

Podcast Essentials - Adobe Audition

[Start of recorded material 00:00:00]

Monica: [00:00:07] Hey, I'm Monica Hunter-Hart, I'm a recent graduate from Columbia's School of International and Public Affairs in the Journalism School. I focused on audio production and podcast production when I was at the Journalism School, and on Monday I'm going to be starting a position at [Meekay Asia? 00:00:22] helping them to build their first podcast from the ground up. I love podcasting, audio production, and I'm just really excited to be here with you all.

[00:00:33] So, today we're going to be talking about Adobe Auditions, specifically. There are a lot of different digital audio workstations or DAWs that you can record podcasts with. There's Audacity, which is wonderful, because it's free. There's going to be a tutorial for that later in this series. There's also ProTools which is the industry standard. And then there is Adobe Audition, which is also very commonly used in the industry. It's cheaper than ProTools, much cheaper, and it's also typically considered easier to use than ProTools. It's a really great option for when you're starting out in the world of podcast producing. And that's why we are focusing on that today. It's also the one that people typically use at Columbia. So if you're a student, that's probably what you'll be using.

OK. So, I'm going to share my screen. OK. Can you all see Audition right now?

Speaker 1: [00:01:35] Yes.

Monica: [00:01:35] Thank you. OK. Excellent. So, just gonna minimize this. All right, so, basically, let's just talk about setting up the audio interface. So, just want to save this, OK. So, when you're first starting out, you want to make sure that you have your output and your input all set up ready to go. So you're going to go up to Adobe Audition up here, preferences, and then audio hardware.

[00:02:09] Can you actually, sorry, Michelle, are you able to see all of these things when I'm doing it? Or can you only see part of the screen?

Michelle: [00:02:13] I was just wondering but, I think that it's a little bit cut off on the sides. And it also, like, once you'd clicked preferences, I could see the modal, but I wasn't able to see that menu up at the top when you started out.

Monica: [00:02:28] OK, well, in that case, I'm going to try sharing my entire screen. So, just give me one moment, sorry. Let's see. OK, so, just going to come back here. Share the entire desktop. There, we've got Audition right here. So, if you go up here, to preferences, audio hardware, you just want to make sure that your default input and output are [unintelligible]

00:03:16] here. So, right now, we're looking at basically, my sound is coming out of Zoom, so Zoom Audio Device. Default input is just the built in microphone. I'm basically going to assume throughout this tutorial that we are not reporting directly into Adobe Audition. You definitely can, and I'll show you a little bit of that, but in general, at Columbia, people use task cam recorders or other external recorders and then physically kind of manually input the files that they've already recorded. So I'm just going to [interference 00:03:54] just the built in zoom audio device.

[00:04:00] So, great, we're good to go. Saving that. Stop share for a moment. The sharing might be a little clunky-er than I had hoped. And I apologize. I also did want to mention, that if you do not have access to a fancy external recorder like a task cam, you can record on a smartphone, which definitely makes recording, a pretty acceptable quality overall. Either an Android or a smartphone, you just want to hold the smartphone about at four inches-ish, which is about four fingers length away from your mouth, and ideally, at some sort of an angle, like this, or like this, so that you're not breathing directly into it and you're minimizing plosives in your speech. Plosives are hard consonants like P's and T's and K's.

And then, if you do that, then it actually works pretty well and that's actually just something that [00:05:00] often people are instructed to do if you are an interviewee, and you are recording with a podcast professional then you are told to use your smartphone to record yourself and then send the audio file in, so it definitely works. Yes, definitely works if you don't have task cams.

So, to Audition here. So, overall, you have two basic work views. You have the [waved form? 00:05:29] view, and you have the multitrack view. So, let's just save a session here. Let's call it Adobe Audition Tutorial 1022. Great. So, let's minimize this a little bit and make this a little bigger. So, this is the basic work view that you're going to be looking at most of the time here. You have your tracks here, you can always add more. I'm zooming out a little bit. But this [00:06:00] is your general, these are your list of tracks, over here you have the files that you're inputting, and let's see. Now, anytime that you're making a change to an audio file that you've imported in this multitrack view, it's not actually making a change to the audio file itself. So that's really important.

So, basically, if you're in this wave form, and you're recording directly into it, let's also just, let's call this Adobe Audition Wave. Let's just do a little quick recording here. Just say something like hello everybody, thank you so much for being here today. Stop. Now, if we want to listen to that – so, we want to come back to the multitrack, we now have this waveform over here, and we can pull it [00:07:00] in.

So, if you're in the waveform view, we're taking an up-close look at the audio file itself. And if we want to make changes in here, everything is going to be permanent. We're affecting the original file. If we don't want

to do that, then we want to make sure that we're making the changes here in the multitrack tool. So, basically, if you're looking at the multitrack view, you're basically Adobe Audition is drawing on the data from the audio file, rather than the audio file itself, and I'll talk a little bit more about that later.

OK, so at any time, you can just switch between these two views up here on the top left. And if you double click on a file, then you're going to go automatically into the waveform view. OK. So, let's go back to just the recording for a second. So if you go up to file, new, audio file, uh, all right, so, [00:08:00] we're actually going to change this to 44100, that's just typically what we use. So, just call this again Adobe 2. Stereo is if you are trying to – it's a little more complicated. You're basically getting a surround sound effect, and it enables panning, so that you can put kind of position different audio files around a pan. So, it sounds like you're listening to, for example, a lawn mowing on the right, and then maybe a dog barking on the left. I don't recommend that in general, just because it's just not necessary for basic podcasting, unless you're really trying to set a scene and you want a listener to feel fully immersed. But again, it's just generally not necessary.

So, we're going to do file name, sample [unintelligible 00:08:54] and then pitch 32. Let's just do one more little recording here [00:09:00]. Testing one two three. Great. And just make sure it's saved. Come over here and we can, oops, there we go, pull that over here. So, perfect. So sometimes you're going to get a message that comes up like this. It says the sample rate of the inserted file does not match the sample rate of the session. That basically just means that the 44100 hertz rate that we set the session to is not actually what this is recorded at. And if that does happen, then you can just click OK, and it will automatically make a copy that's at the correct hertz rate.

It's only if you're recording at like a really different hertz rate. A very, very low one, or a very, very high one when it would actually make a difference when you're copying it over. So, you can just hit OK and that will be fine. All right. Great. So, let's see, yeah, you have some basic controls down here. You know, play, fast forward, rewind, record. [00:10:00] This is pretty useful, loop playback. So if you want to listen to just a particular part of your recording over and over, and kind of make up a minute edit then you can just put that on. So when it's blue, [unintelligible 00:10:14] is selected.

If I want to, let's, so, actually, let's say I want to loop there's something, like, here, I can click I on the keyboard, and then click O, over here, so, I for in, O for out, and that just selects a little portion of this, with the play head somewhere in here, and then you can just hit play once the loop is selected. And then you can just [crosstalk 00:10:47] yes. And so, this is, as I just mentioned, the play head, this is basically just your cursor. But

wherever the play head is, is where the playback will start, whenever you hit the spacebar.

[00:11:00] OK, good, good, good. All right. So, let's see. If you are in the waveform view, I think this is maybe the last thing I want to say about this. You can actually highlight a portion, and then just actually delete it. I don't recommend that, again, because you're permanently deleting it. And then, in the waveform view, you can also just do the same sort of highlight, loop, zoom in, zoom out, [unintelligible 00:11:33]. But I think we're pretty much gonna stay in the multitrack from here on out. All right. Down here you can see the timing of the recording, the timestamps. So over here, now you're at zero, zero, as you move on, you've got two seconds, four seconds, etc., and then the exact time that's down here, but we don't actually need [unintelligible 00:11:52] minimize it so we've got a better view.

[00:12:00] So, yes, so, as we just said, this is basically it's all operating chronologically. So if I put this down here, it's layered on top of each other. You're going to hear both at once. Which is not what we want in this case, but let's just listen [crosstalk 00:12:13] one really great tool is this magnet up here. It's called toggle snapping. If you turn that on, then you're able to have the audio that you're moving next to another clip just automatically snap so that it fits right at the end there. So if you don't have it on, then, I hope you can see this. Let me zoom in just a little bit more. If you don't have that on, then you have to be as precise as you can be, and really zoom in if you want to be right next to it, right like that. But if you do have it on, [00:13:00] then it just automatically does that for you. So toggle snapping can be great, unless you're not actually trying to have that [crosstalk 00:13:05] right next to each other.

All right. So, a couple notes on zooming, and all right, so, zooming can be done just on your keyboard with the classic shortcuts option plus or option minus on a Mac. You can also do that down here [crosstalk 00:13:28] I think we might have another if you could mute yourself if you don't have a question, that would be helpful. OK, and then you can click this button down here. This is very useful. Zoom out full. So if you are zoomed in, and then you [unintelligible 00:13:55] so you can see all the tracks at once and each of the tracks in their entire lengths, you can [00:14:00] hit that, and then suddenly they're looking at the whole project at once. Very useful.

OK, and then this is yet another way to be zooming in and out. You've got this little toggle bar at the top. You can drag the edges of it, and then you can also drag it from side to side if you want to quickly move throughout your project. And this is just really useful if you have a long podcast for 30 to 40 minute podcast, then if you're just trying to scroll all the way through it, it's quite tedious, so, this can be very useful.

So, yes, we just recorded a little bit of audio directly into Audition. But, let's use the method that I am pretty sure most of you are going to be using, which is bring in audio manually when you have already recorded outside of Audition. So there are two ways to do that. So, let's pull this over here. So this right here is the input audio file button. [00:15:00] That's the way that I typically like to do it. I'm going to do this. Monica pre-recording. All right, so we've got this over here. Yeah, you can just literally hit this and then just find it on your computer. There's also a media browser option down here, where you can also go through and find it. This is just for me a little bit more annoying, because it's, you have to sort through all the folders, and I've never been very good at that. But the nice thing about the media browser is that if you do have the files listed down here, you can hit play down here and then test it, so you can make sure you're inputting the right file.

But either way works. So, let's see. I'll minimize this again. And zoom out just a little bit. And I'll pull in this other recording. So now I'm going to mute this top track, so that we can just [00:16:00] hear this track. Great. Another thing you can do is solo a track, and that automatically basically mutes the other tracks, and then makes it so you're just hearing the solo track. And if you want to solo two tracks at once you can, and then you'll only be listening to those two tracks.

So you've got, yes, M for mute, S for solo, R for record, so you can also record directly in the multitrack view. You can just hit record, make sure that you're selecting the track that you want to record onto, and then hit this record button down here. But, let's not do that. OK, so here we go, we've got this new file that I already imported. As soon as you do import a file, whether it's a file that you've just recorded, or you're importing a previously existing file, they're all going to appear [00:17:00] up here in this list. So this is basically your list of files that you're using currently in Audition.

OK. So, let's see. All right. So, I also just wanted to make sure, we already saved this at the beginning, so I just want to make one quick note before we get too far along to just say that when you are saving an Adobe Audition file, it's saved as a dot SESX. That's just a session. An Audition session. And every time that you open a pre-existing Audition session, Audition is going to be drawing on the original audio files, and that's not what you have in here. This is just data that's about the original audio file. It's not the original audio file itself unless you are going into the waveform.

[00:18:00] Basically, you want to make sure that you're saving all of your audio files and the Audition session, the dot SESX all in the same place. Luckily, Audition does already save copies of the files that you're inputting in its own folder as soon as you save a dot SESX folder. But, sometimes, I just, you know, have forgotten to copy a file over, or I lose it or something happens, and I just personally like to make sure that I'm

saving them all in the same place, because audition needs to know where to find those files.

So if you're moving the session at some point, or you're moving some of the files, it might not know where to draw them from. You might get some sort of message popping up that says you know, Audition can't locate this file, and then you might need to kind of go [unintelligible 00:18:55] to that. So just basically [00:19:00] make sure that you're saving [unintelligible 00:19:02].

So, all right, so selecting clips, just gonna minimize this a little bit. Yeah, selecting clips, you can just do by literally clicking on them. If you want to select more than one clip at once on a Mac, you can just hit command, and then you're getting multiple at once, I'm pretty sure on a PC it's just the shift key. Quick side note here. Editing in Audition becomes a million times easier if you use keyboard shortcuts. Sometimes there isn't a pre-existing shortcut for something that you want. So if you find yourself kind of up in the menu bar a lot, then you probably want to be adding some shortcuts.

So, let's see, let me stop the share now and share the full screen. [00:20:00] So, let's see. Edit keyboard shortcuts is the way to do this. Open this edit folder. And here you can see all of the keys that already have shortcuts associated with them. So, I've already mentioned the I and the O for example. But then there are a couple of shortcuts that I really highly recommend you add. My favorites are this and this. So, they don't have to be W and E, that's just what I did. But, I have Was select clips to end of selected track, and I've got E select clips to end of the entire session.

Let's close this here. OK. So, let's see, so if I'm going to hit W, if I select this track and [00:21:00] hit W, then that's going to select all the clips that are in this track, from the play head to the right. Super, super useful if you've already edited things to the right to be right next to each other, they're exactly where you want them, but, then you realize, oh, you have some dead space in the beginning, and you want to pull them all together to one side. Or, E, let's see, yes, and you select all the clips in the entire session to the right of the play head, and then you can drag those all in together.

So, again, this is really great if you have dead space at the end, dead space at the beginning, or a couple of seconds in the middle. Anything like that. Two really, really useful commands. You can also just get there by going to edit select clips to end of session or start of a session, but, it's just so much easier to do keyboard shortcuts. And if you're editing a lot, then it's really going to save you time to have your favorites kind of just easily at your fingertips.

[00:22:00] So let's see. Talking about that [unintelligible 00:22:06] so yeah, as soon as you've selected a clip, you can just literally drag it around to wherever you want. Very easy. Now, we should talk about cutting clips. So, let me zoom in a little bit. So if you want to just shorten the end of a clip, you can hover your mouse right over the end, until you see that little sign. That little bracket, and then just pull it to the side. So that's great if you're just trying to do something on the edges. But oftentimes, you're not wanting that. You're wanting to cut out something in the middle.

So, let's say I wanted to cut out some of this. Yeah, let's say I wanted to cut out the thank you so much for being here today. So, there's a few different ways [00:23:00] that we can do this. You can use the razor tool. So, to activate the razor tool, you can either click the R key on your keyboard, and you'll see that this little razor guy just popped up here. And then you can literally just click where you want it cut. You could also click the razor yourself up here. And if you click it twice, then you're actually going to be cutting all tracks vertically up and down from the point that the razor is. So, even though it looks like I'm cutting down here, it's actually cutting all of them at once. So that can be useful, but you just want to be careful that you have the right razor on.

So, you can use the razor. I actually don't love to use the razor personally because when you are [unintelligible 00:23:55] with a touchpad or a mouse. [00:24:00] So, you know, if I'm kind of, I might [unintelligible 00:24:03] my hand a little bit accidentally and then it's just not exactly where I want it. I prefer personally to just drag the play head to exactly where I want it and then when I'm totally sure, you know the play head is not going to move on its own, then I can do, let's see, I can go up here, edit, clip, and then split. And then it splits the clip right there. Or you can just do, if you're on a Mac, command K, nice and easy. So either way works. There are about a million ways to be splitting this. So let's just, let's get out this thank you so much for being here today.

So I'm just going to make sure that the track that I want is selected, pull this right to the end, command K for selecting, and then [unintelligible 00:24:57] to just before [00:25:00] I say the thank you. All right, at this point, you can actually see the waveform, and you can see where the [makes noise] in the thank you starts. And so you can pull that right down to about there. Playback, the play head there, command K, and let's zoom out, let's take this and just drag it and let me just put those together and see what that says.

All right, yeah, so, it's a pretty abrupt transition, so, I would maybe want to fade the end of that, or add a breath there or something to disguise that. Yeah. In any case, I wouldn't probably actually want to be cutting out the thank you [unintelligible 00:25:57]. But that's how you cut [unintelligible 00:26:03]. OK. So, let's see. I also wanted to just draw a little bit of attention to this track down here. This is the kind of master track. Looking at the entire mix at once.

So, it's basically Auditions output, where all of the tracks are mixed together. So you can make changes to the whole mix at once here if you like, if you're going to add effects to this track, but, it's a little bit more advanced, so we're not going to get into that today. But in case you're wondering, that's what that is. There's a default six tracks when you open up new session, but we can easily make a new track, just [unintelligible 00:26:54] [00:27:00] insert, OK, [unintelligible 00:27:08] just a second, but in any case, yes, you can make as many new tracks as you want. There's no limit.

OK if you want to group clips together, you can highlight them together, and then right click and then scroll down here to groups. Group clips. And then that, you can see that they've changed color. And now, you can drag those together. So let's zoom out here. Yeah, so you can see that those are now highlighted together, and then, when you drag them around, they move consecutively. So, that's great. You can also ungroup them by doing the same thing. Great. And now it's back to normal.

[00:28:00] It's also really useful to name tracks, just to keep track of things, so you could say, for example, you know, narration, interview one, that's just as easy as double clicking here and then just adding a name. You could also add track colors. That's really helpful for grouping things. I often like to have for example all the different takes of my narration in one color, and then maybe all of my interviews with one person in a different color. You can't always put everything onto the same track. Sometimes you are trying to work with different takes, and so yeah, color coding can just be really useful.

All right, so, let's talk a little bit about [unintelligible 00:28:54] so, there are a couple different ways to look at volume. This down here [00:29:00] is the level of the entire master. So, that's the, when all the tracks are together, what is the volume that's being outputted. It's measure in dB, or decibels. So, it goes all the way up to zero here. You never want it to get that far. Once it's at zero, the audio is starting to clip, and getting distorted. That basically means that the sound is overwhelming the system, and going beyond the maximum possible amplifier level. And you'll know that that's happening, because the meter will turn red. So, let's just listen to how loud this is.

All right, so, yeah, so that didn't get all the way up to zero, but it got pretty close. So that's a little too much. We don't want it to be quite that loud. So, we want to adjust that. I'm just going to make this track a little bigger here, and zoom in a little more. So, we can actually just drag this circle right here. There's one on every track. And we can just drag that to the right if we [00:30:00] want it to be louder, and drag it to the left if we want it to be softer. OK, so, let's drag it a little to the left, see if that helps us avoid the red. It's a little better. Let's do it a little bit more. Yeah, there we go. OK, so that's a pretty healthy level.

In general, you want to be aiming for between negative 12 and negative 9 dB. If you are too quiet, then it's just not going to be comfortable for the listener. They're not going to be able to hear what you're saying. So that's obviously something to avoid. But you're also losing some sound quality when it's too quiet. And if it's quiet and then the listener increases the volume, then you're going to get a distorted version. So, yeah, so again, between negative 12 and negative 9 is generally a good thing to shoot for.

So, if you're just going to look at how much [00:31:00] just looking at the volume on one particular track, then you can actually just look right here. This is the volume meter on an individual track. Right now, we have this one track soloed. So, it's going to be the same as this master level down here. So let's see. Yeah, so again we just, yeah. It's not moving up to red now. but let's remove the solo and let's look at this one for a second. Let's sol this one. OK, and now what's going to happen if we do both of them together? I'm guessing it'll be a little too high down here.

So actually, that's not too bad. That's not too bad. So, we can keep that as is. But, yes. So, the easiest way to adjust volume is to start by, you don't really want to do major volume adjustments, until you have all of your tracks basically inputted into the mix. If you start to mix [00:32:00] once you, before you have everything, then you might just need to do it all over again, because you don't really know necessarily what the range, what the highest ends you have are going to be, what the lowest ends are going to be. You don't really know how to adjust them to each other. So it's really good to kind of just put all the files in before you're making the volume adjustments.

And then once you do have everything in, then yes, adjusting first with these little toggles, these circular toggles is good. And then, when you want to get more precise, because sometimes, you don't want to adjust the volume of the entire track. You might just want to make a little adjustment here or there. For example, let's see, solo this [unintelligible 00:32:56] for a second. [00:33:00] All right, so, you can hear in this track a fan noise going on in the background, let's listen to that.

So, I don't actually recommend, I'm going to show you in a few minutes how to get that noise out of there, or at least minimize it in the track as a whole. And that's really what you want to do. But, let's say, just so that I can show you, that you just want to minimize the volume while that's going on in the track. So, yeah, so, you highlight this, you have this yellow line, and that's basically showing you what the volume looks like the whole way through. So right now, it's just steady, because we haven't made any changes.

If you want to [unintelligible 00:33:53] and then you hit this dot here, you click again [00:34:00] over here, and then I like to do a little click in the middle, and then you can just drag to either increase or decrease the

volume. So these are just little markers. All right. If you wanted to, let's see, I'm gonna take those away. If you were just gonna drag this up or down, you could also change the volume on the entire clip like that. But, because we just want to take away the fan noise right here, we're gonna do it this way. So I like to put two little guys next to each other here and here. And then drag it down let's just say, yeah, 14.7 and a half. And then match it over here. And so then it's basically going to sound like a slow, like, actually, a very quick fade here, and then a quick fade in on the other side. So let's see what that sounds like.

[00:35:00] So, I hope you can hear. Yeah, the volume as a whole kind of quickly lowers and then quickly comes back again. Now, that does sound pretty unnatural, because that's just not how audio works in the real world. We don't have suddenly all the noise in the universe just decreasing and then coming right back. So that's why this isn't such a great method for taking out background noise. But if you have a lot of layered tracks, sometimes you can use that. But it's much more subtle. Or if you're trying to fade in and out, that can be useful. If you want to just get rid of one little plosive, for example, if you have a really loud like P or if someone just shouts really loudly, if you have very short, isolated, unfortunate noises that you don't want that can be a great method to get rid of them.

Or if you want to make a subtle change and let's say there's just a couple of words, maybe I'm mumbling [00:36:00] right here. Let's say I have, we felt like I was mumbling there, and we wanted to just increase the volume a little bit, we could just do that. Just pull that up a little bit. Subtlety is generally your friend, because you don't want anything to sound unnatural. So I just pulled it up a little bit there, and I wonder if it'll even be noticeable.

Yeah, OK, so that is going to [unintelligible 00:36:32] a little too noticeable. So, if I were actually going to be doing this professionally, I would make it a much smaller change. But in any case. All right, let's take that out. So, fading. You can either fade just like I said, with creating these markers at the end, instead of dragging them down. Now, you can see here, there's no, once you get to the bottom, then the volume is completely gone. [00:37:00] So if you want to completely fade out, by the end of the clip, then you can do that.

But, there's also this very handy little automatic fade so it's this top box on the right, and you can just pull it, this is just much more precise as you can see, if I were going to try to mimic this nice curve with all these markers, it would just take a long time to do that manually. So, let me just zoom out a little bit. So this is generally the easiest way to do a fade in or out. And if you kind of just pull your cursor up and down, and hold it, then you can adjust the shape, make it really nice and slow if you want, or really quick.

All right. Zoom out. [00:38:00] So let's see. Yeah, so you want your volume throughout the mix to be consistent. That's the most important

thing. You never want the listener to be suddenly overwhelmed with an increase in noise in the way that you maybe would if you were listening to a piece of orchestral music or something, or mixing an orchestral piece. But definitely not if you're just doing a podcast. OK. Let's see. Oh, yes, and I also wanted to point out that if you solo this track, if you drag clips together like this and they overlap, it creates an automatic crossfade. So let's just listen to that.

So we have [00:39:00] all right, so, there's no reason you would actually want these to be crossfaded, if we do that, then you'll hear a little bit of fade in and out. So, if we didn't do that, then you'd hear a really abrupt cut out. So let's see. Let's drag this down, for example. Solo this. So if we just completely layered them on top of each other like this, it would sound pretty weird. You're often going to, you might hear sort of a clip when the one track ends and the other one begins. And so that's another reason to have a fade. Because a fade can minimize that clipping or like little pop sound, which definitely sounds unnatural.

So sometimes you want to edit someone's speech a little bit so that they sound cleaner. [00:40:00] Maybe you want to get rid of some ums, some breaths. Something like that. That's definitely acceptable. But you never want to do it too much. Or else the person will sound unnatural. So humans do take breaths in before they speak, and they're often audible. And even though we're not registering the fact that we're hearing them when we're listening to someone, we will often notice that something is wrong if too many of them are taken out.

So, really, you should just learn to trust your own ear when you're listening to something, when you're editing it. And then decide whether or not it sounds natural to you. So don't over-edit. A good rule of thumb is to only edit out what are often called leading breaths. Or the breaths that occur immediately before you or anyone else starts speaking, and then also editing out the quote, unquote trailing breaths or the breaths that are immediately after the last words spoken by you or anybody else. So I think we have some of those in here. Let's see.

[00:41:00] Yeah, OK. That's a big leading breath. So, let's actually really zoom in on this guy here. That's a leading breath. So if we wanted to cut that out, we could come down here, maybe I'll hit R for razor, just clip that, hit T to get out of razor, and then cut that. Great. So, that sounds very abrupt, especially because we do still have some of the background noise. So I'm just going to hold that a little bit further back, and I'm gonna add a little fade, and hopefully that sounds better.

Yeah, that's fine. That's fine. So, that's [unintelligible 00:41:53] OK. Zoom out here. All right. Sometimes you can actually use breaths [00:42:00] to make cuts in other places sound smoother. So, let's say that you're cutting out part of the middle of a diatribe that someone went on, because it's just too long. Sometimes that cut can sound really abrupt. I

was talking about earlier, how there can be a little clip sometimes when you're cutting things together. But if you copy and paste a breath that was elsewhere in the mix to, right where that cut is happening, and maybe make a little extra space, you can often smooth it out and make the cut pretty much impossible to notice. So, again, this is a situation where you'd want to use your ears and just go with whatever sounds the most natural.

I like to sometimes, if I'm finding a really good audible breath that sounds particularly natural, that doesn't sound too weird, just to cut it and kind of put it towards the end of the recording, or something, and just keep it for later in case I want to draw on it. And then yeah, you also might want to cut out ums, and uhs, and other filler hesitation words. So feel free to do that, but you don't generally want to clean up someone's speech so much [00:43:00] that it no longer sounds authentic. Especially if you are producing a work of journalism, ethically, you want to try to preserve someone's voice as much as possible.

All right. So, yeah, let's zoom all the way out here. OK, so, I'm going to just really briefly touch on EQ and compression. EQ is equalizer. These are a little bit more advanced editing tools, so I don't want to get too much into it. But, basically, equalizer is adjusting the volume of individual frequencies in a mix. And a frequency is just a pitch. So, for example, the human voice, especially, yeah, so my female human voice is going to be kind of centered around [unintelligible 00:43:58] often [00:44:00] especially a low rumbling wind, might be focused on a much lower frequency, and so if you want to minimize one of those sounds, you can actually do that by going in and adjusting EQ.

Compression is the process of reducing the dynamic range of a sound. So that is basically used to make certain sounds in a mix pop, and feel more forward and present. And really, what I'm just going to say in terms of using these right now, is just that you can mess around with the presets. It's too much to get into in this short tutorial. But let's say, yeah, all of these are under effects. Let's say we want to add a bit of an effect to my voice here. Let's loop it. So, let's click I, and [unintelligible 00:44:45] right here. And O right here. Drag the play head, now, this is just going to be looping. Now it's looping.

So let's just try out [00:45:00] – actually, I'm going to show you the [unintelligible 00:45:04] later so let's just try out a compression. So amplitude and compression. Amplitude, so, again, because this is ultimately talking about volume, amplitude, maybe this is getting a little bit too into the nitty gritty physics of the sound, but amplitude is the size of the sound wave, and volume is really the perceived loudness, and amplitude is what it actually looks like. So we're coming down to amplitude and compression.

Let's go to multiband compressor. OK, so, let's listen to just how it sounds normally. So, let's say we wanted to pop in some sort of way. Yeah, so

they have all these great presets, and you can literally just try out different things. Let's see what broadcast sounds like [00:46:00]. Yeah, so you can see, it sounds louder. And then if you were compressing, often, what you're trying to do when you're compressing, is you're amplifying just one specific part of the sound and then if you have other sounds going on, this is a little bit more relevant to music mixing, but if you have, maybe a violin that maybe you want to make the higher end of that pop, and maybe if you have a saxophone and you want to make the alto range pop a little bit. And you're just trying to create space, carve out space in the mix so that each one of them can feel present and heard.

But it can get a little complicated. So if I were you, I would just mess around here. Heavy guitar, that's going to sound interesting on a voice. Let's try. OK, excellent. So that's that. So again you can just kind of mess around with things and then whatever [00:47:00] sounds right to you is probably going to sound right to a listener. Now, let's talk about minimizing background noise. So we have this piece of tape here that has this really ugly band sound in the background and also kind of a hissing. Let's see. So, again, this is going to be different than just trying to minimize the volume as a whole, because this is when you know, you have this fan sound going, it's particularly noticeable here when I'm not speaking, but it's happening the whole time, so you can't just reduce the volume to get rid of it.

So, let's see. Basically, one term that you should know is called a noise floor. That's kind of the shushing that you hear when you turn up any recorder loud enough. Literally every recorder has a noise floor. That's just the sound that naturally happens when the recording is [00:48:00] going on. It's not even the sound of a room. That's a room tone. It's just literally the sound of the recorder itself. There's just, processing itself makes a little bit of a noise.

There's some ways to get rid of that. There's ways to get rid of the room tone. The room tone in this case is the fan. So there's a couple things we can do here. So let's try what's called a parametric equalizer. So, this is a form of EQ. So it's good if you want to remove a sound that is consistent and is occurring on the high end or the low end and you want to minimize it. So let's just put I here, O here, so you can hear just a little bit of just the fan and a little bit of the fan and the voice. Let's go to [00:49:00] effects, filter and EQ, parametric equalizer.

So, if we wanted to take out the high end, then we could do LP, so that's a low pass filter. If we wanted to take out the low end, we could come down here and select HP, it's a low pass filter. So in this case, there is a bit of a hissing going on, which would be on the higher end frequency wise, but there's also this fan, and that's really what I'm concerned about. So we're going to focus on the HP. So, I'm highlighting that.

So, basically, you just pull this HP, this high pass filter, and the more that we pull it, the more that it's going to take out the frequency range. So if it's down here [00:50:00], then all of these frequencies are [unintelligible 00:50:03] lower end, this is the higher end. But if we're pulling it this way, then we're removing everything to the left. So let's see what happens when we're [unintelligible 00:50:12].

So, you might want to turn up your volume a little bit here. But if you do, then you should be able to hear that there is a decrease in the noise of that fan. You can still hear it, unfortunately, and there's no, via this method, there's no way to completely get rid of it, because the frequencies of the fan are interacting with my voice. But you are minimizing it, which is great. If you pull this up too much, then you're actually going to hear some distortion in my voice, because, yeah, we're minimizing some of the frequencies that are also in my voice. So, let's try it here [00:51:00] for example. This is probably a little better.

OK, so right here you're not hearing much of an adjustment to my voice but you're definitely hearing a reduction in the noise of the fan. But if we were gonna put it up here, yeah, you can hear that now, we're only really hearing the higher frequency range of the voice, and that sound very strange and distorted. So, let's just put it around here. We can keep that on. If we wanted to, we could also take it off. Pause that. Great. So that's that.

When you're adding an effect on a track, it's going to come down here. And there's gonna be this nice little list. So we did this compressor earlier, let's actually turn that off. [00:52:00] We don't need that anymore. In fact, we can either turn it off if we want to be able to turn it back on. Or, we can just delete it, [unintelligible 00:52:13] great. You can also drag them around.

OK, so, we've got the equalizer, let's actually take that off as well. Now, another method that you can use, if you just want to generally minimize background noise, so you're not trying to just take out something on the low or the high end, but you're just generally wanting to minimize it, so, let's DeNoise. So we're coming back up to effects, noise reduction restoration, and DeNoise. OK. Actually, we want to have this [unintelligible 00:52:49] here again and move that.

All right. DeNoise. [00:53:00] So now we can literally just drag this around. Down here, we're not taking anything out. So you can still hear it a bit, and then the more that you drag it up, the better it is. So that's a really great method, and that's again, just generally helping to get rid of that background noise.

Another great one is, this is kind of the best option if you want the software to automatically – you know, I'm gonna actually drag the volume up, so hopefully that'll be easier for us to hear when [unintelligible 00:53:38] let's see. Yes, that's much better. OK, if you want the software

to automatically adjust a noise problem, adjusting it to your unique problem your unique circumstances, then there's one great kind of [unintelligible 00:53:53] again [00:54:00] [unintelligible 00:54:00] O looping, come up here to effects, noise reduction restoration, oh, I'm sorry, for this, we actually need to be in the waveform mode.

So, danger, danger, it's going to make permanent changes to the waveform, but it also can be very useful. So, let's do that. Effects, noise reduction and restoration, and we're actually going to capture noise print. Actually, I'm sorry, let's do just this portion right here that has the background noise. Effects, capture noise point, and now, Audition is remembering this background noise, and we're going to go back here, and we're going to go [00:55:00] to Noise Reduction process, and it's actually going to automatically take out that print.

OK, yeah, so you can really only just hear a tiny little hiss at the top. The fan is actually pretty gone. So if we want to now select all, which you can do with command A, you come up here, and make that happen on the entire track. So, apply. Yeah, so you really can't hear that fan anymore. My voice doesn't sound 100% natural, so I wouldn't necessarily recommend – you can basically toggle it. If you, right here, it's at 100%, so we're reducing that noise print by 100% but I wouldn't necessarily recommend that. Let's try 70.

[00:56:00] Yeah so you can just move this around, you can move this around as well until it sounds natural. OK, so, yeah, so we've just permanently changed that recording. So, again, just good to be very careful, and yet another reason to make sure that you have originals of all of your files saved nearby, so that if you do make a change, you can bring back in an original version.

So, the unfortunate truth is that there's really only so much you can do to remove background noise that exists in a recording without reducing the quality of the sound that you actually want in there. So the best thing you can do is, just to make sure that you're getting the cleanest possible audio when you're recording in the first place. OK, so just a couple little terms and things I want to make sure you know about.

Scene of the incident if you're making a podcast, you might be asked [00:57:00] to choose quote, unquote selects. Selects are basically a few minutes of material from the audio that you have that are particularly compelling or demonstrate a point. So, they can be used to show a boss or someone else an update on what you've been working on. Maybe show some of the best of the material you have. Maybe pitch a story. Something like that. It's a series of clips that are often unrelated. You're not creating a scene, you're just choosing the best of what you have. It's a few minutes of the best of what you have. So, that's really, if you're creating selects, you're just going to be putting all of your favorite tracks together, making sure it's not too long, and then exporting those together.

Another thing you might be asked to make when you're creating a podcast is a rough cut. And that's basically a rough draft of a podcast. So it's a version that can then be edited and critiqued. It should ideally be about the length that you're aiming podcast to be [00:58:00] when it's finished, and have as much of the final material that's going to be in that podcast as possible. Depending on how much time you have to produce the podcast, you may make one rough cut or several. Or if you're under a severe time crunch, you might not even have time to make one at all. But it's basically just a rough draft.

OK, so, let's learn how to export. So, you want to be careful if you have any muted tracks, then those are going to be muted when you export. Same with soloing. So right here, I want everything to just be captured naturally. So nothing is muted, nothing is soloed. So we're just going to go to file, export, multitrack mixdown. Now right here, I don't have anything in particular selected, so we're just going to do the entire [unintelligible 00:58:56] [00:59:00] or another great option, if I for example have a lot of dead space at the beginning or end, then I can hit I, and hit O, and then export just the time selection, and then it's just going to export that thing that I've selected right there.

So you can export it as wave, it's slightly better quality, MP3 is often perfectly fine. And you can literally just hit OK, and then it's going to export it, so we can now find that on my computer. You don't want your file to be too big. If you're putting a podcast onto the internet, you probably want it to be a maximum of about 25 megabytes. Otherwise it'll be difficult for people to stream and to download.

Just a couple of other things. It can be helpful to save new versions of your piece every time you've done kind of a significant [01:00:00] round of editing, just in case you want to go back and change something. So I just like to save, you know, tutorial one, tutorial two, etc., as long as you have enough space on your computer for that. Another term that you should probably know, two terms, actually, you might hear people talking about acts and tracks. Acts stands for actualities. That's tape that you have recorded personally, and then tracks, like, in the field, so anything, you know, you're doing an interview in the field, or even an interview on the computer, just something you've recorded that's setting a scene. And then tracks is your own narration. So, hello, and welcome to this podcast, that's tracks. So if you're tracking something, you're adding narration.

You might hear the term axi-tracksy. If the sound kind of becomes too homogenous, and you're just kind of switching between narration and [01:01:00] then like a field tape, just a field tape, or maybe a question and an answer in a really kind of rigid, three seconds of this, three seconds of that, then your ear kind of gets bored, and people say that it's too axi-tracksy, and you might need to make the soundscape a little bit less homogenous.

You also might want to add music so, there are lots of different royalty free music sites. This one right here is something that I downloaded earlier from [unintelligible 01:01:33] dot sessions. That's a great place to get royalty free music. Highly recommend. They don't have a ton of options, but they do have a fair amount. This is called Sonatine Bureaucratique, and then again we can just drag that in. Let's see. I'm gonna hit W here, drag these all to the side. Let's say for example that I want the music to come in and that I want it to [01:02:00] decrease when the narration is going to start, I'll just pull this down a little bit. Let's see how that sounds.

OK, so, yes, music can often really spice things up and yeah, prevent the texture from becoming too homogenous. One rule of thumb is you don't really want to have any particular music queue going on for more than 90-ish seconds, or else the listener can start to kind of, again, your ear gets bored, and you start to zone out. It's a general rule of thumb, it's not necessarily the case but it can be [unintelligible 01:02:47].

OK, so, I think that's everything I wanted to get through, happy to take [unintelligible 01:02:56] at this point. Yeah, please let me know if you have any [01:03:00] questions.

[End of recorded material 01:03:06]