Harnessing the power of artificial intelligence to combat the global burden of hearing loss

Nicholas A. Lesica[1], Nishchay Mehta[2], Joseph G. Manjaly[2], Blake S Wilson[3], Fan-Gang Zeng[4]

[1] Ear Institute, University College London, UK

[2] Royal National ENT Hospital & University College London Hospitals, UK

[3] Departments of Surgery, Biomedical Engineering, and Electrical and Computer Engineering, Duke University, USA

[4] Center for Hearing Research, Departments of Anatomy and Neurobiology, Biomedical Engineering, Cognitive Sciences, Otolaryngology - Head and Neck Surgery, University of California Irvine, USA

Abstract

Hearing healthcare currently meets only a small fraction of the global need. A labor-intensive care model limits access to services, especially in low- and middle-income countries, and existing therapies and assistive devices provide only limited benefit. Artificial intelligence has the potential to address these problems and transform hearing healthcare through advances in a number of areas: the fusion of disparate data to improve diagnosis and treatment; the automation of basic services to increase access, efficiency, and safety; and the development of new sensory devices that support the full richness of the human experience with or without restoration of hearing. Technology developers, together with clinicians and patients, should act urgently on this opportunity to develop a new model of hearing healthcare and create a world in which hearing loss is no longer a leading cause of disability.

Introduction

Hearing loss affects more than 500 million people worldwide and imposes an economic burden that is approaching $1 trillion annually (Wilson et al., 2017; World Health Organization, 2017). For individuals, the consequences can range from disrupted language development to declines in mental health, employment opportunities and quality of life (Livingston et al., 2020; Marschark et al., 2015; Tomblin et al., 2015). The scale of the problem continues to grow as demographic shifts increase the most vulnerable sectors of the population: older people and those in low- and middle-income countries (LMICs).

Current hearing healthcare, which is heavily reliant on specialized equipment and labor-intensive clinician services, is failing to cope. As a result, approximately 80% of the people worldwide who need treatment, mainly those in LMICs, are not receiving it (Barnett et al., 2017; Orji et al., 2020). The recent COVID-19 pandemic has further exposed the inadequacy of the current model. The need for social distancing has forced services to be greatly reduced or halted altogether and their return to normal remains uncertain.

Artificial intelligence (AI) has the potential to provide the technological advances required to transform the current service model. AI has already revolutionized related areas such as computer vision, automatic speech recognition and natural language processing, but has not yet made a significant impact on hearing healthcare. Hearing was once at the forefront of technological innovation in medicine. The ***cochlear implant (CI)***, for example, which enables the perception of sound through direct electrical stimulation of the auditory nerve, has provided hearing to almost one million people. It remains the most successful neural prosthetic in terms of both performance and penetration (Wilson and Dorman, 2008; Zeng et al., 2008). But in recent years, innovation has stalled; interest from academic researchers has waned and market failures that limit competition have allowed industry to become complacent (National Academies of Sciences, 2016; President's Council of Advisors on Science and Technology, 2016).

In this Perspective, we describe the most urgent and important hearing healthcare problems and the potential of AI to provide solutions. We first give a brief overview of the auditory system and its disorders. We then focus on three areas -- clinical inference, automated care, and assistive devices --- in which AI can have immediate impact. Finally, we discuss the mutual benefits of ongoing collaboration between the AI and hearing research communities with the potential to create a future in which hearing loss is no longer a barrier to human communication or fulfillment.

The auditory system and its disorders

The auditory system is a marvel of engineering. Its combination of microsecond temporal precision, sensitivity over more than 5 orders of sound magnitude, and flexibility to support tasks ranging from sound localization to music appreciation is still without parallel in other natural or artificial systems. It is because the auditory system is so powerful that its dysfunction is so devastating; through millennia of biological and cultural evolution, we have come to rely heavily on hearing for functions ranging from threat detection and navigation to communication and entertainment.

The remarkable performance of the auditory system is achieved through a complex interplay of biomechanical and neural components that implement operations such as signal conditioning, filtering, feature extraction, and classification in interconnected stages across the ear and brain (**Figure 1A**). But many of these components are delicate and the function of the system is susceptible to disruption at any of its stages from a number of different causes such as: genetic mutation; damage from noise exposure or toxic drugs; degradation through aging; or disruption of associated sensory or cognitive systems.

To diagnose hearing conditions, a wide range of data is often collected in an attempt to provide insight into the status of each processing stage (**Figure 1B-D**). The function of the ear itself is tested through mechanical and acoustic measurements, typically with the intent of isolating the source of any ***deafness*** (a loss of sensitivity to low-intensity sounds; the term "hearing loss" is more general and covers a range of different conditions). Electrophysiological measures of neural activity and imaging of brain structures are used to assess the early stages of the brain where several common conditions, such as ***tinnitus*** (the constant perception of a phantom sound, often a ringing), are thought to originate. For the most complex conditions, such as ***auditory processing disorders*** (e.g. difficulty understanding speech in noisy environments despite "normal" hearing), subjective measures such as psychoacoustic and cognitive tests are used.

Despite this wealth of data, the diagnosis and treatment of hearing conditions is often problematic. The primary difficulties arise from the multifactorial nature of the conditions and our limited understanding of their mechanistic underpinnings. A given condition (deafness, tinnitus, etc.) can be associated with dysfunction in many different processing stages, and dysfunction in a given processing stage can be associated with many different conditions. For the most complex problems, it is often difficult to disentangle sensory and cognitive components, and objective measures that can differentiate between different underlying causes do not yet exist (Liberman et al., 2016; Parthasarathy et al., 2020).

These difficulties are precisely why AI -- particularly machine learning (ML) -- has the potential to be so valuable in hearing healthcare. One of the main advantages of ML is a capacity to identify patterns in complex data that far exceeds that of conventional statistical techniques. Clinicians have achieved remarkable success in the diagnosis and treatment of hearing conditions through ad hoc processes, but these approaches are fundamentally limited. With the help of AI to assimilate vast amounts of disparate data, consider them in their full complexity, and infer the optimal course of action using technology that is widely available, hearing healthcare could be fundamentally transformed.

Figure 1 | Measures of auditory structure and function

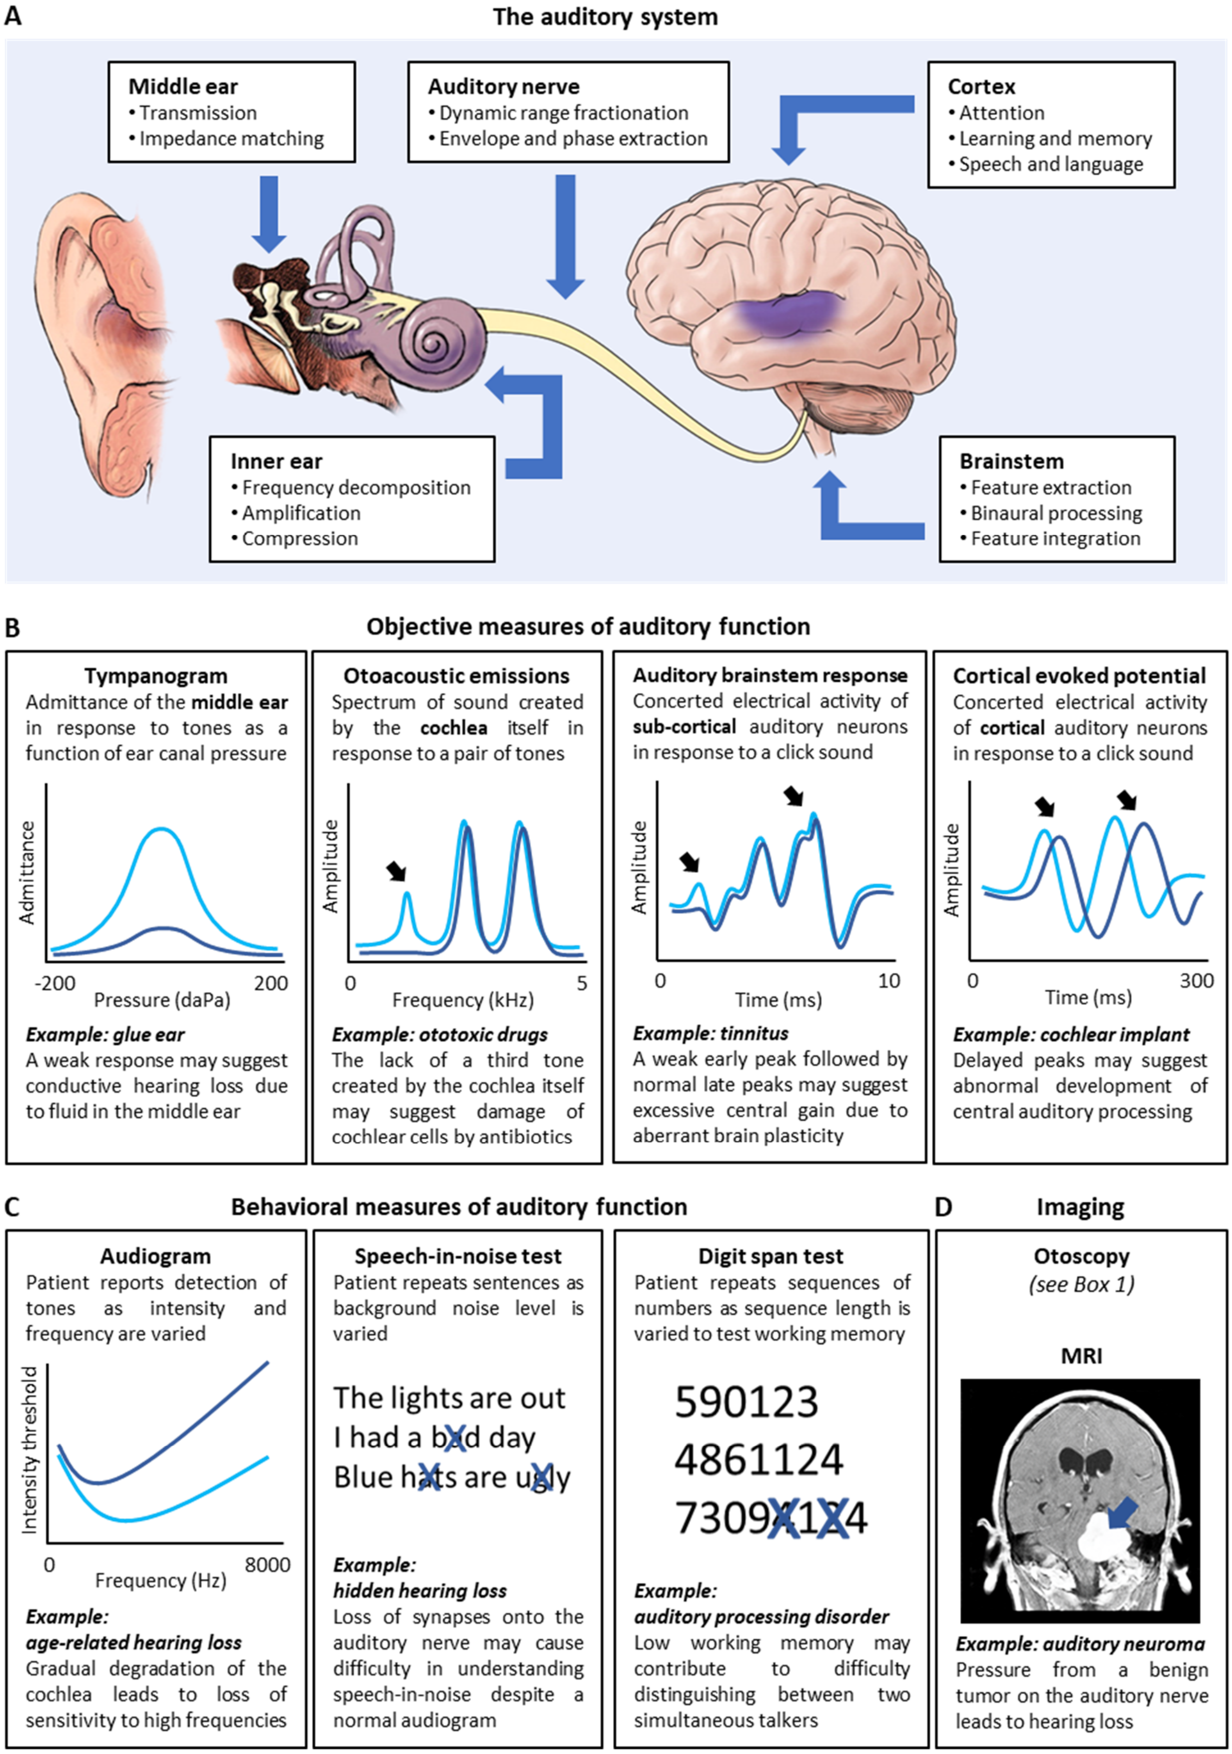

Figure 1 | Measures of auditory structure and function

(A) The major processing stages of the auditory system. Sound that enters the ear canal causes vibrations of the ear drum. These vibrations are transmitted by the ossicle bones in the middle ear to the fluid-filled cochlea in the inner ear. Hair cells in the inner ear amplify and transduce motion of the cochlear fluid into electrical signals that are sent to the brain. These signals are processed by several specialized pathways in the brainstem and the resulting information is integrated in the cortex to produce a coherent auditory experience. Some of the key functions performed at each processing stage are indicated in the boxes. Image modified from (Bance, 2007) (permission requested). (**B)** Examples of objective measures used in hearing assessment. Each panel describes one measure and provides a schematic illustration of the associated results from a patient with (dark blue) and without (light blue) a hearing condition. Key differences are indicated by the arrows. (**C)** Examples of subjective measures used in hearing assessment. **(D)** Examples of imaging used in hearing assessment.

Clinical inference

Clinical inference is a core problem in all medical disciplines. It generally involves the use of diagnostic information about patients and their symptoms to infer underlying causes and predict benefits of different treatment options. The potential for AI to improve clinical inference in hearing has already been recognized and has led to recent efforts in a number of areas such as: estimation of risk from industrial noise measurements (Zhao et al., 2019); classification of deafness from genetic markers (Shew et al., 2019); and identification of specific cochlear damage from hearing tests (Chang et al., 2019).

The most focused work so far has been on the diagnosis of conditions in the ***middle ear*** (the space behind the ear drum that contains the vibrating ossicle bones that transmit sound from the ear canal to the cochlea; see **Box 1**) (Livingstone et al., 2019; Myburgh et al., 2018; Viscaino et al., 2020). Middle ear infections are common in children; in fact they are most frequent reason for children to visit the doctor, take antibiotics, and have surgery (Rovers et al., 2004). Initial diagnoses are typically made based on images of the ear drum, but the classification of these images by clinicians is highly variable. Existing ML technologies (e.g. deep convolutional neural networks) have already been applied to similar problems in other areas of medicine such as classifying skin lesions (Esteva et al., 2017) or identifying diabetic retinopathy (Gulshan et al., 2016). Initial work has shown that similar approaches could be used to identify some of the most prevalent middle ear conditions with high accuracy.

In one recent study, for example, image classifiers were used to identify six different middle ear conditions (Cha et al., 2019). Transfer learning from a database of 10,000 ear drum images was used to modify a number of publicly-available, pre-trained convolutional neural networks and the two top performers (Inception-V3 and ResNet101) were used to create an ensemble classifier that combined the outputs from the two networks. Several variations of the network configuration were explored such as adding an additional hidden layer before the final classification layer or including hand-designed modifications of image color channels as a preprocessing stage. The accuracy of the best classifier reached 90% and there were indications that performance might increase further with a larger database.

Although these initial results are promising, much more work is needed to develop practical applications that can be widely deployed in clinics and, more importantly, in remote settings without specialist resources. Processing power is unlikely to be a constraint; the accuracy of a classifier based on the MobileNet-V2 network, which has only a few million parameters, nearly matched that of the top performers, suggesting that on-device inference should be possible if needed. But robustness will be a challenge. The images used for the initial study were taken by specialist staff using expensive imaging equipment. For wide deployment, applications will need to be accurate despite variations in images taken by non-specialists using inexpensive cameras (Jayawardena et al., 2019).

It should also be possible to go beyond simply automating image-based diagnosis to provide detailed, personalized recommendations for treatment by fusing images with other patient data (Binol et al., 2020). For

Box 1 | AI for middle ear conditions: improved diagnosis and treatment of glue ear

Images of the ear drum taken from an otoendoscope within the ear canal. Left: normal, with the white ossicle bones that connect the ear drum to the cochlea visible behind the ear drum. Middle: glue ear, with fluid behind the ear drum. Right: with grommet in place

Middle ear conditions affect 10% of all people, primarily as young children (Schilder et al., 2016). Despite their prevalence, diagnosis remains problematic: accuracy has been estimated at 50% for non-specialists and 75% for specialists (Pichichero and Poole, 2001). Inaccurate diagnosis leads to a number of problems, including the prescription of unnecessary antibiotics with consequences for the patient and the planet (Lannon et al., 2011). But most of the people with middle ear conditions live in LMICs with no access to hearing healthcare and, thus, their conditions go untreated.

AI-assisted care for middle ear conditions could bring dramatic improvements in both efficacy and accessibility. Consider, for example, otitis media with effusion, also known as ***glue ear***, which is characterized by a build-up of thick fluid in the middle ear that prevents the transmission of sound from the ear drum to the cochlea. The majority of cases are mild and self-resolving, but children with persistent glue ear face developmental challenges with significant long-term consequences (Bennett et al., 2001; Roberts et al., 1986).

Early detection

Detection of glue ear can be difficult; it is common in toddlers who may be unaware of or unable to describe their hearing loss, so the condition may go unnoticed until it has led to delays in the development of speech and language. However, recent studies have demonstrated that ML can be used to identify glue ear from ear drum images (see text). If these technologies can be developed into applications that can be incorporated into routine health checks, they may enable the early detection of glue ear and help to avoid months of developmental disruption in both high-income countries and LMICs.

Prediction of time to resolution

Even after diagnosis, there is uncertainty regarding the appropriate course of treatment. The most common treatment is the insertion of grommets (ventilation tubes) into the ear drum to drain the fluid, a surgical procedure performed under general anesthesia with risk of damage to the ear. Since many cases resolve spontaneously, surgery is not usually performed until after several months of "watchful waiting". The development of AI with the ability to fuse ear drum images with other information about patient history, genetics, etc. and predict time to resolution could help to avoid both unnecessary treatment and unnecessary waiting. In LMICs where surgery may not be readily available, children with cases that are not likely to resolve quickly could be provided with bone conduction headbands (which transform sound into vibrations of the skull that reach the cochlea without passing through middle ear) to restore hearing and maintain development.

Prediction of surgical benefit

Even for those who do receive grommets, the uncertainty continues. The grommets typically remain in place for 6-9 months before being ejected by the healing ear drum, by which time the underlying condition has resolved and hearing has returned to normal. However, the condition recurs in approximately 30% of patients and repeat surgery is usually needed, with further risk of damage to the ear (Browning et al., 2010). AI that is able to predict which cases are unlikely to be resolved by traditional grommets would allow for the use of alternative approaches from the start.

more complex conditions, this fusion will be essential to disentangle the complex interactions between the ear and brain. But making optimal use of the disparate data collected for each patient will be difficult. For some conditions, there is little agreement on best criteria for diagnosis (for an example of ongoing debates, see discussion of auditory processing disorders (Iliadou and Kiese-Himmel, 2018; Neijenhuis et al., 2019)). And even when diagnosis is straightforward, the best treatment may not be clear.

In the case of tinnitus, for example, there is no widely accepted treatment, largely because the underlying mechanisms remain poorly understood (Shore and Wu, 2019). There are some treatments that seem to be effective in some cases, such as cognitive behavioral therapy, but there is currently no method for predicting which treatment might be most beneficial for a given patient other than trial-and-error (Baguley et al., 2013; Cima et al., 2014). By applying ML to the full complement of patient data, it may be possible to predict which treatment will provide the most benefit for a given patient, even without a detailed mechanistic understanding of the underlying problem.

Assembling the datasets required to make the best use of AI will be a challenge. Patients are often served by specialists across multiple healthcare sectors, with each holding vital pieces of information. Even in high-income countries, the systems used to capture, store, share and analyze condition-specific data are largely inadequate. Unlocking and integrating these data will be critical for the precision “phenotyping” required to recommend individualized treatments based on specific underlying conditions, as well as to optimize the testing of new treatments in clinical trials and to make the most of the valuable data that these trials produce (Robinson, 2012).

Efforts are underway either to join existing hearing datasets (NIHR, 2020) or create new disease or treatment registries for analysis (Sing Registry, 2020; Yung et al., 2005). The initial success of the efforts will depend on the alignment of incentives across the different aspects of clinical practice. But as clinicians and administrators begin to see the benefits of data sharing for the rapid translation of research into improved care, more resources will be allocated to building the necessary data infrastructure. It is critical to ensure that the resource allocations faithfully reflect the global burden of hearing loss to avoid potential biases (Gianfrancesco et al., 2018). For example, technologies developed based on data from high-income countries may not be appropriate for use in LMICs where different conditions are prevalent.

Automated Care

At present, nearly all hearing healthcare services -- from initial screening and consultation through to follow-up and rehabilitation -- are provided in-person by highly trained staff using specialized equipment. This “high-touch” model restricts care to places where the required resources are readily available, thus excluding many LMICs, as well as remote locations in high-income countries (WHO, 2013). The accessibility of high-touch care has been further reduced by the COVID-19 pandemic; even in places with the required resources, vulnerable patients may be unwilling or unable to seek in-person care and staff may be unable to provide it safely.

In recent years, there have been efforts to improve accessibility through remote care in which a clinician provides services to a patient in a different location over the internet (Swanepoel et al., 2010). While this approach does have the potential to improve accessibility and efficiency, the scalability of any model that continues to rely heavily on specialist staff, equipment, and facilities is ultimately limited. Fortunately, many of the most common basic services in hearing healthcare can be readily automated.

One example is the measurement of an ***audiogram*** (the standard clinical test for deafness; **Figure 1C**). In a standard measurement, a clinician presents tones at different frequencies and intensities, generally according to a prescribed protocol, in order to determine the patient’s sensitivity threshold for each frequency. The patient signals their perception of a tone either verbally or by pushing a button or touchscreen.

The automation of this process in standard clinical conditions (i.e. with medical-grade earphones in a sound-proof chamber) is straightforward, and a recent study demonstrated that ML-based approaches can provide more comprehensive measurements in less time than the standard approach (Barbour et al., 2019). Gaussian process regression was used to estimate the probability of a listener detecting a tone across a range of different frequencies and intensities, with the frequency and intensity for each tone presentation chosen to maximize the expected decrease in the posterior variance of the estimates. Despite placing few constraints on audiogram shape (other than continuity in frequency and monotonicity in intensity), this approach was able to provide an estimate of sensitivity thresholds across the full range of audible frequencies in less time than is typically required to measure thresholds at only six discrete frequencies. This additional information may be critical for differentiating between different hearing conditions.

Another example of a basic service that can be readily automated is the "fitting" of a CI (the fine tuning of the device's free parameters after the patient has fully recovered from surgery). In the standard approach to fitting, a clinician establishes the dynamic range of electrical stimulation by adjusting the current emitted by each electrode on the device while asking the patient to report their subjective sensation. The clinician then programs the device with a sound-to-current mapping based on the measured dynamic range (often through an ad hoc process (Vaerenberg et al., 2014)) and makes adjustments as the patient reports their perception of sound. Proof-of-concept studies have established that an automated fitting using Bayesian networks can recommend maps that are comparable to those chosen by clinicians (Battmer et al., 2015; Buechner et al., 2015; Meeuws et al., 2017) and that the entire fitting process can, in principle, be done by the patient themselves without the need for a clinician (Meeuws et al., 2020).

The challenge comes when attempting to automate services in non-clinical settings, e.g. remote locations in LMICs, where neither the specifics of the equipment nor the environment can be guaranteed. For CI fitting, the safety of the patient must be considered; inappropriate electrical stimulation can cause non-auditory sensations or even pain. For audiogram measurement, the results must be robust to uncontrolled acoustic conditions (Sandström et al., 2020). AI can potentially help by allowing the problem to be framed as audiogram inference rather than audiogram measurement. Given a sufficient training dataset of paired audiograms measured under non-ideal and ideal conditions (perhaps supplemented by data augmentation) along with calibration routines to determine background noise levels, earphone properties, etc., it may be possible to infer the true audiogram from non-ideal measurements.

Assistive Devices

There are not yet any biological treatments for most forms of hearing loss, so care is generally limited to the provision of assistive devices (**Figure 2**). For profound deafness, the only available option is to provide direct electrical stimulation of the auditory nerve through a CI. For partial deafness, a ***hearing aid (HA)*** may be able to help the ear process sound naturally by providing suitable amplification. In both types of devices, sounds are received by microphones and digitized for signal processing that is customized for each listener through a fitting process as described above. The processed signals are then used to generate sound (for HAs) or current (for CIs).

The signal processing in hearing devices improved rapidly during their early development but in recent years progress has been stagnant (Lesica, 2018; Wilson, 2015; Zeng, 2017). The use of AI has thus far been limited to secondary functions such as automated adjustment or fall detection. But AI has the potential to dramatically improve the primary signal processing in hearing devices and, in particular, to address the most common problem reported by device users: difficulty understanding speech in a setting with multiple talkers or background noise (the so-called "cocktail party" problem). Many attempts have been made to address this difficulty, such as using directional microphones to isolate sound sources in front of the listener or denoising

Figure 2 | AI for the hearing devices of the future

A

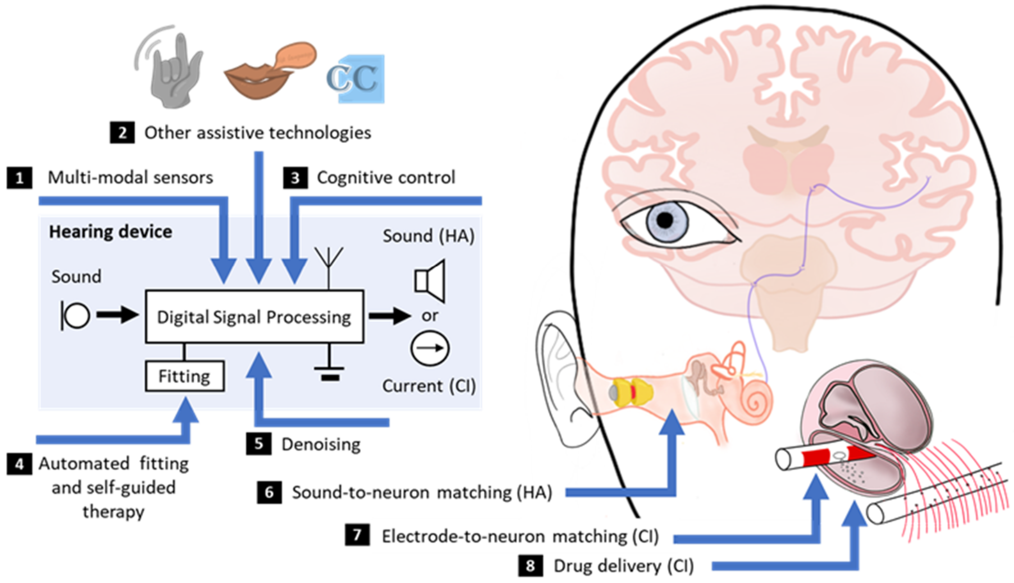

B

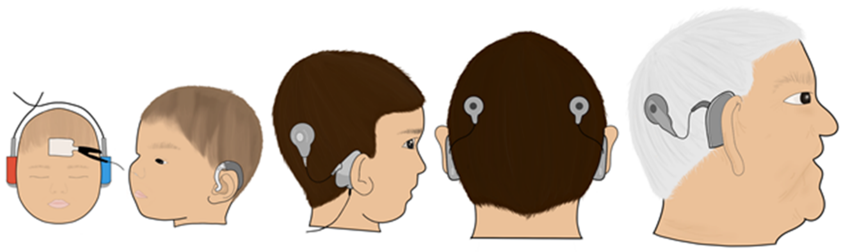

Screening and diagnosis

Now:
Initial screening reveals deficit and periodic testing is used to monitor progression. HAs are provided until loss becomes severe and a CI surgery is scheduled.

- Decisions are made with high uncertainty about the nature of the deficit and the relative benefits of different treatments

With AI:
Data from hearing tests are fused with genetics, imaging, biomarkers and other patient information to:

- Determine the cause of the deficit
- Predict the time course and ultimately severity of its progression
- Predict the relative benefit of different HAs and CIs
- Determine the optimal time for intervention with respect to development, mental health and lifestyle

Devices and implantation

Now:
Current hearing devices work well for speech in quiet, but not for speech in noise or for complex sounds such as music.

- HAs restore audibility but sounds remain distorted
- CIs produce diffuse electrical stimulation and lack intraoperative feedback or biological therapy

With AI:
Improved HAs or CIs are integrated with other technologies 2 to:

- Provide task-specific processing of incoming sounds e.g. speech denoising (see text) 5
- Provide multi-modal augmented-reality and closed-loop cognitive control (see text) 1 3
- Optimize sound-to-sound (HA) or sound-to-current (CI) mapping to match normal or near-normal neural activity (see text) 6 7
- Determine optimal timing and dose of associated biological therapies e.g. for nerve regeneration 8

Fitting and therapy

Now:
Fitting of CIs by clinician through an ad hoc process. Support for development through speech and language therapist.

- In-person service is labor intensive and decisions are made with high uncertainty surrounding optimal device settings or therapy progression

With AI:
Automated device adjustment and self-guided therapy to:

- Provide personalized service, especially for special cases e.g. autism or dementia
- Provide continuous monitoring of device status and warnings of malfunction
- Provide ongoing optimization of device settings to account for brain plasticity
- Provide secure, always-on therapy with continuous adaptation to developmental or degenerative needs 4

Figure 2 | AI for the hearing devices of the future

(A) The key elements of future hearing devices. Current hearing devices use a microphone to pick up sound, which is amplified and filtered before being digitized for signal processing; the processing parameters are fixed after fitting in an audiologist's office; the processed digital signals are converted to either an analog signal delivered to a speaker in hearing aids (HAs) or an electrical signal delivered to electrodes in cochlear implants (CIs) (bottom-left insert). (**B)** Examples of how AI could transform the experience of a deaf person throughout their entire life. The boxes indicate the current state-of-the-art (Now) and the potential for improvement (With AI) in screening and diagnosis (left), devices and implantation (middle), and fitting and therapy (right).

sounds through simple filtering based on the typical low-order statistics of speech, but the benefits of these features in real-world listening conditions are limited (Cox et al., 2014; Humes et al., 2017).

Recent work has demonstrated that approaches based on deep neural networks can be used to dramatically improve understanding of speech in noise for HA users. In just a few years, this work has progressed rapidly from separating the voice of a known talker from steady-state noise to separating multiple unknown talkers in reverberant environments (Wang and Chen, 2018), which required solving the "permutation problem" of assigning a network output for each sound source in the input when the true number and composition of sources is unknown. This feat was achieved by processing incoming sound in two stages: a frame-by-frame source separation and dereverberation stage that uses a U-net architecture (an encoder-decoder convolutional neural network with skip connections) followed by a temporal convolutional network that connects sources across frame sequences (Liu and Wang, 2019). When this processing is used to denoise speech, the performance of HA users in recognition tasks can match or even exceed normal levels (Healy et al., 2020). Similar approaches are also being developed for CIs and have produced promising initial results (Goehring et al., 2019; Lai et al., 2018).

Currently this processing must be performed offline, so further work is needed to achieve similar performance with causal networks that can run in real-time. Ideally, the computational and power requirements would also be reduced to allow the processing to be run directly on a hearing device (Wu et al., 2019). Alternatively, provided that the latency can be sufficiently reduced, at least some of the processing could be performed on a device-connected smartphone or even in the cloud. But if the delay between either the perception of sounds and their associated visual cues or the production and perception of the user's own voice becomes too large, it can be severely disruptive. A latency of 10-20 ms is generally considered tolerable by hearing device users (Goehring et al., 2018).

Separating different sound sources is only the first step toward helping listeners overcome difficulties understanding speech in noise. The second step, which is even more challenging, is determining which sources a device should amplify and which it should attenuate. In some situations, e.g. a single talker in a background of continuous fan noise, it may be obvious which source is of interest. But in others, e.g. a room full of multiple talkers, a source that is of primary interest one minute may become a distraction the next.

To address this problem, efforts are underway to bring hearing devices under "cognitive control". When a listener is attending to a particular sound source, the fluctuations in their brain's neural activity track the fluctuations in the amplitude of the attended source (Mesgarani and Chang, 2012). Thus, the source of interest at any given time can be inferred from correlations between recorded neural activity and possible sources of interest. Initial studies suggest that recordings that are sufficient to identify the source of interest can be obtained from a single electrode within the ear canal, which could easily be integrated with a hearing device (An et al., 2020; Fiedler et al., 2017; O'Sullivan et al., 2015).

Another promising approach is to move beyond hearing devices per se toward a more comprehensive augmented reality (AR) system (Brown Jaloza, 2020). If current trends continue, systems of integrated wearable and associated devices with a variety of multi-modal sensors of external and internal signals will soon become common. Such systems would provide powerful platforms to support deaf people (see **Box 2**). For example, to support speech understanding, AR glasses could provide eye tracking to aid inference of the current sound source of interest along with real-time speech-to-text captioning for instances when auditory perception fails. Such multi-modal systems would also open up new possibilities for treating conditions typically associated with deafness, such as tinnitus, memory loss, or dizziness. There is much work to be done before the full potential of assistive devices can be realized but, with the power of AI, there is little doubt that each of the individual technical challenges can be overcome; the real difficulty lies in integrating the various technologies to provide a seamless user experience.

Outlook

The current model of hearing healthcare improves the lives of millions of people every year. But it is far from optimal: children with middle ear conditions are triaged to "watchful waiting" while their development is disrupted; people with tinnitus are subject to treatment by trial-and-error, often with little or no benefit; the deaf are provided with devices that don't allow them to understand speech in noise or enjoy music. And those are the lucky ones: most people with hearing conditions live in LMICs with little or no access to treatment or support of any kind.

Despite the potential for AI to produce dramatic improvements, it has yet to make a significant impact. We have described opportunities for AI to reshape hearing healthcare and have outlined a vision for a future where AI not only transforms the diagnosis and treatment of hearing conditions but also supports those who live with deafness to engage with the world on their own terms. But for this potential to be realized, there are challenges that must be overcome.

The prevailing business model of hearing healthcare will need to change. Regulations that restrict the manufacturing and distribution of hearing devices have created a highly-concentrated market in which consumers are forced to choose from only a small number of manufacturers and service providers (National Academies of Sciences, 2016; President's Council of Advisors on Science and Technology, 2016). But action is finally being been taken to remove unnecessary barriers and increase competition (Warren and Grassley, 2017), leading to the recent introduction of ***personal sound amplification products (PSAPs)***, which can be sold over the counter and provide comparable benefit to standard hearing aids at a fraction of the cost (Brody et al., 2018; Cho et al., 2019; Humes et al., 2017). These developments have paved the way for full market disruption through AI-powered devices that provide even more benefit along with automated services.

Our limited understanding of the auditory system and its disorders presents a more fundamental challenge, but even this can be overcome through close collaboration between AI and hearing researchers. Initial attempts to make direct comparisons between deep artificial neural networks and biological neural networks in the brain have already provided insights into the computational mechanisms that underlie hearing and generated new hypotheses to be tested experimentally (Fontan et al., 2020; Huang et al., 2018; Kell et al., 2018; Keshishian et al., 2020; Schrimpf et al., 2020). To realize the full potential of this approach, AI researchers will need to develop new tools that are able to match the full complexity of auditory processing, with parallel computations across multiple timescales from microseconds to minutes (Picton, 2013), integration of inputs from other sensory and cognitive modalities (Atilgan et al., 2018), and flexibility to perform a number of qualitatively different tasks (Bregman, 1994).

Hearing researchers can contribute to this development by sharing their understanding of how these complexities are handled by the biological neural networks in the brain. The resulting tools will have the potential to be transformative not only for hearing, but also for other domains in which multi-scale, multi-modality, and multi-task capabilities are critical. Hearing researchers can also share experience gained through a long history of wearable device design (Levitt, 2007) to help hardware developers meet the needs of modern

Box 2 | Supporting multiple normals through AI

Hearing healthcare is focused on treating deafness, but this outcome is not always feasible or even desirable. The majority of those who are currently profoundly deaf will never have their hearing restored simply because of limited surgical capacity. And while CIs have clearly improved the lives of many deaf people, they would not necessarily do so for everyone. Some people may have a specific condition that cannot be addressed by a CI. Others who could have their hearing restored may prefer to remain deaf; not all people with hearing loss view it as a problem to be fixed (National Association of the Deaf, 2020).

While AI can certainly transform restorative treatments for deafness, it's impact could be even larger for those who remain deaf. Much of the disability associated with deafness arises from the fact that hearing is currently required for engagement in society. AI has the potential to bring about a new societal model with support for "multiple normals," in which alternative modes of engagement are readily available (Friedner et al., 2019). But the inclusion of deaf people in the development of new technologies at the earliest stages is critical to ensure that the results are matched to user needs (Hill, 2020).

Supporting informed decision making

The benefit that an individual receives from a CI can vary widely. Given that a CI also has downsides -- significant upfront and ongoing costs, risks and complications associated with surgery, continued dependence on associated support and services, etc. -- decisions about whether to undergo implantation can be difficult. Accurate predictions of benefit would be a great help; unfortunately, such predictions are not currently available. Attempts to explain variation in CI outcomes through traditional approaches have been largely unsuccessful (Zhao et al., 2020). But efforts to apply ML to the problem have produced promising initial results.

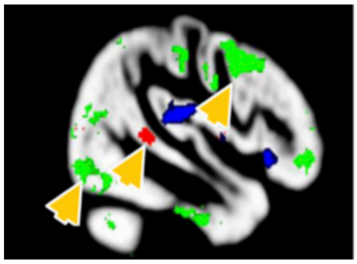

A brain image indicating areas (red and green) where pre-implantation morphology was predictive of CI benefit, such as occipital and pre-frontal cortices, and areas (blue) that were impacted by deafness but were not predictive of benefit, such as primary auditory cortex. Image from (Feng et al., 2018) (permission requested).

In one recent study, a support vector machine classifier was used to predict improvements in speech perception in children after implantation (Feng et al., 2018). The inputs to the classifier were morphological measures of neural preservation from MRI images in higher-level auditory and cognitive regions. Based on these image data alone, the correlation between the classifier prediction and the actual benefit observed 6 months after implantation approached 0.5. With further development to build predictive models that fuse image data with other measures of auditory structure and function (see **Figure 1**) and other patient data, much more accurate predictions may be possible.

Supporting hearing-optional communication

It is becoming increasingly easy to imagine a world in which deafness is not a disability, as AI is already making many settings more inclusive. In higher education, for example, much of the content is delivered as structured communication from teacher to students through technology platforms on which accessibility features are now readily available; standard software, such as Powerpoint, has the capacity to provide captions in multiple languages in real-time during ongoing presentations. The recent switch to remote learning because of COVID-19, which requires all communication between teachers and students to be routed through technology platforms, provides an opportunity to make accessibility features part of standard leaning models by default.

A translation device with stretchable sensor arrays on each finger attached to a wireless circuit board on the wrist. Image from (Zhou et al., 2020) (permission requested).

Supporting alternative modes of unstructured social communication is more challenging, as many deaf people communicate through signed rather than spoken language. But technologies for real-time automated translation can potentially bridge this gap. One recent study demonstrated the potential for a glove-like device that tracks finger movements to enable translation from American Sign Language to English (Zhou et al., 2020). This technology required the coordinated development of hardware that is comfortable, durable, and flexible and associated software to classify signals from the device using support vector machines. Though the overall accuracy of the system in this initial study was 98%, the vocabulary was limited to only 11 gestures, so much more work is needed to enable use of the full complement of gestures as well as integration with facial and other movements. Applications based on such technology have the potential to support natural communication not only between deaf and hearing people but also between deaf people from different countries, each of which has its own unique signed language.

wearables: low-latency, always-on processing; miniaturized, low-power platforms; and durability and comfort for long-term, frequent use.

Ongoing collaboration between AI and hearing researchers would create a win-win situation for both communities and also help to ensure that new technologies are well matched to the needs of users (Davies-Venn and Glista, 2019; Lindsell et al., 2020). With coordinated and bold efforts, we could together spark another technological revolution that would dramatically transform hearing healthcare. With the power of AI, a world in which hearing loss is no longer a disability, even for those who continue to live with it, is finally within reach.

Acknowledgements

We are grateful to Shievanie Sabesan, Dani Sive, and Andreas Fragner for their help with this work and to Krystal Cachola and JoAnne Gu for the artwork in Figure 2.

References

An, W.W., Pei, A., Noyce, A.L., and Shinn-Cunningham, B. (2020). Decoding auditory attention from single-trial EEG for a high-efficiency brain-computer interface. In 2020 42nd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC), pp. 3456–3459.

Atilgan, H., Town, S.M., Wood, K.C., Jones, G.P., Maddox, R.K., Lee, A.K.C., and Bizley, J.K. (2018). Integration of Visual Information in Auditory Cortex Promotes Auditory Scene Analysis through Multisensory Binding. Neuron *97*, 640-655.e4.

Baguley, D., McFerran, D., and Hall, D. (2013). Tinnitus. The Lancet *382*, 1600–1607.

Bance, M. (2007). Hearing and aging. CMAJ *176*, 925–927.

Barbour, D.L., Howard, R.T., Song, X.D., Metzger, N., Sukesan, K.A., DiLorenzo, J.C., Snyder, B.R.D., Chen, J.Y., Degen, E.A., Buchbinder, J.M., et al. (2019). Online Machine Learning Audiometry. Ear Hear. *40*, 918–926.

Barnett, M., Hixon, B., Okwiri, N., Irungu, C., Ayugi, J., Thompson, R., Shinn, J.B., and Bush, M.L. (2017). Factors involved in access and utilization of adult hearing healthcare: A systematic review. The Laryngoscope *127*, 1187–1194.

Battmer, R.-D., Borel, S., Brendel, M., Buchner, A., Cooper, H., Fielden, C., Gazibegovic, D., Goetze, R., Govaerts, P., Kelleher, K., et al. (2015). Assessment of "Fitting to Outcomes Expert" FOX™ with new cochlear implant users in a multi-centre study. Cochlear Implants Int. *16*, 100–109.

Bennett, K., Haggard, M., Silva, P., and Stewart, I. (2001). Behaviour and developmental effects of otitis media with effusion into the teens. Arch. Dis. Child. *85*, 91–95.

Binol, H., Moberly, A.C., Niazi, M.K.K., Essig, G., Shah, J., Elmaraghy, C., Teknos, T., Taj-Schaal, N., Yu, L., and Gurcan, M.N. (2020). Decision fusion on image analysis and tympanometry to detect eardrum abnormalities. In Medical Imaging 2020: Computer-Aided Diagnosis, (International Society for Optics and Photonics), p. 113141M.

Bregman, A.S. (1994). Auditory Scene Analysis: The Perceptual Organization of Sound (MIT Press).

Brody, L., Wu, Y.-H., and Stangl, E. (2018). A Comparison of Personal Sound Amplification Products and Hearing Aids in Ecologically Relevant Test Environments. Am. J. Audiol. *27*, 581–593.

Brown Jaloza, L. (2020). Inside Facebook Reality Labs Research: The Future of Audio.

Browning, G.G., Rovers, M.M., Williamson, I., Lous, J., and Burton, M.J. (2010). Grommets (ventilation tubes) for hearing loss associated with otitis media with effusion in children. Cochrane Database Syst. Rev.

Buechner, A., Vaerenberg, B., Gazibegovic, D., Brendel, M., De Ceulaer, G., Govaerts, P., and Lenarz, T. (2015). Evaluation of the “Fitting to Outcomes eXpert” (FOX®) with established cochlear implant users. Cochlear Implants Int. *16*, 39–46.

Cha, D., Pae, C., Seong, S.-B., Choi, J.Y., and Park, H.-J. (2019). Automated diagnosis of ear disease using ensemble deep learning with a big otoendoscopy image database. EBioMedicine *45*, 606–614.

Chang, Y.-S., Park, H., Hong, S.H., Chung, W.-H., Cho, Y.-S., and Moon, I.J. (2019). Predicting cochlear dead regions in patients with hearing loss through a machine learning-based approach: A preliminary study. PloS One *14*, e0217790.

Cho, Y.S., Park, S.Y., Seol, H.Y., Lim, J.H., Cho, Y.-S., Hong, S.H., and Moon, I.J. (2019). Clinical Performance Evaluation of a Personal Sound Amplification Product vs a Basic Hearing Aid and a Premium Hearing Aid. JAMA Otolaryngol Head Neck Surg *145*, 516–522.

Cima, R.F.F., Andersson, G., Schmidt, C.J., and Henry, J.A. (2014). Cognitive-Behavioral Treatments for Tinnitus: A Review of the Literature. J. Am. Acad. Audiol. *25*, 29–61.

Cox, R.M., Johnson, J.A., and Xu, J. (2014). Impact of advanced hearing aid technology on speech understanding for older listeners with mild to moderate, adult-onset, sensorineural hearing loss. Gerontology *60*, 557–568.

Davies-Venn, E., and Glista, D. (2019). Connected hearing healthcare: the realisation of benefit relies on successful clinical implementation.

Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., and Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature *542*, 115–118.

Feng, G., Ingvalson, E.M., Grieco-Calub, T.M., Roberts, M.Y., Ryan, M.E., Birmingham, P., Burrowes, D., Young, N.M., and Wong, P.C.M. (2018). Neural preservation underlies speech improvement from auditory deprivation in young cochlear implant recipients. Proc. Natl. Acad. Sci. U. S. A. *115*, E1022–E1031.

Fiedler, L., Wöstmann, M., Graversen, C., Brandmeyer, A., Lunner, T., and Obleser, J. (2017). Single-channel in-ear-EEG detects the focus of auditory attention to concurrent tone streams and mixed speech. J. Neural Eng. *14*, 036020.

Fontan, L., Cretin-Maitenaz, T., and Füllgrabe, C. (2020). Predicting Speech Perception in Older Listeners with Sensorineural Hearing Loss Using Automatic Speech Recognition. Trends Hear. *24*.

Friedner, M., Nagarajan, R., Murthy, A., and Frankfurter, R. (2019). Embracing Multiple Normals — A 12-Year-Old Boy in India with a Cochlear Implant. N. Engl. J. Med. *381*, 2381–2384.

Gianfrancesco, M.A., Tamang, S., Yazdany, J., and Schmajuk, G. (2018). Potential Biases in Machine Learning Algorithms Using Electronic Health Record Data. JAMA Intern. Med. *178*, 1544–1547.

Goehring, T., Chapman, J.L., Bleeck, S., and Monaghan*, J.J.M. (2018). Tolerable delay for speech production and perception: effects of hearing ability and experience with hearing aids. Int. J. Audiol. *57*, 61–68.

Goehring, T., Keshavarzi, M., Carlyon, R.P., and Moore, B.C.J. (2019). Using recurrent neural networks to improve the perception of speech in non-stationary noise by people with cochlear implants. J. Acoust. Soc. Am. *146*, 705.

Gulshan, V., Peng, L., Coram, M., Stumpe, M.C., Wu, D., Narayanaswamy, A., Venugopalan, S., Widner, K., Madams, T., Cuadros, J., et al. (2016). Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. JAMA *316*, 2402–2410.

Healy, E.W., Johnson, E.M., Delfarah, M., and Wang, D. (2020). A talker-independent deep learning algorithm to increase intelligibility for hearing-impaired listeners in reverberant competing talker conditions. J. Acoust. Soc. Am. *147*, 4106.

Hill, J. (2020). Do deaf communities actually want sign language gloves? Nat. Electron. 1–2.

Huang, N., Slaney, M., and Elhilali, M. (2018). Connecting Deep Neural Networks to Physical, Perceptual, and Electrophysiological Auditory Signals. Front. Neurosci. *12*.

Humes, L.E., Rogers, S.E., Quigley, T.M., Main, A.K., Kinney, D.L., and Herring, C. (2017). The Effects of Service-Delivery Model and Purchase Price on Hearing-Aid Outcomes in Older Adults: A Randomized Double-Blind Placebo-Controlled Clinical Trial. Am. J. Audiol. *26*, 53–79.

Iliadou, V., and Kiese-Himmel, C. (2018). Common Misconceptions Regarding Pediatric Auditory Processing Disorder. Front. Neurol. *8*.

Jayawardena, A., Waller, B., Edwards, B., Larsen-Reindorf, R., Anomah, J.E., Frimpong, B., Gina, A., Netterville, J., Saunders, J., and Basura, G.J. (2019). Portable audiometric screening platforms used in low-resource settings: a review. J. Laryngol. Otol. *133*, 74–79.

Kell, A.J.E., Yamins, D.L.K., Shook, E.N., Norman-Haignere, S.V., and McDermott, J.H. (2018). A Task-Optimized Neural Network Replicates Human Auditory Behavior, Predicts Brain Responses, and Reveals a Cortical Processing Hierarchy. Neuron *98*, 630-644.e16.

Keshishian, M., Akbari, H., Khalighinejad, B., Herrero, J.L., Mehta, A.D., and Mesgarani, N. (2020). Estimating and interpreting nonlinear receptive field of sensory neural responses with deep neural network models. ELife *9*, e53445.

Lai, Y.-H., Tsao, Y., Lu, X., Chen, F., Su, Y.-T., Chen, K.-C., Chen, Y.-H., Chen, L.-C., Po-Hung Li, L., and Lee, C.-H. (2018). Deep Learning–Based Noise Reduction Approach to Improve Speech Intelligibility for Cochlear Implant Recipients. Ear Hear. *39*, 795–809.

Lannon, C., Peterson, L.E., and Goudie, A. (2011). Quality Measures for the Care of Children With Otitis Media With Effusion. Pediatrics *127*, e1490–e1497.

Lesica, N.A. (2018). Why Do Hearing Aids Fail to Restore Normal Auditory Perception? Trends Neurosci. *41*, 174–185.

Levitt, H. (2007). A Historical Perspective on Digital Hearing Aids: How Digital Technology Has Changed Modern Hearing Aids. Trends Amplif. *11*, 7–24.

Liberman, M.C., Epstein, M.J., Cleveland, S.S., Wang, H., and Maison, S.F. (2016). Toward a Differential Diagnosis of Hidden Hearing Loss in Humans. PLoS ONE *11*, e0162726.

Lindsell, C.J., Stead, W.W., and Johnson, K.B. (2020). Action-Informed Artificial Intelligence—Matching the Algorithm to the Problem. JAMA *323*, 2141–2142.

Liu, Y., and Wang, D. (2019). Divide and Conquer: A Deep CASA Approach to Talker-Independent Monaural Speaker Separation. IEEEACM Trans. Audio Speech Lang. Process. *27*, 2092–2102.

Livingston, G., Huntley, J., Sommerlad, A., Ames, D., Ballard, C., Banerjee, S., Brayne, C., Burns, A., Cohen-Mansfield, J., Cooper, C., et al. (2020). Dementia prevention, intervention, and care: 2020 report of the Lancet Commission. The Lancet *396*, 413–446.

Livingstone, D., Talai, A.S., Chau, J., and Forkert, N.D. (2019). Building an Otoscopic screening prototype tool using deep learning. J. Otolaryngol. - Head Neck Surg. *48*, 66.

Marschark, M., Shaver, D.M., Nagle, K.M., and Newman, L.A. (2015). Predicting the Academic Achievement of Deaf and Hard-of-Hearing Students From Individual, Household, Communication, and Educational Factors. Except. Child. *81*, 350–369.

Meeuws, M., Pascoal, D., Bermejo, I., Artaso, M., Ceulaer, G.D., and Govaerts, P.J. (2017). Computer-assisted CI fitting: Is the learning capacity of the intelligent agent FOX beneficial for speech understanding? Cochlear Implants Int. *18*, 198–206.

Meeuws, M., Pascoal, D., Janssens de Varebeke, S., De Ceulaer, G., and Govaerts, P.J. (2020). Cochlear implant telemedicine: Remote fitting based on psychoacoustic self-tests and artificial intelligence. Cochlear Implants Int. *21*, 260–268.

Mesgarani, N., and Chang, E.F. (2012). Selective cortical representation of attended speaker in multi-talker speech perception. Nature *485*, 233–236.

Myburgh, H.C., Jose, S., Swanepoel, D.W., and Laurent, C. (2018). Towards low cost automated smartphone- and cloud-based otitis media diagnosis. Biomed. Signal Process. Control *39*, 34–52.

National Academies of Sciences, E. (2016). Hearing Health Care for Adults: Priorities for Improving Access and Affordability.

National Association of the Deaf (2020). Community and Culture – Frequently Asked Questions.

Neijenhuis, K., Campbell, N.G., Cromb, M., Luinge, M.R., Moore, D.R., Rosen, S., and de Wit, E. (2019). An Evidence-Based Perspective on "Misconceptions" Regarding Pediatric Auditory Processing Disorder. Front. Neurol. *10*.

NIHR (2020). Hearing health – Health Informatics Collaborative.

Orji, A., Kamenov, K., Dirac, M., Davis, A., Chadha, S., and Vos, T. (2020). Global and regional needs, unmet needs and access to hearing aids. Int. J. Audiol. *59*, 166–172.

O'Sullivan, J.A., Power, A.J., Mesgarani, N., Rajaram, S., Foxe, J.J., Shinn-Cunningham, B.G., Slaney, M., Shamma, S.A., and Lalor, E.C. (2015). Attentional Selection in a Cocktail Party Environment Can Be Decoded from Single-Trial EEG. Cereb. Cortex N. Y. N 1991 *25*, 1697–1706.

Parthasarathy, A., Hancock, K.E., Bennett, K., DeGruttola, V., and Polley, D.B. (2020). Bottom-up and top-down neural signatures of disordered multi-talker speech perception in adults with normal hearing. ELife *9*, e51419.

Pichichero, M.E., and Poole, M.D. (2001). Assessing diagnostic accuracy and tympanocentesis skills in the management of otitis media. Arch. Pediatr. Adolesc. Med. *155*, 1137–1142.

Picton, T. (2013). Hearing in time: evoked potential studies of temporal processing. Ear Hear. *34*, 385–401.

President's Council of Advisors on Science and Technology, (PCAST) (2016). Aging America & Hearing Loss: Imperative of Improved Hearing Technologies.

Reinfeldt, S., Håkansson, B., Taghavi, H., and Eeg-Olofsson, M. (2015). New developments in bone-conduction hearing implants: a review. Med. Devices Auckl. NZ *8*, 79–93.

Roberts, J.E., Sanyal, M.A., Burchinal, M.R., Collier, A.M., Ramey, C.T., and Henderson, F.W. (1986). Otitis Media in Early Childhood and Its Relationship to Later Verbal and Academic Performance. Pediatrics *78*, 423–430.

Robinson, P.N. (2012). Deep phenotyping for precision medicine. Hum. Mutat. *33*, 777–780.

Rovers, M.M., Schilder, A.G., Zielhuis, G.A., and Rosenfeld, R.M. (2004). Otitis media. The Lancet *363*, 465–473.

Sandström, J., Swanepoel, D., Laurent, C., Umefjord, G., and Lundberg, T. (2020). Accuracy and Reliability of Smartphone Self-Test Audiometry in Community Clinics in Low Income Settings: A Comparative Study. Ann. Otol. Rhinol. Laryngol. *129*, 578–584.

Schilder, A.G.M., Chonmaitree, T., Cripps, A.W., Rosenfeld, R.M., Casselbrant, M.L., Haggard, M.P., and Venekamp, R.P. (2016). Otitis media. Nat. Rev. Dis. Primer *2*, 1–18.

Schrimpf, M., Blank, I., Tuckute, G., Kauf, C., Hosseini, E.A., Kanwisher, N., Tenenbaum, J., and Fedorenko, E. (2020). Artificial Neural Networks Accurately Predict Language Processing in the Brain. BioRxiv 2020.06.26.174482.

Shew, M., New, J., Wichova, H., Koestler, D.C., and Staecker, H. (2019). Using Machine Learning to Predict Sensorineural Hearing Loss Based on Perilymph Micro RNA Expression Profile. Sci. Rep. *9*, 3393.

Shi, L.-F., Carney, L.H., and Doherty, K.A. (2006). Correction of the peripheral spatiotemporal response pattern: a potential new signal-processing strategy. J. Speech Lang. Hear. Res. JSLHR *49*, 848–855.

Shore, S.E., and Wu, C. (2019). Mechanisms of Noise-Induced Tinnitus: Insights from Cellular Studies. Neuron *103*, 8–20.

Sing Registry (2020). The Genetic Sensorineural Hearing Loss Registry.

Swanepoel, D.W., Clark, J.L., Koekemoer, D., Hall, J.W., Krumm, M., Ferrari, D.V., McPherson, B., Olusanya, B.O., Mars, M., Russo, I., et al. (2010). Telehealth in audiology: the need and potential to reach underserved communities. Int. J. Audiol. *49*, 195–202.

Tomblin, J.B., Harrison, M., Ambrose, S.E., Walker, E.A., Oleson, J.J., and Moeller, M.P. (2015). Language Outcomes in Young Children with Mild to Severe Hearing Loss. Ear Hear. *36 Suppl 1*, 76S-91S.

Vaerenberg, B., Smits, C., De Ceulaer, G., Zir, E., Harman, S., Jaspers, N., Tam, Y., Dillon, M., Wesarg, T., Martin-Bonniot, D., et al. (2014). Cochlear implant programming: a global survey on the state of the art. ScientificWorldJournal *2014*, 501738.

Viscaino, M., Maass, J.C., Delano, P.H., Torrente, M., Stott, C., and Cheein, F.A. (2020). Computer-aided diagnosis of external and middle ear conditions: A machine learning approach. PLoS One *15*, e0229226.

Wang, D., and Chen, J. (2018). Supervised Speech Separation Based on Deep Learning: An Overview. IEEEACM Trans. Audio Speech Lang. Process. *26*, 1702–1726.

Warren, E., and Grassley, C. (2017). Over-the-Counter Hearing Aids: The Path Forward. JAMA Intern. Med. *177*, 609–610.

WHO (2013). Multi-country assessment of national capacity to provide hearing care (World Health Organization).

Wilson, B.S. (2015). Getting a decent (but sparse) signal to the brain for users of cochlear implants. Hear. Res. *322*, 24–38.

Wilson, B.S., and Dorman, M.F. (2008). Cochlear implants: A remarkable past and a brilliant future. Hear. Res. *242*, 3–21.

Wilson, B.S., Tucci, D.L., Merson, M.H., and O'Donoghue, G.M. (2017). Global hearing health care: new findings and perspectives. The Lancet *390*, 2503–2515.

World Health Organization (2017). Global costs of unaddressed hearing loss and cost-effectiveness of interventions: a WHO report (Geneva, Switzerland).

Wu, J.-Y., Yu, C., Fu, S.-W., Liu, C.-T., Chien, S.-Y., and Tsao, Y. (2019). Increasing Compactness of Deep Learning Based Speech Enhancement Models With Parameter Pruning and Quantization Techniques. IEEE Signal Process. Lett. *26*, 1887–1891.

Yung, M., Gjuric, M., Haeusler, R., Van de Heyning, P.H., Martin, C., Swan, I.R.C., Tange, R.A., Huy, P.T.B., and European Otology Database Project Group (2005). An international otology database. Otol Neurotol *26*, 1087–1092.

Zeng, F.-G. (2017). Challenges in Improving Cochlear Implant Performance and Accessibility. IEEE Trans. Biomed. Eng. *64*, 1662–1664.

Zeng, F.-G., Rebscher, S., Harrison, W.V., Sun, X., and Feng, H. (2008). Cochlear Implants: System Design, Integration and Evaluation. IEEE Rev. Biomed. Eng. *1*, 115–142.

Zeng, F.-G., Tran, P., Richardson, M., Sun, S., and Xu, Y. (2019). Human Sensation of Transcranial Electric Stimulation. Sci. Rep. *9*, 15247.

Zhao, E.E., Dornhoffer, J.R., Loftus, C., Nguyen, S.A., Meyer, T.A., Dubno, J.R., and McRackan, T.R. (2020). Association of Patient-Related Factors With Adult Cochlear Implant Speech Recognition Outcomes: A Meta-analysis. JAMA Otolaryngol Head Neck Surg *146*, 613–620.

Zhao, Y., Li, J., Zhang, M., Lu, Y., Xie, H., Tian, Y., and Qiu, W. (2019). Machine Learning Models for the Hearing Impairment Prediction in Workers Exposed to Complex Industrial Noise: A Pilot Study. Ear Hear. *40*, 690–699.

Zhou, Z., Chen, K., Li, X., Zhang, S., Wu, Y., Zhou, Y., Meng, K., Sun, C., He, Q., Fan, W., et al. (2020). Sign-to-speech translation using machine-learning-assisted stretchable sensor arrays. Nat. Electron. 1–8.