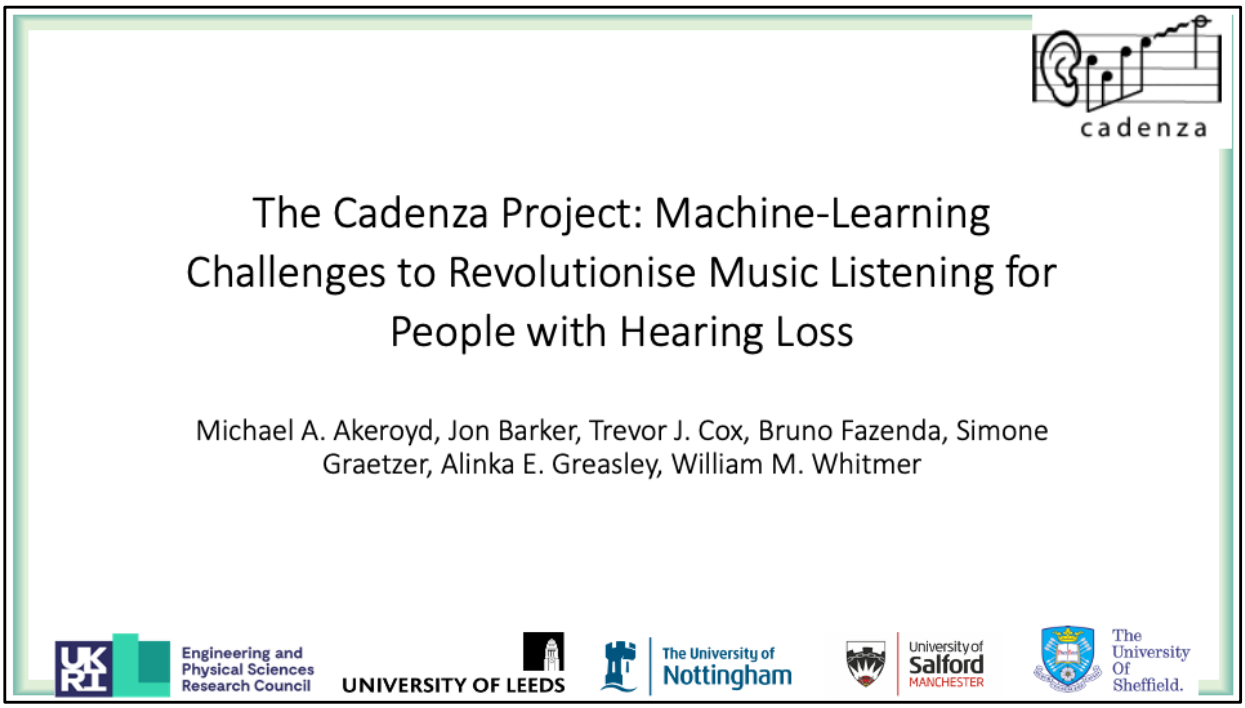

We've now got a session on the new Cadenza project, which is all about trying to improve how we process and mix music so it works better for people with a hearing loss.

The methodology we're going to use is machine learning challenges. I'll explain more about that a little later on.

I'm Trevor Cox. And with me also is Alinka Greasley. The project involves four different universities, a range of project partners from BBC R&D to RNID, and the project is funded by EPSRC. The project starts in about a month's time.

Task for you

cadenza

- Question:
  - What issues do people with hearing loss have with listening to music?
  - With or without hearing aids
- You'll be in break-out groups
- Please capture your thoughts on the padlet
  - https://padlet.com/tjcox1/9h7ixitl5j7gfdyn
- You have 10 minutes!

Before outlining the project, I've got a task for you. And the question I want you to answer is what issues do people with hearing loss have with listening to music? You might have direct experience of this yourself, or you might know about the experiences of others. That could be with or without hearing aids. And I'm going to put you into breakout groups through the platform we have here. But as you chat, please catch your thoughts on the padlet using this link or the QR code. It is important to capture your thoughts because it will help shape what we do in the project. You have 10 minutes

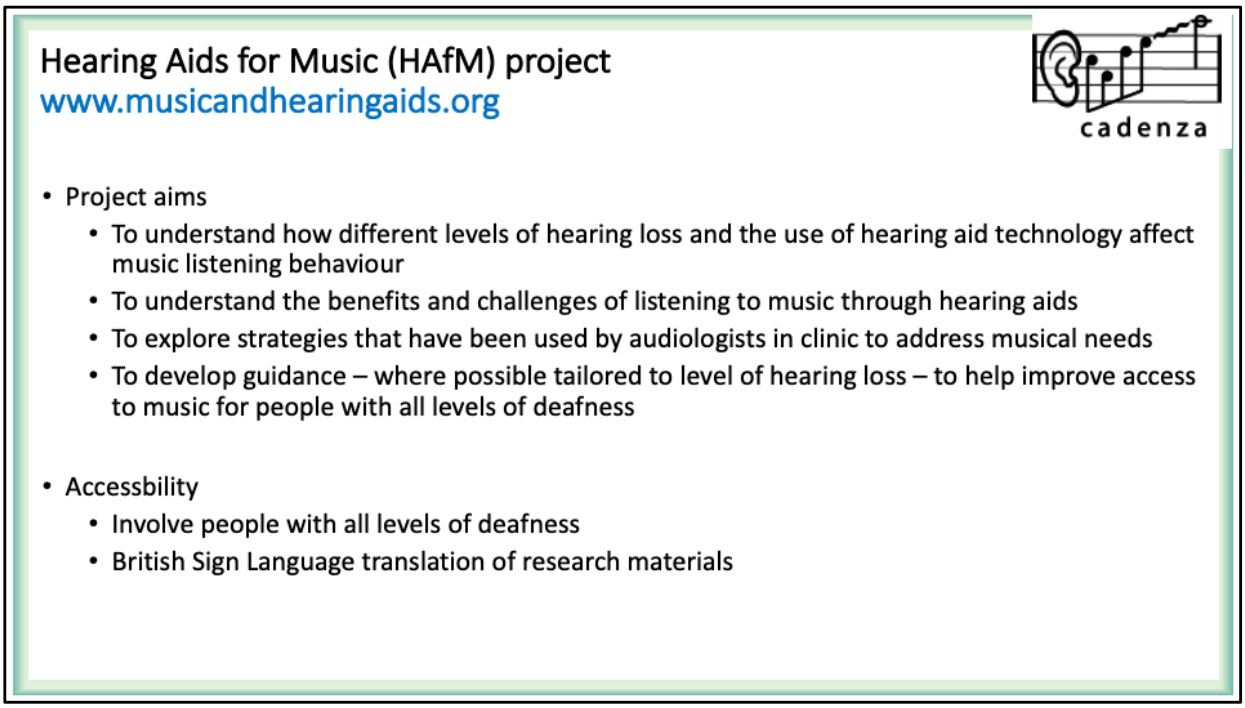

I'll now pass over to Alinka Greasley from Leeds University, who is going to talk a bit about the project on Hearing Aids for Music.
Thanks, Trevor.
So I'm going to say a few words about the Hearing Aids for Music project which started in 2015. This project was a collaboration between the University of Leeds and Sheffield Teaching Hospitals which explored how hearing loss and the use of hearing aid technology affects music listening behaviour. At that time there had been quite a bit of published literature on music perception with a cochlear implant, but very little on hearing aids and this

project sought to provide new data on the benefits and challenges of listening to music through hearing aids. We engaged hearing aid users and audiologists in the research with an explicit aim of developing guidance to help improve music listening experiences and access to music.

Accessibility was at the heart of the project, and we ensured that people of all levels of deafness could participate, by providing sign language interpretation where needed, including translating our national survey questions into British Sign Language.

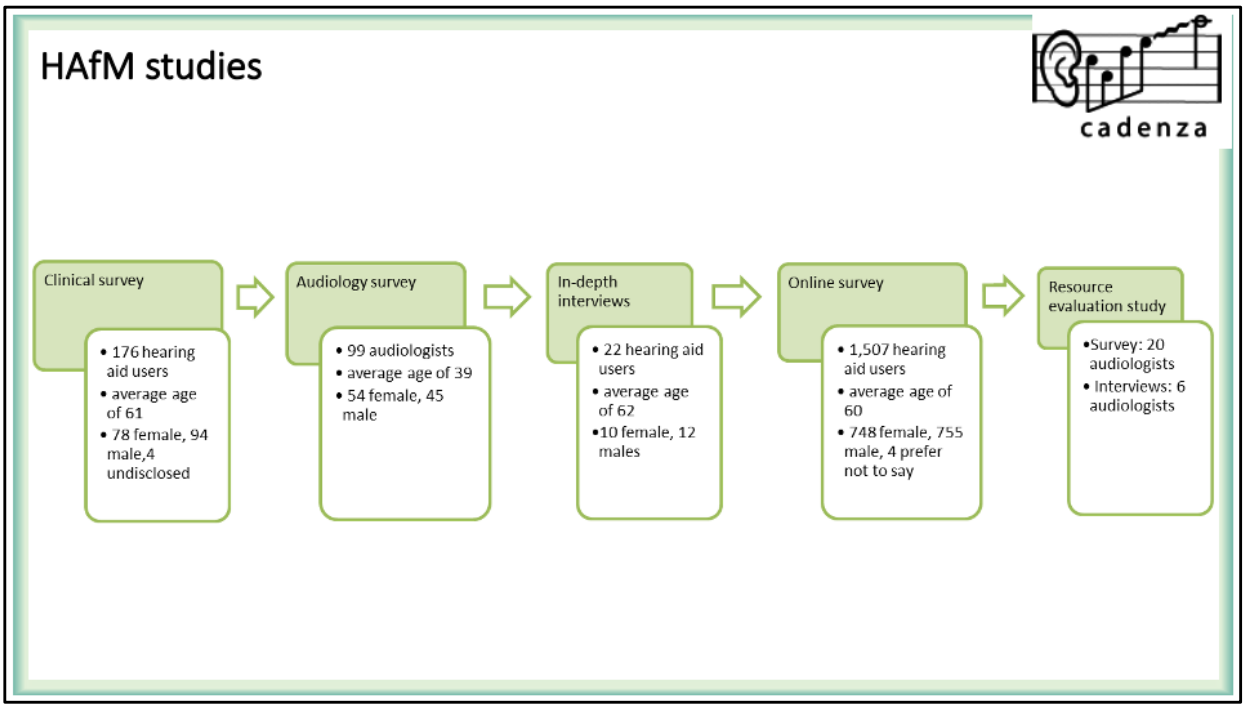

We carried out a series of five studies. This included a clinical survey asking patients whether they had experienced any problems with music listening and whether they felt this affected their quality of life, whether they had discussed music with their audiologist and the outcomes of these discussions. A survey among audiologists asking them to reflect on their training level and background, their experiences of discussing music listening issues, their experiences of optimising hearing aids for music listening (and their perceived confidence and ability to do so). An interview study which explored the experiences of 22 hearing aid users in greater detail. And a national survey which explored patterns in music listening experiences according to levels of hearing loss among over 1,500 hearing aid users. We used the results to develop resources which we subsequently asked audiologists to reflect on in an evaluation study.

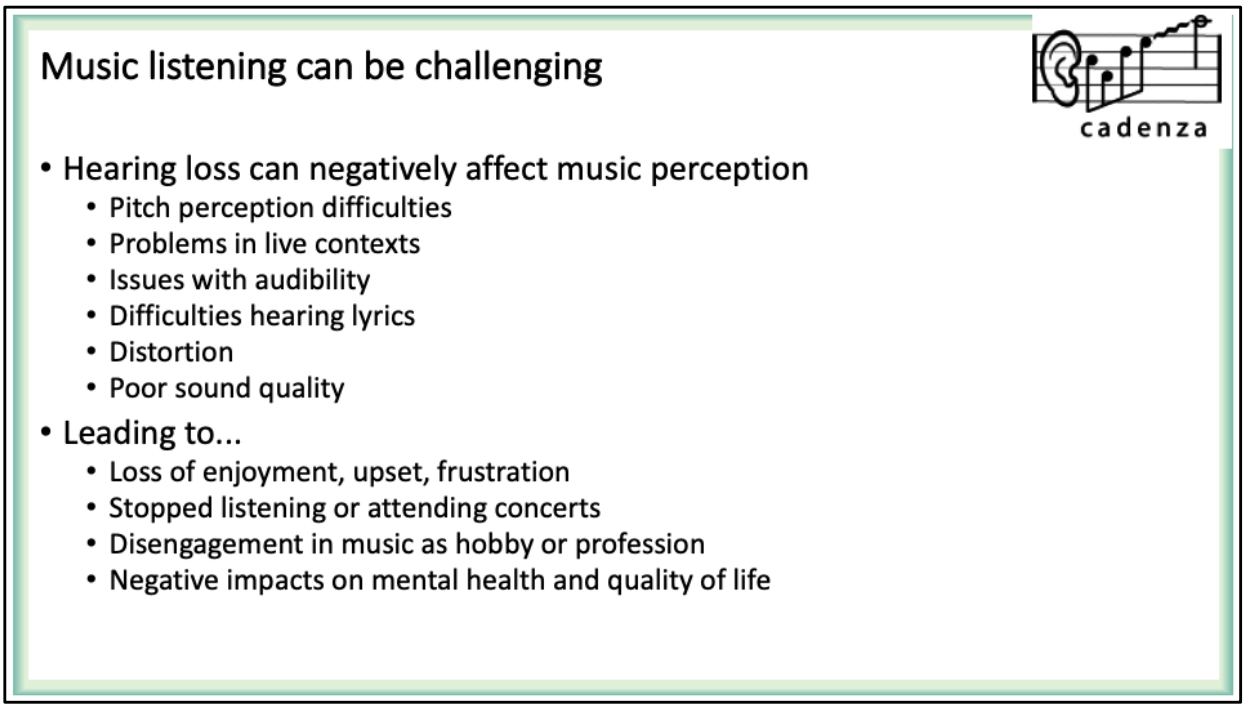

So what did we find? Listening to music can be challenging for people with hearing loss. Our participants reported a range of problems including problems with pitch perception making it harder to follow melodies or hear out harmonies, difficulties hearing lyrics, problems with audibility, with some reporting that quiet passages were too quiet where loud passages were distorted, and issues with poor sound quality. In some cases, this led to a loss of enjoyment of music or stopping listening or attending concerts. In other cases, people were disengaging with music altogether, with associated negative effects on their health and wellbeing.

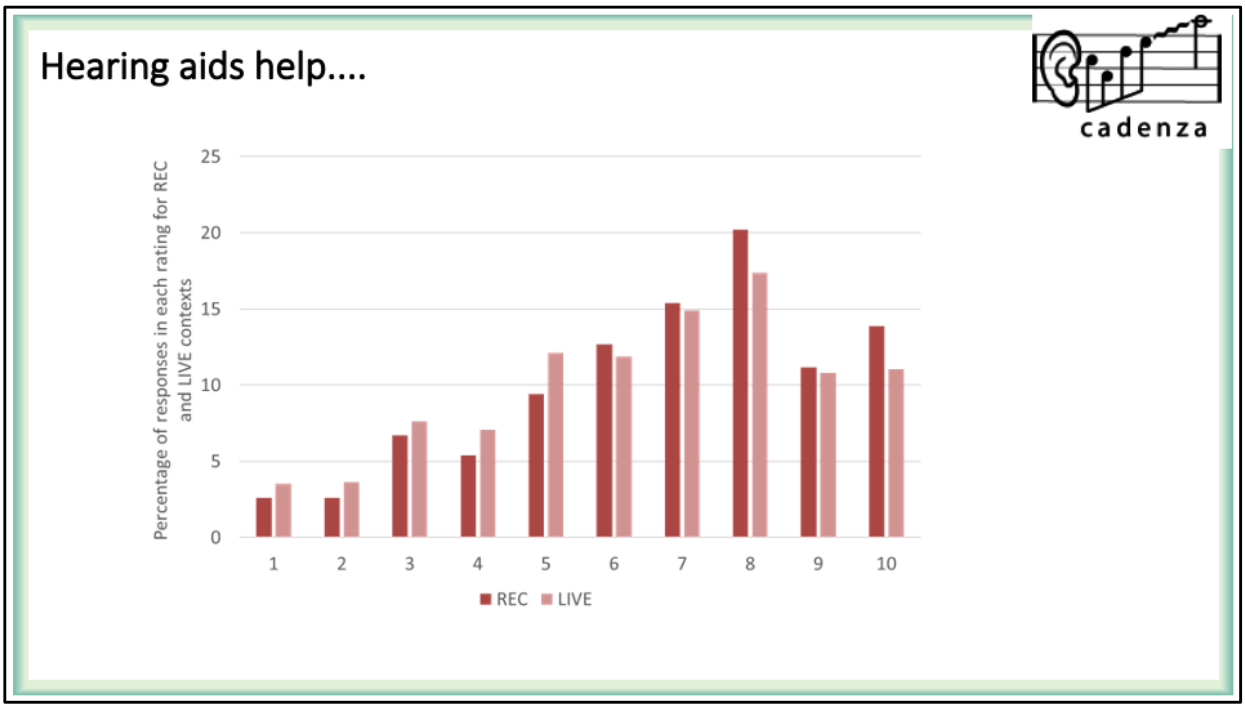

Hearing aids were reported to help overall, as we can see here from this graph of the overall ratings of helpfulness of hearing aids for listening to music in live and recorded contexts. The modal response was 8 out of 10 for rated helpfulness. However, the average response for both contexts was 6 out of 10. Clearly there is room for improvement in music processing within hearing aids.
As one respondent put it *"[music] seems to be an unsolvable problem in otherwise excellent hearing aids"*

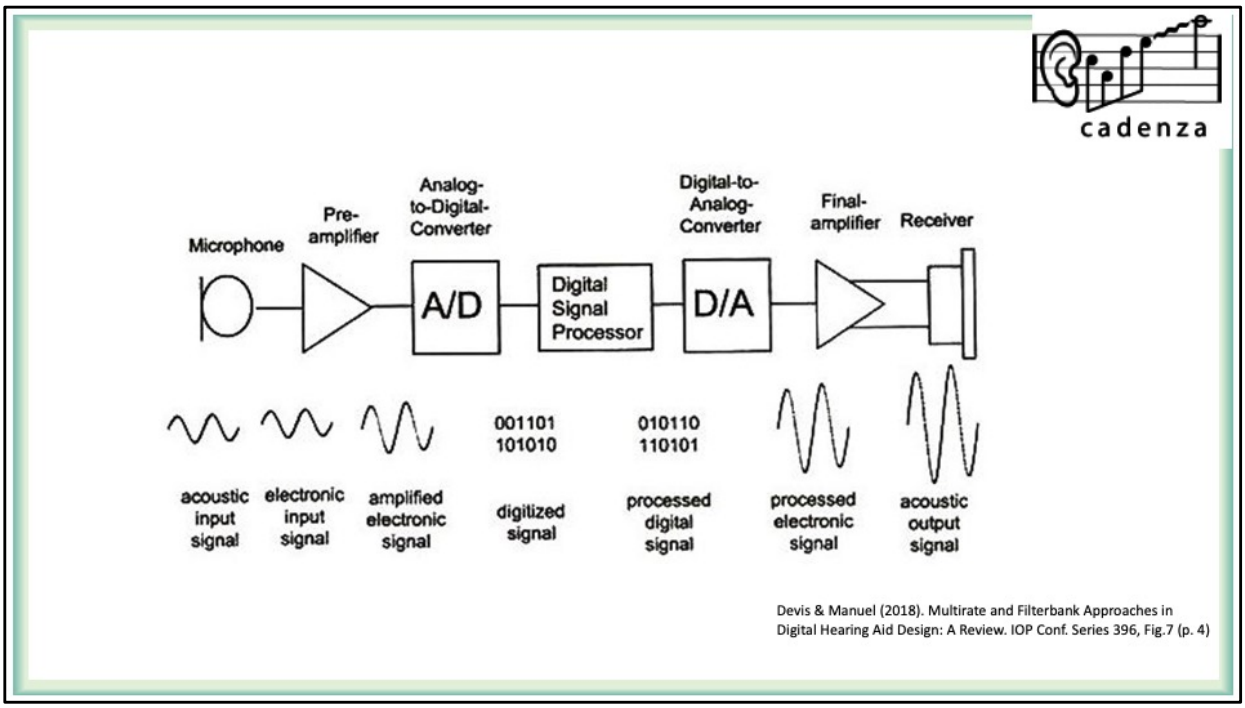

So why are hearing aids poor for music?
It's worth explaining briefly how digital hearing aids work. Sound enters the microphone as an acoustic input. The analog to digital (A/D) convertor transforms the continuous microphone signal into a series of numbers. Digital circuitry processes the input series of numbers to create the output signal, as another series of numbers. The D/A convertor transforms the output series of numbers back into a continuous signal for the receiver to emit. (NB: Older hearing aids were analog and didn't have this digital circuitry.) There's a lot of clever stuff that happens in the middle box, the Digital Signal Processor. This is where a person's hearing loss level is taken into account (different gains are applied at different frequencies), and various DSP strategies are employed. Over time, hearing aid manufacturers have adapted DSP strategies to improve speech perception, since a primary function of hearing aids is to enable communication, and with good effect. Compression makes quiet sounds audible, but loud sounds comfortable for the listener. Noise reduction suppresses noise that maintains more or less constant sound pressure level over time (e.g. noise in a car or train), making it easier to have conversations. Feedback reduction suppresses unwanted acoustic signal that feeds back from the receiver to the microphone.

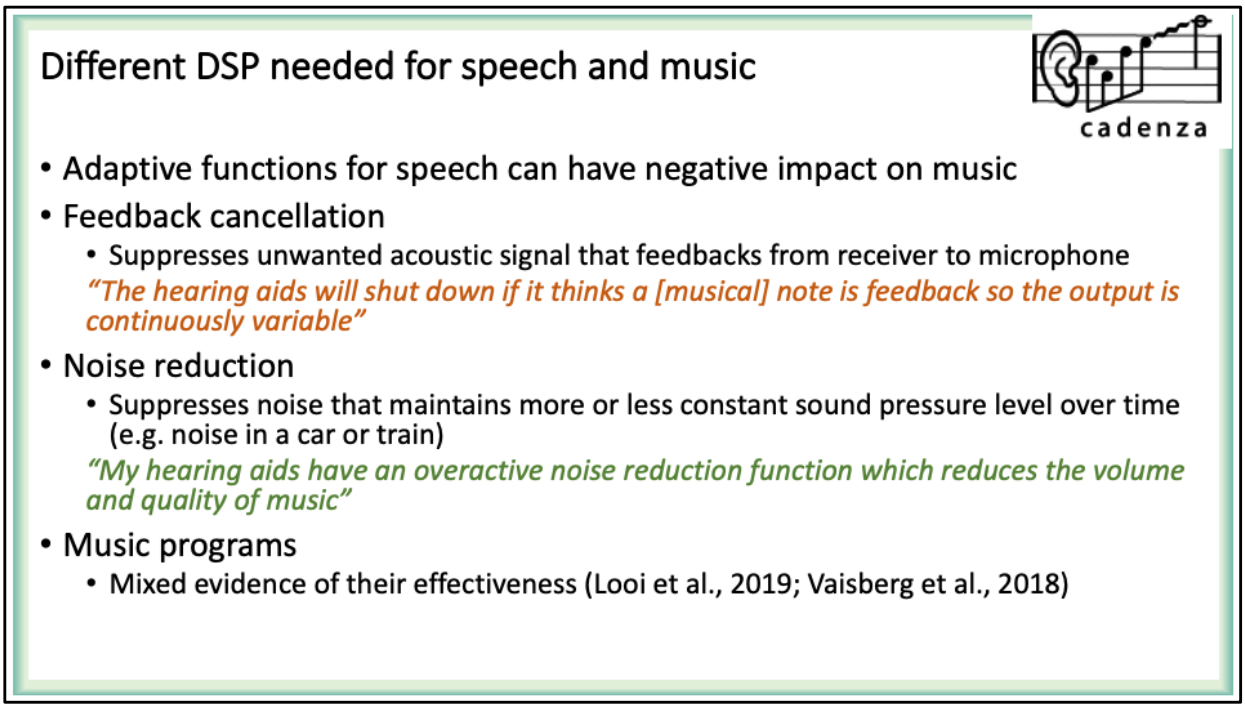

However, these DSP strategies for speech perception can have a negative impact on music perception. Feedback reduction can mistakenly analyse pure tones as feedback and suppress. Similarly, hearing aids can misidentify music as noise and suppress it. As one man summarised, *"all these extra things that have been added to improve speech recognition…have nothing to do with music and have actually got in the way"(M,82 yrs.)* It is perhaps no wonder that audiologists emphasise that turning off adaptive functions for speech is necessary for good musical outcomes. However it's not just the adaptive functions. Compression strategies can also cause problems, particularly when music goes from very quiet to very loud quickly (as might be expected in a symphonic work). Manufacturers have tried to address the issue of poor sound quality for music, and have implemented music programs, but increasing evidence shows that these are not effective. It is necessary therefore to arrive at some novel signal processing solutions for music for people with hearing loss for use in hearing aids.

Feedback

- Question:
  - What issues do people with hearing loss have with
  - With or without hearing aids
- https://padlet.com/tjcox1/9h7ixitl5j7gfdyn

We will feedback what was on the padlet here.

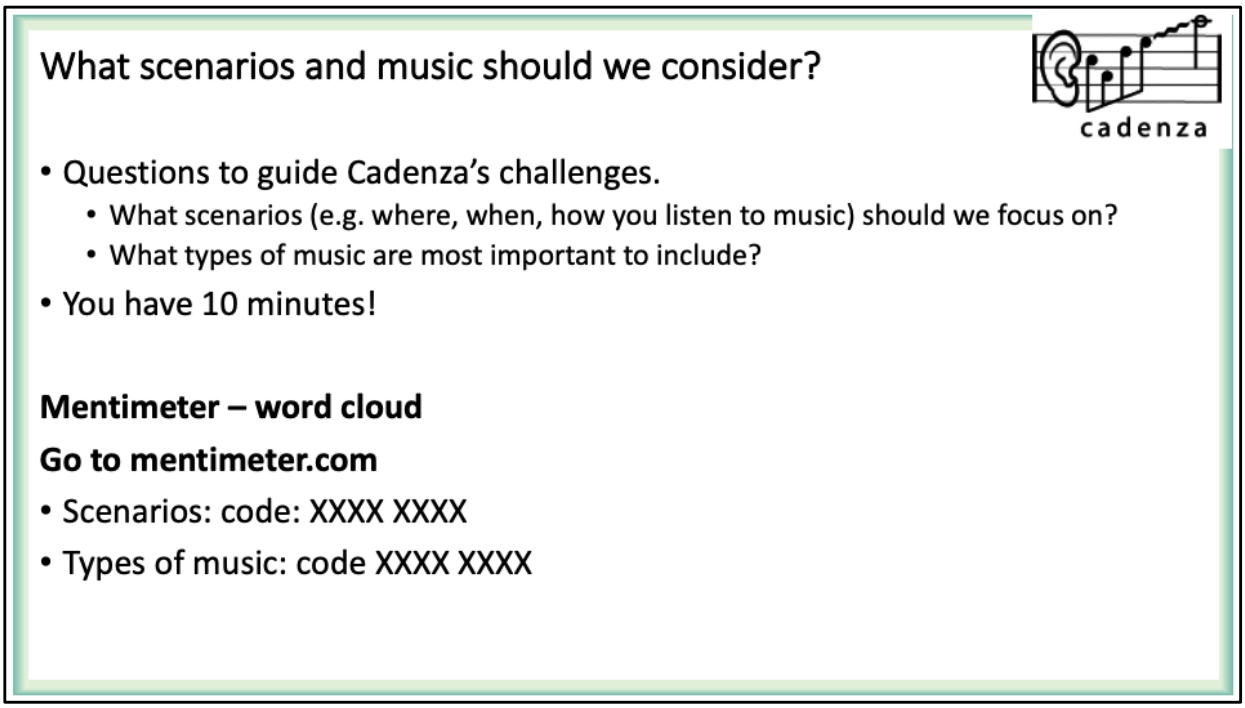

One of the things we have to decide in Cadenza is what listening scenarios and music to consider for remixing or reengineering so it sounds better for those with a hearing loss.. So here is a chance for you to give us your input on this. What should we focus on? There are lots of places where we consume music: gigs, at the home, over headphones, in the car, in the kitchen? So where, when and how you listen to music? Which ones are the most important to focus on? And also what types of music are most important to include? And this could be genre of that, but it could be also be other aspects: is it solo instrument? Is it about a band or orchestra? With or

without lyrics Think as broad as you want to. And what we've got a couple of word clouds for you to fill up. Please follow the links.

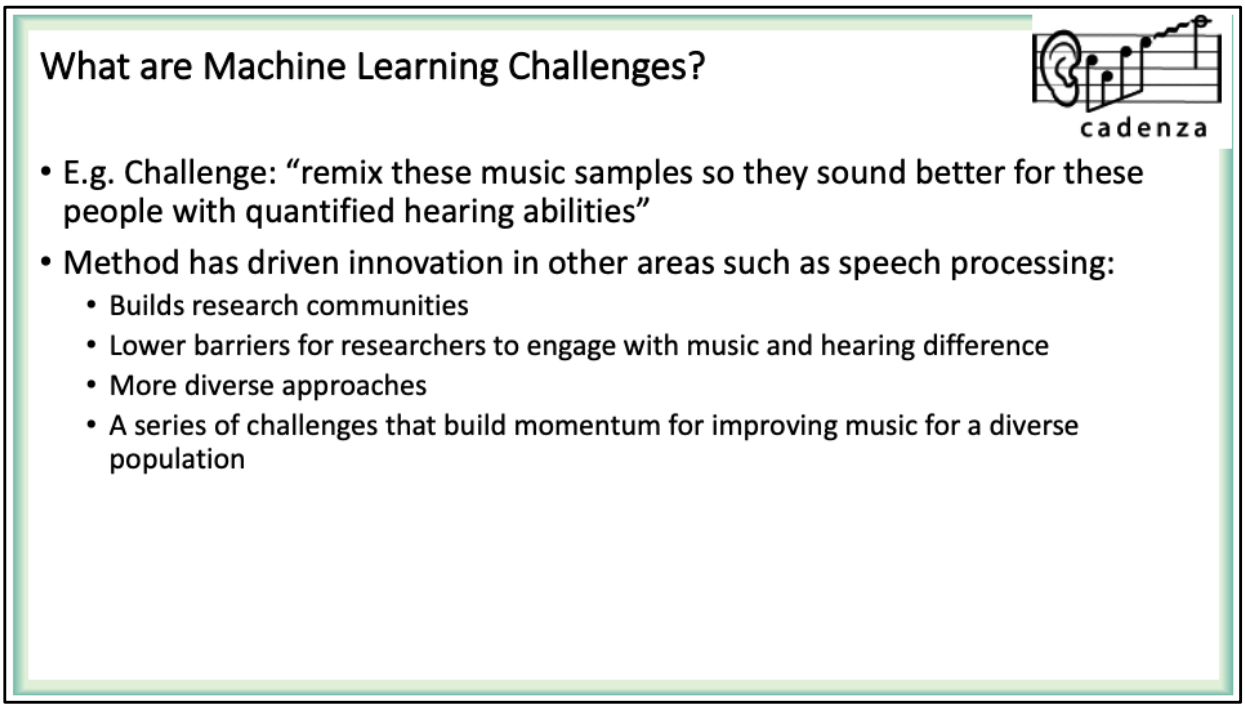

Let's get back to the project. I guess for many people in the room, this term machine learning challenges needs explaining. So what are these? Well, essentially what we do is we challenge teams from around the world to try and improve music. We ask them to remix and/or process a set of music samples so they sound better for those for people with quantified hearing abilities. We give them the music samples to work with and we give them various tools to help them do it. So this could be being done on a hearing aid, but it could also be being done on a consumer device like a Bluetooth loudspeaker, or it could be being done back in the studio.

What is a challenge methodology? In normal research projects, you have a researcher sat in your university trying to come up with new algorithms to process music. In a challenge project, what you do is set competitions and people around the world produce new algorithms and compete.

Why this approach? Well, first of all, it builds research communities. If I just work on it in Salford with my lone researcher, then that's a community of just Salford. By actually getting lots of people competing, we lower barriers for researchers to engage with music and people with hearing difference. And that is vital because much machine learning ignores diverse hearing. And that's one of the big aims of Cadenza, as a project, is to expand the number of people who consider hearing loss when they're thinking about processing music.

Another advantage is we also get more diverse approaches, if 20 teams enter, you get 20 algorithms. That's more than we can generate within one university in a project. And what also happens is we use a series of challenges, Cadenza is four and a half years long to fit three challenges in, and that builds momentum and builds that community. So the challenges and research can continue run by the community after the funding ends.

So that's why we're using this kind of technique, and it's been shown to work really well in other areas. If I take, for example, the Chime challenges, which are about automatic speech recognition, or turning audio into text, those have driven significant advances. And there's lots of areas of signal processing doing similar, challenges are now a fairly standard way of advancing a field.

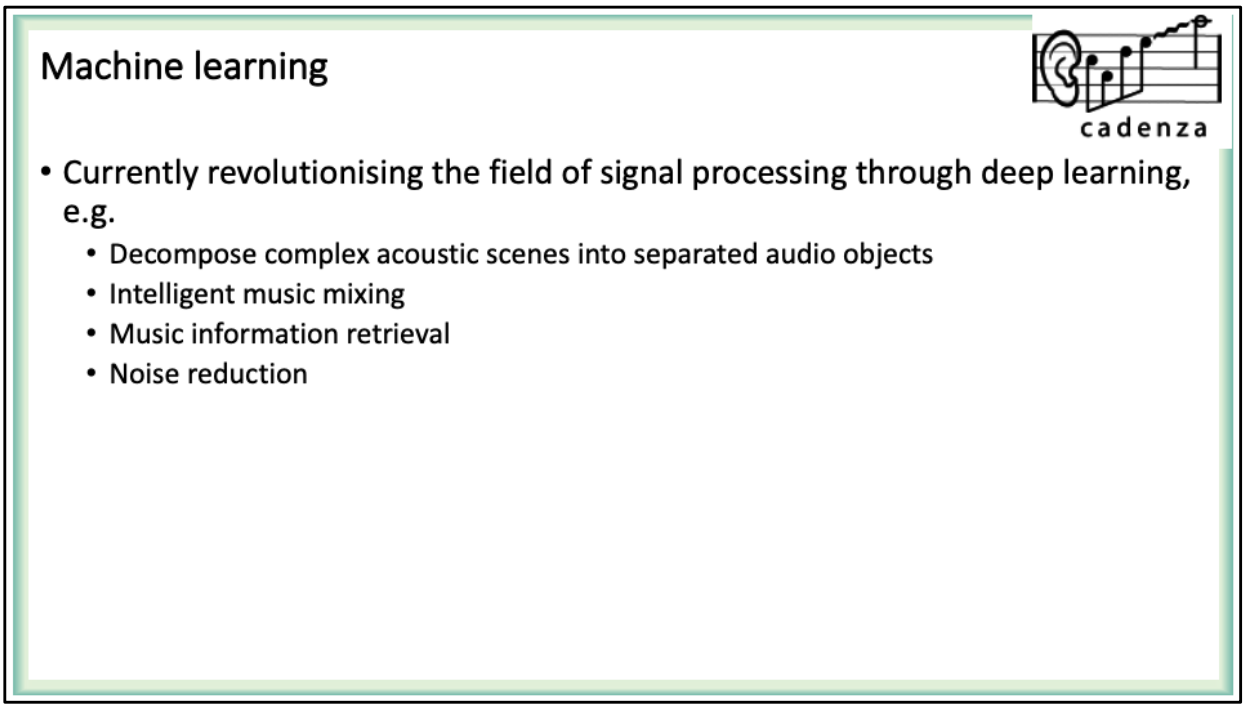

Why machine learning? Well, because we think there's an opportunity to revolutionize what we do with music, because we're seeing radical changes in other areas but mostly not related to people who might have a hearing difference. That signal processing would be done through deep learning, which is the zeitgeist at the moment.

To take some examples, we see lots of techniques to do things like take complex acoustic scenes, it could be an outdoor scene and decompose it into separate audio objects. There's a car, there's the sound of a bird, that

kind of thing. Well, that's the sort of thing we do when we listen to music. We have intelligent music mixing already. We have music, information retrieval, giving some music, extract information from it using machine learning techniques. And we have a whole bunch of noise reduction techniques. And why might not be relevant? Well, we don't always listen to music in quiet conditions so that noise reduction might be useful for listening in the car, for example. We're seeing a revolution in machine learning. And that's why we think it's the time is right to run a series of challenges so people with hearing difference aren’t overlooked.

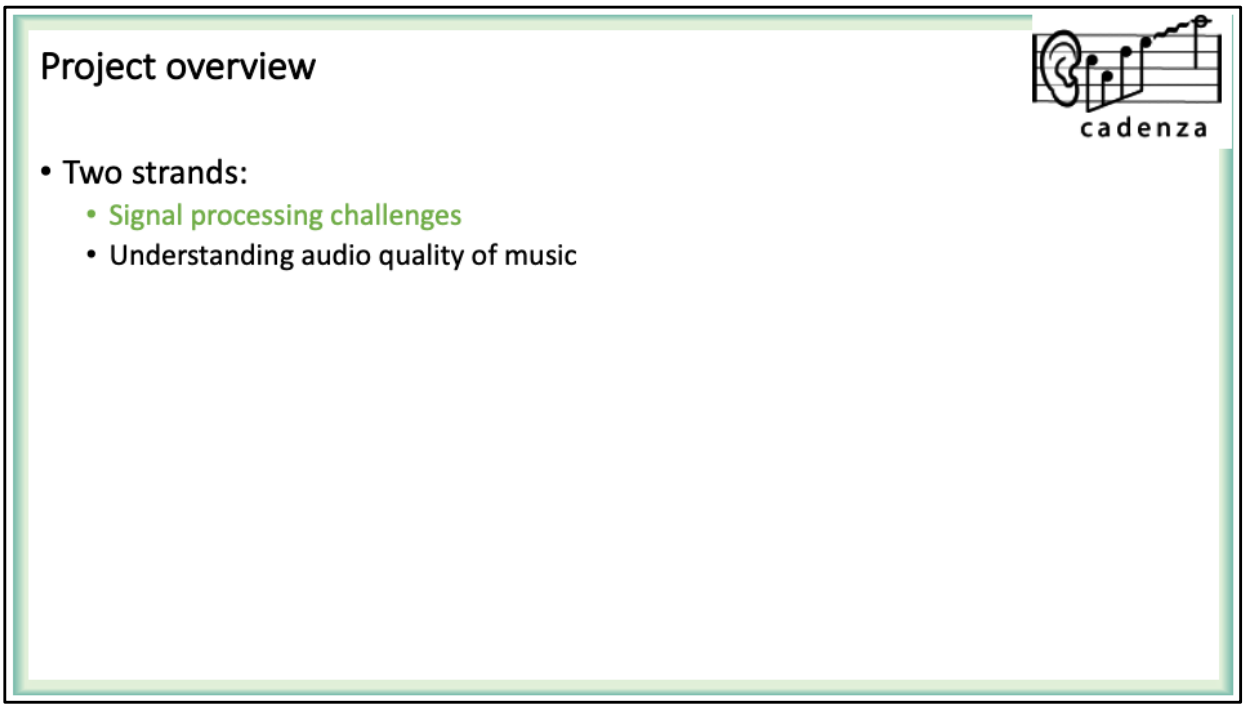

There's actually two parts to the project. There's the actual challenges, which is what I've been kind of talking about. But there's also understanding what we mean by better quality music. And I'll come back to that towards the end because it's a very important strand of our research.

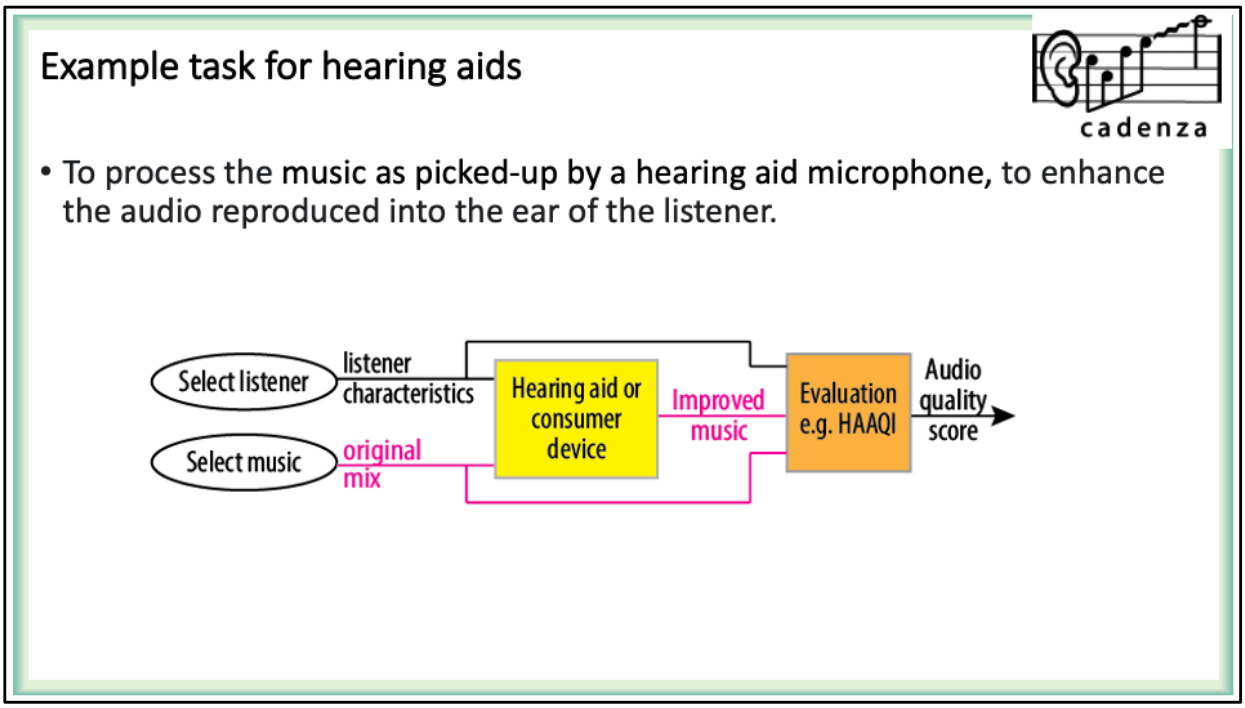

Let me take a possible scenario of improving the processing of music on a hearing aid. The music is picked up by the hearing aid a microphone. We want to enhance the audio reproduced into the ear of the listener. This could be a very typical behind the ear hearing aid. The diagram is a very simple sort of signal processing schematic.

We have the original music. We select a listener. Everything has to be bespoke to the listeners hearing characteristics.

And then we have the music being processed by this hearing aid or consumer device in the yellow box. And this is actually what entrants are going to work to try and improve. So there should be improved music and that's actually what will be piped into the listener.

Now we've got to evaluate whether it works or not. And what we have here is actually a metric you may not be familiar with HAAQI, the hearing aid audio quality index. I'll talk a bit more about this in a moment. But basically, this is a way of scoring the quality of the audio. So the challenge is basically make a better yellow box.

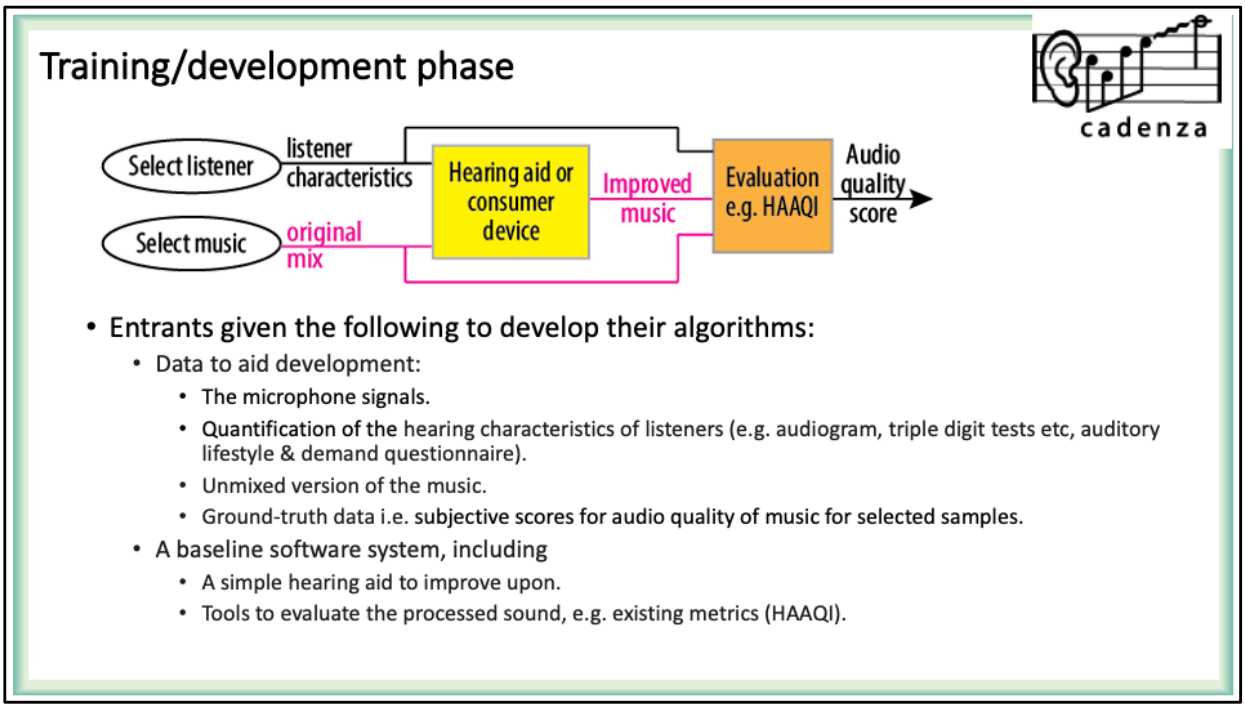

There's three phases to a challenge. Currently we're in the planning and preparation stage, which actually is a lot of work!
What is on this slide is the second phase when it is open for competitors. This a training and development phase where people are making the algorithms and we have to give them lots of information.

So we've got data to develop these algorithms. First of all, the microphone signals on the hearing aid, in this case the quantification of the characteristics listener. At the most simple scale this could be audiograms, but

there is other stuff, like super threshold testing, like triple digit tests, there's questionnaires that you give to people who use hearing aids.

Then we give him other things like the unmixed version of music. So what was it originally and maybe also given ground truth data? Some scores where people have listened to some mixed music and scored it for how good it is.

And the other important thing, alongside all this data to develop things, is we give them a baseline system. That's a piece of software which basically solves the task in a simple way. We make it simple because people are trying to beat it. So if we make a very good system, people get put off. It also gives them a software framework to build upon so they don't have to change everything, and just work on what they have expertise on.

The baseline will be a simple hearing aid to improve upon, maybe openMHA. And some simple tools to process the sound, for example. It could be the objective metric HAAQI.

People are given this data and software, they have many months to play around and try and come up with better algorithms.

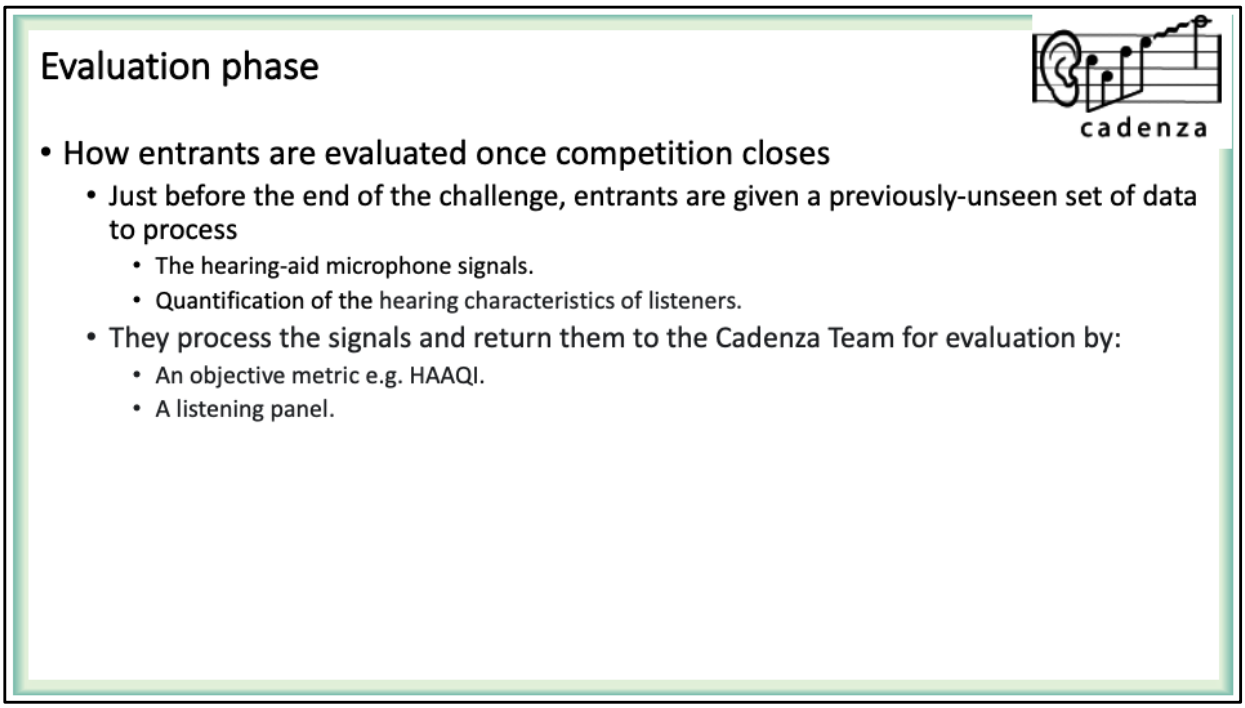

Then the challenge closes and then we have the evaluation phase.
Now what happens just before that, maybe a month before it closes, is entrants get a new set of data to process. This is previously unseen data, but is similar to what they have seen in the training and development. This would be a set of hearing aid microphone signals and the quantification of a set of listeners. But different pieces of music and different listeners. Entrants send us the processed audio, and then the Cadenza team evaluate it.

We do that by two methods. The first one would be an objective metric e.g. in round one HAAQI. And the second method would be a listening panel of people actually listening and scoring these things

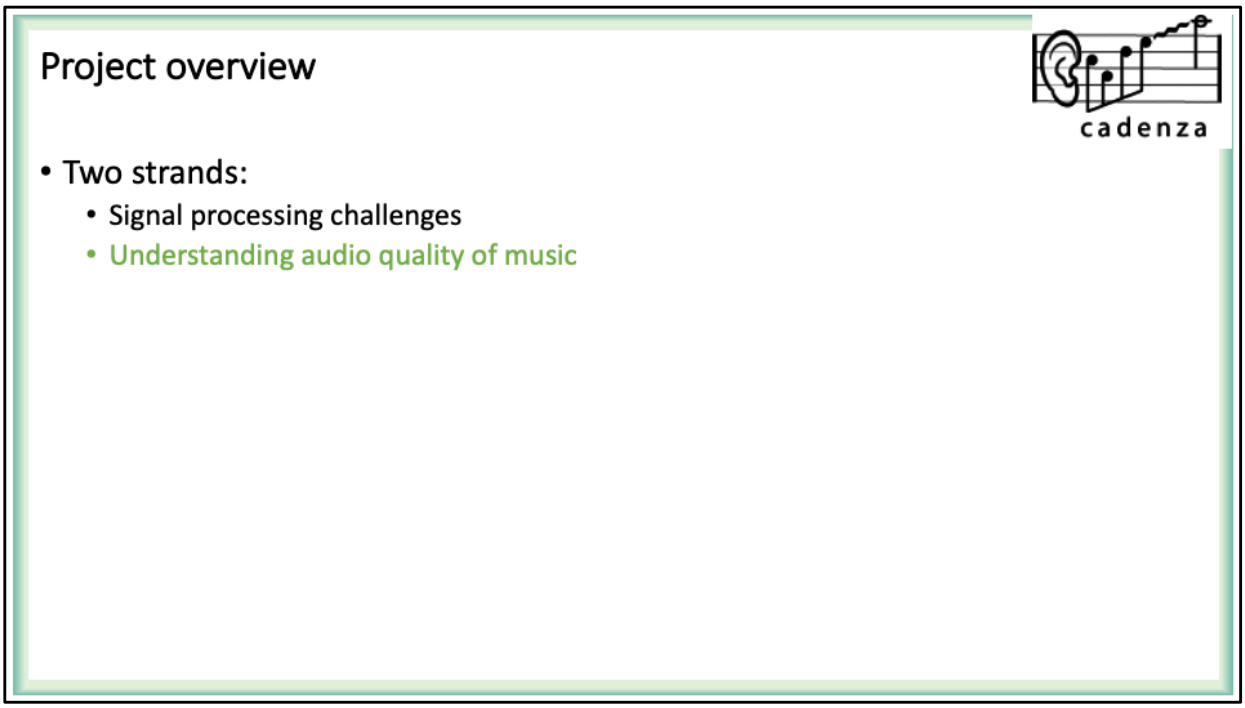

That brings me on to the second strand of the project, which is about this understanding the audio quality of the music.

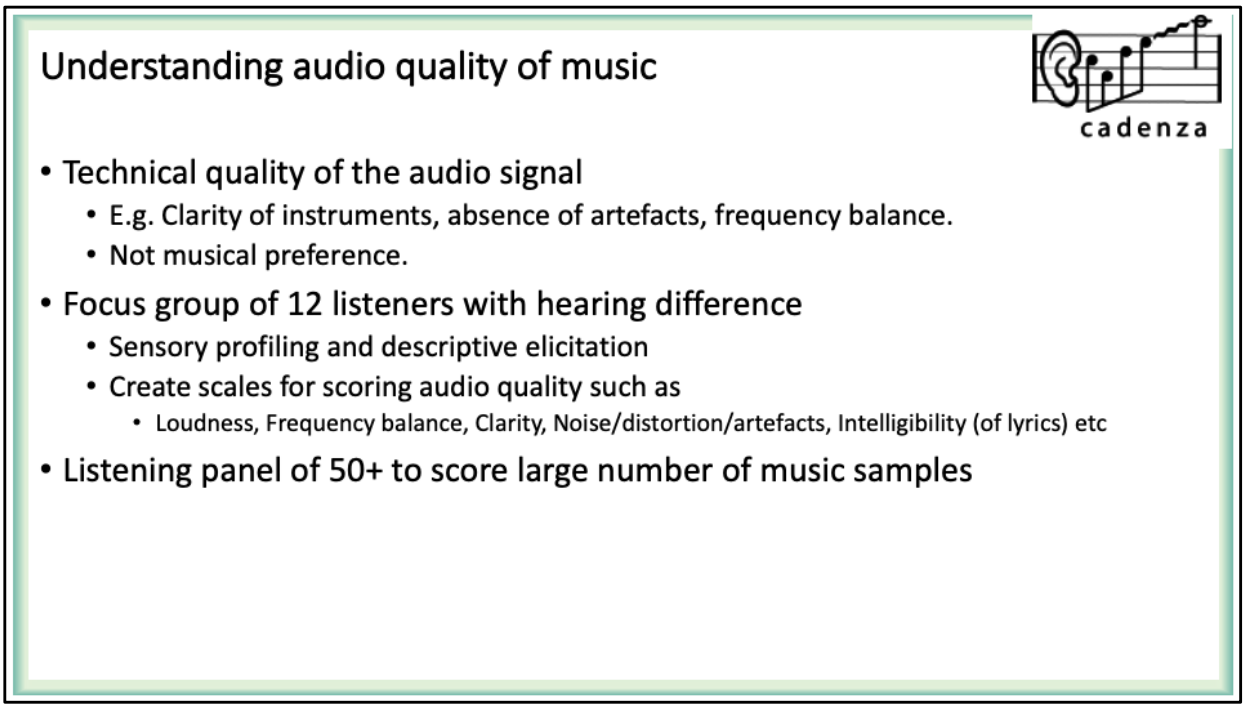

So what we're talking about here is about the technical quality of the audio signal. We're not talking about musical preference here. It's quite hard to remove musical preference, but we're going to try and minimize the effect of this. What we're talking about is things like, can you pick out new instruments out of the mix? Are there artifacts? Hearing aids often add artifacts into the signals. Is the frequency balance right? I mean, is it too muffled or is there too much bass, that kind of thing? Can you pick out the lyrics?

It's those sort of features that we're looking for rather than anything to do with musical preference. Now we're going to start off by using this metric HAAQI because it exists and we can get going with our challenges, but we know it has weaknesses. One of the weaknesses is it scores music with one number, which is very convenient, but we know that music listening is much more complicated and it's multidimensional.

So we need to develop a better way of scoring it. And so alongside running our first challenge, we're going to be developing a better metric to use in future challenges. To do this we'll use a standard process that actually started in food science to do with

taste, but it has been applied to audio by various people. First of all, you start with a focus group who listen to examples of music processed through hearing aids and they discuss what can they hear, what's good and what's bad, and they try and come up with a consensus of the key features which differentiate these different samples.

That could be lots of different things. Take absence or presence of artifacts for example. What we want them to do is come up with a scale which can then be used to score that aspect of audio quality. They need to define it. And then they set up a scoring scale that other people can use. We will have a whole set of those scales based around a lot of different aspects of music listening.

This might be loudness, frequency balance, clarity, artifacts, noise, intelligibility of lyrics. There's lots of things it could be.

And then the actual music samples provided by the challenge entrants will be scored by a listening panel. 50 or so people with hearing difference. That's how we're going to evaluate the entries.

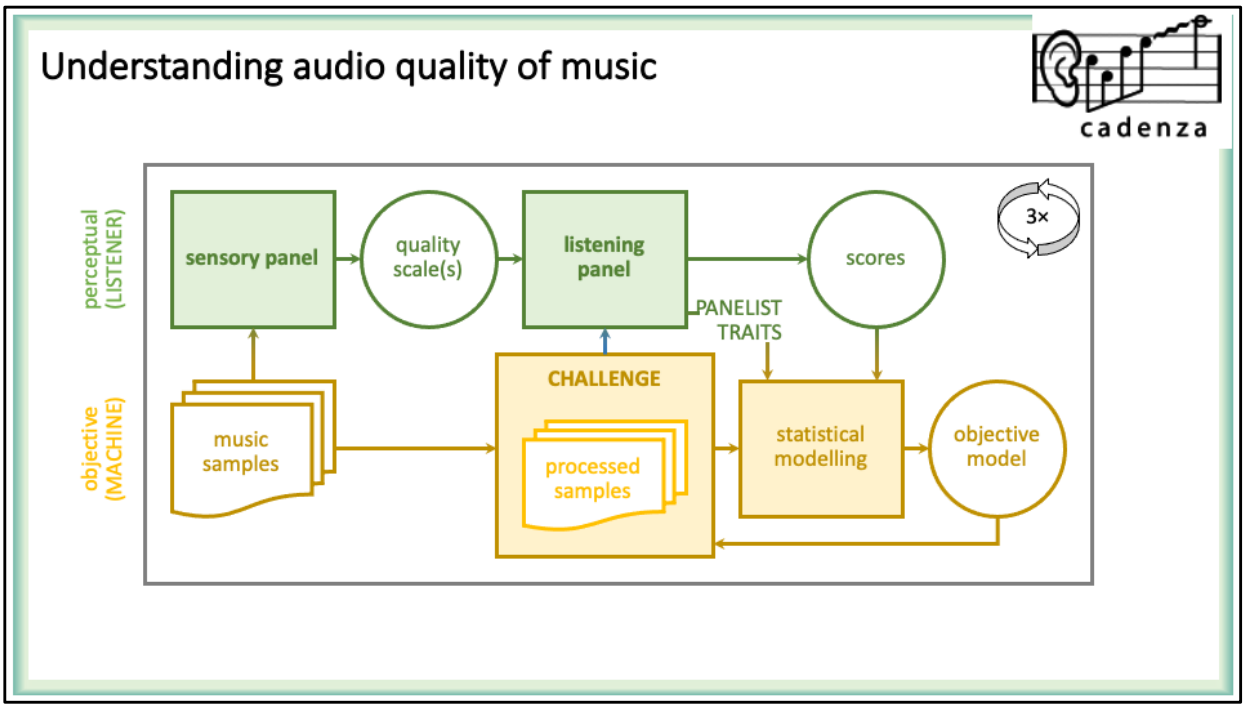

That's the idea of the project. And this diagram kind of takes us through the different phases. So we're going to start with the sensory panel or focus group. You take some music samples and they develop some quality scales that can be used by a listening panel. The challenge entrance produce lots of interesting ways of processing and remixing music that our listening panel evaluate. And we go around this loop improving how we understand and score the audio quality of music. And we actually have three challenges. So we go around this loop three times.

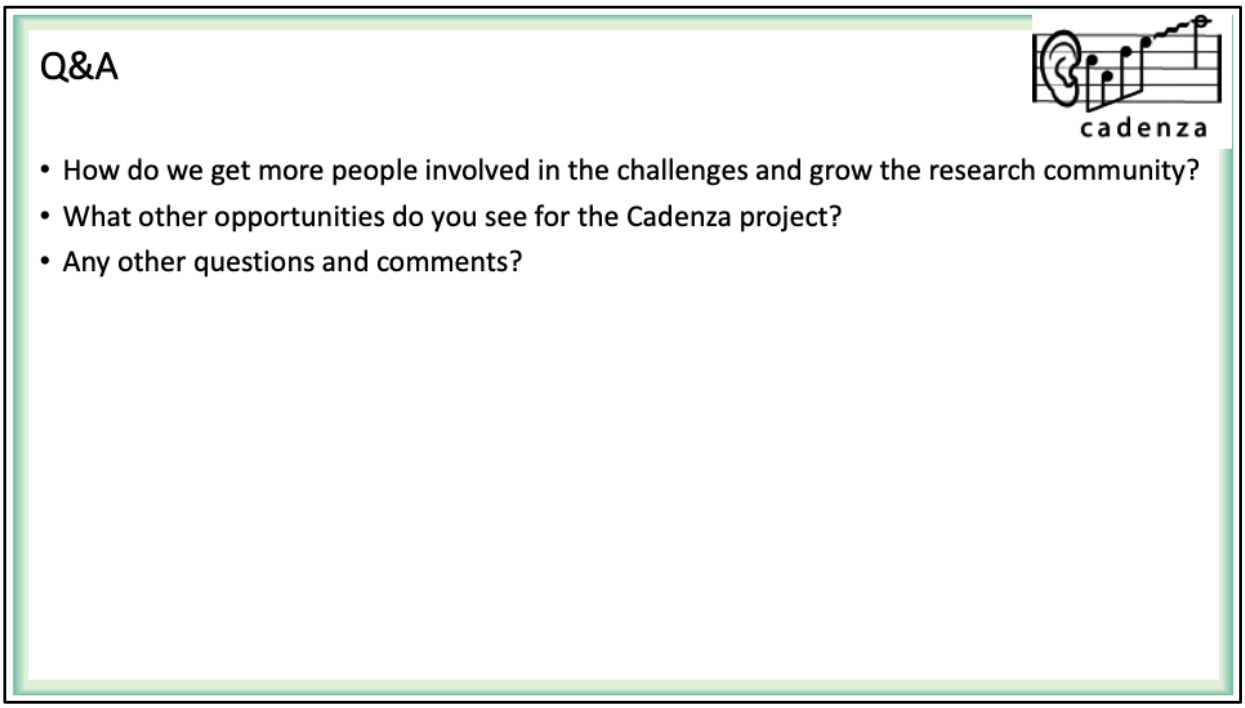

So that's the outline of the project. As I said, it starts on officially the 1st of July. First challenge will be early in 2023. Sensory profiling will start early in the autumn. That's kind of roughly the timescale we're working to. We wanted to talk about it now, even though it hasn't started yet, to get your input to shape the work. We got some already from the padlet and wordclouds. But now is a chance for more. It could be just you've got questions and comments you want to make about the project, and that would be great. But the question is I'm really interested. What opportunities do you see for the cadenza project? Maybe there is something you thought

of I didn’t mention we should definitely do. And the other thing we always have a challenge is we want to grow this research community so more people in signal processing are involved in people who have hearing different. So how do we get to those people? Publicity is really important.

We're just going to do this to plenary, by the way. So questions in the chat box, or raise a hand to comment.

Thank you very much.