The Basics of Drishti

A short instruction manual for the importation and manipulation of image files in Drishti: Volume Exploration and Presentation Tool

Written and Compiled by Bailey Lovett, Summer Research Scholar in the Research School of Biology of the Australian National University

Based on the Youtube Help Videos by Ajay Limaye of the Australian National University, Developer of Drishti



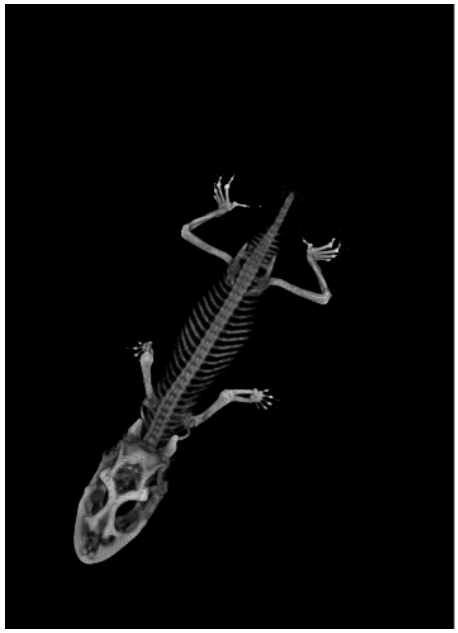

Contents

Importing an Image File into Drishti

To import an image file into Drishti you will need to use the DrishtiImport application.

1. Open the DrishtiImport application and you should see a window like this:

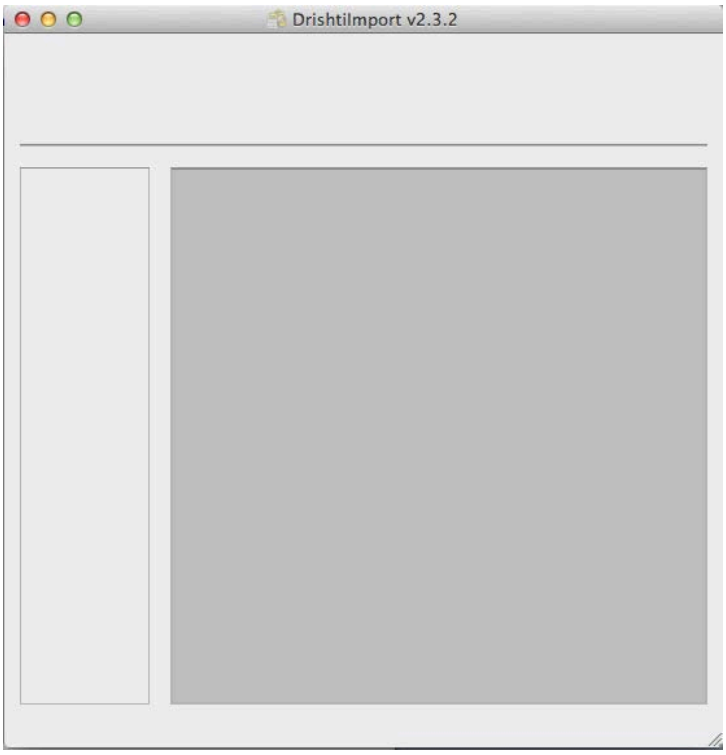

2. Simply drag and drop your file or folder (if you are using the DICOM format you will have to drag the folder in as each slice of the image has its own file)
3. The application will then prompt you with a window asking which format you would like to import the images as. Select the appropriate image format your imported files are in and click OK. The application should then load the files into the bottom image box.

4. Scroll through the files using the slider on the left hand side of the image box to check that all of your images are there.

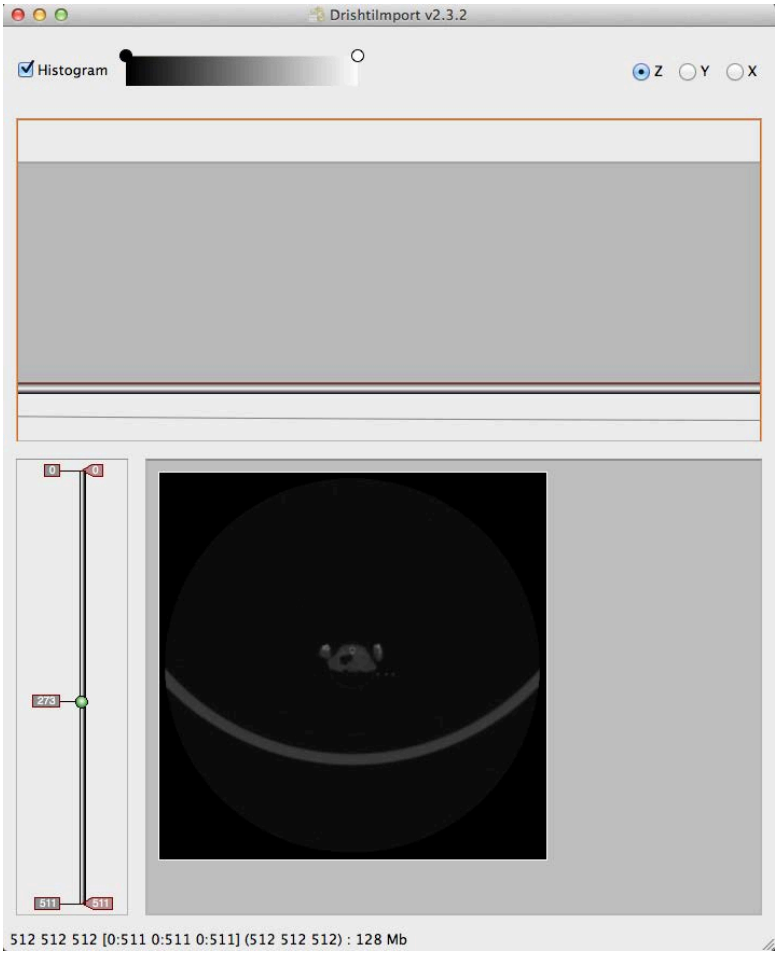

5. After ensuring that all of your images have been imported into the application, go to Files > Save As. A series of windows will then pop up:

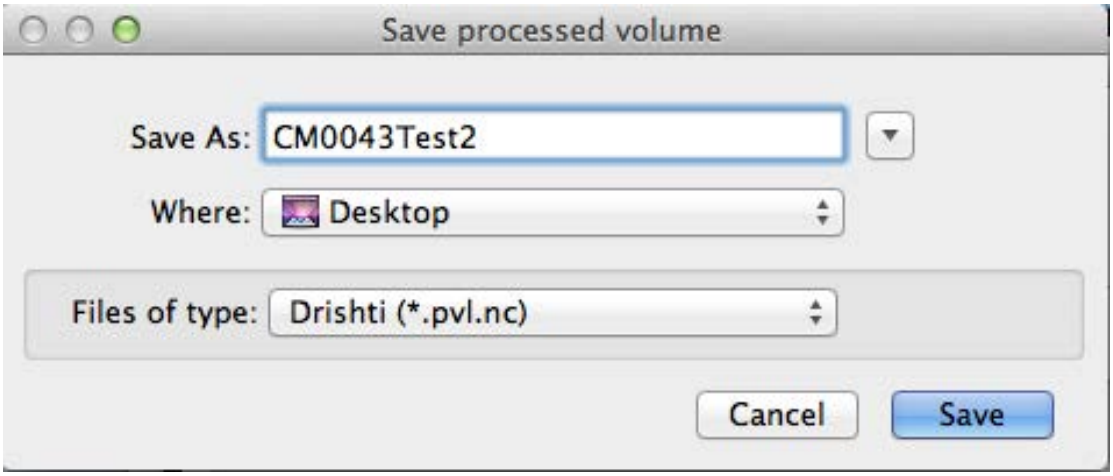

Save your imported image file under an appropriate name and select the destination you want the file to be saved into. Leave the "Files of type" as the default "pvl.nc", which is selected when the window appears, as this is the file type that is able to be loaded into the DrishtiRender program.

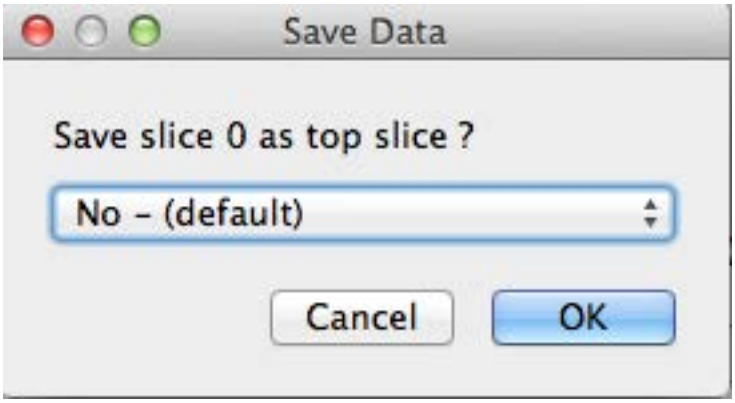

You can either choose to select No for saving slice 0 as the top slice, as selecting Yes will invert the image series, which may be unnecessary for the image series you are working with. However, with some image series, by not inverting the scan (not selecting Yes), all of the image parts become mirror images of each other. When faced with this situation, select Yes for saving slice 0 as the top slice.

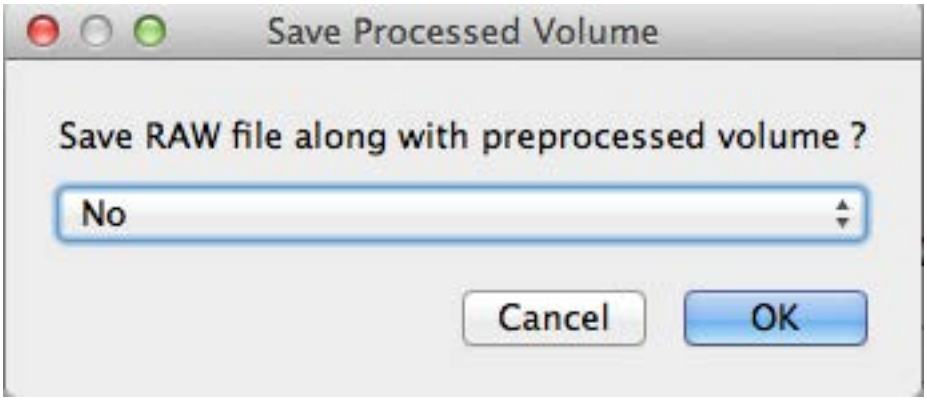

Select "No" for saving a RAW file along with the preprocessed volume, as you already should have the original image file so a RAW file is not required.

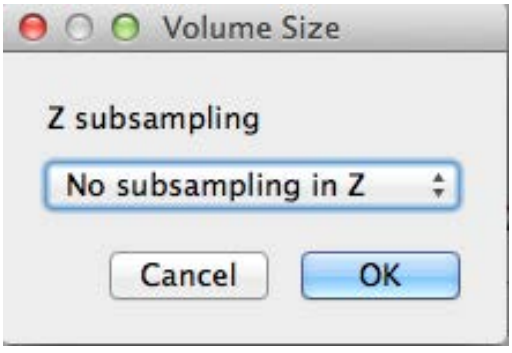

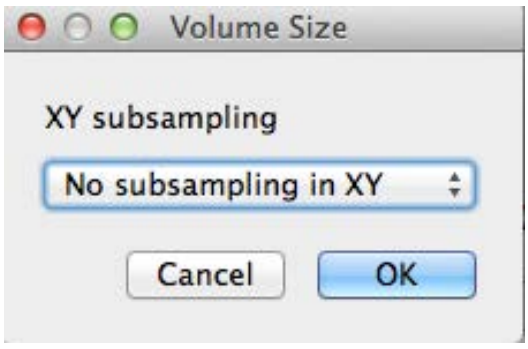

Select "No" to Z and XY subsampling if your computer is able to run image manipulations in Drishti easily and smoothly. If your computer processor is rather slow, you may want to consider subsampling, particularly if your image is large, as this will enable speedier operation in Drishti.

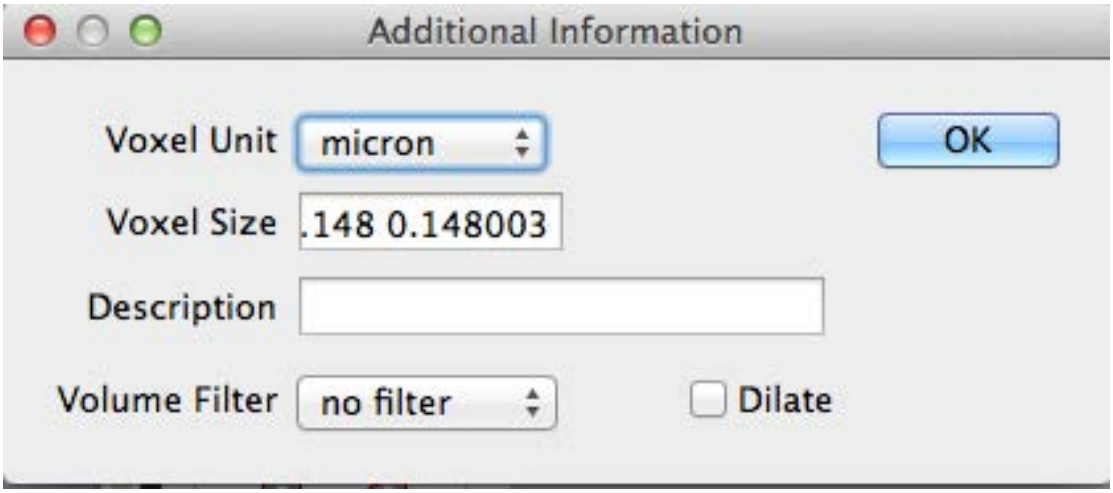

The above window is very important, particularly if you are planning on measuring aspects of your image, as the information you enter under “Voxel Size” and “Voxel Unit” will determine the scale of the image in Drishti and how accurate measurements you make of the image will be. Once you have entered the appropriate information, click OK and the image will be saved.

Opening an Image in Drishti

DrishtiRender is the main program you will use to manipulate the image files you have imported. There is an additional program known as DrishtiPaint, which is used for volume masking, however you can most likely do all you need to in DrishtiRender.

1. Opening the Drishti download folder “drishti-2.3.2” and double clicking “drishti” will give you access to DrishtiRender.

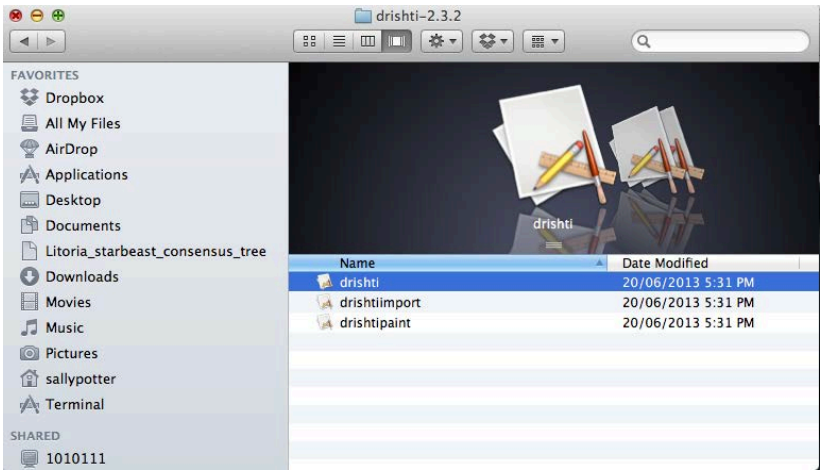

2. The DrishtiRender window should then open up and look like this:

3. To load your image (in its pvl.nc format) into the image window, select File > Load Volume > Load 1 Volume, then select your image from its destination folder. The image should then load into the program and you should be able to see a 3D representation of your file on the left hand side of the program window. (Note: When DrishtiImport saves your file in the pvl.nc format, it will also save a copy that has a "pvl.nc.001"; do not load this file into DrishtiRender, only load the file with "pvl.nc". If you want to move your data around between folders however, you will have to move both files as the pvl.nc file is a header file and is linked to the actual volume data which is stored in the pvl.nc.001 file).

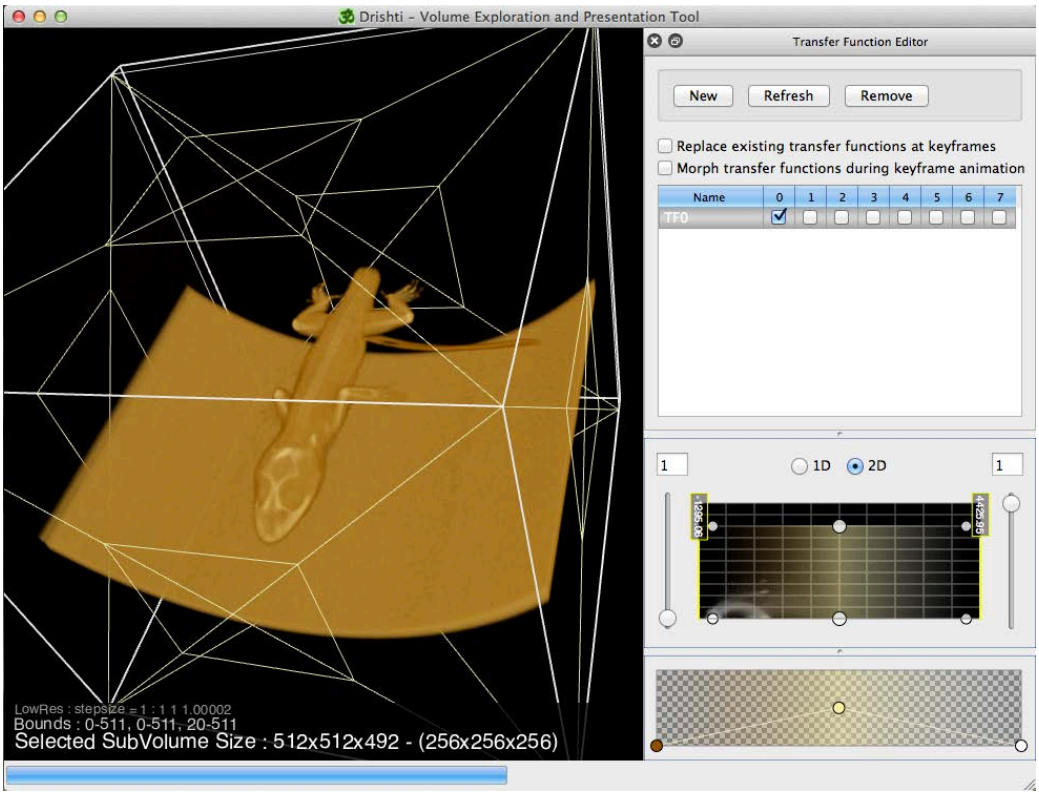

Cropping an Image in DrishtiRender

Now that you've loaded your image into Drishti, the first thing you want to do is crop the image to your liking. As you can see in the above image, there is a box surrounding my image. As I only want the lizard image by itself (i.e. without the curved surface below it), I will crop the image. To crop, simply left click and drag the crossed square (which should light up red) on each side of the box in and out until you have cropped all the unwanted parts out of your image. You can rotate the box to make sure you have cropped from all necessary sides by left clicking on the box and dragging in any desired direction. You can also move the entire image in any direction by right clicking and dragging, which helps to center the image in the field of view. To zoom in or out, simply scroll using the mouse wheel. Now, as you can see, I'm just left with my lizard image.

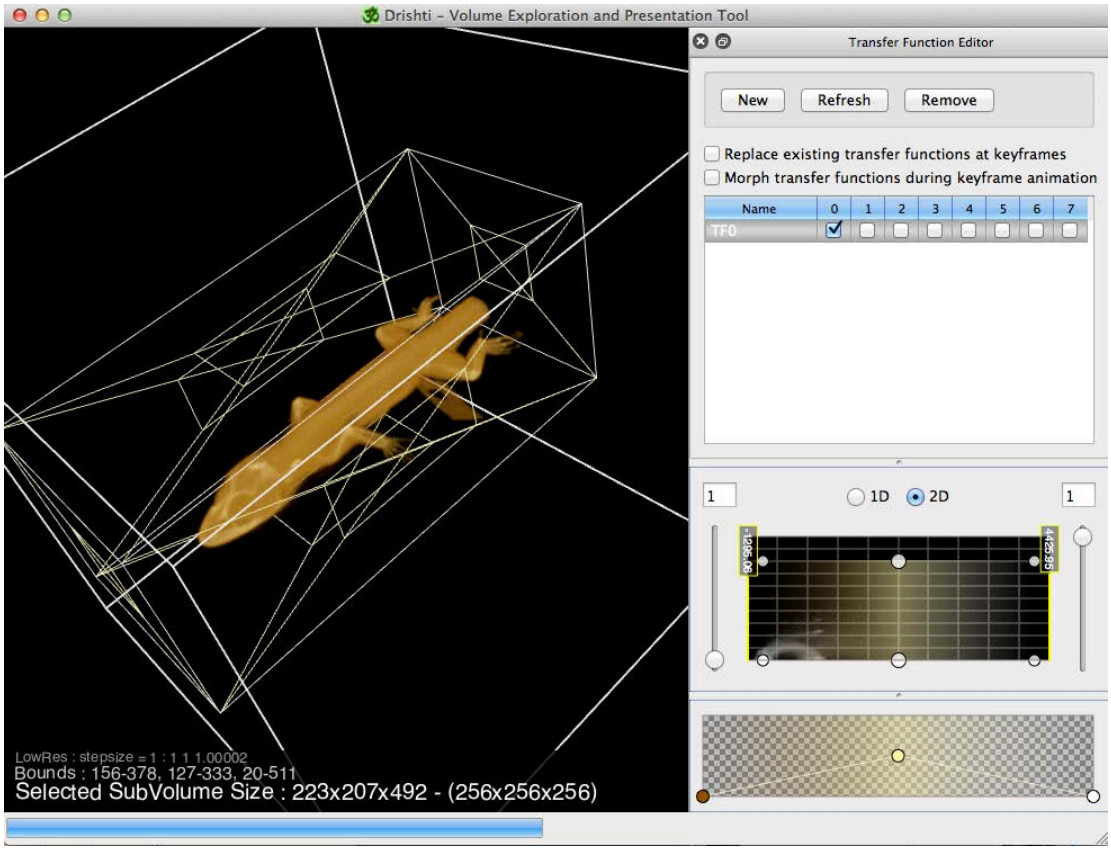

Switching to High Resolution

When you first load your image into DrishtiRender it will be in a low-resolution format, and obviously you are going to want to see the image in high resolution. Once you are satisfied with your cropping, switch to high-resolution by either going to the Toggle tab and selecting “Hires Mode”. You can easily switch between high-resolution and low-resolution mode by pressing Fn + F2 (Mac) or F2 (Windows). You can also remove the Bounding Box around the image by pressing the “b” key. Now you’re ready to start manipulating your image!

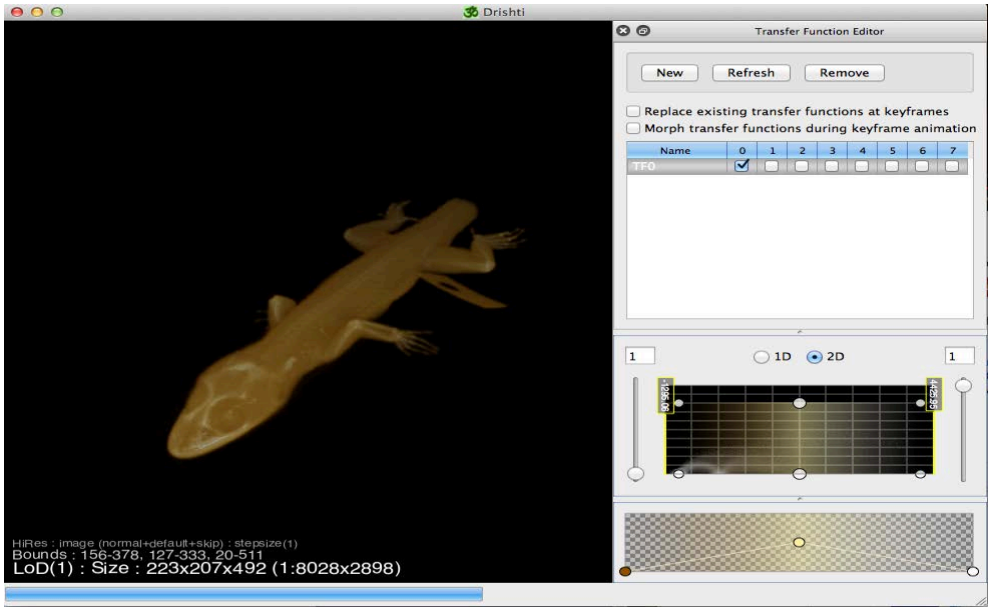

Adjusting Image Lighting & Shading in Drishti

Opening the Shader Widget

To change the lighting and shading aspects of your image, go to View and select Shader Widget. The window for this application should then pop up on the right side of your screen (at the bottom of the Transfer Function window). Click anywhere on the window and drag upwards to enlarge it, so that your workspace ends up looking like this:

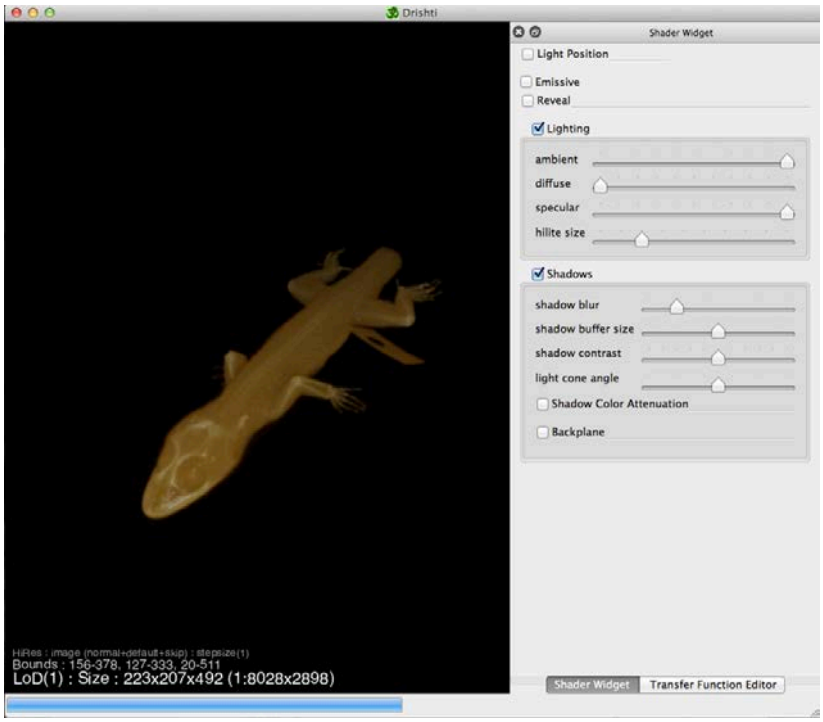

Changing the Light Position

This is very easy to change and can be done by selecting the light position box in the Shader Widget and then clicking and dragging the light position target circle until the desired light position is achieved. If you want the light to be behind the object, check the box labeled "-Z". If you want to center the light position, simply press the space bar. Following changing the light position, you are free to adjust the lighting parameters (i.e. ambient, diffuse, specular, hilite size) until your desired lighting is achieved.

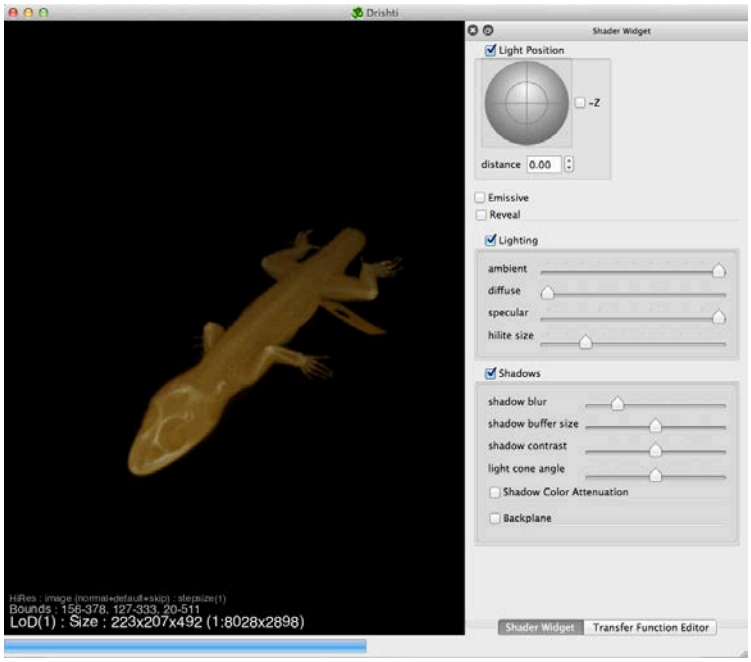

Changing the Shadow

The "Shadows" option should already be ticked on the Shader Widget and includes a range of parameters that you can change which will alter the shadowing of your image. First, you must press the "1" key to turn on shading on the image. From here, you can adjust parameters such as the shadow blur (changes the sharpness of the shadows), shadow buffer size (changes the quality of the shadows), shadow contrast (changes how dark the shadow is on the image), and light cone angle (used to change the angle of the light). You can also change the colour of the hue of the shadow using the "Shadow Colour Attenuation" option.

Backplane Option in the Shader Widget

Before using the option, the background colour must be changed so you can see the effects of using the option. This can be achieved by selecting the "drishti" tab at the top of the window and selecting Preferences. Click on the button labeled "Background Colour", and change to a colour which will make the shadows easy to see (i.e. white). The Backplane option allows you to bring the shadows to the back, with the shadow scale option changing the prominence of the shadow and the contrast changing how contrasted the shadow is to the image.

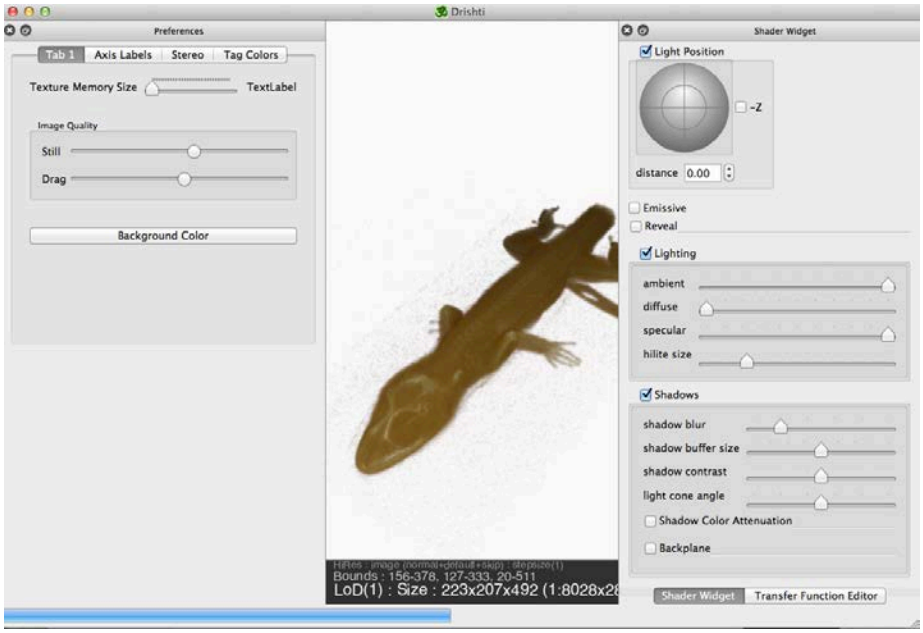

Using the Reveal function in the Shader Widget

This function is useful when working with images which have internal components that you want to observe and measure etc. Upon opening this function, you are given three options under the heading “type”. The most effective function for viewing the interior components of the image is to use the “Peel” option, which makes the top surface transparent; the “Keep Inside” option gives you a line-like drawing effect, while the “Use Shadows” function basically brings back the surface layer. You can adjust the degree to which any of these three options is effective by changing the frontface and backface sliders, where the difference between the sliders governs the amount of rendering on the image. The mix slider changes the degree to which the surface layer and interior layers are integrated. When looking inside an object it is favourable to have the ambient and diffuse sliders under the “Lighting” option high.

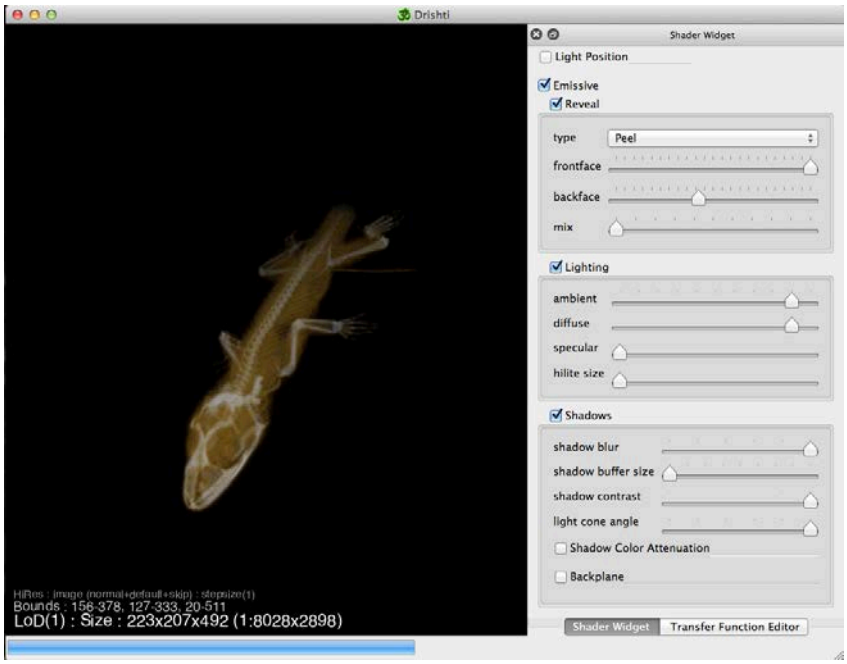

Using Transfer Functions in Drishti

Transfer functions are very useful in Drishti because they ultimately allow you to isolate different components of your image (including internal ones), enhancing your ability to accurately measure and manipulate the image.

Basic Introduction to the Transfer Function Box

The area of your window you will be using the most for manipulating transfer functions on your image will be the small multi-coloured box on the bottom right hand side of your window:

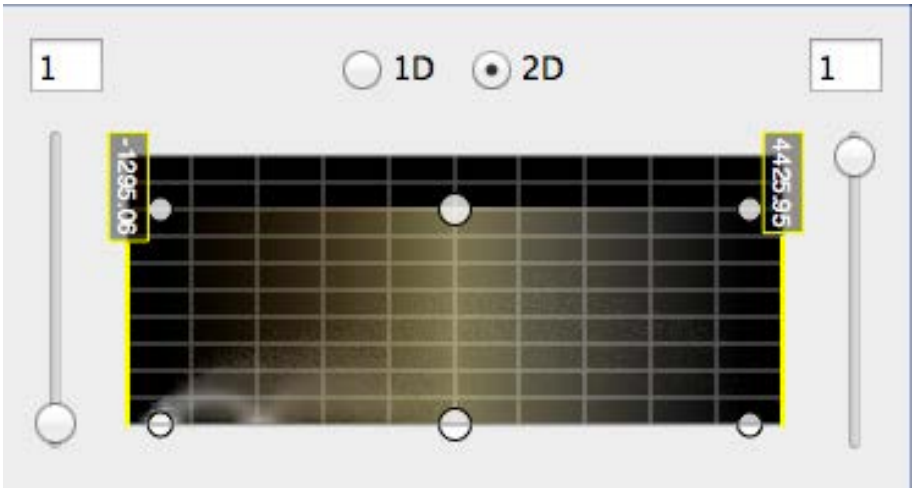

In order to manipulate your image within this box, there are several basic functions that must be known:

1. To remove the small measurement boxes from the two vertical sides of the box, hover over one of them with your mouse and press the spacebar and they should disappear.
2. To remove the grid from the transfer function box, simply press the “g” key
3. To shift one of the points (small round balls) in the box just left click and drag.
4. To shift points that are opposite each other (e.g. the two down the bottom which appear to be reflections of one another) at the same time, hold down the Shift key, left click with the mouse and drag.
5. To add a point to the transfer function, left click inside the box where you would like the point to be placed.
6. Use right click to remove a point.
7. To move the entire transfer function (i.e. all the points at the same time), hold down the middle mouse button (scroll wheel) and drag the function to where you desire.
8. Hold down the Ctrl key and press the “z” key to undo the previous action in the box. You can keep doing this until the transfer function is set back to default if you wish to start again (i.e. you can undo every operation you have done previously).

9. To redo an operation, hold down the Ctrl key and press the “y” key. Works in the same fashion as the undo function in the sense that you can redo all previous operations that you may have undone.
10. The two large points in the middle of the transfer function shape form the “spine” of the transfer function. To move the spine horizontally only, hold down the “r” key, left click and move the mouse to where you desire. To move the spine vertically, simply hold down the Shift key, left click and move the mouse.
11. If you need help remembering what keys/mouse click combinations you need to do to achieve a particular operation, you can get help by hovering your mouse in the white transfer function box (one that contains the title “Name” followed by a series of numbers), holding down the Ctrl key and pressing the “h” key. Then click on HotKey + Mouse Help to bring up the list of possible functions:

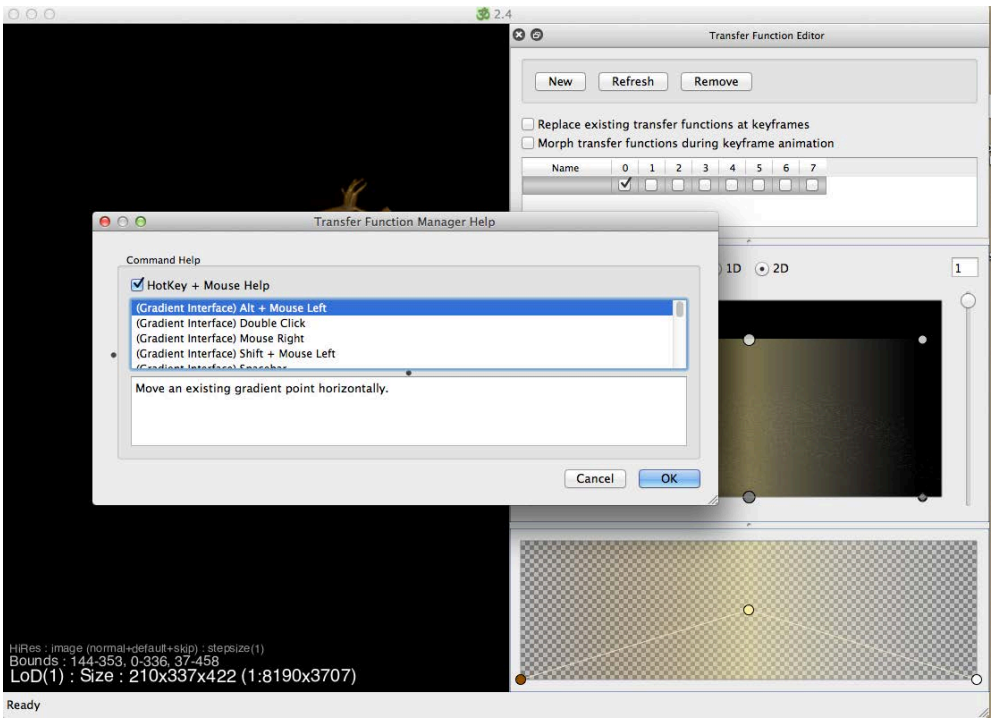

Interpreting the Colours in the Transfer Window Box

It is not very apparent in this particular image dataset, but in the transfer function box you will notice a small white arch (possibly two arches). These are essentially the phases within the image and will help to determine which parts to manipulate so that you can observe what you need to in order to carry out your measurements etc. If you click the “1D” button you will be able to see a graph-like display of where the phases are (peaks in graph correspond to different phases):

The arches essentially correspond to components of your image that are of different densities (i.e. differences in voxels). For example, in terms of my lizard image, the skin of the lizard is bound to have a different density than the bone, therefore the two arches you can observe in my transfer function window correspond to the bones and the skin or overlying tissue:

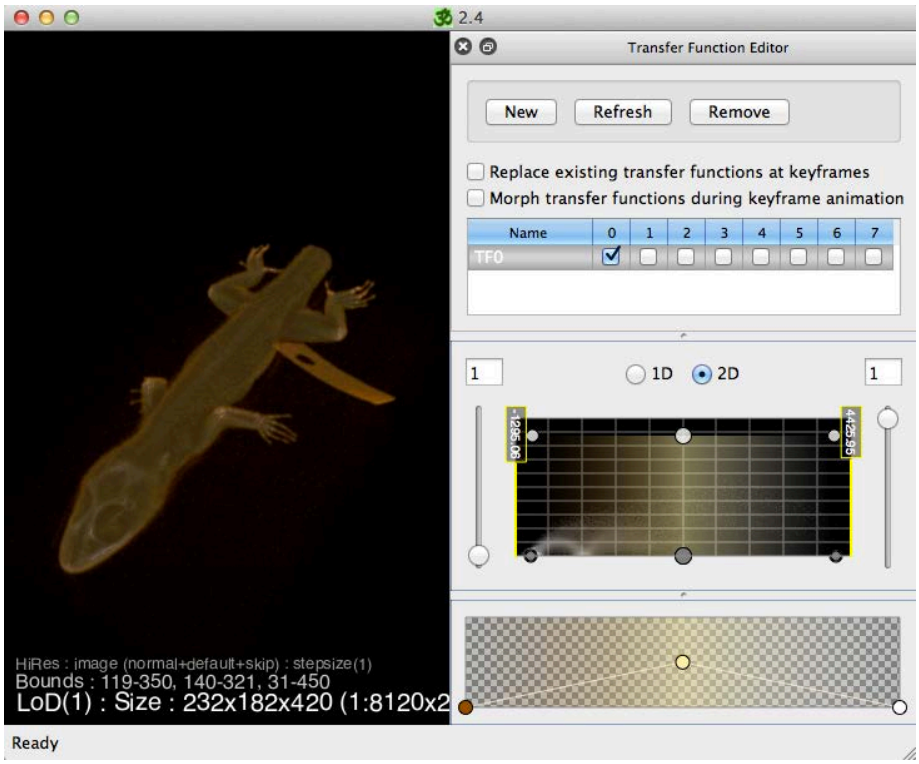

In the above image, you can observe that the transfer function area encompasses both of the arches (box includes both the defined arch (skin) and also the second faint arch (bones). As a result of this, both the skin and bones are included in the image on the left, even though the surface tissue obscures the bone. In the below image, you can observe that the transfer function area only encompasses the second arch, which means that only the bone is observed in the left image, as the surface tissue has been excluded:

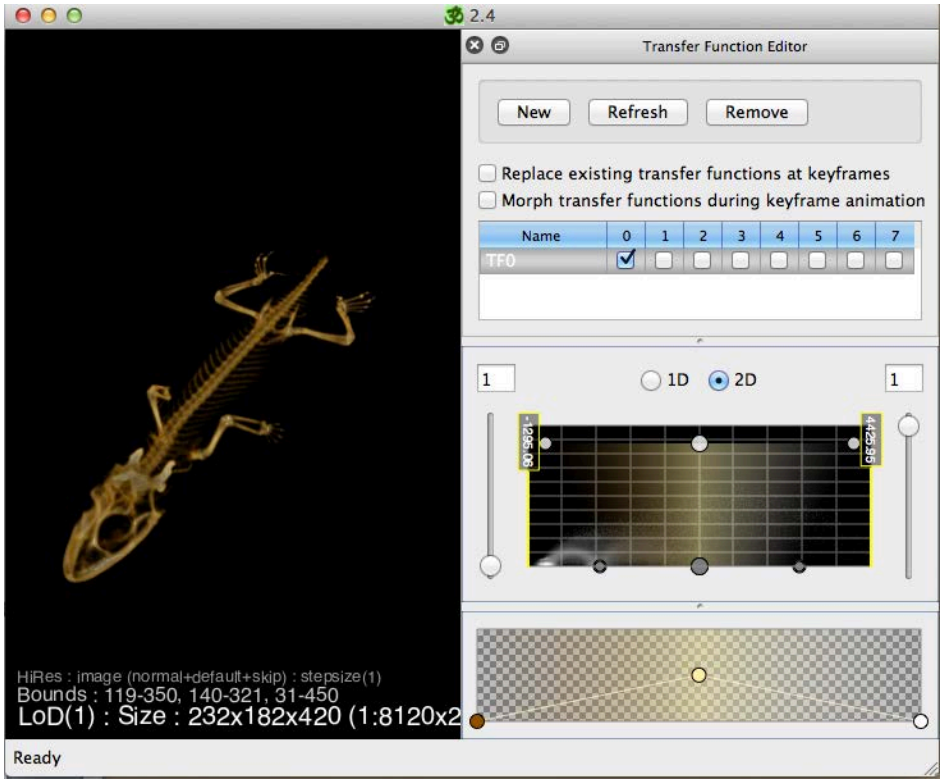

Meaning of Values in Boxes at the Top of the Transfer Function Box

On either side of the top of the transfer function box are two numbers, one in each box (they both have the default value of 1). The values of these numbers control the opacity of the image, with the right hand box controlling the opacity of the top part of the spine and the left hand box controlling the opacity of the lower part of the spine. So by setting the right hand box to 0, the opacity of the top part of the spine will be at its lowest point, and by setting it to 1, the opacity of the top part of the spine will be at its highest point:

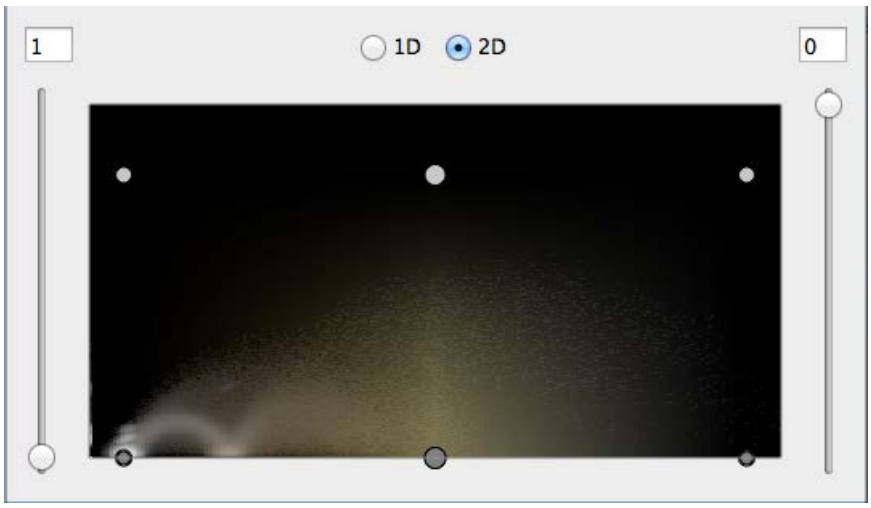

The same occurs for when the opacity value in the left hand box is changed, except that this will alter the opacity of the lower part of the spine instead of the upper:

The operation changes slightly when both opacity boxes are set to 0.5. This essentially gives an opacity gradient that is highest in the center of the spine and decreases in as it radiates to the outer areas of the box.

The point to remember about this aspect of the transfer function window is that the progression always moves in a linear fashion (i.e. from bottom to top). So no matter which way you position your points, the first set of points (i.e. those closest to the left and to the bottom) will always be the lowest and therefore correspond to the left hand opacity box, while those closest to the right and to the top will correspond to the right hand opacity box.

Function of the Sliders on Either Side of the Transfer Function Box

These sliders essentially provide boundaries for the opacity of the image, allowing you to get a more precise range for the opacity of the image (as opposed to relying on the values in the boxes). Testing out different values in the boxes and adjusting the sliders will enable you to determine the best degree of opacity for observing the components you intend to quantify/manipulate on your image.

Manipulating the Colour Interface

The colour interface is the one that is below the main transfer function, or plotting, interface. It allows you to manipulate the colours shown on the image, which correspond to the colours in the main transfer function box/interface. This commands for this interface work in the same way as those in the main transfer function interface (i.e. left click to add a point, Ctrl + z to undo an operation, etc.):

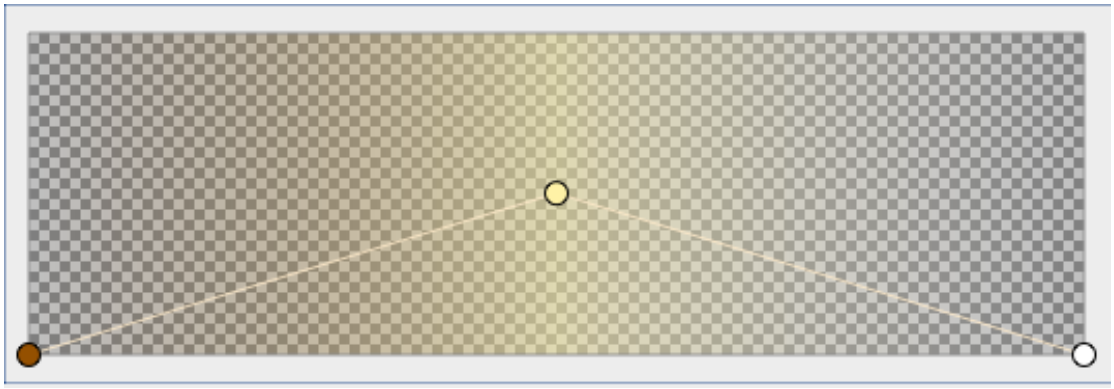

Below is some basic information about the workings of the colour interface:

1. Moving a point/colour (each point corresponds to a colour) up or down will change the opacity of that colour on the image
2. You can remove the original middle point and any added points but you cannot remove the two endpoints (white and brown dots)
3. To change the colour of a point, double click it and select the colour you wish from the colour wheel window which pops up (this will also change the colour on your image):

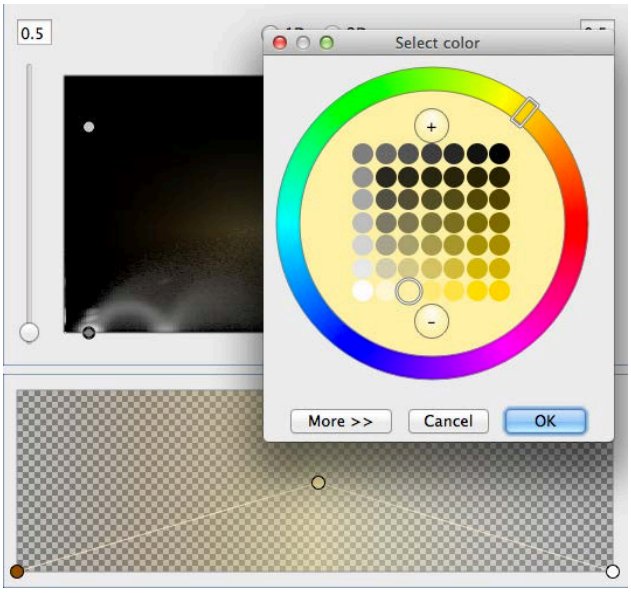

4. To save a colour combination/range you have created, press the "s" key and name your created colour range. Next time you press the spacebar over the

colour interface, a Colour Gradient window will appear and you will be able to select your created colour combination.

Naming Different Transfer Functions

For this you will be using the white transfer function window:

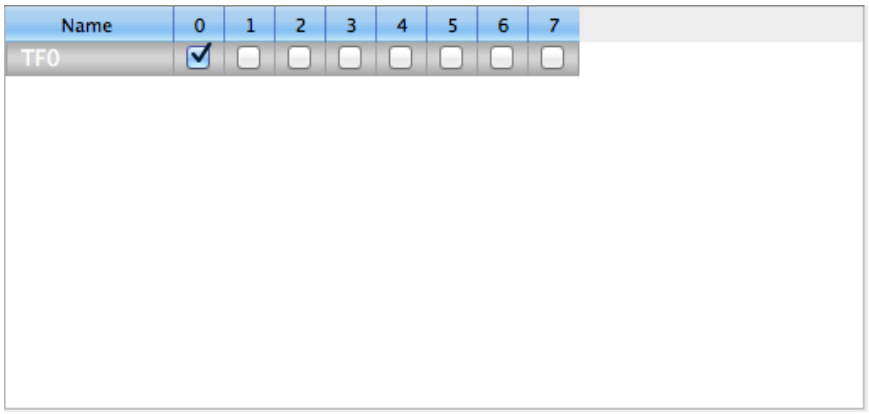

To name a function, simply hover over the Name part of your current function (i.e. where it says TF0), left click the mouse and press the spacebar. Then name the function according to what part of the image you are highlighting with that function (e.g. bones). A new transfer function can then be created underneath the first one (i.e. TF1) by clicking the "New" button, thus allowing you to differentiate between parts of the image using colour.

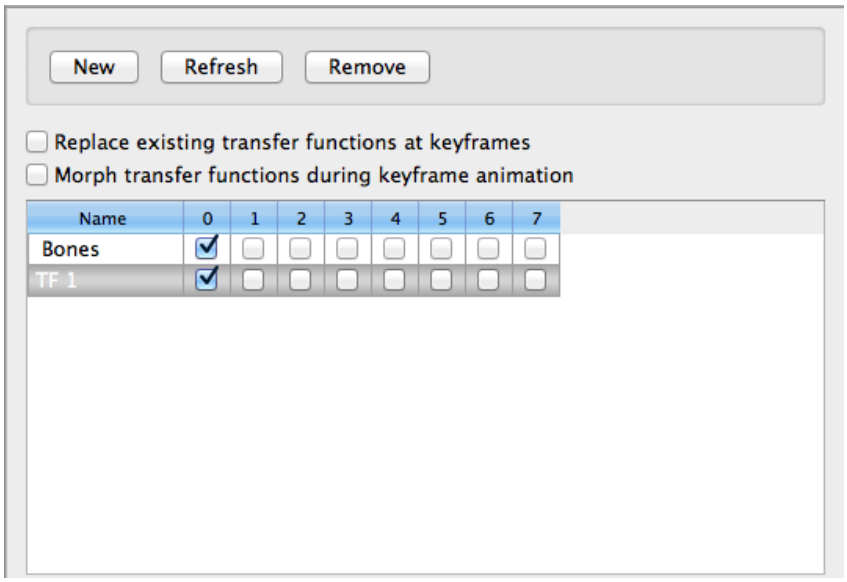

In a previous section of this part of the help manual, I discussed how it was possible to show different components of an image by moving the transfer function to encompass particular arches that represented different densities. I will now couple this with using the "New" function to enable you to switch between components of different densities without having to move the transfer function plot around in its window. I will use my skin-bone lizard example to demonstrate how it is done:

1. Create two transfer functions, one labeled “Skin” and the other labeled “Bone”:

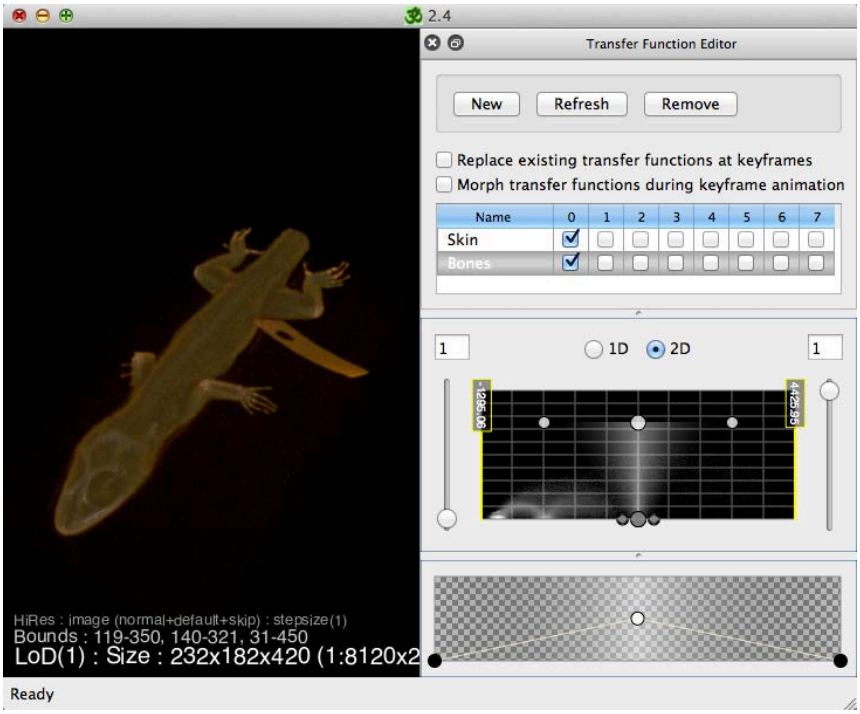

2. Starting with the “Skin” function, shift the transfer function window so that it encompasses both the bone and skin arches, or reduce the transfer function box and then, by holding down the middle mouse button, shift the box so that it only encompasses the first (skin) arch. This will serve as your “Skin” function:

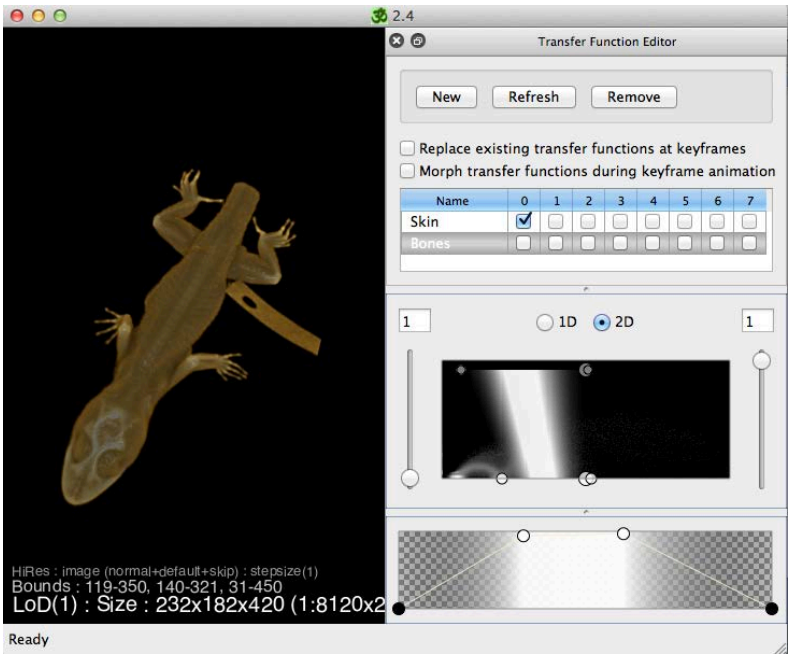

3. Now click on the “Bones” transfer function and adjust the transfer function box so that it only encompasses the second (bone) arch. This will serve as your “Bones” function. You can change the colours of this function in the colour interface by using the methods described earlier, this will enable you to see the bones better when the skin is replaced over them (when you have both transfer functions active):

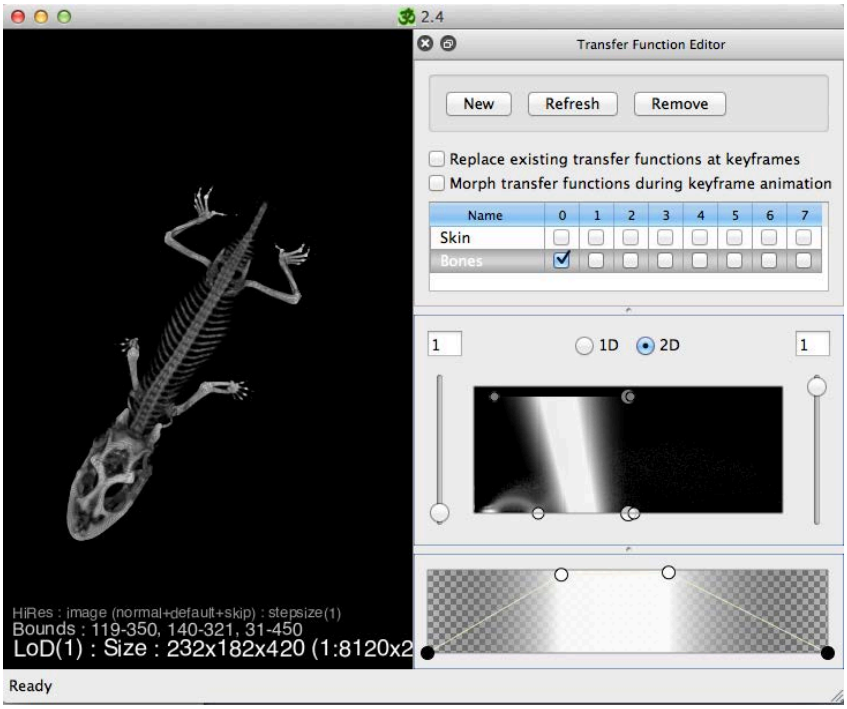

4. Now you should be able to select and deselect the tick boxes in each transfer function to display the components that you want to look at on your image. For example if I deselect “Skin” and select “Bones” I will only be able to see the bone structures, and if I deselect “Bones” and select “Skin” I will only be able to see the surface/skin tissue. If I select both functions I will still only be able to see “Skin”, because the surface tissue will obscure the bones underneath. It may take a bit of playing around the get the functions how you want them but in terms of transfer functions this is what is often required to get the image representation you require.

Carving Images in Drishti

Carving is a useful feature of the Drishti software as it allows you to remove unwanted parts of your image or components of your image that are obscuring others (i.e. surface features blocking the observation of internal components). This is a short section detailing the simple process of carving out portions of your image.

1. First set up your image in low resolution mode (check down the bottom left hand corner of the image window; if in high resolution tap F2 (or fn + F2 if you are using Mac) to change to low resolution. You can use high resolution for carving if you want, using low resolution just makes the process a lot faster.
2. Next, hover over your image and press the spacebar; this should bring up the Command Help window. Type "mop carve" into the Command String box and press OK:

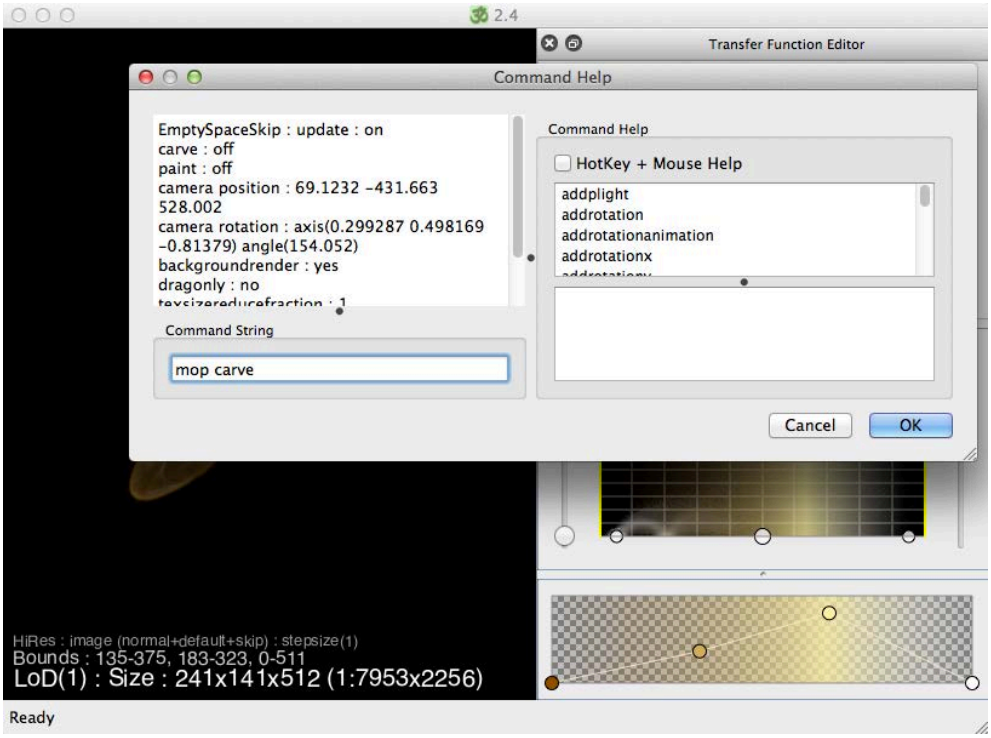

3. You will now see a small line of text in the top left-hand corner of the image window that reads, "mop (carve: 10:2). These are the default settings for the carve function, where 10 is the default radius of the carve tool and 2 is the voxel decay.
4. Tapping the up arrow will increase the radius of the carving tool; tapping the down arrow will decrease it.
5. Tapping the left arrow will decrease the number of voxels in the decay, while tapping the right arrow will increase the number. If the decay is 0, then the carve tool will punch big holes in your image rather than giving a smooth carve, therefore it is best to use a high value for the voxel decay if you are working with intricate structures.

6. Now you can start carving! By holding the Shift key down and left click and drag with the mouse, you can begin to carve out features of your image, as shown below (I went a bit overboard):

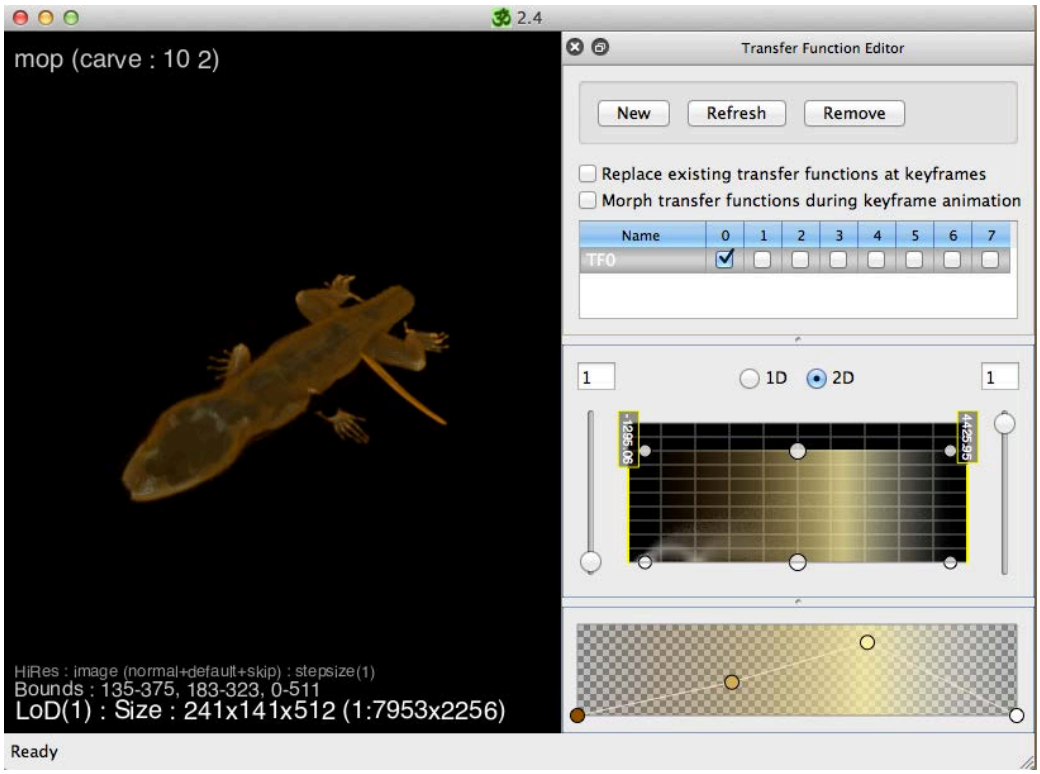

7. To undo the carving you have just made (redraw the image), simply hold down the Shift key and middle click and drag with the mouse to fill the carved areas back in.

It will take some playing around with the carve radius and voxel decay to get your carve right (as you can see the voxel decay was probably not correct when I tried to carve out the surface of the arm of my lizard). However apart from that the process of carving an image is very easy and as you may observe you can see some of the inner structures of my image.

Using Clipping Planes in Drishti

When asked to measure the height of the head and pelvis of one of my specimens, I ran into problems using the path point method, as there were natural lesions in the head and pelvis which caused the points to be placed in the wrong location and the surfaces I was trying to measure were very convoluted, thus giving me inaccurate measurements. Then I discovered clipping planes, which essentially allow you to produce a cross section of the image (i.e. cut down the middle of the body exposing a flat plane surface), which enables simpler measurement of the height of particular structures.

To add a clipping plane to your image to enable the measurement of tricky surfaces/structures, simply follow the steps below:

1. Have your image set up in one of the "normal planes" (i.e. either vertical or horizontal as shown below) by rotating the image to the general angle and then double clicking the left mouse button to position it fully on the plane:

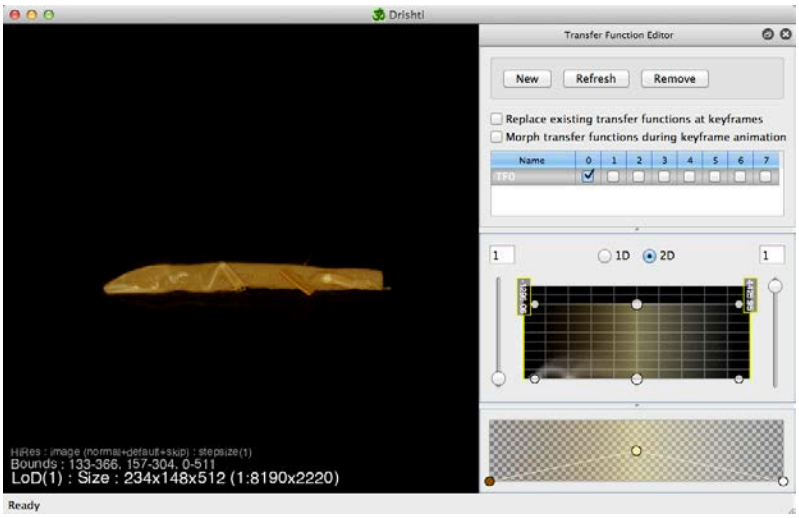

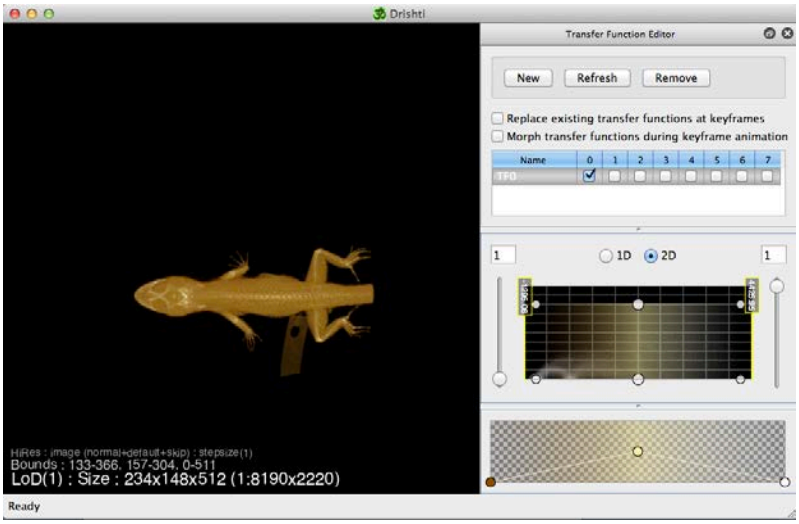

I will be using the second position when running through this process as it allows me to see what I am doing more clearly, but feel free to use the other plane if you find it easier to work with.

2. Next, you want to add your clipping plane to the image. To do this, simply press the “c” key when hovering the mouse over the image window. The clipping plane should appear and automatically should remove one half of your image, depending on the orientation of the image:

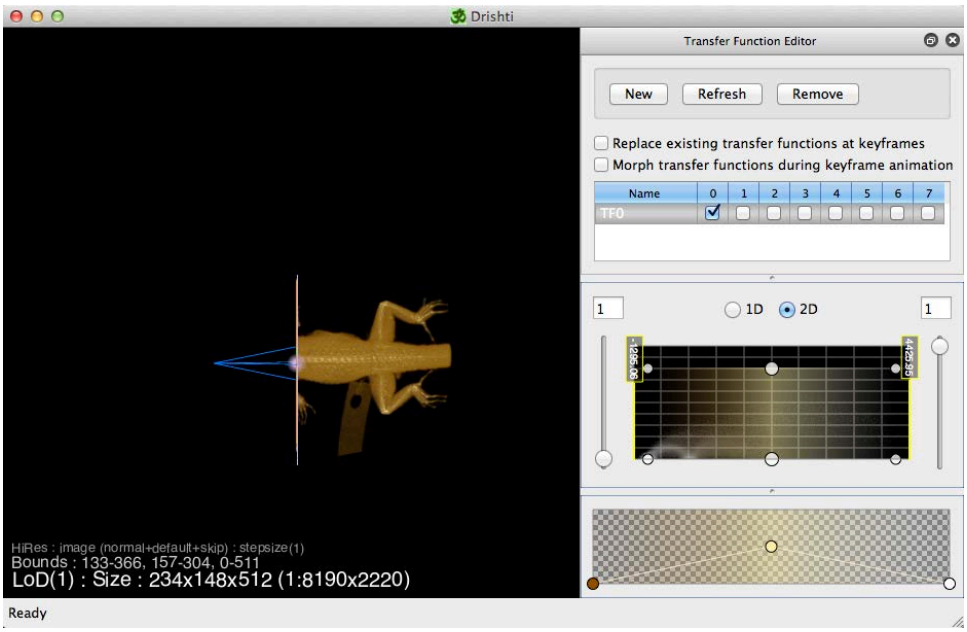

If you rotate your image slightly you will see that there are 4 different coloured shapes comprising the clipping plane; a large purple square, a green diamond, an orange diamond and a blue pyramid. These are your clipping plane axes:

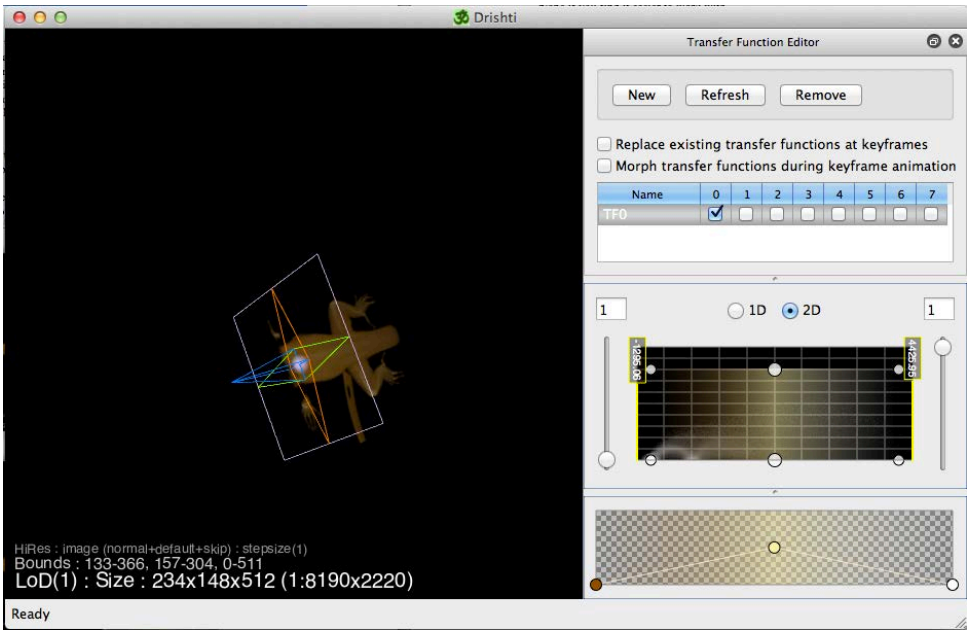

3. Now you will use your axes to rotate the clipping plane on the image to achieve the cross section of your image that you desire. You can click on any of the axes and drag the clipping plane into the correct position:

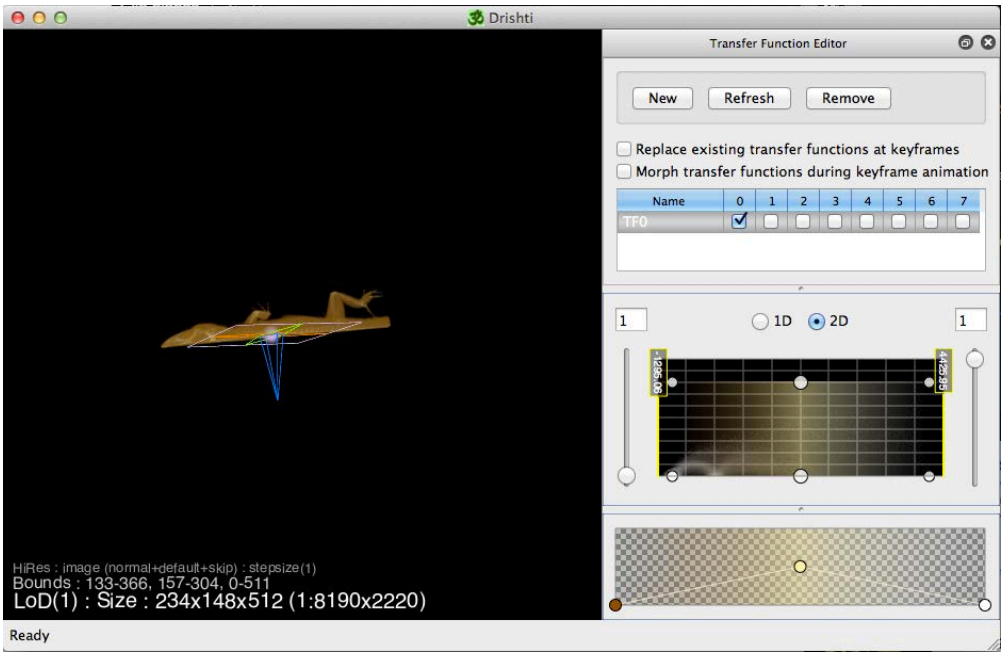

4. Conversely you can rotate it a particular number of degrees using the command window, which will appear when you press the spacebar while hovering your mouse over the clipping plane. When the command window appears, type "rotate" into the command string:

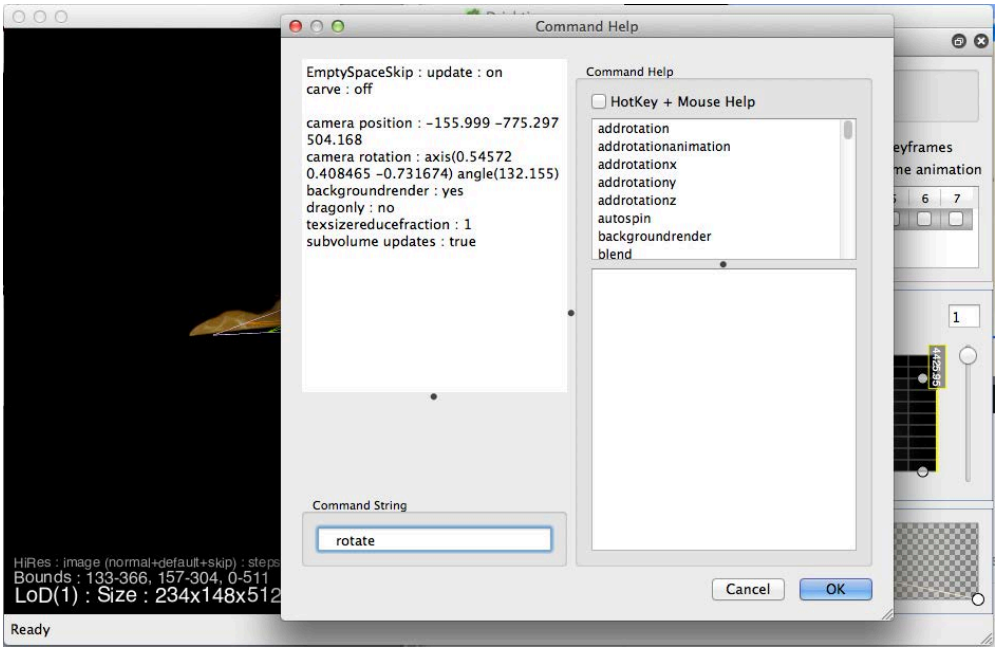

Below the "rotate" function on the Command Help list there is another function called "rotate (a|b|c). Click on this function as this is the one you will be using to rotate your clipping plane:

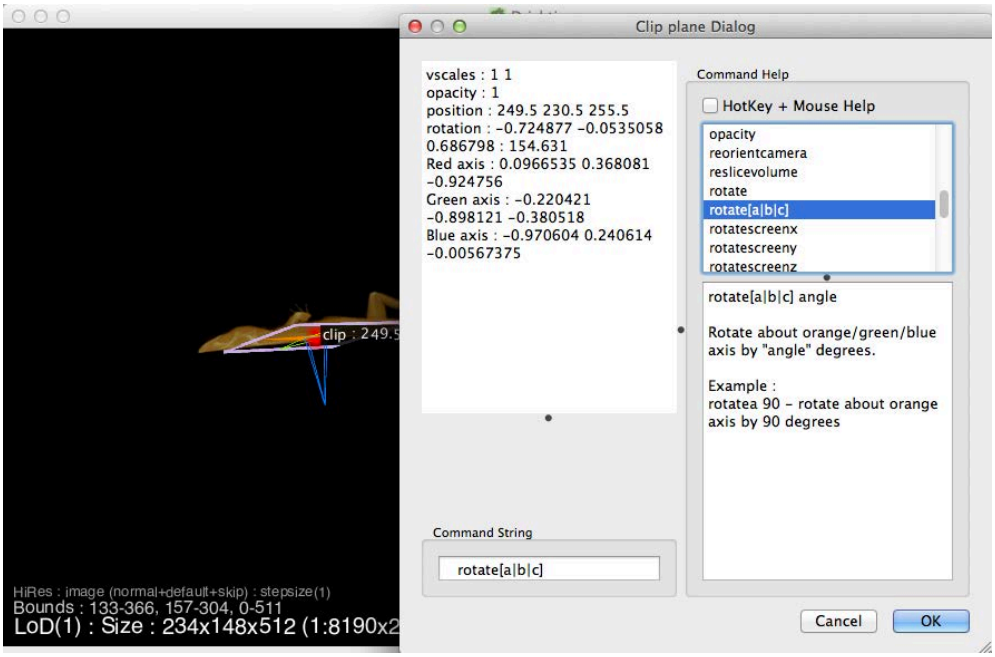

There are three axes that you can rotate the clipping plane about through the command window: the “a” axis (orange), “b” axis (green) and “c” axis (blue). The orientation you want your clipping plane to be in will determine which axis you will specify in the command window. The “a” and “b” axes are most useful to use in my case, however if you are wanting a horizontal cross section through your image (rather than a vertical one as in my case) then the “c” axis will be helpful to use. Specify your axis by entering either the letter “a”, “b”, or “c” directly after the “rotate” command (make sure the function is set to “rotate (a|b|c)” and not to “rotate”). Following your axis specification, you must specify the number of degrees to which you would like to rotate your clipping plane to. It is usually best to rotate in increments of 90 degrees. For example, if I have my clipping plane originally set in the position it is in the image above, and I type “rotatea 90” into the command window, I will end up with a plane which looks like this:

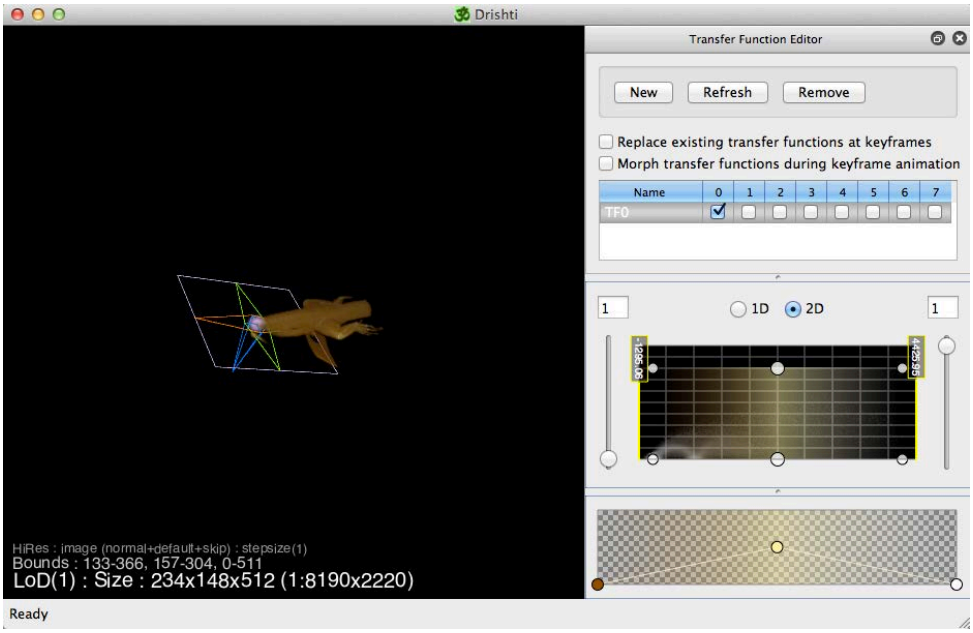

My image is on a slight angle because I was messing around with clicking and dragging the clipping plane with my mouse before I began using the command window to adjust its position (best not to do this if you want your clipping plane position to be accurate). It will probably take a bit of trial and error with rotating the axes to get the result you desire, but once you have worked out the axes you need to rotate and how much to rotate it by, the process should become very simple and allow you to get measurements of those tricky structures.

5. If you want to shift the "c" axis, which controls how far the clipping plane is vertically positioned on the image, simply right click and drag the blue axis up and down the image until you have achieved the desired vertical cross section. As shown below, I have moved the clipping plane on top of the image and have dragged down the blue axis so that the top half of the image has been removed:

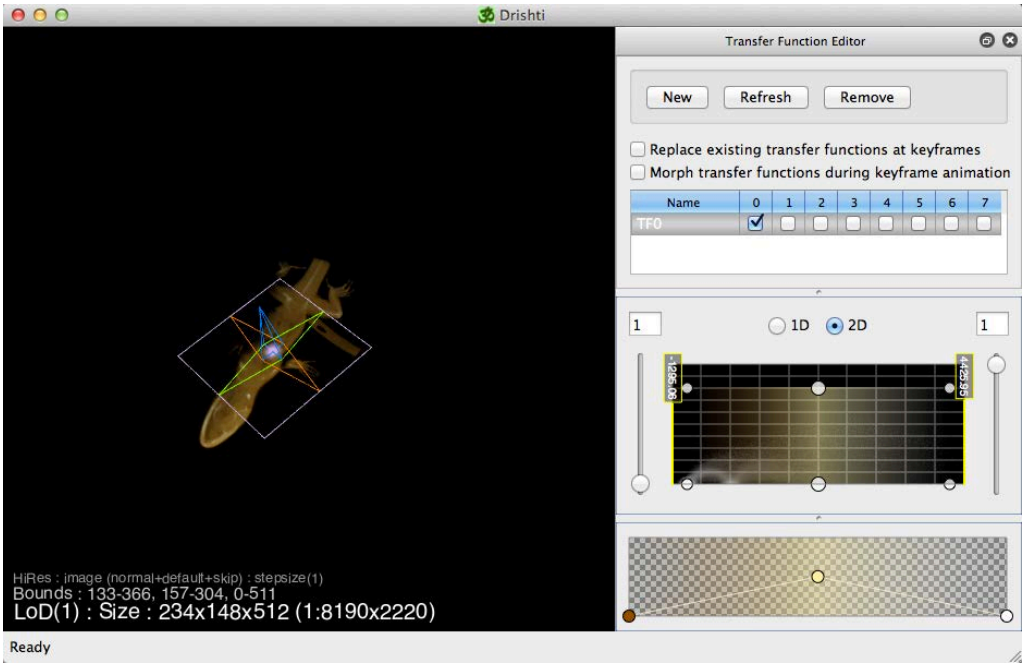

6. To remove the clipping plane from the image window, simply press the "v" key, which will prevent it from being visible on your image. To get rid of the clipping plane entirely, hover over it with your mouse and press the Delete key. Your image will return to its original condition (no cross sections).
7. Once you have prevented the clipping plane from being visible on the image, you can take measurements from the cross section you have created, as shown below:

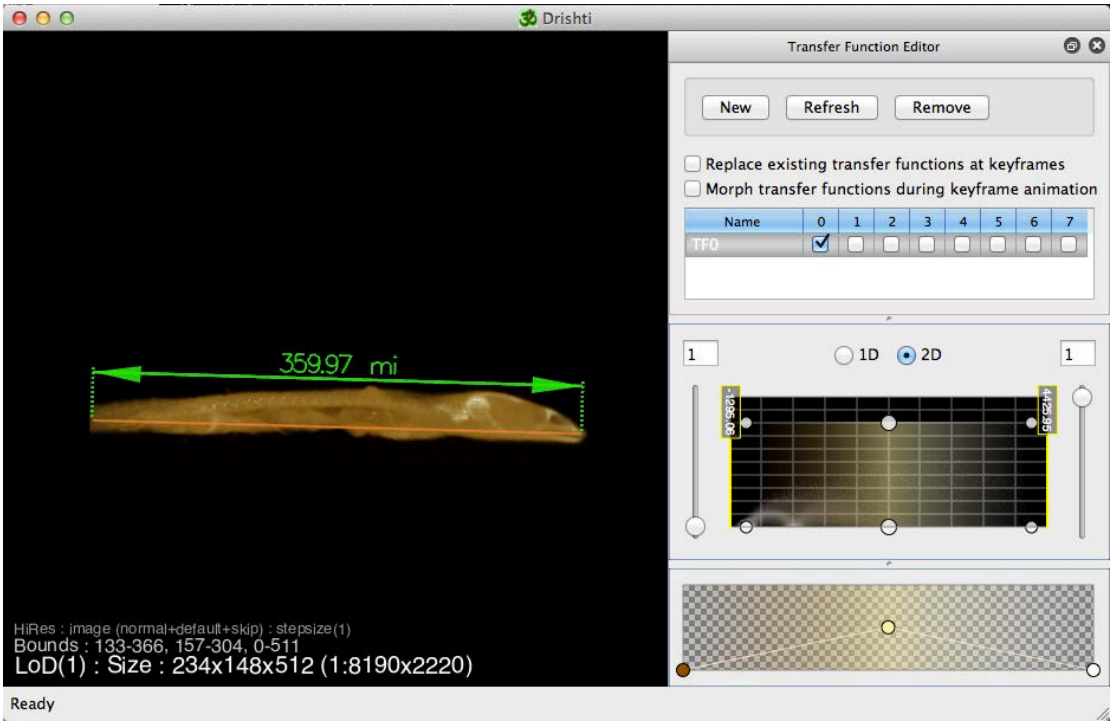

In conclusion, the clipping plane feature has allowed me to create cross sections of my image so that I can measure tricky structures on my specimen without the significant issue of added points ending up in the wrong place on my image, ultimately saving me a lot of frustration!

Taking Length & Angle Measurements in Drishti

Length and angle measurements on an image are very easy to take on Drishti, thanks to the point-and-path method that the program allows for.

Taking a Length Measurement (Straight Line)

1. Firstly, you must indicate on your image which part you would like to take a length measurement of, which can be achieved by adding points onto the image. When adding points to take a measurement, it is best to orientate the image face on to you, which is achieved by rotating the image so you have a bird's eye view of it and then double clicking the left mouse button to ensure that it is directly face on. This is crucial for getting an accurate length measurement, as the program will put the point on the surface of the image which is closest to the viewer, so if your image is not orientated correctly you will end up with points in random places. To add a point, simply hold down the Shift key and left click on one end of the part you want to measure. The point should glow green, indicating that this is the point that is currently active (will turn blue following this but will go green again if you left click on it). Now do the same at the other end of the structure you are trying to measure, so that you are left with 2 points, one at each end of the structure.

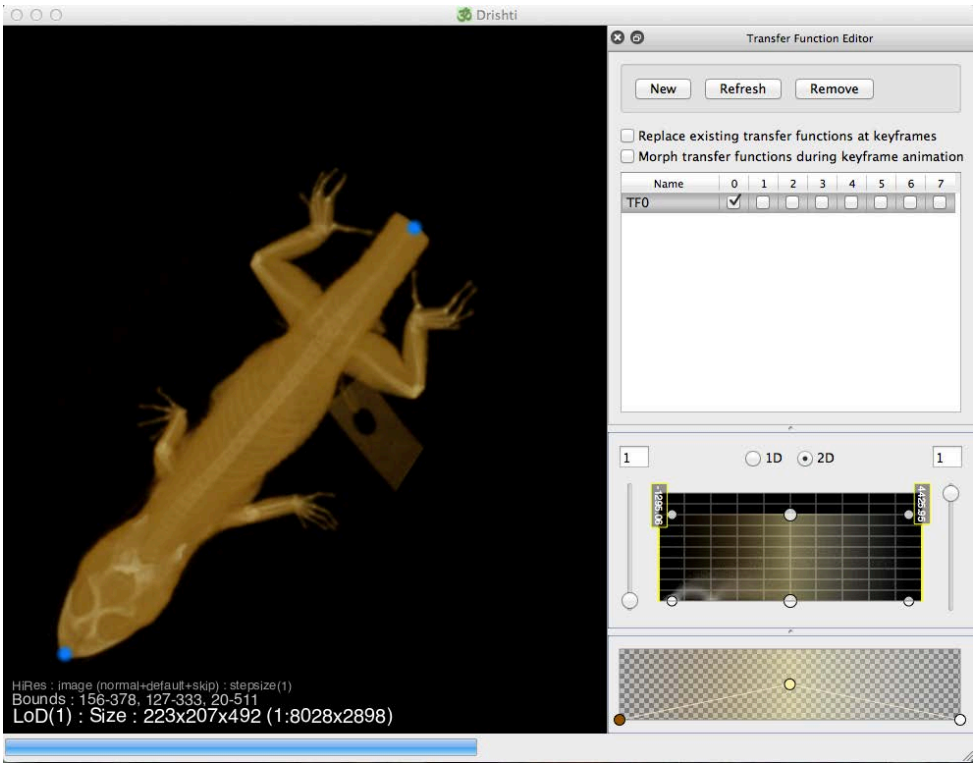

2. Now we need to draw a path, which can be achieved by tapping the space bar when the mouse is hovering over the image, which brings up the command window. Type "path" into the command window and press Enter. The program will then draw a path between the two points (I have shifted the path

slightly off the image in my case as the dorsal part of the lizard is curved so the path has disappeared into it).

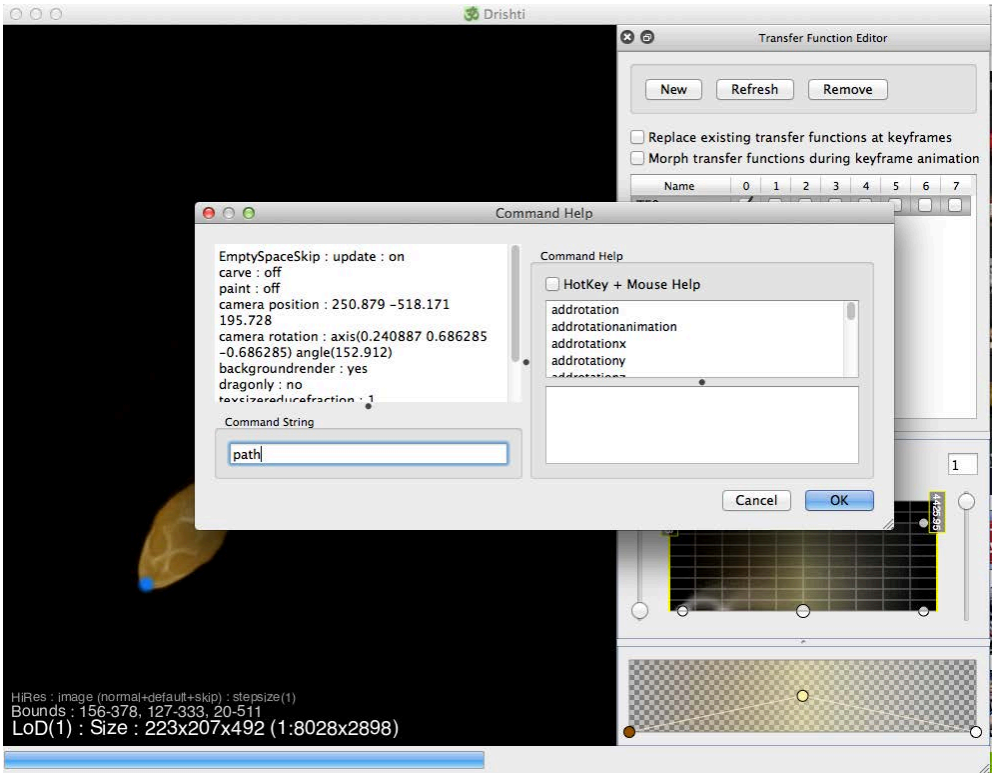

If you hover your mouse over the path, you will be able to see how long the path you have drawn is, and hence the length of the structure.

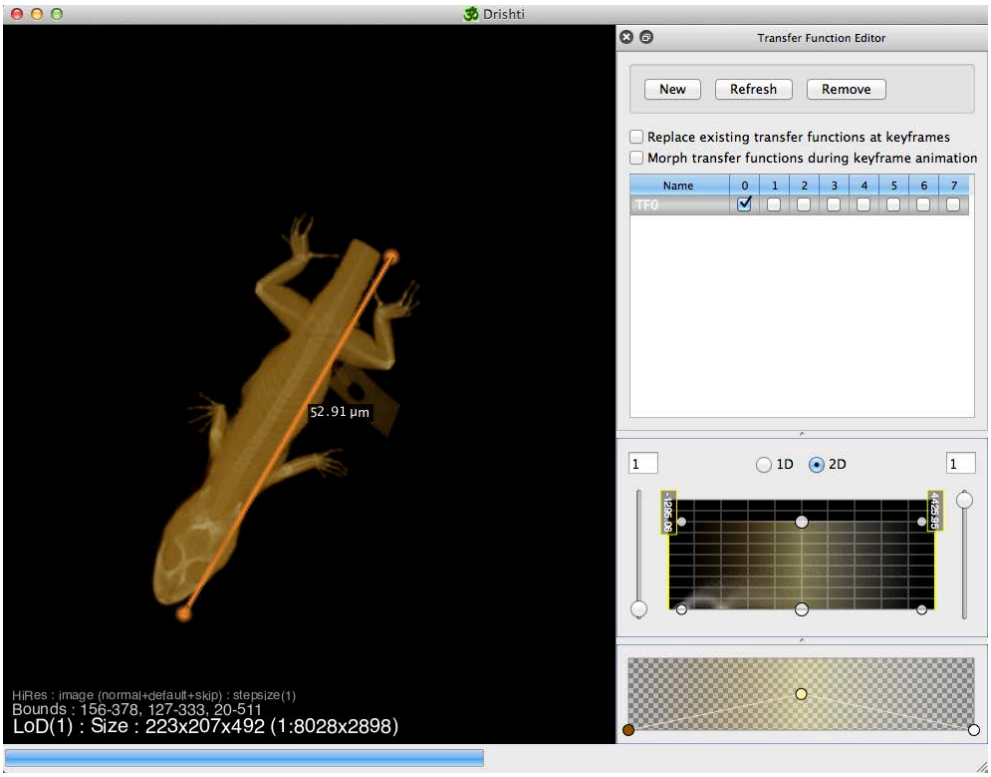

3. To place a permanent measurement on the object, hover over the path until it glows brightly then tap the spacebar. The Path Parameters window will then pop up. Select “display path length” to display the length of the path you have drawn. You can also change the colour of the measurement text and the distance it is from the image in this window, as you may observe. The units the measurements are taken in are the ones you specified when you first saved the image from DrishtiImport.

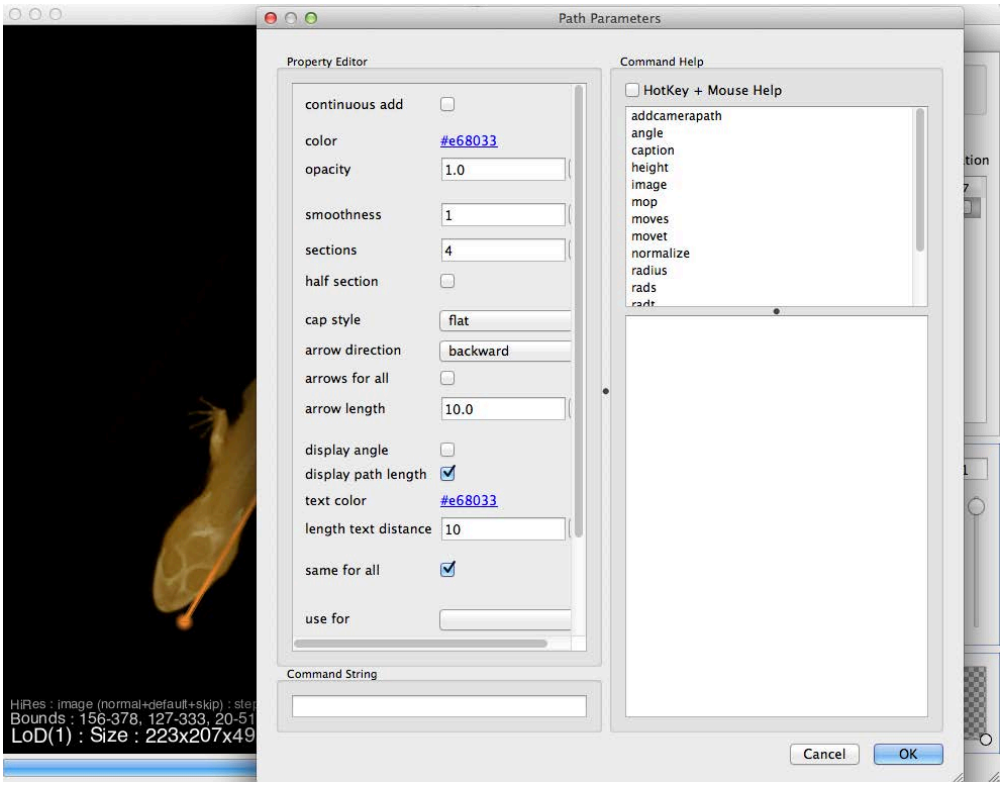

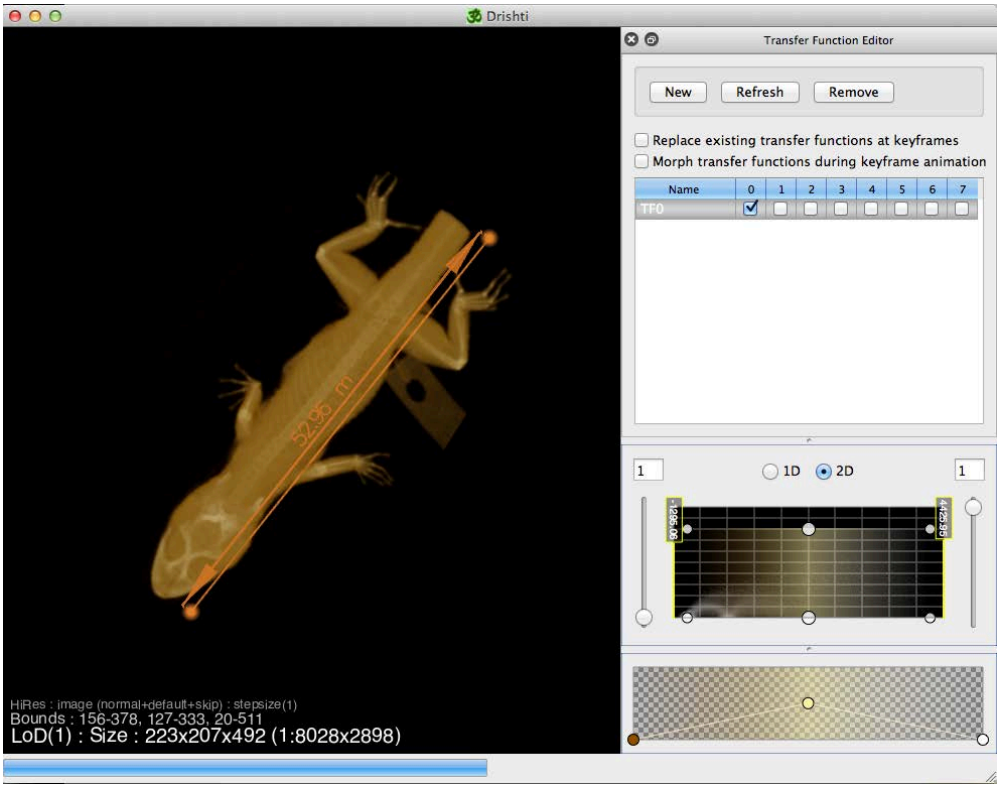

4. From here, you can pretty much make as many straight measurements as you want on your image, although it is best to rotate the image around as you are making the path to make sure that the points are even and positioned correctly.

Taking a Length Measurement (Curved Line)

Sometimes the structures that you want to measure on your image are curved or unable to be measured accurately with a simple straight line. In order to obtain measurements of structures such as these, additional points must be added, along with the use of the "path smoothness" option.

1. Similar to the straight-line method, add points to the image using the Shift key and the left mouse button, after ensuring that the image is appropriately positioned to allow maximum accuracy of the measurement. This time, add as many points as you deem fit for gaining an accurate measurement of the length of the structure.

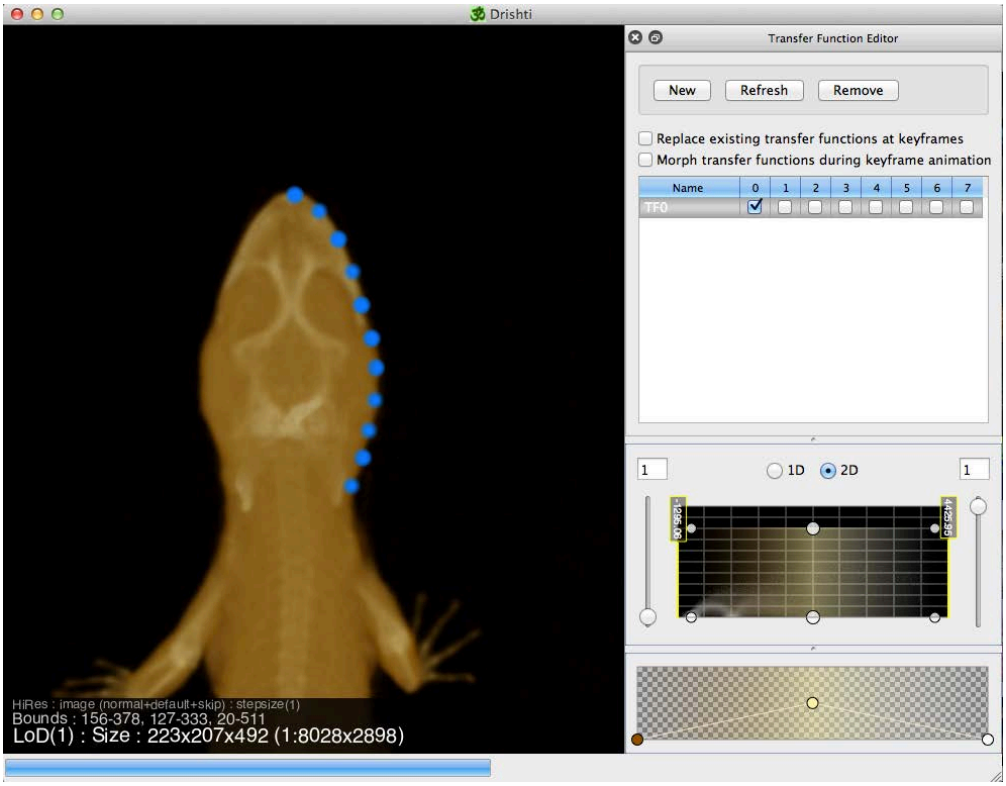

2. Now, just as you did before, add a path to connect all of the points together.

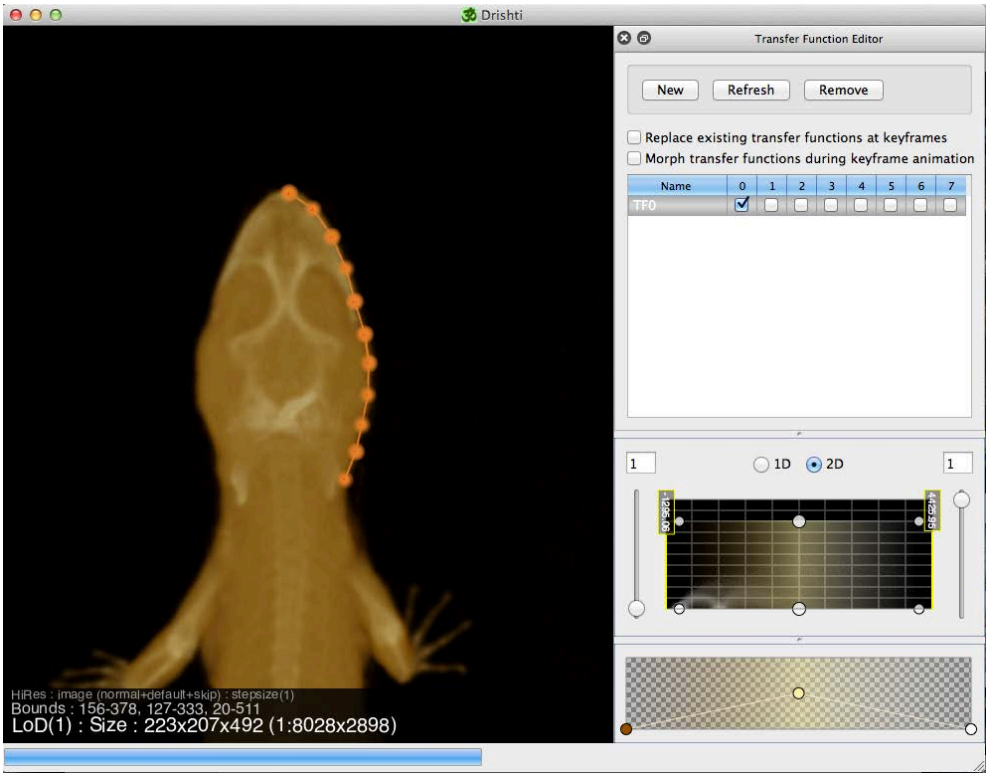

3. Now that you've added your path, you can get a more accurate measurement of the structure by shifting points that are not in the right place (by clicking and dragging).

4. Additionally, you can use path smoothness. Hover over the path and tap the spacebar to bring up the Path Parameters window. Increase the path smoothness value as much as desired to obtain a path that most closely fits the structure you are trying to measure. Select the “display path length” box and your measurement should appear on the structure.

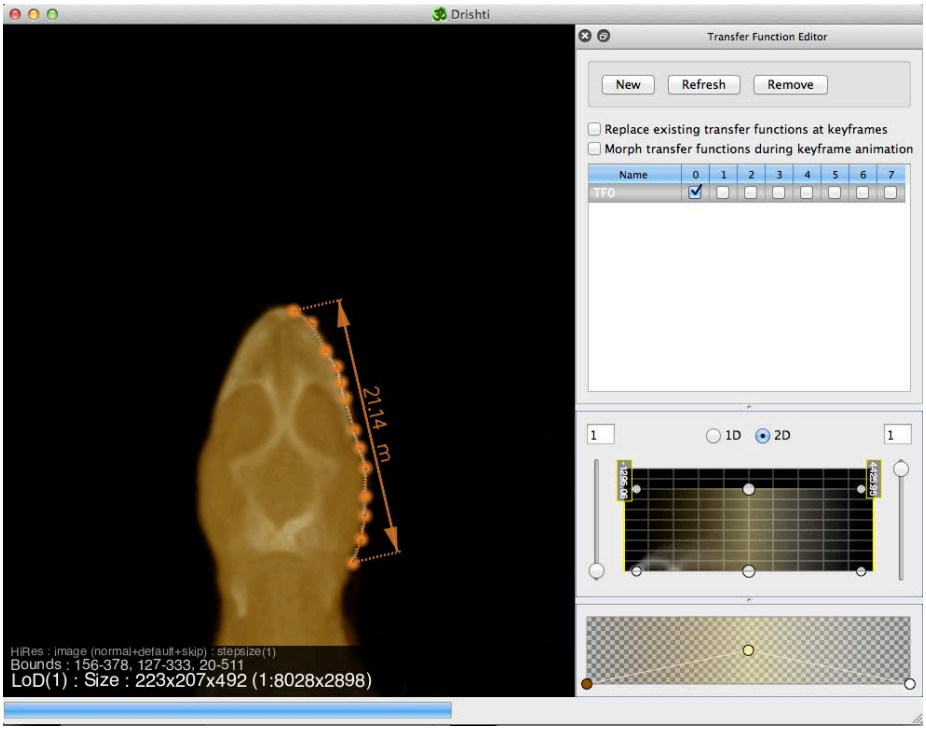

Taking an Angle Measurement

1. First, add points and draw a path between the points:

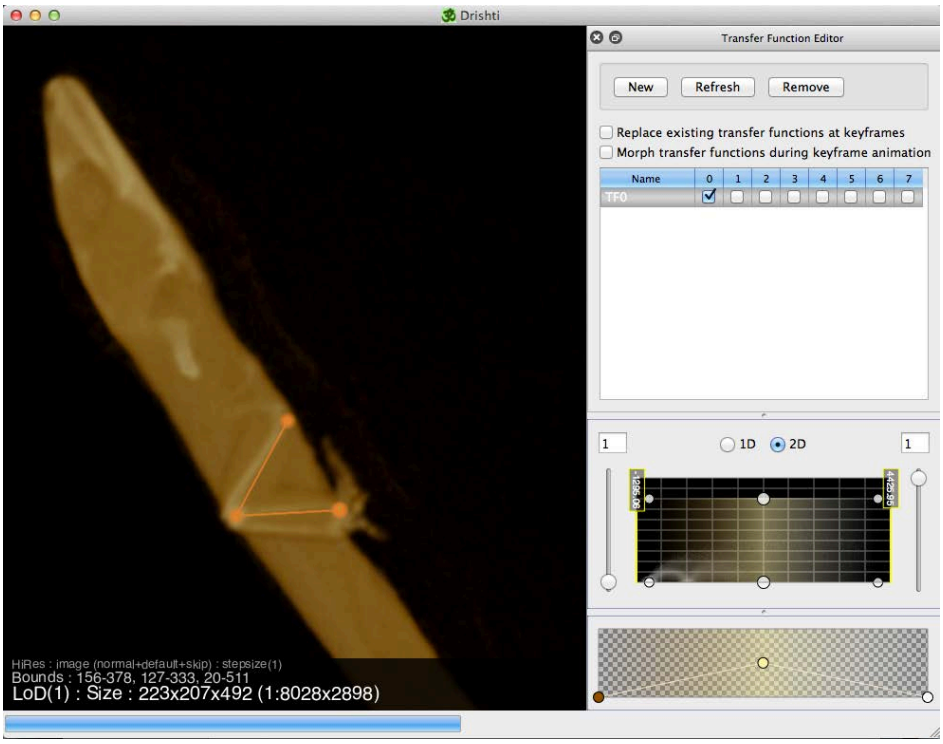

2. Then hover over the path until it lights up and press spacebar to bring up the Path Parameters window. Select “display angle”. The angle of the path will then be displayed.

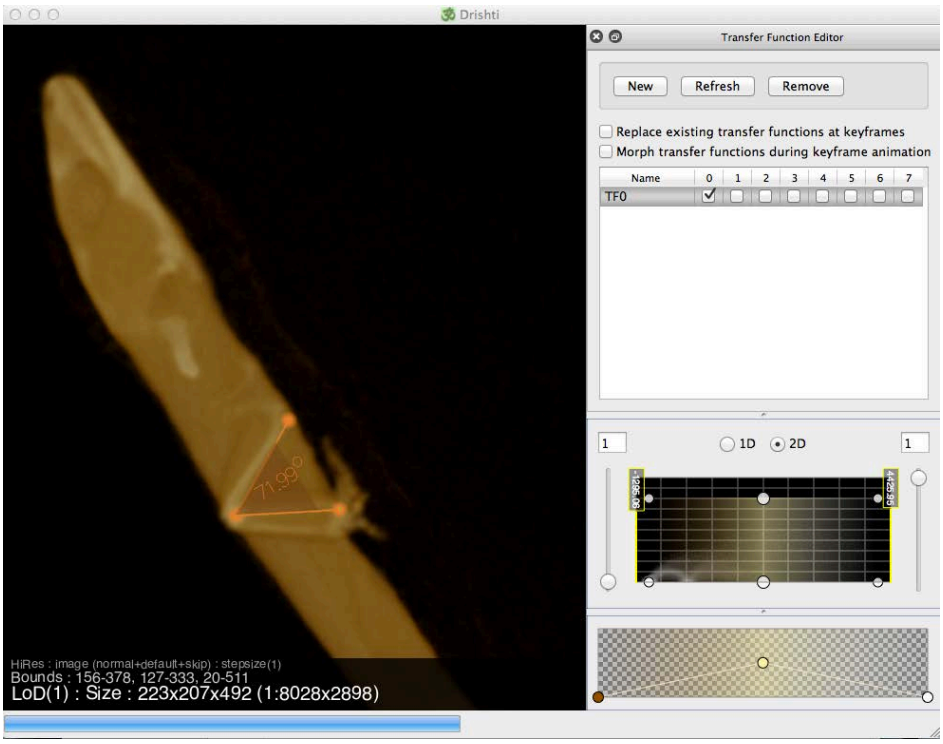

Now you should feel confident enough to take length and angle measurements from your own images!

Additional Help

All of the information I have detailed in this short manual has been collected from both the Drishti website and from the many help videos created and shared by the developer of Drishti, Ajay Limaye, of The Australian National University. All credit regarding the information detailed in this manual should be attributed to him, as all I have done is collated his information into a brief summary. Below is the list of references I used to compile this manual and links to helpful pages and videos.

Main Drishti Website (contains help pages, Drishti downloads:

http://sf.anu.edu.au/Vizlab/drishti/

Google Code Site for Drishti (can download latest version 2.4 from here):

http://code.google.com/p/drishti-2/

GitBut Site for Downloading the Latest Drishti Source Code:

https://github.com/AjayLimaye/drishti

Online Help Page for Drishti (contains help documents for each of the Drishti programs and a list of links to help videos for carrying out the range of operations Drishti offers):

http://sf.anu.edu.au/Vizlab/drishti/help.shtml

Link to Basic Help Videos (created by Ajay Limaye and contains all of the videos I used to create this manual):

http://www.youtube.com/watch?v=3nmwOpjn54w&list=PLC511E6135E402898

Drishti & ANU Logos: Sourced from Google Images

Front Cover & Manual Pictures: Sourced from own dataset using screenshot function

Note: The images in this manual were generated from the version 2.3.2 of Drishti; the latest version, Drishti 2.4, is now available and is recommended for use over the 2.3.2 version (instructions in this manual still apply to the latest version).

I trust that this manual was helpful, and if you have any questions about anything mentioned within it feel free to contact me on email:

bailey.lovett@hotmail.co.nz