

Preparing data for open access: a guide to basic workflows

Introduction

When it comes to our language documentation collections, we all want to show our best work – but what happens when some of our best recordings involve a bystander stepping into the frame or a speaker revealing a piece of sensitive information? It can be disappointing to restrict access to an otherwise unproblematic recording over just a few seconds of sensitive information or a bystander hovering in the background. This guide walks you through some simple audio and video editing techniques for redacting sensitive information and blurring faces or figures to make as much of your data as possible ready for open access. We will cover the following:

- 1) An overall workflow for processing audio and video data, evaluating if it needs editing, and evaluating whether it is worthwhile to invest the time to edit it
- 2) How to redact sensitive information in audio using brown noise in Audacity
- 3) Blurring faces or figures using a Gaussian blur mask in Adobe Premiere Pro

To follow along with this guide, you will need to install the following software:

- Adobe Premiere Pro (2020 version)
- Audacity (any version from the last few years should work)
- A media player of your choice (e.g., VLC, Windows Media Player)
- Transcription software of your choice (e.g., ELAN, Praat)

Audacity is free, but Premiere Pro is not. If your institution does not offer it, it can be paid for on a monthly subscription basis, which is generally affordable. If you have a funding, you may even be able to get reimbursed for the expense – it can't hurt to ask! If your budget is very tight, you will want to make sure you triage your data to identify priorities for video editing (see following section) and try to block out a month where you will have enough time to invest in getting your editing done.

A basic workflow for preparing data

A basic workflow for preparing open-access ready data involves the following four broad steps:

1. Extract WAV audio and concatenate and encode videos, if necessary, and make any global audio changes (e.g., panning single channel stereo to mono), using your normal workflow
2. Triage recordings to identify portions that need redaction or editing, the extent of the edits that will be needed, and any relevant time stamps
3. Make redactions to audio
4. Make edits to video

In this guide, we will not go over the types of tasks listed in step 1, as these tasks vary based on different equipment and other factors, but these should be completed first before making any audio or video edits. This will save time in the long run.

This guide also does not cover how to make decisions regarding what to redact or blur – these choices are dependent on your IRB protocol, the desires of the project participants, and the ethical and social norms in their communities.

Triaging recordings

The first step is to identify which recordings will need to be edited. In reality, this step starts while you are still in the field:

- Keep a notebook on hand to jot down the presence of an interloper or note that the recording was blissfully uninterrupted can help save time later.
- As much as possible, try to watch and review recordings with participants while you are still in the field to identify any sensitive information that you may not understand.
- As you segment and transcribe, use an additional tier to note of portions that will need editing.

For any recordings that you have not evaluated in the field or during transcription, you will need to sit down and watch and listen to them. As you watch, jot down the approximate timestamp ranges for any problematic portions. For recordings that you are confident are “clean”, you can speed up this process some by increasing the playback speed in VLC or ELAN.

As you evaluate your recordings, here are some things to keep in mind:

- Does the recording need audio editing, video editing, or both?
 - Audio editing tends to be much faster than blurring faces
- If the recording needs audio editing, how much and how frequently?
 - If the overall content of the recording concerns something sensitive, it may be best to just restrict it
 - If the sensitive information is prominent, central to the discourse, or repeated often, it may be best to just restrict it
 - If the sensitive information is just a few tid-bits here and there (e.g., a name or a 2 second offensive epithet), redacting the audio is a good way to make it acceptable for open access.
- If the recording needs faces blurred, how many interlopers are there and how long are they in the frame?
 - Blurring is easiest when there is only one subject; the more subjects that need to be blurred, the more work it will take.
 - Blurring fastest when the subject(s) are only in the frame for a short duration (e.g., someone walking past behind the speaker).
 - For videos where there are interlopers for long periods of time, it is faster to blur faces that are stationary than faces that are moving around a lot.
 - By “stationary”, I mean **really** still – someone who is swaying or turning their head side to side will require frequent manual adjustments.
 - When subjects are very far in the background, facial features are often grainy, and it may not be necessary to blur. In most cases, we personally know the person in the background and they are recognizable to us based on their clothing and the context more than the actual features of the face. Zoom in on their face and ask yourself if it is actually recognizable by itself.

I recommend starting with the easy cases first (e.g., redacting single words or phrases or blurring a single subject walking by), then gradually working up to the more complicated ones. You will get faster with experience!

Redacting audio with Audacity

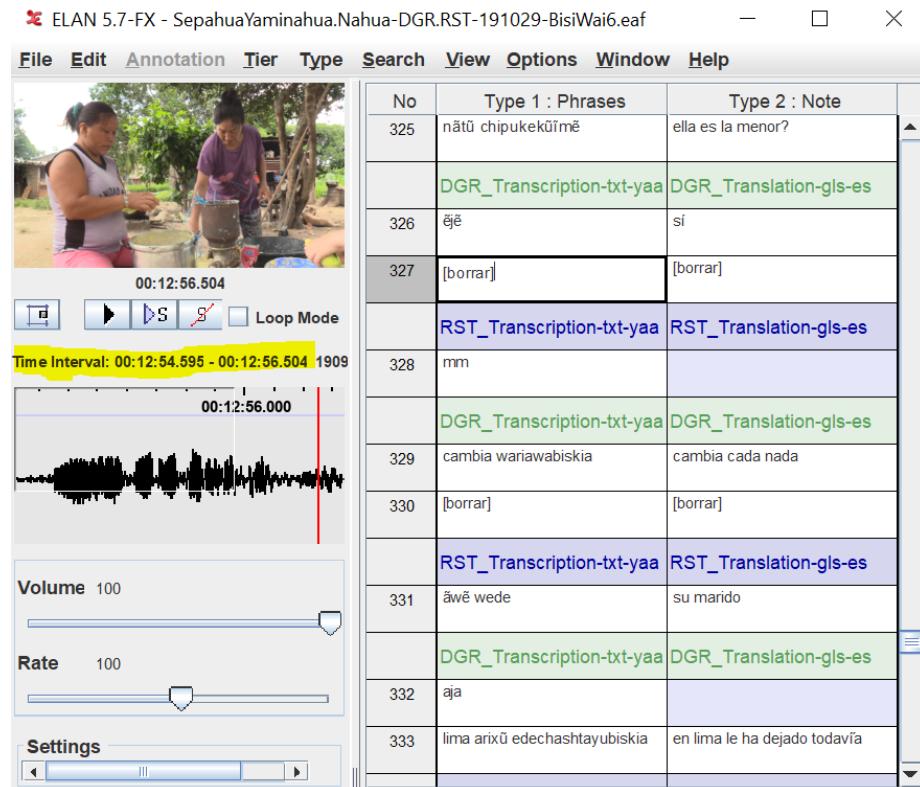
It is generally easy and fast to redact audio that has already been segmented, transcribed, and evaluated (as described in the preceding section). Keep in mind that if the redactions are so frequent or lengthy that they interfere with the ability to understand the content, that it may be best just to leave it as-is and restrict access, or, in the most extreme cases, to delay archiving until you can review the recording fully with the speaker(s).

To start, you will need a .wav audio file.¹ You should already have already listened to and evaluated this file for portions that need to be redacted. The first step is to identify the start and end times of the portions to be redacted. This is very easy in ELAN. As long as you have labeled segments that need to be redacted consistently, as in this example, you can use the search function of ELAN to locate them. The time intervals are displayed in the search results:

FIND							
An annotation on any tier in the set that matches string borrar							
>	Nr	Tier	Annotation	Parent	Child	Begin Time	End Time
	1	DGR_Trans...	[borrar]		[[borrar]]	00:12:54.595	00:12:56.504
▶	2	DGR_Trans...	[borrar]		[[borrar]]	00:12:59.315	00:13:01.064

I usually write the timestamps down in a list on a notecard for quick reference.

If you need to check something, clicking the results will take you to the relevant part of the recording.



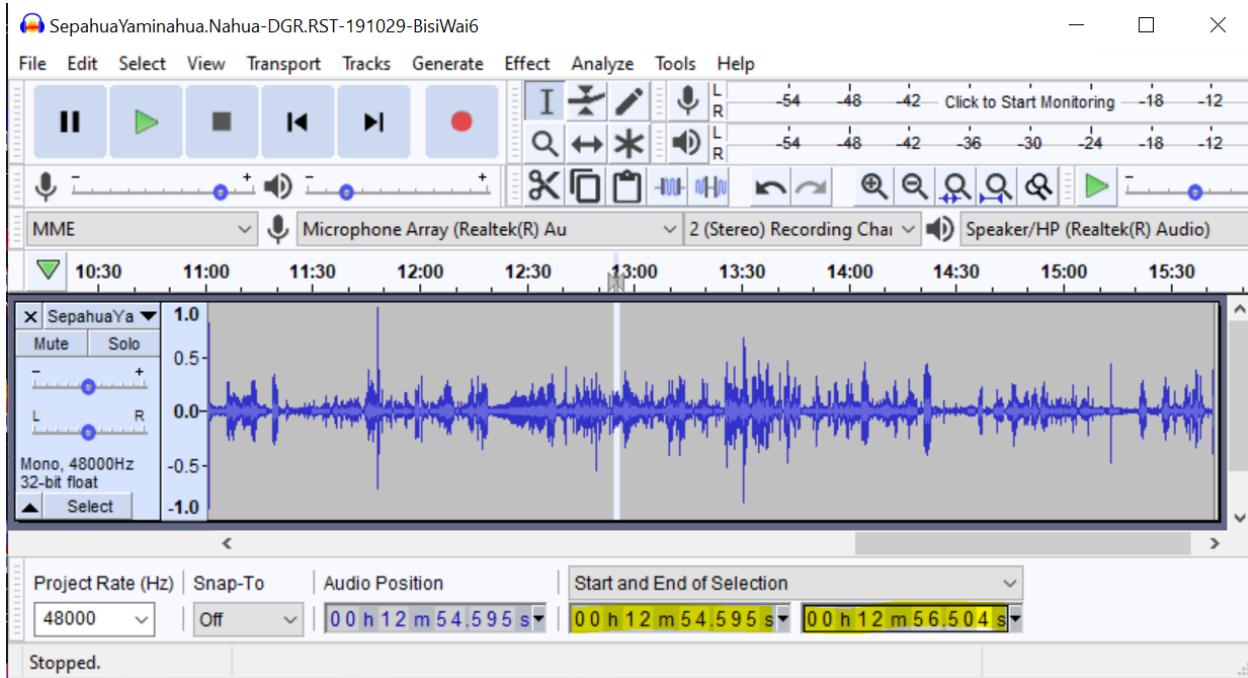
No	Type 1 : Phrases	Type 2 : Note
325	nātū chipukekūmē	ella es la menor?
	DGR_Transcription-txt-yaa	DGR_Translation-gls-es
326	ejé	sí
	[borrar]	[borrar]
	RST_Transcription-txt-yaa	RST_Translation-gls-es
328	mm	
	DGR_Transcription-txt-yaa	DGR_Translation-gls-es
329	cambia wariawabiska	cambia cada nada
330	[borrar]	[borrar]
	RST_Transcription-txt-yaa	RST_Translation-gls-es
331	āwē wede	su marido
	DGR_Transcription-txt-yaa	DGR_Translation-gls-es
332	aja	
333	lima arixū edechashtayubiska	en lima le ha dejado todavía

In some cases you may need to select a single word or short phrase from within a longer segment.

To do this, click and drag the cursor in the waveform box. An info box with the time interval of the selection will pop up.

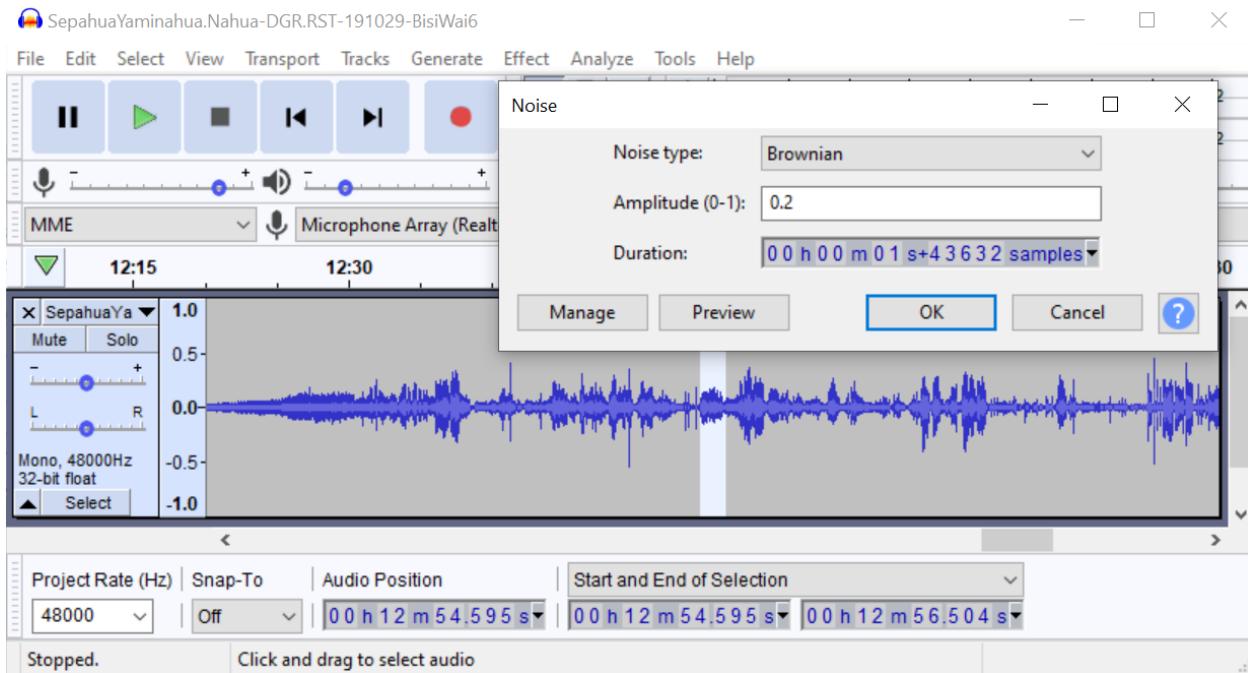
¹ You can extract the audio from a video file and subsequently replace the audio with the modified track using FFmpeg.

The next step is to open the .wav file in Audacity. Just click and drag the file to the shortcut to open. Once it is open, type the start and end times into the selection box at the bottom of the window:



Use View -> Zoom -> Zoom In to get a better view of the selection.

To redact the audio, click on the “Generate” menu -> “Noise”, then select Brownian from the drop-down menu. Using the amplitude of the selected waveform as a guide, take a stab at the amplitude and enter that into the box.



You can preview the noise to see if it seems excessively quiet or loud. My advice on the amplitude is to always estimate on the low side, or even go a little lower if the selection is longer than .2 seconds. After you have generated the noise, play the recording to see how you like it. You can undo and make adjustments if needed.

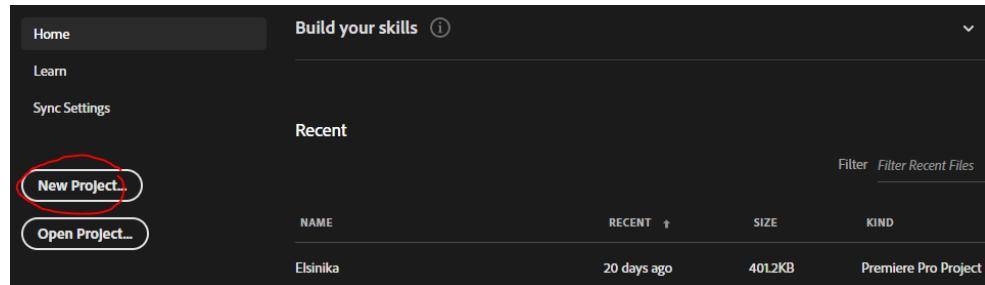
The Generate menu provides several different options including Noise, Silence, and a Tone. Personally, I find the tone to be a bit jarring – it sounds like “bleeping out” curse words on network reality TV, and calls attention to the fact that something has been obscured. Silence is a better option, but it can sound like there was a problem with the audio cable. My preferred option is to redact using noise roughly calibrated to the amplitude of the surrounding speech. This sounds like a puff of air into the microphone – for redactions of brief segments, it does not disrupt the “flow” of the discourse, and a casual listener may not even notice it. Brownian noise is perceptually lower pitched than the other noise options, white and pink, but these are worth experimenting with.

Repeat the process for each segment that needs to be redacted. When you are done, use File -> Export -> Export as WAV to export the redacted recording. I usually give my redacted .wav files the same file name as the original, but add the suffix “-redacted” to distinguish it. If the audio file corresponds to a video, use FFmpeg or whatever software you usually use to change the audio track on the video that you plan to archive. You’re done!

Blurring faces in Adobe Premiere

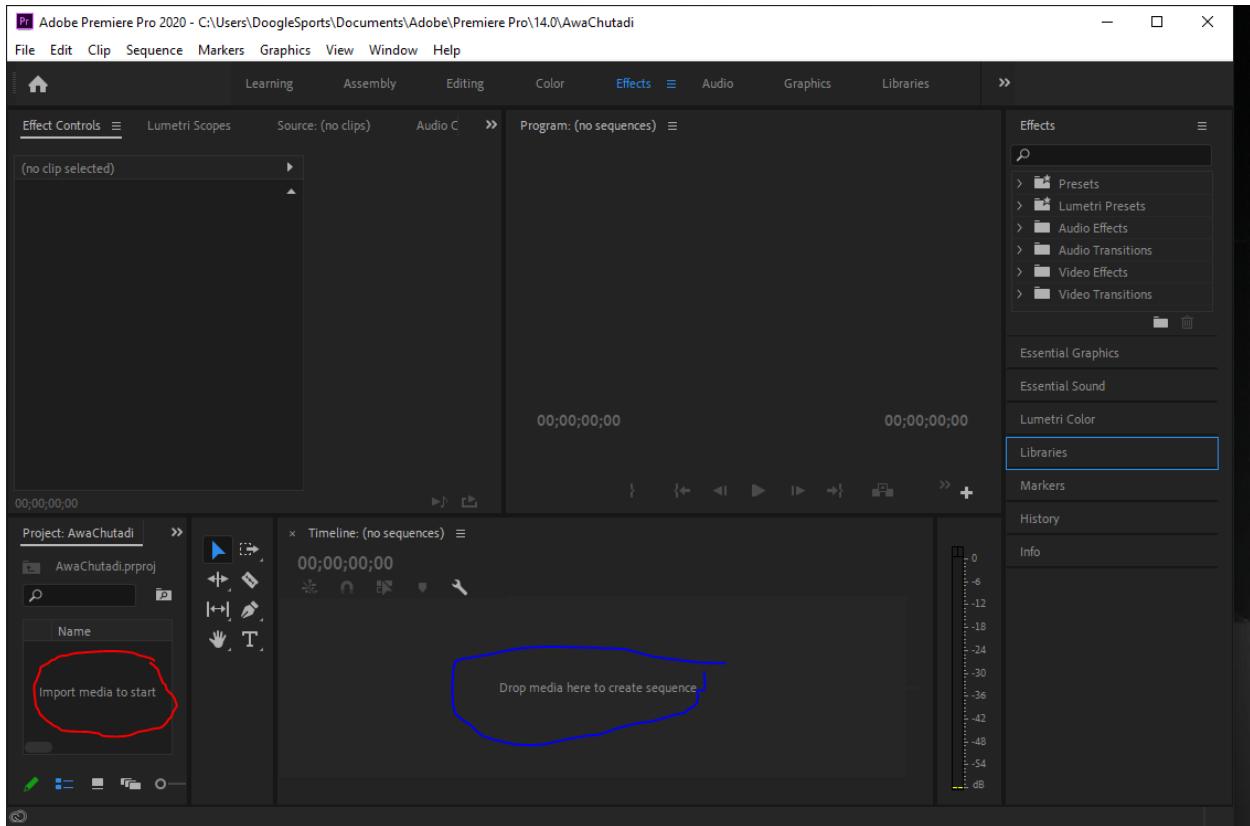
To blur faces in Premiere, you will need a video file that has already undergone any audio processing or redaction that is necessary. You can also go back and do audio redaction later, but it will streamline your workflow if you do it first.

To start, open Premiere and open a New Project:

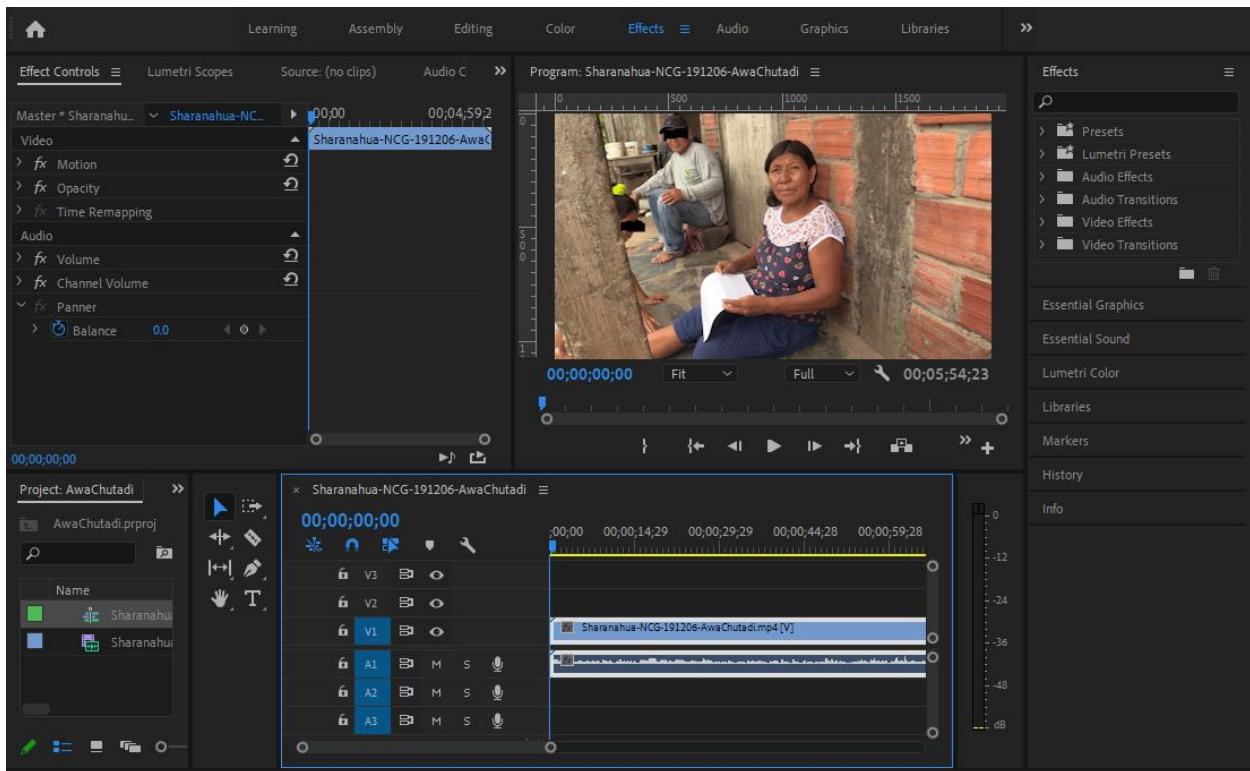


Give the Premiere project a descriptive name -- this could be identical to your video file name, but could also be a shorthand version. The Premiere project just saves your editing progress -- you will give the resulting, finished file a “real” name when you export at the end of the process. The file path is similarly used just for the project save file. The default settings below the name and file path are just fine and you need not change them.

Starting a new project will open a window with lots of widgets. Drag your video file into the box that says “Import media to start” (circled in red). Then, click and drag it from that box to the area labeled “Drop media here to create sequence” (circled in blue) in the Timeline box to the right.



It may take a few moments for the video to load in the Timeline. When it has, the window should now look like this:



Before we get into the actual blurring process, I want to introduce some terminology:

Effect: a type of modification we are making to the data, in this case a blur

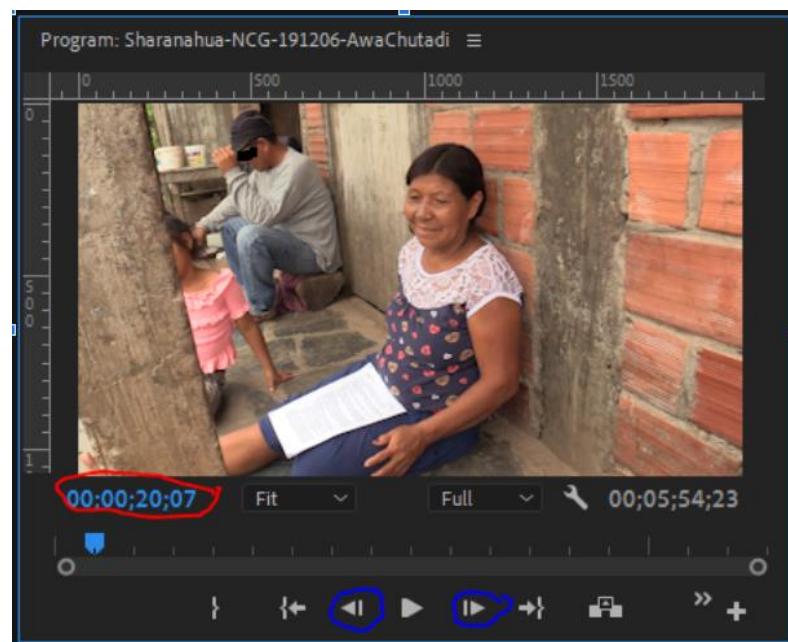
Mask: a shape that defines the area where an effect should be applied

Clip: a temporal segment of the video where a mask will be applied.

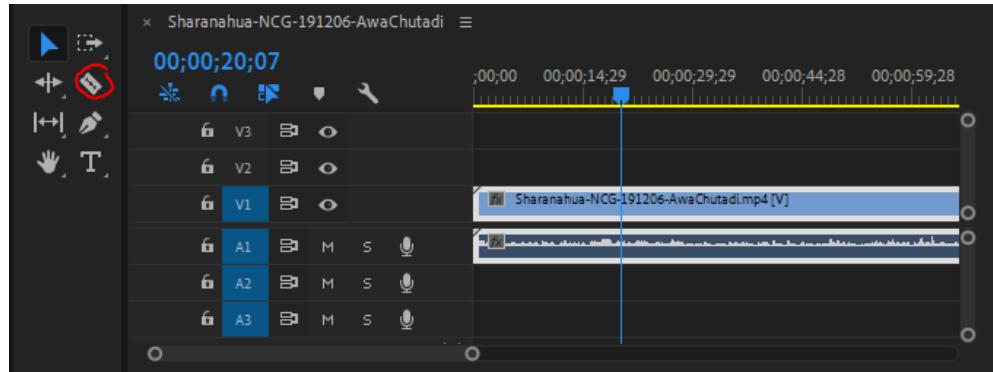
The first task is to create clips. If you need to blur a face that is present and visible for the entire duration of the video, then the entire video is your clip and can move to the next step. In most cases, however, you will have a face that only appears for a few seconds or minutes. Masks span over clips, so we need to define the specific parts of the video where we want them to apply. This not only makes the editing work simpler, but it speeds up processing.

Using the timestamps that you have already identified as a guide, start marking when faces appear and disappear in the recording. If you have multiple faces, start with the one that is present for the longest durations -- I will talk about how to work on a second face later.

To find the precise spot you want to clip, click on the playback position time (circled in red below) and enter the time you want to start (or end) a clip. Then use the frame forward/backward buttons (circled in blue) to make adjustments to the precise point where you want the mask to begin or end.



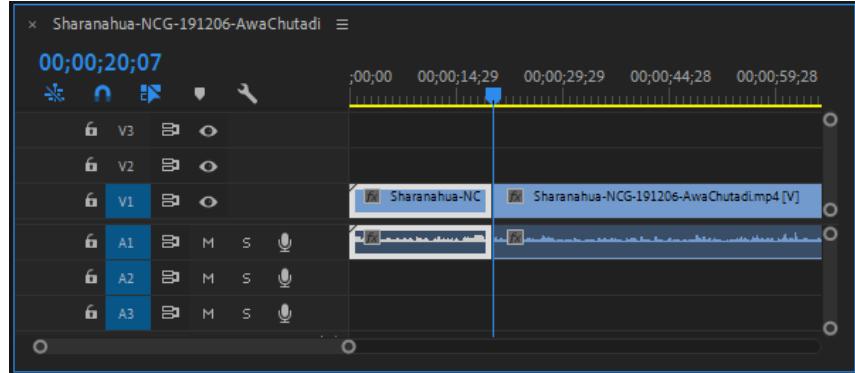
Then, to the left of the timeline box below, select the razor tool (circled in red).



The cursor will now have a razor shape when you are in the Timeline box. Position the razor over or very close to the vertical blue line indicating the playback position, and click. Even if you are not perfectly

centered on the blue line, Premiere will automatically position the start/end of the clip to that point as long as you are close.

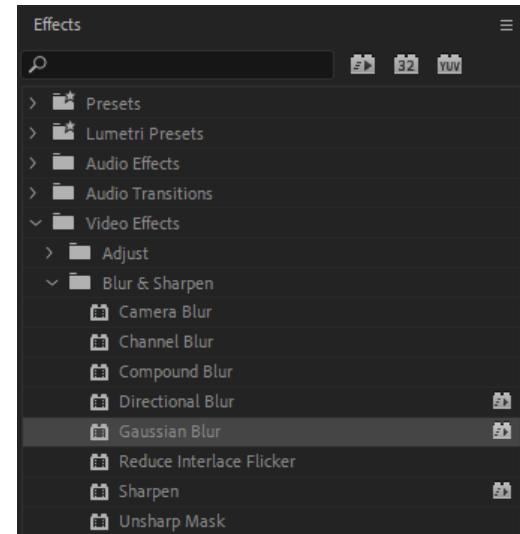
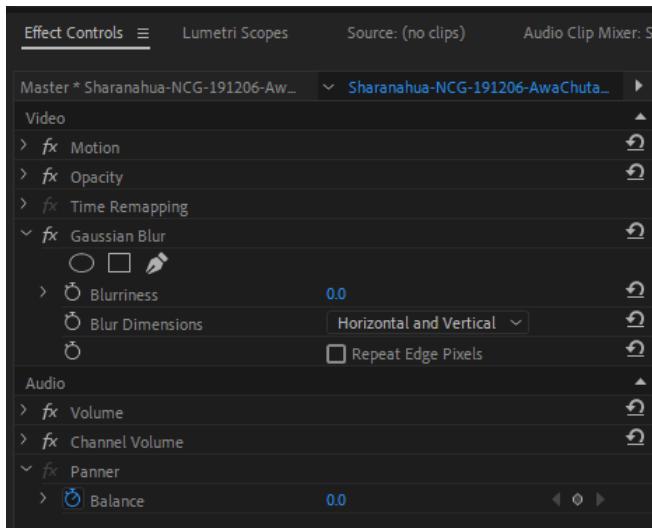
You should now have two clips:



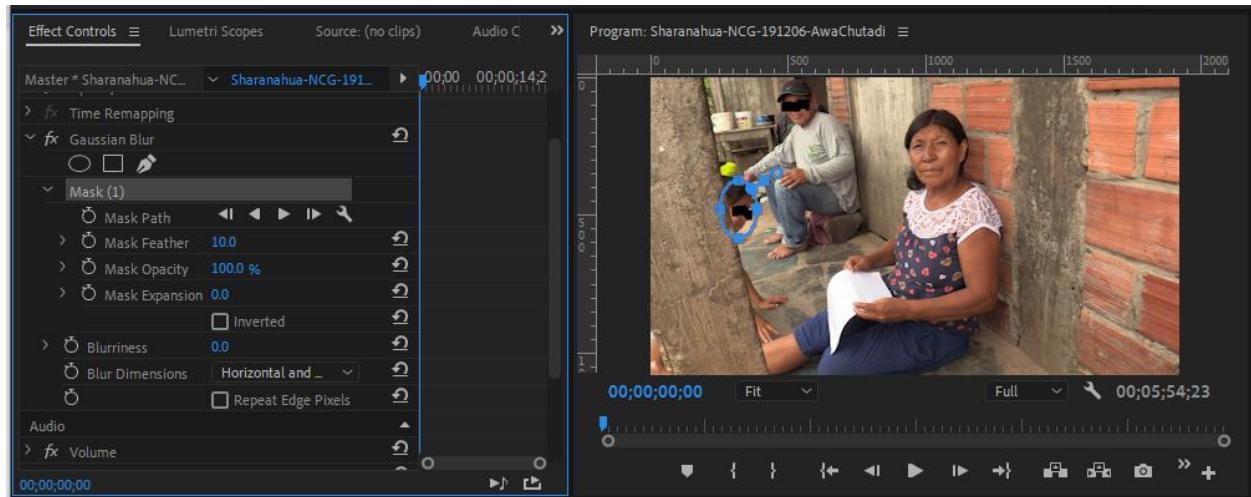
You can create and blur one clip at a time, or create all your clips and blur them all later. It doesn't really matter. Remember -- we are doing just one face at a time!

When you're done, position the playback to the beginning of the clip that you want to edit. On the far right "Effects" sidebar, click Video Effects -> Blur & Sharpen -> Gaussian Blur.

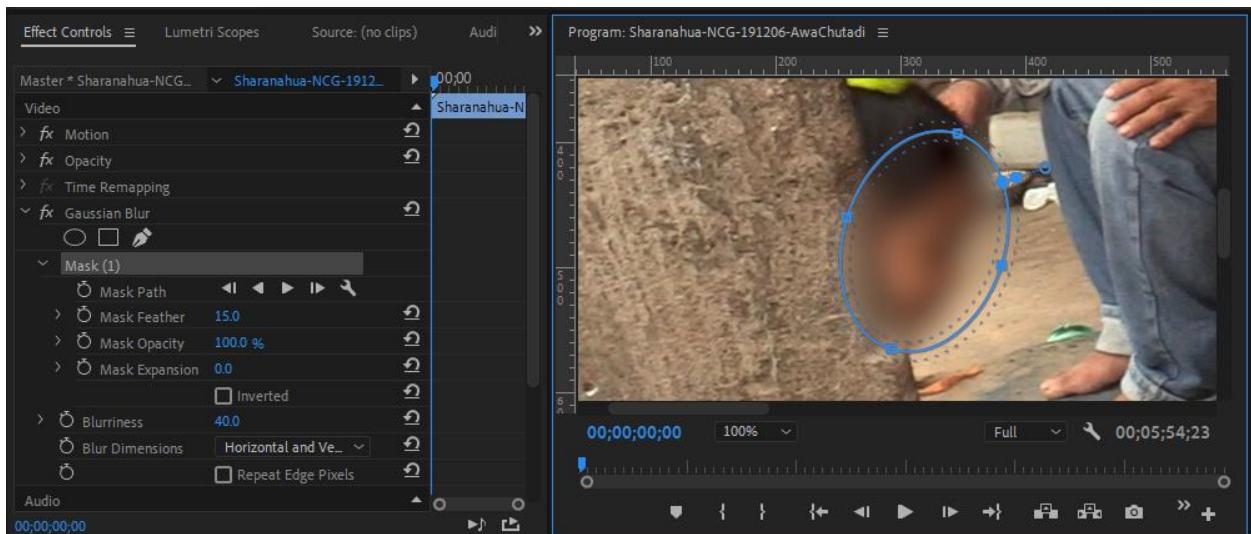
This will cause the effect to pop up on the far right pane under effect controls.



Clicking the oval icon under "Gaussian Blur" will create an oval shaped mask on the video preview. Click and drag it into position and use the anchors to adjust the size and shape.



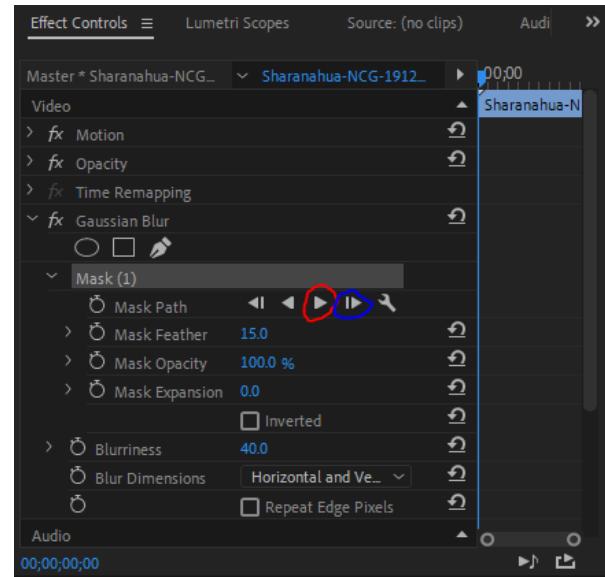
Next, adjust the properties of the mask in the Effect Controls. For faces that are well-lit and in the foreground, you will need to use a higher Blurriness. For figures in the far background, 20 is usually good, but for well-lit figures in the foreground you may need to go as high as 40 or 50. The mask feather adjusts how gradual the transition between blurred and non-blurred areas is, giving it a softer look. For figures far in the background 5 or 10 is usually good, and for figures that are very close in the foreground, 20 is the highest I go. For this part, I usually zoom the video in to full size to get a good look at the effect. Make sure there are good margins around the area you are blurring; this will help the program track the position automatically. If it is too tight, there will be a higher likelihood of it getting out of position.



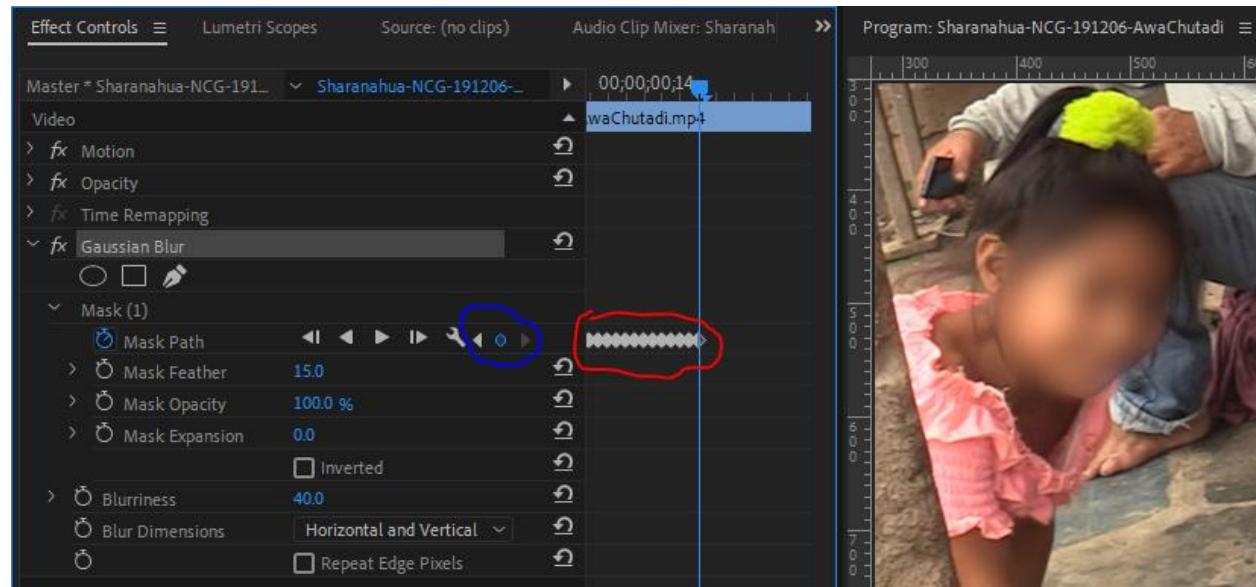
For this face, I used a relatively high blurriness of 40, and increased the mask feather to 15 so that the transition space is wider. You can click the effect name “Gaussian Blur” to make the blue mask outline disappear, then click on Mask (1) to make it reappear if you need to adjust it. You can also rename the mask by right clicking/control clicking on “Mask (1)”.

Now, because this figure is moving, we need to make the mask track across the movement. To do this, click on the forward arrow button ("Track selected mask forward"), circled in red.

You can also track forward one frame at a time, using the arrow circled in blue. There are also corresponding options to track backward. Tracking forward will automatically analyze the face and Premiere will try to change the shape, size, and position of the mask to stay on the selected area. When faces are pretty stationary and stay rotated in the same orientation, the program does well; when faces are moving and turning the program can have tracking errors. You will need to monitor the automatic tracking. If the mask gets uneven, click stop and reposition it.

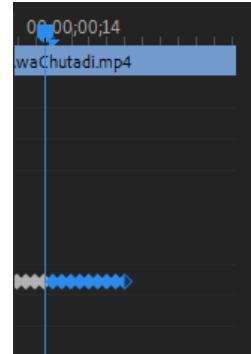


As you track forward (or backward), Premiere creates what are called "keyframes". A keyframe is basically a reference point that the program can compare to when it attempts to reposition the mask. Premiere shows you each place a keyframe is created in the box to the right of the effect controls. They are circled in red. To delete an individual keyframe, you can click the blue circle (circled in blue), which will turn white. The arrows on either side let you go forward and backward through the keyframes you have created for a given clip



If you have a situation where you were not closely monitoring the mask tracking and you need to delete a group of keyframes, you can do this by clicking and dragging a box around them which will select them.

You can use move the blue playback position line to the last place where the mask was correct, then select all the keyframes after that. Press the delete key to delete them. You can then adjust the mask and continue tracking forward.



As you work, don't forget to periodically save your project! The auto-save does not always save your work.

Some factors that can cause the mask to get out of position:

- Rotation or tilting of the face
- Motion of the face behind other objects or motion of other objects (e.g., gestures in the foreground) in front of the face
- Rapid motion
- Mask margins that are too tight

Portions of video that have these complications do better if you make the mask around them larger. But if the figure is in the foreground, it might be better to just fine-tune the tricky parts by stepping through frame by frame. Going frame by frame is also a good option when you want to preserve a gesture that crosses over a face in the background.

When the face is very stationary or is on the margins of the frame (and therefore less noticeable so you can do a wide mask), you can also set a keyframe for the mask position, then skip the playback position forward to any motion. The mask will stay in the same position from the last keyframe assigned to the next keyframe -- even if these are minutes apart. Doing this cheat instead of actually tracking forward is way faster because the program doesn't have to step through each frame.

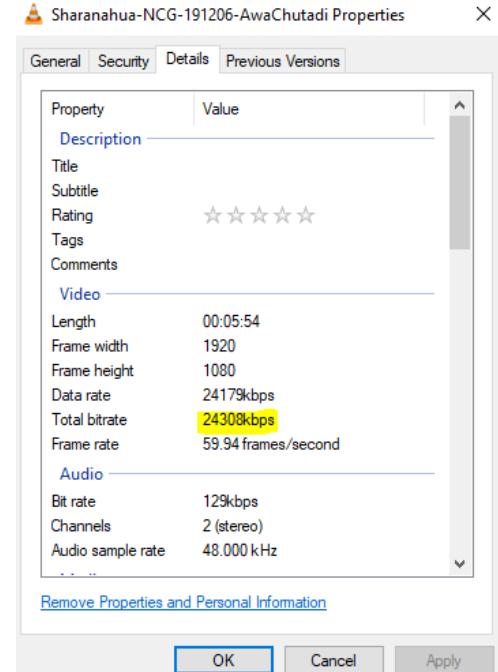
If you have parts of the video where there is more than one face that needs to be blurred at a time, start with only blurring the face that is visible for the longest duration. If the durations are equal, or just pick one. When you are done blurring it, then you can move onto the next face.

If the durations are the same, you can just click the oval icon under the effect to create a new mask. I find it useful to minimize the first mask under Effect Controls to avoid accidentally changing its values. Adjust the blurriness and feather for the second mask, then proceed as described above. If the durations are different, you will want to get the razor tool out again and make more clips in the Timeline at the bottom. Once you have already defined the mask and keyframes for it, clipping within a span of the mask will not mess it up. (This is why you want to start with the longest duration!)

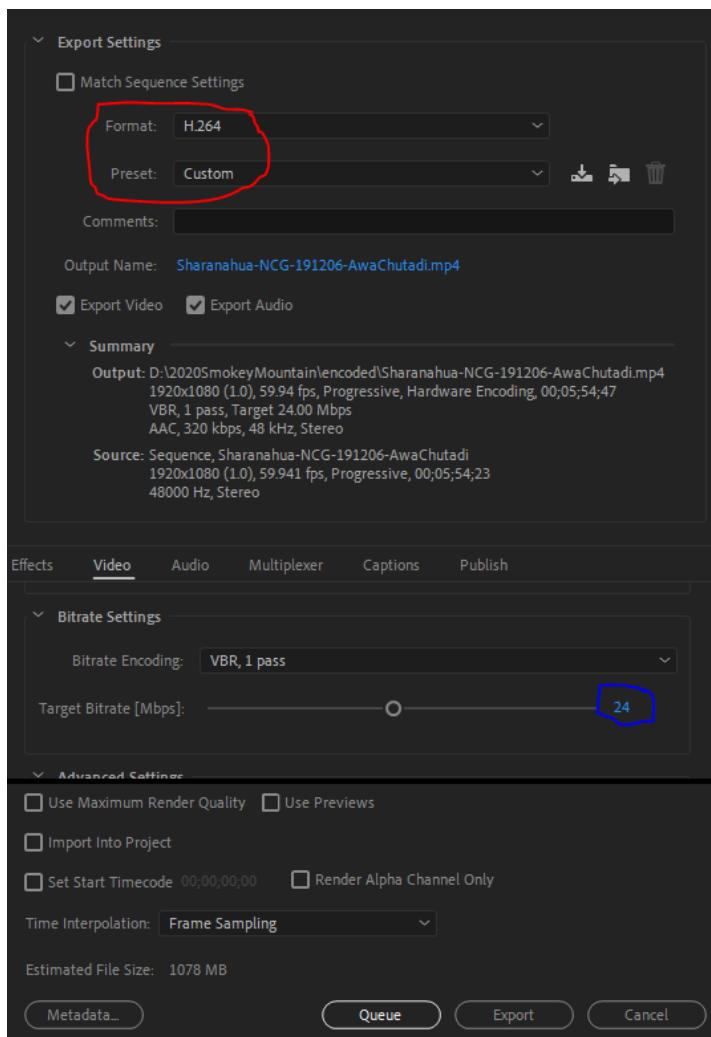
Right-clicking/Control clicking on the mask name under Effect Controls will give you the option to copy the mask. This is useful when a span of video that needs a single mask has been split by clips.

When you have finished applying all masks, you can play the video in the window to see how it turned out. When the edits are done, click into the Timeline area to activate it (a blue outline will appear), then click File -> Export -> Media.

Check the properties of your original for the bitrate information:



Use the settings as displayed below and circled in red. Enter the video bit rate in the area circled in blue. The bit rate does not have to be perfectly exact, I just go with the nearest Mb.



Check the audio tab (seen in the image to the left, next to "Video" above "Bitrate Settings" and adjust bitrate settings of the audio to match your original.

Click Export. It will take a few minutes, but soon your video will be ready! Remember to test it by playing it with audio to make sure there weren't any problems!