Aural diversity in acoustics research

Bill Davies

We begin with a bit of context. When we teach acoustics students about hearing we greatly emphasise normal hearing. Here for example is a plot of the normal hearing threshold that I show my students. Sound pressure level is plotted vertically against frequency horizontally and the curved line represents the lowest sound of all that we can hear if we've got unimpaired hearing. Most concepts and designs in acoustics are based on the idea of normal hearing.

I also tell my students that human hearing can be modelling a staged process. So here we see a simple block diagram. There are four main blocks. The acoustic signal goes in at the bottom, to the ear. The bottom block represents hearing, where your cochlea converts the acoustic signal to neural impulses. The next two blocks are about processing the signal at a variety of levels of sophistication, especially parsing it into different sound objects and then further processing just one of those sound objects such as a speech signal. The block at the top is where we determine some kind of response to the sound; this might be an emotional response like annoyance or a preference or some other discrimination.

Let’s see what the literature has to say about diversity in hearing.

There are many papers on differences in hearing; mostly on one specific hearing impairment or another. This is a nice example of pure-tone audiograms showing diversity in age-related hearing loss. The data are plotted as the difference between the person’s threshold and the normal hearing threshold on the vertical scale and frequency on the horizontal axis. The data are divided into age groups so we see that, as expected, people 60-69 years old have a much bigger loss at high frequencies than people 20-29 years old. What’s interesting is that even within those fairly narrow 10-year age bands, we see a significant range around the mean values. This represents individual differences in threshold.

Hearing diversity is not all about impairment, however. We can also find evidence of diversity in quite basic mechanisms like pitch detection in unimpaired listeners. Here are some data from one of my research students. He measured pitch detection in amateur and professional musicians. He found diversity even in how people detect pitch – whether they rely more on the fundemental or the higher harmonics. And also a continuous range of ability in detecting relative pitch changes and even in absolute pitch assignment. So, one might even say that pitch is a multi-dimensional perception.

Next we turn to diversity in auditory processing.

One big topic in this area is the clinical condition called auditory processing disorder. This is a diagnosis that can be given when you have a normal hearing threshold but you have difficulty in processing speech. Here is a fairly typical result from the literature. This graph shows performance in speech processing as the percentage of words understood on the vertical axis. The bar on the left with the arrow shows children without a diagnosis of APD, the control group, getting nearly 80% of the words correct for this particular experimental condition. And we can compare them with the black bar which shows a group of children with a diagnosis of APD – they only get 60% correct. In the next set of bars we make the conditions worse by turning up some background noise. Now the control group children drop to just over 70%, while the APD kids drop further to under 50% correct. One can find a broad range of outcomes for different subtypes of APD in different conditions. Evidence of diversity in auditory processing.

Another interesting area in the auditory processing literature looks at auditory experiences of autistic people. Here's an interesting result showing differences in processing speech prosody. Prosody refers to patterns of stress, intonation, pitch and so on that are found in speech. They often convey extra information besides the language elements of the speech, such as the emotional state of the person speaking. Autistic people are sometimes thought to have difficulty with detecting prosody correctly, and this has been hypothesised to stem from auditory processing disorder. This paper compared autistic and non-autistic teenagers on a task where they had to evaluate several aspects of a person speaking, including things like manners and social emotional state. Interestingly, they found no significant difference in the proportion of correct answers, but they did find that the autistic kids were significantly slower to respond. The authors suggest that the autistic kids sometimes take longer not because of auditory processing problems, but because of the additional cognitive load of working out the emotional state manually from cues. This interesting result highlights a more general problem in sometimes distinguishing APD from cognitive difficulties.

Now we turn to diversity in responses.

My first example is a very common kind of response that we want to measure in acoustics and that's annoyance to road traffic noise. Here's an example of a meta-analysis. There are 20 or so different individual studies on that one graph. Each of those points actually represents an average from a from an individual study and you can see that there's a significant range there. The vertical axis is the percentage of people who said they were highly annoyed and the horizontal axis is the level of the noise that was annoying those people or not. Of course there will be many factors that could possibly explain those differences – the studies are of different locations, perhaps the weather was different, maybe the local football team had just won, and so on. But there are almost certainly are also significant individual differences going on as well.

My next example is one of my own papers, on soundscape perception. These respondents were given a collection of sounds and asked to mix the soundscape of a place in Manchester they all know well. The graph plots the pleasantness score they gave for each individual sound on the vertical axis against the level they wanted to mix it at on the horizontal axis. What's interesting about this is that the graph shows that they agreed quite closely on their preference for each of these individual sounds but where they differed quite a lot was on what sound level they thought each of these sounds should be in the soundscape.

My final example of diversity in response is from a paper on concert hall preference. Tappio Lokki found his participants split into two sub-groups, as shown on this factor plot. The first group, labelled G1, is more closely associated with clarity and definition, and in opposition to reverberance. So they want a hall to sound very clear and detailed. The other group, G2 are all about bass and loudness. (I'd be in G2.)

After looking at lots of these papers (and more I don't have time to show you) I thought that we do have a bit of a tendency in acoustics to go straight for the mean value and to want to use that to establish our relationships and use it for our design and once we've done that we tend to forget about the individual scatter in the results which can be quite significant. Here are three example plots showing significant amounts of scatter. For all three of these, the average values were used to develop equations and I think perhaps the scatter was then forgotten. Perhaps this is hiding aural diversity.

Is aural diversity found in the acoustics literature? I think we can say there are a lot of papers representing different aspects of aural diversity in acoustics, although you have to dig around quite a lot and there are hardly any papers on the topic of aural diversity per se. The great majority of papers are just on one specific hearing condition or one specific aspect. On the other hand if we think about our own individual aural experiences then what medics call comorbidity is common. For example, I have age related hearing loss and also noise induced hearing loss but then I'm also autistic so I have at least some elements of auditory processing disorder. I could also quite easily put myself in a preference group for concert hall acoustics or soundscapes. This made me think that (of course) aural experience is multidimensional and yet I'm one person so can I represent myself as a point in some multi-dimensional space?

Other fields have had significant success by representing important concepts as a multi-dimensional space. Perhaps the best example of this is face-space which is used by researchers in the visual perception of human faces. The idea of face-space is that it is a multidimensional space that represents all aspects of face perception. Each of us maintains a face-space inside our heads and we use it to evaluate each face we encounter. Face-space has been successfully used to explain several important results in facial perception.

Face-space inspired Jane Conway to invent mind-space. This is a conceptual space that represents every mind that you encounter in your life.

Inspired by these examples, I tentatively suggest aural-space. This is a multi-dimensional space of subjective aural experience. Each person's aural experience would be a single point in aural-space. The real aural-space is likely to be a high-dimensional space. Here I'm showing an example 3D reduction with toy examples labelling the dimensions. We don't yet know what the dimensions will be, expect that they will represent significant aspects of subjective aural experience. The actual dimensions could be determined empirically.

Each person would be a single point. Two people with similar aural experiences would be close together. A third person with very different aural experience would be some distance away. A group of people with similar experience, an auraltype, would be a cluster of points. Any vector in auralspace, say between two clusters or two individual points, would be a measure of aural divergence along that axis. So one could generate measures of sural diversity on different dimensions. Where would normal hearing go? In the centre, of course. In this way, we would be able to build up a fair representation of the complete extent of aural diversity, without losing the useful concept of normal hearing.

I think the aural-space concept has several advantages. It represents the most important aspects of subjective aural experience. It represents the full range of individual differences. It represents common experiences fairly as clusters. It retains normal hearing as a central cluster while fairly representing other auraltypes and individual divergence. It also has at least one pretty significant disadvantage. We currently have a complete absence of data to populate and scale aural-space! So we can't yet say what the actual dimensions will be or where the clusters and points will be within it. But if we like the idea enough then it would be possible to obtain that information by experiment.

To conclude, then, we can say that normal hearing dominates acoustic thinking. Aural diversity is to be found in the literature even if it is a little bit hidden, but there is evidence for diversity at all stages of the human hearing chain. Finally, I proposed aural-space as a conceptual framework to try and encapsulate all of this.