Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 2791–2801
https://doi.org/10.54941/ahfe1008503

OpenIN X: A Dual-Path Wearable Auditory Interface for Divided-Attention Support in Unilateral Conductive Hearing Loss

Yutong Cai[1] and Lei Xia[2]

[1]Jiangnan University, Wuxi, Jiangsu, China
[2]Tongji University, Shanghai, China

ABSTRACT

People with unilateral conductive hearing loss frequently perform two listening tasks at once: extracting dense factual content from a talker in front of them, and remaining responsive to events outside that conversation. Devices that optimise the first task—amplification, remote microphones, closed in-ear receivers—can make the second harder, because the ear carrying the improved speech signal is also the ear that would otherwise monitor the room. This paper presents OpenIN X, a dual-path wearable auditory interface that treats the situation as an information-allocation problem rather than an amplification problem. Semantically dense target speech is routed to a closed in-ear module on the intact side; brief environmental events are routed, after one fixed category-blind processing chain, to an open bone-conduction module worn on the affected side, where the ear canal may be unavailable to an in-ear fit. We report the design rationale, the integrated wearable design, and a within-participant feasibility protocol in which normal-hearing adults under acute unilateral attenuation complete twelve continuous service-counter scenes across four output conditions. Speech comprehension and environmental-event detection are reported as separate outcomes and never combined into a composite score. A worked analysis of simulated data shows that, at pilot scale, per-condition rates move in coarse steps, so the pilot can establish feasibility but not effect sizes.

Keywords: Wearable auditory interface, Bone conduction, Unilateral conductive hearing loss, Divided attention, Accessibility, Information allocation, Auditory ergonomics

INTRODUCTION

A person collecting a parcel at a service counter must hold several things at once: the time after which collection becomes possible, the document to present, and the counter to go to. While the clerk delivers those details, a chime may sound at the entrance or someone may call out from behind. Following the conversation and remaining responsive to the room are not the same activity: the first requires sustained processing of semantically dense material, the second requires noticing that something has changed and deciding whether it warrants a turn of the head.

For a listener with unilateral conductive hearing loss these demands converge on a single usable channel. Congenital aural atresia is representative, since the affected canal may be absent or too narrow for a conventional

Received September 26, 2026; Revised October 26, 2026; Accepted October 31, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

in-ear receiver. Siegbahn et al. (2021) found that adults with unilateral congenital ear canal atresia showed impaired horizontal localisation and poorer recognition of speech in competing speech, so the difficulty extends beyond reduced audibility on one side.

Existing support pathways address parts of the problem and often trade one part against another. Delivering a selected talker to the intact ear improves speech-in-noise recognition (Sakamoto et al., 2020), and bone conduction provides benefit without depending on the affected canal (Volgger et al., 2022). Rerouting, however, is a trade rather than a restoration: for unilateral severe-to-profound sensorineural loss, Kitterick et al. (2016) found a modest speech-in-noise advantage when speech was rerouted alongside a comparable disadvantage when noise was, with no established localisation benefit, and in unilateral mixed or conductive loss bone-conduction implantation improved localisation in only three of five recipients (Asp and Reinfeldt, 2018).

A second cost sits on the other side of the same decision. The intact ear is both the best route for a high signal-to-noise speech stream and the listener's only reliable route to the room, and occluding it degrades precisely that monitoring: hearing protection reduced auditory situational awareness while increasing listening effort (Smalt et al., 2020), and earmuffs and double protection impaired the localisation of reverse alarms while head protection alone had little effect (Laroche et al., 2022). A device that optimises the conversation on the intact ear can therefore remove, without announcing it, the channel through which the room was being monitored.

We therefore treat this as an allocation problem rather than an amplification problem. Listening in adverse conditions draws on a limited pool of cognitive resources (Pichora-Fuller et al., 2016), and dual-task paradigms are the established means of exposing that competition (Gagné, Besser and Lemke, 2017). OpenIN X applies two decisions in sequence. Anatomical suitability determines placement: a closed in-ear module on the ear with usable air conduction, an open bone-conduction module on the side whose canal is unavailable. Information density determines content: dense target speech goes to the in-ear output, and brief environmental events, after one fixed category-blind processing chain, go to the bone-conduction output.

Three claims are deliberately withheld. Affected-side placement describes a wearing configuration, not an acoustically isolated pathway, since bone-conducted vibration propagates through the skull; the design does not restore binaural hearing or localisation; and the processing chain does not decide whether an event is dangerous—it makes events more likely to be noticed, not understood. Within those limits the paper contributes an auditory information-allocation strategy for divided attention in unilateral conductive hearing loss, an integrated wearable design that embodies it, and a within-participant feasibility protocol that measures speech comprehension and environmental-event detection as separate outcomes.

BACKGROUND AND RELATED WORK

Unilateral conductive hearing loss is usually described as restricted transmission through the affected outer or middle ear, but in complex environments the listener must also separate target speech from competing sound and stay responsive to events on the impaired side. Asp, Jakobsson

and Berninger (2018) observed acute reductions in spatial listening and speech recognition after inducing unilateral attenuation in normal-hearing adults, which establishes acute occlusion as a controlled, if incomplete, method: it cannot reproduce congenital anatomy or long-term adaptation, so its evidence speaks to the interface rather than to the clinical population.

Environmental sounds share no single defining frequency region: their identification depends on spectral distribution, temporal envelope, periodicity, and onset structure, and many remain identifiable after substantial spectral restriction (Gygi, Kidd and Watson, 2004). Salience is likewise multidimensional and context dependent (Huang and Elhilali, 2017), and bone-conducted speech loses intelligibility as higher frequencies attenuate (Toya et al., 2023). Re-encoding events as learned symbolic tones would discard what makes a sound identifiable, and leaving the stream unprocessed would ignore the transducer's limits, so OpenIN X applies one fixed chain to every event: a bell should still sound like a bell.

Bone conduction has been used in wearable and mixed-reality systems precisely because it adds an auditory layer while leaving the ear canal open (Lindeman, Noma and de Barros, 2007), and deaf and hard-of-hearing participants have expressed clear demand for wearable sound-awareness technology (Findlater et al., 2019). What has received less attention is the assignment of high-density target speech and processed environmental audio to two separate wearable outputs within one concurrent listening task.

THE OpenIN X DUAL-PATH DESIGN

OpenIN X treats anatomical asymmetry as a design constraint rather than a deficit to be corrected (Figure 1). The intact side carries a closed in-ear module delivering target speech at an improved signal-to-noise ratio; the affected side carries an open bone-conduction module that leaves the canal unoccupied and presents environmental events after fixed processing. The competing background is not routed to either module, since it defines the adverse condition against which the allocation is evaluated.

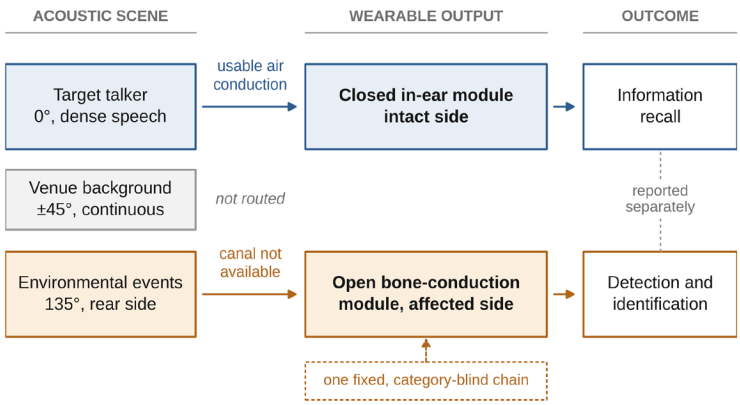


Figure 1: The OpenIN X information architecture. Placement follows anatomical suitability; content follows information density. Speech and environmental outcomes are measured and reported separately.

The environmental processing chain is deliberately category-blind: it does not identify an event as a bicycle, an alarm, or a human call, and it does not translate events into learned tone patterns. The same offline chain is applied to every event in three stages—band-limiting and spectral shaping to the usable response of the selected transducer at its wearing position, dynamic-range compression to a fixed presentation level, and, where needed, onset emphasis that sharpens an event's attack without changing its spectral identity. The parameters are set once, when the transducer is calibrated, and held constant across all conditions and event categories.

The design board in Figure 2a records how this allocation became a wearable form, starting from the question sketched at its top left: how an open-ear band and a true-wireless in-ear bud could become one device. Side-view form studies around the pinna reject a sharp transition in favour of a continuous curve; perspective and top-view head-envelope sketches place the band behind the head with an ear unit at each end; and the in-ear (TWS) unit is drawn as a separate part joining the ear unit beside the open (OWS) part. Figure 2b–c shows the resulting design as worn: the band passes behind the head, an open unit labelled OPEN PART rests in front of the ear, and an in-ear unit labelled IN PART sits in the concha. The renderings show one end of the band carrying both units; in use, the allocation in Figure 1 assigns the in-ear unit to the intact ear and the open unit to the affected side.

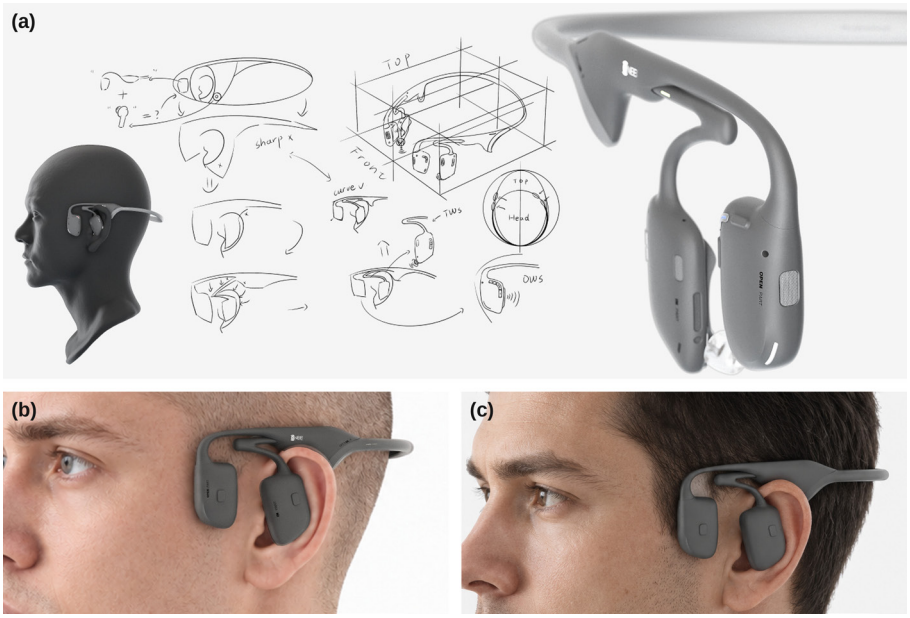


Figure 2: Design development and resulting form. (a) Design board: concept question (open band + in-ear bud), form studies (sharp rejected, curve adopted), perspective and head-envelope construction, in-ear (TWS) and open (OWS) parts, digital fit on a head model, and a close-up rendering. (b, c) Design renderings of the integrated form as worn.

The disclaimers stated above fix the outcome language used throughout: detection and identification, not situational awareness. Behavioural data come from commercially available in-ear and bone-conduction devices chosen for stable, controllable output, while the renderings in Figure 2 show how the allocation can be organised as one wearable; they are evidence of form and integration intent, not of acoustic performance.

FEASIBILITY EVALUATION: METHOD

Design and Participants

The study uses a within-participant design with four output conditions: C0, a baseline with neither module active; C1, in-ear speech support only; C2, bone-conduction environmental support only; and C3, the dual path with both active. Each participant completes all four in a counterbalanced order. Because condition order is counterbalanced while material group follows session round, difficulty differences between groups load onto practice position rather than onto condition, preserving the interpretability of condition contrasts. Complete counterbalancing requires a participant count that is a multiple of four, so the pilot sample of adults with self-reported normal hearing is planned in multiples of four, and the occluded side is recorded for each participant. All participants give written informed consent before testing.

Simulated Unilateral Attenuation and Apparatus

Acute unilateral air-conduction attenuation is produced with an earplug or earmuff on the simulated affected side, and the opposite ear is treated as the intact side. Following Asp, Jakobsson and Berninger (2018), the manipulation is called acute unilateral attenuation, not simulated congenital aural atresia; the attenuation it achieves depends on fit and is not intended to reproduce a particular hearing level.

A commercially available closed in-ear earphone is worn on the intact side and a commercially available bone-conduction headset on the simulated affected side. Both stay in place for the whole session and only the playback channels change between conditions, so refitting adds no between-condition variance. Active noise cancellation, transparency, spatial audio, adaptive equalisation, and other optional processing are disabled. Playback synchronisation is checked before each session, and because no outcome depends on sub-second timing, wireless latency does not need independent calibration.

Participants sit at the centre of a quiet room (Figure 3). A target-speech loudspeaker is placed at 0° and 1.5 m, two venue-background loudspeakers at −45° and +45°, and an environmental-event loudspeaker behind the simulated affected side at 135° (mirrored to 225° for a left-ear occlusion), all at 1.5 m and near ear height. Target speech, combined background, and environmental events are each calibrated in dB(A) at the head position. A calibration block before formal testing sets the speech-to-noise ratio so that performance avoids floor and ceiling, after which all levels are held constant.

MATERIALS

Twelve continuous audio scenes are used, grouped into four material groups (A–D) of three scenes each, so that one group is presented per condition. Each scene runs approximately 56–57 s over a continuous venue background of indistinct conversation, footsteps, and equipment noise, and depicts a different public-service exchange, such as a document-review counter, a hospital outpatient desk, a bank counter or a pharmacy.

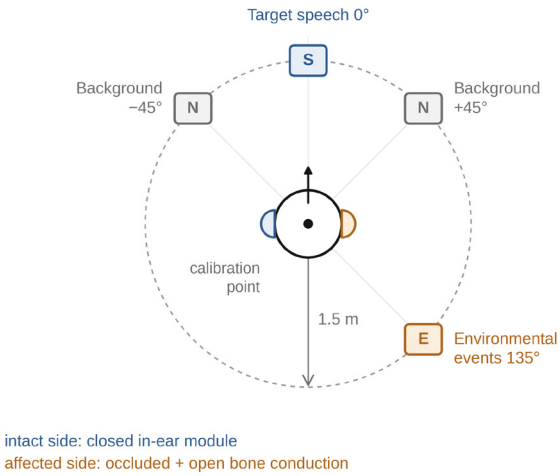


Figure 3: Top view of the listening room, drawn for a right-side occlusion. Environmental events arrive from behind the simulated affected side—the direction the occluded ear would ordinarily cover.

Within each scene the counter clerk speaks three times, and each turn carries one scored information unit drawn from time, quantity, required document, amount, and location. Three environmental slots are defined at fixed positions inside the three clerk turns. Each slot either contains a single brief event of approximately 1.2–1.6 s, or contains no event beyond the running background. Events belong to four categories: alarm or warning tone, bell-type sound, human call, and vehicle or movement sound. Every material group contains five event-present and four event-absent slots and includes all four event categories, with one or two instances of each.

Procedure

After a guided practice scene, participants complete the twelve formal scenes in four blocks of three, one block per condition. In every condition the front loudspeaker presents the dialogue and the rear loudspeaker the original environmental events. When in-ear speech support is active, a synchronised high signal-to-noise copy of the dialogue is also presented through the intact-side earphone; when bone-conduction support is active, the processed event is also presented through the bone-conduction device. Participants are told to attend to the dialogue and to remain aware of the surrounding environment, but not how many events will occur.

Each scene plays once without interruption and is followed immediately by nine items on a tablet: three five-alternative target-information questions, three six-alternative environmental-slot questions, and three seven-point ratings of target-speech clarity, environmental-event detectability, and overall task load. Both item types include a 'cannot determine' option; the environmental items additionally include an explicit 'no environmental event' option. Presentation and answer times are logged for every item.

Event performance is therefore measured by retrospective forced choice rather than a key press during the scene—a deliberate choice with a known cost. Consumer wireless devices add uncalibrated latency that would contaminate any reaction-time measure; forced choice avoids it and also captures whether an event was identified, not merely noticed, at the cost of temporal resolution and added memory load. The logged interval between item presentation and answer is therefore reported as answer-selection latency, never as reaction time.

Measures and Scoring

Target-information accuracy is the proportion of scored information units answered correctly, out of three per scene and nine per condition. Each item offers five alternatives including 'cannot determine', which is never the keyed answer and is scored incorrect, so chance performance is approximately 20 per cent.

Environmental-slot responses are scored as a signal-detection problem. Naming any event category on an event-present slot is a hit and on an event-absent slot a false alarm; answering 'no environmental event' is a miss or a correct rejection. Each participant contributes thirty-six slots, nine per condition, of which five are event-present and four event-absent. Hit rate is hits over event-present slots and false-alarm rate is false alarms over event-absent slots; category identification is computed over detected slots only, so a failure to identify is not confused with a failure to notice. A 'cannot determine' response is a declined decision, so it is excluded from both denominators and reported separately as an uncertainty rate; counting it as a correct rejection would deflate the false-alarm rate asymmetrically.

The three seven-point ratings are reported as condition means with standard deviations. The single task-load rating is a lightweight proxy only; it is not the NASA Task Load Index and is not comparable to published NASA-TLX values, and a full multidimensional workload instrument (Hart and Staveland, 1988) is planned for the confirmatory stage.

ANALYSIS PLAN

Given the sample size, analysis is descriptive and no confirmatory significance claim is made. Each outcome is reported as condition means and standard deviations, with participant-level values for the key contrasts so that a group mean is never the only evidence offered. Four planned within-participant contrasts organise interpretation: C1 − C0 and C3 − C2 estimate the contribution of in-ear speech support, C2 − C0 and C3 − C1 that of

bone-conduction support, and C3 is compared with C1 for speech retention and with C2 for event retention. The two outcome families are never collapsed into a composite score, which would conceal the trade-off the design exists to address.

EXPECTED OUTCOMES AND ILLUSTRATIVE ANALYSIS

If the allocation works as designed, three patterns should appear. In-ear support should raise target-information accuracy, placing C1 and C3 above C0 and C2; bone-conduction support should raise the event hit rate, placing C2 and C3 above C0 and C1; and the dual path should approach C1 on speech while retaining most of C2's event detection. A dual path that costs substantial speech accuracy, or that fails to lift detection above C1, would count against the allocation.

To show how these patterns will be read, and what a sample of this size can resolve, we simulated a dataset for four hypothetical participants. Item-level responses were drawn on the study's real answer keys from assumed rates that encode the predicted pattern, using a random seed fixed in advance, and were scored with the study's analysis script. Table 1 is therefore a worked example of the reporting format, not a finding.

Even with the predicted pattern built in, the simulated estimates are coarse. Target-information accuracy was 55.6% in C0, 97.2% in C1, 52.8% in C2 and 75.0% in C3, and the event hit rate was 47.5%, 46.3%, 95.0% and 100.0%. One simulated participant scored 22.2% on speech in C2 while the others scored 55.6–66.7%, and the hit-rate standard deviation in C1 reached 32.5 points, because each hit rate rests on five slots and moves in 20-point steps. At this scale the pilot can show whether the protocol runs and whether the predicted pattern appears at all, but not how large any effect is.

Table 1: Worked example from simulated data for four hypothetical participants (not observed): condition means, standard deviations in parentheses. *n = 3 (one simulated participant detected no events).

Measure	C0 Baseline	C1 In-ear	C2 Bone Cond.	C3 Dual Path
Target-information accuracy (%)	55.6 (18.1)	97.2 (5.6)	52.8 (21.0)	75.0 (5.6)
Event hit rate (%)	47.5 (9.6)	46.3 (32.5)	95.0 (10.0)	100.0 (0.0)
Event false-alarm rate (%)	8.3 (16.7)	0.0 (0.0)	12.5 (14.4)	18.8 (12.5)
Category identification, given hit (%)	70.8 (34.4)	61.1 (9.6)*	77.5 (20.6)	85.0 (10.0)
Speech clarity (1–7)	3.2 (0.4)	5.4 (0.2)	3.3 (0.7)	5.8 (0.9)
Event detectability (1–7)	2.9 (0.5)	2.3 (0.4)	5.8 (0.2)	5.3 (0.5)
Task load (1–7)	5.2 (0.2)	4.3 (0.2)	4.8 (0.4)	4.5 (0.3)

DISCUSSION

The evaluation asks whether target speech and processed environmental audio support two distinguishable functions when delivered through separate wearable outputs. A useful dual-path result does not require C3 to outperform every condition on every measure: C1 may well produce the highest target-information accuracy, since it optimises one task and leaves the other unsupported, and C2 the highest hit rate for the mirror-image reason. C3 is informative when it approaches C1 on speech while retaining much of C2 on events—a balanced profile across two concurrent activities, not a maximum on one metric.

Three design implications follow for accessibility-oriented hearing wearables generally, independent of the particular result. First, environmental awareness deserves to be an explicit design target with its own measure rather than an assumed by-product of better speech access; the hearing-protection literature shows how reliably it is lost when one channel is optimised in isolation. Second, an unused or unusable ear is a placement opportunity, not only a deficit: it can host an output that would be unacceptable on the ear carrying the primary task. Third, a category-blind chain is a defensible default for a safety-adjacent channel, avoiding the learning cost of symbolic codes and the failure modes of on-device classification, at the price of leaving interpretation to the user.

Several limitations bound what this study can establish. Acute unilateral attenuation cannot reproduce congenital anatomy or the listening strategies of the target population, and the sample supports feasibility and participant-level description only. The environmental stream is processed offline and synchronised to known event times, and retrospective forced choice measures detection and identification without temporal resolution. With five event-present and four event-absent slots per condition, hit and false-alarm rates move in steps of 20 and 25 points, so small differences should not be over-read. In the present materials the pattern of occupied slots is identical across groups and several categories recur at fixed positions, so a participant could learn the regularity across rounds; re-randomising occupancy and category while keeping the five-to-four balance corrects this. Finally, the integrated design is shown in renderings and supplies no evidence about comfort, fit stability, or acoustic equivalence with the evaluation apparatus.

Future work should recruit participants with genuine unilateral conductive hearing loss, compare the evaluation apparatus against integrated OpenIN X hardware, and implement real-time environmental acquisition. Loudness-matched controls are needed to separate the allocation strategy from the contribution of added level, and evaluation should extend to comfort, user control, and sustained use.

CONCLUSION

OpenIN X combines a task-oriented auditory information-allocation strategy with a wearable embodiment for listeners whose affected ear canal is unavailable to a conventional in-ear fit. Rather than treating unilateral

conductive hearing loss as a threshold to be corrected, it treats the listener's situation as a divided-attention problem, assigning dense speech and brief environmental events to two outputs matched to each side's capacity.

The protocol establishes a transparent chain from output function to behavioural outcome, keeps speech comprehension and environmental-event detection separate at every stage, and states in advance what an early-stage study of this size can and cannot support. The next steps are to complete the normal-hearing pilot and then to test the allocation with listeners who have unilateral conductive hearing loss, the population the design is for.

ACKNOWLEDGMENT

The authors thank the participants for their time and contribution to this study.

REFERENCES

Asp, F., Jakobsson, A.-M. and Berninger, E. (2018). "The effect of simulated unilateral hearing loss on horizontal sound localization accuracy and recognition of speech in spatially separate competing speech", Hearing Research, 357, pp. 54–63.

Asp, F. and Reinfeldt, S. (2018). "Horizontal sound localisation accuracy in individuals with conductive hearing loss: Effect of the bone conduction implant", International Journal of Audiology, 57(9), pp. 657–664.

Findlater, L. et al. (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. New York: ACM, pp. 1–13.

Gagné, J.-P., Besser, J. and Lemke, U. (2017). "Behavioral assessment of listening effort using a dual-task paradigm: A review", Trends in Hearing, 21, 2331216516687287.

Gygi, B., Kidd, G. R. and Watson, C. S. (2004). "Spectral-temporal factors in the identification of environmental sounds", The Journal of the Acoustical Society of America, 115(3), pp. 1252–1265.

Hart, S. G. and Staveland, L. E. (1988). "Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research", in Hancock, P. A. and Meshkati, N. (eds.) Human Mental Workload. Amsterdam: North-Holland, pp. 139–183.

Huang, N. and Elhilali, M. (2017). "Auditory salience using natural soundscapes", The Journal of the Acoustical Society of America, 141(3), pp. 2163–2176.

Kitterick, P. T., Smith, S. N. and Lucas, L. (2016). "Hearing instruments for unilateral severe-to-profound sensorineural hearing loss in adults: A systematic review and meta-analysis", Ear and Hearing, 37(5), pp. 495–507.

Laroche, C. et al. (2022). "Effect of hearing and head protection on the localization of tonal and broadband reverse alarms", Human Factors, 64(7), pp. 1105–1120.

Lindeman, R. W., Noma, H. and de Barros, P. G. (2007). "Hear-through and mic-through augmented reality: Using bone conduction to display spatialized audio", in Proceedings of the 6th IEEE and ACM International Symposium on Mixed and Augmented Reality. Piscataway: IEEE, pp. 173–176.

Pichora-Fuller, M. K. et al. (2016). "Hearing impairment and cognitive energy: The Framework for Understanding Effortful Listening (FUEL)", Ear and Hearing, 37(Suppl. 1), pp. 5S–27S.

Sakamoto, Y. et al. (2020). "Effects of FM system fitted into the normal hearing ear or cartilage conduction hearing aid fitted into the affected ear on speech-in-noise recognition in Japanese children with unilateral congenital aural atresia", The Journal of Medical Investigation, 67(1–2), pp. 131–138.

Siegbahn, M., Engmér Berglin, C., Hultcrantz, M. and Asp, F. (2021). "Adults with unilateral congenital ear canal atresia – sound localization ability and recognition of speech in competing speech in unaided condition", Acta Oto-Laryngologica, 141(7), pp. 689–694.

Smalt, C. J. et al. (2020). "The effect of hearing-protection devices on auditory situational awareness and listening effort", Ear and Hearing, 41(1), pp. 82–94.

Toya, T., Kobayashi, M., Nakamura, K. and Unoki, M. (2023). "Methods for improving word intelligibility of bone-conducted speech by using bone-conduction headphones", Applied Acoustics, 207, 109337.

Volgger, V. et al. (2022). "Audiological results and subjective benefit of an active transcutaneous bone-conduction device in patients with congenital aural atresia", European Archives of Oto-Rhino-Laryngology, 279(5), pp. 2345–2352.