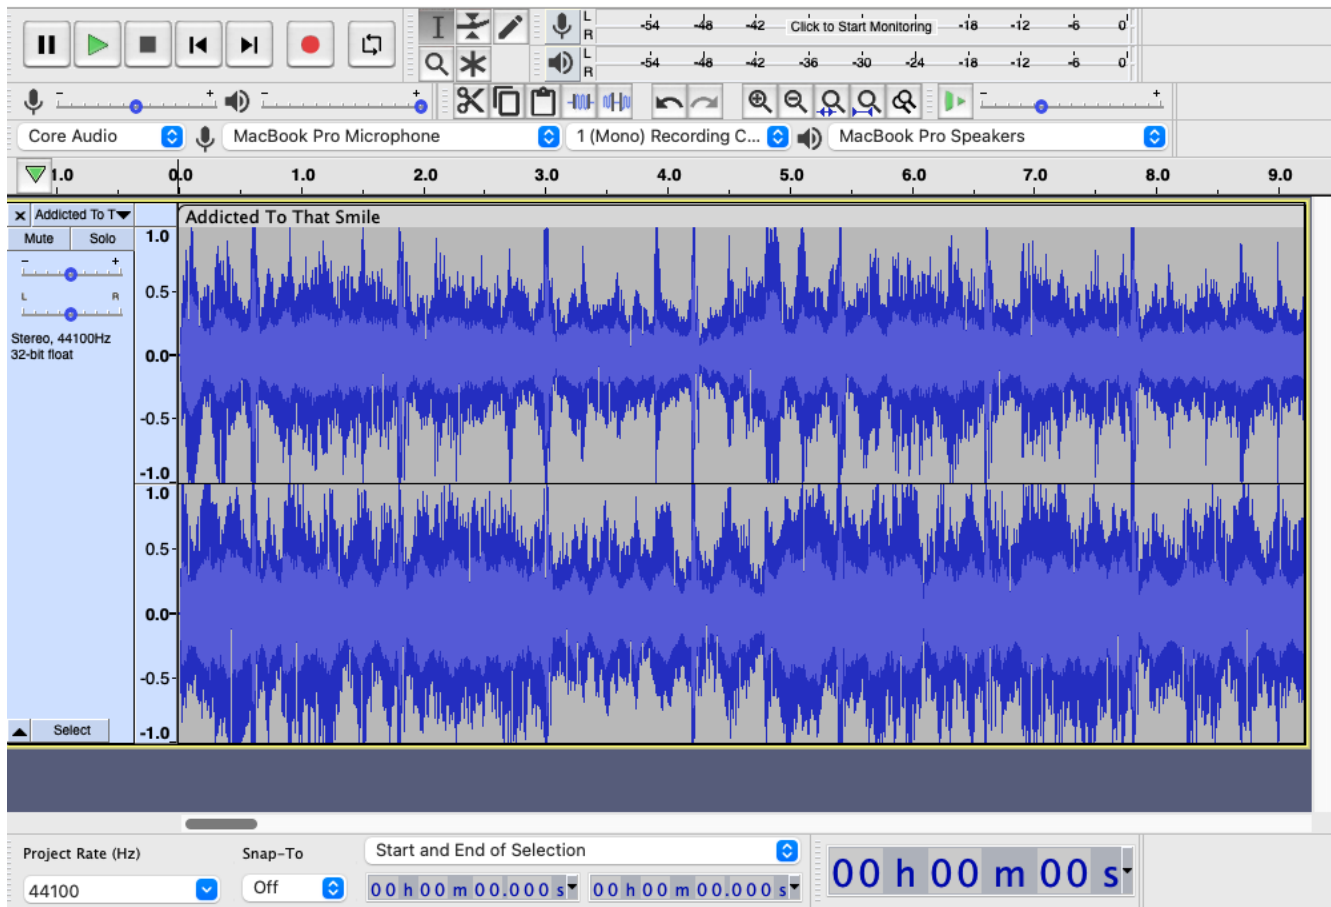


Audacity 3



In this unit we will investigate:

- the Audacity interface, controls, and track control panel
- importing and playing audio tracks
- editing techniques: selection, cut, copy and paste; labelling; amplification and other special effects; trimming, rearranging, and time shifting
- noise removal; voice removal
- varying tempo, speed, and pitch
- normalization, equalization and compressor
- recording sound, and overdubbing
- creating a music mix and exporting
- theory: sound, file formats, sound quality, recording
- glossary.


These activities use Audacity version 3 which is a free download.

For directions on how to access Audacity 3, set it up, and download the sound files used in these activities see the notes at the [end](#) of these activities.


Quick reference




Transport toolbar (playback controls)



Edit toolbar



Tools toolbar



Selection tool

- Left click – start point for playback
- Left click and drag – select part of track; play only selection
- Shift + left click and drag – extend selection
- Shift + arrow key (i.e., left or right arrow ⇐ ⇒ on keyboard) – extend selection
- Ctrl / Cmd + Shift + arrow key – reduce selection
- Left click in title bar and drag – time shift.

Zoom tool

- Left click – zoom in closer; click and drag to zoom in on the area selected
- Right click – zoom out

Envelope tool

- Left click – set control point
- Left click and drag – move control point.

Draw tool

Note: will only work when zoomed in to see individual samples.

- Left click – set amplitude for sample
- Left click and drag – alter amplitude for sample; set multiple samples.

Multi tool

- Left click – as for selection tool
- Right click – as for zoom tool
- When zoomed in to individual samples the left mouse button will act the same as the draw tool.*

Recording/export settings

Quality required	Recording		Export bit rate	
	Sample rate	Sample size	mp3	uncompressed
Voice only	44.1 kHz	16 bit	45-85 kbps	16 kbps
Music	44.1 or 96 kHz	16 or 24 bit	170-210 kbps	16 kbps
Very high	192 kHz	32 bit	320 kbps	32 kbps


Activity 1 – Introduction to Audacity

Audacity is a free, easy-to-use audio-editing program. With Audacity we can record our own sounds and music, or we can work with existing audio files. In this introduction to Audacity we will import an *MP3* music file and use it to investigate the interface.

For these activities you may use your own audio files, or you can download the named files from [here](#).

1. a Start Audacity.
- b From the menu choose *File > Import > Audio* and select the file *Addicted To That Smile.mp3*.

A waveform view of the track will appear, with a track control panel at its left. (We will look at the track control panel in detail shortly.)



This is a *stereo* recording with two channels. The top waveform is for the left channel, the lower for the right, i.e., the sounds that will come out of the left and right speakers or headphones respectively. If you look closely, you will see these are not identical. (In a *mono* recording the L and R channels are the same.)

To see the wave forms more clearly use the *Zoom in* button on the *Edit* toolbar to magnify them. 

- c Use the green *Play* button to listen to the music. This is on the playback controls at the top.




As the music plays, set the output volume to an appropriate level. To do this move the output volume slider near the top of the screen to the right of the *Edit* toolbar.



Note: you may have to also adjust your computer's volume or even disable mute if it has been set.

Unless you have headphones, the sound should just be comfortable enough for you to hear *without distracting others*.

- d Above the track is a time scale. This track is just over 2 minutes long. A playback indicator, shown as a green triangle above a vertical bar, moves along the track as it plays.



- e Experiment with the playback controls of: *Pause; Stop, Start* and *End*. (We will use *Record* and *Loop* in later activities.)

Note: the *Pause* button acts as a toggle switch; one click will pause, a second click will restart the track.

2. As a track plays information about the sounds is displayed on screen.
 - a Play the track again and note, at upper right, the levels of sound output indicated as a green bar, one for each track. This is a stereo recording. The L indicates the sound level that will come out of the left speaker or headphone, and the R, the right speaker or headphone.




This device is known as the VU (volume unit) meter. It shows how loud the current sound is playing at for each speaker, and how loud it has just been. A red *clipping indicator* will appear at the right-hand end to show if the sound is so loud it has become distorted.

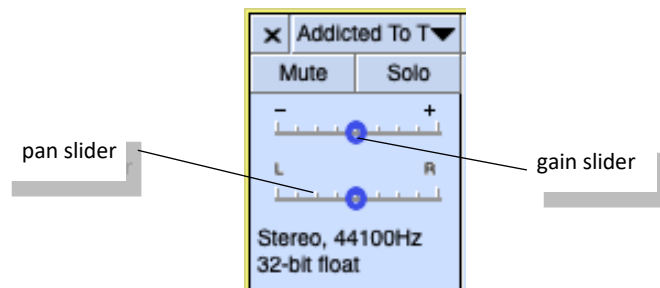


By default this VU meter will be horizontal and show levels in decibels, e.g. this track plays around -12 dB to -6 dB+.

Note: the meter above this one is used when recording to indicate input (recording) levels. We will look at this later.



- b The VU meter shows the overall sound being output to the two speakers. To vary the *balance* of volume from the speakers drag the pan slider on the track control panel at left.




Pan right will mean the volume from the right speaker/headphone will be increased, and that from the left decreased.

The slider above the pan slider is used to increase or decrease “gain” i.e., signal strength. Increasing the gain too much will cause distortion, but on the other hand, if a track is distorted, reducing the gain can limit the distortion.

Experiment with each of the pan and gain sliders while the music is playing.

3. To work with tracks we need to be able to select just parts of them. To do this we use the *Selection* tool. This is chosen by clicking on the I bar on the *Tools* toolbar just to the right of the playback controls.
 - a Check the *Selection* tool is active (darker) and click anywhere on the track. A fine line (the playback indicator) will appear over the waveform.
 - When you press *Play* the music will start from the point where you clicked.
 - b Use the mouse and click and drag on the track. This will select just part of the track.



When you click the *Play* button now, just the selected part plays.

If you click the *Loop* button and then *Play* the selected part will repeat over and over.

- c To extend a selection click at the edge of the selection, hold down the *Shift* key, and drag the end of the selection with your mouse. This can be done at the beginning or end of the selection.



For finer control in extending a selection use *Shift* + the keyboard arrow keys \leftarrow \rightarrow . The left arrow will add to the beginning, while the right arrow will extend the end of the selection.

To shrink such a selection use *Ctrl + Shift* (Windows) or *Cmd + Shift* (Mac) + the arrow keys.

Note: Once part of a track is selected, we can use the *Fit selection in window* button on the edit toolbar to just work on this part.



- d To cancel a selection click anywhere on the track outside the selection.
 - e To select a whole track click in the grey control panel at left, or double click on the track itself.
4. In addition to the *Selection* tool there are four other tools to assist in working on tracks. We will briefly look at these now.


- a The *Zoom* tool is used to see the track information more closely.



When the tool is selected a left mouse click will zoom in, and a right mouse click will zoom out.

A click and drag with the *Zoom* tool will zoom in on the area covered by the dragging action.

If you zoom in closely enough, you can see dots on the track. These are individual sound samples, as will be explained later.



A quick way to zoom right out is to use the *Fit project to width* button.




Note: when using the zoom options the time duration of the track does not change, we are just closer to it. If necessary, when zoomed in, you can use the scroll bar at the bottom to see different parts of the track.

- b The *Envelope* tool is used to adjust volume in specific spots.




When this tool is selected a thick blue line will appear above and below the track.


Clicking on this line will add control points.




These control points can be dragged vertically to alter the volume of selected parts of a track.
To remove a control point simply drag it outside of the track and then release it.

- c The *Draw* tool is used to modify existing sounds. This tool only works when you are zoomed right in and can see the individual sound samples. 

Each sample can be moved vertically altering its volume, or if you click and drag horizontally you can sketch an altered sound wave.

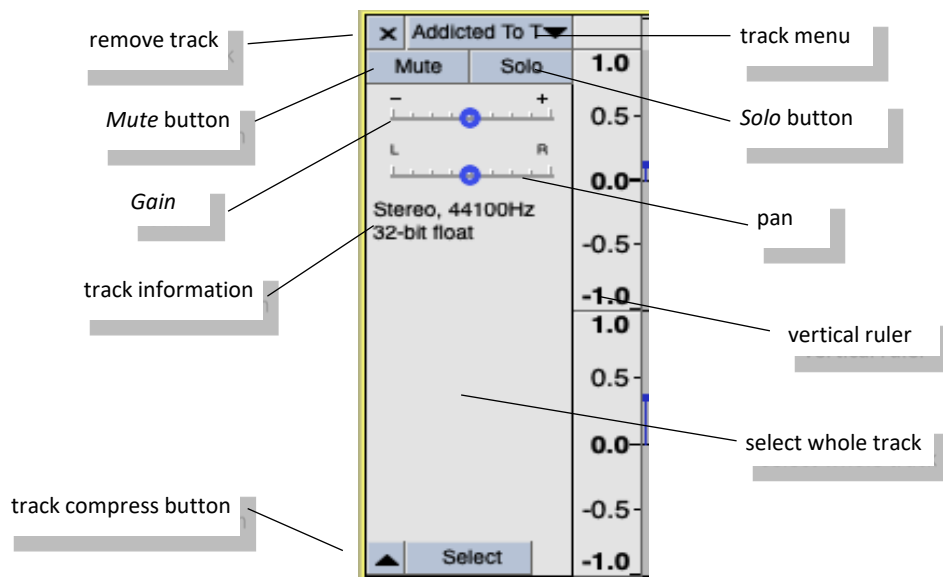


- d The *Multi-tool* can do a variety of the tasks of the other tools. When selected a left mouse button acts the same as the *Selection* tool and a right button acts as a *Zoom* tool (use both buttons to zoom in, right to zoom out). 

If zoomed into sample level the multi tool acts like the *Draw* tool.

Note: earlier versions of Audacity used a *Time Shift* tool. This has been replaced by a dragging action as will be shown in the next activity.

5. The final item we will look at is the track control panel.



Looking at each of these controls, clockwise from top right, they are:

- track menu – click on this to rename a track, to display the waveform differently, etc.
- solo button – play only this track; this is useful when you are working on several tracks at the one time, and want to listen to just this one
- pan (channel balance) – vary the volume of sound going to the left and right speakers/headphones
- vertical ruler – this is used as a guide to waveform levels; experiment with this by left and right clicking in the ruler and by click-and-dragging in it
- select whole track – click anywhere in the grey middle part to select the track; (alternately you can double click in the track itself to select it); to select multiple tracks use *Shift+click* in the control panel of each; to select all tracks use *Ctrl +A* (Windows) or *Cmd + A* (Mac)

- track compress button – clicking this will shrink the track to take up less space; this is useful if you are working on multiple tracks at one time; (you can also drag the bottom of a track up or down to make the track wider or narrower vertically)
- track information – here this is a stereo track, recorded at a rate of 44 100 Hz, in 32-bit format
- gain – increase or decrease the signal strength, i.e., the intensity of the sound
- mute button – silence this track (i.e., when you have several tracks loaded at once)
- remove – delete this track from the project.

Note: some of these controls are only used when working with multiple tracks, e.g., when creating a music mix.

6. Close and exit without saving the project.


Follow up – Sound waves

Audacity is sound manipulation software. We will use Audacity to record, input, edit, mix, and encode sound files. To gain an appreciation of what we are manipulating it is important to understand what is involved with sound, and sound waves.

Sound


Sound is just vibrations in the air. When a drum is hit, the drum skin vibrates up and down causing the air around the drum skin to move back and forth. The diaphragm of a speaker works in a similar way. Magnetic impulses cause the diaphragm to move and it, in turn, causes the air near it to move in pulses. These movements, or vibrations, are passed on through the air until they reach our ears.

The human ear drum picks up the vibrations and we experience sounds.



Source: en.wikipedia.org/wiki/Image:Processing-of-sound.svg


The sound vibrations are usually drawn as a wave pattern.




These waves have amplitude and frequency.

Amplitude and frequency

Amplitude is the wave height. The higher the wave the louder the volume of the sound.




loud sound




soft sound

Amplitude is measured in decibels (dB).

Frequency is how many vibrations (waves) there are in a second. The more vibrations per second the higher the pitch of the sound.



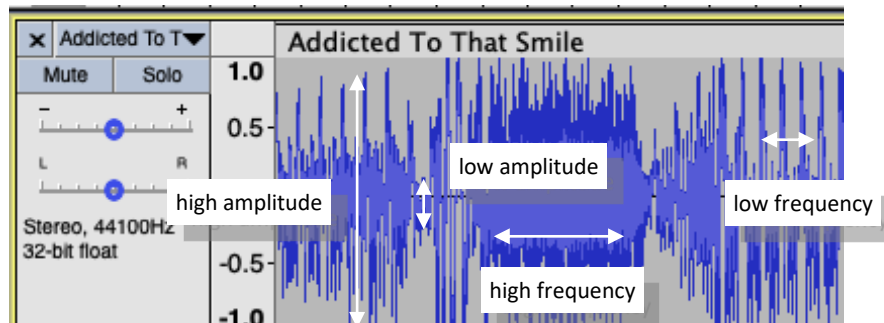
high pitched sound



low pitched sound

Frequency is measured in Hertz (Hz). One Hertz is one vibration per second.

When sounds in Audacity are shown in the (default) waveform view, the amplitude and frequency of all the sounds that make up the track are displayed.



The degree of amplitude (loudness) can be seen on the vertical ruler, while the frequency (pitch) can be measured against the timescale along the top of the track.

With Audacity we can manipulate the amplitude and frequency of sounds as we shall see in upcoming activities.

I/O devices

The playing and recording of electronic sound depends on sound vibrations in the air. This is achieved using I/O (input/output) devices. The main output devices are speakers or headphones, while most input is by microphone.

A speaker or headphone makes sounds by using magnets attached to a diaphragm. The magnet is turned on and off by very rapid electronic pulses. The movement of the magnet causes the diaphragm to vibrate at the required amplitude and frequency. The movement of the diaphragm, in turn, makes the vibrations in the air that we hear as music or other sounds.



Microphones work like speakers, but in reverse. Vibrations in the air cause a diaphragm attached to a magnet to move. The magnet moves in an induction field creating electronic pulses that can be recorded.


Note: in addition to recording using a microphone, electronic sounds can be created directly on a computer or digital music device (e.g., electronic organ) and input directly in a digital format.

Activity 2 – Editing

In Audacity we work in what is called a *project*. A project is an organizational device to hold all the tracks, labels, edits, etc., for a given task.

In this activity we will use an Audacity project to edit some short sound files.

1. a Either start Audacity or, if already in the program, click on *File > New*.
This will start a separate Audacity project.
- b From the menu choose *File > Save Project > Save Project As* naming it *homer*.
This project will hold a series of soundtracks from Homer in *The Simpsons* TV program.
Note: when you save a project Audacity will automatically add *.aup* to the end. Audacity project files are only recognised by Audacity. They will not play in any other audio program.
- c From the menu choose *File > Import > Audio* and load the file *doh.wav*. (If prompted to make a copy of the file choose *Yes*.) This is a mono track, i.e., just one channel.



Play the track (it is only 1.6 seconds long). As the playback indicator moves along the track note how the height of the wave corresponds to the volume of sound. You can also distinguish where individual word sounds occur.

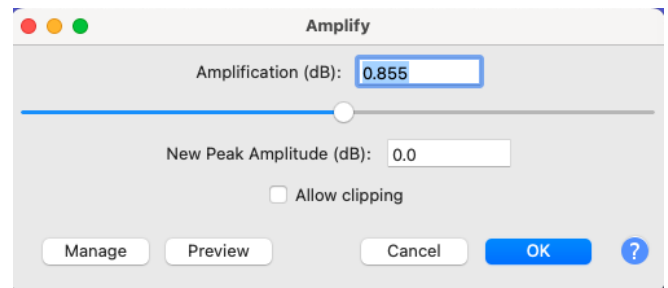
If you wish click on the *Loop* button then *Play* to repeat the track indefinitely.

2. The simplest changes we can make to tracks are to apply special effects.


- a Select the track (double click on the track), and then from the menu choose *Effect > Amplify*.
Amplification will automatically work out the loudest the track can be made without distortion and display this value.

Accept the values given and click *OK*.

Note how the waveform has changed and how the sound is louder when played.



Amplification is limited to the range of the waveform (by default -1 to +1) unless clipping is allowed. (Clipping will lead to distortion.)

- b Click on *Undo* and listen to the sound again. The change has been removed.
Note: we can use repeated undo to back up changes until when a project was last opened. 
- c Select just part of the track and amplify it.
Play to hear the change and then undo.
- d We can also apply negative amplification to quieten a track or selection. Try this.
- e Experiment with some of the other effects. (Remember you will have to select all or part of the track before you can choose an effect.)

You might like to try *Change Pitch*, *Echo*, *Reverb*, *Reverse*, *Change Speed*, and others.

Play and then undo each effect as you go.

Note: Some effects, such as *Fade In* or *Fade Out* work with just a selection of a track, while others, such as *Wahwah* work best with music.

3. Labels can be used to identify parts of tracks such as where a key change occurs, or specific selections such as a guitar solo or song chorus.

a Use your mouse to select just the *d'oh* part of the sound.

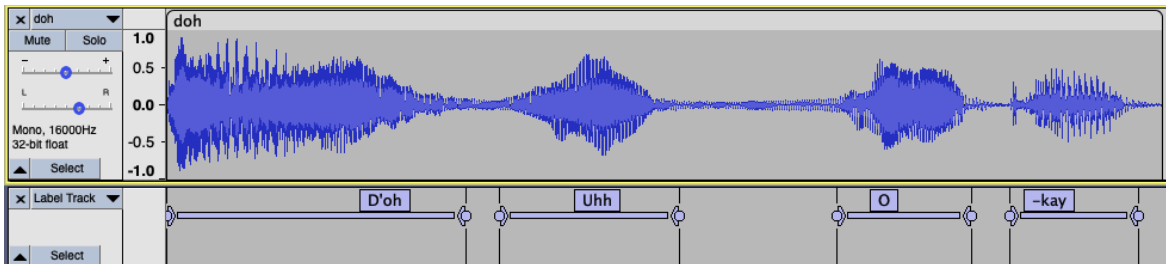
From the menu choose *Edit > Label > Add Label At Selection*.

Label this section *D'oh* (or similar).

Note: the shortcut for adding a label is *Ctrl+B* (Windows) or *Cmd + B* (Mac).

b Label the other parts of the track as – *Uhh*, *O*, and *-kay*.

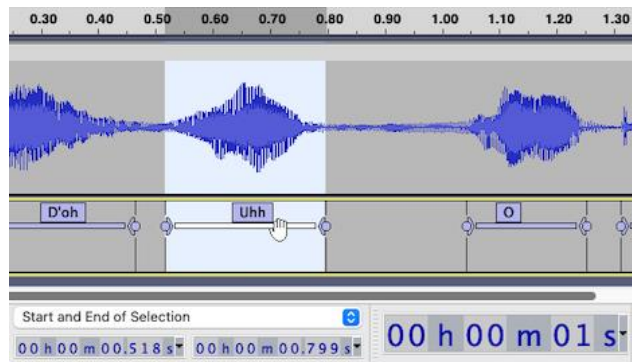
Click on each label in turn to see or play the part selected.



A label track has been added with a clear indication where each sound occurs.

Labels can be very useful when editing a long track. They can indicate required sections of a song or important parts of a recorded interview. They can also be used to select named parts of a track to work on.

Note: as you click on a label to select a part of a track, its start point, end point, and length (in seconds) is shown in the display in the lower part of the screen.



This selection shown above starts 0.518 sec into the track, ends at 0.799s, and is 0.1 sec long.

c Re-save.

4. Our next step is to investigate some of the basic copy and paste techniques.

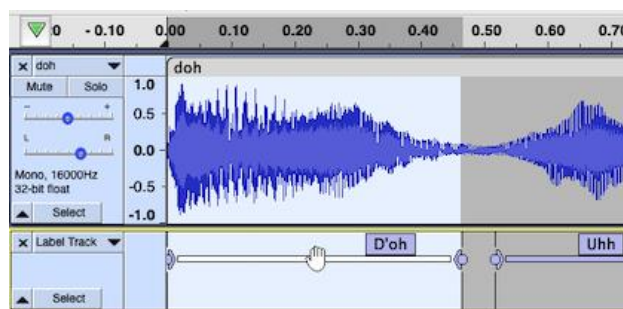
Note: the shortcuts for cut copy and paste are:

Windows: *Ctrl+X*, *Ctrl+C* and *Ctrl+V* or Mac: *Cmd+X*, *Cmd+C* and *Cmd+V*

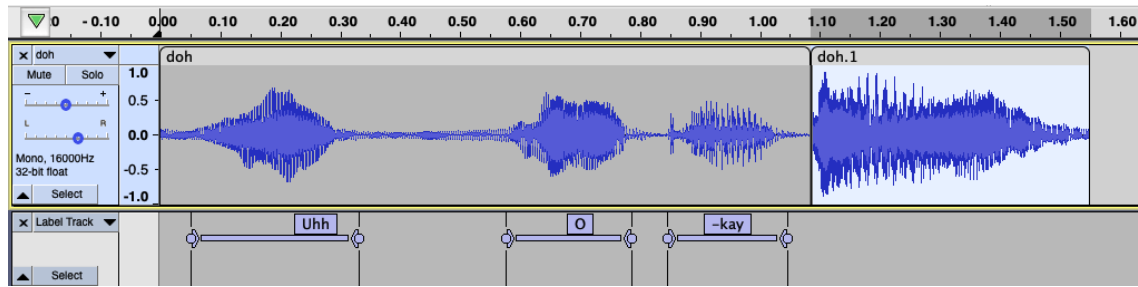


cut copy paste

a Click on the line under the label to select the *Doh* sound only.



- b Cut this sound out from the track using the *Cut* button on the *Edit* toolbar. Click once at the end of the track, and then the *Paste* button.



The *Doh* has been moved to the end.

De-select the moved part (*click away from the track*) and then play the track from the beginning.

- c Undo twice to return to the original sound.


This time *Copy* the *Doh* sound and *Paste* it at the end. Play the whole track to hear the *Doh* in two places now.

- d Undo to finish.


Note: labels do not move when the sounds are shifted. As such labels are only of limited use in annotating a track that is going to be rearranged.

- 5. Apart from simple cut, copy and paste, at times we need to isolate parts of a track to either move them somewhere else, or to work on just a particular section. An isolated part of a track is called a *clip*.

- a Click on the *Uhh* label to select just this sound.
- b From the menu choose *Edit > Clip Boundaries > Split* (or *Ctrl+I / Cmd+I*).



A vertical line now appears before and after the *Uhh* sound, and the track has been split into three separate clips. Each of these three can now be manipulated individually. (We will see ways this can be done shortly.)

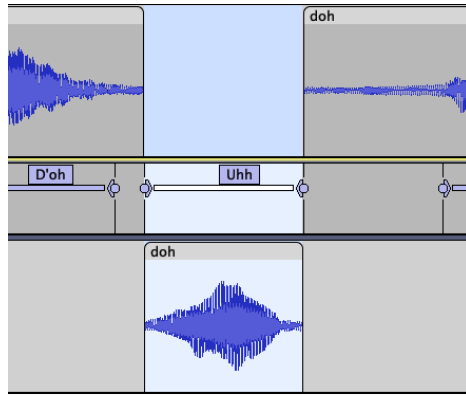


To join the clips together again use your mouse to click on each of the vertical lines in turn, and the clips will be merged back into one track.

- c Sometimes it is necessary to take a section of track and move it to a separate track.

Again select the *Uhh* sound.


This time choose *Edit > Clip Boundaries > Split New* (or *Ctrl+Alt+I / Cmd+Alt+I*).



The original track has now been split into two tracks. The new track has just the *Uhh* sound, while there is a silent gap on the original. This is the best way of isolating a specific clip onto its own track.


Play each track to check this. As you do experiment with the *Mute* and *Solo* buttons on the track control panels.

Note: when a track is split like this the second track carries the same name as the original. If you wish to rename the second track, use the menu on the track control panel (black down arrow) and choose *Name...*



d Undo to restore the *Uhh* sound back to the original track.

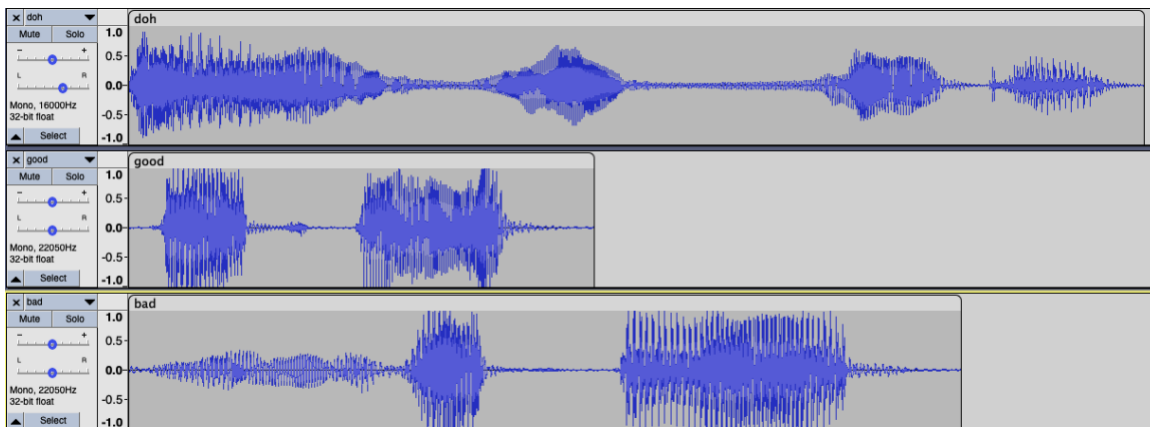
Again click on the *Uhh* label to select the sound and experiment with the following two buttons on the edit toolbar. (Click *Undo* after trying each.)




e Delete the label track (click the at top left of its control panel) as we will not use it further in this project.

6. We are now ready to work on multiple tracks at the same time.

a Import *good.wav* and *bad.wav* (*Project > Import > Audio*). These are also mono tracks.




- b Use the *Mute* and/or *Solo* buttons on the track control panels to play each track separately.
- c At the moment if we press *Play* all tracks will run at the same time. We will now rearrange them to play one after the other. (In earlier versions of Audacity this was called time shifting.)
Move your mouse pointer into the white bar just above the *Doh* track.




The pointer changes into a hand shape.

Use this to drag the track to the right until it lines up with the end of the *Good* track. A vertical yellow line will appear as it lines up exactly.



Next drag the *Bad* track so that it lines up with the end of the *Doh* track. (If necessary, use the *Zoom* tool right click to zoom out to see more of the tracks.)




- d Click on the *Start* button to return to the beginning and then *Play* to hear the three sounds in sequence. (Make sure *Mute* or *Solo* are not selected for any track.)
- e Rearrange the sounds into a *Bad-Doh-Good* sequence.
- f Re-save the project.

Note: You can move tracks up or down by using the menu on a track control panel or by simply using your mouse to drag the whole track up or down. This will not change the order in which the tracks are played but can be useful if you are working on many tracks at once by bringing two closer together.


- 7. As we link separate tracks it is sometimes effective to fade one out and fade the next in. (The following assumes the tracks are in the order *Good-Doh-Bad*).
- a Drag the tracks to overlap the *Doh* with the *Good* sound so the *Doh* begins before the *Good* ends. Play.

- b Use the *Selection* tool to select the part of the *Doh* sound that overlaps the *Good* sound. From the menu choose *Effect > Fade In*. Note how waveform changes. Rewind and play. Note how the overlap is smoother.
- c Drag the *Bad* sound to overlap the end of the *Doh* sound.



overlap

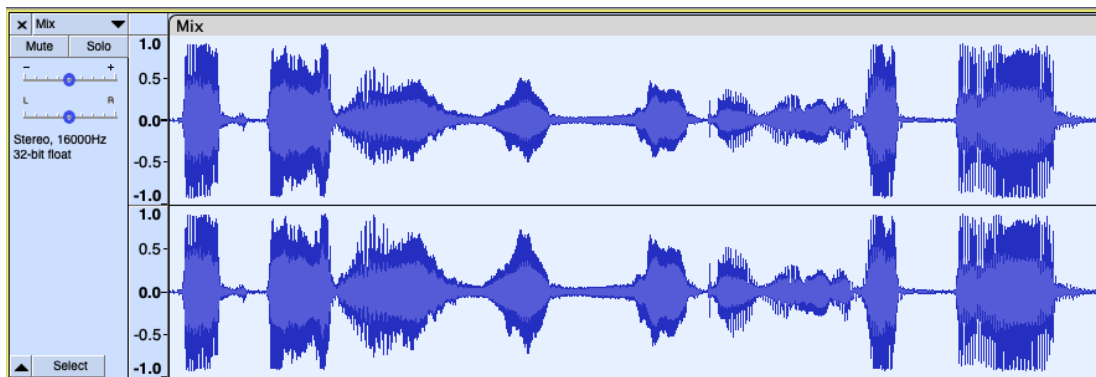
Fade out the end of the *Doh*, and fade in the start of the *Bad* sounds where they overlap.



fade in

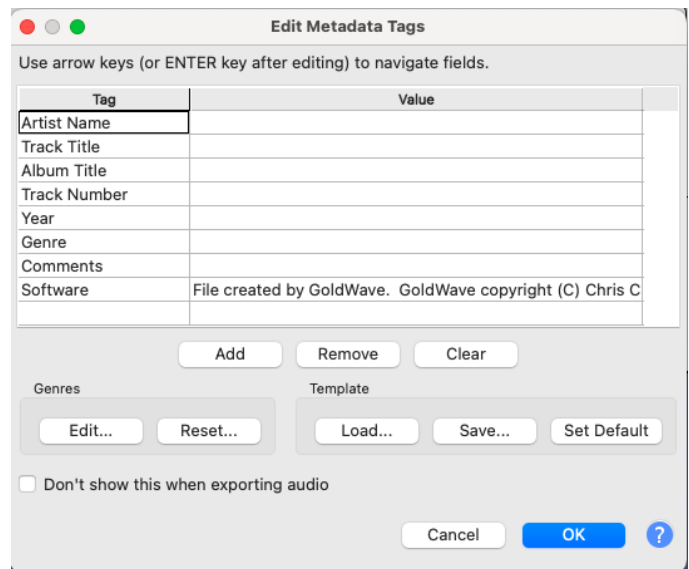
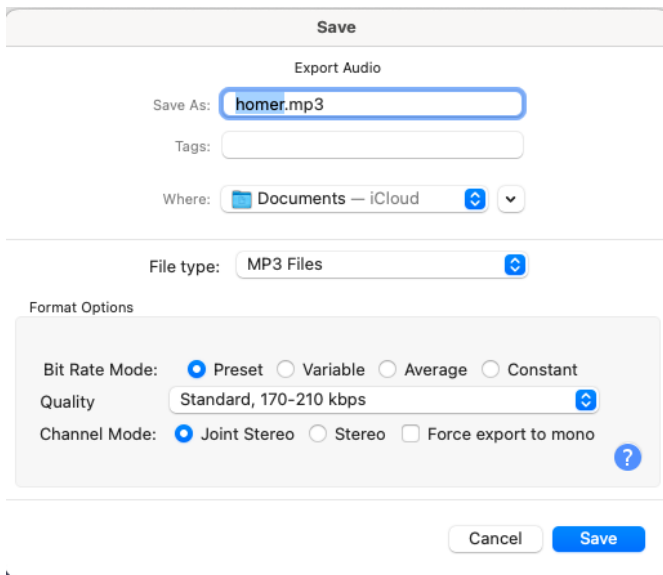
Note: as a rule, *fade out slowly – fade in quickly*.

- d Play to hear what it sounds like now and then re-save the project.
8. We have finished our editing and are now ready to export this sound as a new track.
- a Select all tracks (*Ctrl+A* or *Cmd+A*) and click on *Effect > Normalise*, accepting the default settings. This will adjust the sounds out so that they blend together at around the same volume.
 - b Again select all tracks and then choose *Tracks > Mix > Mix and Render*.



- c Rewind and play the mixed track.

- d We will now export this track as an mp3.
From the menu choose *File > Export > Export as mp3*.



Choose a suitable name and folder, accept the default settings, and then click *Save*.

You will be given the option of adding metadata ¹ to the file. It is not necessary to include it in this case so just press *OK* and the file will be saved to the chosen folder.

- e Save and close the project.

Exit and then find and run *homer.mp3*.

This file will now play in any computer, media player or mp3 player, or it could be shared with other users by email, Facebook, Twitter, or other social media.

9. Start a new project and load some other sound files to experiment with the tools described above.

¹ Metadata is additional information about a sound file that shows up when the file is played on a digital device.

Follow up – Sound files

The sound clips we worked with in the last activity are relatively trivial, but they demonstrate the basic processes needed to manipulate sounds.

In the Activity we worked with *wav* clips and exported our final mixed track as an *mp3*. WAV and MP3 are just two ways of recording sounds.

MP3

MP3 is an abbreviation for MPEG Audio Layer III. It is a standardised format for recording audio in a digital format suitable for play back on devices such as computers and media players.

MP3 files are typically relatively small. This is because the data in them is compressed taking up much less space for storage, transmission, or streaming. They are not as good a quality as some other formats because they involve the removal of aspects of sound that are outside of the range of human hearing. Despite this they still give a fair reproduction of the original sounds for most listeners.

AIFF and WAV

Alternatives to MP3 are AIFF (Audio Interchange File Format) for Apple devices and WAV (Wave Audio File) for Windows computers. Both are uncompressed digital file formats that keep the original quality of the recording, but with the result of a much larger file size to transmit. For example, AIFF uses about 10Mb to store one minute of stereo audio, as compared to about 2 Mb for a comparable MP3 file.

Other formats

In addition to MP3 and WAV, Audacity can also work with other less common recording formats such as OGG (Ogg Vorbis), WMA (Windows Media Audio), and MIDI (Musical Instrument Digital Interface). For a full list use the drop-down arrow next to *File type* under *File > Export > Export Audio*.

Note: some of these formats, including WAV, WMA and MP3, are covered by patents.

Acquiring music files

To get music into Audacity you will need to either record or import it.

Be aware that copyright restrictions apply to most commercially available music.

Music can be digitised ready to use in Audacity either by:

- direct input from a device such as a tape or CD player
- direct input from an electronic musical device such as an electric organ that has a MIDI (Musical Instrument Digital Interface) connection
- ripping from CD or DVD
- download from the internet.

Recording music from CDs or other sources is beyond the scope of these activities, but Audacity has tutorials on how to do this [here](#).

Music ripping is the process of using a dedicated program to copy digital audio directly from CD or DVD to computer.


Downloading is importing music that someone else has previously recorded. Most downloadable music is in the MP3 format. There are numerous sources that offer such files such as app stores, or from peer-to-peer or bit torrent sites. Music files are also available through social network sharing. Again, copyright restrictions may apply.

Apart from music a wide range of sounds can be downloaded from sites such as [Freesound](#).

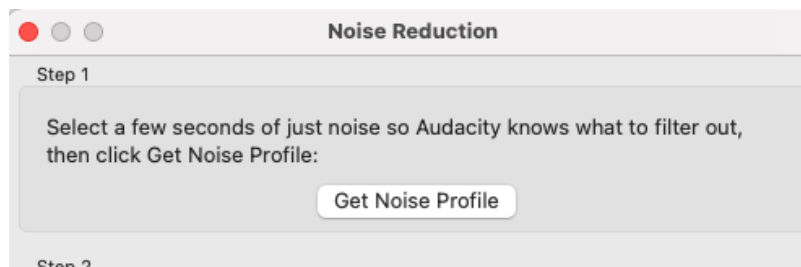
Activity 3 – Refining the mix

In this activity we will look at some more involved editing techniques.

1. The first thing we will investigate is noise removal. This is used to “clean up” tracks that have a persistent sound such as a hissing or humming in them, for example on a recording made from a vinyl record or copied from a tape.
 - a Start a new project and import the audio file *play it.wav* from the folder *Casablanca*. (This folder contains a series of sound clips from the old Humphrey Bogart movie.)
 - b Play the clip and note the background hiss on this recording.
Removing the hiss involves two steps. We must firstly identify the hissing sound, then remove it.
 - c To identify the hiss, select a part of the track with no talking on it, e.g., from around the 3 second mark:



Next from the menu choose *Effect > Noise Reduction* and click on the *Get Noise Profile* button.

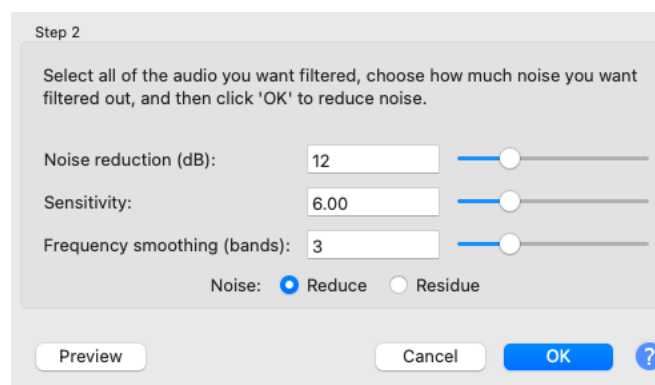


This will identify just the hissing we wish to remove.

- d Now to remove this hiss from the track.

Double click on the track to select the whole clip.

Again, choose *Effect > Noise Reduction* but this time click on the *Preview* button in the Step 2 part of the dialogue.



The clip will play with the sounds associated with the hissing removed.

This will sound a bit flat as, in addition to the hissing, other similar frequencies have been removed.

- e Experiment by sliding the *Noise reduction* slider left or right, hitting the *Preview* button after each move. This will decrease or increase the amount of noise removed.

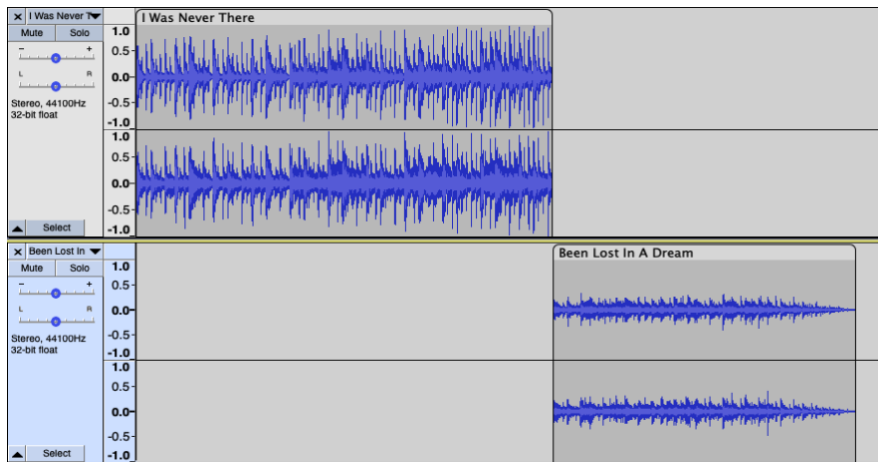

When you are happy that enough hissing has been removed, without destroying too much of the integrity of the original track, click *OK*.

- f Play the track without the hissing and see how it has improved.
- g Import another clip and again remove the noise from this track without reducing the sound quality too much.

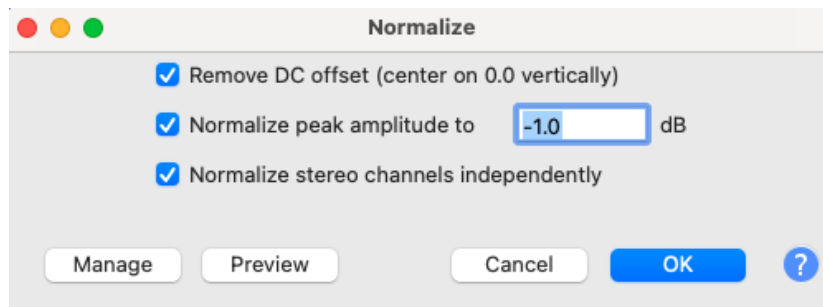
2. Normalization is a process whereby the gain of a track is increased to improve the overall strength of a sound signal without distorting it. It does this by adjusting the maximum amplitude (volume) of each track to a pre-set level. (If necessary, normalization will also centre each track around 0.)

Normalization is also useful if you want to mix recordings from two different sources or recorded at different volumes.

- a Import the two tracks *I Was Never There. mp3* and *Been Lost In A Dream.mp3*.
- b Play each track separately (use the *Mute* or *Solo* buttons) and note how the volume differs between them.
- c Drag one track to just after the other.



- d We will now normalise the two tracks. Select both tracks (*Ctrl+A / Cmd+A*) and then from the menu choose *Effect > Normalize*. Tick all three checkboxes and click then *OK*.




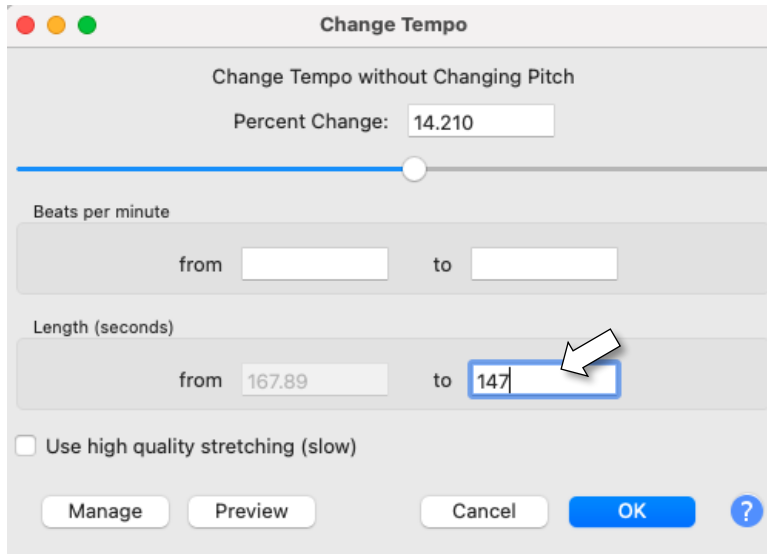
Play the sounds from near the end of one to the start of the second. See how the volume has been smoothed from one track to the next.

(Since the transition from one to the other is rather abrupt you might like to overlap them and then fade one out and fade the other in.)

Normalisation is usually the final step taken before a group of tracks are mixed so that they blend better.


3. At times we may want to vary the tempo, pitch or speed of a track. For example, we might want to make a song or part of an audio track a little shorter or longer to fit in with other events such as synchronizing with video.
 - a Import a music track and listen to a part of it to hear what it sounds like
 - b Select the whole track and then from the menu choose *Effect > Change Tempo*.

Shorten the length by about 20 seconds



Click *OK* and, after processing, note how the song now takes less time to run.

- c Play the tune and note how it runs with a faster beat, but otherwise still sounds the same. *Change Tempo* speeds up the beat of the music without affecting the pitch.
 - d *Effect > Change Speed* will also shorten or lengthen a track, but with this effect the pitch will change (though you might want this to happen).
 - e Finally *Effect > Change Pitch* will change the tone without affecting the tempo. This effect can be used to transpose a song by changing the key, which can be very useful for music students.
4. Equalization is an effect that works similarly to a graphics equaliser on a stereo system. By boosting or reducing frequencies in different ranges you can alter the overall “feel” of a sound, e.g., more bass, or more treble.
 - a Import a track and listen to it, paying particular attention to the range of sound.
 - b Select the track, and then choose *Effect > Graphic EQ*.

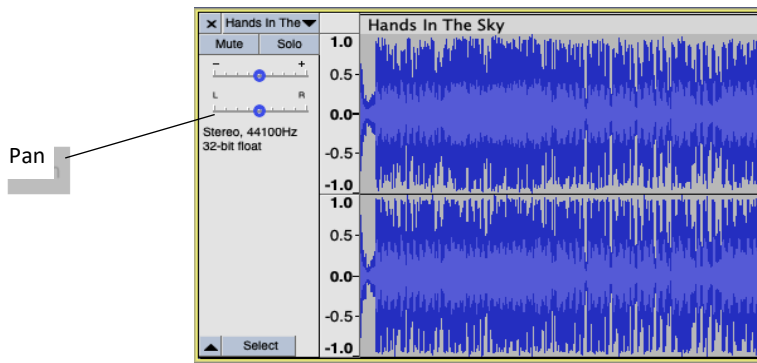


Using the sliders, you can increase or decrease the gain (volume) of individual frequencies by up to 20 dB. The frequencies to the left are the base (low) notes, those to the right the treble (high).

- c Use the *Preview* button to hear what a change sounds like.
 Note: The *Flatten* button will reset all frequencies back to their original setting. *Invert* will swap positive with negative gains.
- e Once you like the sound as previewed, click on *OK*. Audacity will go through the entire track adjusting each of the frequencies as set.
- f For more control over equalization, you can investigate the *Filter Curve EQ* effect [here](#).
 The use of equalization is a matter of taste and preference.

5. If for any reason you wish to work on just one channel of a stereo recording, we can split the two into separate tracks.

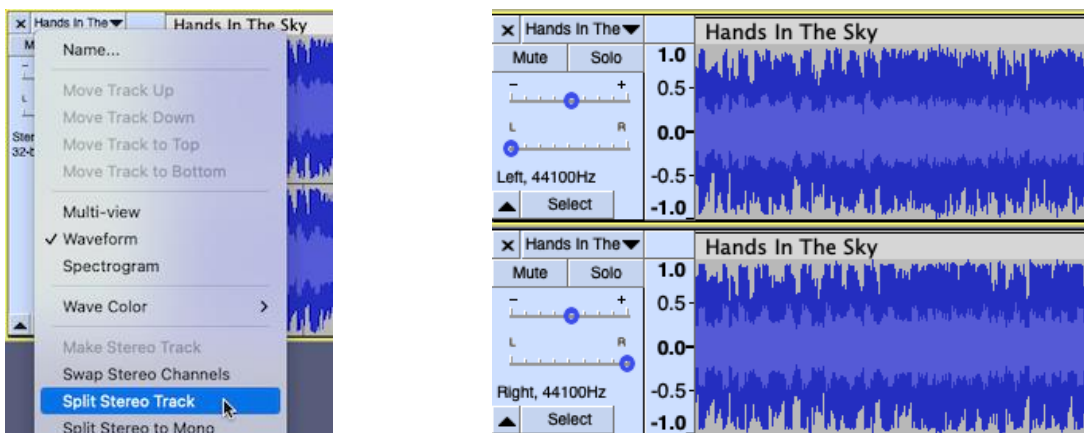
- a Import the file *Hands In The Sky.mp3*.
- b Play the song.



As it plays slide the pan control all the way left or right to hear the differing sound coming from each speaker and how the VU meter indicates the relative output to each speaker.



- c Click on the menu down arrow in the track control panel and choose *Split Stereo Track*.



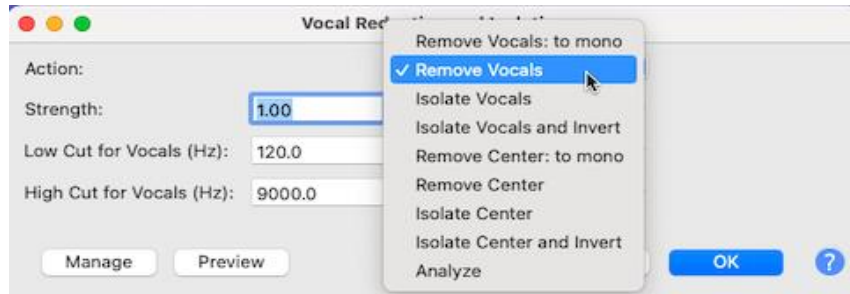
Each channel is now a separate track and can be worked on or used independently.

- d To re-join the tracks (or to join any two mono tracks) click on the down arrow on the track control panel of the *upper* track and choose *Make Stereo Track*.

6. In the last step both tracks on the song *Hands In The Sky* have voice and instrumentation. Sometimes it is useful to remove the vocal part and just leave the instrumentation behind. This can be done for purposes of practice, to make a backing for karaoke, or to just have an instrumental version of a song.

The process of voice removal is not suitable for all tracks (only so called “centre panned vocals”) but will work with the sample track.

- a If necessary re-import the file *Hands In The Sky.mp3*.
Play a short section to hear both vocals and instrumentation.
- b From the menu choose *Effect > Vocal Reduction and isolation*



In the pop-up menu that appears choose *Remove Vocals* from the drop-down list, accept the defaults, and press *OK*.

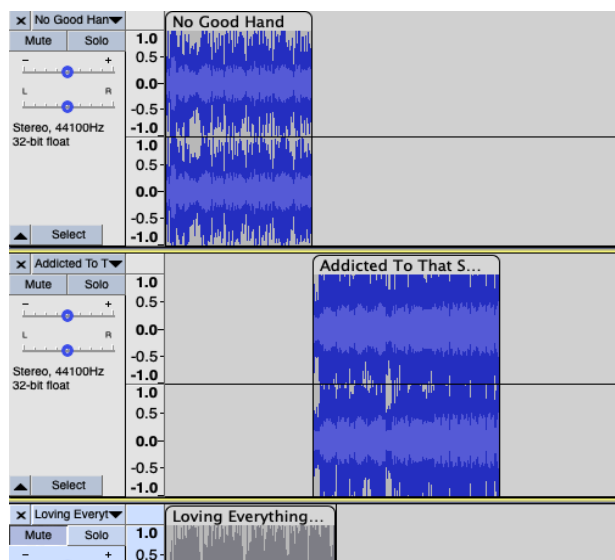
- c Play the track and see how (almost) all the vocals have been removed.
- d Repeat with a track of your choice.
If necessary, vary the frequencies as suggested in the vocal remover dialogue.
You may also like to experiment with the other options in the drop-down list.

7. To bring all this together and to finish our investigation of editing techniques we will look at creating a music mix.

A music mix is where you import a range of musical tracks of your choice and then combine them into one performance piece. A music mix is more involved than just playing one track after another as the songs should be blended to move smoothly from one into the next.

While a music mix can consist of many tracks, for this activity we will limit ourselves to mixing just three tracks. The processes used however could easily be employed to create an extended mix of 10, 15 or more tracks.

- a Start a new project and import three songs to mix together.
- b For the purposes of this activity select clips of about 15-20 seconds from each track and discard the rest. Make each clip a logical section, e.g., with a neat start and end (words or music).
- c Drag the second clip to just after the first.



Mute the third clip for now and play.

Note how the change from one clip to the next is abrupt.

- d The simplest way to blend from one clip into another is to fade one out and fade in the other.

Select the last 8 seconds of the first clip and then *Effect > Fade Out*.


Select the first 3 seconds of the second clip and then *Effect > Fade In*.

Play, and note how the change is now not as abrupt.

Remember the rule of thumb is, fade out slowly-fade in quickly.

- e Now drag the second clip to overlap the upper clip.

Again, play and see how the first song blends more naturally into the second.



- f Next drag the third clip to just after the second. Overlap them and fade out-fade in the two clips.

The three tracks now are mixed together moving smoothly from one to the next.

- g For better mixing change the tempo of the overlapping parts.

Say the one song has a rapid beat while the next is slower. Even with a fade out and a fade in the transition can sound abrupt or forced.

To overcome this, as the first clip fades out, slow its tempo down (*Effect > Change Tempo*). Next, for the time the second fades in, make it a little quicker, before it reverts to its original tempo. In this way the change is less noticeable.

To illustrate this, say the first clip has a tempo of 40 beats per minute, while the second is at 20. In this case adjust the end of the first, and the start of the second to 30 beats per minute.

Modulating between changes in tempo can be tricky but is worthwhile. When done well it provides a smoother, more professional sounding transition.

Note: the most sophisticated mixing is where you utilise a transitional key as one song modulates into the next. This is done using *Effect > Change Pitch* during the transitional time but requires a knowledge of keys and a good musical "ear".

- h Select all the clips and normalise them.

If you wish equalise to get the "feel" you want for the music.

- i To finish use *Tracks > Mix > Mix and Render* and then *File > Export*.

- j Experiment by adding more clips, or full-length tracks to a music mix.


Follow up – Sound quality

The sound quality of a digital recording is dependent on three factors, the sample rate, the sample size, and the export rate. We can set values for each of these in Audacity. We will look at each of these factors in turn but first we must appreciate what is meant by a sound sample.

Sampling

As we saw earlier, a recording is made when a microphone captures sound. If the microphone is connected to a computer, components inside the computer (e.g., its sound card) convert the electrical pulses to digital values.

To do this the original sound wave is sampled, or read, at regular points. If you zoom right in on a waveform in Audacity you can see the (millions of) individual samples, each represented by a dot on the end of a line.



The sample is a value, represented by a number, of one point on the original sound wave. By taking all the samples, one after the other, the original sound can be reproduced. In this way an analogue sound wave in the real world can be represented by numbers (digits) in a computer.

Once a sound is digitised in this way it can be edited and manipulated by programs such as Audacity.

Sample rate

The sample rate is the number of times per second the sound in the original waveform is read. For example, 1 000 samples per second is a rate of 1 kHz. The higher the sample rate the more accurately the original sound can be represented, i.e., a more faithful reproduction of a greater range of frequencies.

As a rule, the sample rate should be at least twice the original frequency. For example, human speech is in the 16-20 kHz range, i.e., the vibrations produced by human vocal cords when speaking are generally between 16 000 and 20 000 per second. For voice recording a sample rate of 44.1 kHz is usually employed. This rate is also adequate for general music recording.

Even higher sample rates are available (96 or 192 kHz) for high quality music reproduction. These can be used, but there is a penalty. The higher the rate the bigger the resultant sound file, often without a noticeable improvement in the playback quality of the music.

Sample size

The sample size (or format) is the number of bytes used to represent a sound. A larger sample size gives better dynamic range, i.e., louder louds and softer softs. A 16-bit sample size is fine for normal output; it is about CD quality. If you are looking for high quality audio or for more flexibility in editing, use a 32-bit sample size.

Sound export

The export rate (or bit rate) is the number of digits used to represent sounds, for example when we are ready to export music from Audacity. The higher the bit rate the better the quality but, again, the larger the file size.

A bit rate of 45-85 kbps is adequate for voice to be heard on low-range equipment, e.g., headphones or computer speakers. Increase the bit rate to 170-210 for music. For high quality output, for example to be played through a stereo system, set a high bit rate, e.g., 320 kbps.


Note: see the [Follow up](#) at the end of the next activity for suggested recording and export values.

Activity 4 – Recording sound

So far, we have imported the tracks we wish to use, but Audacity also has the facility to record sound. The way this is most used is to record voice, although music or sound effects can also be recorded directly into Audacity.

Note: to input sound into Audacity you will require a suitable microphone connected to your computer. Check the microphone is working before starting this activity.


1. We will begin by setting the input level for recording. You can either record your voice, or alternately set the level using music that is playing, e.g., on a radio or CD.
 - a Connect a suitable microphone to the computer and start a new project in Audacity.
 - b Click on the recording VU meter at the top to *Start Monitoring*.




Monitoring is using the microphone for input but is not making a recording. It is, in essence, “listening” to the current sounds and outputting the sound levels on the VU meter.

- c Speak into the microphone or move the microphone near to music that is playing.

If the microphone is working a green bar will appear. This will show the volume currently being read, the peak volume in the last three seconds, and the maximum volume that has been input.




Again, a red *clipping indicator* will appear at the right-hand end of the VU meter to show if the level is so high the sound will be distorted.



Note: even though a L and R channel are displayed on the meter, standard microphone recording is mono only. (It is possible to use stereo microphones, or to have multiple input microphones, but that is beyond our investigation of Audacity.)

- d Move the input volume slider left and right and note how this affects the levels in the VU meter.



- e Adjust the input volume slider so that the peak input volume is about -6 dB.



The best level to record is between -12 dB and -6 dB; do not exceed -2 dB. (You can click and drag to make the VU meter wider if you wish.)

This is now ready to record.

- f To stop monitoring click on the recording VU meter.

Note: out of interest you might like to investigate the VU meter options. To do this click on the microphone at the left of the meter and choose *Options ...*

For example, you could have a vertical display or use a linear rather than decibel scale. (If using the linear scale, set recording maximum in step e above to 0.5.)

2. Now that we have used monitoring to set the recording level, we can make a first recording. (The steps below assume you are making a voice recording. If you wish replace the voice with music.)

a To start we will make a test recording.

Click on the *Record* button and speak a sentence into the microphone. Speak close to the microphone (about 10cm / 4 in.), but not too close (to avoid distortion).




When finished click the *Stop* button.

b Press *Play* to hear the recorded sentence.


c The recording will show as a mono track (unless you have a stereo microphone, or a sound card which can input more than one channel).



Look at the recorded waveform – there should be no clipping visible. Clipping is where the volume of the sound is louder than Audacity can record. When this happens the tops and bottoms of the recorded wave are chopped off ("clipped").




good recording



recording with clipping

Clipping results in a distorted sound. If the level was set correctly in step 1 above, there should be no clipping.

Note: to see where clipping takes place use *View > Show Clipping* from the menu.



The vertical red lines indicate anywhere clipping has occurred. Use *View > Show Clipping* again to turn this off.

d Normalise the test recording to maximise its volume.

- e Finally mix and export the test recording (*Tracks > Mix > Mix and Render; File > Export*).
Listen to the recording on the device it is to be played on. In this way you can check the levels and quality chosen are suitable for your equipment.
If all is well, you can continue with the current settings. If not, you may have to alter Audacity's default settings. (See *Follow up* at the end of this activity.)

3. Now that we have finished our test, we can make a full-scale recording. Firstly, we will make a simple monologue.

- a Delete the test track recorded in the previous step.
- b Get ready to make the full-scale recording. If recording voice it is best to have a script to follow, or a book to read from. (If using a music source, have it ready to play.)
- c Press the *Record* button and make the recording. *Stop* when finished.
Save the project to a suitable folder.
- d Listen to the recording.

Use the editing techniques described in [Activity 2](#) to remove mistakes, unwanted silences, overlong pauses, hesitations, transient sounds (coughs, chair creaks), or other distractions.

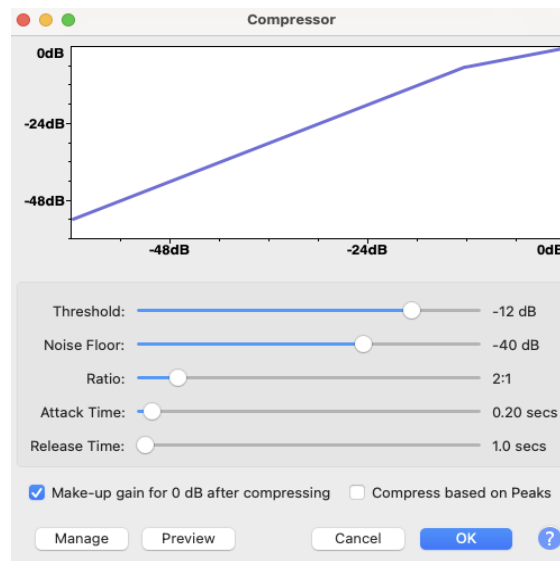
This step can take some time. Use labels if necessary to help identify sections of the recording.

Re-save as you go.

Note: after each edit, select before and after where you made the edit (e.g., cutting out a cough), and listen to make sure the flow still sounds natural.

- e At this stage rather than normalise we will level out the varying voice levels in the recording. (Do not do this step for the recording of music.)

Select the entire track, then choose *Effect > Compressor*. Accept the defaults and click on *OK*.



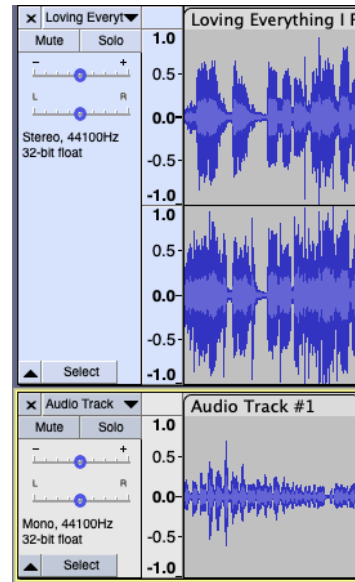
Listen to the result. The various different volumes of voice have been smoothed out to a relatively constant volume.

To achieve this the *Compressor* makes the loud parts quieter, then it amplifies everything. The result is that the quiet parts end up at the same volume as the rest of the track.

If the quiet parts are still too quiet, undo and then try again with a *Threshold* setting of -18 dB. If on the other hand, if the voice sounds unnaturally squashed, undo and try again with a *Threshold* setting of -6 dB.

- f To finish, save your project, then mix and export.

4. In addition to a monologue you can use Audacity to overdub. An overdub is where speaking is added to an existing track, e.g., adding a voice over music.
 - a To do this, first set levels as describing in step 2 above.
 - b Next load in the track to be overdubbed.
 - c From the menu choose *Transport > Transport Options > Overdub*.
 - d Holding down the *Shift* key press the *Record* button and record 4-5 seconds of voice while the music plays.



An audio track will be added below the music track.

- e *Play* and check that the voice track can be heard over the music.

If necessary, alter volumes. (It is easier to make the music track quieter.)

Test again until you have a good blend of voice and music.
- f When ready continue as for normal recording, adding your voice track as an overdub to the music track. Save, normalise, mix, and export to finish.

Note: you can also overdub drums or other sounds onto music tracks.

Follow up – Making a recording

Using a microphone

As indicated in step 2 above when using a microphone to record voice speak close to it, holding it about 10-15 cm or 4-6 inches away, but not too close to avoid distortion or “popping”.

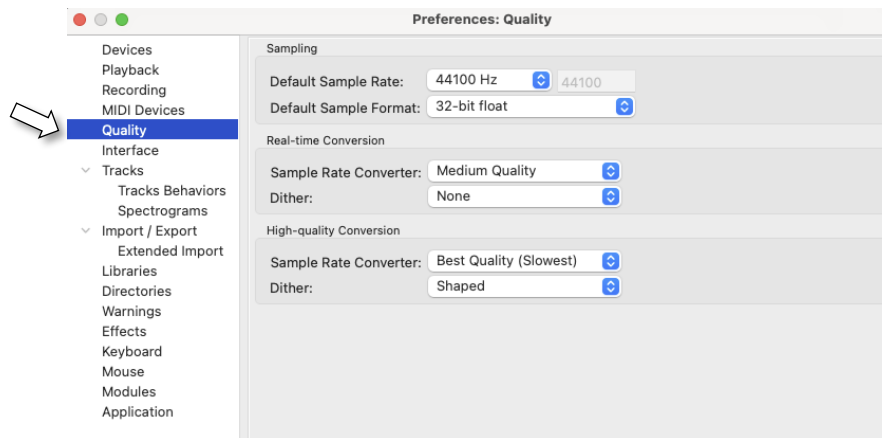
Here are a few other hints for voice recording:

- location – position the microphone so that it picks up a minimum of ambient (surrounding) sound such as air conditioning, computer fan, or other noises
- placement– keep your head still and place the microphone so that you are talking past it, e.g., level with your mouth but a little off to the side
- voice – speak in a normal, conversational tone of voice, but also speak clearly and enunciate carefully; it may help to speak a little slower than normal.

Recording settings

On the menu use to access the preferences settings. (Windows: *Edit > Preferences*; Mac: *Audacity > Preferences*)

Click on the *Quality* tab.



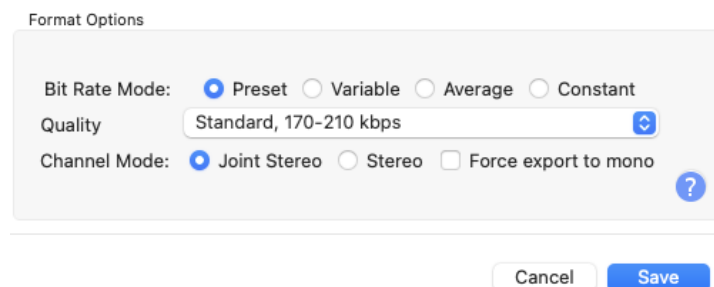
The default sample rate and sample formats are 44100 Hz and 32-bit float. These are adequate for recording music. You can increase the sample rate to 96 kHz or even 192 kHz for high quality recording of music, but these rates result in a much larger file size. The higher sample rates also may be excessive in situations in which playback is to be on low-range equipment such as headphones. In this case a much larger file size will result in no discernible improvement in playback quality.

For voice recording a 16-bit sample size is sufficient and will result in a smaller file size. You can also change the *Channels* under the *Devices* tab to *1 (Mono)* if you are using a standard microphone.

Export settings

Load a soundtrack into Audacity, and then choose *File > Export > Export as mp3*.

In the dialogue that appears under *Format Options* you can see the default quality is 170-210 kbps:



This is good for CD quality. To reduce file size, you choose *Variable* as the *Bit Rate Mode* and drop this to 45-85 kbps for a voice-only recording.

Cancel without saving and exit Audacity.

In summary the settings are:

Quality required	Recording		Export bit rate	
	Sample rate	Sample size	mp3	uncompressed
Voice only	44.1 kHz	16 bit	45-85 kbps	16 kbps
Music	44.1 or 96 kHz	16 or 24 bit	170-210 kbps	16 kbps
Very high	192 kHz	32 bit	320 kbps	32 kbps

Creating a podcast

A podcast is downloadable audio file such as a radio show. You can create and distribute your own podcasts of voice, sound, or music-based projects. The podcast can be played on an iPod, a similar mp3 player, on a computer, or can be uploaded to a blog.

The process for recording a podcast is:

1. Lay down, record, or import the tracks you wish to use.
2. Record the intro to the podcast, links between tracks, and the conclusion (aka outro).
3. Trim the tracks of unwanted audio and clean up any noise or transient sounds (such as coughs).
4. Sequence the tracks into a logical order, adding intro, links, etc., as well as any music to be used.
5. Smooth the transitions between tracks.
6. Normalise the tracks or use the compressor.
7. Make a final check by listening through the entire recording and then correcting any problems.
8. Mix and encode.
9. Share and listen.

Glossary

- AIFF (Audio Interchange File Format) – Macintosh default sound format (uncompressed)
- amplify – to increase signal strength (usually measured in decibels)
- amplitude – the height of a wave; in representing sound the higher the wave the louder the sound
- analogue – storage of information in a continuous (non-digital) form
- audio – recorded sound; this includes music, voice, noises, effects or any other sounds
- background music – music played behind a dialogue or narration to add atmosphere
- bit rate – the number of bits per second used to encode a signal; the higher the bit rate the better the quality
- channel – audio communication means; in stereo the left and right channels can carry differing signals; in mono the left and right speakers can be set at differing volumes
- clip – a small selection from a longer track
- clipping – distortion that results when input volume exceeds the capabilities of the equipment used
- codec – program to compress a file to a smaller size
- compression – reducing the size of a file to save storage space or transmission time; achieved using a compression algorithm or codec
- credits – a spoken list of names at the end (or beginning) of a podcast acknowledging activities undertaken
- cross fade – a transition where one track fades out as the next fades in
- decibel (dB) - a measure of the relative loudness of a sound (this is a logarithmic measure)
- dialogue – spoken parts in a recording
- digitised – audio converted into a form that can be used by computer; usually represented as numbers composed of 0s and 1s
- duplicate – isolate a part of a track to relocate it or to experiment on it
- edit – to make changes to a track; includes splitting, joining, deleting, re-arranging, and mixing
- effects (FX) – created or changed sounds used to add colour to, or to enhance a recording
- encode – convert a track or tracks to a digital format such as mp3 or wav
- envelope – a broad outline of the general limits of the waveform of a track
- equalise – to increase or lower the level of selected frequencies; used to modify output sound to suit individual taste
- export – to take all of the audio in a project and save it as a file in *wav*, *mp3*, or other format
- fade – a gradual decrease or increase of volume
- fade in/fade out – increase/decrease of volume over a selection of a track
- Foley – incidental sound effects added to a recording such as footsteps, door slam or phone ring
- format – the way chosen to store data on computer e.g., *wav*, *mp3*, etc.; also sample size
- frequency – how often something occurs; in waveform the closer waves are together (more frequent), the higher the pitch
- Hertz (Hz) – a measure of frequency; 1Hz is one event per second (usually vibrations)
- I/O – input/output; generic term for such devices
- ID3 tag – meta-information such as title, performer, track data or other information about a sound file, to be stored in the file itself
- iPod – a brand of portable digital media player
- import – to upload audio from a memory device and convert it ready to be used by a computer
- input – to load information into the computer; audio is input from a sound file or by a recording device such as a microphone
- interview - a recorded dialogue between (usually) two people where questions are asked to obtain information or to determine a point of view
- intro – an introduction or lead into a track, usually spoken
- label/label track – a text identifier lined up with a distinct section of a waveform
- LAME encoder – third party dynamic library used by Audacity in converting a file to *mp3* format
- link – a joining track added between two other tracks to provide continuity, usually spoken
- microphone (also mike or mic) – a device to convert voice or sounds into electrical signals
- MIDI (Musical Instrument Digital Interface) – a protocol (standard) used so that electronic musical instruments can communicate with each other and with computers
- mix – to merge two or more tracks into one
- mono – a recording with only one distinct channel of sound
- monologue – a single-voiced recording

MP3 – MPEG-1 Audio Layer 3; a popular form of compressed digital encoding for audio tracks

mp3 player – a portable device designed mainly to record, store and play mp3 tracks

narration – a voice added over the top of an audio track or other presentation

noise removal – the process of taking out any unwanted, persistent sounds from a recording

normalise – the process of correcting any vertical offset of a track and adjusting the maximum amplitude to a preset level; normalization is usually carried out just before mixing

OGG (Ogg Vorbis) – patent-free, audio encoding and streaming compression format similar to mp3

output – to produce information from a computer; audio is output to a sound file or speakers

outro – (opposite of intro) a closing or conclusion to a track, usually spoken

overdub – to add a voice over to an existing track

overlap – where one track is positioned so that it begins before another track is finished; both tracks will be heard at the same time

pause – to temporarily hold up the playback of a track

pitch (or tone) – the frequency of a sound; the higher the frequency the higher the tone

playback indicator – device to indicate on a timeline the current play position of a track

podcast – a downloadable audio file; usually made available for others to listen to

popping – distortion that occurs when plosives (p, b, etc., sounds) are spoken too close to a microphone

producer – person responsible for the infrastructure of a recording production

project – a device to hold in one place all imported tracks, recordings and edits for a particular task

record (v) – the process of inputting and storing audio

recording (n) – pre-input audio stored in some useable format

rip, ripping – the process of copying digital audio (or video) from one form to another, usually CD to computer file

sample – the part of an analogue signal selected to be recorded digitally

sample rate – the number of samples taken each second, measured in Hertz (Hz)

sample size (or format) - the number of digits used to represent a sample

script – the story line and spoken parts of a story, or the text of a narrative

selection tool – device used in Audacity to choose a track or to mark a section of a track

signal – transmitted values such as audio output

silence – no sound; the process of replacing a selection of a track with no sound

sound bite – a short selection taken from a longer audio track; used to emphasise or make a point

sound card – a computer circuit board used for the input, processing and output of audio

split – to cut a selection from one track and place it on a new track; the cut part is replaced with silence

speaker – a device to convert recorded electrical signals into audible sound

stereo – a recording with more than one channel of sound

sting – a short, catchy musical phrase used to distinguish between sections of a program or as a characteristic identifier for the program

storyboard – a series of panels outlining the sequence and major changes of action or plot in a production to be recorded

timeline – a chronological representation of audio in a recording

time shift – to move the start point of a track relative to other tracks

track – a single audio recording

transient – a random unwanted sound on a track such as a chair creak or a cough

trim – to delete all except for the selected part of a track

uncompressed – a file still at its original length

voice over – a narration added to an existing track

vox pop (Latin: voice of the people) – informal interview with random people to extract general opinions

VU (volume unit) meter – an analogue display of volume levels averaged out to indicate perceived loudness

WAV (Wave Audio File) – Windows default sound format (uncompressed)

waveform – a method of visually representing sound as a curved line with peaks and troughs

WMA (Windows Media Audio) – a compressed sound format

writer – person who prepares the script for a production

zoom tool – a device used by Audacity to enlarge or shrink the detail that can be seen

Download and set up Audacity 3

Download

Audacity is free and can be downloaded from: audacityteam.org/download.

Click on the link for your operating system and the download should start automatically. Follow the on-screen instructions for installation.

Set up

Before Audacity can be used with some files it is necessary to instal the *LAME* encoder² and the *FFmpeg* suite³. (These are third party applications and for copyright reasons cannot be included in the initial download of Audacity.)

These two are available from:

Windows: lame.buanzo.org/-lamewind/.

TO DOWNLOAD Lame and FFmpeg for Windows, click links BELOW:

The Audacity(r) QA Team suggests users download the ZIP version instead of the .EXE or .DMG (for Mac) versions. If you use the installers, and Audacity does not detect LAME, download the ZIP option, extract the files inside to a well known folder, then open Audacity, go to Library Preferences and configure it to search on the well known folder you extracted the files to.

NOTICE: MacAfee and at least one other antivirus flags this site as unsafe, which is a false positive. FILES I HOST ARE SAFE, but always check your downloaded files with www.virustotal.com.

RECOMMENDED Installer Package for Windows: [Lame_v3.99.3_for_Windows.exe](#) - (SHA256 SUM [here](#))

ZIP OPTION:
[libmp3lame-win-3.99.3.zip](#) (Issues? Some help [HERE](#))

If you need or want a newer version of Lame, because of the performance improvements with newer AMD and Intel processors, here is [v3.100 in a ZIP file](#), known to work with Audacity 2.x. You can also find [3.98.2](#) that is known to work with 1.3.x

You can also download the standalone [lame v3.100 exe by clicking here](#) - ([Virustotal Scan here](#), dont mind the one false positive with Qihoo-360 antivirus)

FFmpeg for Windows and latest Audacity versions
FFmpeg RECOMMENDED ZIP OPTION: [ffmpeg-win-2.2.2.zip](#)

FFmpeg 2.2.2 Binary for Windows, compatible with Audacity 2.0.6 and later (please update, or use v0.6.2 below)
[ffmpeg-win-2.2.2.exe](#) - (SHA256 SUM [here](#))

(FFmpeg is not for MP3 export for that you need LAME, see above)

FFmpeg for older versions of Audacity
FFmpeg 0.6.2 Binary compatible with Audacity 1.3.13 to 2.0.5 only, on Windows:
[FFmpeg_v0.6.2_for_Audacity_on_Windows.exe](#) - (ZIP version - [here](#))

² Used to export MP3 files from Audacity.

³ A suite of libraries and programs for handling video, audio, and other multimedia files, and streams.

Mac: lame.buanzo.org/#lameosx64bitdl

FFmpeg and LAME on macOS / Mac OS X

64-bit libraries for Audacity 2.3.1 and later

Audacity 2.3.1 on Mac is now a 64-bit application and thus requires 64-bit libraries for LAME and Ffmpeg.

It is recommended to use the .pkg installer versions (ZIP version is provided as alternative, PKG is better than DMG).

64-bit LAME for Mac OS X

Mac OSX 64 bit LAME mp3 library for Audacity 2.3.1 and above:

RECOMMENDED download: [lame_64bit_osx.pkg](#)

ZIP Alternative zip download: [Lame Library 64bit MAC OSX for Audacity 2.3.1 and above.zip](#)

For installation instructions please see [LAME installation](#) in the Audacity Manual.

64-bit FFmpeg for Mac OS X

Mac OSX 64 bit FFmpeg library for Audacity 2.3.1 and above:

RECOMMENDED download: [ffmpeg_64bit_osx.pkg](#)

Alternative zip download: [FFmpeg64bit MAC OSX for Audacity 2.3.1 and above.zip](#)

For installation instructions please see [FFmpeg installation](#) in the Audacity Manual.

Again, download and follow the on-screen instructions for each.

Audio files

The audio files referred to in these activities can be accessed from:

edit.net.au/files/Audacity/Audacity_resources.zip

This is a *zip* (compressed) file. Click on the link to download the file.

When done open *Windows Explorer* or *Mac Finder* to your *Downloads* folder and double click on the file to access its contents.

The files will need to be made available to any device that that is running these activities.

The music files in the download are sourced from:

freeplaymusic.com

and are licensed for in-classroom use.



Further copyright-free sample files can be accessed from this source.

Additional help

[Audacity manual.](#)

[Audacity wiki.](#)

