for me.*” (P7).

The other indicator styles received more mixed feedback. The feedback for both *Edge indicator* (\bar{x} = 0.25) and *Circle + line* (\bar{x} = -0.125) was split between four negative and four positive participants. Some participants were concerned that *Edge indicator* was distracting and not sufficiently discreet: “*I do prefer the tool be as*

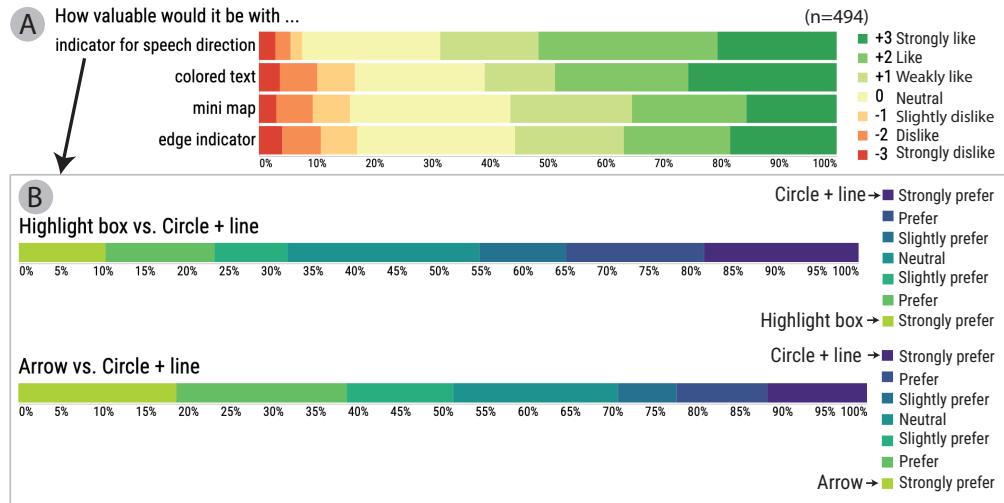


Figure 13: Participants' preferences for different visualization techniques in the online survey. A) Results indicating how valuable the specific indicator would be for the user. B) Preferences for the specific indicators for speech direction.



Figure 14: Examples of seven visualization scenarios that participants experienced in the in-person study.

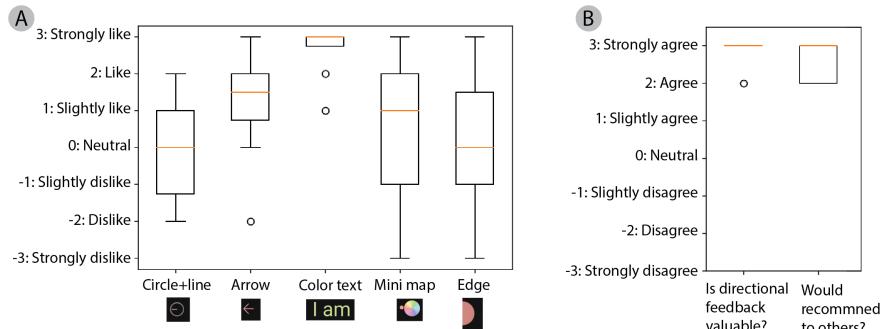


Figure 15: Boxplots of results of the in-person study. A) Participants' preferences for different visualization techniques. B) Overall opinions about augmented mobile ASR application.

discrete as possible and would perhaps choose to avoid bright colored things moving around since this would be eye-catching and this kind of attention is often undesired" (P3) and "Indicator moving around the edge was distracting and causing a bit of eye strain" (P2). On the other hand, another participant found this style particularly useful: "the color dot moving to the speaker direction worked REALLY well"

(P1). For *Circle + line*, some participants struggled with its legibility: "If the analog direction indicators were larger (and translucent, or set behind)" (P8) and "The lines in a circle were a bit slower and not as accurate (buggy)" (P5). The *Mini map* was rated positively by five participants and negatively by three. The most favorable participant stated: "this is also great for environmental awareness for those with

single-sided hearing or no hearing at all." (P3) and a participant who disliked the *Edge indicator* commented: "*steady map in the corner worked a bit better* (P5)".

Overall, all participants agreed with the value of directional feedback ($\bar{x} = 2.88$, seven Strongly agree:+3 and one Agree:+2) and would recommend these features to other users of captioning technology ($\bar{x} = 2.63$, five Strongly agree:+3 and three Agree:+2): "*I really liked that almost immediately I could tell that there was a speaker change, so that as soon as the text started to show up, I could better contextualize that text as attributed to a new speaker.*" (P1), "*I'm very happy to see this tool being developed, it's a great addition to other speech recognition tools!*" (P3), and "*This prototype is definitely a life changer and I strongly believe that it will improve the quality of access to communication with speakers for many users*" (P6).

6.4 Discussion

Consistent with the large-scale survey, the value of the diarization and localization features was immediate to all users. The participants were asked if directional guidance would be valuable in their mobile transcription experience. All eight users agreed. Also, all eight users would recommend this feature to mobile captioning users.

While the large-scale survey helped inform our testing and exclude conditions (e.g., *Highlight box*), the lab study allowed us to more rigorously evaluate the techniques in a realistic scenario. This difference became significant for the *Edge indicator* and *Mini map*, where issues, such as discreetness and distracting aspects, became evident during live usage.

The results suggest that the combination of *Colored text* and *Arrow* would meet the preferences of most users, thanks to the balance of directional encoding and clarity. The arrow has redundant benefits too, since colored text might not always be reliably visible depending on lighting and screen conditions (e.g., strong sunlight, or dim display) and might also not be usable for colorblind users. The mixed feedback for other techniques indicates that the interface may also benefit from mechanisms that would allow users to customize the visualization style. Such customization could also apply to rendering properties, such as color, transparency, and line thickness, as some participants found *Circle + line* particularly difficult to interpret. In both the large-scale survey and the in-person lab study, the *Arrow* was preferred over *Circle + line*. Through more customization options and extended usage in their daily lives, participants will be able to provide more nuanced feedback about these techniques.

7 Limitations and future work

In this work, the microphone array topology was designed for experimentation and integration with mobile phones. In the future, we plan to explore other form factors such as smart glasses and smartwatches, where multiple microphones could also be used. To do so, we plan to create a microphone array platform for rapid prototyping where microphone topology can be easily reconfigured and customized with a user-friendly UI and software library.

Localization needs to be accurate in everyday scenarios that may suffer from external noise and interference sources. We plan to add machine learning approaches to improve the noise robustness of

our current localization, which is based on classical linear signal processing.

In our user studies, errors due to background noise were identified as a major issue with mobile ASR. In future work, background noise could be suppressed by using a steerable beamformer, which would benefit from more microphones. In our experimentation with classical beamforming techniques like filter-and-sum beamformer [7] and Minimum Variance Distortionless Response (MVDR) [64], we achieved a few dB SNR (signal-to-noise ratio) improvement with our microphone geometry. With the recent advances in neural beamformers [41, 65], which are trained beamformers, higher SNRs could be possible.

We are interested in scaling up the user evaluation, both with a larger set of participants and over longer time periods and in ecologically valid settings. We therefore plan to conduct longitudinal user evaluations with frequent users of mobile ASR to further advance our understanding of the usability of this approach in the wild.

8 Conclusion

Motivated by a foundational large-scale survey with 263 users frequent users of captioning technology, this work demonstrates how microphone array signal processing can supplement ASR by measuring and visualizing the spatial dimension of audio. Although ASR technology has greatly improved, diarization and localization features are not commonly available. To investigate the potential of enabling mobile captioning for group conversations, we implemented a low-latency 360° localization algorithm that can run on general-purpose low-power microcontrollers, and a custom sound perception hardware solution with four microphones.

Our technical evaluation of localization and diarization with the SpeechCompass microphone array demonstrates benefits over the traditional single-microphone configuration and pure machine-learning-based approaches.

We introduced a mobile captioning app that uses sound localization to diarize and visualize speech directions for group conversations. The integration with our embedded software and hardware brought new capabilities to mobile ASR, including sound source location, speaker diarization, and user control of the diarization, allowing the suppression of specific speech directions. Using the developed mobile phone application, we conducted an in-person study with frequent users of mobile ASR technology to gather feedback on the techniques and different visualization styles, and their potential for improving the captioning of group conversations. All the participants found the diarization, localization, and visualization features to be useful, and particularly appreciated the combination of a directional arrow and colored text.

This work demonstrates that low-power microphone array processing can be integrated into new and existing mobile devices, thereby leveraging audio's natural spatial properties to enhance audio experiences and the understanding of speech.

In the future, we hope that our approach will inspire the widespread adoption of advanced microphone arrays that natively unlock the potential of spatial sound processing and perception in mobile and wearable devices.

Acknowledgments

We thank Sagar Salva and Dmitrii Votintcev for their ideas on prototypes and interaction designs. We are grateful to Pascal Getreuer, Richard Lyon, Alex Huang, Shao-Fu Shih, and Chet Gney for their help with algorithms. We also thank Shaun Kane, James Landay, Malcolm Slaney, and Meredith Morris for their feedback on this paper. We appreciate the contributions of Carson Lau for the phone case mechanical design and Ngan Nguyen for electronics assembly. Finally, we thank Mei Lu, Don Barnett, Ryan Geraghty, and Sanjay Batra for UX research and design.

References

- [1] Chandykunju Alex, Kevin Martin Jose, and Arun Joseph. 2013. Sound localization and visualization device. In *2013 IEEE Global Humanitarian Technology Conference (GHTC)*. 262–264. doi:10.1109/GHTC.2013.6713692
- [2] Android. 2022. Introducing Live Transcribe. <https://www.android.com/accessibility/live-transcribe/>. Accessed 2022-03-26.
- [3] Android. 2022. SpeechRecognizer API Documentation. <https://developer.android.com/reference/android/speech/SpeechRecognizer>. Accessed 2022-10-25.
- [4] Xavier Anguera, Chuck Wooters, and Javier Hernando. 2007. Acoustic beamforming for speaker diarization of meetings. *IEEE Transactions on Audio, Speech, and Language Processing* 15, 7 (2007), 2011–2022.
- [5] ARM. 2022. CMSIS DSP Software Library. <https://www.keil.com/pack/doc/CMSIS/DSP/html/index.html>. Accessed 2022-05-12.
- [6] Ava. 2022. Ava Captioning Solution. <https://www.ava.me/>. Accessed 2024-12-10.
- [7] Jacob Benesty, Jingdong Chen, and Yiteng Huang. 2008. Conventional beamforming techniques. *Microphone array signal processing* (2008), 39–65.
- [8] Larwan Berke, Khaled Albusays, Matthew Seita, and Matt Huenerfauth. 2019. Preferred Appearance of Captions Generated by Automatic Speech Recognition for Deaf and Hard-of-Hearing Viewers. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI EA '19). Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3290607.3312921
- [9] Larwan Berke, Christopher Caulfield, and Matt Huenerfauth. 2017. Deaf and Hard-of-Hearing Perspectives on Imperfect Automatic Speech Recognition for Captioning One-on-One Meetings. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility* (Baltimore, Maryland, USA) (ASSETS '17). Association for Computing Machinery, New York, NY, USA, 155–164. doi:10.1145/3132525.3132541
- [10] Rachel Boll, Shruti Mahajan, Jeanne Reis, and Erin T. Solovey. 2020. Creating Questionnaires That Align with ASL Linguistic Principles and Cultural Practices within the Deaf Community. , Article 61 (2020), 4 pages. doi:10.1145/3373625.3418071
- [11] Danielle Bragg, Nicholas Huynh, and Richard E. Ladner. 2016. A Personalizable Mobile Sound Detector App Design for Deaf and Hard-of-Hearing Users. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility* (Reno, Nevada, USA) (ASSETS '16). Association for Computing Machinery, New York, NY, USA, 3–13. doi:10.1145/2982142.2982171
- [12] Hervé Bredin, Ruiqing Yin, Juan Manuel Coria, Gregory Gelly, Pavel Korshunov, Marvin Lavechin, Diego Fustes, Hadrien Titeux, Wassim Bouaziz, and Marie-Philippe Gill. 2020. pyannote.audio: neural building blocks for speaker diarization. In *ICASSP 2020, IEEE International Conference on Acoustics, Speech, and Signal Processing*.
- [13] Janine Butler, Brian Trager, and Byron Behm. 2019. Exploration of Automatic Speech Recognition for Deaf and Hard of Hearing Students in Higher Education Classes. In *Proceedings of the 21st International ACM SIGACCESS Conference on Computers and Accessibility* (Pittsburgh, PA, USA) (ASSETS '19). Association for Computing Machinery, New York, NY, USA, 32–42. doi:10.1145/3308561.3353772
- [14] Chao Cai, Henglin Pu, Peng Wang, Zhe Chen, and Jun Luo. 2021. We Hear Your PACE: Passive Acoustic Localization of Multiple Walking Persons. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 2, Article 55 (jun 2021), 24 pages. doi:10.1145/3463510
- [15] Caluá de Lacerda Pataca, Saad Hassan, Nathan Tinker, Roshan Lalitha Peiris, and Matt Huenerfauth. 2024. Caption Royale: Exploring the Design Space of Affective Captions from the Perspective of Deaf and Hard-of-Hearing Individuals. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 899, 17 pages. doi:10.1145/3613904.3642258
- [16] Artem Dementyev, Pascal Getreuer, Dimitri Kanevsky, Malcolm Slaney, and Richard F Lyon. 2021. VHP: Vibrotactile Haptics Platform for On-Body Applications. In *The 34th Annual ACM Symposium on User Interface Software and Technology (Virtual Event, USA) (UIST '21)*. Association for Computing Machinery, New York, NY, USA, 598–612. doi:10.1145/3472749.3474772
- [17] Hoang Do, Harvey F Silverman, and Ying Yu. 2007. A real-time SRP-PHAT source location implementation using stochastic region contraction (SRC) on a large-aperture microphone array. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*, Vol. 1. IEEE, 1–121.
- [18] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T Freeman, and Michael Rubinstein. 2018. Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation. *arXiv preprint arXiv:1804.03619* (2018).
- [19] Hakan Erdogan, John R Hershey, Shinji Watanabe, and Jonathan Le Roux. 2015. Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 708–712.
- [20] Leah Findlater, Bonnie Chinh, Dhruv Jain, Jon Froehlich, Raja Kushalnagar, and Angela Carey Lin. 2019. Deaf and Hard-of-Hearing Individuals' Preferences for Wearable and Mobile Sound Awareness Technologies. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300276
- [21] Israel D Gebru, Siley Ba, Xiaofei Li, and Radu Horaud. 2017. Audio-visual speaker diarization based on spatiotemporal bayesian fusion. *IEEE transactions on pattern analysis and machine intelligence* 40, 5 (2017), 1086–1099.
- [22] Abraham Glasser, Kesavan Kushalnagar, and Raja Kushalnagar. 2017. Deaf, Hard of Hearing, and Hearing Perspectives on Using Automatic Speech Recognition in Conversation. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility* (Baltimore, Maryland, USA) (ASSETS '17). Association for Computing Machinery, New York, NY, USA, 427–432. doi:10.1145/3132525.3134781
- [23] Steven M. Goodman, Ping Liu, Dhruv Jain, Emma J. McDonnell, Jon E. Froehlich, and Leah Findlater. 2021. Toward User-Driven Sound Recognizer Personalization with People Who Are d/Deaf or Hard of Hearing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 2, Article 63 (jun 2021), 23 pages. doi:10.1145/3463501
- [24] Google. 2018. Google Surveys Methodology. http://services.google.com/fh/files/misc/white_paper_how_google_surveys_works.pdf. Accessed 2022-03-22.
- [25] Beth G Greene, David B Pisoni, and Thomas D Carrell. 1984. Recognition of speech spectrograms. *The Journal of the Acoustical Society of America* 76, 1 (1984), 32–43.
- [26] François Grondin and François Michaud. 2019. Lightweight and optimized sound source localization and tracking methods for open and closed microphone array configurations. *Robotics and Autonomous Systems* 113 (2019), 63–80.
- [27] Ru Guo, Yiru Yang, Johnson Kuang, Xue Bin, Dhruv Jain, Steven Goodman, Leah Findlater, and Jon Froehlich. 2020. HoloSound: Combining Speech and Sound Identification for Deaf or Hard of Hearing Users on a Head-Mounted Display. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility (Virtual Event, Greece) (ASSETS '20)*. Association for Computing Machinery, New York, NY, USA, Article 71, 4 pages. doi:10.1145/3373625.3418031
- [28] Jahn Heymann, Lukas Drude, and Reinhold Haeb-Umbach. 2016. Neural network based spectral mask estimation for acoustic beamforming. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 196–200.
- [29] Pulsar Instruments. 2022. Decibel chart – decibel levels of common sounds. <https://pulsarinstruments.com/news/decibel-chart-noise-level>. Accessed 2022-07-27.
- [30] Yusuf Isik, Jonathan Le Roux, Zhuo Chen, Shinji Watanabe, and John R Hershey. 2016. Single-channel multi-speaker separation using deep clustering. *arXiv preprint arXiv:1607.02173* (2016).
- [31] Dhruv Jain, Leah Findlater, Jamie Gilkeson, Benjamin Holland, Ramani Duraiswami, Dmitry Zotkin, Christian Vogler, and Jon E. Froehlich. 2015. Head-Mounted Display Visualizations to Support Sound Awareness for the Deaf and Hard of Hearing. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 241–250. doi:10.1145/2702123.2702393
- [32] Dhruv Jain, Leah Findlater, Jamie Gilkeson, Benjamin Holland, Ramani Duraiswami, Dmitry Zotkin, Christian Vogler, and Jon E. Froehlich. 2015. Head-Mounted Display Visualizations to Support Sound Awareness for the Deaf and Hard of Hearing. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 241–250. doi:10.1145/2702123.2702393
- [33] Dhruv Jain, Kelly Mack, Akli Amrous, Matt Wright, Steven Goodman, Leah Findlater, and Jon E. Froehlich. 2020. HomeSound: An Iterative Field Deployment of an In-Home Sound Awareness System for Deaf or Hard of Hearing Users. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3313831.3376758
- [34] Vahid Ahmadi Kalkhorani, Anurag Kumar, Ke Tan, Buye Xu, and DeLiang Wang. 2024. Audiovisual Speaker Separation with Full- and Sub-Band Modeling on the

Time-Frequency Domain. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 12001–12005. doi:10.1109/ICASSP48485.2024.10446297

[35] Yoshihiro Kaneko, Inho Chung, and Kenji Suzuki. 2013. Light-emitting device for supporting auditory awareness of hearing-impaired people during group conversations. In *2013 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE, 3567–3572.

[36] Ellington Kirby, Seoyoon Park, Yan Wang, and Yingying Chen. 2016. HearHere: Smartphone Based Audio Localization Using Time Difference of Arrival: Demo. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking* (New York City, New York) (*MobiCom '16*). Association for Computing Machinery, New York, NY, USA, 509–510. doi:10.1145/2973750.2985625

[37] Charles Knapp and Glifford Carter. 1976. The generalized correlation method for estimation of time delay. *IEEE transactions on acoustics, speech, and signal processing* 24, 4 (1976), 320–327.

[38] Raja S. Kushalnagar, Gary W. Behm, Aaron W. Kelstone, and Shareef Ali. 2015. Tracked Speech-To-Text Display: Enhancing Accessibility and Readability of Real-Time Speech-To-Text. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility* (Lisbon, Portugal) (*ASSETS '15*). Association for Computing Machinery, New York, NY, USA, 223–230. doi:10.1145/2700648.2809843

[39] Ahmet Köse, Aleksei Tepljakov, and Sergei Astapov. 2017. Real-time localization and visualization of a sound source for virtual reality applications. In *2017 25th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*. 1–6. doi:10.23919/SOFTCOM.2017.8115577

[40] Bowon Lee, Amir Said, Tom Kalker, and Ronald W Schafer. 2008. Maximum likelihood time delay estimation with phase domain analysis in the generalized cross correlation framework. In *2008 Hands-Free Speech Communication and Microphone Arrays*. IEEE, 89–92.

[41] Bo Li, Tara N Sainath, Ron J Weiss, Kevin W Wilson, and Michiel Bacchiani. 2016. Neural network adaptive beamforming for robust multichannel speech recognition. (2016).

[42] LibriVox. 2022. Alice's Adventures in Wonderland by Lewis Carroll (Version 2). <https://librivox.org/alices-adventures-in-wonderland-by-lewis-carroll-4/>. Accessed 2022-07-12.

[43] Hong Liu and Miao Shen. 2010. Continuous sound source localization based on microphone array for mobile robots. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 4332–4339.

[44] Richard F Lyon. 2017. *Human and machine hearing: extracting meaning from sound*. Cambridge University Press.

[45] James C Makous and John C Middlebrooks. 1990. Two-dimensional sound localization by human listeners. *The journal of the Acoustical Society of America* 87, 5 (1990), 2188–2200.

[46] Microsoft. 2022. Translator. <https://translator.microsoft.com/>. Accessed 2022-03-26.

[47] MiniDSP. 2022. USB Mic array. <https://www.minidsp.com/products/usb-audio-interface/uma-8-16-usb-mic-array>. Accessed 2022-05-15.

[48] Kai Morich. 2022. usb-serial-for-android. <https://github.com/mik3y/usb-serial-for-android>. Accessed 2022-11-09.

[49] Pius Kavuma Basajabaka Mugagga and Simon Winberg. 2015. Sound source localisation on Android smartphones: A first step to using smartphones as auditory sensors for training A.I systems with Big Data. In *AFRICON 2015*. 1–5. doi:10.1109/AFRICON.2015.7331970

[50] Alex Olwal, Kevin Balke, Dmitrii Votintcev, Thad Starner, Paula Conn, Bonnie Chinh, and Benoit Corda. 2020. *Wearable Subtitles: Augmenting Spoken Communication with Lightweight Eyewear for All-Day Captioning*. Association for Computing Machinery, New York, NY, USA, 1108–1120. <https://doi.org/10.1145/3379337.3415817>

[51] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. LibriSpeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 5206–5210.

[52] Tae Jin Park, Naoyuki Kanda, Dimitrios Dimitriadis, Kyu J Han, Shinji Watanabe, and Shrikanth Narayanan. 2022. A review of speaker diarization: Recent advances with deep learning. *Computer Speech & Language* 72 (2022), 101317.

[53] Nissanka B. Priyantha, Anir Chakraborty, and Hari Balakrishnan. 2000. The Cricket Location-Support System. In *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking* (Boston, Massachusetts, USA) (*MobiCom '00*). Association for Computing Machinery, New York, NY, USA, 32–43. doi:10.1145/345910.345917

[54] Ashutosh Saxena and Andrew Y Ng. 2009. Learning sound location from a single microphone. In *2009 IEEE International Conference on Robotics and Automation*. IEEE, 1737–1742.

[55] Matthew Seita. 2020. Designing Automatic Speech Recognition Technologies to Improve Accessibility for Deaf and Hard-of-Hearing People in Small Group Meetings. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI EA '20*). Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3334480.3375039

[56] Deep Sleep. 2022. Rain Sound. <https://www.youtube.com/watch?v=13EL6Mgeocc&t=3448s>. Accessed 2022-07-12.

[57] David Snyder, Daniel Garcia-Romero, Gregory Sell, Alan McCree, Daniel Povey, and Sanjeev Khudanpur. 2019. Speaker recognition for multi-speaker conversations using x-vectors. In *ICASSP 2019-2019 IEEE International conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 5796–5800.

[58] Speaksee. 2022. Speaksee microphone kit. <https://speak-see.com/>. Accessed 2022-10-25.

[59] Hassan Taherian, Ashutosh Pandey, Daniel Wong, Buye Xu, and DeLiang Wang. 2024. Leveraging Sound Localization to Improve Continuous Speaker Separation. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 621–625. doi:10.1109/ICASSP48485.2024.10446934

[60] Hassan Taherian and DeLiang Wang. 2024. Multi-Channel Conversational Speaker Separation via Neural Diarization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 32 (2024), 2467–2476. doi:10.1109/TASLP.2024.3393726

[61] Giuseppe Valenzise, Luigi Gerosa, Marco Tagliasacchi, Fabio Antonacci, and Augusto Sarti. 2007. Scream and gunshot detection and localization for audio-surveillance systems. In *2007 IEEE Conference on Advanced Video and Signal Based Surveillance*. IEEE, 21–26.

[62] Quan Wang, Carlton Downey, Li Wan, Philip Andrew Mansfield, and Ignacio Lopez Moreno. 2018. Speaker diarization with LSTM. In *2018 IEEE International conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 5239–5243.

[63] Lucinda Wilder. 1975. Articulatory and acoustic characteristics of speech sounds. In *Understanding Language*. Elsevier, 31–76.

[64] Yuteng Xiao, Jihang Yin, Honggang Qi, Hongsheng Yin, and Gang Hua. 2017. MVDR algorithm based on estimated diagonal loading for beamforming. *Mathematical Problems in Engineering* 2017 (2017).

[65] Yang Yang, George Sung, Shao-Fu Shih, Hakan Erdogan, Chehong Lee, and Matthias Grundmann. 2024. Binaural Angular Separation Network. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1201–1205.

[66] Jiahui Yu, Chung-Cheng Chiu, Bo Li, Shuo-yiin Chang, Tara N Sainath, Yanzhang He, Arun Narayanan, Wei Han, Anmol Gulati, Yonghui Wu, et al. 2021. Fastemit: Low-latency streaming asr with sequence-level emission regularization. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 6004–6008.

A Appendix

A.1 Localization algorithm

The following GCC-PHAT formulation was used:

$$G(f) = \mathcal{F}^{-1} \left(\frac{X_1(f)[X_2(f)]^*}{(|X_1(f)[X_2(f)]^*|)^{-0.3}} \right), \quad (1)$$

where $X_1(f)$ and $X_2(f)$ are the Fourier transforms of the two microphone signals, $[]^*$ denotes complex conjugate, and \mathcal{F}^{-1} is the inverse Fourier transform. $G(f)$ is the resulting cross-correlation. We use partial normalization to the power of -0.3 since it provides more robustness to noise by giving less weight to delays around end-fires, which are more likely due to noise. Partial normalization deviates from the original GCC-PHAT as it uses full normalization to weigh all delays equally.

The time delay between the two microphones is extracted from the cross-correlation in the following way:

$$\Delta t = \frac{\text{argmax}(G(f))}{f_s}, \quad (2)$$

Where f_s is the audio sampling frequency (44.1 kHz), and argmax is the index of the maximum peak in the cross-correlation, corresponding to delay in samples.

To convert the time delay into azimuth angle (in the microphone plane), the microphone spacing needs to be known to calculate the maximum delay:

$$\Delta t_{\max} = \frac{\Delta d}{c}, \quad (3)$$

where $c = 343\text{m/sec}$ is the speed of sound, and Δd is the distance between the microphones.

Assuming far-field sound waves, we can use a simple formula to calculate the azimuth angle. The far-field approximation assumes planar sound waves and is valid approximately if the microphone

is a meter or more away from the source.

$$\theta_{azimuth} = \cos^{-1}\left(\frac{\Delta t}{\Delta t_{max}}\right) \quad (4)$$