The Case Of Riddler's Robot Musical Death Trap

School of Batman,
28/06/2018, Podcast transcription

School of Batman is a figshare podcast, where we ask academics to fight crime using their research. From neuroscience to linguistics, academia is helping Batman save Gotham!

CG Chris George
HW Harry Whalley

CG Welcome to the School of Batman, a podcast where we ask academics to help Batman fight crime using their research. I'm your host, Chris George, amateur scientist and professional Batman enthusiast. In today's episode we'll be discussing the case of Riddler's Robot Musical Death Trap. We're pleased to be joined by Harry Whalley, who has a PhD in Composition from the University of Edinburgh and is currently a course leader in Music Composition and Technology at the University for the Creative Arts. Harry we like to start the podcast with a little bit of background on yourself, your journey and how you got to where you are and some details about your research as if you're at a dinner party explaining what you do.

HW I started with a background in jazz played by ear and played a lot of funk and jazz and blues as a teenager and then that lead me to study Jazz at Middlesex University where I was very lucky to have Malcom Edmiston as a lecturer in Composition and that's where I really realised that composing and arranging music was what I loved best. I taught in a school for a little while and then met Nigel Osborne who is a really interesting composer and a bit of a polymath at a conference to do with the holy grail of all things. I immediately applied to study a masters in Composition at Edinburgh, under him, then I did a PhD there. And this relationship between composition, improvisation and technology has been part of my interest, really all along. My PhD was based on the book, Gödel, Escher, Bach by Douglas Hofstadter, thinking about how music, mathematics and art can be used as analogies to more complicated and philosophical ideas. So I don't want to commit the heresy of paraphrase about that book but it's been a wide journey taking in things from different disciplines all the way.

CG Right and so you say you started your journey by playing by ear, so for the none musicians listening, what does playing by ear mean?

HW So playing by ear is where you aurally understand the music and then respond to it. I suppose it's in contrast to sight reading or reading notation. It's how many music's in the world are communicated. If you think about folk music or Indian classical music or singer-songwriters, you apply your understanding of music theory and your instrument. It's quite an intuitive way of working in some respects but it also takes a lot of practice which can be formal, you can formalise that process as well as just learn how to do it yourself.

CG I do think a lot of people will be surprised that many musicians have no real concept of musical theory or how to sight read or maybe even know what they're playing. Has any of your work touched upon how that interpretation works with how can somebody improvise, play with other musicians, without real concepts of musical theory?

HW Well I've done a lot of workshops as it happens with various groups of people from school children to young adults who are using music because of its social cohesion rather than perhaps intrinsic having an interest in it. I've also worked even in a prison, working with some people there who

were writing songs about their experiences which was interesting. And it fascinates me how with all those very desperate groups of people, everybody has their own way of translating what they hear into a concept that they understand. So in the case of somebody writing songs that are very meaningful for them, the focus is on the lyrics and perhaps a melody is really just an extension of the words and the shape of the phrases of the text. Working with children, sometimes it's very gestural. It's the physical aspect of playing is what they remember. You hit this bit then you move over here and play that and then you follow this. So it's quite a physical thing. Other people can think quite conceptually. It's really amazing how so many people think about what's ultimately just pitches and rhythms and instrumentation in so many different ways and the music theory that we, standard western notation or common notation, is one highly formalised way of thinking about music. Which of course because it's more or less universal, means that we can communicate to one another with it very easily but in some ways it's has a lot of advantages over other types of communication in that it can be very precise. But in some respect it just happens to be the one that is formalised and understood.

CG So I think it's probably time that we move onto your story. And before we start our story we like to ask our guests if they have a favourite Batman actor that they would like to represent their Batman in their stories. So do you have one?

HW I do yes, I think it's got to be, it's a little bit meta but Del Boy or David Jason in the Only Fools and Horses episode where he's also the king of Peckham in that episode which is really nice counterpart to Bruce Wayne.

CG I think you're our first guest that's chosen Del Boy.

HW Well hopefully not the last.

CG Awakening with a foggy head, a sickeningly familiar shade of green blurs our hero's vision as the room begins to take shape. Vison comes to focus, the Batman awakens in a cage inside a Riddler puzzle room, a strange helmet strapped to his head and surrounded by a ragtag robot orchestra. Speakers burst to life and fill the room with rhythmically complex dissonant music, quite unlike anything Batman has heard before. This is followed by the disembodied voice of the Riddler announcing the challenge, 'make the robot orchestra reproduce the music with a hundred percent accuracy or die'. A sticky situation he's in again here.

HW I think it's a very sticky situation. So I'm not sure if everybody so far has managed to help Batman but this might be the last time because it's a very difficult challenge.

CG Yes so we're in a Riddler puzzle room. So first of all we're hearing some very strange music and some very difficult music so I think you have an example of how that music might sound.

HW Yes so I worked on a piece a few years ago, which is by chance similar to this story, where an EEG head set which is an Electroencephalogram, was used to control in this case an artificial neural network using the brain, if you like, as an interface with a computer for a musical end. I'm by no means the first person to have tried this but it's interesting how so much of our experience with instruments is through breath and through the hands and to take that away and just interface directly with some sonic material, through a brain interface, is quite a startling thing. So I think the first thing that Batman is going to have to come to terms with is how disorientating this is going to be. I'm not sure, he said the helmet is strapped to his head but I don't know if Batman himself is tied down or not or whether he's able to also conduct these robots when he tries to make a plan to escape.

CG Yes as the keeper of the story it's only the helmet but he's in a cage.

HW So there are a few parts in this scenario. First Batman has to code the music himself which as we were talking about earlier in terms of playing by ear is related to ear training. I think Batman, as the character Bruce Wayne, I know he's highly trained in martial arts so I wouldn't be surprised if he had some piano lessons enforced on him as a child, so maybe he's done a bit of this before?

CG Of course he had a piano to get into the bat cave.

HW Of course he had a piano to get into the bat cave that's right. So this might just be one of his skills that he's going to need to really break out of here. So he'll use his martial art ear training skills to understand and take apart and analyse what he's hearing. So the first thing that he'll want to think about is what pitches are being played or what are the notes that are being played and there are two different ways that people tend to do this. One is with what's known as relative pitch which is where if you hear one note and you know what it's called then you can figure out what the next pitch is. So if you hear a C and then you hear a D, if you know that the first one was C then you can figure out well that's one whole tone higher so that's what that is. And that's something that you can learn but there's also absolute pitch which is quite rare in the population and that is where if you hear a note, even without any context, then you can discern what the pitch is.

CG And where does absolute pitch come from? Who defined that?

HW I'm not sure where the term comes from but it's interestingly more common in tonal languages, such as Mandarin, where intonation isn't just related to things like the end of a sentence if it's a question we tend to intonate upwards, but forms more part of a fabric of the language. So it's more common in people who have tonal languages. The mystery is not so much why some people have absolute pitch or perfect pith but why most people don't because in a way relative pitch is a more complicated thing to do so with absolute pitch the sound as it enters the ear canal it vibrates a membrane and eventually it gets down to the cochlear and the cochlear has these little hairs on it that vibrate in sympathy, they're in a fluid but they vibrate in sympathy to the vibrations that are going on. And there is a direct relationship between frequency and how that neuron fires. So if you hear a pitch at 440 cycles per second which we would call A then the cochlear fires an electrical impulse, it turns it into a rhythm if you like at 440 cycles per second and so there is a one-to-one relationship between the firing of the neuron and the pitch which means that it's always the same channel that is being stimulated when you hear that pitch. So it's weird that we wouldn't be able to learn that. The mystery is not so much why some people have it because we can look at the mechanics of how the ear works and see that it should be a fairly straight forward thing to learn but it's more why some people don't, or most people don't.

CG And with the idea of that the idea of C as a subjective, creative idea there are intrinsic natural properties to that note of C right that would still exists if humans weren't around? Is that the right way of thinking about it?

HW Yes absolutely. So there's a difference between the pitch which is a class. So if we have 440 cycles per second which is A just below middle C and you hear 441, then they're different frequencies but they're both basically still A. In fact you could choose 441 to be your baseline for A, there is a bit of variation in how you want to tune it. But the frequencies is just the measured physical number of cycles per second.

CG And do different countries or tonal languages have an influence? Do countries that have different scales and different intervals and different things like that, do they have a higher or lower frequency of people who can have absolute pitch?

HW So they do with the tonal languages have a higher proportion. I think there might also be, there is a higher proportion in people who have, if you go to a conservatoire for example, there'll be a higher proportion there. So it's hard to know sometimes what the cause and effect is and also in other words if you have perfect pitch and that's been discovered at a young age then you might be encouraged to pursue your career in music further. It's definitely by no means necessary thing to have, in fact it can be a little bit of a hindrance for some people. I don't have it but if something is being transposed, so for example if you're playing something in the key of C but it's been transposed or people are transposing using a transposing instrument or for some other reason and it's a bit like someone saying this is red and they're pointing at blue, it would be an uncomfortable translation to have to make whereas when you have relative pitch you just say alright well red is the new blue and you just carry on.

CG So we've discussed one aspect of what he's going to have to pick apart, he's going to have to pick apart the aspect of pitch. What else does he have to look at?

HW So of course there's rhythm. So pitch, there is two ways to think about pitch. One is in the frequency domain so that's saying this frequency is persisting over time and you can use things like fast Fourier transform to make spectrographs and see those images and get an idea of which frequencies are persisting over time. And the other way to think of pitch is almost as a rhythm. So they're two scales of the same thing I like to think of it in that way. So a rhythm is an event or a number of events which are happening over time and at the very smallest scale we can think of a frequency, so say we take A of 440 cycles a second and that might be…

if we slow it down and then an octave higher is 880 cycles a second, so you get a ratio of two to one. So these two pitches being played at the same time, you can think of as proportions that you also hear as a rhythm. So an octave is the ration of two to one so rhythmically it is…

or the ratio of three against two is, in terms of pitch we hear that as a perfect fifth…

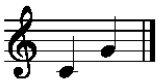

but the rhythm of three against two is…

So rhythm is all to do with proportions. At the very micro level like that which becomes pitch but also over greater periods of time. So a rhythm has a relatively short period, in other words it repeats quite often so our expectation of that rhythm, we know what's going to come next and we feel quite comfortable with that. And ratio's that are more irrational than that or have less simple rations than three against two or two against one, become very hard for us to predict and that's the case both in terms of rhythm and pitch. So a simple ratio like two against one or three against two as a pitch will be a very constant interval, like a fifth or an octave or a third, were as a dissonant interval in pitch is the same as an irrational ratio in rhythm and both those things, so simple ratios in pitch and rhythm in a way create simpler music and more complicated rhythms and pitches create more dissonant and complex rhythms. So I think with Batman in this cave, the

complexity of the sound that he's hearing is going to mean that the ratios that he's going to have to analyse are going to be difficult to hear and not these simple ratios of fifths and octaves or two against one or three against two.

CG Listening to that has inspired a question that might come out quite rambley so do bear with me. So humans have a frequency range of sounds we can hear but that's not the whole frequency range much like with light, there's different areas of light and different animals can hear different frequency ranges. So is that just because that's all our ears have evolved to be able to hear? If we could bypass the ears and directly stimulate the brain with this brain machine interface that you were saying, could we then start to hear for example music for bats, and extend that range of notes that we were able to hear and have a whole new range of music available to us? Is that in theory possible?

HW So our frequency for our ears is roughly 20 hertz, at the lowest 20 cycles a second, at the lowest range. Really that's just barely audible it's more like a rumble that you feel. Up to 20,000 but that's really with people who are young and have very good hearing, it drops off as you get older as the ear deteriorates and the hairs in the cochlear become less sensitive. In a way we've had this already in reverse in that we used to have super human hearing, if you are 40 and you can remember what it was like to have hearing when you were 10, you used to be able to hear much higher frequency. So we've already experienced this. The question with the bats is whether our brains are so closely evolved to the inputs that they're expecting that they might try to translate it into the perceptions that we already have. So, I think an analogy to this would be if we were to develop infrared eyesight, maybe our brains don't have a way to separate that from our perception of colour already. So, we can translate infrared into the visible spectrum, and in the same way that you can transpose frequencies above our hearing range, or below our hearing range, into our hearing range. But the question is really, is it possible for the brain to then interpret that in a way that is different to what it has evolved to make… I don't know, I was going to say make us privy to. It's a bit of a difficult question of perception there, but…

CG More philosophical than I thought.

HW Yeah. David Eagleman has looked at this a little bit, in terms of mapping one sense onto another, looking at things like touch. So, if you wear a suit that can manipulate your back, so you can feel pressure, different areas of the back as a sort of bitmap, then over time your brain can learn to translate that into a sense. So, there are multimodal translations between the senses, and of course there are also people with synaesthesia who have this naturally, where pitch perhaps, and colour, are perceptually related. So, if you hear a particular pitch, you might, in a mind's eye sense, sense a colour, or a taste, or so on. There's lots of different types of synaesthesia. So, the perception of music, or sound, is more than just the analysis of rhythms and pitches, because we have bodies, and we have senses, and our brains make sense of it in a way that is not as reductionist as the way I described earlier with pitch and rhythm. So, going back to Batman, one of the questions that he's going to have to really think of, is how, if he hears this music, in what way is 100% accuracy of the music? What does that really mean, because it's getting mediated through him, then back out to robots. So, there's lots of different types of translation that go on, and in the case of music, that includes the way that music and sounds affect us, and our physiology.

CG So what kind of technology would this strange helmet have to be to allow Batman to be able to get to the degree of granularity that's going to be necessary to reproduce this, and give those input signals to that robot orchestra? What are we looking at there?

HW The obvious example of this would probably be an EEG, or electroencephalograph, which is quite coarse in the amount of information it's able to get from the human brain. There are more evasive

technologies, like neuro-link, and also of course FMRI, which can image the brain, but those are very large machines that haven't yet been shrunken down to the size of a helmet. But I suppose it could be an MRI of some sort. But whatever the case, the truth is that the amount of detail that any kind of current computer-brain interface is able to get at the moment, is not really enough to have the amount of complexity required to play instruments. But it can give basic commands, which you can train. So, I imagine in this case, Batman is going to be listening and analysing the music, perhaps with some sort of computer aid to help him. I think it might have to be a bit of a cyborg moment for Batman, where he uses both his abilities as an accomplished pianist, as we decided that he was earlier, and as well as someone who uses technology to aid him. And I think this relationship between computer interfaces and the human is where he'll be able to succeed the most.

CG So going back to some of your researches there, I think for the purposes of this discussion we can assume some quite advanced technology. The Riddler is a very smart man, and maybe has access to technology that isn't available yet. But looking at the technology that is, how was the EEG and neural network that you were working with, how did that work?

HW The EEG was basically looking at how much, if you like, global activity a particular musician had at any one moment during a performance, and that was then fed to an artificial neural network that was developed at Plymouth[1] University, which had 1,000 neurons in its model, and any time one of those neurons fired, it would play back a little sample, or a little grain if you like, of sound from something that it had been listening to. And what it was listening to were the musicians on stage. So, this had the effect that as the performer had to concentrate more on what was going on, then the artificial network would increase its activity, and so the audience would be able to hear a cloud of sounds around that, related to how much the performer was concentrating.

CG And who is the performer in this instance? Is that the person with the EEG on?

HW That's the person with the EEG on. Yeah, so in this… the first time it was performed it was with Pete Furniss, who is a clarinet player. I actually tried it once with myself as a composer on a stage with the EEG, and I projected the graph onto a projector screen for the audience to see, so they could see… it was basically between 0 and 100, you could just see this wiggly line move up and down. And as I sat on stage and looked up at this graph to make sure it was working, it was flatlining at 100, and I thought 'Something has gone wrong, there's a connection broken.' But it was just because I had so many nerves at the beginning of a performance that I was just maxing it out. So, I had to tell people that I'm going to relax and calm down for a while before we can start.

CG Right. So, when we have a master like Batman with intense control of his body and his mind activity, we'd need a very probably quite complex version of the technology that you were working with, that neural net, so maybe with… you had 1,000 neurons, you'd need significantly more than that. But you say… so what kind of brain states, and what kind of physical activity would he have to do, to enable such nuanced information to be sent to that robot orchestra?

HW Yeah, well one of the fascinating things is that the type of task that a musician who is trained in that task, when they start there's a huge amount of brain activity. So for example, if you're able to sight-read, and you start sight-reading, then the amount of global activity peaks very quickly. But then as you continue in that task, the level decreases over time, and can actually go below baseline. So, the amount of activity and the task are related, especially at the beginning when you start that task, and this is known as task switching. And you might be familiar with it when you're driving, then you're distracted by something that you're not able to return quickly to driving, and

[1] The recorded version mistakenly said Portsmouth University

you're distracted for a while as you change the radio. So, if you had an EEG on then, you'd see that as you went from the task of driving to the task of trying to turn one radio station to another, that there'd be a big spike in activity as you go between those tasks. So, there's quite a big cost in switching from task to task. So, in the piece with Pete Furniss, the compositional intention was that the artificial neural network would be fed by him, and therefore he would need to have a high level of activity for as long as possible throughout the performance. And so, the way to achieve that was to make the performer switch different types of musical task, so one moment he would be sight-reading, and then the next moment he would be interpreting graphic notation, then the next moment he might need to gesture, then another moment he would have to play something via memory, and then back to gesture, and so on. And every time we performed this, he would have a slightly different version of the score. And this was so that the whole thing didn't become the single task of playing by memory. So, it's a very taxing thing for him to do. You could also train these sorts of systems' facial movements, because they need to be able to understand when the electrical signal that they're receiving is coming from the face, rather than the brain, in order to cancel it out. And therefore, you can also use that in these interfaces to act as another level of control, so in particular maybe raising the eyebrows or grinning, and that kind of stuff, can also be used as inputs.

CG So back to Batman and his final escape. Is it reasonable to visualise this as some kind of 0-100 scale, that by Batman doing, say, a variety of different martial arts, and switching between all the martial arts he knows, mixed with facial combinations, mixed with tapping his head and rubbing his belly at the same time, he could manipulate that 0-100 gradient in a way that would allow him to feed that neural network, in such a way to recreate that music? Is that a fair visualisation?

HW I think in a DC comic land, that would be fantastic, that would be a great thing to do. The other part of this that we didn't look at is of course how the robot orchestra is playing. There's a history of things like… actually I was at the Museum of Mechanical Sound in Utrecht recently, where they have machines that play violins, and this kind of thing. So maybe the neural network is connected to these physical robots that then have to translate the neural net into gestures. So, then we've got these different levels, or layers, of mediation. So, we have the music, the complex music, that goes from the speakers, so it's an electrical signal into an acoustical signal into Batman's ears, then they're converted back from acoustic to electrical again as they go into his brain. And then that's interpreted by an organic neural network, i.e. Batman. Then the EEG headset converts that electrical signal back into a digital signal, which is then an artificial neural network, which would be a digital representation of a biological system, perhaps. And then that feeds attenuators and motors and whatever the robot orchestra needs, so that's translated into a physical movement which then plays an instrument, which creates vibrations in the air, and eventually we're back out into acoustics again. So, there's all these different layers of translation - biological interpretation; digital; acoustic. And I think a lot of music, it goes through these different types of mediation all the time. As I was talking about a little bit earlier with the way that sound and music affects us, it's all about these different types of translation, and at different levels and layers and interpretation, and symbolism. And it's a really interesting, complex thing. But yes, at the same time it's really simple. It's just vibrations. It's a bit paradoxical how complicated it is, and how simple it is, at the same time.

CG That's all for now. Thanks for listening, and thanks to Harry for joining us.

HW Bye!

https://soundcloud.com/schoolofbatman/the-case-of-riddlers-robot-musical-death-trap-harry-whalley