

Using Layers and Masks in Affinity Photo

By Andy Mills LRPS

Introduction

Post-processing digital images on a computer is an essential part of modern photography. Post-processing is used to help the photographer 'realise' the original intent.

It is of course acknowledged that many digital cameras do allow in-camera manipulation of the JPG image but that is somewhat limited in scope. Some cameras even allow optional applications to be purchased and installed to do further manipulation such as double exposure, time-lapse, selective colour control, star-trails, bracketing, etc. the most expansive type of post-processing is done on a computer outside of the camera.

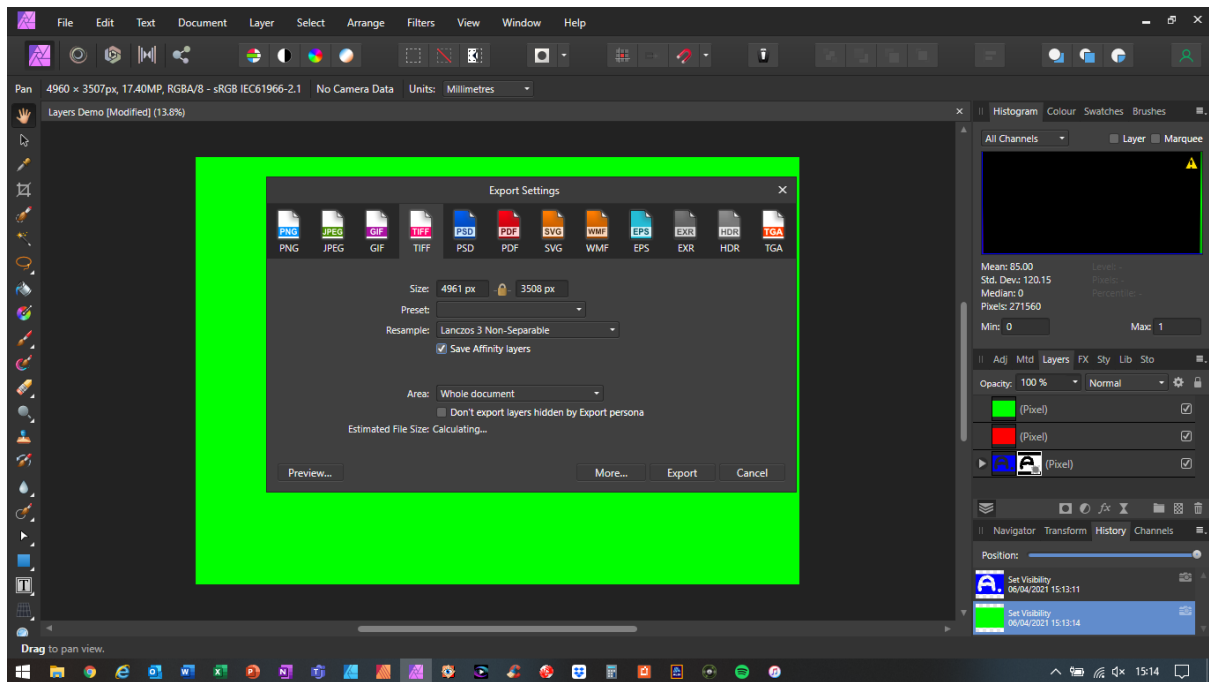
Affinity Photo is one of the most powerful and advanced digital photo post-processing tools available. Don't think its low cost means it has low functionality – the opposite is true! Affinity Photo is a British product trying to compete and win a market share in a market dominated by American company Adobe. Serif, the British company who created the Affinity suite of software, has thrown down the gauntlet to the competition in true style.

Simple editing of images can be done in a destructive manner. 'Destructive' simply means the change can't be un-done. However, Affinity Photo has a number of features built-in which help prevent manipulations being irreversible. Most of the editing actions create a new layer. The user can choose to 'merge' the change into the image which consolidates the layers. Another feature which then still allows editing actions to be un-done is the 'Undo' option. Affinity Photo maintains a history of editing actions which can be un-done in reverse order. That means you can click Undo many times to get back to an earlier status. It is widely regarded as good practice to retain the layers used to edit an image, particularly if you want to go back and refine any adjustments later on.

Affinity Photo also has the ability to store a 'snapshot' of the image at any time. That helps the user to save a snapshot at various stages of editing to create recovery points that are easy to return to if necessary. This is another way of being able to return to an earlier step in post-processing. Note that the snapshots are lost when the file is closed.

The image and all its layers can be saved to a file using the .afphoto format so that the computer can be turned off and the editing project continued at a time in the future by opening the .afphoto file.

Affinity Photo does allow the layers to be saved in a TIFF file if you tick the box for "Save Affinity Layers":



Note: If an image is exported to a JPG, GIFF, PNG etc format the image is flattened and the layers are lost.

In addition to creating a layer for each editing action, the user can combine multiple images to create artistic effects or create composite images such as sky replacement or high-dynamic range, as done by the photographers in the Victorian era in the mid-1800s by using multiple negatives.

Therefore, managing layers is at the heart of post-processing digital images. Affinity Photo is as powerful at handling layers as any other competitive product.

Computer Requirements

Post processing software uses the graphics processor in a computer. This means that a computer with a powerful graphics card will give faster and better results than an old or low-cost computer. Similarly, the computer screen is important because that is how the user visualises the image. Therefore, the best computer for digital image post-processing is likely to be a desktop with a powerful graphics card and separate high-resolution screen and adequate RAM. A good laptop with external screen is a fair compromise. The larger the number of layers the larger the computer RAM requirement. Other than RAM size there is no limit to the number of layers that can be created and adjusted, then saved in the .afphoto file.

The Affinity Photo user should consider how the image's attributes are controlled all the way from the screen, through the post-processing software (Affinity Photo in this case) and out to the printer or output file. Adjusting and calibrating the screen colours and luminance allows the user to see an accurate representation of colours and luminance. This allows the Affinity Photo software to be used to adjust the image to the correct (desired) effect. Control of the printer to ensure the final printed image is a faithful representation of what is desired and seen on the computer screen is dependent on the type of paper used and the translation of the colours seen on the screen to the printer inks. This is managed by something called ICC Profiles which configure the printer. ICC Profiles are available for various printers and types of photo paper. Adjustments made to the digital image via layers will be futile if the computer screen and printer outputs are not controlled too. If the computer output is a digital image which is intended to be projected, then the projector would need to be calibrated and controlled too.

Affinity Photo has been developed to run on modern personal computers, either Apple Mac or Apple iPad or Microsoft Windows. The latest Affinity Photo version and the computer requirement to run it is available on the Affinity website: <https://affinity.serif.com/en-gb/photo/full-feature-list/>

Understanding Layers

There are several different types of layer available in Affinity Photo as follows:-

- **Pixel layer** = A layer for raster images which are comprised of pixels that you can edit
- **Image layer** = An un-editable image that has been imported or placed as a layer
- **Vector layer** = A layer for lines and other vector content
- **Fill layer** – a layer for filling with a solid or gradated colour
- **Live filter layer** = A filter that affects the layer below it
- **Mask layer** – A mask that allows the user to allow or block the layers below the mask from visibly affecting the layers above it
- **Adjustment layer** = A layer that makes adjustments to layers beneath it

When opening the Affinity Photo application directly from the Affinity Photo icon, it starts with no image loaded. A new blank image can be created by using the File - > New menu control.

If you configure your computer to use Affinity Photo as the preferred image editor you can click on an image file and it will open Affinity Photo then load the image as a sequential set of actions. Affinity Photo can open and edit various file types, such as JPG, TIFF, GIFF, PNG etc including various RAW image formats used by the popular camera manufacturers. It will even open Adobe DNG, PDF and PSD files.

If opening a JPG, GIFF, TIFF, PNG etc image files Affinity Photo starts in the **Photo Persona**.

If opening a RAW image file Affinity Photo starts in the **Develop Persona**.

An Affinity Photo Persona is the 'personality' of the user interface - the user interface layout giving a relevant 'look and feel'. This is simply a screen layout and relevant set of adjustment controls for manipulating that type of file.

The types of Persona are:-

- **Photo Persona** for editing and adjusting images
- **Liquify Persona** is used for pushing pixels around
- **Develop Persona** for converting RAW files into editable digital images
- **Tone Mapping Persona** is used for adjusting colours, contrast etc to create a range of high dynamic range images to choose from, including an HDR black and white option
- **Export Persona** is used for outputting files and parts of files in various formats

Affinity Photo allows the user to easily switch between Persona to make available the required tools. This includes converting a JPG image back into a format suitable for adjustment in the Develop Persona.

Most image adjustment are done in the Photo Persona. An image layer would need to be converted to a pixel layer before it can be edited in the Photo Persona. This conversion is done by using the Rasterise control. If the user tries to adjust an image layer the Affinity Photo application will automatically convert it into a pixel layer and issue a warning to advise the user.

The Layer drop-down menu is the main way to create a new layer. A new layer is generated when the adjustments in the Adjustments Panel on the right are used. The layers are listed in the Layers Panel on the right of the screen. The layer at the top of the list is the dominant layer and covers the layers below. However, the layers below can be made visible via masks or opacity adjustments as described below.

How Affinity Photo handles layers is probably best learnt by a new user doing some experimentation and practice as follows. Rather than beginning with using layers of real digital images, it is suggested a few plain coloured layers, such as the Red, Green and Blue primary colours, are created and used for experimentation.

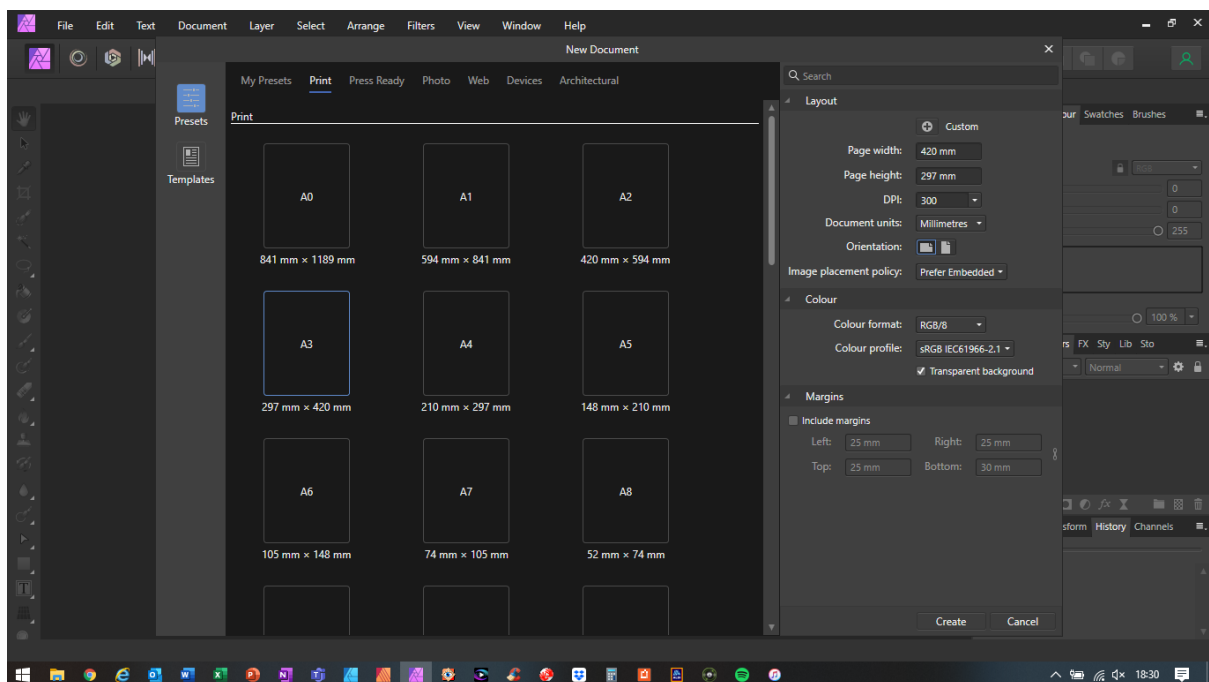
Open the Affinity Photo application without an image being loaded.

Click on File -> New...


In the window that opens select Print then select A3

Then Select Landscape format and tick the box that says transparent background.

You should then have a screen that looks like this:



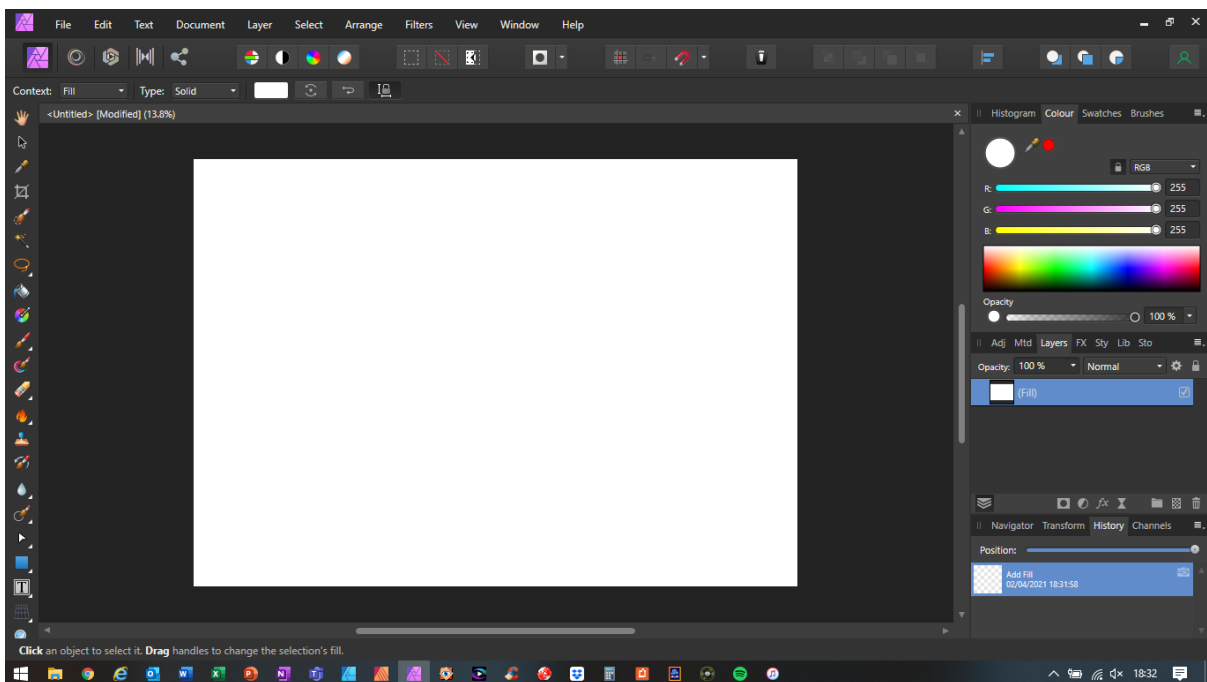
Click Create:



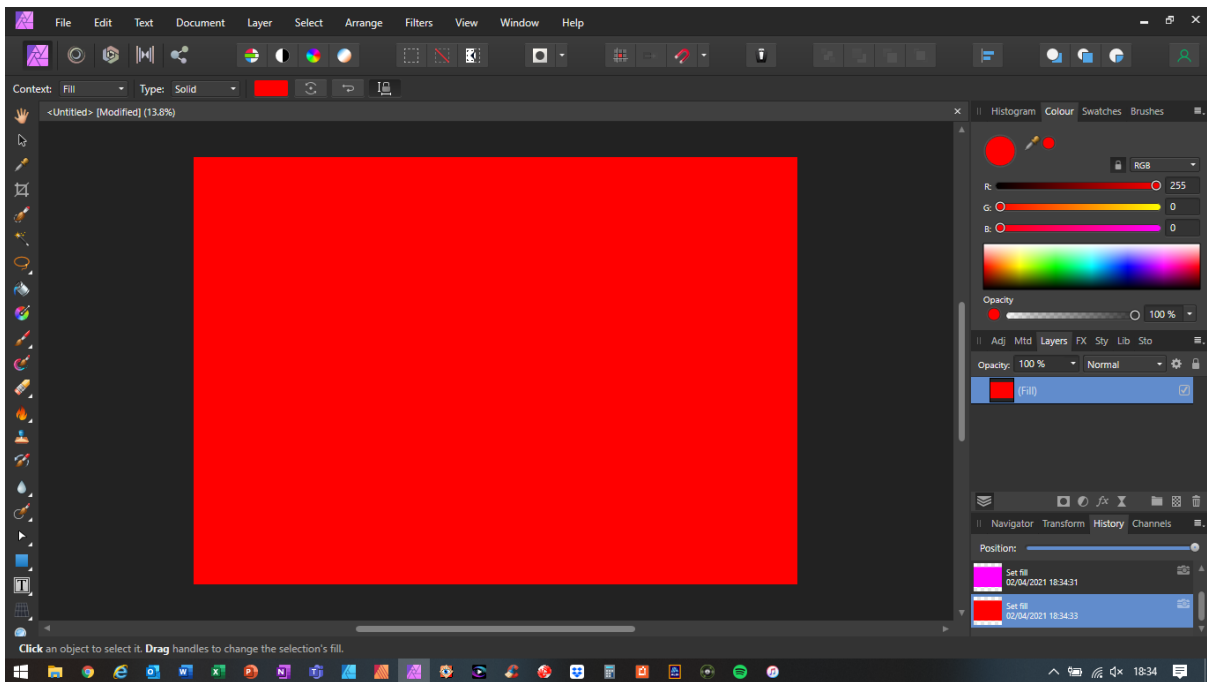
You will see that the Layers Panel on the right is empty.

Go to the Layer drop-down menu then click on New Fill Layer

If the colour sliders on the right are set to white you will see a screen like this:




You can then adjust the G (green) and B (blue) sliders down to zero (0) just leaving the R (red) slider at 255. This will make the fill layer pure red as follows:



Create another Fill Layer and fill it with Green.
Create another Fill Layer and fill it with Blue.

Your screen should then look like this, with three layers showing in the Layers Panel on the right:




Because the blue layer is at the top of the list in the Layers Panel, that is visible.

Note that all layers have taken on the A3 dimensions used to create the transparent background.
Also note that the background is not a layer, it is just the canvas upon which you are working.

You will see that all three layers have a tick in the box in the Layers Panel. If you un-tick all three layers you will see the transparent background which is the canvas and shown as a checkerboard pattern, as seen above before you created any Fill Layers.

If you click on the name area of each layer, where it says (Fill), you can type in a name for the layer. You will see in the following image that the layers have each been given a name based on its colour:




If you tick the box for the red layer the image will be seen as red, instead of the checkerboard background.


If you then tick the green layer it will hide the red and just show green.

If you then tick the blue layer it will appear blue because this is the dominant top layer.

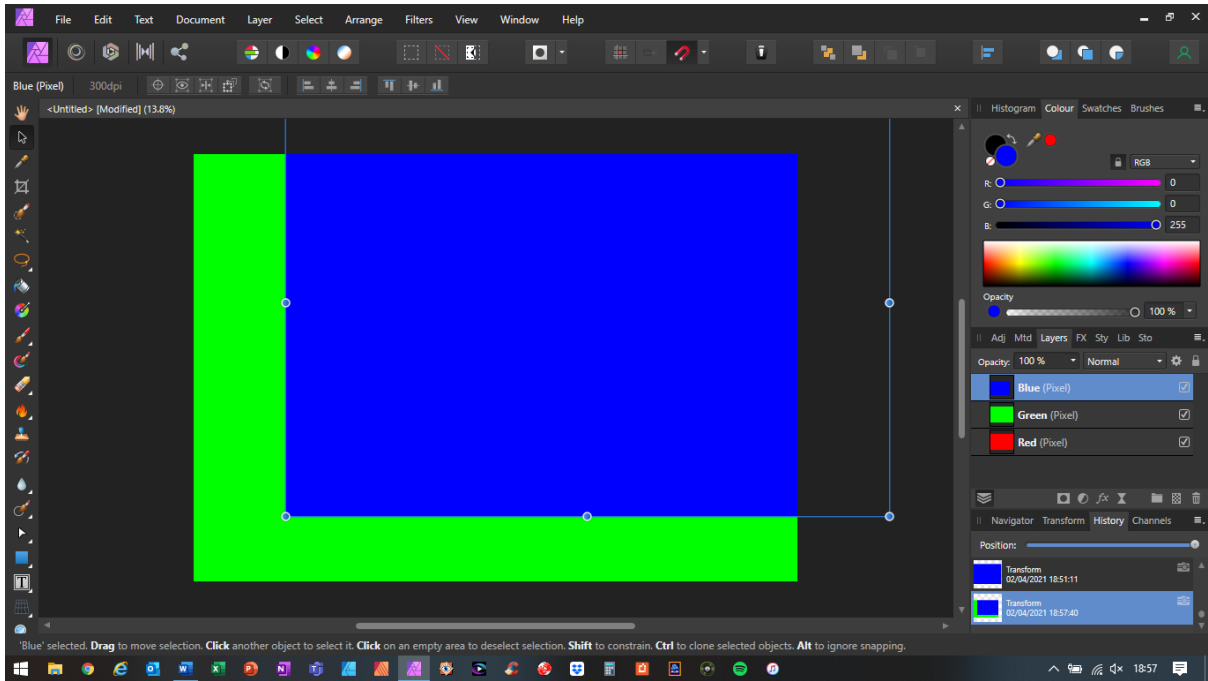
The Fill Layers are not fully editable, you can only change the colour content, so to give full control for editing and adjusting the three layers we created, you will need to convert them to pixel layers. This is easily done by clicking on a layer in the Layers Panel then going to the Layer drop-down menu at the top of the screen and clicking on Rasterise. The names area of each layer in the Layers Panel then says (Pixel), as seen in the following screenshot:



If one layer is selected in the Layers Panel on the right (e.g. blue as seen above) and the Move Tool is selected in the Tools Panel on the left you can move that layer, as shown in the screenshot below:



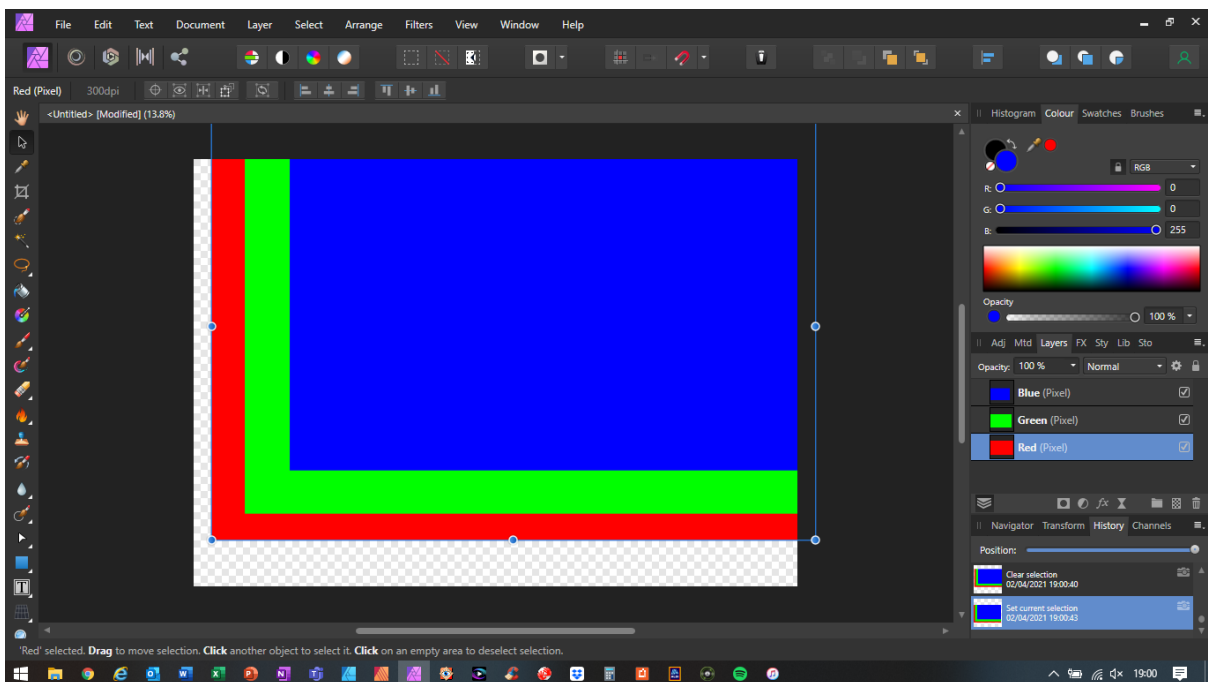
Eight little circles denote that the layer is not locked in position and can be moved. Try moving the blue layer slightly so it appears as follows:



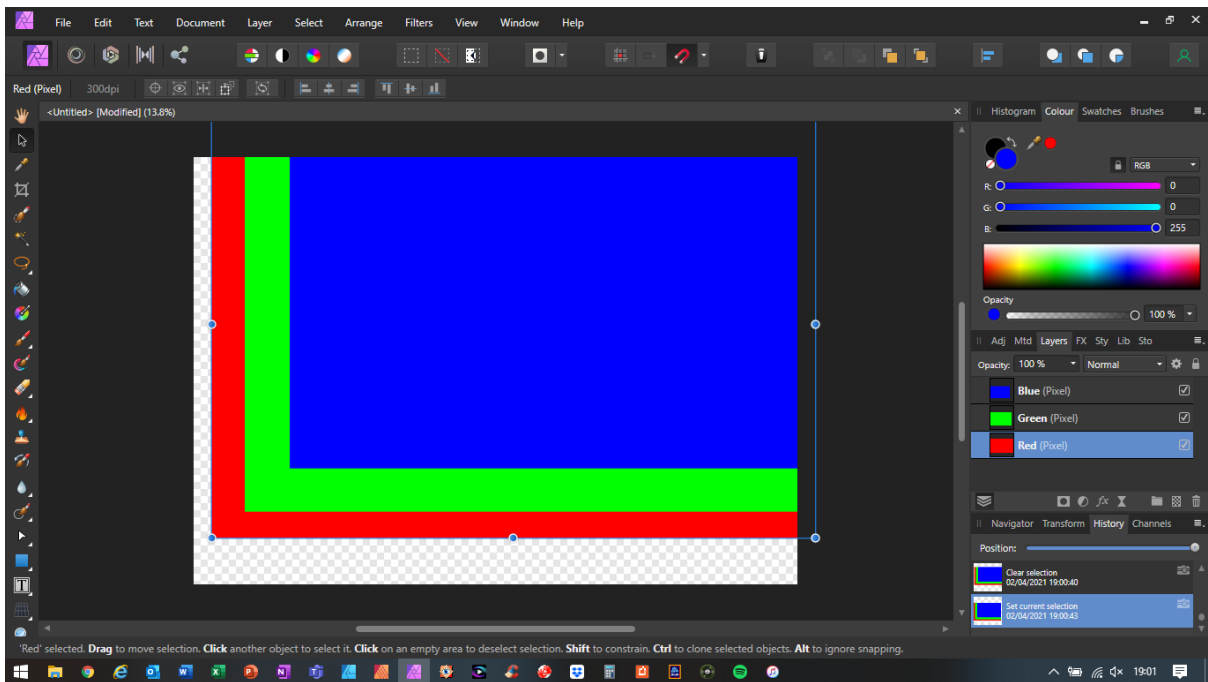
Moving the blue layer then reveals the green layer below.

Try moving the green layer next. This reveals the red layer at the bottom.

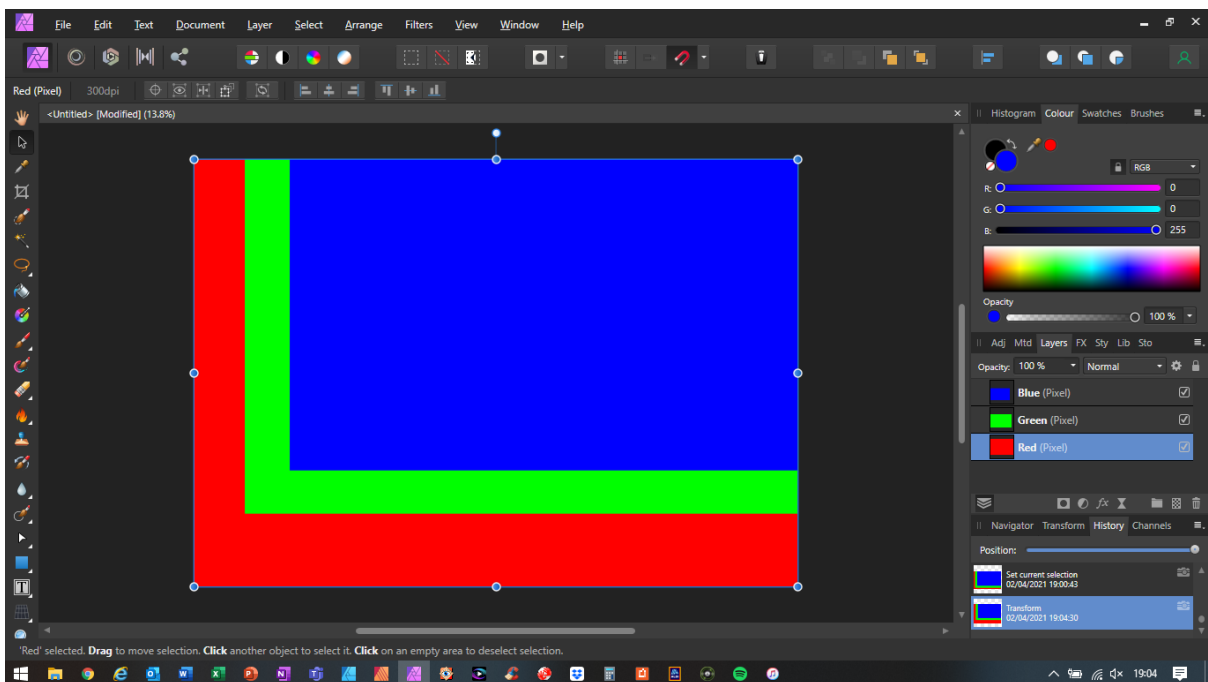
If you move the red layer too, that reveals the checkerboard background again:



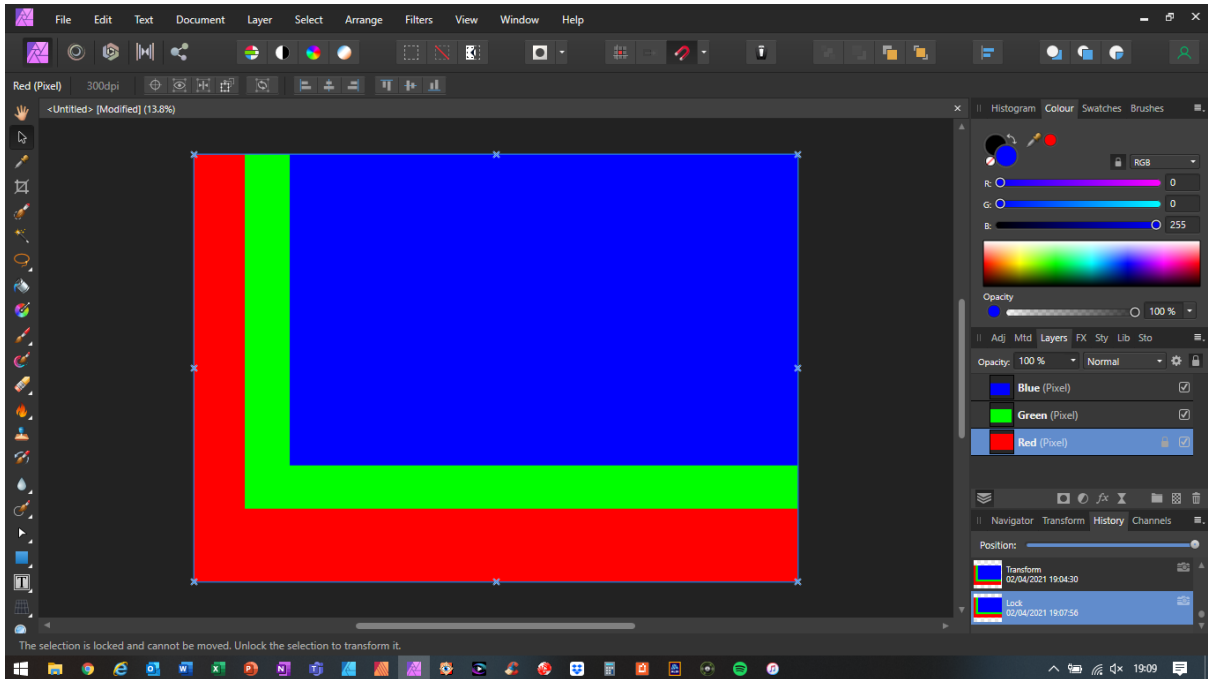
The layers can be easily re-aligned to the background by turning on the Snapping Tool (looks like a red horseshoe magnet) in the top Control Panel:



Move the red layer back down to cover the background.
Notice the red and green alignment lines that briefly appear to show alignment:

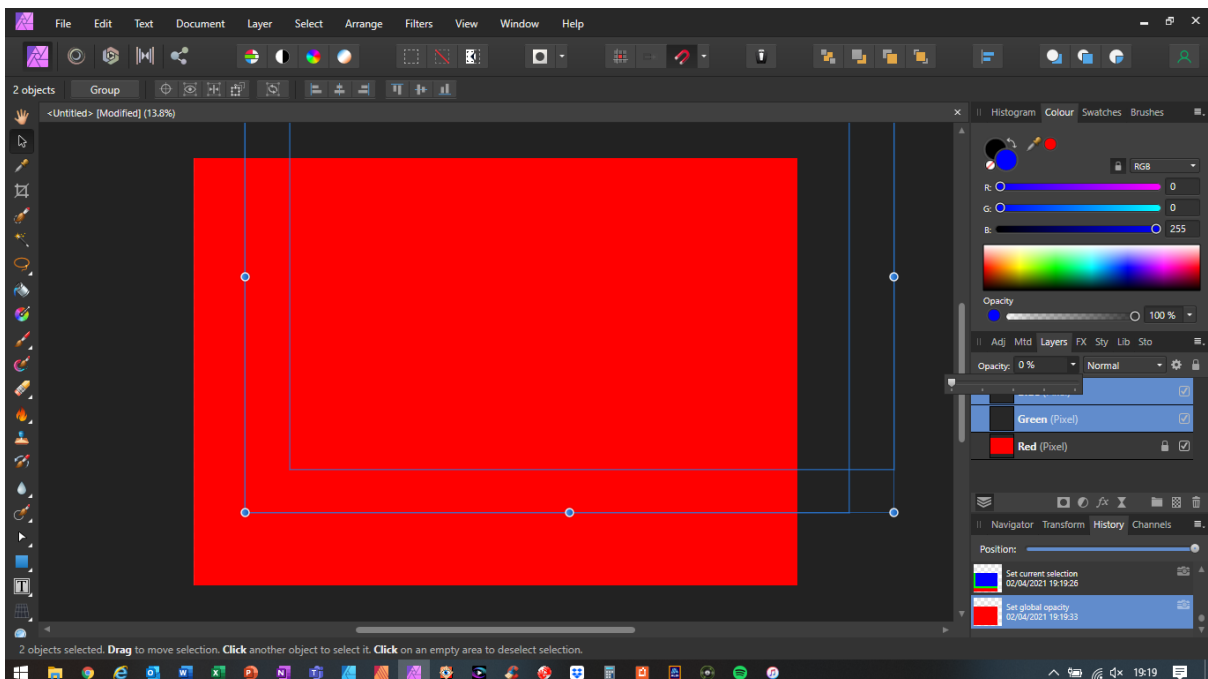


The position of a layer can be locked to prevent movement by selecting the layer in the Layers Panel, on the right of the screen, then clicking on the little padlock symbol at the top of the Layers Panel. A padlock symbol then appears on the locked layer, next to the tick, and the little circles around the red image change to little crosses, as shown in the following screenshot:

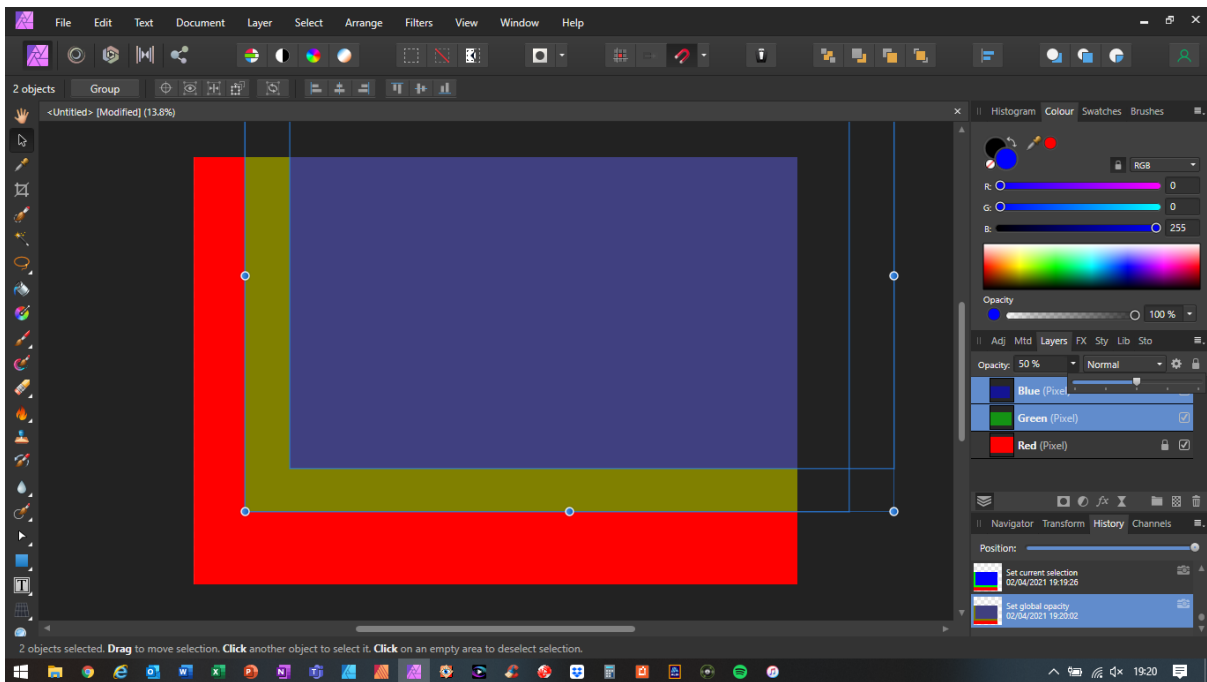


With the blue and green layers off-set and the red layer locked, as shown in the above screenshot, the blending of the layers can be experimented with as follows.

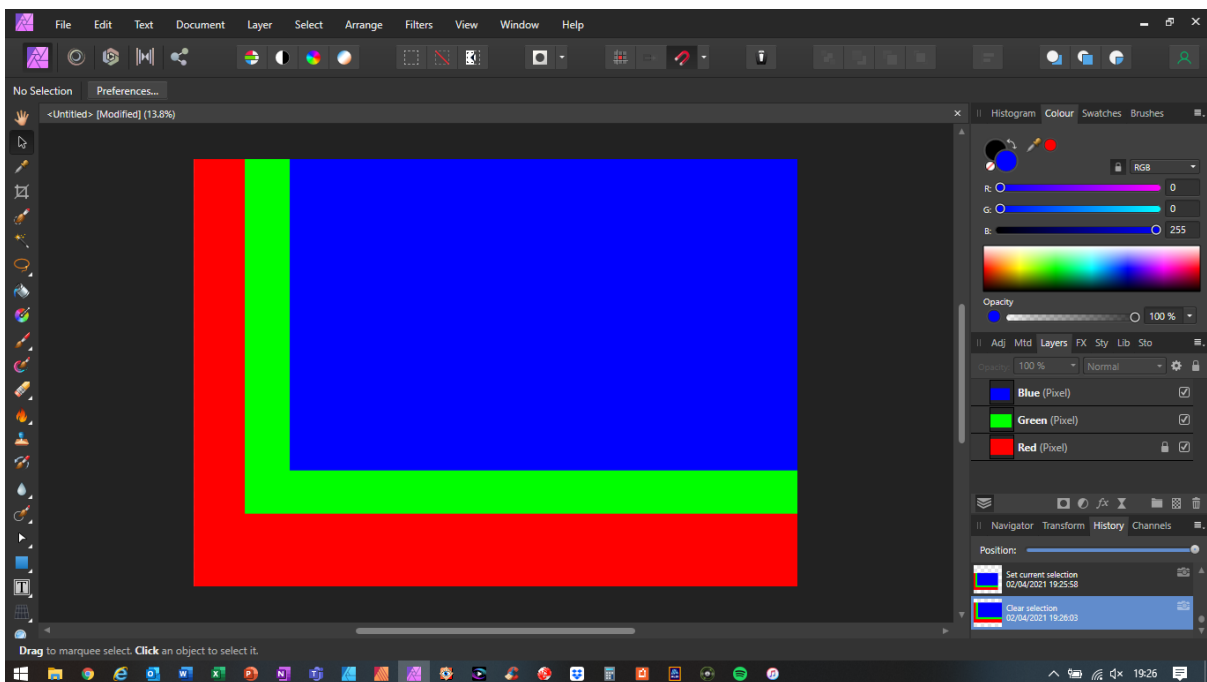
Select just the blue and green layers in the Layers Panel on the right, then move the Opacity slider (at the top of the Layers Panel) all the way down to 0%. This makes the blue and green layers transparent thus showing just the red layer:



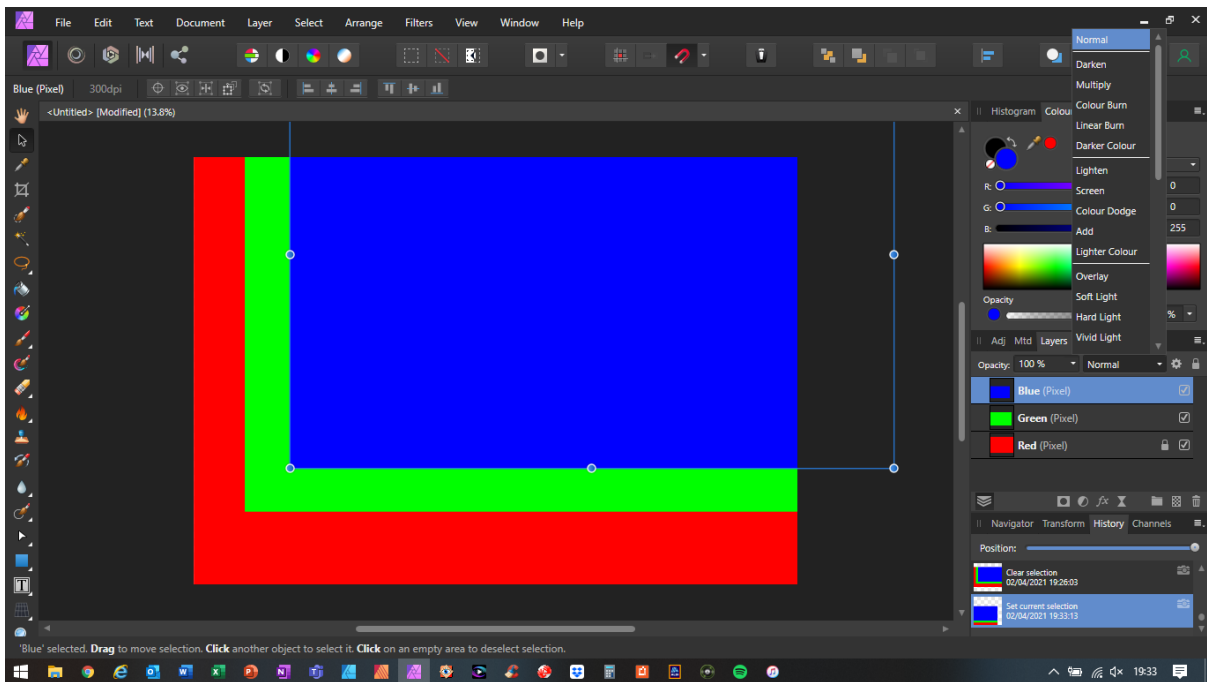
Move the Opacity slider back to 50% and notice how the colours from all three layers blend to form different colours where they overlap:



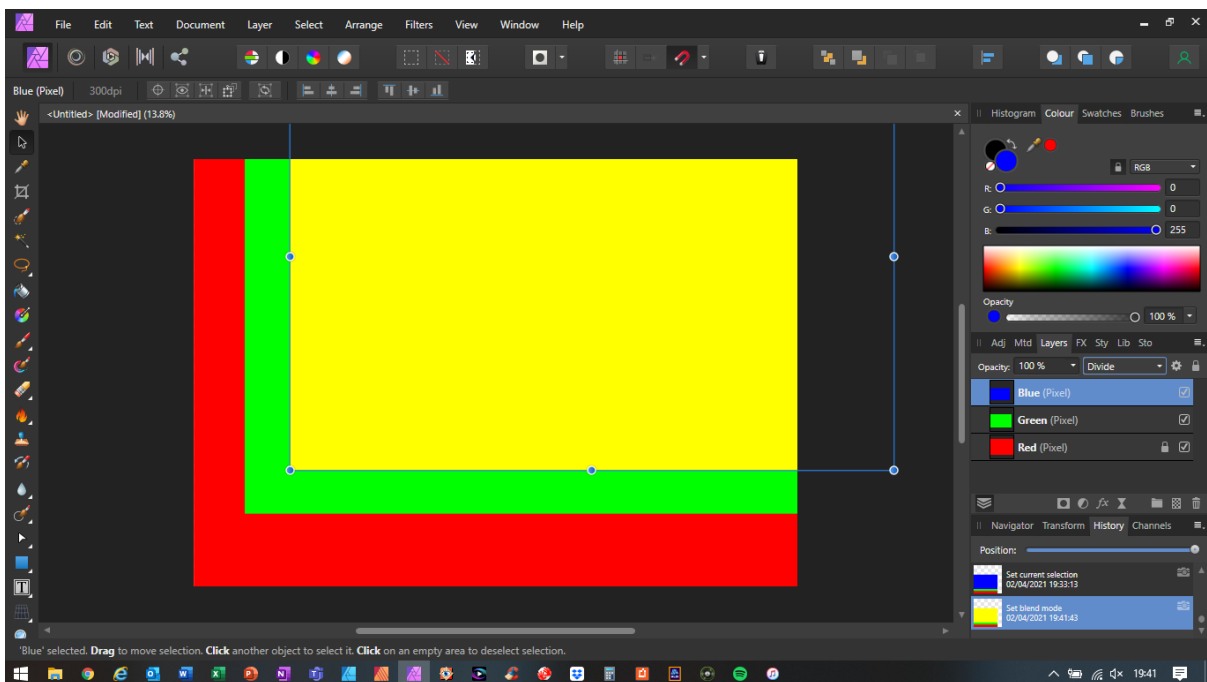
Returning the Opacity slider back to 100% restores the R, G and B colours to the layers respectively.



Affinity Photo has a number of pre-defined blend modes to combine layers. Select one layer in the Layers Panel then use the dropdown menu at the top of the Layers Panel (it is set to 'Normal' as the default). Clicking on that menu reveals the list of pre-defined blend modes available. Moving the mouse pointer over each option gives a preview of its effect.

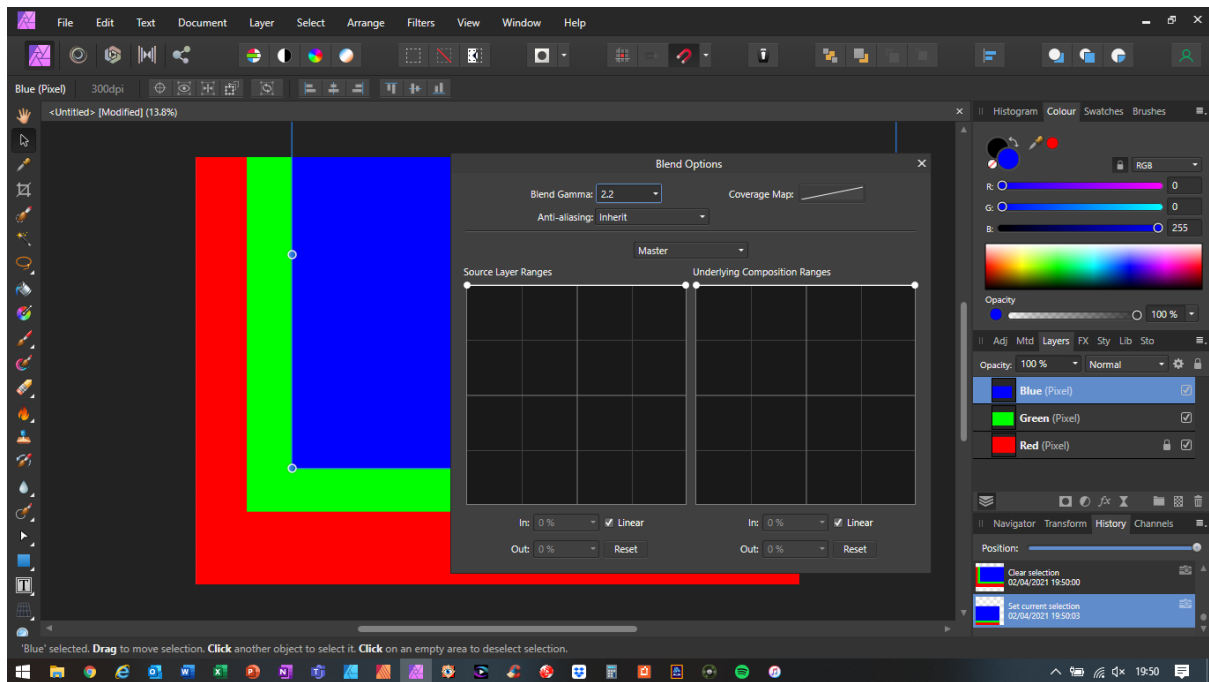


If you choose the blend mode called 'Divide', the area of the blue layer that overlaps with the green layer gives you yellow (which is correct, when you blend green and blue light you get yellow):



Note: It is suggested you return the blend mode menu to 'Normal' before proceeding to the next step.

In addition to those pre-defined blend modes, Affinity Photo has another very powerful way to blend layers which is not available in other popular software. This is accessed via the little cog symbol next to the blend mode menu at the top of the Layers Panel, after a layer is selected. It reveals two graphs, as follows:



The graphs are like Histograms, with black on the left and white on the right. The colour channels can be selected individually or all of them are acted upon when 'Master' is selected.

The graphs have nodes which are little white circles, if not selected, or a black circle when selected. If the 'Linear' option is enabled, at the bottom of the graphs, the nodes can be moved to form straight lines between nodes, and additional nodes can be created by clicking on the line. This allows a series of vertical, horizontal and diagonal lines to be drawn on the graph to create various blend effects.

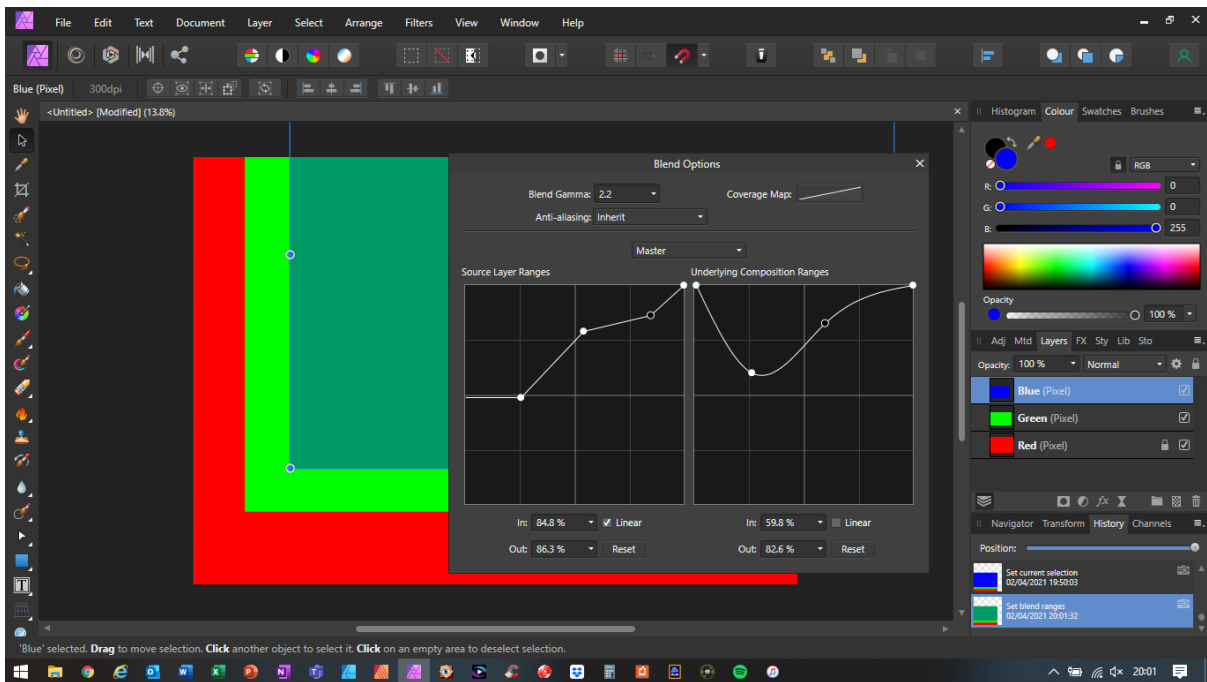
By un-ticking the 'Linear' option the straight lines become curves. This gives a more graded variance across the luminosity range.

The left graph acts on the selected layer to change its luminosity, or individual colours if one of the colour channels are selected.

The right graph acts in a similar way but on the layers below the selected layer.

Any combination of blends can be created by forming linear or curved graphs on the current layer and lower layers.

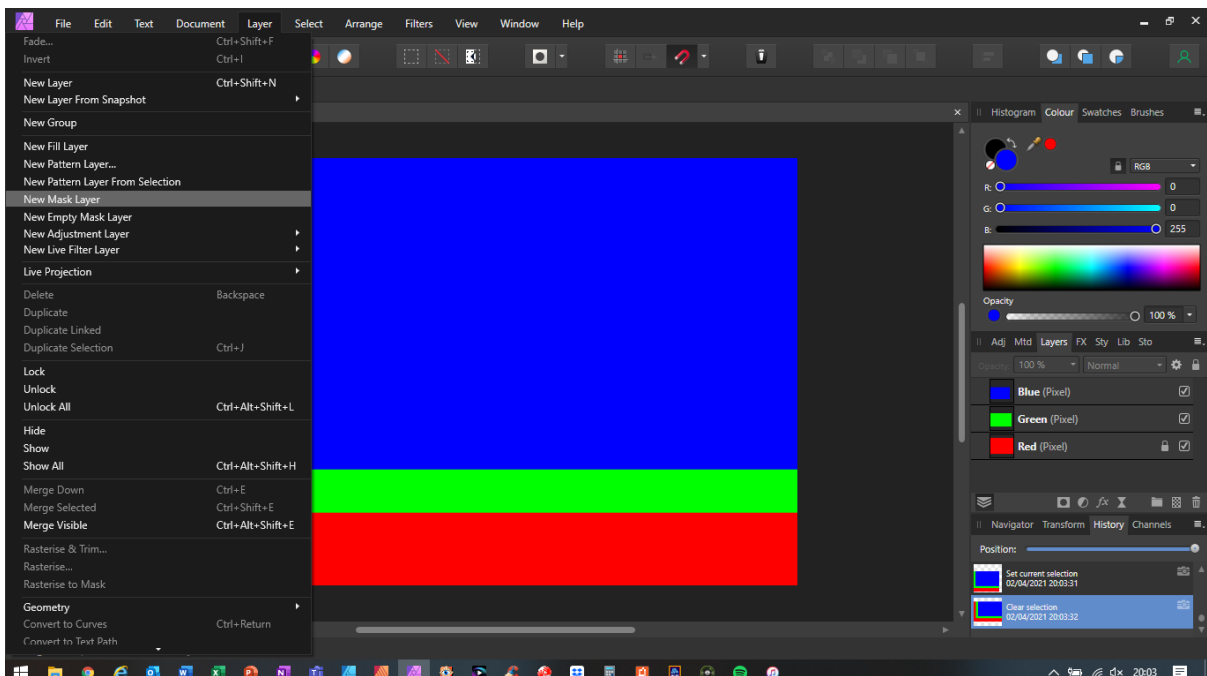
Experiment with the graphs to become familiar with their effects and their great potential to control how layers are blended:



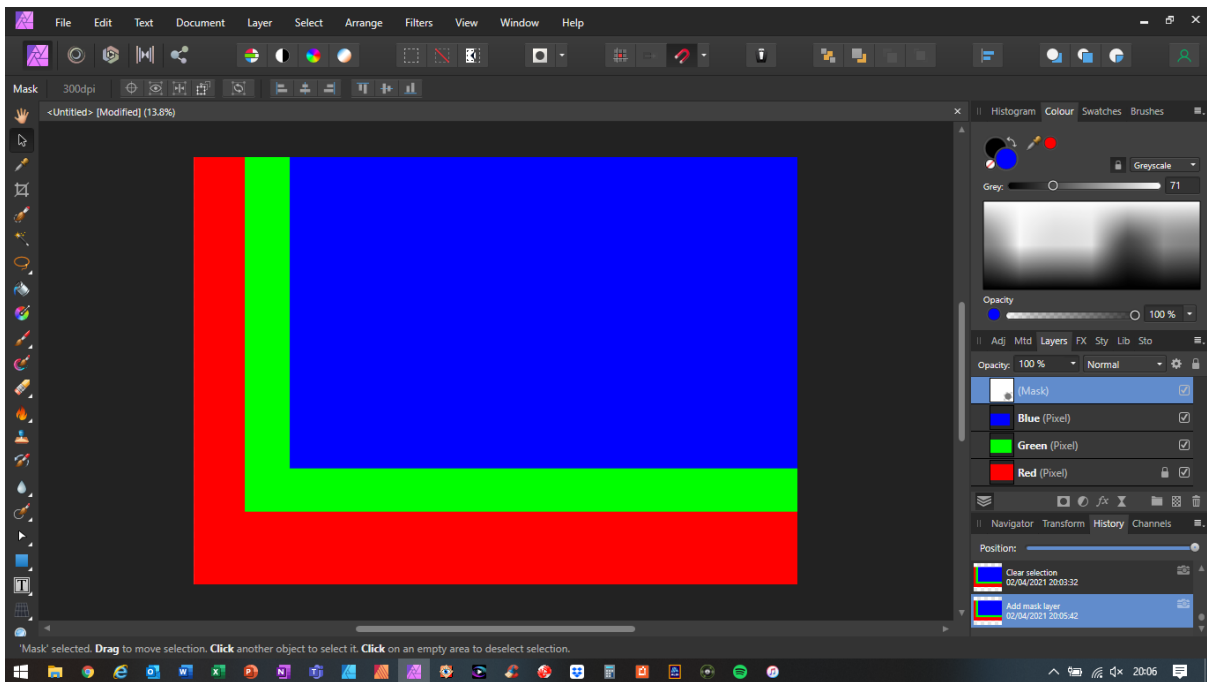
Clicking on the reset button restores the graph to normal (flat at the top hence not affecting the layer)

Understanding Masks

With the three layers visible, go to the Layer drop-down menu at the top of the screen and select New Mask Layer:



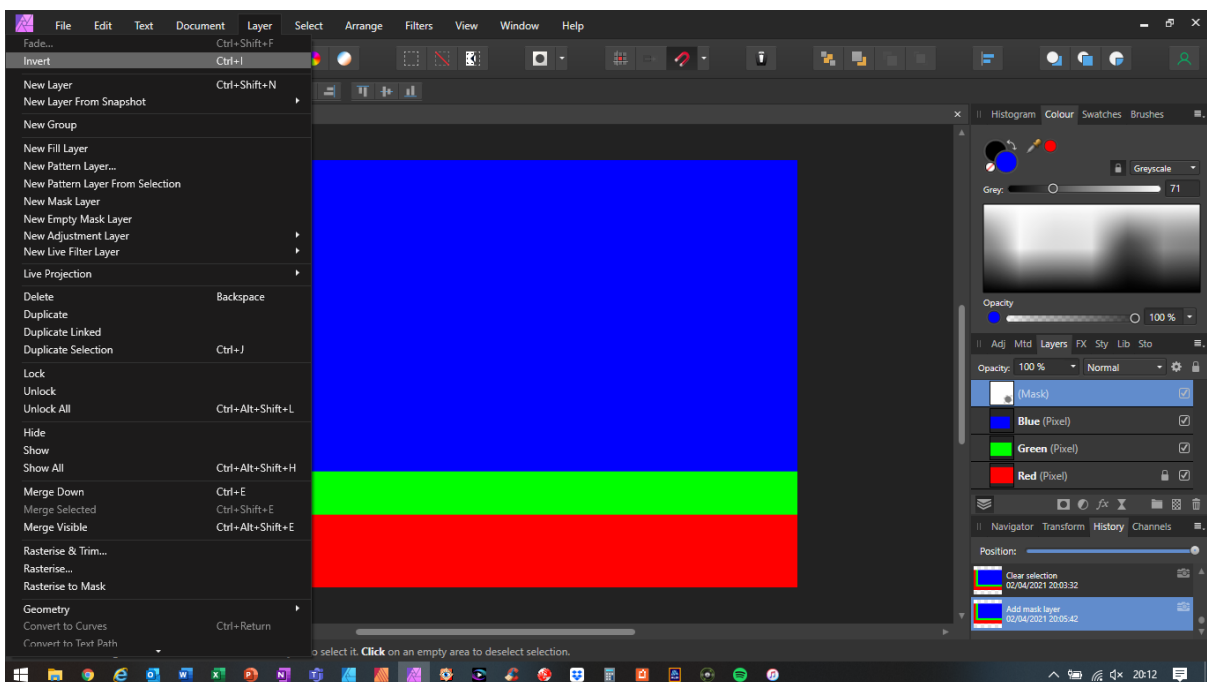
The new layer will appear in the Layers Panel on the right. If no layers are selected the new Mask layer will appear at the top:




A mask acts to control the luminosity from the layers below. If you make the mask black it will not allow luminosity (light) to be visible from the layers below. If the mask is white, it will allow all colours of light from the layers below to be visible. You can adjust the mask to be any shade of grey between black and white. Using grey rather than pure black or white has the effect of controlling the amount of luminosity passing through the mask.

Note: black is no light and white contains all colours of light.

A mask can be easily inverted from white to black, then from black back to white again. When a layer is selected in the Layers Panel the Invert control is available in the Layer drop-down menu at the top of the screen:



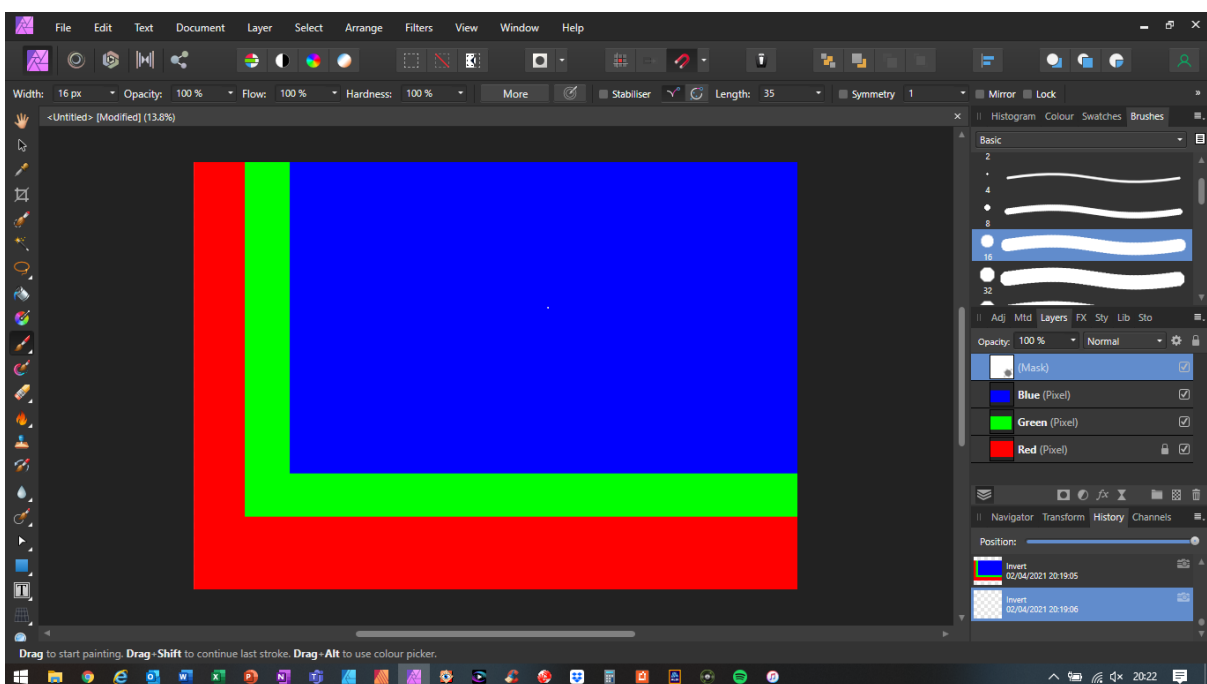
If you do invert the white mask black it will make all layers below it invisible, thus leaving just the checkerboard background visible:




Rather than inverting the whole mask to block all layers below it from being visible, you can paint an area of the mask to black-out just a part of the mask. To try this out, invert the mask again so it becomes white, hence all layers are visible through it. You can do this by going to the Layer drop-down menu and clicking in the Invert control, or you can use the keyboard shortcut of Ctrl-I (on a Windows computer) or Cmd-I (on an Apple computer).

Select the Paint Brush Tool from the Tools menu on the left of the screen.

Go to the Brushes Panel in the top-right of the screen and select the 'Basic' set of brushes and scroll down to find a brush style to suit your choice. In this example the Basic 16 brush has been selected:

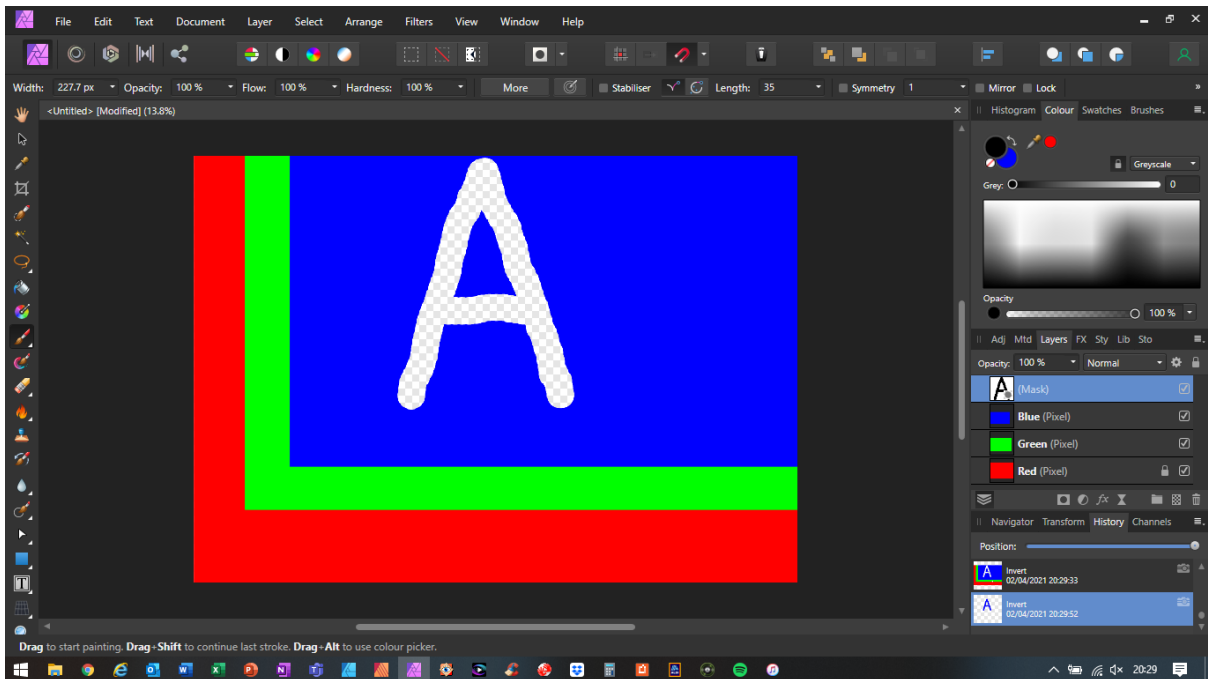


Using the right-hand square bracket on the computer keyboard, increase the brush size to something suitable:

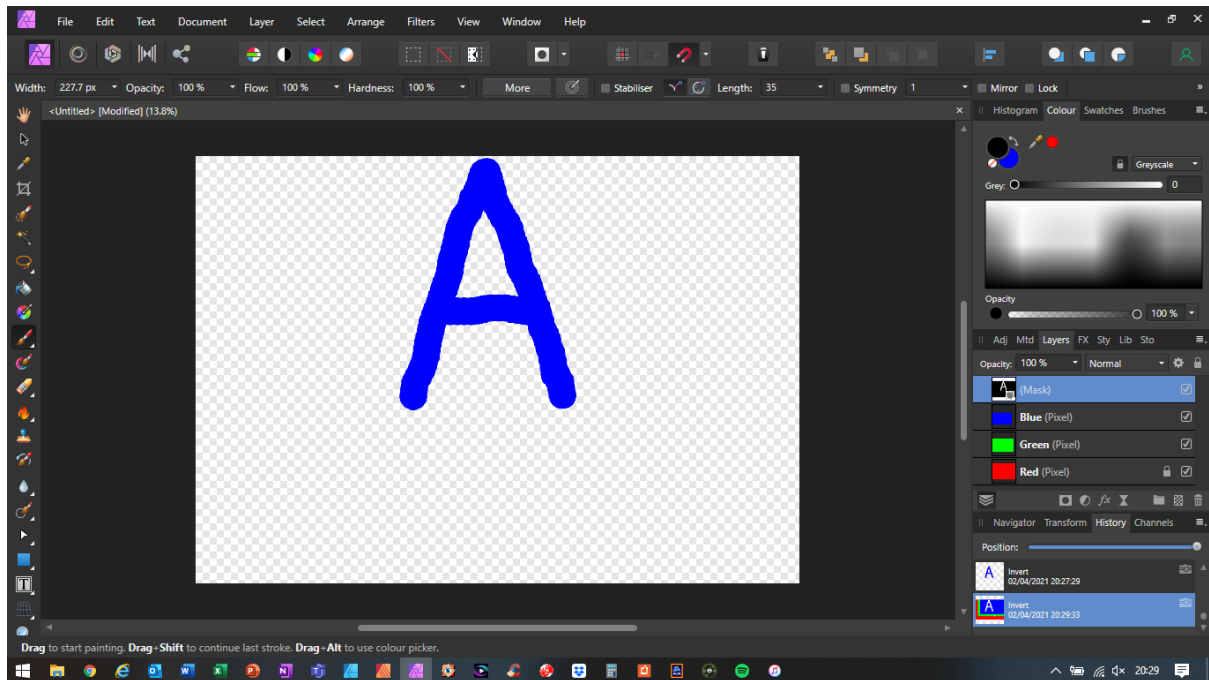


Draw a shape or area onto the main screen area as follows:

In this example the letter A has been written onto the mask. The checkerboard background is visible where the letter A is written because the paintbrush defaulted to black and the black letter A blocks light from all layers below.



If you now invert the mask again you will see that the letter A is blue and all other layers are not visible so the checkerboard background is visible, apart from the blue letter A.



Hint: Instead of creating a New Mask Layer which is created as a white mask, you can create a New Empty Mask layer which is created as a black mask.

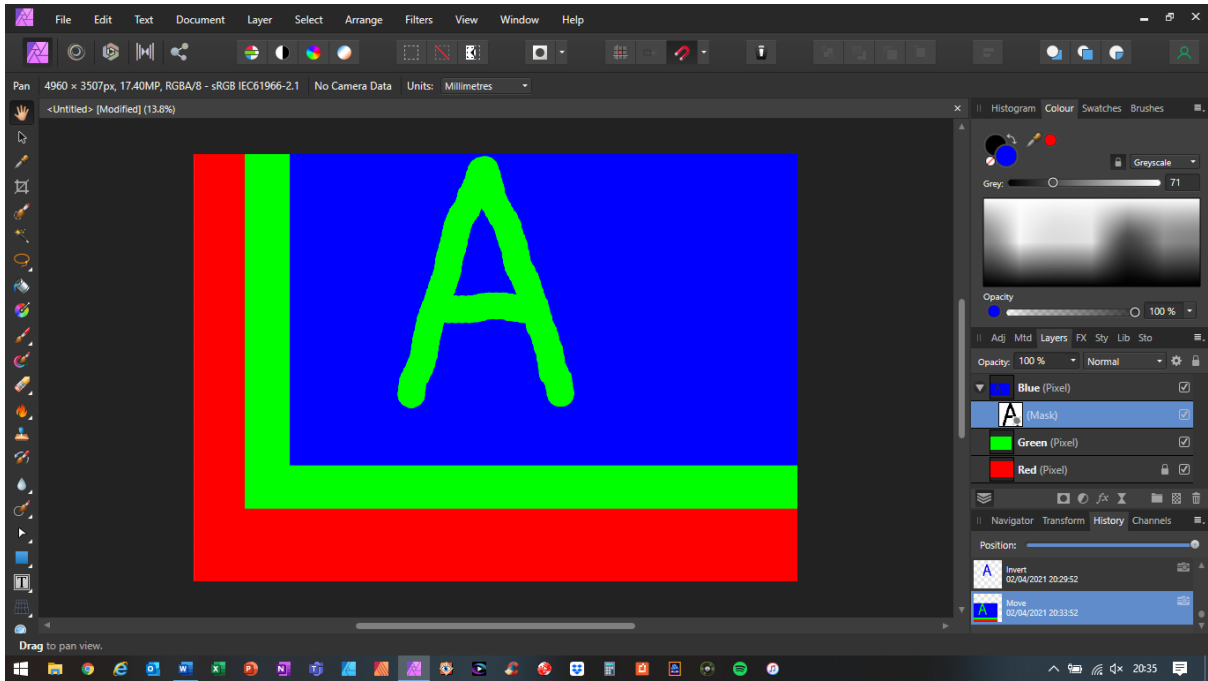
Hint: You can also create a new mask layer by clicking on the Mask Layer icon at the bottom of the Layers Panel – it looks like a white rectangle with a black circle in the middle. This creates a white mask layer but if you hold down the Alt key (or Cmd key) while you click on the icon it will create an inverted mask which is black.

Hint: You can change the colour of the mask to any shade of grey and you can change the brush to any shade of grey. This affects the luminosity of the mask hence the strength of the layers below.

Using Layers and Masks Together

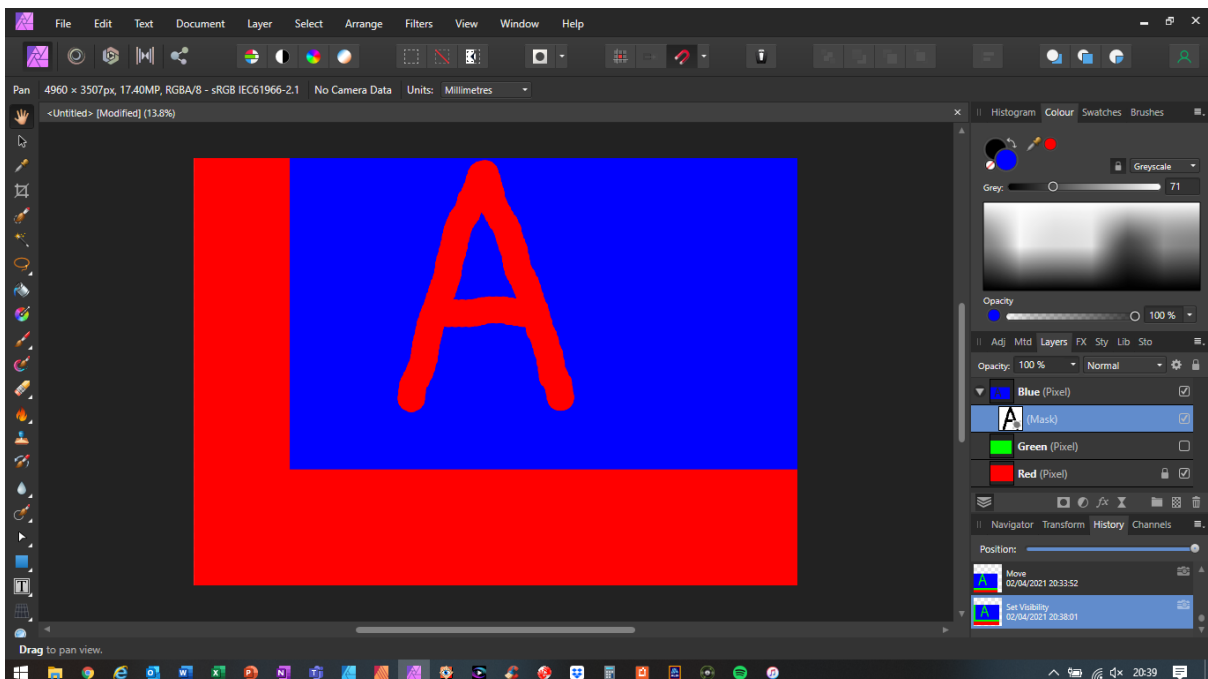
The layers and masks can be nested so there are Parent layers with Child layers.

To move the mask layer to become a Child of the blue layer you simply drag the mask into the blue layer in the Layers Panel. When it is nested you get a little triangle on the left end of the layer in the Layers Panel:




When the mask layer is moved into the blue layer to form a child layer, it only acts upon the blue layer. The black letter A in the mask blocks the blue light from the blue layer, allowing the layer below to be visible, which in this example is the green layer.

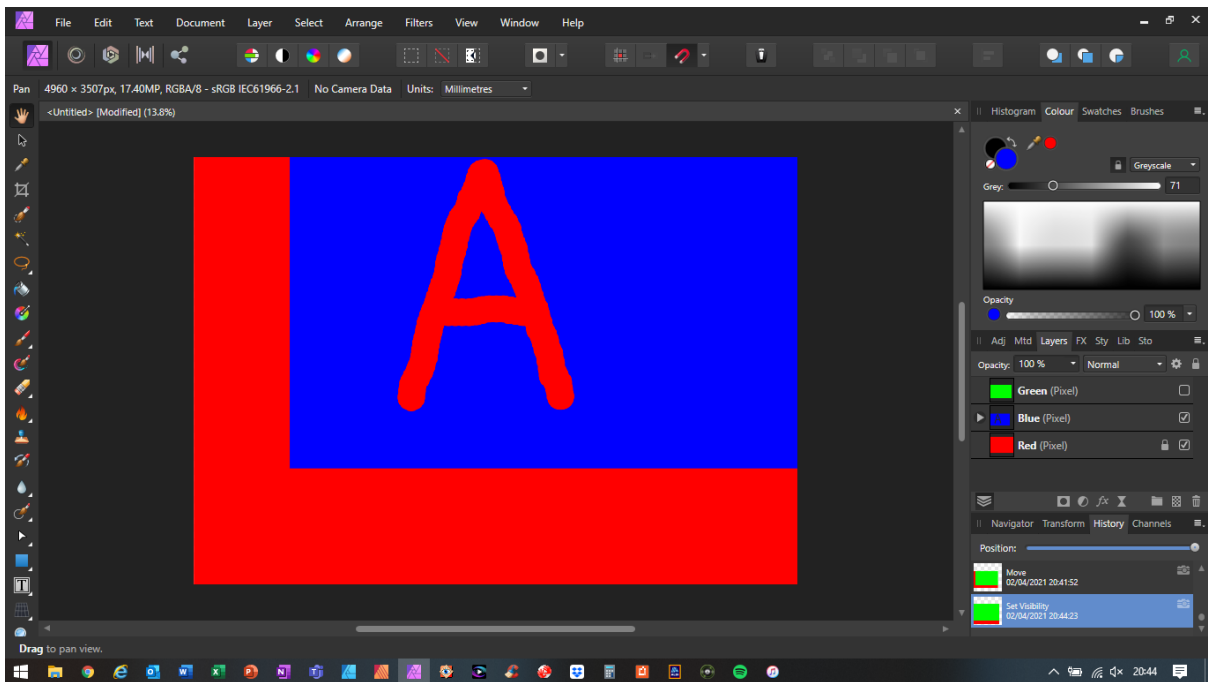
If you un-tick the green layer, leaving just the red and blue layers visible, you can see the red layer through the letter A drawn on the mask:



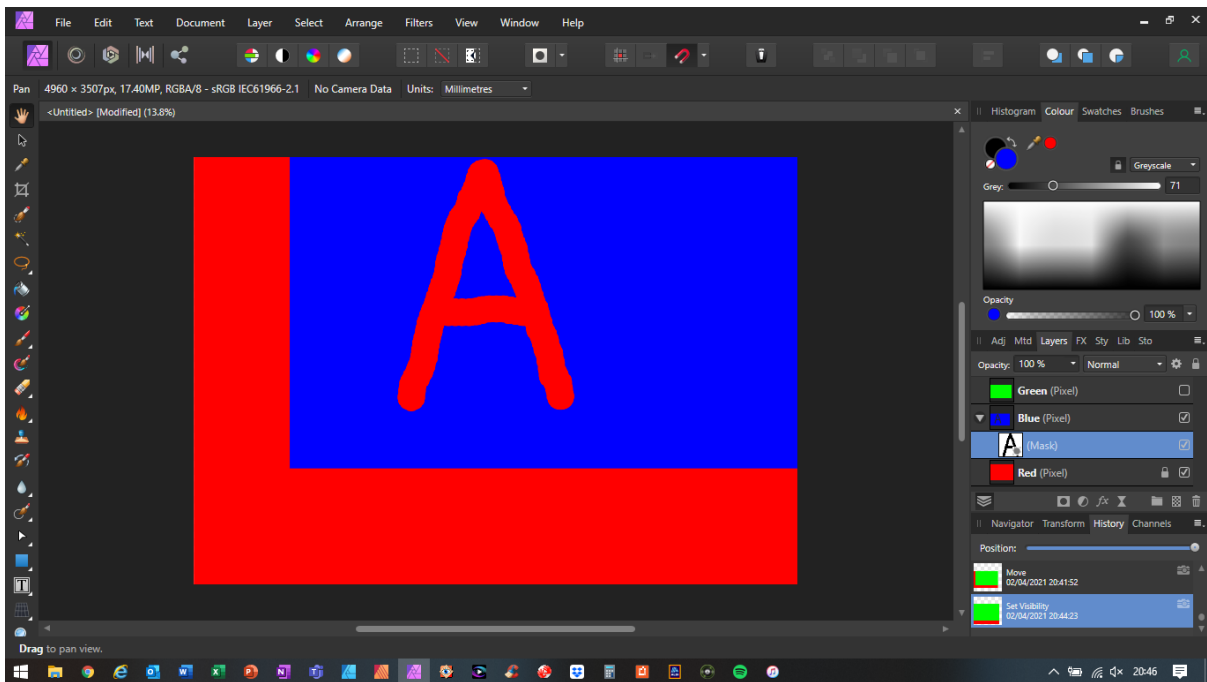
The layers, seen in the Layers Panel on the right, can be moved up or down the stack of layers to change their priority, as seen in the following screenshot:



You will see that the green layer has been moved to the top and this now covers the blue layer. The blue layer with its child mask layer is still there so it can be revealed by unticking the green layer:

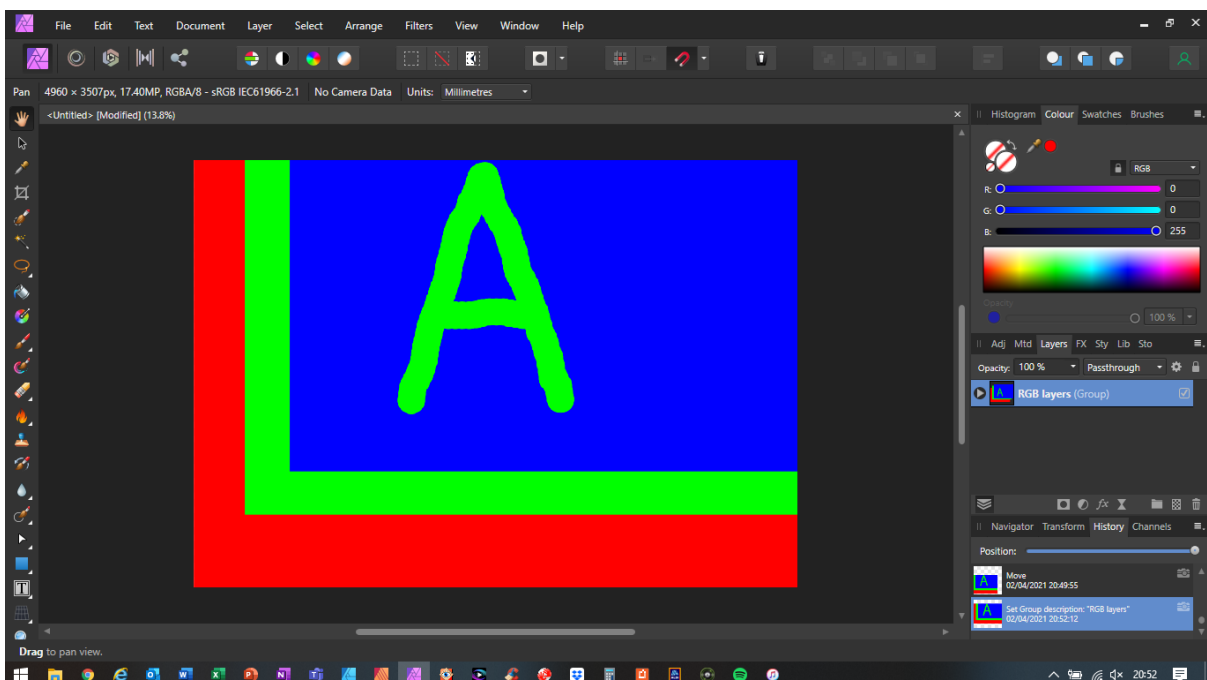


By clicking on the little triangle symbol on the left of the blue layer you can hide or unhide its child layer:



When working with a large number of layers you can group multiple layers to tidy-up the number of layers visible in the Layers Panel. To group multiple layers you select all that you want to group then either go to the Arrange drop-down menu at the top of the screen and click on the Group control; alternatively just click on the little folder symbol at the bottom of the Layers Panel.

In the following screenshot you will see I have moved the blue layer to the top then I selected all layers then grouped them. The Group is in the Layers Panel and it can be re-named if you wish:

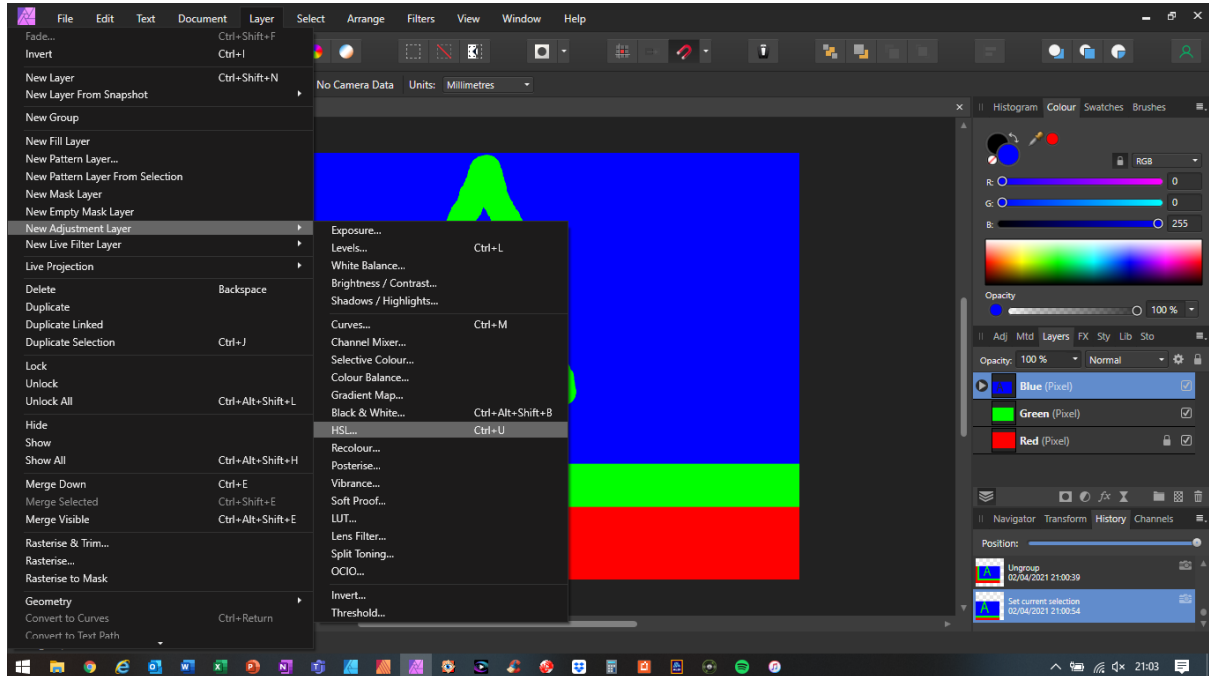


In the screenshot above you will see I have renamed the Group so it is now called 'RGB layers'

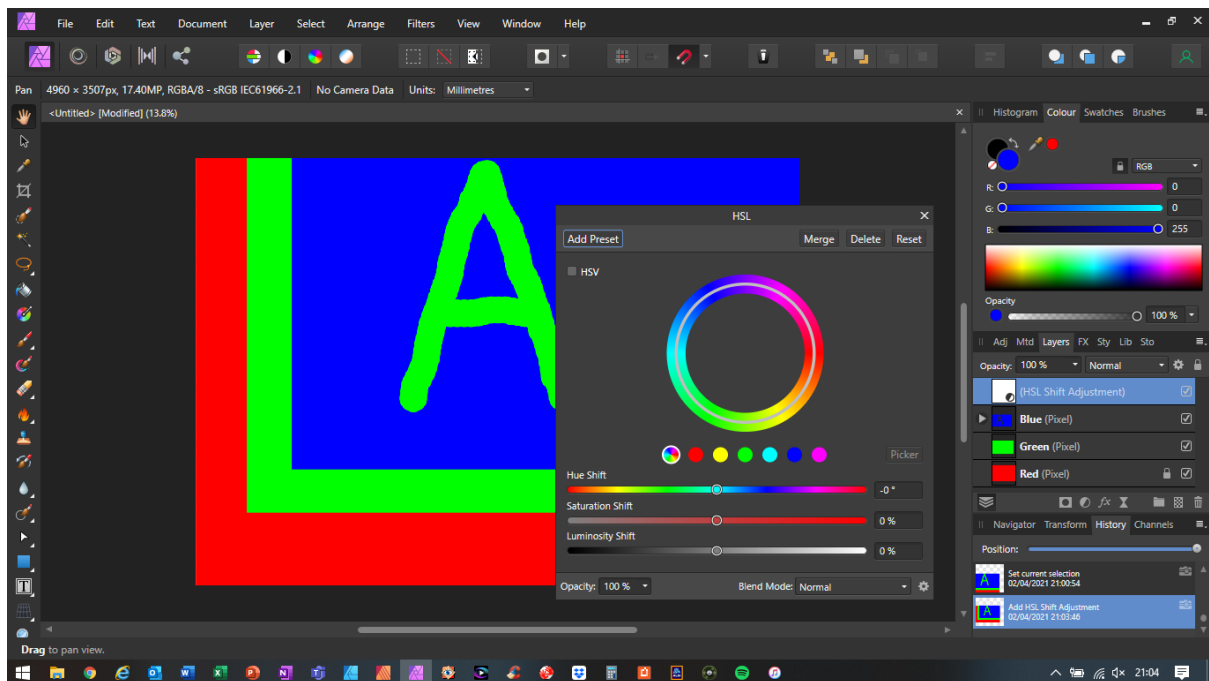
Note: It is possible to ungroup the grouped layers too, using the ungroup control in the Arrange drop-down menu.

Using Adjustment Layers

To create an Adjustment Layer, go to the Layer drop-down menu then click on New Adjustment Layer and move your mouse to the right to select the type of adjustment you want to create, such as HSL adjustments:



The HSL adjustment panel opens:



Moving the saturation slider or Luminosity slider affects all layers because the new HSL adjustment layer is added at the top of the layers. This HSL adjustment layer can be moved down to a lower position if desired. If, as described above, the HSL adjustment layer is moved into a pixel layer, as a Child, it will only affect that pixel layer (as was demonstrated for moving the mask layer above).

The HSL adjustment panel can be closed leaving the layer present (by clicking the little cross in the top right corner for windows computers or the red dot in top-left for Apple computers).

Note: If you click 'Delete' the layer will be deleted and the adjustment lost. Alternatively, if you click on 'Merge' the adjustment will be merged into the pixel layer and the adjustment panel closed so further adjustment can be made with that panel. Therefore, it is best to simply close the HSL adjustment panel, so it remains as an adjustment layer, so do not merge or delete it (unless that is what you require).

You can also experiment with creating other types of layer, such as the New Live Filter Layer for adjusting other attributes of pixel layers, such as Blur, Sharpen, Distort, etc.

Conclusion:

In the above examples I have used plain Fill Layers with the Red, Green and Blue primary colours. This makes it easy to see the layers in the screenshots and the effects of blending, masking and adjustments.

Painting on a mask with a shade of grey, rather than black or white, is a way to control mask luminosity. By changing the brush colour from black to white allows any black brushing to be undone.

These methods for creating, moving, nesting, grouping, masking, merging and blending layers enable the photographer to create a limitless number of original creative designs. Learning how to handle layers and masks in this way opens the door to many more possibilities too.

Andy Mills LRPS