Hello and thank you for coming. My name is Alfie Cliff, and I am a graduate of BSc Music Technology from Birmingham City University.
Today I'm going to be discussing my final year project, in which I wrote a dissertation exploring localisation through bone conduction audio and built a "surround sound" bone conduction headset to establish whether placing a 5.0 array of bone conduction transducers around the head would allow participants to effectively localise sound objects through the device.
The purpose was exploring new and potentially more effective ways of creating a perceived spatial audio field.
The rationale was to make spatial audio experiences accessible to a wider range of hearing differences particularly those with difference related to the middle or outer ear, which would benefit users of cochlear implants or other bone conduction devices.
Why does this need to be explored? Well, in recent years, technology which intends to immerse its users in the experience has been increasingly sought after, the prime example of this being virtual reality, however other examples include gaming and surround sound cinema.
As this technology becomes more popular, solutions must be found to ensure immersive audio experiences are open to as many people as possible.
When carrying out research for this project, I could not find any previous literature that explored the possibilities of a surround sound bone conduction device, and there was little literature which attempted to improve or study the effectiveness of localisation through bone conduction devices, even though many people rely on these devices to hear. So, this motivated me to carry out the project.
Before I begin, I would like to thank Professor

Andrew Hugill for inviting me to talk here today, and my project advisor, Dr Honghao Leo Wang, for providing valuable input while this project was underway.
So, I'm sure most people here are familiar with bone conduction but allow me to cover some key details.
To make this clear; I am going to use the term 'air conduction' throughout this presentation to refer to acoustic hearing, as it makes sense in this context.
People who typically perceive sound through air conduction use their tympanic membrane; also known as the eardrum, which reacts and vibrates according to disturbances in the airwaves, the cochlea then receives the information and transmits it to the brain.
Alternatively, bone conduction sound, as described by a 2008 study from Furuichi, is perceived through vibrations that stimulate bone structures, and these vibrations travels directly to the cochlea without the use of the ear drum, making it a useful listening method for people that are unable to use their ear drum or auditory canal.
However, In the context of localisation, bypassing the middle and outer ear poses a key issue when attempting to translate air conduction localisation techniques to a bone conduction device.
Localisation is described as the psychoacoustic ability to determine the direction of audio.
A well-established method that our brains use to do this is known as The Duplex Theory, which, according to a 2016 study from Hartman, states that the brain can determine the direction of a sound source on a left and right plane by distinguishing (sometimes very small) differences in sound levels, known as interaural level differences (ILDs), and the delay in time of arrival in one cochlea relative to the other, known as interaural time differences (ITDs).

This ability can be translated to bone conduction audio relatively easily, if time and level differences are present in a bone conduction device, the brain will identify the direction of the sound being produced.
However, that only tells you left and right in terms of sound direction. To accurately localise air conduction sound from all directions and distances, a method known as head-related transfer function is used, where, as described in a 2000 study by Begault; the brain identifies several spectral cues such as filtering from the outer ear, the amplitude of sound waves picked up by each ear drum, and the amount of reverberation on a sound in a particular space, and utilises these cues to identify the location of a sound source in any direction.
Air conduction listening methods such as surround sound are effective because they anchor people’s ability to localise by producing the sound in the space. But spatial audio experiences have also been replicated through binaural devices, by artificially adding head-related transfer function cues to audio sources. How can we translate these methods to bone conduction audio if it bypasses the outer ear? Well, I will start to explain how I tried to answer this.
It’s important to note that localisation is not a natural ability, the brain learns how to use our own hearing to establish the direction of sound sources by creating neural associations between head-related transfer function cues and the direction these cues correspond to. This learning process is presented in a very interesting study by Hofman, where participants had their localisation cues modified with outer ear moulds that had a different shape to their own. They initially lost their ability to accurately localise sound sources, but with extended use, they managed to learn to use the new ear moulds

to localise. After this, the participants did not lose the ability to localise through their own ears again, they could adapt between localisation cues.
Now onto the project. Before I began the dissertation, I suddenly had an idea – to build a device which positioned an array of bone conduction transducers around the head and used multichannel mixing techniques to attempt to recreate a spatial audio field. I searched for current research related to bone conduction and localisation and could not find anyone who had done this before. It occurred to me that this may be because it doesn't work or make sense, and this feeling developed the more research I carried out. But I ultimately decided to go ahead with the project regardless, to establish the viability of this idea.
So, I created a bone conduction headset with 5 transducers, placed in the same relative position to a 5.0 surround sound loudspeaker set up, and then created a testing method which compared people's ability to localise through both the air conduction and bone conduction listening methods.
By carrying out a test on both methods, an evaluation of the effectiveness of the device could be carried out in comparison to the current air conduction listening method.
So, this is the headset I built. Firstly, I'd like to say that this project took place amid covid-19 lockdowns, so I very much consider this a prototype, which I built from home using supplies I could order online and build in my university home.
The XLRs could be plugged into the output board at the studio, and this would reproduce surround sound audio from Pro Tools through to the transducers inside the headband.
The headband would produce sufficient pressure for participants to perceive the audio outputted through

each transducer. To compensate for perceived level differences in each transducer, I would play white noise through them, slowly bring the level up from silence, and the participant would indicate when they could perceive the stimulus. These levels would then be kept consistent to each other throughout the main test.

So, as mentioned, compensation tests were carried out first.
Then, moving stimuli would be played. This would be 5 seconds of white noise which would start either directly forward or directly behind, and then rotate clockwise or anti-clockwise around the head and end up back at the starting position. The participants had to identify the starting position, as well as the rotation direction, and then rank the how difficult it was to localise through each method, on a scale of 1 to 5, 5 being extremely difficult. This was to see whether participants had initial awareness of the spatial field.
Then, the stationary stimuli were produced, which was the longest portion of the test. This consisted of 3 seconds of sine tones ranging from 500Hz to 15kHz, each panned to random positions around the head. 48 tones in total would be played, and participants would have to identify the directions. The purpose of this was to determine if there was a difference in localisation ability depending on frequency. This was then followed by some white noise impulses, to test localisation ability for stimuli with a wide frequency range.
Each test was repeated through the device and a loudspeaker surround sound setup, where participants were instructed to keep their head still to ensure a fair comparison.
Unfortunately, I was not able to test this on a wide

range of participants due to time pressures and social distancing restrictions, but 3 participants aged 21 to 22 were tested on, who reported no hearing differences. Only one had previously used a bone conduction device, and very briefly. There are obvious limitations to this, which I will cover later.

Given what we know from research and the circumstances of the test, it was obvious that participants were going to localise audio much more accurately through air conduction listening, as they have already developed familiarity with air conduction localisation cues.

In the moving stimuli test, participants correctly identified the starting position and rotation direction of the audio, but all participants stated it was significantly more difficult to localise through the bone conduction listening method, even participant 1, who gave entirely correct responses through the bone conduction device but gave one incorrect response when listening through loudspeakers.

In the stationary stimuli test, participants identified the direction of audio at a rate of 71.25% correct responses to the air conduction stimuli, compared to 25.69% correct responses to the bone conduction stimuli.

As suspected, the participants were also much more capable of identifying simply whether the stimuli were coming from the left, right or central position. There were 93.75% correct responses to the air conduction stimuli on this plane, compared to 66.66% correct left to right responses to the bone conduction stimuli.

There were no significant findings regarding localisation of different frequency ranges.

The most interesting aspect of this test was an

observation made by the participants and me. As the testing progressed, the participants began to give quicker and more confident responses, they were not necessarily more accurate responses in all cases, but it is possible that the participants gained some perspective of a spatial field.
While this test showed that this headset did not immediately produce audio which was accurately localised by the participants, the research suggests the idea could still be effective.
From what we know about localisation, it was reasonable to hypothesise that the participants had begun the process of learning to localise through the bone conduction headset, and as mentioned earlier, the study from Hofman has shown that we can adapt to new localisation cues. So, in this case, the characteristic differences perceived between each transducer may be recognised by the brain, as the sound which travels through these transducers will be filtered depending on the bone it is travelling through to reach the cochlea and the density of the tissue dividing the transducer from the bone.
With enough use, and if the user has the real audio position affirmed to them after they have heard it, to create a neural association between different sound characteristics and directions, perhaps localisation would be possible, and effectiveness could be increased. However, this must be established with further research and testing.
The primary limitations of this research were the scale of the testing and the general build quality of the device.
The limited number and diversity of the participants could be vastly improved by gathering participants of all age groups and with a wide range of hearing differences, this would really help establish the effectiveness of the device, particularly when it is

used by the people that it aims to cater to. The testing methods would also be improved in the ways discussed – longer usage to give participants the chance to learn to localise through the device and having the direction of audio stimuli affirmed to people during the test may be pivotal in establishing the neural associations required to localise the audio.
The device would hugely benefit from a greater focus on design, making it easier to use, and if this new design applied greater and more consistent pressure to the transducers, it would also make the headset more effective by improving audio clarity for the user.
The limitations set out the main issues which must be addressed if future work is carried out. The idea could be followed and tested again with those key improvements to really establish the effectiveness of a device of this nature.
However, I feel that the viability of this work is ultimately to be decided by anyone who seeks to benefit from using the device, which I aim to discuss with people here today.
Thank you all for taking the time to listen to me speak. Please feel free to ask any questions.