Podcast Essentials - Audacity

[Start of recorded material 00:00:00]

Eva: [00:00:09] Let me get my screen share started. I'm just going to start with the PowerPoints. And then I'll go into Audacity and then I might bounce back and forth a little bit. But yes, so sound editing and Audacity is part of the Podcasting Essentials Workshop Series. My name is Eva. The first name is confusing but you can just call me Eva. I am going to do a few things today. [00:00:38] It won't strictly be in perfect order because a lot of times with audio, you know, you end up going back and forth a little bit. You might need to edit and then change your levels, and then go back and edit again. So you know, but we'll do what we can to make it linear.

[00:00:58] So we're going to start by talking about recording and importing and other ways that you get sounds into Audacity. Then we'll go over balancing the sound, removing noise in Audacity, assembly, layering, sort of building other sounds around just the main sound that you have. And then I'll talk about effects a little bit, not too much. I have to like hold myself back, because I really like doing that. And then we'll go over just making sure that your files are managed well towards the very end. [00:01:26] It's like not as exciting but – I'm really curious what sort of projects people are working on or thinking about, or how much experience we have.

[00:01:38] I usually ask people like on a scale of one to five, is five is like I'm an expert and I should be teaching this workshop, and one is like I've never even used a phone, like how do we feel about our own skill level. Like are we feeling like we're at a two, feeling like we're at a four? No judgment, just curious what you guys are interested in, if you're interested in like fiction podcasting or, you know, narrative nonfiction, something like that. Good. [00:02:09] OK, cool. So we have some familiarity. We're not feeling like experts. That's totally OK. We don't feel like an expert every day. But yes, OK, we're at four 00:02:19. Good for you. That's great. OK. That will help me sort of think about what I'm talking about.

[00:02:28] Let's start by talking about recording and other ways that you get your into Audacity. Basically there's three ways. You can use a microphone or a mixer to record directly into Audacity, which is what's really great about Audacity, that you can both record and edit in it. You can also import existing or previously recorded sounds into your project. So whether that's a sound file that you downloaded off a website like FreeSound or somewhere else where [00:02:56] where they have like creative comments, licensed sounds, or if you've recorded using a different setup and you want to, you know, then import it onto the computer, you can do that.

[00:03:08] Audacity can actually generate sounds for you as well. So sometimes that you might want to use that is if you want to generate white noise or if you wanted to add near silence to your piece, which we'll talk about in a little bit. Or if you just want to add a simple tone like a standard A4. That's a tone that we talk about a lot in audio. It kind of sounds like a dial tone. But let me – great. Let me switch my share to actually [00:03:39] to actually sharing Audacity. I'm just going to pull that up. And has anybody used Audacity before or are you guys just sort of starting from scratch in that way? Which would be totally fine. It's actually pretty fun that we are doing Audacity today because Audacity put out a really big update yesterday.

[00:04:07] And a lot of it has changed. So if you've used Audacity before yesterday, it's a whole new application. Used it before but maybe not fully, to its full extent. For sure, that makes sense. OK, so I'm using monitors, which is why it seems like I'm looking down. That's because I'm looking at Audacity. My other monitor has my presentation. So for starters, when you open it, it looks like this. [00:04:36] I'm sharing my whole screen, which I know people don't always do, but that's because there's some important things they're going to use in these menus that aren't captured if I only share Audacity. Great. So I'm – my input is a stereo microphone that I have hooked up separately.

[00:04:56] It looks like – trying not to knock things off my desk. It looks like this. You may have seen it before. It's a pretty commonly used podcasting microphone. So I keep it on my desk so that I don't have to use the sound in my webcam. But I could also use the sound in my webcam. I highly recommend never ever using your built in microphone. In podcasting a lot of times people toss around the phrase ABB, and that means anything but built-in. [00:05:24] So if somebody, if you're interviewing somebody and they're like, I don't have a really fancy microphone, I don't know what I should be using, you can tell them anything but the built-in microphone. Even just a simple pair of headphones will sound a lot better and will isolate the sound a lot better than just that built-in microphone.

[00:05:41] So that's that. I'm using that. It's a stereo microphone so I'm going to take full advantage of that, and I'm going to record in stereo. If your microphone can only record in mono, that's OK too. But if it can, you might as well give it the opportunity to record in stereo and you can always cut that down later if you feel that the file's too big. And my output right now, ordinarily yours should be the stereo microphone. Right now it's the Zoom audio device, because I want you guys to be able to hear it. So super-basically, recording into Audacity is this record button. [00:06:09] So I'm going to press it and I'm going to speak into my microphone and you can see that the levels are bouncing up here in that green bar, and that's telling me that it is recording and the WAV form is being generated.

[00:06:23] And I'm going to press stop because I am all set with that recording and I don't need any more sound from that. So let me just play that and you guys can let me know if you can't hear it or if you can hear. So I'm going to press it and I'm going to speak into my microphone and you can see. Hopefully you can hear that. Because if you can't, we're going to have a real problem with doing the effects. [00:06:48] OK, great. So that's pretty simple. I'm looking at this and I can see from these WAV forms that they're not very tall. So if I were going to record again, I might adjust the microphone sensitivity here. So that's the recording volume. I tend to have mine pretty low because I have a pretty loud voice in general. So I had mine way down here as my default setting.

[00:07:08] But maybe I would move it up. I'm going to delete this. And I'm going to try again. And now I'm seeing some really confident healthy peaks here. And you can see it's not really blowing out. At the top I'm not getting any red or yellow indicators at the top of the window. So I know that it's doing well. Great. So those look really healthy. Want to see them going up and down like that in a pretty solid manner. In a chunky manner there. [00:07:38] So it's just that simple to record into Audacity. You just want to make sure that you press stop and not pause. And then it'll start – it'll go back to the beginning and that will be an entire track.

[00:07:50] Now if I wanted to add more tracks, I can go up here to import, import audio, pretty straightforward. It'll add another track using – let's see. I have somewhere in here examples. We'll get to this again later with file organization but I keep everything. Raw audio. Let's see. I'll just add this piece of music. There we go. You can see first of all, it's much longer than what I just recorded, so now the scale of the window has blown out to a full four minutes. [00:08:20] Instead of the previous like 15 seconds that we were working with. And you can see that this is much, much louder. So we're going to have to deal with that later. We're not going to do it at this exact moment, but when we get to talking about layering and levels, we'll fix that.

[00:08:37] So then the last matter is generating. It's just how it sounds. You can generate, let's see, a noise. It'll give you a few. White, pink, Brownian. Don't worry too much about the other two. White noise is the only one you'll probably use. And it just basically gives you 30 seconds of sort of basic white noise. It's very loud. You'll want to turn that down too. But here, let me turn my playback volume down and play what that sounds like for you. Yes. [00:09:09] So we can see that that blew out a little bit. But we hear the sound of white noise. It's kind of the sound of an air conditioner. It's similar to the concept of white light, where white light is a lot of frequencies of light. It's many, many, many colors on top of each other to the point where you can't discern one from the other.

3

[00:09:26] White noise is very similar. It's many, many frequencies, one on top of the other, to the point where you can't discern other frequencies. Or you can't pick them out individually. But you can also generate silence, which can be helpful. If you just want to keep your clips spaced out and you don't want them to, you know, accidentally run into each other, and you don't want this emptiness in the middle, you can generate silence, the default 3.5 seconds. Just a little bit of silence. And so while this is [00:09:56] called silence, it's actually near silence. So there's no way for it to generate sheer silence. So instead it's generating like the smallest possible amount of white noise. Just a very, very, very small amount. So I don't even know if we'll be able to hear it. But if I put it all the way up – you can't really hear it.

[00:10:23] It's near silence. You shouldn't really be able to hear it. But it just is data basically that's filling in the gaps in other data, which can be really helpful. So if I zoom back out, I realize I zoomed in without telling you where the zoom in button in. It's right there. It's also Control One. But I can just grab this guy right at the top and then it would be helpful if my silence, you know, was in between these two, and that way I know I have a bumper between them and they'll never run into each other, because I put a piece of data in [00:10:53] between them that acts as like a book mark. So now I'm going to hop over back to my PowerPoint. Do we have any questions about recording straight into Audacity? I mean of course there's many nuances that go into podcast recording. But just in terms of the mechanics of the program, does that seem pretty straightforward?

[00:11:22] I'll wait until I see. But I'll move on until I see questions in the chat. Feel free to ask them any time and I'll answer. So levels. Shaping your sound to the best possible quality. So you can have as much sound as you want, but if it's very unbalanced, it's not very helpful to you. So a very, very brief science lesson. I promise not to take too long on this, even though it is pretty interesting. So the most common units of measurement you'll see when we talk about recording sound for podcasting [00:11:51] are Hertz and decibels. So decibels are what are involved in the levels, which is why I'm bringing it up. So decibels are actually a tenth of a Bel, which is named after a bell, and it's used to describe sort of the experience of hearing a sound.

[00:12:10] And it expresses a relationship between the intensity of the sound and the amount of area that the sound is covering. So when you're very close to something, you're experiencing a higher decibel than you are if you're further away, even though the sound hasn't changed at all. It's just traveled further or expanded across more space, so it's not as intense. And then Hertz is frequency. I think that one is something we're pretty familiar with. It's also called the pitch of a sound and it's literally just peaks in the wavelength per second, [00:12:40] because sound is a wave. It's many wavelengths. It's alternations of – it's alternating compressed particle experiences called compressions, and decompressed particular experiences called rarefactions.

4

You don't need to know that to use Audacity. I'm just a geek. But the relationship between decibels and Hertz is important to understand, that there is very little relationship, I guess is really what I should have said. It's important to understand that those two are very distinct and that Hertz doesn't really change. The pitch and the frequency of something doesn't really change, with the exception of – speaking of Hertz and frequency, I'm just going to close that.

[00:13:20] I don't know if you can hear, but there's an alarm outside my building. But so – so organic things, they're Hertz values may change, so my voice, for example, I've given a range of Hertz for my voice here. Because it's not always going to be the same. But it's always going to be a pretty similar decibel experience if you're at the same distance from me. If you're at a further distance from me, it's going to be very different. [00:13:47] So some things have consistent decibel values. For example, thunder. Because we're never experiencing it really any closer than we are now. We can't get much closer to the clouds in order to have a more intense experience of the sound. So it'll always be around 100 to 120 decibels.

[00:14:05] A fire alarm will always be a very high decibel rate because it's very close to where we are. That's the purpose, to let us know to evacuate. And it crosses the threshold of what we call the feeling threshold with decibels, which is around 125 or 130 decibels; you start to actually start to experience a painfulness from the sound because these alternations of compressions and rarefactions are physically affecting a series of little bones in your [00:14:35] ears called auditory [ocibles 00:14:36]. Again, not something you need to know to use Audacity but just geeky and fun. And at a certain point it becomes too much for them and that's around the area of a fire alarm. So A4, very common note, it's a reference note that's used often in sound editing.

[00:14:49] It's 440 Hertz, which means that it has – experiences 440 peaks in the span of a second. But it's almost never the same decibel rate. Because depending on what instrument is producing it, depending on how close you are to the instrument, something like that, it will always change. So while the frequency of a sound might not change, the decibels will always change, depending on how close you are when you're recording it and how close you are when you're hearing it. So it's important to think about decibels when we edit because you want things to sound [00:15:21] balanced and you want them to sound like they have a real relationship to one another, they way that they do in real life. It wouldn't make sense if you were listening to a podcast and the music was a lot louder than the narrator speaking.

[00:15:31] It would be very difficult to understand the podcast. So the decibels that we are going to see displayed on the screen in Audacity are threshold values. They're not exactly the same as measuring decibels of a

sound that you're experiencing in real life. Because those decibel values would be on the positive. But the decibel threshold of Audacity is zero. And so everything underneath it is – the threshold being the last amount that it can measure. So the most that [00:16:02] Audacity can possibly comprehend, it will call that zero decibels as its threshold, and from there you're trying to give it an amount of information that it can comprehend. So the absolutely loudest, anything that you ever record or have in your podcast or whatever soundscape you're creating should be at that negative one value. So that means it's one decibel under the maximum threshold that Audacity can understand.

[00:16:29] You don't ever want it to get closer than that. You probably don't even really want it to be there, but if you have an explosion in the middle of your podcast, that's the absolute loudest anything in your audio should be. Then underneath that, about six to 12 decibels below the threshold, that's the normal loudness for narrative audio. So for your main dialogue, your main narration, if you are analyzing a film or music, you might [00:16:58] want to have that at the same volume. Underneath that, so about six to 12 decibels beneath the threshold of your narration, then you'll have a range for like textural and transition sounds. So that might be the stinger for the network of the podcast, like the podcast network that you work at, or it might be something like bird's chirping in the background.

[00:17:23] You wouldn't want those to overpower the narration either. And then at the very, very bottom, you want a good 18 to 20 decibels between your narration and any background music or room tone. So now let's take a look at trying to actually balance those in Audacity. So we're not going to get into assembly yet. That's just the next slide that I have. OK. Back into Audacity. So here we have [00:17:52] my audio that I recorded. Then I had some white noise that I generated. And then I have this song, which I often use. It's the theme song for a podcast that I edit. So it's always on my computer. So right now we can see that this is all completely sort of like out of whack, I guess. We need almost everything to come way down, and we might even need our track that we like to come way up.

[00:18:22] So I'm highlighting our narration track. And I'm going to just – oop, I zoomed out. Sorry. Zooming in. Great. So let's listen to what this sounds like by itself. So I'm going to press solo and that mutes all the other tracks. And I just listen to this. OK. So I can see that it's between 12 and18. So it could probably stand to be amplified a little bit. Let's gather it. Let's amplify it just a little bit. [00:19:52] So amplification decibels, seven is going to be too much. Let's give it five. Because you can always add more. You don't have to do it all. You know. This – there were go. They look even chunkier. OK, great. So at most you can see I'm crossing into sort of like the yellow and orange and I'm not ever really crossing into anything red and I'm not getting anywhere near that zero value.

[00:19:28] So I'm happy with that and where it was bouncing around sort of on this monitor. So now let's address this white noise. This clearly needs to come down. Because it's in the same track line as my narration, which I just brought up, I don't want to do something sort of lazy and just drag the gain down and make the whole track quieter. Because that'll affect the other [00:19:59] things in there that I already carefully – carefully made louder. So I could highlight the whole thing and I could try to make it quieter using an effect. There's a lot in here. But I think I actually will do it with the overall gain slider, and I'll do that by making it a new track. So I'm just going to get a new track, add a new stereo track. I tend to type in and search for things rather than look through the menus.

[00:20:41] It's totally up to you if that's not your style and you want to go through the menu. That was under track, new track. And I'm going to bring my white noise down here. And then now I can turn it down all by itself and so the whole track is quieter there. I could also choose to do something called enveloping and that – when you use the envelope tool, when you click, it adds like data points of the levels [00:21:14] that you want the volume to be at, or the amplitude to be at. And then I can move those points down. So now my white noise fades in from essentially zero. It's still quieter. And it's still pretty quiet. Oop, there we go. So maybe something like that. Just like a cone of white noise that gently fades in.

[00:21:38] A time you might use this is for room tone. So there are a few different opinions about room tone. I was recently speaking with someone who also works in podcast editing and they said that they always adhere very strictly to gathering room tone and always putting it underneath or putting some kind of white noise underneath and always having it fade in at the beginning as the basis for the soundscape of their podcast. That's not personally how I do it. It's totally a preference. [00:22:05] It's up to you. For me, I don't mind acknowledging the form as much. So it doesn't bother me to have an opening that's sort of sharply bring you into the world rather than fades gently.

[00:22:18] But depending on the subject or the style of the story that you're trying to present with your podcast, that's going to be an artistic choice that you can make yourself. So I can bring my – I can bring my song down too. You can see that when I bring the overall gain of the track down, it doesn't lower the peak, so it doesn't give me a good visual of how these are equivalent. That's one thing to keep in mind. If you are someone who is really processing the tracks visually, you may find that [00:22:51] moving these gain sliders really tricks your brain a little bit and you are not realizing just how loud or just how quiet different things are. So let me go back to zero on both of these and just do it with the envelope. So just make sure that everything is shorter than my main narration track by a lot.

7

[00:23:19] Let me see if this mixes OK. I almost think that – OK, I can definitely hear the white noise and I can't hear the soundtrack very well. I happen to know that because that's because the beginning of the sound is really, really quiet. So I'm just going to get rid of it. Now let's try to listen to our mix.

[Audio clip]

Eva: [00:23:47] OK. I hear the sound coming up a little bit in the back. I'm going to go back to my enveloping tool and I think I'm going to make the – I just realized you all have to listen to my voice like 100 times. Talk about recording. Because I'm just going to use that same audio sample. All right, how about this?

[Audio clip]

Eva: [00:24:13] I would say that song is still a little quiet. I'm going to bring it up just a little bit more with my envelope. OK, now I can hear it. I don't know if you guys are hearing the exact same thing that I'm hearing, but I can hear the bells, the sound in the background, and I can definitely hear the white noise. Great. So what would be one manner of like balancing and you could choose to do them all in different tracks. You could choose to keep them [00:24:42] in the same track and only use that enveloping tool. And just keep in mind that any time you can always grab the top and shift them around like that. Which you didn't used to be able to do. So that's one of the Audacity features. And you can also now turn them like that, which you couldn't do before.

[00:25:03] But let's talk about assembly a little bit. So I did just do a little bit of assembly while we were doing that. Like I said, it's not going to be like perfectly linear, just because sometimes you have to go back and forth. But if I go to my narrative track, finish my outline – I guess I didn't write anything about that. But basically you wan to make sure earlier in your production process, which I know you guys have been learning a lot about podcasting. [00:25:31] You've gone to all these workshops. But podcasting, you can really save yourself a lot of work later on if you do a lot of work upfront. So I always try to have an outline for the audio that I'm going to be editing.

[00:25:44] Whether that's a script that I already have or it's just a series of bullet points, or it's just a three act narrative structure that I've like designed, you want to make sure that you have something, like the bones of something that you can start fitting these tracks onto. So the main part of your storyline, whether that's dialogue or narration, is going to continue to follow your writing. So the more writing you do, the easier it is to record. The more recording you do, the easier it is to edit. And it sort of, you know, funnels like that. So if you don't spend enough time writing, [00:26:15] and then you sort of haphazardly record, you will have a lot of editing to do to make it sound even sort of normal. So that can be frustrating. So if we go back to our

8

desktop, also, again, any questions, just feel free to pop them in the chat. If we go back to our desktop, if I was really assembling this to the vision of a storyline, I would probably have the music fade up – I'm going to get rid of my white noise because that was just an example, and we don't need to hear white noise anymore.

[00:26:54] I would probably have my music fade up in the beginning and knowing that, I'm going to move my audio to actually start about here. And then I have my extra silence in the beginning as a little buffer again. This is where the silence really comes in handy. Because you really want to have something between multiple audio tracks so that you are never running your narration into itself and so that you're giving it like [00:27:24] a nice amount of time to breathe so that people can understand what you're saying. Because podcasting is often – podcasts are often listened to by people who are like in the car doing their dishes. It's a very different form than a book. You can't just go back to the top of the page and re-read the last few sentences.

[00:27:44] Sure, you could, you know, take your soapy dishcloth off, walk over to your phone, and tap back 30 seconds to re-listen to something you already heard. But most people aren't going to do that. So you want to make sure that you're really making it clear either as many times as you need to or the only time that you do, you're really making it clear. So you want to have these silences where just sort of the information that you've imparted has room to breathe. So let's take a look while we're [00:28:13] assembling and we're working on our narrative track.

[Audio clip]

Eva: [00:28:22] There we go. There was one. So that was the noise of me breathing right there. I'm not 100 sure, you know, you might recognize it if you're not me. I edit my own voice all the time, so I always recognize this little peak as being the sound of me breathing. So the thing about breath is that it often naturally happens in a place where it makes sense grammatically or narratively. We sort of learned, you know, from like socializing with people over a long time or [00:28:52] however long we've been around, when it's appropriate in the clauses and phrases and of our sentences to pause long enough to take a breath. So it's always a good idea to take a look at where those breaths are and see if maybe instead what they need is a small pause. That could be a time to switch to another piece of audio that you have, like an interview.

[00:29:15] It could just be a few seconds of silence. But you almost certainly want to at least silence the breath sound, depending on who you're recording. Some people it's very heavy and obvious. I personally, when I'm editing my own voice, it bothers me to listen to. So I'm just going to press delete and goes right away. But if I just press delete, so does that space, so does that perfect

natural punctuation that follows the, you know, dialectical patterns that we're used as listeners and speakers. So instead what I'm going to do is I'm going to use this tool, which is silence audio selection. So all it does is the place that you highlighted, it replaces it with that silence data, which is technically near silence data. So now what we hear instead of me breathing …

[Audio clip]

Eva: [00:30:15] Oh, I have it on loop here. Sorry. Don't need to loop that. These little guys at the top, that's the section that will repeat after you play it. So there you go. So it went that time, instead it went straight from the first thing I said, it had an appropriate gap, and then it went to the next thing I said, which is perfect. I can see that that is another breath of mine. Eventually with editing you might get used to that. [00:30:47] So I'm going to silence that one too. I can just tell by the way that it looks. Because I see my own breaths in my audio editing all the time. So maybe in that case that would be a time where I would just snip it and – oop, pardon me. Just going to zoom in a little bit. So instead I'm going to just split it and that you can do with either right clicking on it or clicking Command I, and it splits it.

[00:31:24] So now I have the section that I highlighted, the section before what I highlighted, and the section after. So if I decided that after that breath, that meaningful pause in the dialectical pattern, I wanted to pull the rest of this away, and I wanted to have my other narrator or a piece of dialogue or a sample of news radio broadcasts from the time of the event, that I'm writing my podcast about, maybe I would have that there, if it completed [00:31:52] well. So if I wanted to do that – let me just – minus, minus. A couple choices. You could either insert the audio right in there. So this is just a noise. And when you do that it'll like default to making its own track down below. So then you would, you know, bring it back up.

[00:32:21] Won't need that track anymore. And maybe I want that here after the silence and then I'll cut back to my original narration like that. I can see that this is a little shorter in peaks, whatever this is. In that case it's a sound effect that I use often. So I'm going to amplify it just a little bit, maybe like three – again, you can always go back and amplify it again. [00:32:49] It's a lot harder to un-amplify something you've amplified. So with editing and effects I always like to use a gentle touch and sort of like nudge it forward towards where I want it, instead of making big leaps and bounds and then having to use like Control Z a million times.

[00:33:05] There we go. So they look pretty similar. And now we have it all strung together over our sound. I'm going to listen back to it. It's going to be confusing. That's the sound effect.

[Audio clip]

10

Eva: [00:33:24] So that would be how it cut together. It did sound like the levels were all pretty consistent throughout. So that's great. Another thing I might do while I'm assembling my narration track is to do noise reduction. This is a really common effect that's used and we'll go into a couple other effects as well. But let's just start with this one and we'll get to the rest of the effects when we get to them. So noise reduction. It involves two steps. You want to [00:33:54] highlight it. You want to highlight the noise that you need reduced. And then you go into effects and – pardon me, I'm going to start over. First you want to select the noise you don't want. So not the entire track. Just the noise you don't want.

[00:34:17] So let's say I found this background noise. Maybe that was, you know, my air conditioner resetting, or a car, or a fan that's running; something that's relatively consistent and it happens in a lot of places in a track. I would go into Effects, Noise Reduction, and I would click Get Noise Profile. So now it's assessing those frequencies and it's going to – now it understands them, and then I'm going to highlight everywhere where that noise [00:34:46] might appear. So this whole audio track or this whole clump. Go back into my effect and go for Step Two. So how much noise do you want filtered out, and OK to reduce noise. So these three things it gives you. The third one, Never Change. Three is just fine. You don't need to change it.

[00:35:17] This is how much it will carve out. So any time that it finds that sound, how many decibels will it reduce it? The reason that mine opened like this is because these are the values that I always use first, a very small amount, and then maybe later I go back in and do a little bit more. But maybe I want it to cut out just a bit more than that. I want it to cut out five decibels. The sensitivity is how, yes, how sensitive. [00:35:44] It's like how high the standards are for finding that noise. So if something is just a little bit like that noise, is it going to cut it out or not? So that is as much as it can be. It's like, oh, if it's just a little bit like that noise, I'm going to definitely take it out. Whereas 2.5 is more like it has to be pretty darn close to being that noise, and then I will take it out.

[00:36:11] And by I, I'm speaking as the computer. And then, OK. And as you saw, sort of, I'm going to hit Apple C so it un-does, and then Control Shift Apple C. It's Command. It used to be a picture of an apple in that key. And it'll go up and down and you can kind of see that it has pulled a lot of those smaller things off. So that's really helpful if you recorded your narrative somewhere where you couldn't really control the background noise. Like I said, the more time you put into recording, the less time you'll need to edit. [00:36:41] So ideally you would just record somewhere where you had a lot more control over the environment. But that doesn't always happen, especially in New York. Obviously my apartment has background noise, you can hear.

[00:36:57] So let's – now that we've put together sort of like in a row, a few clips of narration and one other clip, I'm going to talk more about putting together multiple types of sounds. So that would be layering, right? So layering sounds is basically like when you want to build the world around the narrative track. And so the narrative track is your dialogue, your own narration, whatever is the main information [00:37:29] that you're conveying. Usually it's a human voice. But I don't know. It's your podcast. Maybe you're going to do something really avant garde. So when you're building a soundscape, you want to consider your main narrative sound, which is dialogue, narration, [and that's 00:37:45] the story.

[00:37:48] The background sounds, which would be room tone – so anything that's constantly part of the recording environment. So like your AC system, it's usually that. I don't even need a second example. It's almost always the HVAC. Every single room tone issue is the HVAC. And then any background music that you're adding. So room tone is a background noise that's already there. Music is background noise that you're adding, that's constant. That's what those two have in common. And then texture, which are sounds that bring life [00:38:17] to the world of the story. So that might be like a bird chirping, it might be a car going by. It might even like a subway. If I were going to record a podcast about my neighborhood – I live pretty close to campus, right where the One Train goes above ground at 125th Street – I might even open my podcast with the noise of the One Train when it goes above ground at 125th Street.

[00:38:42] And that would be a transitioning sound. Because it informs the listener that now we're moving to either a different speaker or were in a different place. And what's really fun is that you can also interruptions, which is what I would call sort of the sixth significant type of sound, you can as an editor, you have a lot of power of authorship, to use interruptions as transition sounds. So something like a train going by, a dog barking, something like that might sort of snap you into the world of somewhere else that – somewhere other than the first scene [00:39:18] that you were in. So if you were narrating about a family, and then you cut to your interview with the mother of family, on the barking of the dog or on the phone ringing in the house – phone ringing I used as the icon for the transitions because it's a very common transition sound.

[00:39:37] If you do interviews over the phone, I highly recommend recording the sound of the – of you waiting for them to pick up or of them calling you and hearing that. Because it's a really wonderful noise that lets the listener know and introduces another character kind of. It's like, OK, here comes the next character. They're calling on the phone. I can't wait to meet them. And also on the other side, sort of in an effort to like [00:40:07] ameliorate the amount of issues that you have, allowing the phone ringing sounds to happen before the audio of a phone interview does wonders to excuse the audio issues

that can happen when you have a phone interview. So a lot of times that sound is very compressed. It actually cuts off a lot of wavelength values. Many of – Hertz – many of which are values that humans speak at. So the phone audio is just not as good. It doesn't have the richness and the fullness of true human speech.

[00:40:43] So adding that phone sound before your interview sort of clues your audience into that fact and it allows that lack of quality to actually be part of the character of the piece instead of just a mistake. So that's like a fun little part about that. So I just highly recommend that. So let's go back to Audacity now and do some effects. I'm going to return to my Audacity share and we're going to look at some things we could do the voice. [00:41:15] We'll just do this for a couple quick minutes and then we'll export our files and then I can do a few questions if you have any. But if this is me speaking …

[Audio clip]

Eva: [00:41:32] Great. Let's do some really weird stuff. So something you can always do is you can change pitch and tempo so you can make – you can stretch it out to be a lot slower, or you can compress it to be a lot faster. Let's change our pitch. So maybe I want to make a sound that sounds like I'm haunting on my own computer. So I speak at a B3 242 Hertz [00:42:00] according to this. That's great. I must have been really calm. My voice was really low. So let's see. So I'm going to do it by semi-tones, so that's like a half a note. And I'm going to do eight semi-tones down. There we go. We saw some of the loss there. So some of these effects are what are called [lossy 00:42:32]. It's sort of a strange word that's only used in audio editing and it means – it has to do with when you're processing a sound how much of the clarity, richness and fullness of the sound do you lose.

[00:42:44] So telephones, very lossy, the way that I was just speaking about before. Some of these effects a little bit lossy.

[Audio clip]

Eva: [00:42:54] So I sounded very edited in that. Like it's almost not low enough to be believably a deep voice. Instead it's like just enough to be sort of uncanny and sound like maybe I'm using a voice changer, which I am. Something like that. Another one that's pretty fun is reverb. This can be used to give sort of [00:43:24] a grandiosity to what you're saying. A lot of podcasting network introductions have reverb on them, and I promise they're not recording their introductions in a cathedral. They're using a reverb effect to make it sound as though – there's just something about that sound that invokes sort of like a common social memory or understanding of the importance of having, of being a person who has access to a giant fancy room that sort of like adds – yes, [00:43:56] it adds value to the speaker.

13

[00:44:01] So a lot of those like stingers at the front that are like, "You're listening to a brand podcast," those have like a little bit of a reverb on them. So a lot of different things that it gives you here. Mostly the room size and the milliseconds that you want it to delay are going to be the only things that you want to affect. So just a little bit of a bigger room and a little bit of pre-delay. That's fine. I'm not even going to turn the reverb down. [00:44:30] Wet sound and dry sound, don't worry too much about it. If you are someone who's a real nerd about sound and you were like, "Oh, that decibel thing was so interesting," you can look into it. Basically it has to do with how much sound was actually measured and how much sound was added by the software as a guess for filling in gaps that it didn't hear.

[00:44:51] There we go. We saw it lose a little. Now …

[Audio clip]

Eva: [00:44:57] So fun. OK. It's because it's Halloween, so I like want to change all those things. So that's effects kind of in a nutshell. There's a lot of other ones. You can compress things. You can distort things. You can add an echo. Things like that. Cubic Curve is really fun. It creates like a very distressing sort of sound. I guess I really am just in the Halloween mood. I keep picking all these creepy effects. But if you were doing a narrative fiction podcast, that would be like a really fun thing to include. [00:45:31] I produced like a radio play reading of an Edgar Allen Poe story and one of the characters, the description of his voice was that it sounded as though like thousands of voices were coming out of his mouth, and I used a combination of odd harmonics and reverb and just track layering with like a slight millisecond delay to create that.

[00:45:53] And it was really fun. OK, let's save our project. That's – none of this will matter if you don't save your project. So something important about Audacity. Saving your project is not the same as exporting the final sound. So when you save your project, it's going to save – see, it gave me the warning. Great. So I'm going to save this as my Demo Project 1029. And I'm going to absolutely make sure that I am saving it in – so I always [00:46:25] save things in a folder I have called Audio Editing, and then the Workshop Demo 1029, Audacity Project. That's how many layers of folders I have. I really, really, really suggest you do the same. It might seem like overkill, but nothing is more frustrating than losing files because you were disorganized or because you were in a hurry.

[00:46:45] Just make a bunch of folders. Label them very clearly, put dates on them if you can. So that's where that went. And so now if I were to export it – oop, I skipped right over that. File export. You have a few choices. I never use OGG. I do use MP3 and WAV or WAV form. WAV form files, they're very, very big. They have a lot of clarity, a lot of depth, a lot of fidelity. However,

they're not accepted upload formats by almost any distribution or [RSS feed 00:47:17]. So that would be a – you want to record things as WAV files but you don't want to export them as WAV files, unless you're not putting them anywhere and they're just for posterity. WAV files are much larger so they're just not accepted by anywhere that you can upload podcasts. Instead I would export it as an MP3 if I were going to upload it, and here it is, Demo Project 1029. I'm not going to push a slash in the file name because that can create problems when you're moving to different things. Audio editing, workshop demo.

[00:47:54] And then I'm even going to create another new folder, and it's going to be called Exports, in the Workshop Demo folder. And that's where I'm going to save it. Quality, it's up to you. Standard is just fine. A lot of places that you're trying to upload to, if you try to give them a much bigger file that's like intense clarity, intense quality, they won't even take it. So standard is just fine, especially for something that you still need to – yes, that's what I want. Those for something you might work on again. [00:48:25] This you can kind of ignore. That's like if you're making music. But this is for podcasting so we are doing it. So there you go. It's exporting the audio and then I can show you in Finder – oop – that it is in there. Editing, Audio Editing, Workshop Demo, Audacity Project.

[00:48:45] And you can see that it comes with the Audacity project file extension and then there's another one. You have to keep these together in the same folder or you absolutely will lose pieces of the project. And I recommend keeping it next to all the raw media that you used as well, because sometimes even if you feel like you're 100 percent done, and you're like I definitely am not going to need those files again, [00:49:15] you're going to need those files again and you're going to be really, really sad that you forget them. So exporting, move all of it. Any time you're moving it, even a little bit, you pick up that entire folder and you know that it's very important. So that's actually all of the slides that I have.

[00:49:37] I was hoping to end a little earlier so I could take more questions. But is there anything in particular that you guys wanted to ask me or that I could go over again that you would find helpful? I really like playing with Audacity, so I have plenty of time.

Michelle: [00:49:59] One question that we get asked a lot personally, because we have a hard upload limit of 25 megabytes on the various podcasting websites, and so we get asked questions about what is an acceptable rate to compress rate to compress your podcast to on export, to preserve a pretty good sound quality while making sure that the file is small enough to fit into the application.

Eva: [00:50:25] That is a really good question. Let me take a look at – 25 megabytes, you said?

Michelle: [00:50:34] Yes.

Eva: [00:50:35] I'm like a little bit – I'm not like, "Oh my gosh, shocked." But a little bit, mostly because I – I'm just going to look at what my most recent export was. Yes, my most recent export for the last podcast that I produced was 62 megabytes. So that is pretty significant. I guess I would say that it's better – I guess I have kind of a non-answer for you, which is that I think [00:51:05] it would be better to do two parts of one episode, have them both sound great, than to compress one episode. So that would be my recommendation, because I get really sad when I listen to something and I can hear that it's missing like the fullness of human interaction, like the voice. I also feel like it can take audiences out of it.

              [00:51:29] So I try to preserve as much fidelity in the human voice part of the podcast as possible. So I might even do something like make it much shorter or I might use lower fidelity additional sounds. Just so that I can like fight to preserve as much of the voice as possible. That would be my answer. Sorry.

Michelle: [00:51:52] Well, it's a limitation that we have because we are using WordPress to do our hosting and we have that small problem. But I love those suggestions about places where you can kind of cut some corners to preserve what's really important about the sound.

Eva: [00:52:06] I also have a recommendation if you want to change doing that and you want to instead embed a podcast player from somewhere else into WordPress, which is what I do. So I actually use Podbean to host, which I can't remember what the upload limit is, but I think it's certainly much more than 25. It's like a limit per month of uploading for the free account and then you can pay for a paid one.

[End of recorded material 00:52:31]