FOVIA

High Definition Volume Rendering Manual for *XStream® HDVR®*

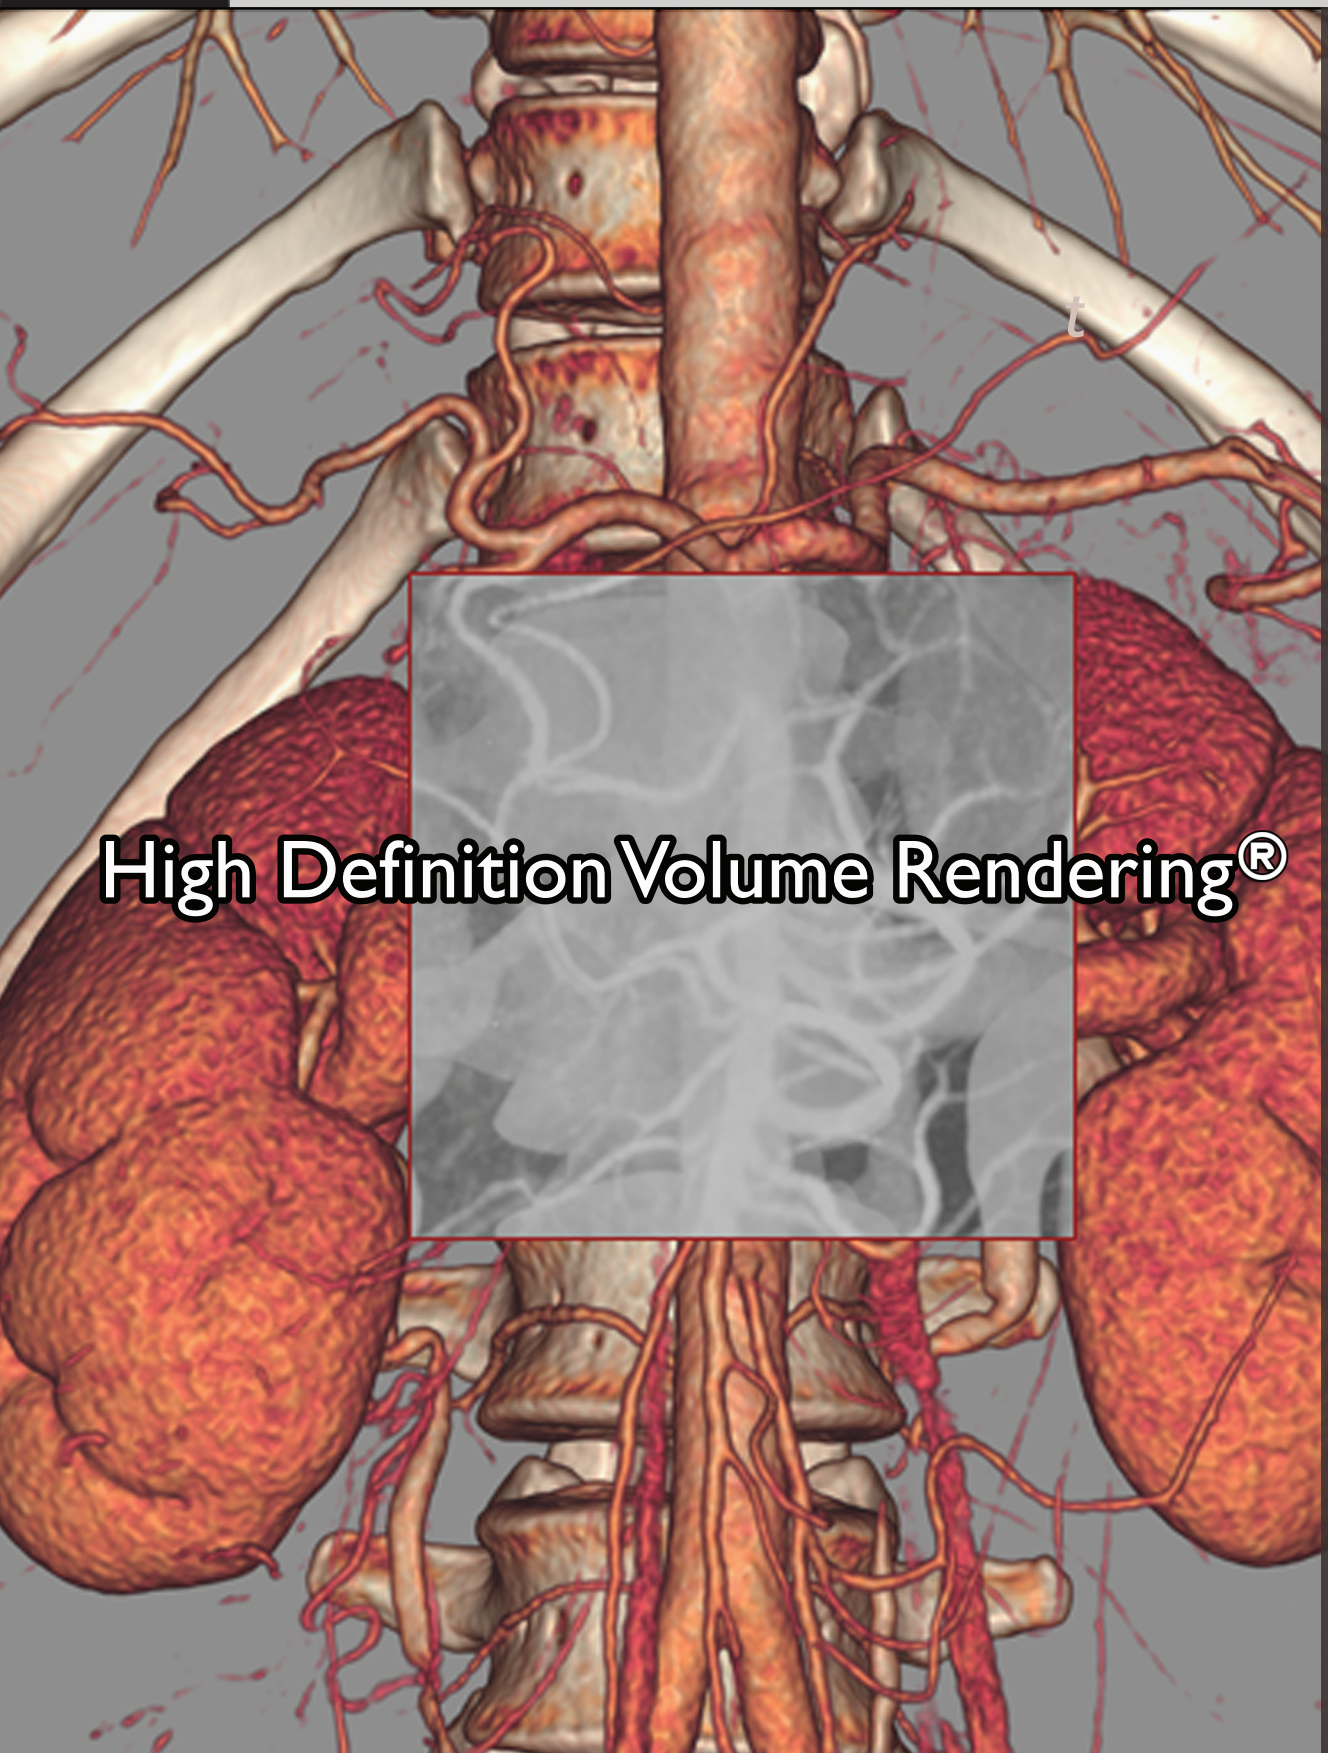

Inside...

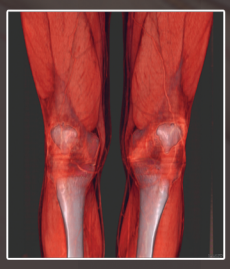

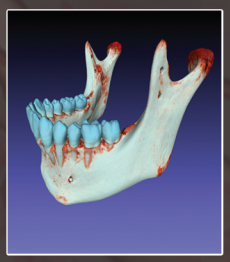

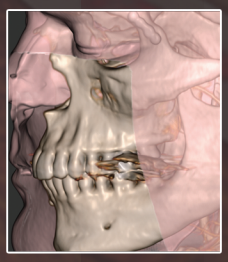

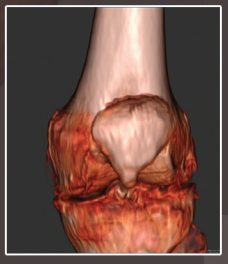

TABLE OF CONTENTS

MANUAL

FOViA

TABLE OF CONTENTS

MANUAL

FOViA

Fovia is strictly an OEM (original equipment manufacturer) supplier whose solutions are extremely easy to integrate. Our High Definition Volume Rendering® software development kit (SDK) gives your customers the quality and performance they demand, and gives you the ease-of-integration and quick time to market you require.

1.1 GETTING STARTED - LAUNCHING THE SOFTWARE

LAUNCH *XStream® HDVR®* USING THE SERVER/CLIENT PATH

Check JRE Version

The Fovia Client-Server Application requires Java to be installed. If you are running a 64 bit operating system, you should ensure that you are running the 64 bit version of the JRE. If you are running a 32 bit operating system, you should ensure that you are running the 32 bit version of the JRE. Attempting to run the 64 bit version of the Workstation on a 32 bit version of the JRE will result in an error, even if you are running a 64 bit operating system.

Check to ensure that you have the correct Java Runtime Environment installed. Launch the application-specific command window as follows:

WINDOWS: From the Start menu search, enter "run command.com" to open a command window.
MAC: From Spotlight, search and launch the "Terminal" application to open a command window.
LINUX/ UBUNTU: Open a terminal window, or switch to console mode.

In the command window, enter "java -version" and confirm that build 1.6 or later is installed. If this results in an error or you are running an older version, please install Java 1.6 or later. NOTE: If the version information includes "Java HotSpot™ 64-Bit Server", this indicates that you are running 64-bit JRE. If the version information includes "Java HotSpot™ 32-Bit Server", this indicates that you are running 32-bit JRE. Close the command window.

Download the Software

1. Create a folder on your desktop and name it Fovia. The ZIP archive file will be saved to this folder.
2. Download the software version appropriate for your operating system and move it into the Fovia folder created above.
3. Use your favorite unzipping utility to decompress the ZIP file and save it in the Fovia folder.
4. When you open up the contents of the decompressed file, you will find two folders: "Server Application" and "Client-Server Workstation Application."
5. Your Fovia representative should have provided a license file called "hdrclic.dat". This file will activate the SDK. Copy the "hdrclic.dat" file into the "Server Application\32-bit" and "Server Application\64-bit" folders.

Launch Application Server

Navigate to the application folder containing server and client. Then use the operating system specific instructions below to launch the workstation application. Be sure to leave open any boxes that pop while launching the application. They assist in the visualization process. Instructions:

WINDOWS

- Navigate to the "Server Application" folder.
- Depending on the Java version specified above, select "Start Server 64-bit.bat" or "Start Server 32-bit.bat".
- Navigate to "Client-Server Workstation Application" folder.
- Depending on the Java version specified above, select "Start Client 64-bit.bat" or "Start Client 32-bit.bat".

LINUX

- Navigate to the "Server Application" folder.
- Depending on the Java version specified above, select "server/server32" folder or "server/server64" and execute "run.sh".
- Navigate to "Client-Server Workstation Application" folder.
- Depending on the Java version specified above, select "client/client32" folder or "client/client64" and execute "run.sh".

MAC*

- Navigate to the "Server Application" folder.
- Depending on the Java version specified above, select "Start Server 64-bit.command" or "Start Server 32-bit.command".
- Navigate to "Client-Server Workstation Application" folder.
- Depending on the Java version specified above, select "Start Client 64-bit.command" or "Start Client 32-bit.command".

*Mac only: When running both client and server on OSX, it is advisable to install and run the script below to increase the allowed shared memory buffer size. The SDK will use shared memory buffers to transfer data if running the client and server on the same machine and a clean install of OSX, the operating system will disallow allocation of shared memory buffers of the size necessary for this feature. Other operating systems allow adequate allocation.

Download and run the Fovia Shared Memory Startup Script (sent from a Fovian representative).

Configuration Settings

A configuration dialog will be displayed (Figure1.1). Select "Local Rendering" if running both client and server on the same system. If running this application in a client-server environment, launch the server application on a separate system and specify its IP address in this dialog. In the latter case, the JPEG option should be checked so the server will send compressed JPEG images rather than much larger uncompressed images.

Click the "OK" button.

The application should now be up and running.

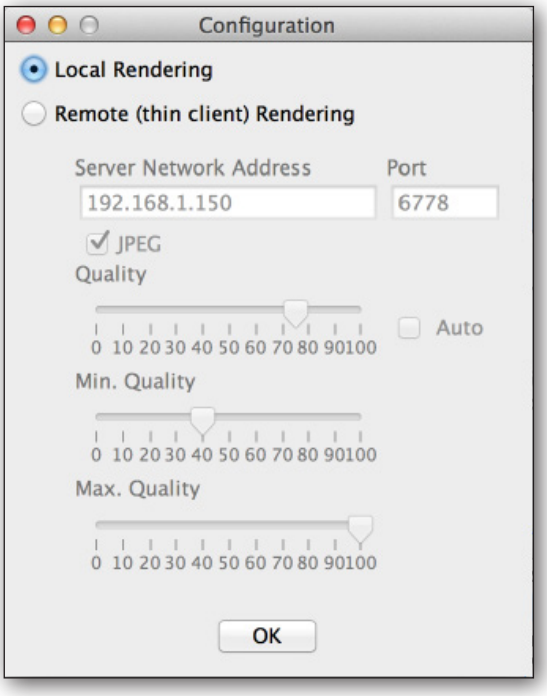

FIGURE 1.1: Configuration dialog

1.2 GETTING STARTED - DATA TYPES

Fovia offers a DICOM data loader that supports any cross-sectional modality in the medical or industrial space. In addition, *XStream® HDVR®* supports SEG-Y for geo-science, and the raw data loader can support just about any type of data. The server API within the SDK provides an interface that allows customers to implement their own loader. When implemented, Fovia's software uses callbacks to your code to parse the slice data. Data types supported include: CT, MRI, CBCT, OCT, MicroCT, PET, SPECT, .Float, SEG-Y, .RAX, and more!

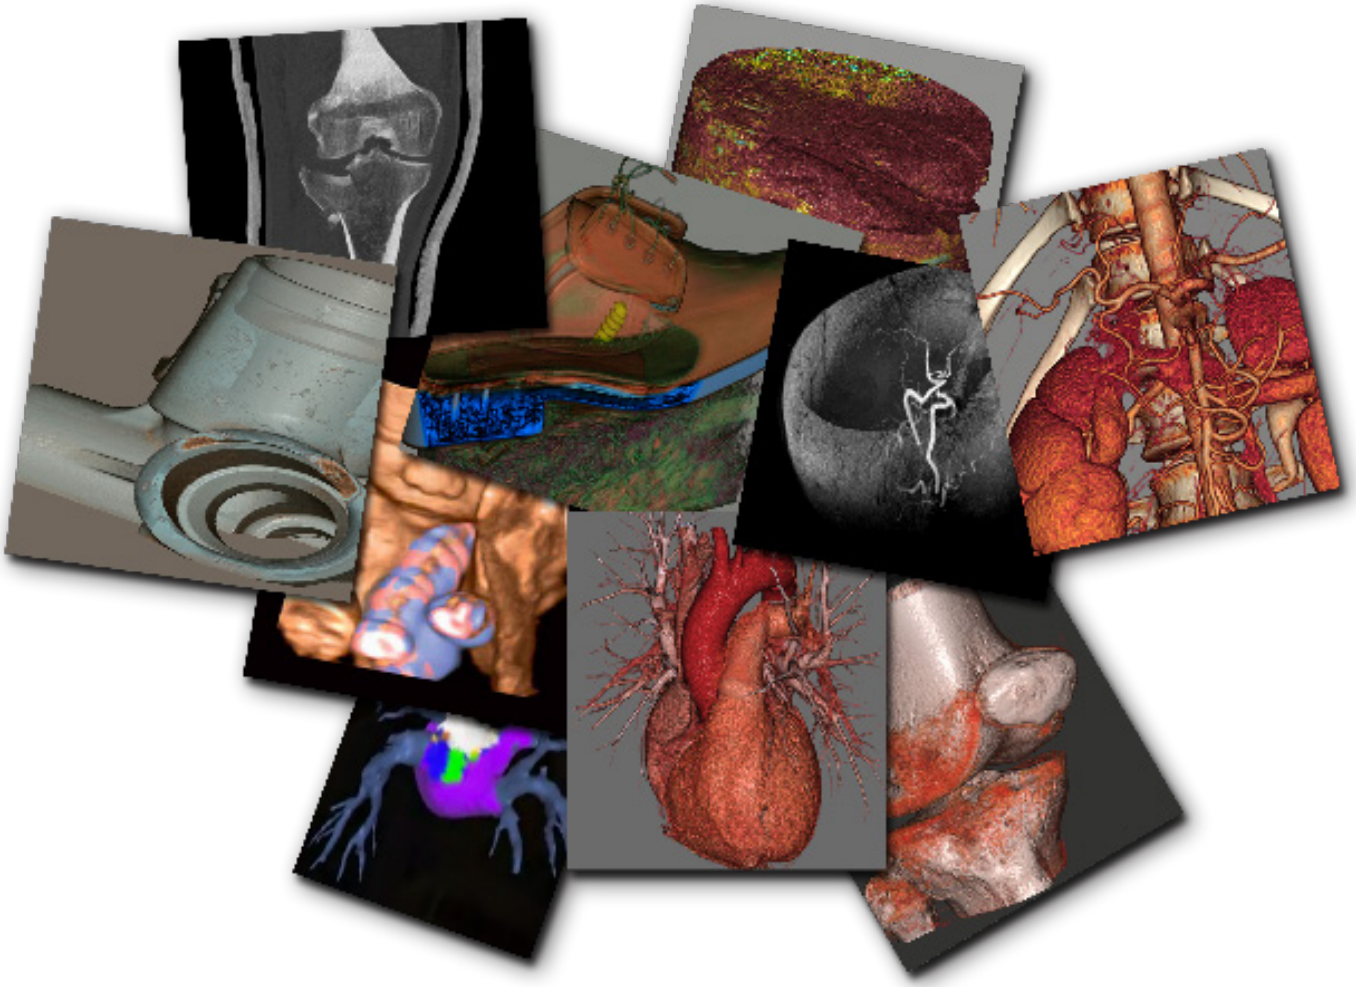

FIGURE 1.2: Data Types

1.3 GETTING STARTED - DATA LOADING

1.3a HOW TO LOAD NON-DICOM DATASETS

To open non-DICOM files using the raw loader, click on the File Drop-down Menu, choose "Open Raw/Segy Data," select the folder containing desired files and click "Open" (Figure 1.3). Alternatively, click on the "Open" button in the Dataset Tab, choose "Open Raw/Segy Data," select the folder containing desired files, and select "Open" (Figure 1.4). A dialogue box will become available (Figure 1.5).

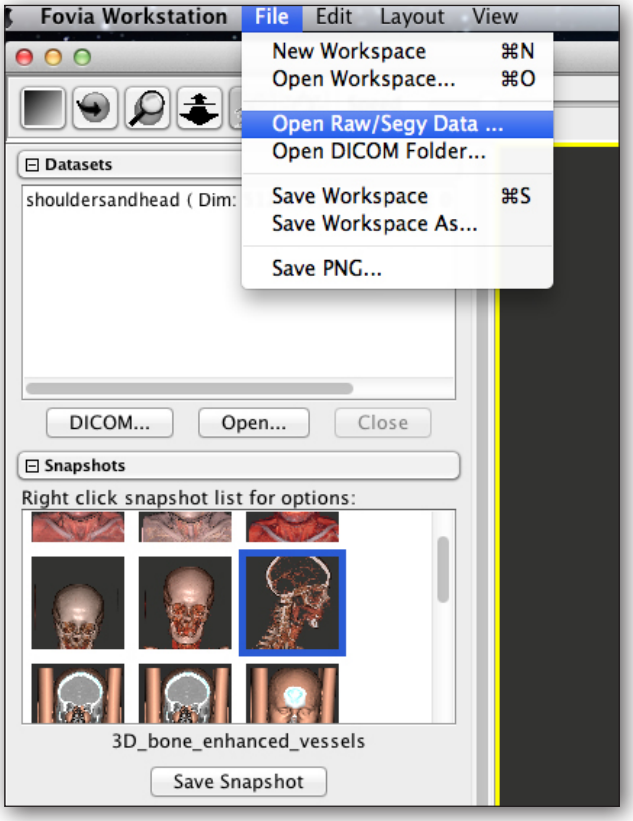

FIGURE 1.3: Non-DICOM data loading from file menu

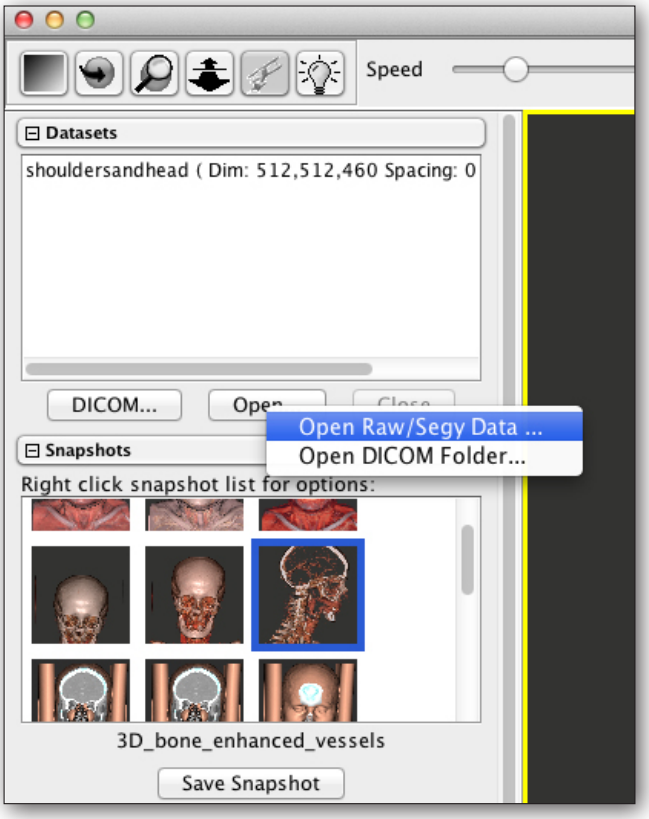

FIGURE 1.4: Non-DICOM data loading from Dataset Tab

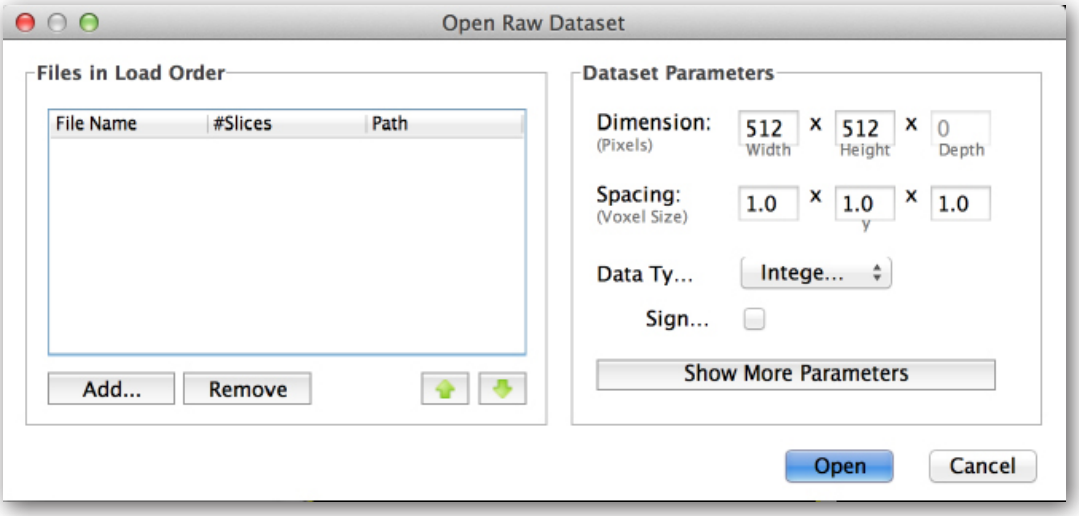

FIGURE 1.5: Raw Data dialogue box

Once the dialogue box is open select "Add" to add your non-DICOM files. Then select the type of file you wish to load from the drop down menu: SEGY, RAW, or "ALL" for all other files.

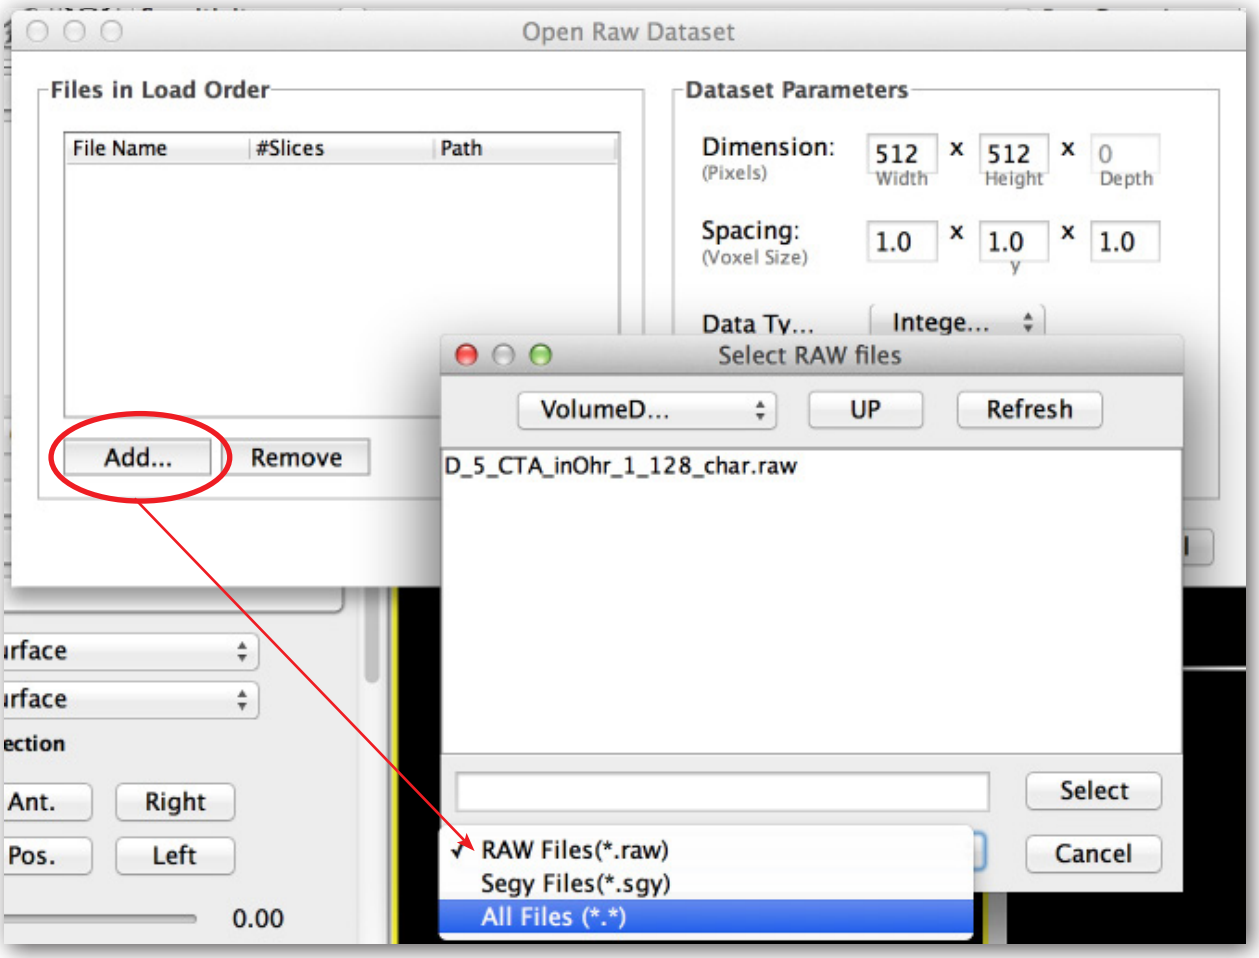

FIGURE 1.6: Raw Data dialogue box

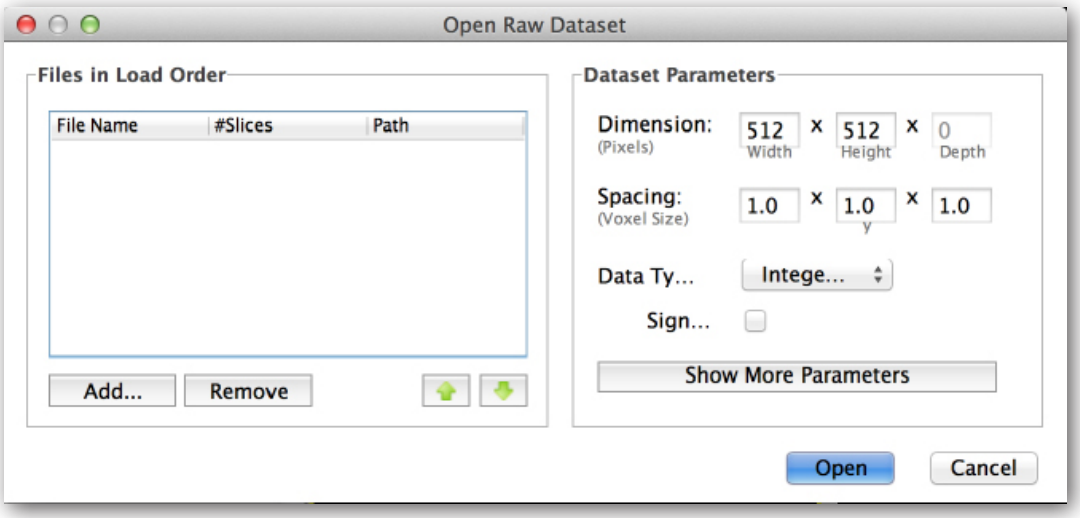

FIGURE 1.7: Raw Data parameters

The default dialogue box provides a few parameters: dimension, spacing, etc. To view all of the parameters, select the "Show More Parameters" button (Figure 1.8). Select "Hide More Parameters" to collapse the extended dialogue box (Figure 1.9).

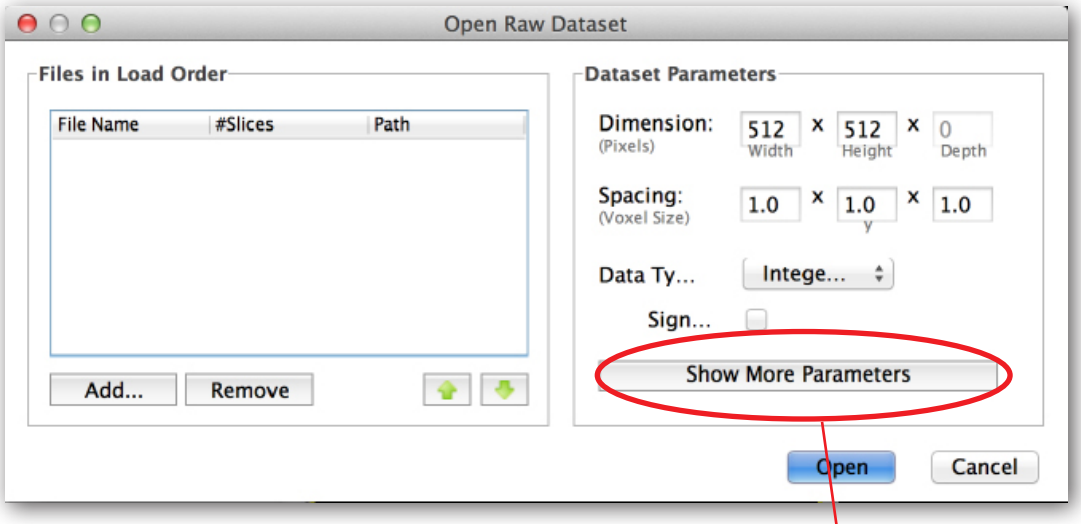

FIGURE 1.8: Raw Data expand parameters

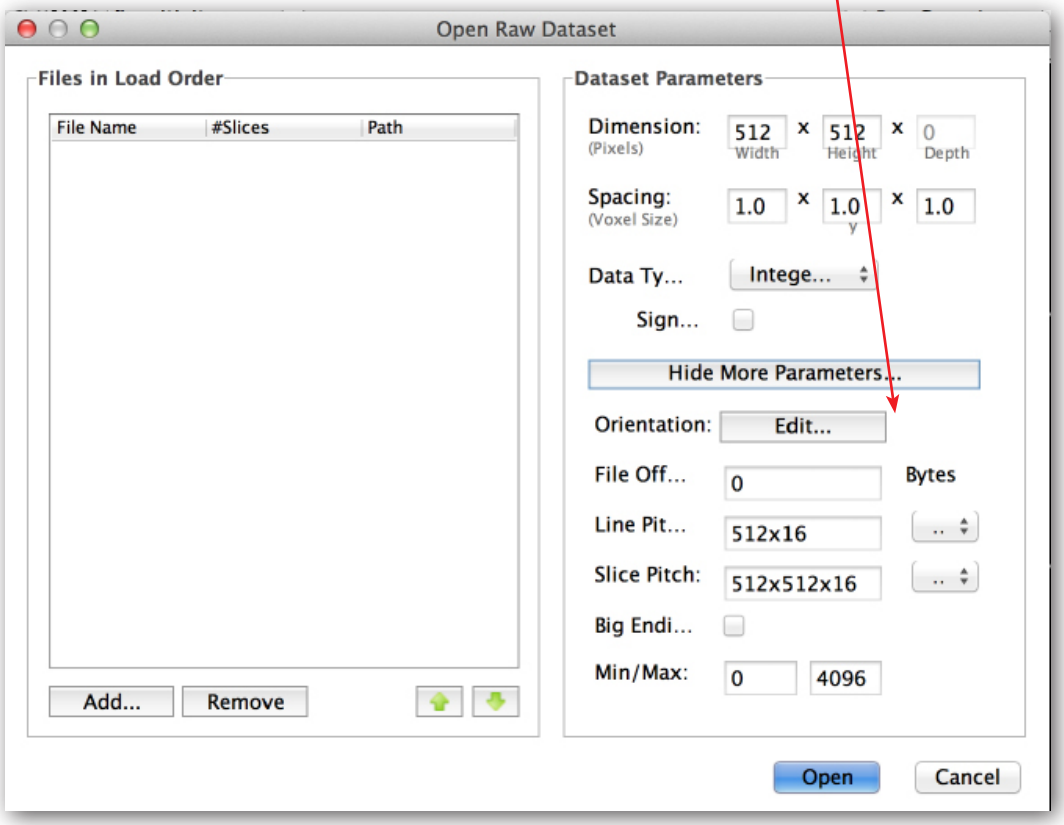

FIGURE 1.9: Raw Data expanded parameters

1.3b HOW TO LOAD DICOM DATASETS

There are several ways to load volumetric DICOM data. To load files from the File Dropdown Menu, click on File, choose "Open DICOM Folder," select the folder containing desired DICOM files and click "Open" (Figure 1.10). To load files using the Dataset Tab, click on the DICOM button, select the folder containing the DICOM files, and select "Open." The DICOM button loads data into the DICOM library. Alternatively, click on the "Open" button in the Dataset Tab, choose "Open DICOM Data", select desired files and click "Open".

To load multiple datasets, simply continue to load datasets as described above. Once all desired datasets have been loaded, each can be assigned its own viewing window. To do this, click on a viewing window to make it active (you will see a yellow outline around the active window), then double-click on the datset in the Dataset Tab window. Learn more about this in the Layout Chapter.

Additionally, Workspace Files can be used to open datasets or recall your work from a previous session. Utilize workspace files to save dataset paths, snapshots and corresponding segmentation.

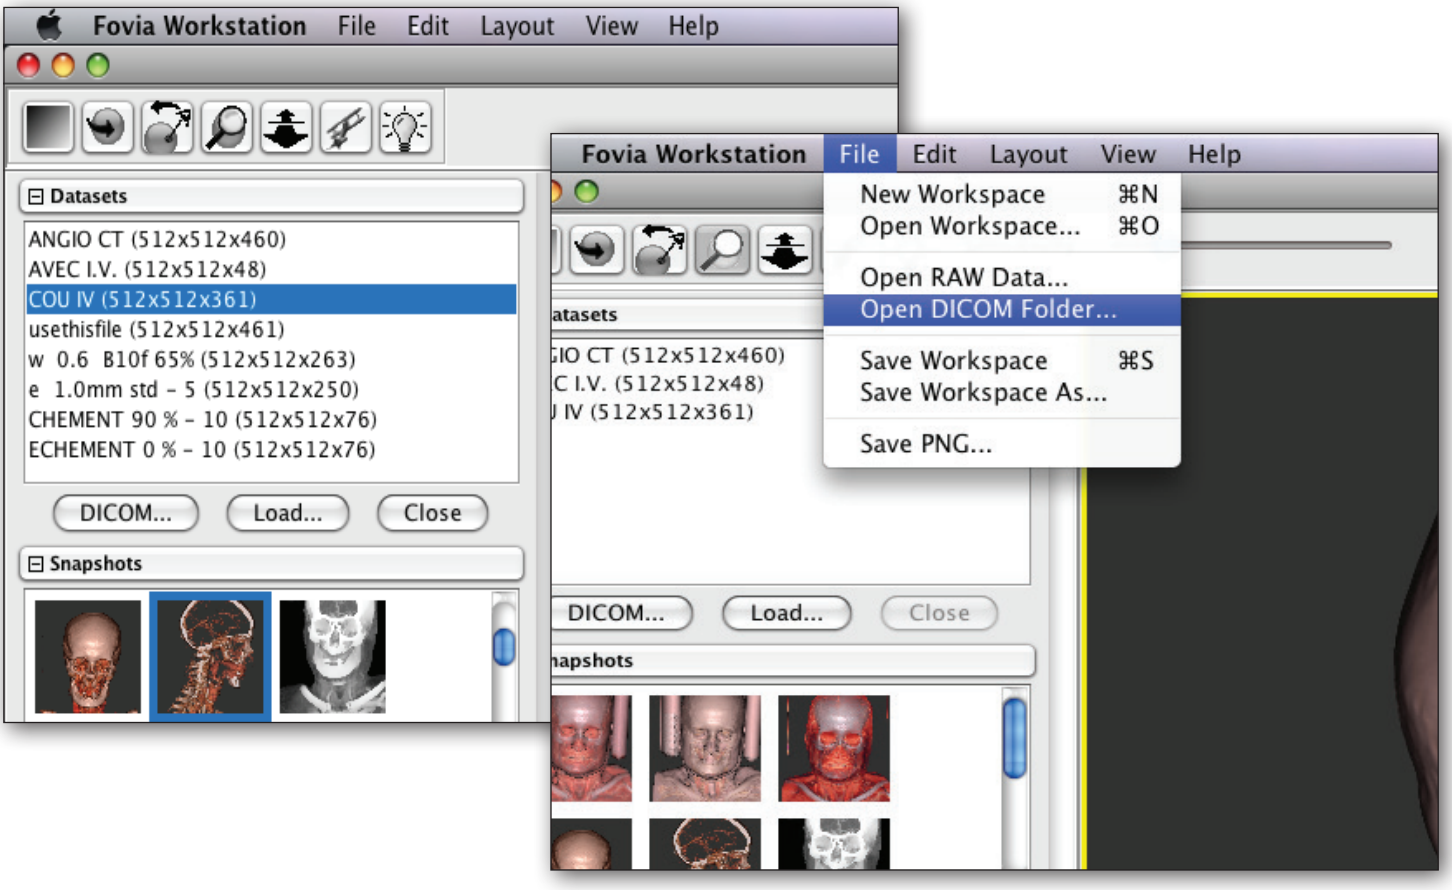

FIGURE 1.10: Opening DICOM files

1.4 GETTING STARTED - RENDERING MODES

To access the rendering types, right-click in the viewing window. The options are listed under "Rendering Types" (Figure 1.11). The three categories include Parallel Brute Force, Parallel Adaptive and Perspective Adaptive. The available options included in each category and their uses are discussed below..

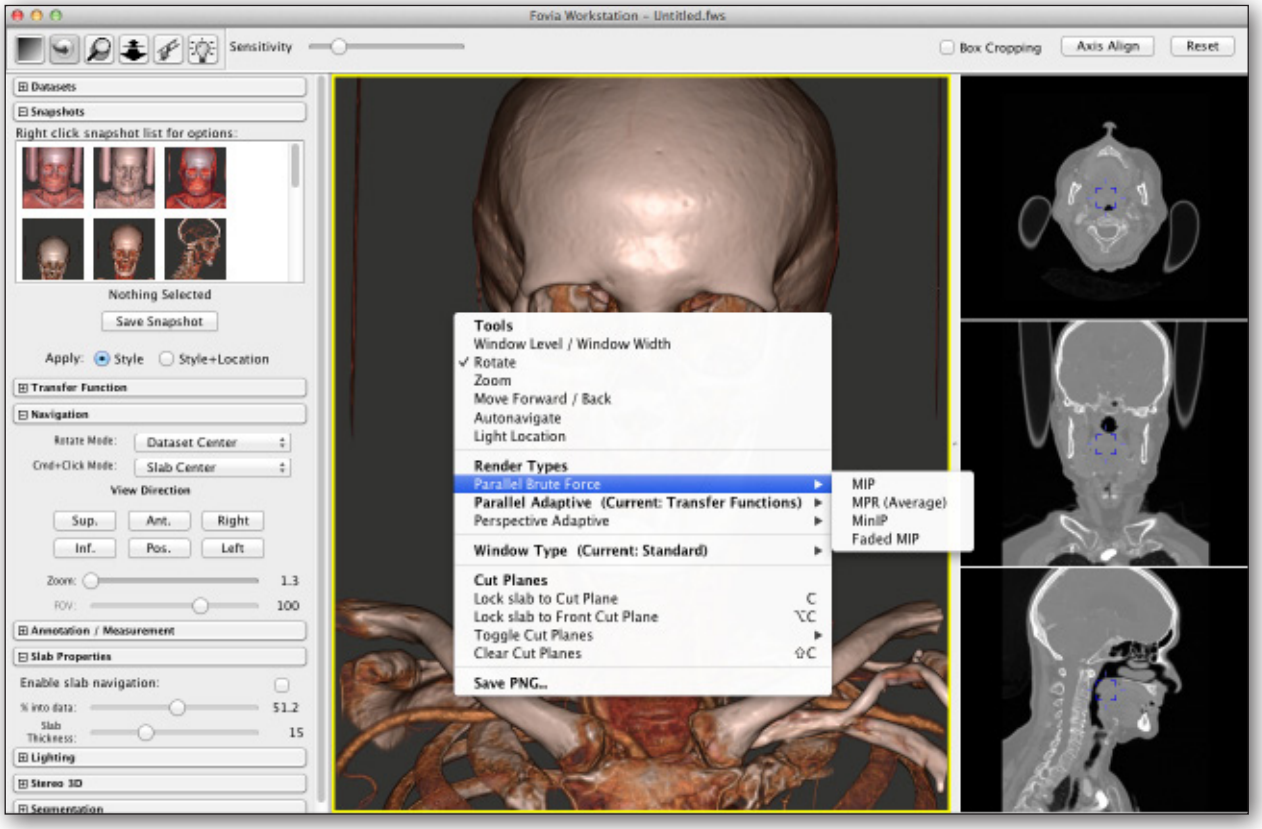

FIGURE 1.11: Rendering Types

RENDERING TYPES

Parallel Adaptive: Parallel Adaptive Mode is a rendering type unique to Fovia. Adaptive rendering modes render faster than brute force modes. Render modes using parallel projection utilize rays that are all parallel to each other. Therefore, all rendered objects appear to be the same distance from the current viewpoint regardless of their actual distance. Parallel rendering modes are useful for taking measurement points across the dataset, because there is a one-to-one relationship between pixels and length. That is, each pixel on the screen has the same absolute size relative to the dataset. Parallel rendering is also usefull if the viewpoint is outside the object and you would like to rotate the object around in front of you.

Perspective Adaptive: Render modes that use perspective projection utilize rays that are emitted from the center of the current viewpoint, and more closely simulate how the human visual systems works. Objects in a perspective projection appear smaller the farther they are from the camera, whereas objects appear larger the closer they are to the camera. Perspective rendering allows the viewer to explore complex regions close to or inside the dataset, and is useful for fly-through or auto-navigation of structures like the colon, vessels, airway and nerve canals.

Parallel Brute Force: In brute force renderings, the dataset is evaluated at sampling points along the ray path through the volume. The density of sampling points depends upon the zoom level for the dataset. At high zoom levels there will be at least one sample point per voxel, resulting in a highly detailed image. The lower the zoom level, the farther apart the sample points will be. If the zoom level is low enough, some small anatomical structures may be missed because they are smaller than the sample point separation. Therefore, brute force modes are ideal for viewing at high zoom levels. Brute force modes are more computationally intensive because they sample more individual data points than adaptive rendering modes do.

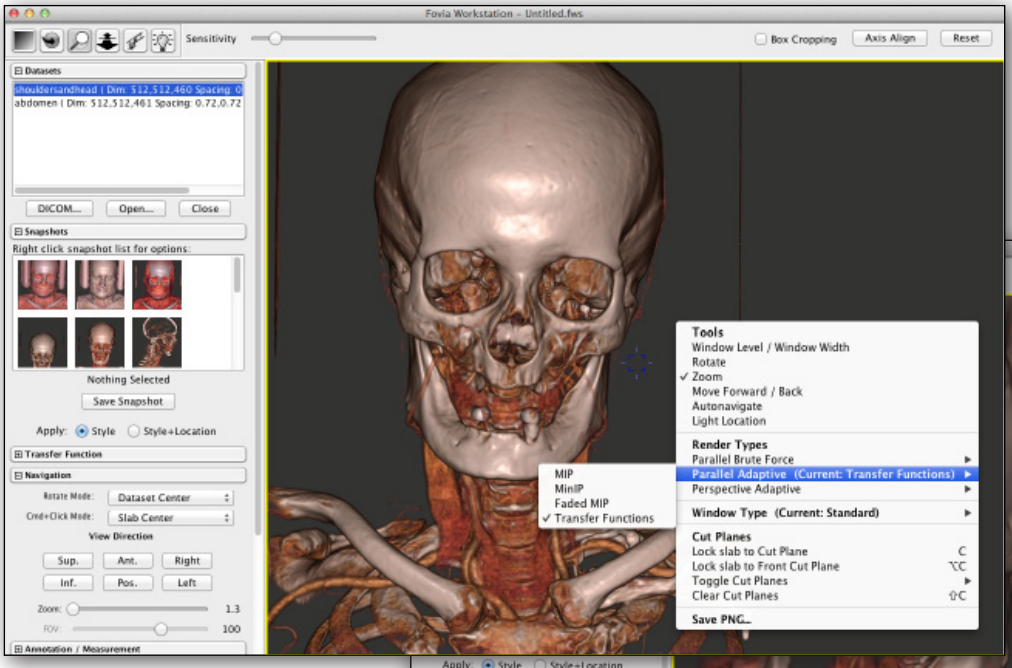

FIGURE 1.12: Parallel Adaptive Rendering Type (unique to Fovia)

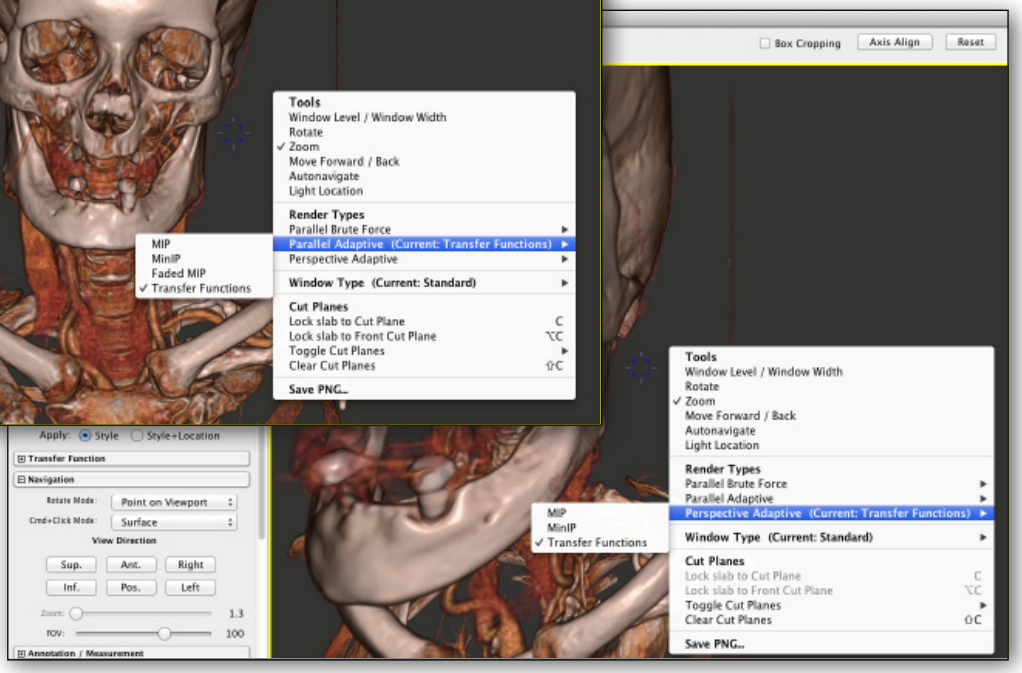

FIGURE 1.13: Perspective Adaptive Rendering Type

RENDERING MODES

Within each render type, there are rendering modes from which to choose. They include:

MIP (Maximum Intensity Projection): A MIP rendering displays only the voxel with the highest Hounsfield value (or other dataset unit) along the ray's direction of travel. The voxel is visualized as a grayscale pixel ranging from black (minimum density) to white (maximum density). This mode is useful for identifying regions of high density, regardless of the structure's position within the dataset. This mode is best suited to visualize bone, contrast-enhanced or other materials with a high attenuation value.

MinIP (Minimum Intensity Projection): A MinIP rendering is a form of MPR rendering in which only the voxels of lowest Hounsfield value (or other dataset unit) are visualized. A MinIP image is calculated in the same manner as an MPR rendering, but the visualized image shows only voxels of minimum attenuation. MinIP images are useful for identifying regions of low relative density.

MPR (Multiplanar Reconstruction - MPR-Average, only available in brute force rendering): An MPR rendering is a grayscale image consisting of rendered pixels that represent the average Hounsfield value (or other dataset unit) of the voxels through which the ray passes. An MPR rendering is generated by casting rays through the volume, accumulating the value of the voxels, calculating the average value of the voxels and projecting the image onto a 2D plane. Brighter pixels correspond to material of high density, and darker pixels correspond to material of low density. MPR-Average modes work best on datasets with slab thicknesses of < 50 mm. As slab thickness increases, the image will blur more due to the averaging that takes place along the ray's traversal through the dataset. If you would like to view individual MPR slices from the orginal scan data, follow these steps:

1. Set the rendering mode to Parallel Brute Force MIP or MPR.

2. Under the Navigation tab, select the Superior view button.

3. Under the Slab Properties tab, set the slab thickness to 0.

4. Individual scan slices can be viewed by adjusting the '% Into' slider, or with the mouse using the 'Forward/Back' mouse control feature.

Faded MIP (Available only in parallel rendering): Faded MIP is unique to Fovia. A Faded MIP rendering is similar to MIP. However, in Faded MIP voxels become darker as they increase in distance from the camera. In a conventional MIP rendering, the voxel with the highest Hounsfield value (or other dataset unit) along the ray path is always displayed. This can cause foreground structures of lower density to become obscured by background structures of high density. Use of Faded MIP rendering allows better visualization of foreground structures in these cases. (Figure 1.14)

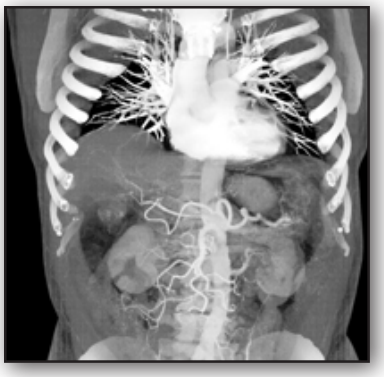

FIGURE 1.14: Faded MIP allows better viewing of foreground structures than traditional MIP

Transfer Function (Available only in adaptive rendering): This rendering mode is an industry standard for working with transfer functions and volume rendering. Transfer Function based rendering modes produce colored 3D volume renderings with translucent, transparent and opaque effects. The transfer function maps a voxel's Hounsfield value (or other dataset unit) to three components: color, lighting and opacity values. The *XStream® HDVR®* SDK supports multiple transfer functions, allowing groups of voxels to be assigned to different transfer functions. These voxel groups can be visualized separately through the use of segmentation techniques.

LAYOUT

The Layout Menu enables the user to organize the workspace in a manner appropriate for their workflow. Different volumes can be loaded into separate windows and windows can be linked as well.

2.1 LAYOUT - FILE MENU

The Layout Drop-Down Menu, located on the top left side of the user interface (Figure 2.1), has five options for the user. The first three options are 2x2, 1x2, and 2x1. "Sidebar" is the fourth option, which provides a submenu that allows the user to control the location of the sidebar (Figure 2.2). The fifth option is "Floating Render Window" that you can move independently while continuing to work in the main environment. The floating render window can be re-sized by selecting the bottom right corner with your mouse and dragging. Double click on any window to display full screen view or double click to return to other views.

FIGURE 2.1: Layout Drop-Down Menu

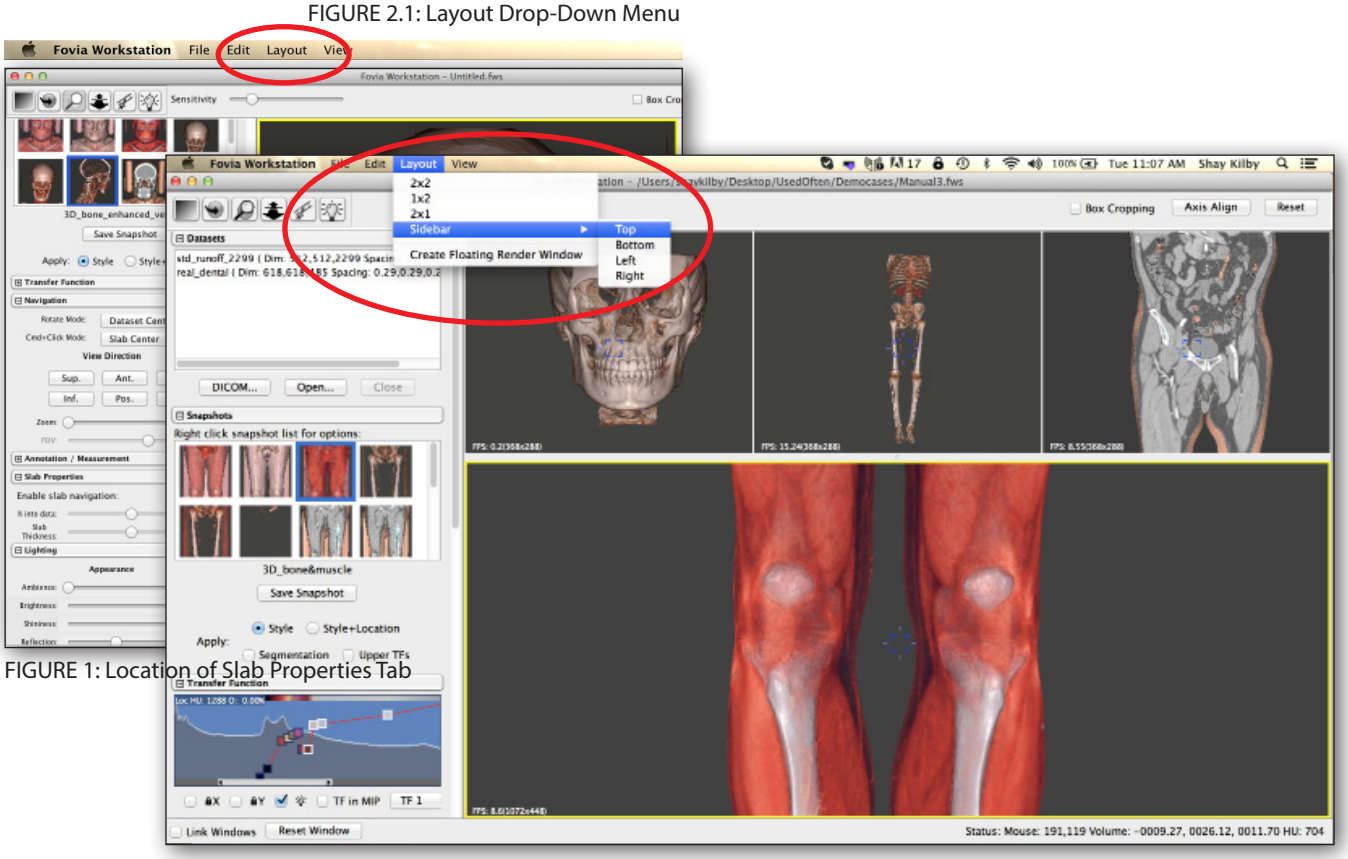

FIGURE 1: Location of Slab Properties Tab

FIGURE 2.2: "Sidebar" allows you to create a "bar" with 3 windows on the top, bottom, left, or right

You can load different datasets into different windows. Simply select the window and then double-click on the desired dataset in the dataset tab window.

You can easily move data sets around in each window using the preset orientations in the Navigation Tab. And, you can apply different snapshots to different windows independently.

2.2 LAYOUT - LINKING WINDOWS

You can link windows in the application easily by selecting the "Link Windows" checkbox at the bottom left of your screen (Figure 2.3).

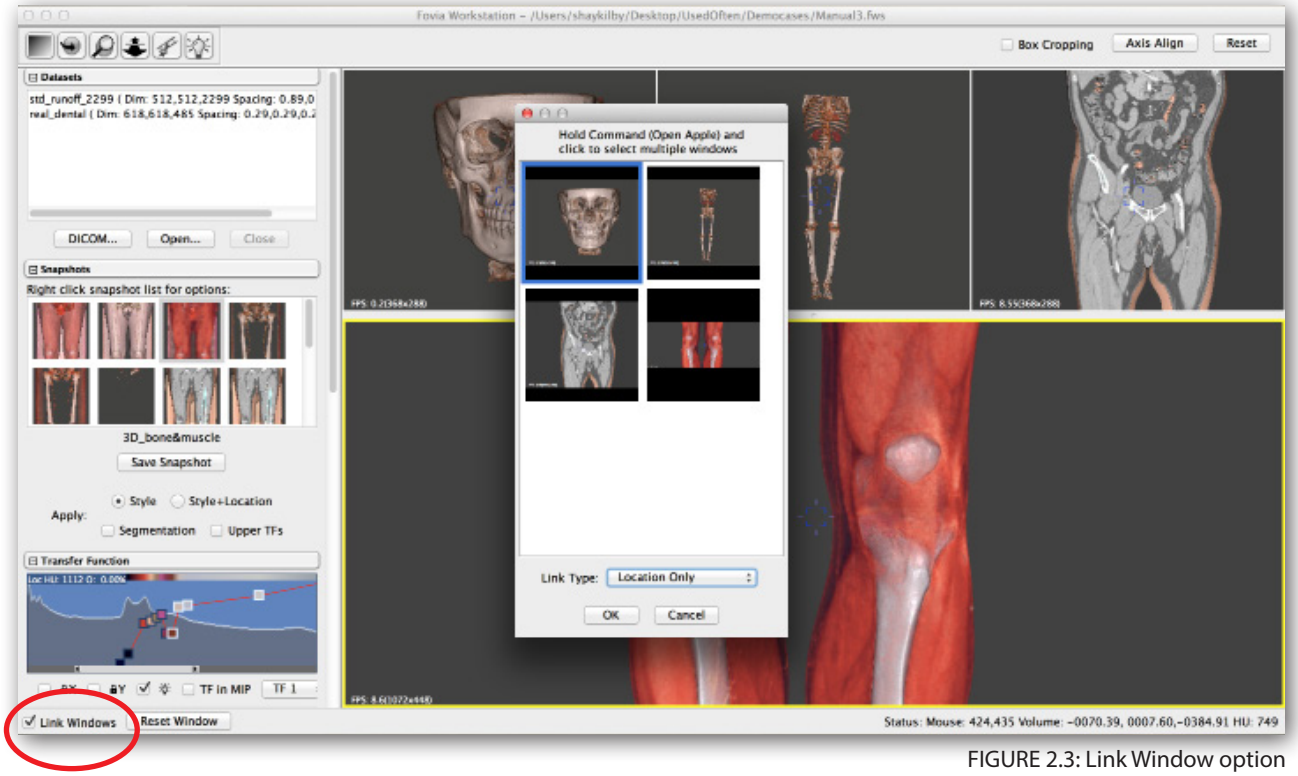

FIGURE 2.3: Link Window option

Using linking windows you can:

- Link by Location Only or by Location & Rotation (Figure 2.4)
- Select windows to link

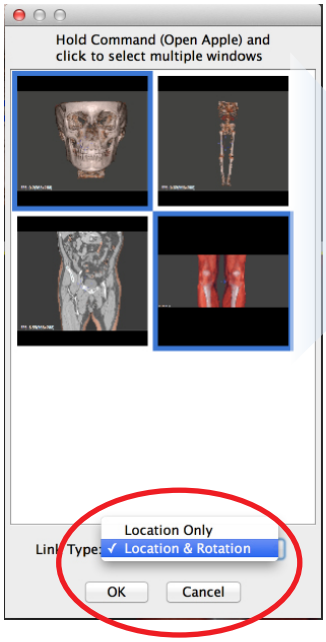

FIGURE 2.4: Link by Location & Rotation

Select any combination of windows to link by holding down the command/control key and left-clicking on the desired windows to link

NAVIGATION

The primary navigation options in the Fovia user interface include the pan, zoom and rotat tools, slab control and fly-through control. Options for selecting these tools as well as detailed descriptions of their respective functionalities are discussed below.

3.1 NAVIGATION - TOOLS & SHORTCUTS

There are multiple ways to select navigation functions. The first method is to select the icons that represent desired functions from the upper left control panel (Figure 3.1). Note that the sensitivity slider increases the sensitivity of the selected tool.

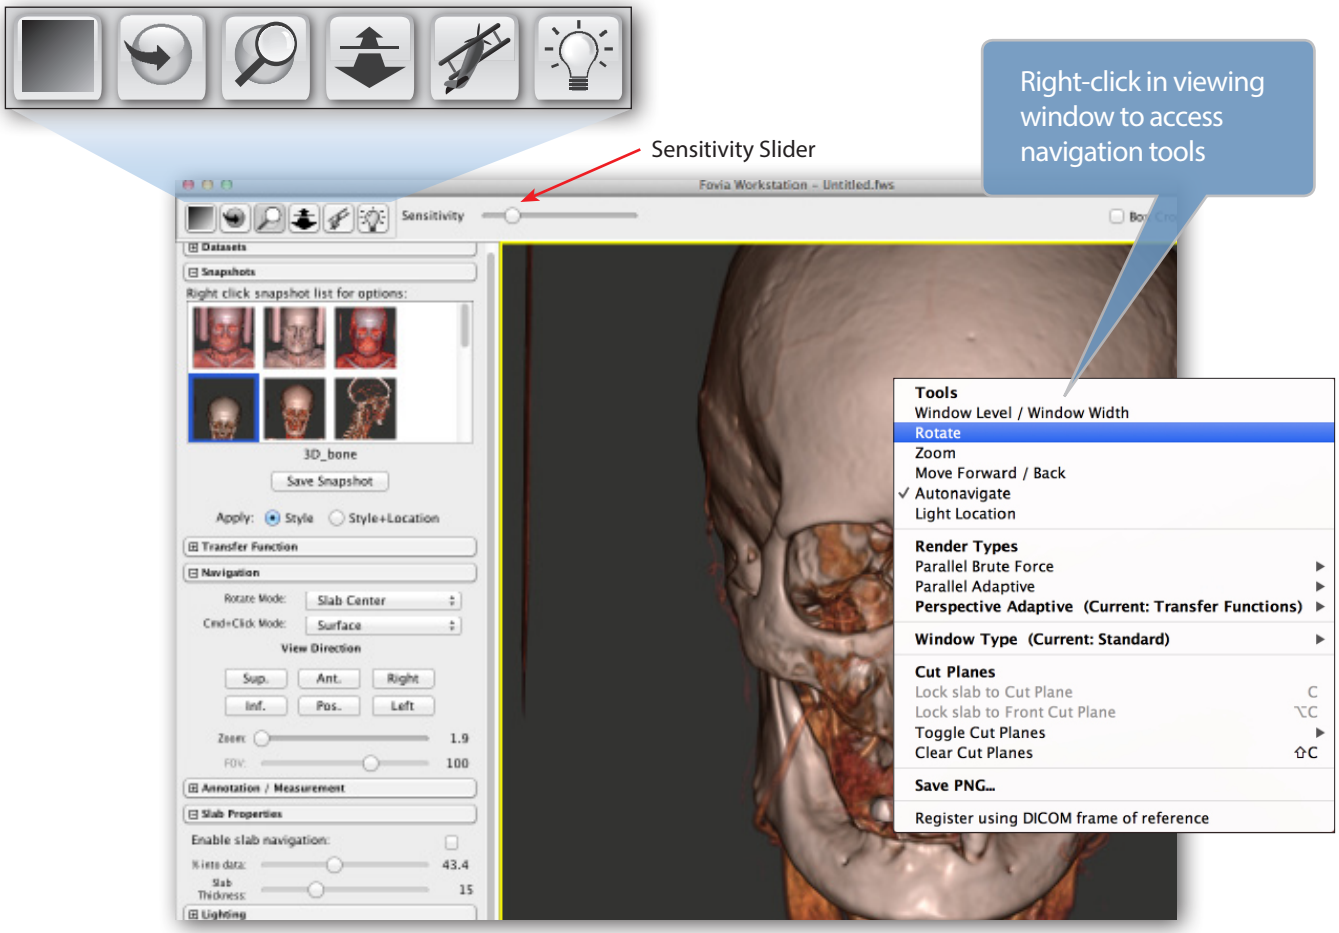

FIGURE 3.1: Location of Navigation Tools

The user can also right-click in the viewing window to access the tools. This opens a menu list from which to select a desired function. Alternatively, the simplest method to navigate volume data is by using the mouse buttons and scroll, in combination with keyboard shortcuts. The navigation tools and their respective shortcuts are listed below.

Window Level/Width Tool
Shortcut: press right and left mouse buttons simultanerously.
Adjusts the entire transfer function on the horizontal axis. If right and left buttons are held down simultaneously (referred to as tracking) and moved up and down in the viewing window, all existing TF (Transfer Function) ranges along the horizontal axis will move, without changing the ranges themselves, thus adjusting what part of the dataset is visualized. Tracking right and left changes the expanse of all TF ranges, allowing one to visualize a smaller or larger "part" of the dataset.

Rotate Tool
Shortcut: Control (Command in Mac) + left-click OR Middle + left mousebutton
Rotates the dataset in the viewing window.

Zoom Tool
Shortcut: Alt + left-click OR
Middle + right mousebutton pressed simultaneously
Zooms in and out of the viewing window.

Pan Slab
Shortcut: Middle mouseclick and drag
Pans through the slabs of the dataset. Slab navigation must be enabled to use this tool. This tool is discussed in greater detail in the Slab Properties chapter.

Fly-Through:
Automatically follows a path through a portion of the dataset via auto-navigation. This tool is discussed in greater detail later in this chapter.

Lighting
Adjusts the location of the light source within the viewing window. This tool is discussed in greater detail in the Lighting Chapter of the manual.

Panning
Shortcut: Right-click and drag
Moves the entire object around the screen. Note that there is no tool in the toolbar for this function; it is accessed solely through its shortcut.

FreeHand Cut
Shortcut: Shift + left-click
Enables the user to cut around desired objects. This tool can also be accessed under the Segmentation Tab. More information on this tool will be provided in the Segmentation chapter. When accessing this tool, a pop-up appears. Click "okay".

3.2 NAVIGATION - NAVIGATION TAB

The Navigation Tab is located on the left-hand side of the user interface (Figure 3.2). This tab includes a number of functions, including Rotate Mode, Ctrl + Click Mode, View Direction, Zoom and FOV.

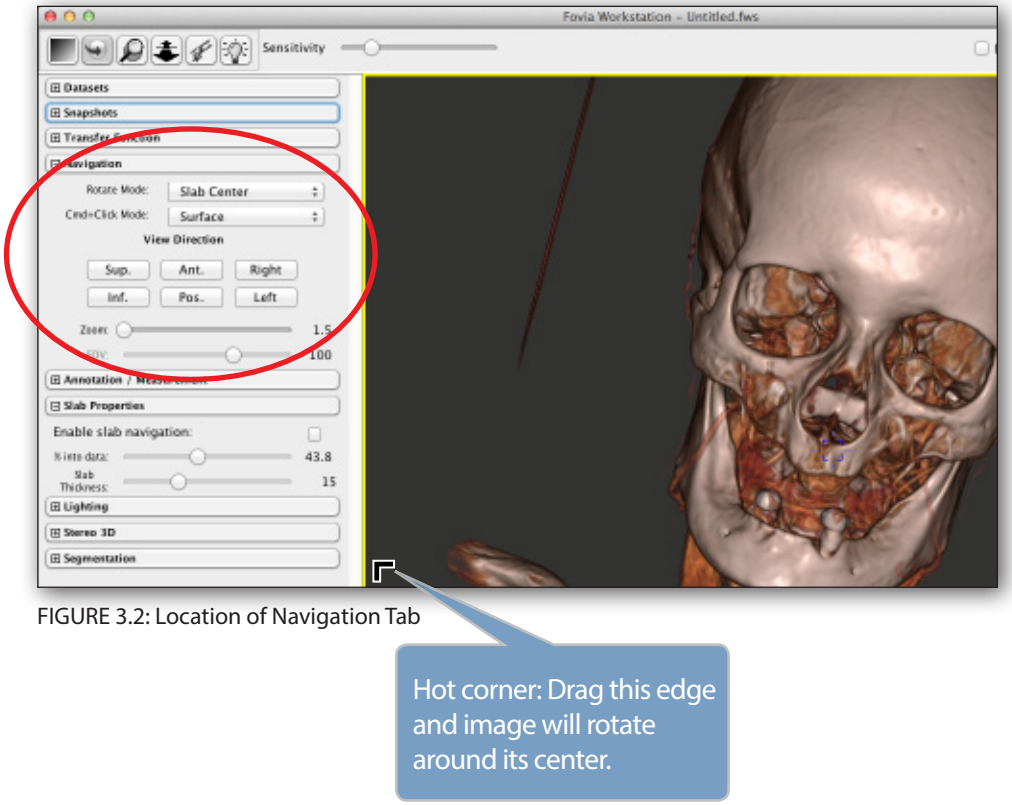

FIGURE 3.2: Location of Navigation Tab

The Rotate Mode determines the axis around which the object will rotate. The options include Surface, Slab Center, Point on Viewport (which the user determines with a mouseclick) and Dataset Center. The Ctrl + Click Mode is similar to the Rotate Mode in that it also determines the axis around which the object will rotate, *but only when accessing the rotation fuction through its shortcut*: Control (Command in Mac) + click. This allows the user to quickly toggle between two points of rotation if necessary. Another rotation method is to drag on the outer corner of the viewing window, which causes the flat/2D image to rotate around its center (rotation around the axis formed by the view direction).

The View Direction sets the view orientation of the dataset (Figure 3.5). Options include Superior (Sup) or Inferior (Inf); Anterior (Ant) or Posterior (Pos); Right or Left (Figures 3.3 and 3.4).

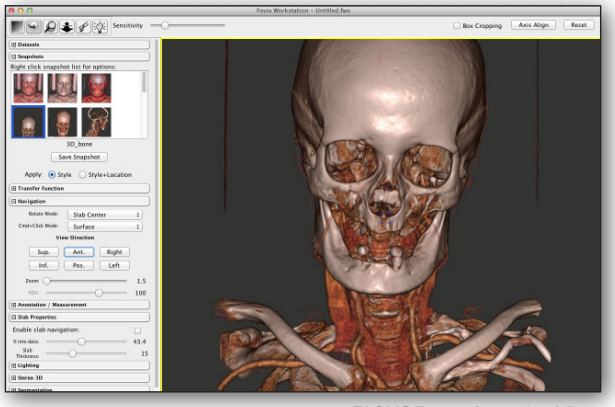

FIGURE 3.3: Anterior View

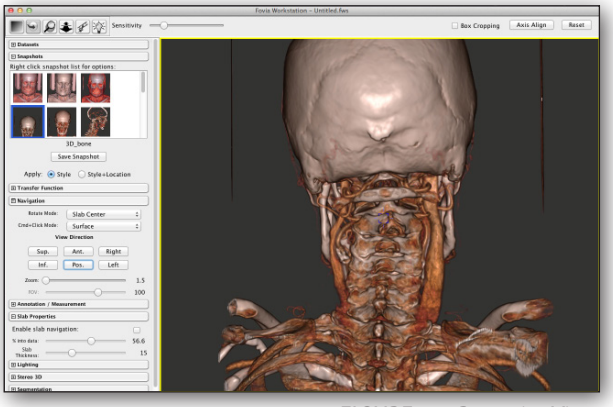

FIGURE 3.4: Posterior View

In addition to the Zoom Tool in the toolbar, the Navigation Tab allows the user to zoom via a slider (Figure 3.6). This function zooms in and out on a scale of 0.1 to 700. Note that the functionality works in Parallel Rendering Mode but not in Perspective Rendering Mode.

The FOV (Field of View) slider, located below the Zoom Slider in the Navigation Tab (Figure 3.7), determines the angle in degrees from the left side of the image to the right side, at middle height. For natural perspective viewing this control should be set at 60 degrees, which is the default. Note that the functionality works in perspective rendering mode but not parallel rendering mode: it will become active once the user switches to perspective rendering mode.

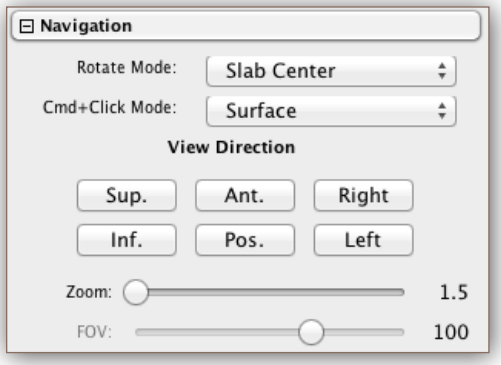

FIGURE 3.5: View Direction

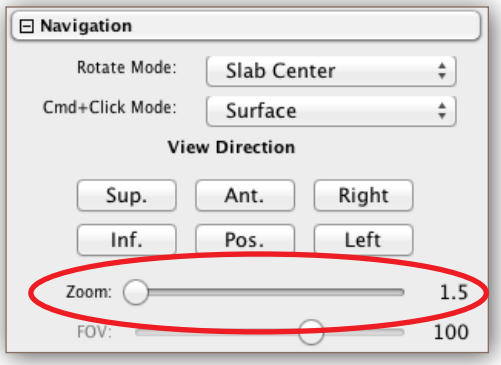

FIGURE 3.6: Zoom in the Navigation Tab

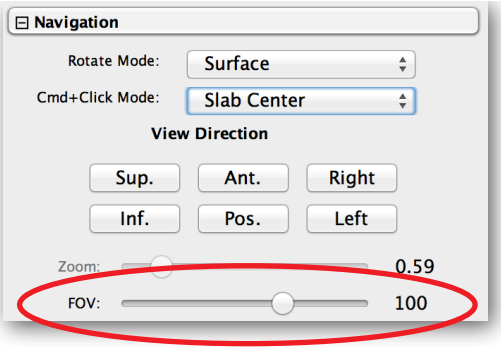

FIGURE 3.7: FOV in the Navigation Tab

3.3 NAVIGATION - FLY-THROUGH AUTO-NAVIGATION

Perspective Rendering allows for auto-navigating internal structures that have airways for virtual fly-through. In medical imaging, perspective visualization allows for endoluminal viewing of blood vessels and other structures such as colons and lungs. Select the icon in the main toolbar to enable auto-navigation mode. To change the Rendering Mode to Perspective for Fly-Through, left-click on the screen and select "Transfer Functions" listed under "Perspective Adaptive" from the dropdown menu (Figure 3.8). Hold down the left mouse-button in the viewing window to begin the fly-through. Path correction can be done at any time via manual navigation. It is important to position the visible blue cross bar (+) at the center of the desired vessel or structure (Figure 3.9). See the Creating Media Chapter (p. 62) for options to record fly-through actions.

FIGURE 3.8: Choosing Perspective Rendering Mode

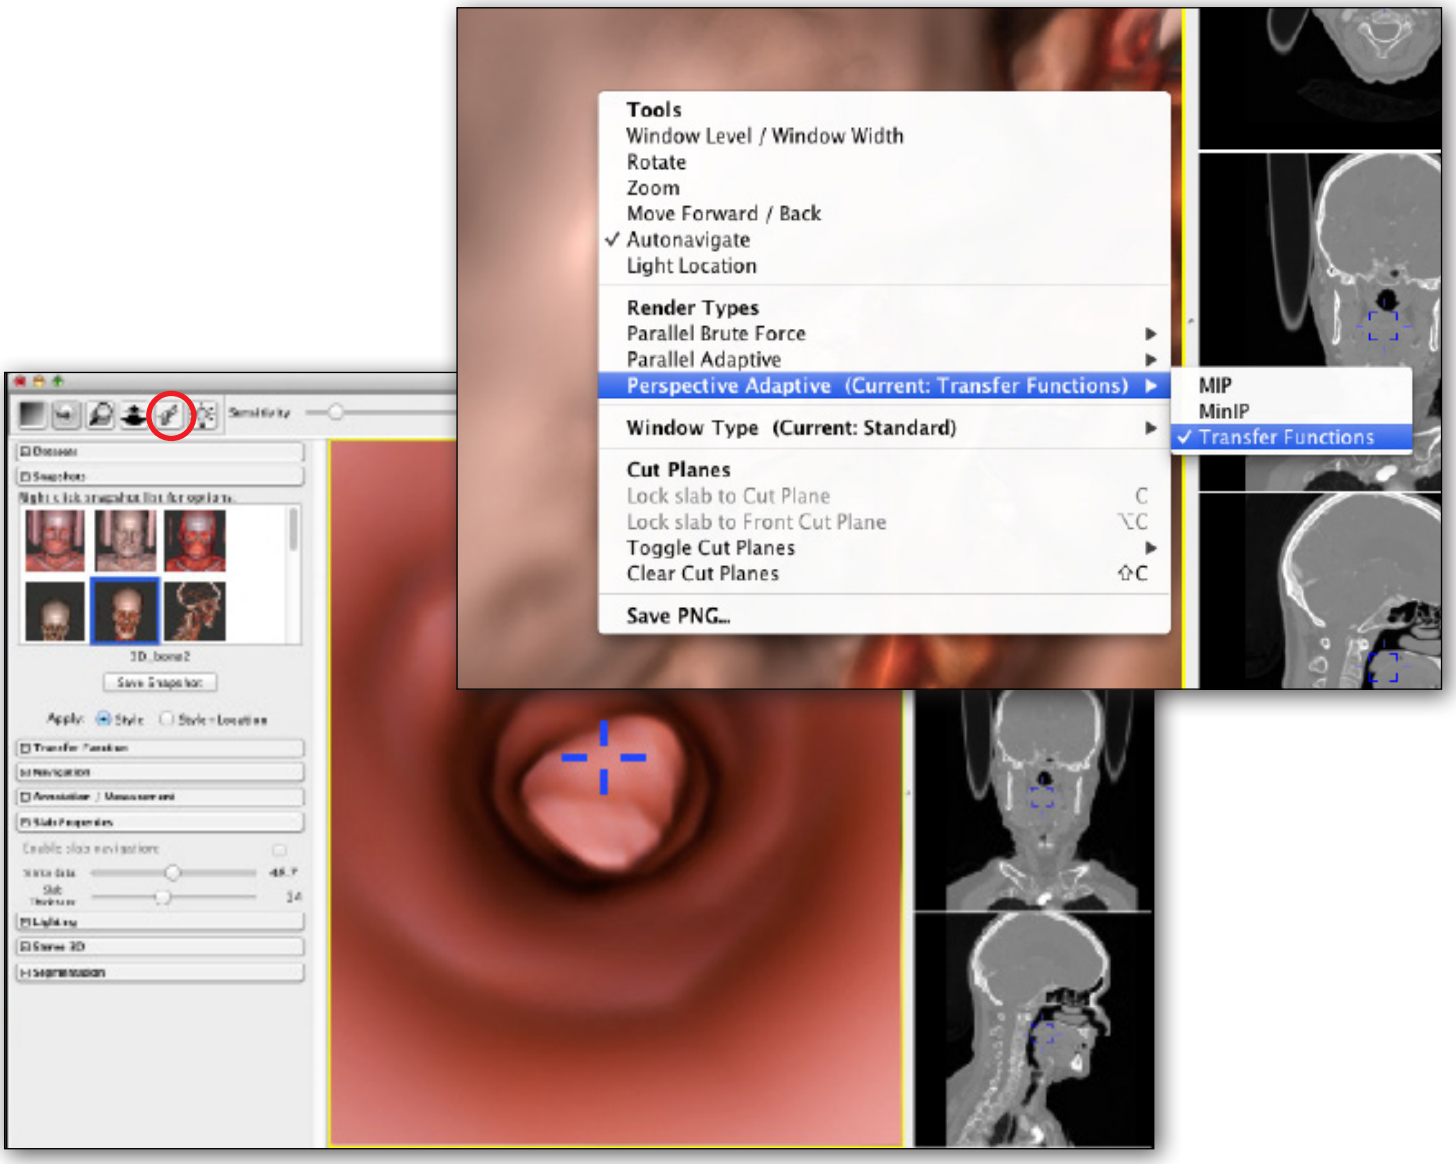

FIGURE 3.9: Location of fly-through, ideal position of object for auto-navigation, in blue crosshairs

The transfer function is a critical tool for exposing the internal structures of volumetric data and fine tuning the image quality. The Transfer Function maps the radiographic visibility of original pixels of volumetric data to corresponding color values and opacities.

4.1 TRANSFER FUNCTION - INTRODUCTION

The Transfer Function Editor is located in the Transfer Function Tab to the left of the screen (Figure 4.1). The Transfer Function works in conjunction with the Snapshots and the Segmentation therefore the corresponding chapters should be referred to as well. The Snapshot Chapter (p. 32) shows the user how to utilize pre-defined Transfer Function settings, and how to save and apply Transfer Function settings. The Segmentation Chapter (p. 39) will show you how to utilize up to 256 separate Transfer Functions (8 available in the workstation). Multiple Transfer Functions can be used for a variety of workflows and allow the user complete control over viewing and manipulating the data.

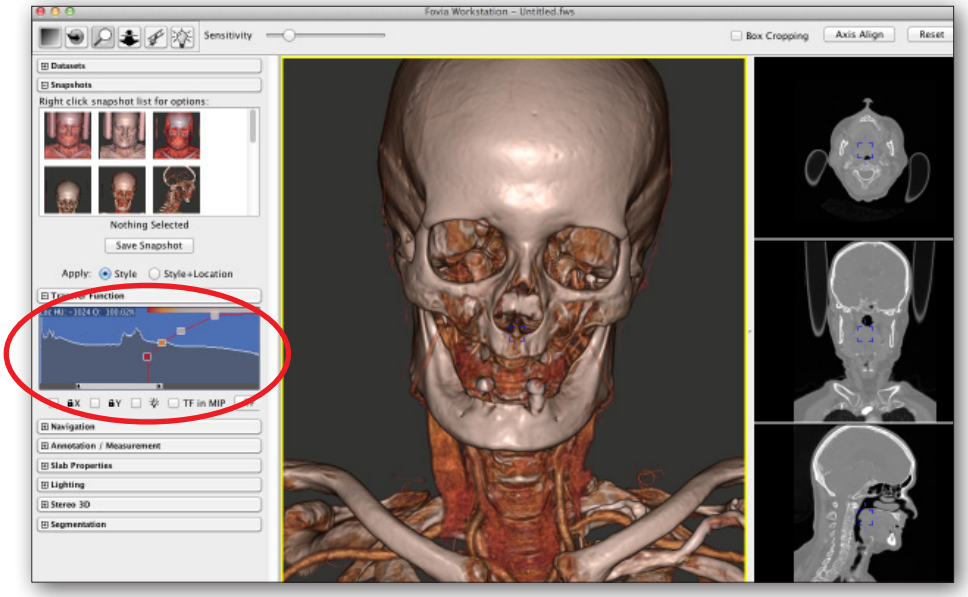

FIGURE 4.1: Transfer Function Location

Within the Transfer Function (Figure 4.2), the gray area is a histogram representing the voxel values in the data set. The height of the histogram at a particular point on the X-axis indicates the number of voxels that exist in the data set at that particular value. In this case, with CT, it represents how many voxels exist at a particular Hounsfield value.

The vertical axis (Y-axis) determines the opacity/translucence of the data set. The horizontal axis represents the scalar field value. The values to the right equal more dense material and values to the left equal less dense material. In a CT scan, values to the right equal dense tissue such as bone, while values to the left equal less dense tissue like soft tissue and skin.

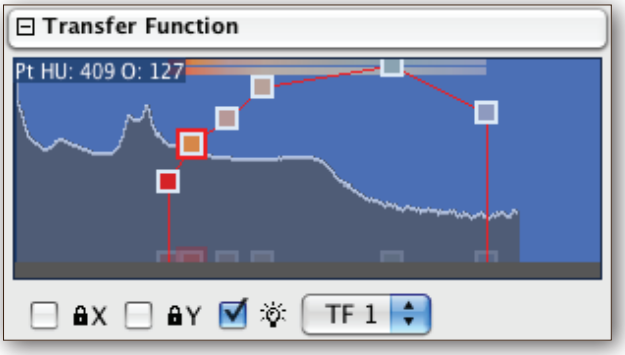

FIGURE 4.2: Transfer Function

Each square control point represents the color map of original pixel values translated into RGB. The section of the dataset that is visible in the viewing window is represented by the area of the histogram that lay between the two bounding control points (Figure 4.3, Control Point A and Control Point C).

The voxel data that has the density represented by the vertical map of the histogram at Control Point B will be represented on screen by the color assigned to Control Point B. The opacity of the voxel data is determined by Control Point B's position on the Y-axis.

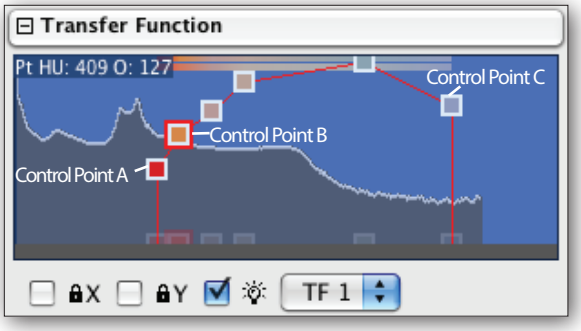

FIGURE 4. 3: Control Point B Lower Opacity - TF

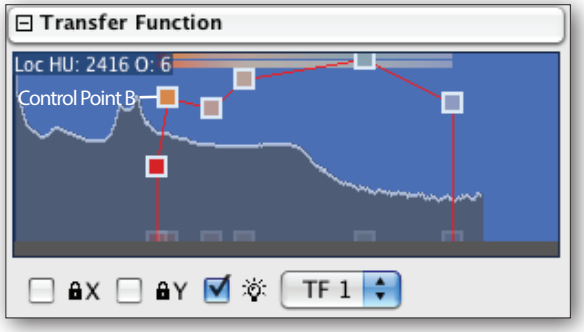

FIGURE 4. 4: Control Point B Higher Opacity - TF

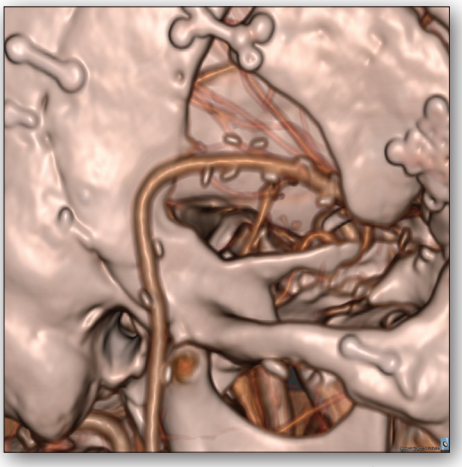

FIGURE 4.5: Control Point B
Lower Opacity - *XStream HDVR*

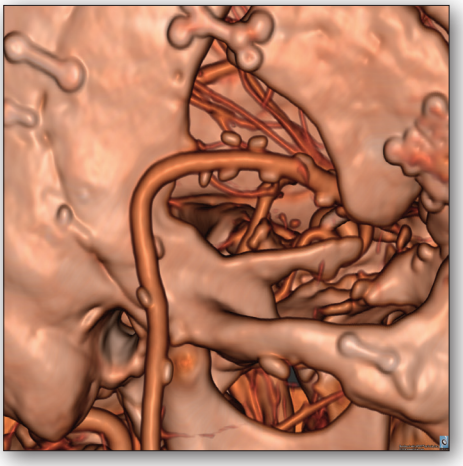

FIGURE 4.6: Control Point B
Higher Opacity - *XStream HDVR*

The result of the vertical movement of Control Point B vertically (shown in Figures 4.3 and Figure 4.4) is reflected in the rendering difference shown in Figures 4.5 and Figure 4.6. By moving Control Point B vertically in the Transfer Function, the artery in the corresponding data set (Figure 4.6) becomes more opaque. In other words, voxels corresponding to the RGB value of Control Point B (orange) are rendered as more opaque. In addition, the overall color changes because Control Point B, with its higher opacity, now contributes more to the final color of the material. Controlling the movement of control points will be discussed in the Navigation Section. Note that these images are from a CTA Runoff.

4.2 TRANSFER FUNCTION - NAVIGATION

4.2a EDITING CONTROL POINTS

This section introduces Transfer Function editing using control points. Control points can be adjusted both individually and as a group. It is possible to add and subtract control points, move control points vertically and horizontally, and change the color of control points.

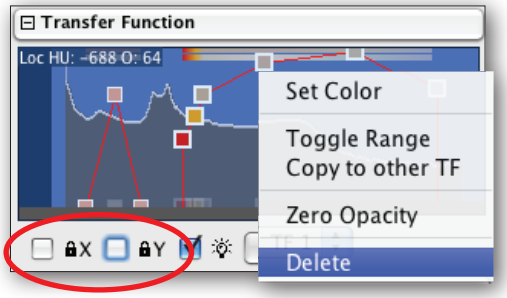

FIGURE 4.7: Locked X-axis

Select a control point by left-clicking on it. To select multiple control points, hold down control (command in Mac OSX) and left-click on multiple control points in succession.

Double-click to add a control point. To delete a control point, use the delete key on the keyboard. Alternatively, the user can delete control points by placing the cursor in the Transfer Function Window and selecting the Delete Function within the right mouseclick drop down menu (Figure 4.7).

As discussed in the introduction, control points can be moved to affect what data is rendered as well as how it is rendered. It is a good idea to adjust control points in one axis at a time. This is facilitated by the Axis Locking Function shown in Figure 4.7 (the figure illustrates the locking of the X-axis). Moving on the X-axis only will change which scalar field values are affected by the selected control point, while moving on the Y-axis only will change the opacity/translucence of data corresponding to the selected control point.

The Zero Opacity Function is available in the right mouseclick drop down menu (Figure 4.8), and sets any selected control points to zero opacity. Below, in Figures 4.8-4.11, the three selected control points on the left are affected.

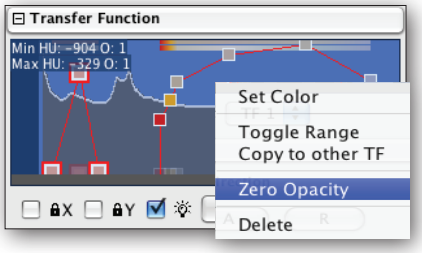

FIGURE 4.8: Zero Opacity

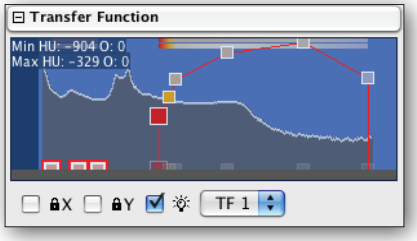

FIGURE 4.9: Zero Opacity Applied to TF

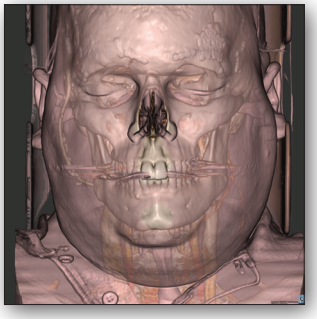
FIGURE 4.10: Original Opacity

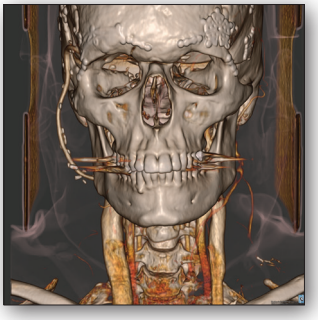
FIGURE 4.11: Zero Opacity

The Color Picker can be used to change or adjust the color of the control points, and thus which colors are rendered to represent the dataset. To open the Color Picker, right click on a control point (Figure 4.12). Within the Color Picker there are three different methods available to choose color. The default option (shown in Figure 4.13) allows the user to choose based on the HSB color system (Hue, Saturation, and Brightness/Value). The second method is to select a color from a Swatch Library. When using the swatch library, the color history is catalogued as colors are selected (Figure 4.14). The third method is to select colors using RGB sliders (Figure 4.15).

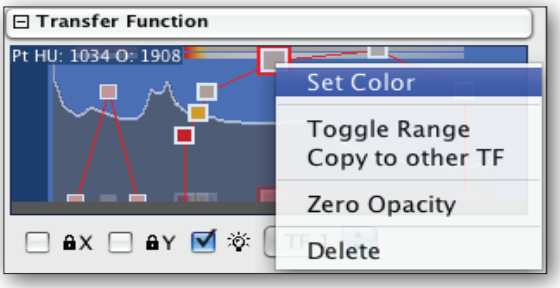

FIGURE 4.12: Set Color

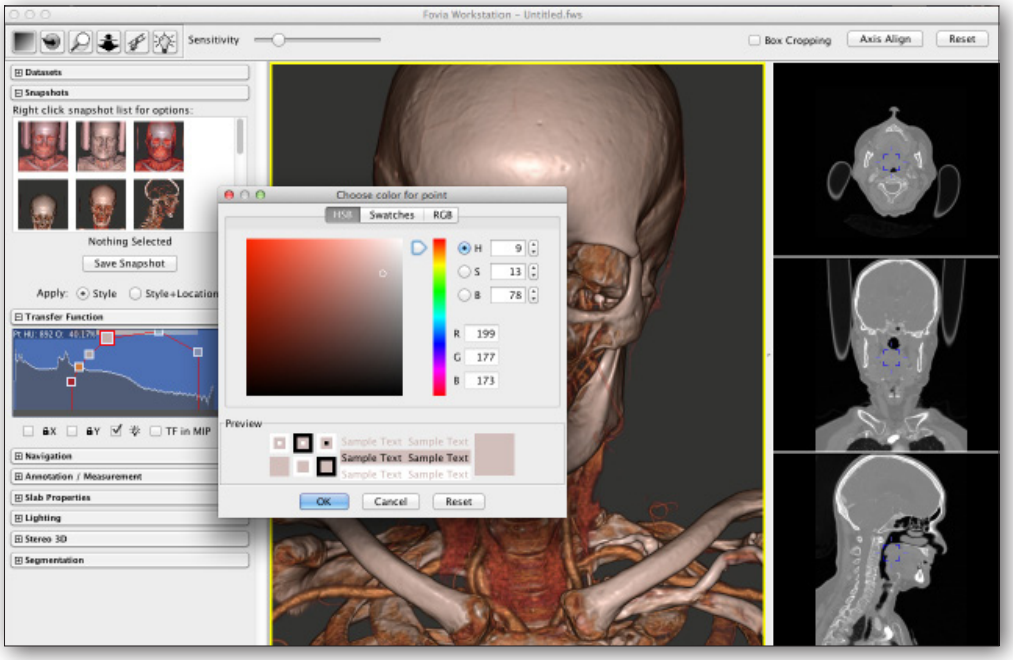

FIGURE 4.13: Set Color - HSB

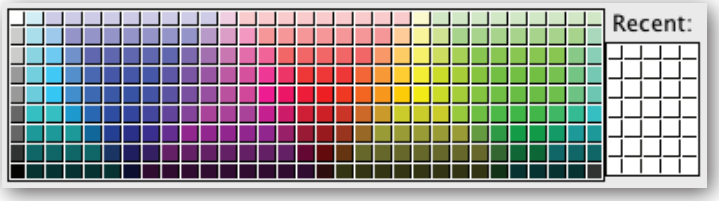

FIGURE 4.14: Color Swatches

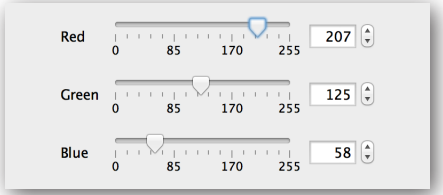

FIGURE 4.15: RGB Sliders

The color values (as well as the opacity values) between control points are populated using linear interpolation. Extreme color changes between control points in conjunction with high opacity may increase the probability of under-sampling artifacts, referred to as zebra patterning. When setting the color for non-transparent surfaces when Phong lighting is on, avoid setting color components to 255. This is to prevent over-saturation of the image in regions featuring specular reflective highlights. (See Lighting Section (p. 68) for more information).

4.2b ADVANCED EDITING - ZOOM, MULTIPLE RANGES, TRANSFER FUNCTIONS WITH SEGMENTATION

To increase accuracy in editing, zoom in and out of the Transfer Function Editor (Figures 4.16-4.18) by using the mouse scroll wheel. Once zoomed in, a scroll bar appears at the bottom of the Transfer Function screen that allows movement horizontally along the Transfer Function.

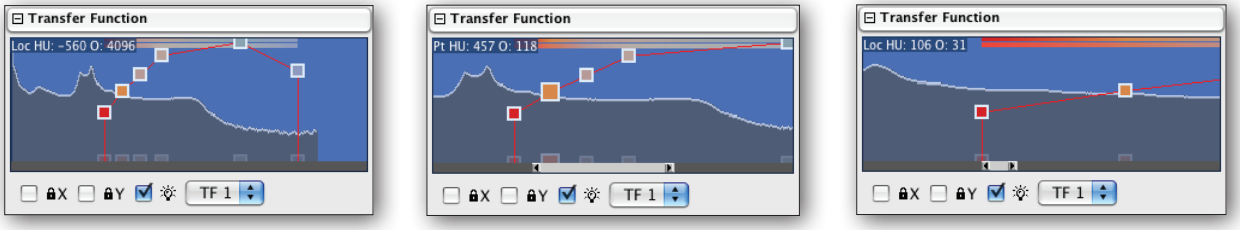

FIGURES 4.16-4.18: Transfer Function Zoom

The Transfer Function may have multiple ranges (as many as eight contiguous segments) and each range may be represented by up to 50 points (Figure 4.19 shows 2 ranges). These ranges can be toggled on and off separately by using the Toggle Range Function in the right mouseclick drop down menu (Figure 4.20). Note that in Figure 4.20 the red line representing the Opacity Transfer Function for range 1 is now black, demonstrating that range 1 has now been turned off (Figures 4.21 and 4.22 show that change). To make a new range, double click within the scalar field.

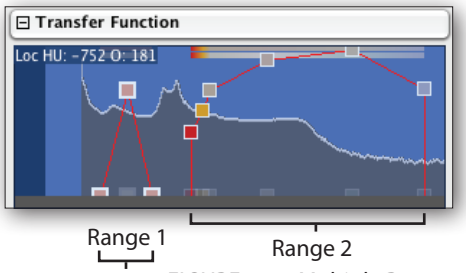

FIGURE 4.19: Multiple Ranges

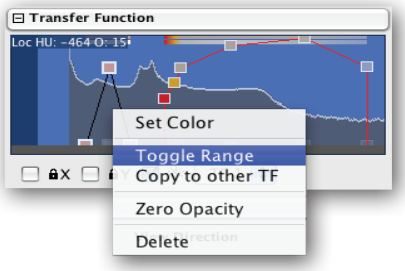

FIGURE 4.20: Toggle Range

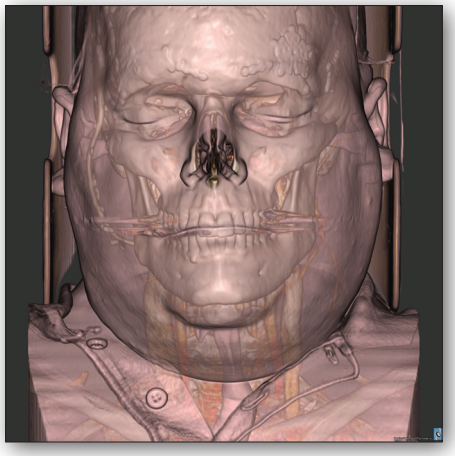

FIGURE 4.21: Multiple Ranges - VR

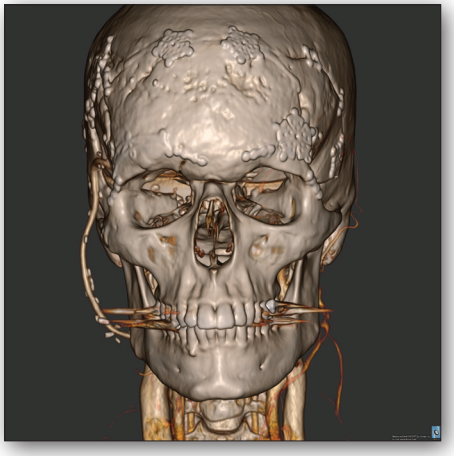

FIGURE 4.22: Toggle Range - VR

4.2a TRANSFER FUNCTION WINDOW LEVEL/WIDTH

The Transfer Function can be moved as a whole, each range can move individually, or a subgroup of selected points can move horizontally, vertically, or a combination of the two. The entire Transfer Function can be moved by left-clicking and dragging anywhere in between the two outermost bounding control points, or by left-clicking on the Color Scaler at the top of the Transfer Function (Figure 4.23) and dragging. When there is more than one range, move each one individually by left-clicking and dragging anywhere in between the two bounding control points of that range, or by clicking on the Color Scaler directly above that range and dragging. To move multiple ranges at once, use the Window Level/Width tool at the top toolbar (to be discussed next). Additionally, a group of control points can be moved by selecting the desired points, left-clicking with the cursor on one of the selected control points, and dragging (Figures 4.24-4.25).

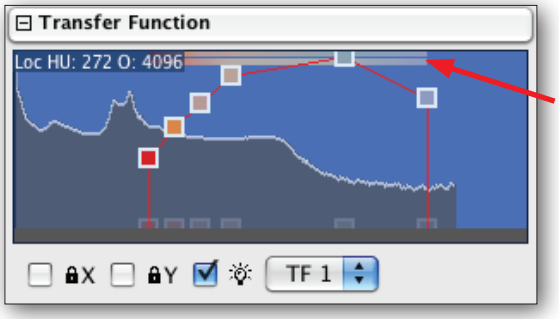

FIGURE 4.23: Location of Color Scaler for movement of Transfer Function

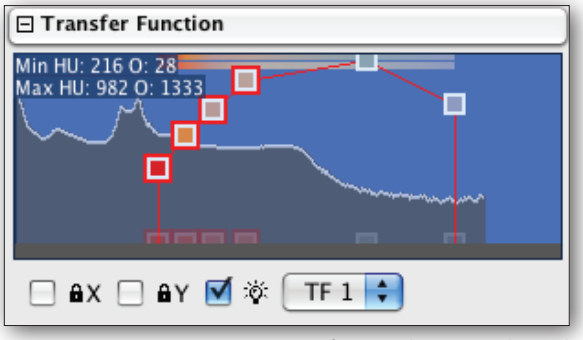

FIGURE 4.24: Group of Control Points selected

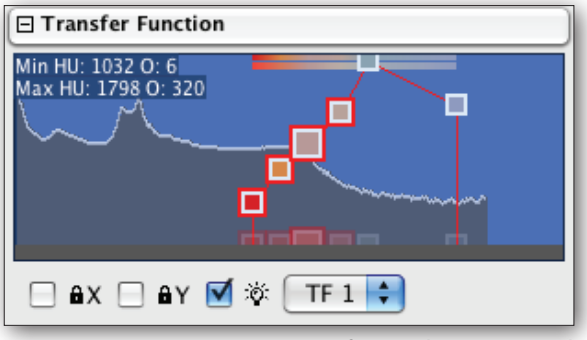

FIGURE 4.25: Group of Control Points moved

In addition to editing the Transfer Function within the Transfer Function Tab, the user may employ a Window Level/Width Tool that can be used to adjust the entire Transfer Function. This tool can be found in the tool bar at the upper left corner of the screen (Figure 4.26-4.27). Dragging the mouse up and down on the viewing screen moves the position of the Transfer Function horizontally on the scalar field, thereby changing which densities are visible on the viewing screen. Dragging the mouse left and right on the viewing screen changes the distances between the control points. In addition, selecting both the right and left buttons on the mouse simultaneously creates a shortcut for this tool.

FIGURE 4.26: Window Level/Width Tool

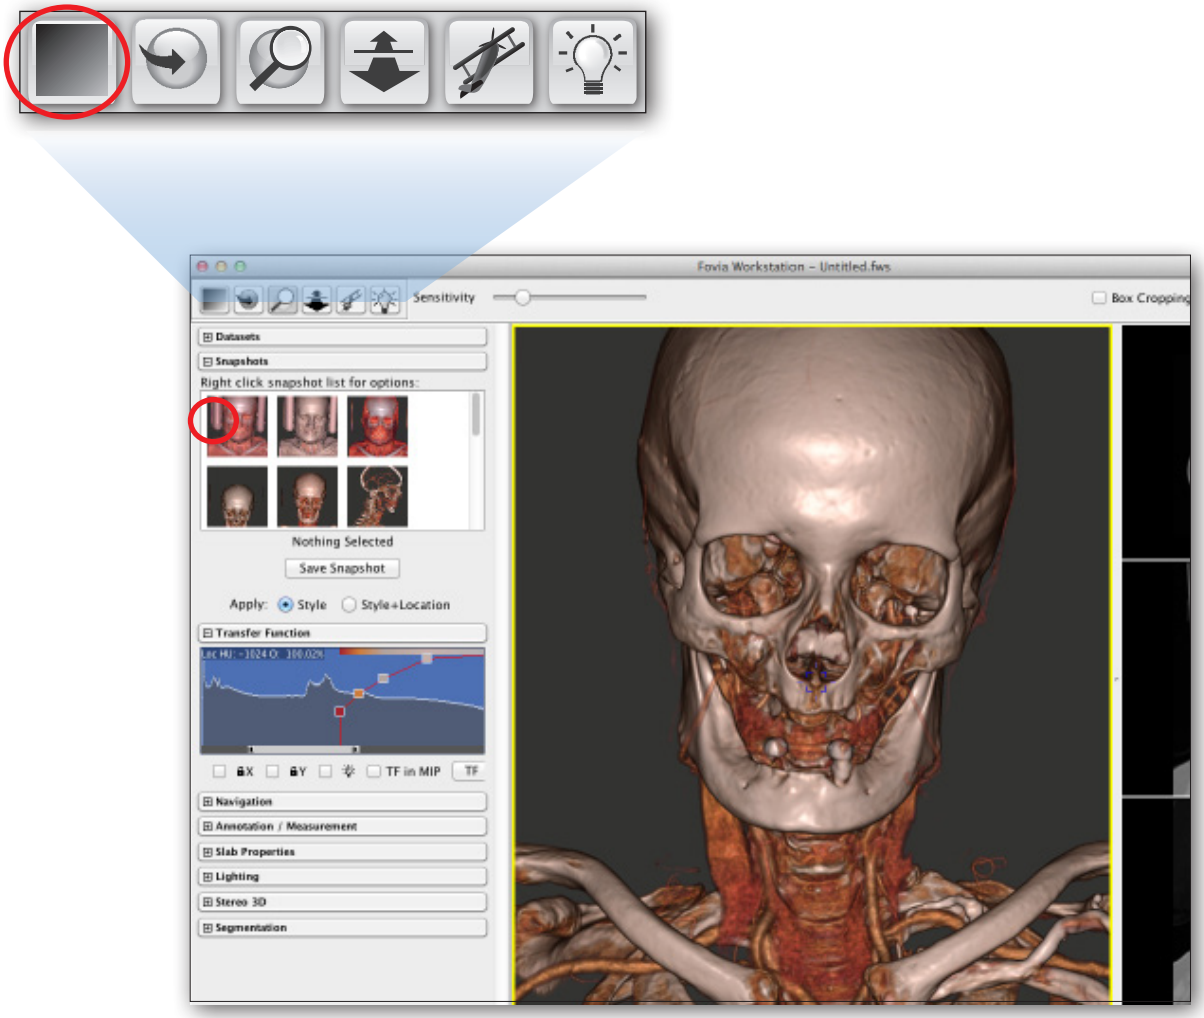

FIGURE 4.27: Location of Window Level/Width Tool

4.3 TRANSFER FUNCTION - TF LIGHTING

Within the software, lighting can be adjusted a few different ways. In this chapter, we will focus on the Phong Lighting Control Tool (Figure 4.28) found within the Transfer Function Tab. Subsequent chapters will discuss how lighting can also be adjusted by using the Lighting Tool found in the main toolbar (Figure 4.29) and how overall lighting controls can be found and adjusted in the Lighting Tab to the left of the screen (Figure 4.30).

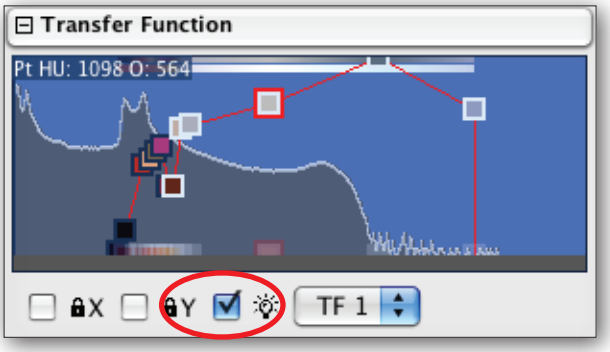

FIGURE 4.28: Phong Lighting Control Tool

FIGURE 4.29: Lighting Control Tool

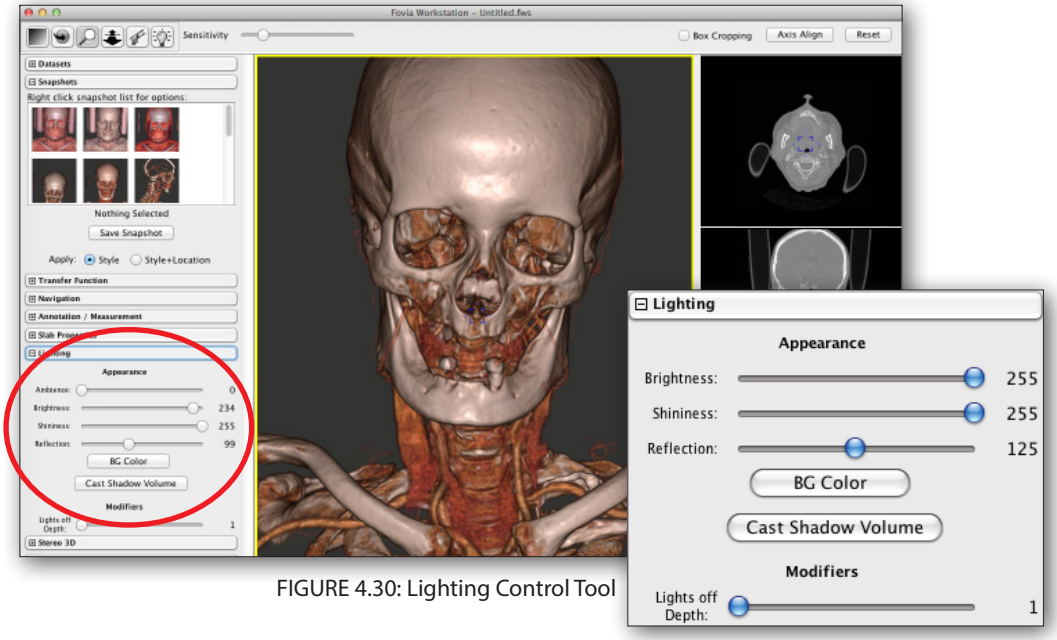

FIGURE 4.30: Lighting Control Tool

The Phong Lighting Control allows you to set lighting ON or OFF for specific data-ranges. This is accomplished by setting the lighting for specific control points or groups of control points. In Figure 4.31 the control points to the left (outlined in black) have lighting turned OFF while the control points to the right (outlined in white) have lighting turned ON. The Phong Lighting Control box is checked for the selected control point (Control Point A) ,indicating that lighting is turned ON.

FIGURE 4.31: Control Point Specific, Lighting Turned ON or OFF

Adjusting the lighting allows the user to see certain surface properties more accurately. The preset "3D_bone&muscle" (Figure 4.32) is a good example of turning the lighting OFF to view semi-transparent tissue (e.g. muscle) and turning the lighting ON for more opaque tissue (e.g. bone). The result of this preset are illustrated in Figures 4.33-4.34 on the following page. CT modality representation of muscle and air or MRI modality tissues are the most common candidates for visualization with the lighting turned OFF. With semi-transparent tissue such as muscle, lighting turned ON can lead to increased noise (Figures 4.35-4.36 on the following page). Lighting should be turned ON for more opaque tissue such as bone, as it helps to provide definition. Notice the lack of definition in the bone when lighting is turned OFF (Figures 4.37-4.38 on the following page).

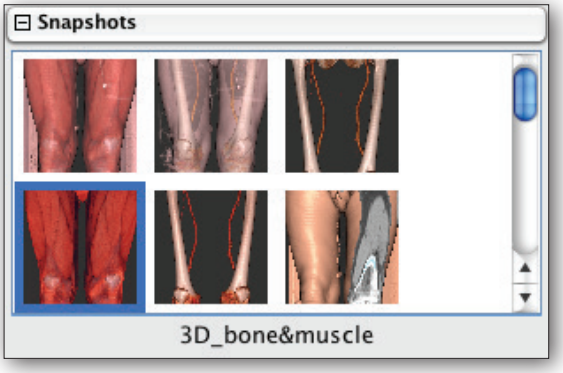

FIGURE 4.32: Preset: 3d_bone&muscle

Each Transfer Function below is rendered as the *XStream HDVR* image to its right. The user can examine these Transfer Functions and images to further understand the use of lighting within the software.

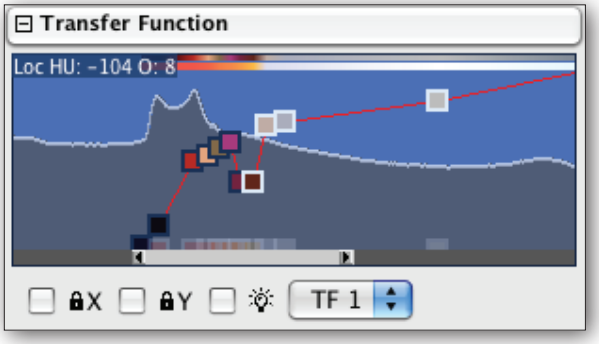

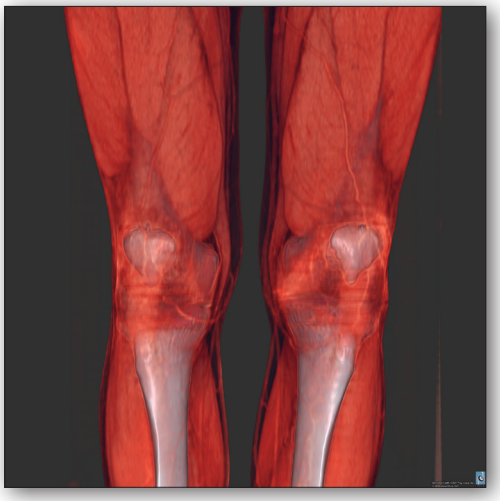

FIGURES 4.33-4.34: Lighting for Muscle OFF and Bone ON

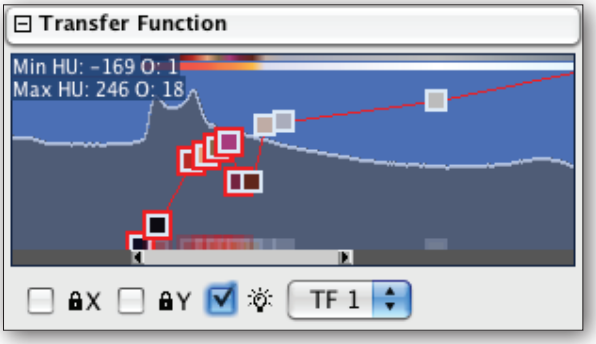

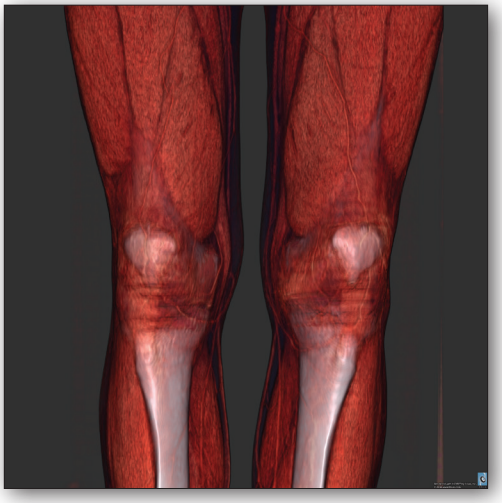

FIGURES 4.35-4.36: Lighting for Muscle and Bone ON

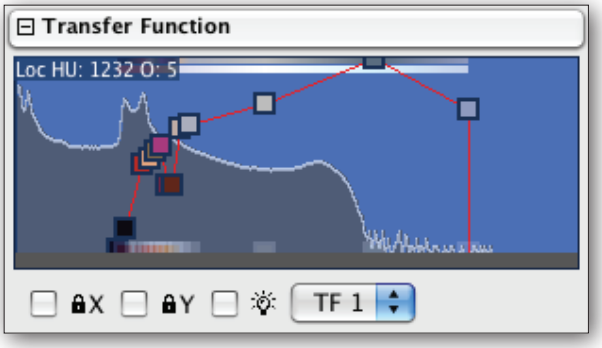

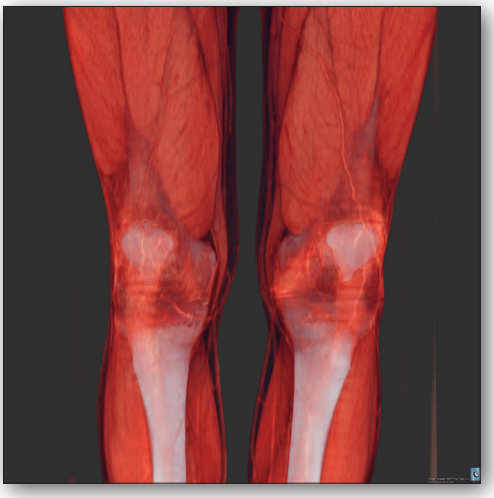

FIGURES 4.37-4.38: Lighting for Muscle and Bone OFF

4.4 TRANSFER FUNCTION - BEST PRACTICES

Fovia's *XStream HDVR* comes with a variety of presets to help familiarize the user with the Transfer Function. A few are demonstrated below (Figures 4.39-4.44).

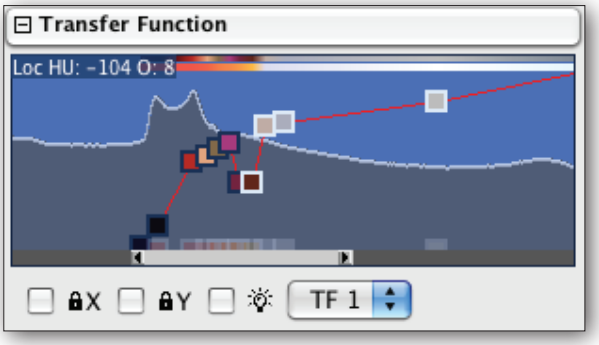

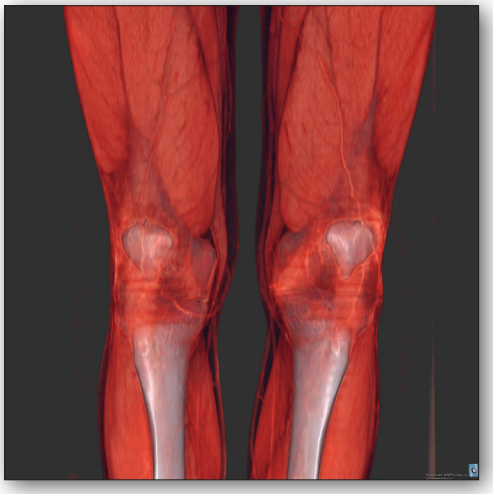

FIGURES 4.39-4.40: 3D bone&muscle

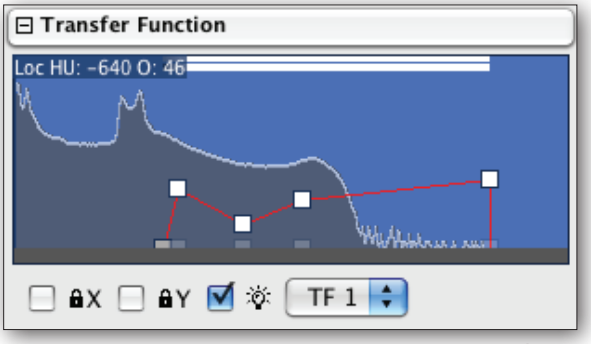

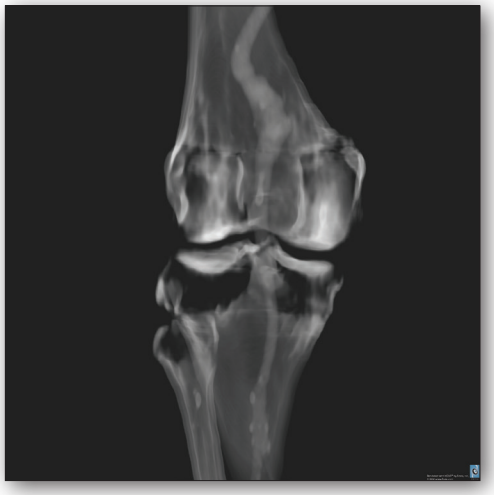

FIGURES 4.41-4.42: Xray_bone

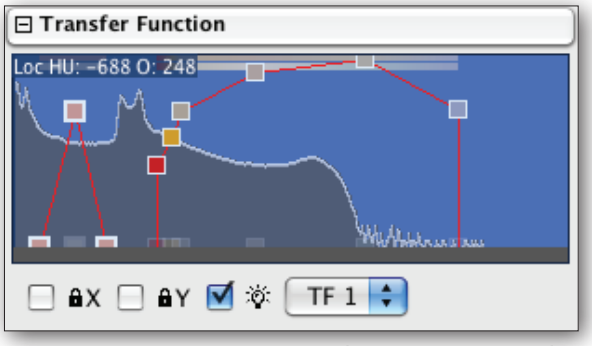

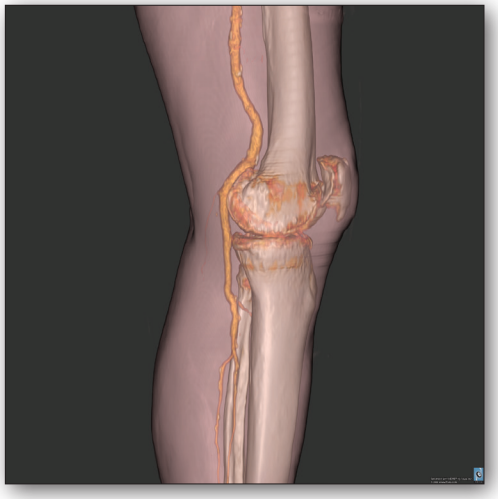

FIGURES 4.43-4.44: 3D_bone&transparent_skin

In addition to the presets, the user may want to create and/or customize the Transfer Function so that it applies more directly to the user's specific viewing needs. The following information provides general tips to the user for improved Transfer Function utilization.

BEST PRACTICE TIP 1: Minimize Undersampling (Figures 4.45-4.48). Rapidly changing color (sharp color gradients in the Transfer Function) in conjunction with high opacity may increase the probability of under-sampling artifacts and create the "zebra" pattern shown Figure 4.48. Therefore, keeping color as uniform as possible and maintaining relatively low opacities will reduce the probability of under-sampling artifacts.

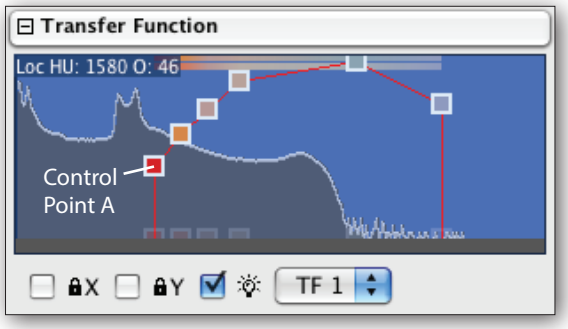

FIGURE 4.45: Control Point A with Low Opacity

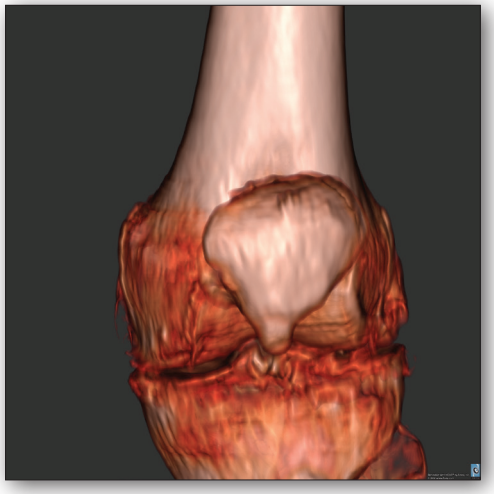

FIGURE 4.46: Minimal Under-Sampling

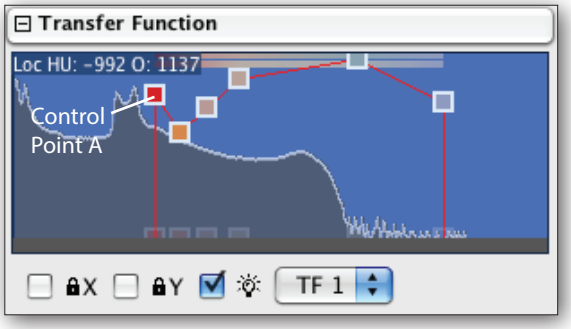

FIGURE 4.47: Control Point A with High Opacity

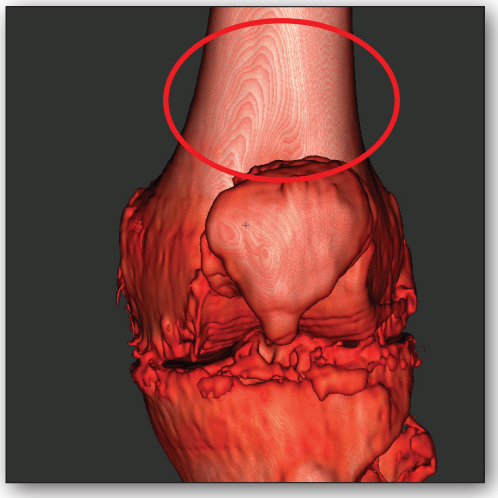

FIGURE 4.48: Increased Under-Sampling and "Zebra" Effect

BEST PRACTICE TIP 2: Control Point Spacing (Figure 4.49). For best interactive performance, try to keep the distance between control points along the X-axis as equal as possible.

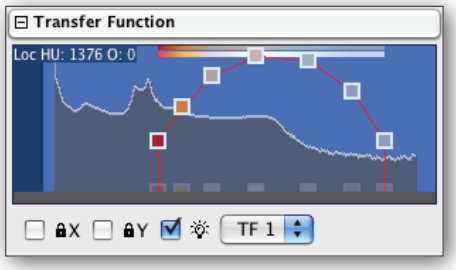

FIGURE 4.49: Evenly spaced Control Points

Customization Tip 3 - Value Range for Skin (Figures 4.50-4.54). The range of values for film-like tissue such as skin is considerably smaller than ranges of values for other types of tissue and will take up only a small portion of the scalar field within the Transfer Function. Because of this small range of values, this tissue has the highest probability of being under-sampled. Therefore these tissues should be represented by Transfer Function ranges that are as wide as possible within "skin section" of the scalar field.

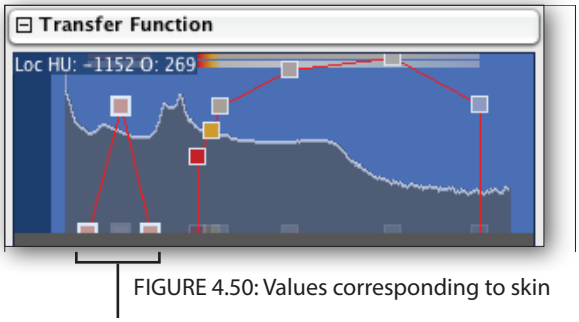

FIGURE 4.50: Values corresponding to skin

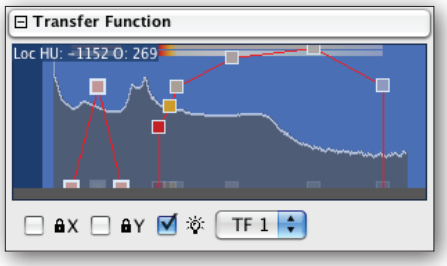

FIGURE 4.51: Wide TF Range of skin values

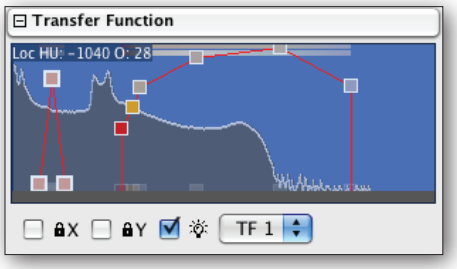

FIGURE 4.52: Narrow TF Range of skin values

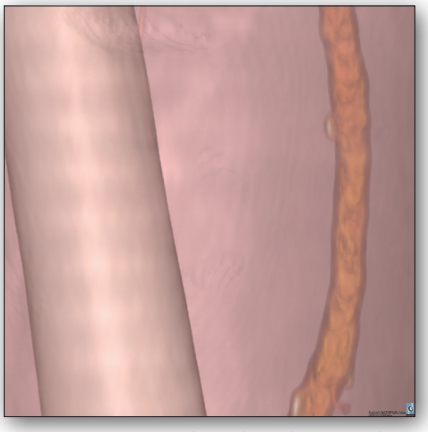

FIGURE 4.53: Reduced Under-Sampling

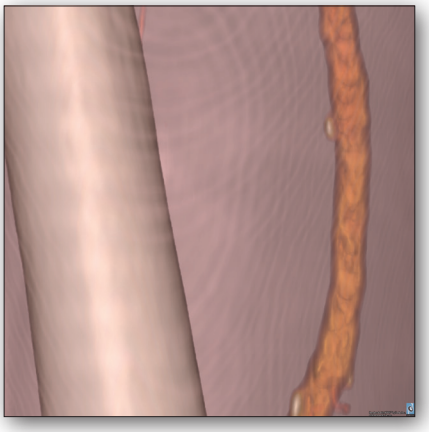

FIGURE 4.54: Increased Under-Sampling

BEST PRACTICE TIP 4: Transfer Function Shape (Figures 4.61-4.66). Begin with a triangular or trapezoid shaped Transfer Function (Figure 4.55-4.56). Next, to add more depth to the rendered image, progress towards a bell shaped curve (Figure 4.57-4.58) by adding additional ascending and descending controlpoints. This will add more depth to the rendered image. These ascending and descending control points will largely affect the tonal range of the rendered image, while the bounding control points will contribute more to the overall color.

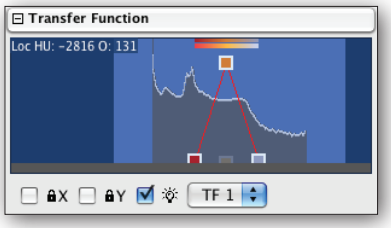

FIGURES 4.55-4.56: Triangle Shape TF

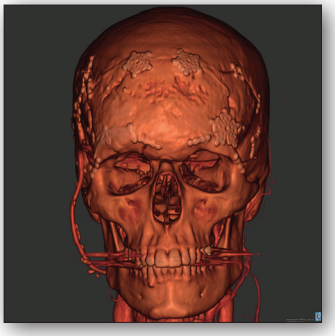

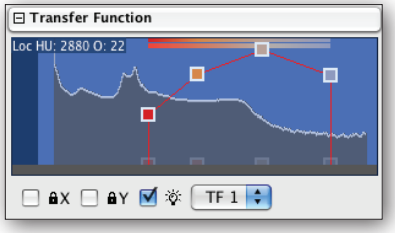

FIGURES 4.57-4.58: Adding Control Points

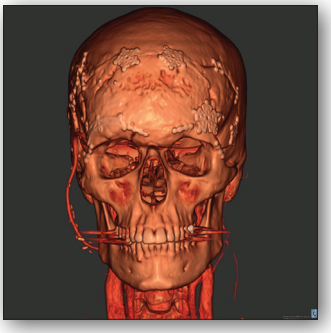

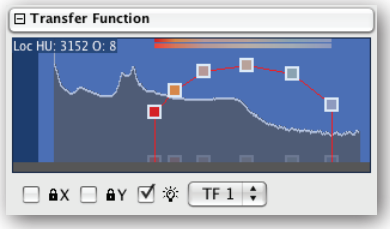

FIGURES 4.59-4.60: Adding Control Points for depth

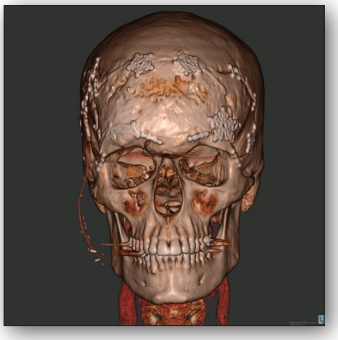

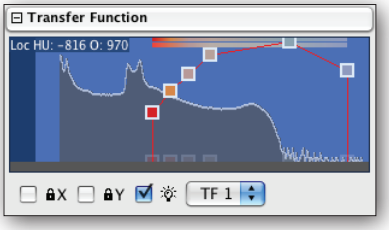

FIGURES 4.61-4.62: Final adjusted TF

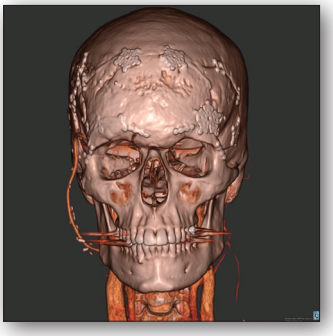

Different colors and transparencies can be assigned to various tissue types in order to define skin, muscle, soft tissue and bone. These assignments can be saved as "Snapshots." A variety of preset snapshots are available to begin the Fovia experience.

5.1 SNAPSHOTS & WORKSPACES - SNAPSHOTS

Default snapshots are provided with the workstation application. These snapshots should appear upon loading a datasets. To import the default snapshots into a workspace, simply right-click in the snapshot window and select "Open Defaults" (Figure 5.1).

DEFAULT SNAPSHOTS:

- 3D_bone&muscle&transparent_skin
- 3D_bone&transparent_skin
- 3D_bone&muscle 3D_bone
- 3D_bone2
- 3D_bone_enhanced_vessels
- heart_fly
- heart_fmip heart_surface
- lung_film PET_start PET_start Xray
- x-ray
- xray_bone
- FHC_3D_MPR
- FHC_3D_SOFT_MPR
- FHC_3D_MPR_bone_expose
- FMIP

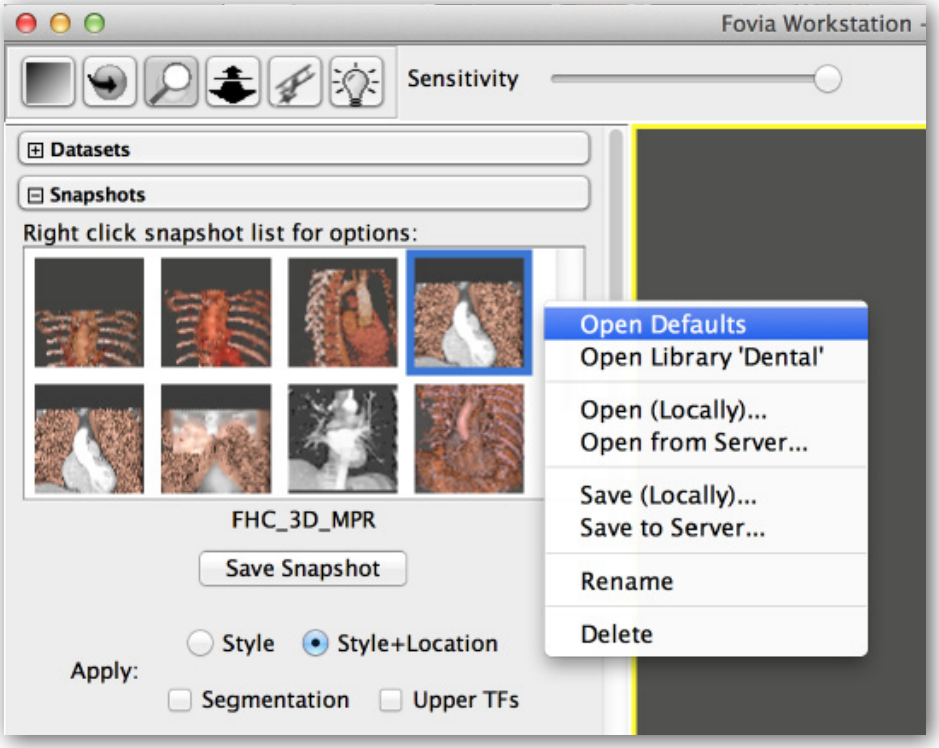

FIGURE 5.1: Open default snapshots

TO LOAD PRESET SNAPSHOTS

- Place cursor in snapshot window.
- Right mouseclick.
- Select "Open Defaults."

TO APPLY A SNAPSHOT TO A DATASET

- Select the viewing window with the targeted dataset for snapshot
- Select "Style" to apply only the settings
- Select "Style + Location" to apply both the settings and the orientation
- Select "Segmentation" to apply snapshot segmentation
- Select "Upper TFs" to apply TFs above TF1 (TFs above TF1 will be discussed in the Segmentation Chapter (p. 39)

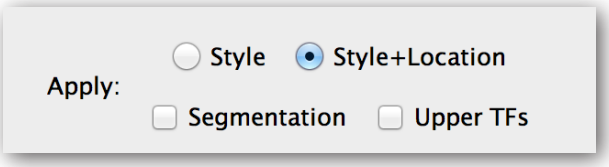

FIGURE 5.2: Snapshot application options

The user can create and save Custom Snapshots, then export them locally or to the server to be imported and applied to other datasets. You can export as a PRS (Standard) or as an XML (Legacy) file. When saving a snapshot to a file on disk, one should note that only XML files can be saved or loaded by the *XStream HDVR* SDK. The SDK does not make use of PRS files, these files are only used in the Workstation application.

TO CREATE A SNAPSHOT

- Select window with desired style, orientation, and/or segmentation
- Click "Save Snapshot" button

TO RENAME A SNAPSHOT

- Right mouseclick
- Select "Rename"
- Type in name

TO EXPORT/SAVE A SNAPSHOT

- Right mouseclick
- Select "Save Locally"
- Select destination folder

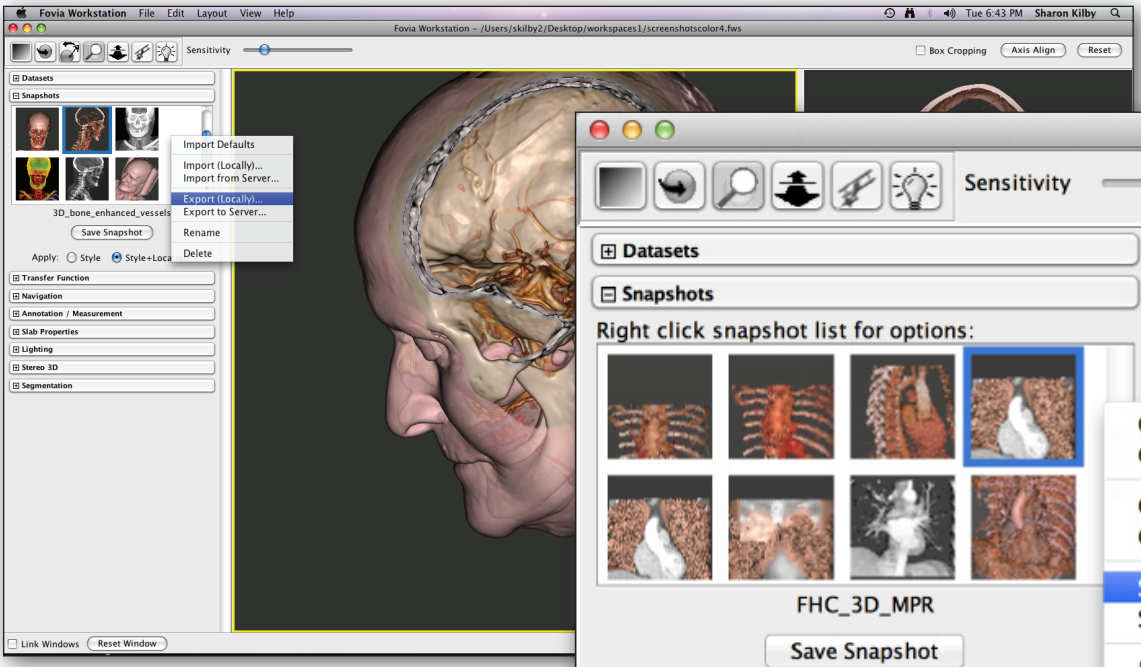

FIGURE 5.3: Export snapshot

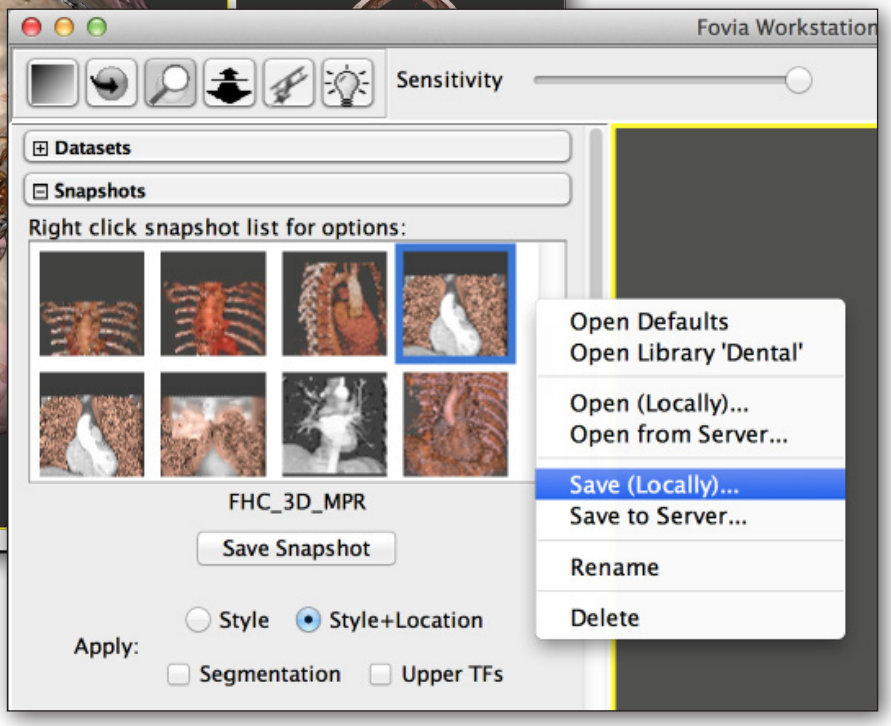

FIGURE 5.4: Export options

5.2 SNAPSHOTS & WORKSPACES - WORKSPACES

Workspaces can be created and saved by the user. Workspace files save all snapshots, data sets paths, and segmentations to a file that can be revisited in another user session.

TO SAVE A WORKSPACE

- In the File Drop-down Menu select "Save Workspace" or "Save Workspace As..."
- Navigate to the appropriate folder
- Click "Save"

TO OPEN A WORKSPACE

- In the File Drop-down Menu select "Open Workspace"
- The user session will be restored

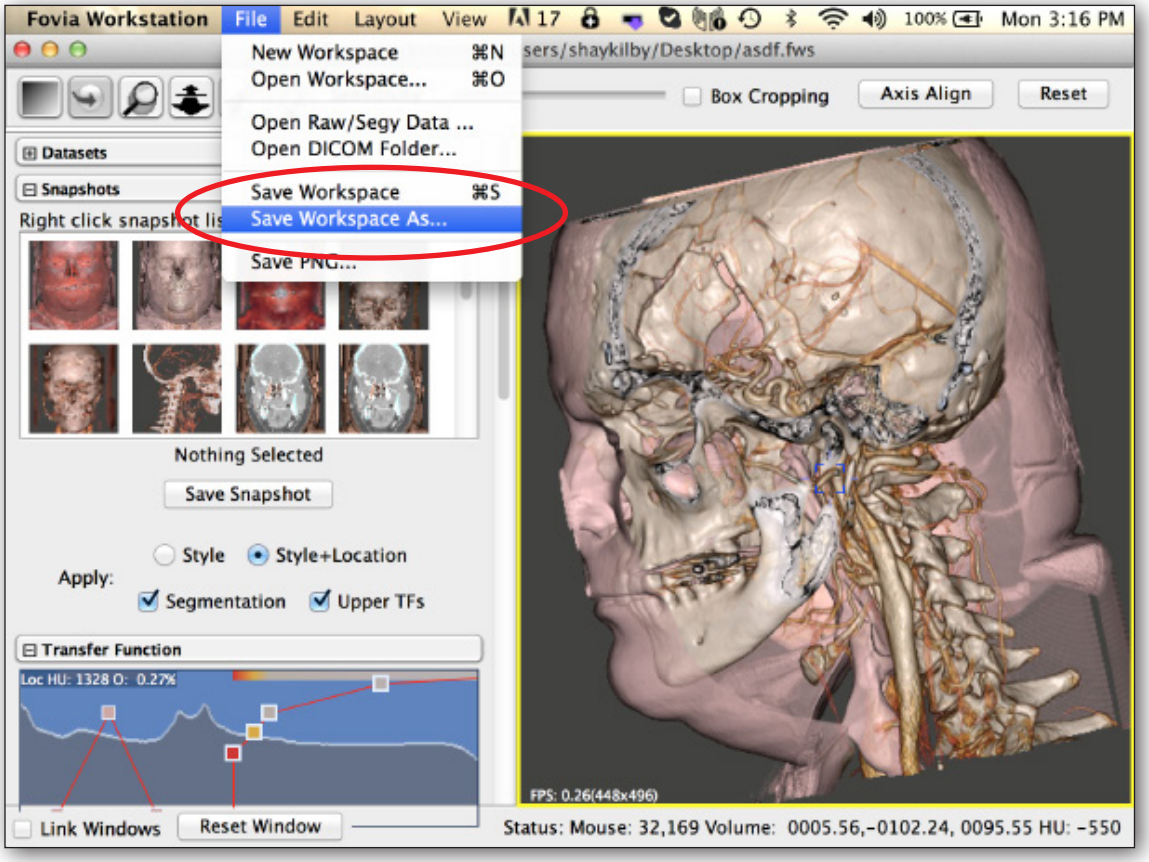

FIGURE 5.5: Saving Workspaces

Annotations are text, lines, or other graphical labels that are drawn on top of the images rendered by the *XStream HDVR* Engine. Annotations can be used to label structural features, as outlines for region selection operations, as measurement lines or for other purposes.

6.1 ANNOTATION - 2D MEASUREMENT

6.1a Arrow Tool

There are two different 2D annotation tools. The first is the arrow tool in the Annotation/Measurement Tab (Figure 6.1). If you are measuring in 2D make sure the "3D" box is unchecked (Figure 6.1). When drawing the arrow you will notice the annotation information appears simultaneously. Also note that you can choose to uncheck "Show Arrow" at the bottom of the annotation tab in order to hide the arrow.

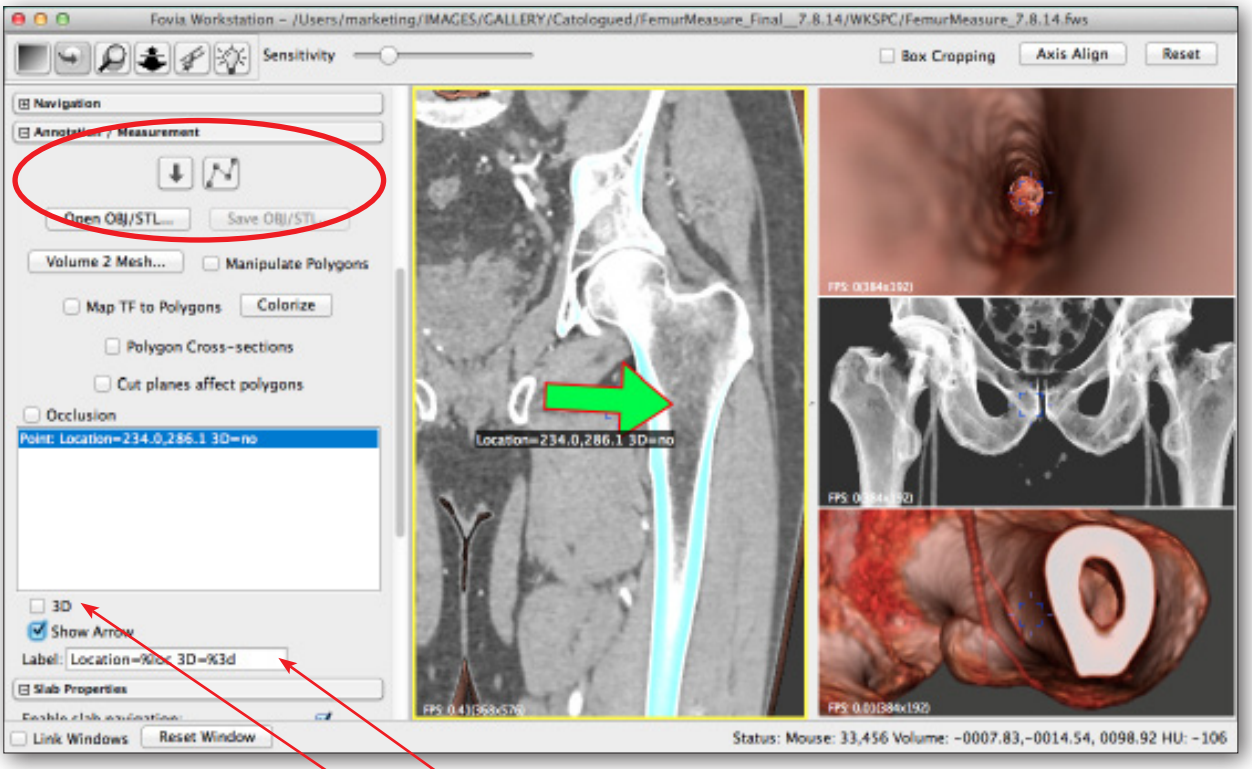

FIGURE 6.1: Annotation Tools

To annotate in 2D make sure the "3D" box is unchecked.

To change the text of the annotation simply type your text in the "Label" box.

6.1b Measurement Tool

The second 2D Annotation tool is the Measurement Tool. Simply click on the measurement button (Figure 6.2) and click and drag across the desired anatomy.

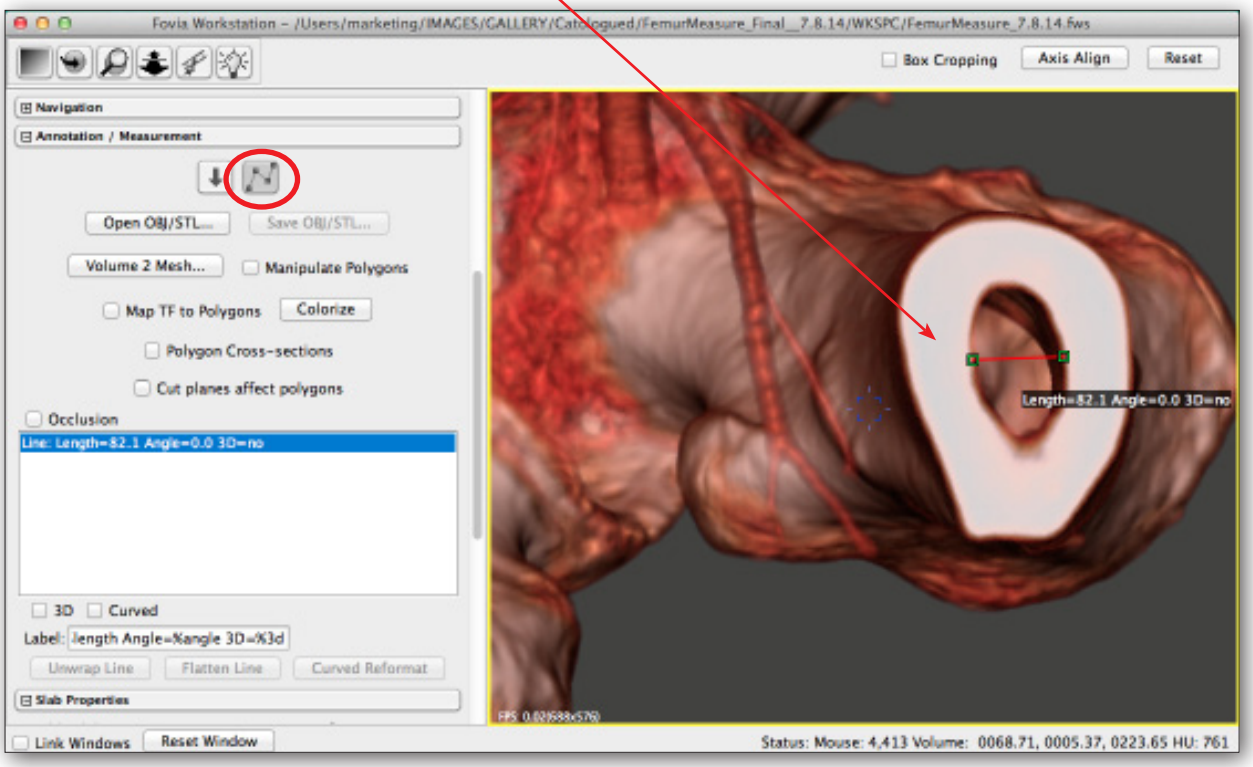

FIGURE 6.2: Create a line

Measurements are given in mm. To create an angle, first create a line, release the mouse, and then left click to create the second line and thus the angle (Figure 6.3). Use the delete key to delete any point on the line or delete the entire measurement in the window of the Annotation Tab. To create a curved line that follows the curve of the anatomy, select the checkbox "Curved" (Figure 6.4).

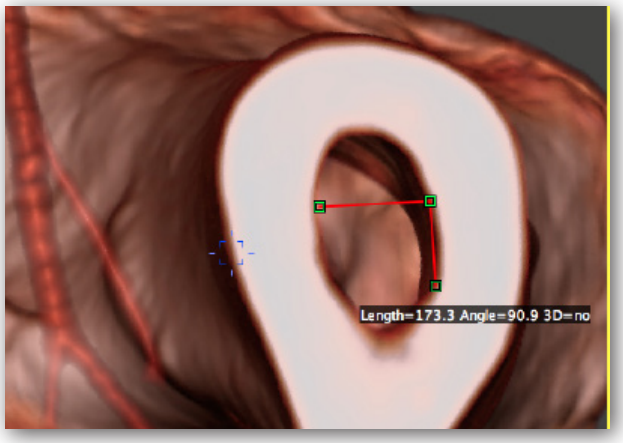

FIGURE 6.3: Create an angle

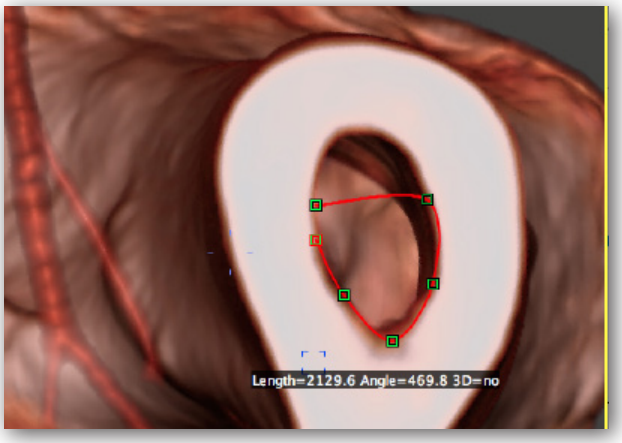

FIGURE 6.4: Curved line

6.2 ANNOTATION - 3D MEASUREMENT

To access the 3D tools, select the checkbox towards the bottom left of the Annotation Tab (Figure 6.5). Once the 3D box is checked, points drawn with the arrow or line tools are in volume space coordinates. If you click on a point on the screen in this mode, the workstation will internally find the volume space coordinate of the voxel at that point. That coordinate is then converted to screen space coordinates by the workstation so that the annotation can be drawn in the right place, and will move when the viewpoint is manipulated. If the 3D box is unchecked, lines and arrows are drawn in screen space pixel coordinates and measurements will display the length in screen pixels, rather than volume coordinates (mm).

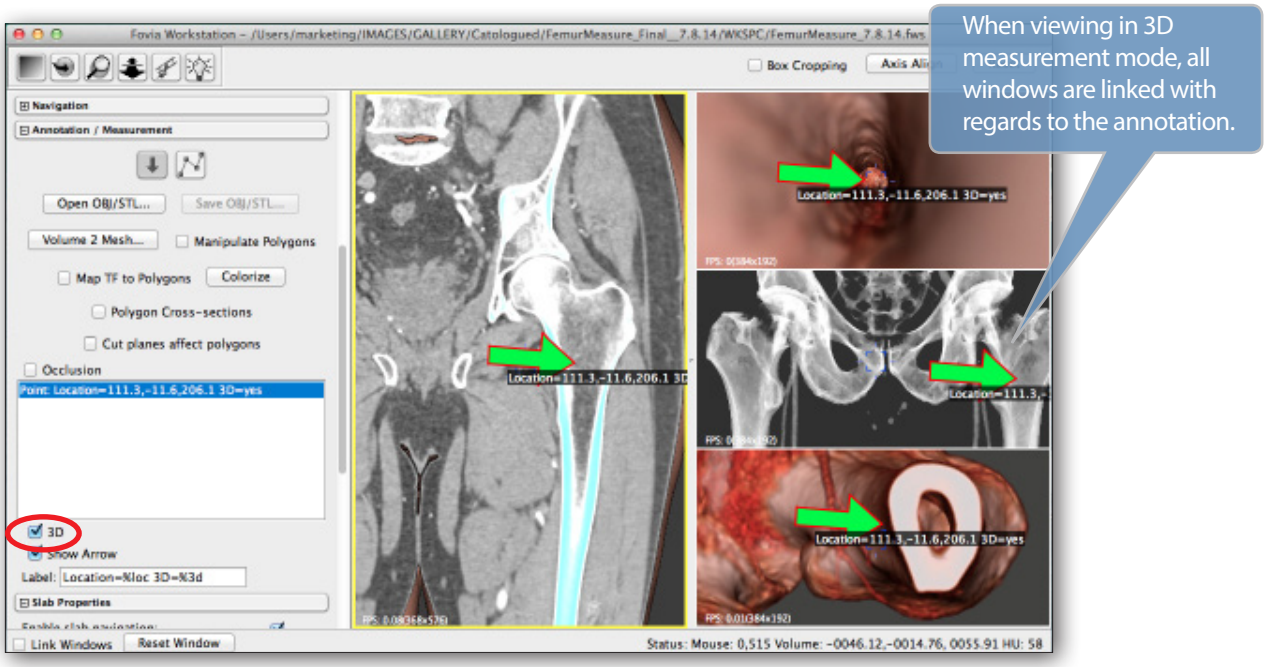

FIGURE 6.5: 3D Measurement

6.3 ANNOTATION - VESSEL TRACE

The workstation supports a feature to automatically computes and generate a path along a vessel or other navigable channel, such as a nerve channel trace, or other path through a channel in the dataset. The path-finding feature specifies start and end points, and other path-finding operation parameters. The path-finding algorithm then sends out tracers from the start and end points that travel through the channel. If the tracers from the start and end points get within a threshold distance of each other, they will connect and a path will be established. The threshold distance and other parameters for the path-finding operation can be customized based on the size and material type of the channel through which a path is being found. (Figure 6.6)

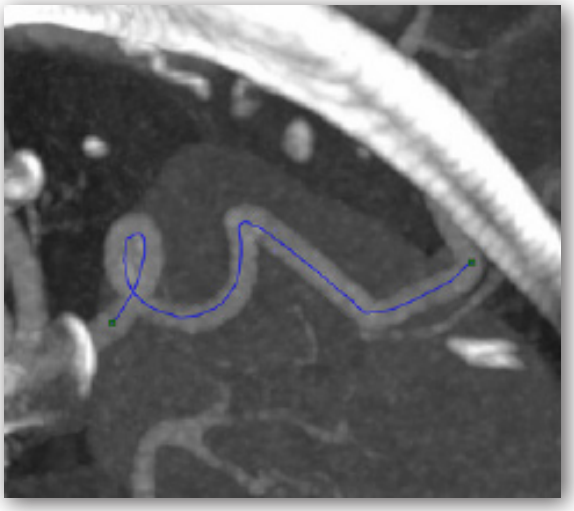

FIGURE 6.6: Path through a curved vessel

For a basic Vessel Trace, mouseclick on a vessel while in 3D measurement mode. Navigate to the next area of interest in the vessel, depress the shift key on your keyboard and left mouseclick on the new area of the vessel. The trace will automatically be created (Figure 6.7). In the bottom of the Annotation/Measurement Tab you can "Unwrap" and "Flatten" the line as well as view in "Curved Reformat."

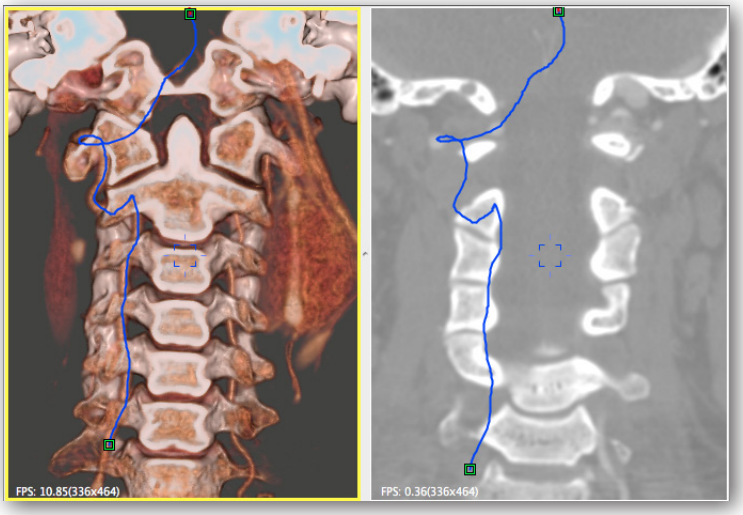

FIGURE 6.7: Vessel Trace

Manual and automatic general purpose segmentation is provided to the user. *XStream HDVR* segmentation allows the setting of Label Values in a Label Volume to segment out different structural components of a dataset to up to 8 different Transfer Functions (256 in the SDK), this method is unique to Fovia.

7.1 SEGMENTATION - OVERVIEW

To work with segmentation capabilities, segmentation first needs to be enabled. To do this, check the "Enable Segmentation" checkbox, located under the Segmentation Tab on the left side of the user interface (Figure 7.1). Upon checking this button, an alert box states, "A memory overhead of 50% of the dataset size will be incurred. Proceed?" For segmentation, our software allocates this memory for a second volume. This volume is called the Label Volume and that Label Volume stores information per voxel on what Transfer Function to display, with up to 256 TF's to choose from. Because of this unique method Segmentation consumes an amount of memory equivalent to 50% of the dataset size, however this technique has a number of functional advantages (explained in the paragraphs below). In addition, the implementation of the label volume within the Fovia *XStream HDVR* rendering engine is highly optimized to incur minimal rendering slowdown even if using up to 8 TF's. This performance is highly unique to Fovia. To proceed select yes in the alert box.

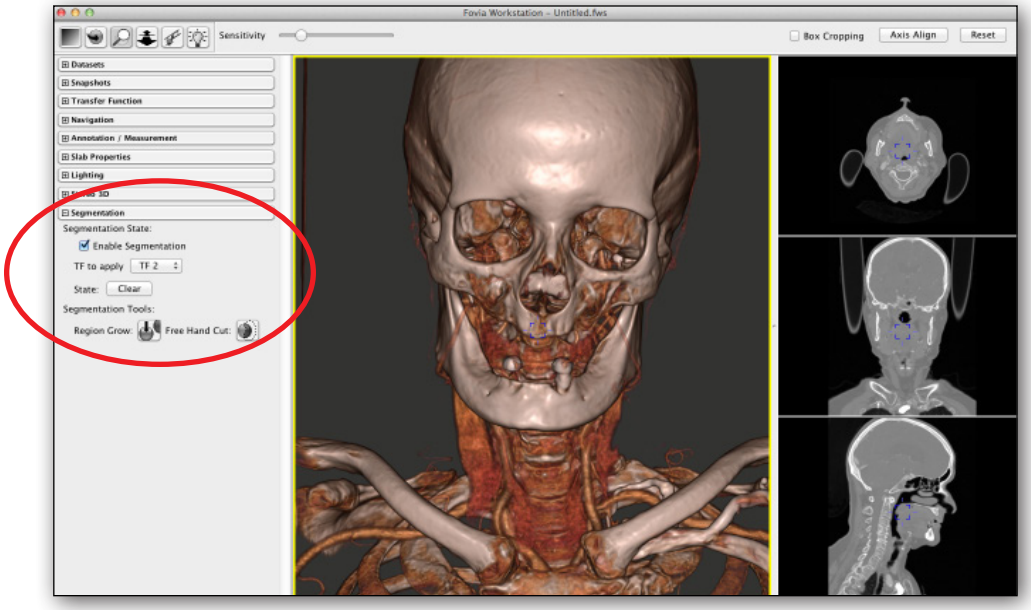

FIGURE 7.1: Location of Segmentation Tab

One of the advantages of using the Label Volume that is internal to the engine rather than altering the original volume data is that labels are assigned in the ancillary label volume during segmentation. This allows the original data to be displayed (e,g, Axial 2D slices without segmentation). Additionally this label volume allows up to 8 (or 256 with an alternate build version) different labels to occur within each volume. Note: Even though you have 8 specified "types" of regions (or 256 types), you can have an unlimited number of regions with these "types" available for use.

The Fovia engine uses the label volume to denote which of 8 (or 256) different Transfer Functions to utilize when rendering each selected area. Because the rendering engine utilizes the label to decide which TF to render, you can create wildly imaginative visualizations. For example one region can be showing skin, another translucent skin and muscle, and yet another opaque bones with no skin or muscle visible and finally something else with an X-Ray type appearance. This flexibility is unique to Fovia *XStream HDVR*.

7.2 SEGMENTATION - TF SELECTION

Once segmentation has been enabled, an active transfer function must be set. The selected transfer function is what voxels will be associated with when segmentation is actually applied to the dataset. You can select your desired Transfer Function in the drop-down menu within the Segmentation Tab, labeled "TF to apply" (Figure 7.2).

TF2-TF8 transfer functions are initially represented by flat colors that can be changed by the user. Once a transfer function is applied to a segment, the color and opacity of the set transfer functions will appear on the selected segment(s) (Figure 7.3). The segment can be manipulated by adjusting the corresponding transfer function(s).

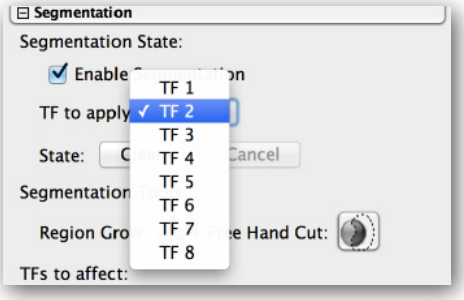

FIGURE 7.2: TF Selection

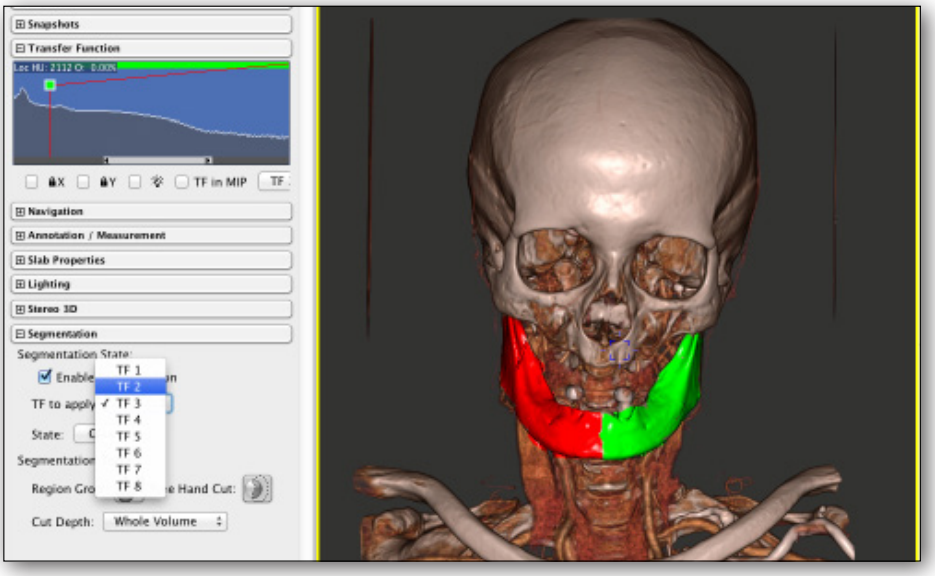

FIGURE 7.3: Mandible segment with two different transfer functions applied

The segmentation tools can be used in a variety of ways. For example, an empty Transfer Function (Figure 7.4) can be used to completely remove segments. This capability is referred to as the "virtual scalpel." To use this function, delete the range in any transfer function by clicking anywhere in the scalar field and pressing delete, and that will be your "virtual scalpel TF" (make sure you select the corresponding TF in your segmentation tab (Figure 7.5)).

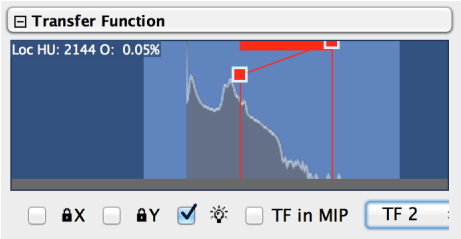

FIGURE 7.4: Transfer Function with a range

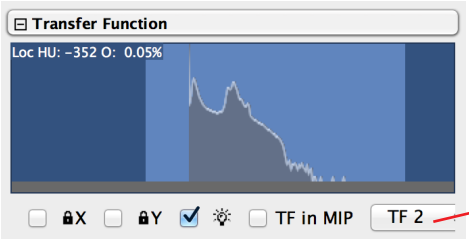

FIGURE 7.5: Transfer Function with range deleted

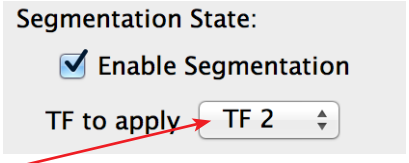

FIGURE 7.6: Corresponding TF selected

As a transfer function is edited, any segments applied to that transfer function will change (Figures 7.7-7.10).

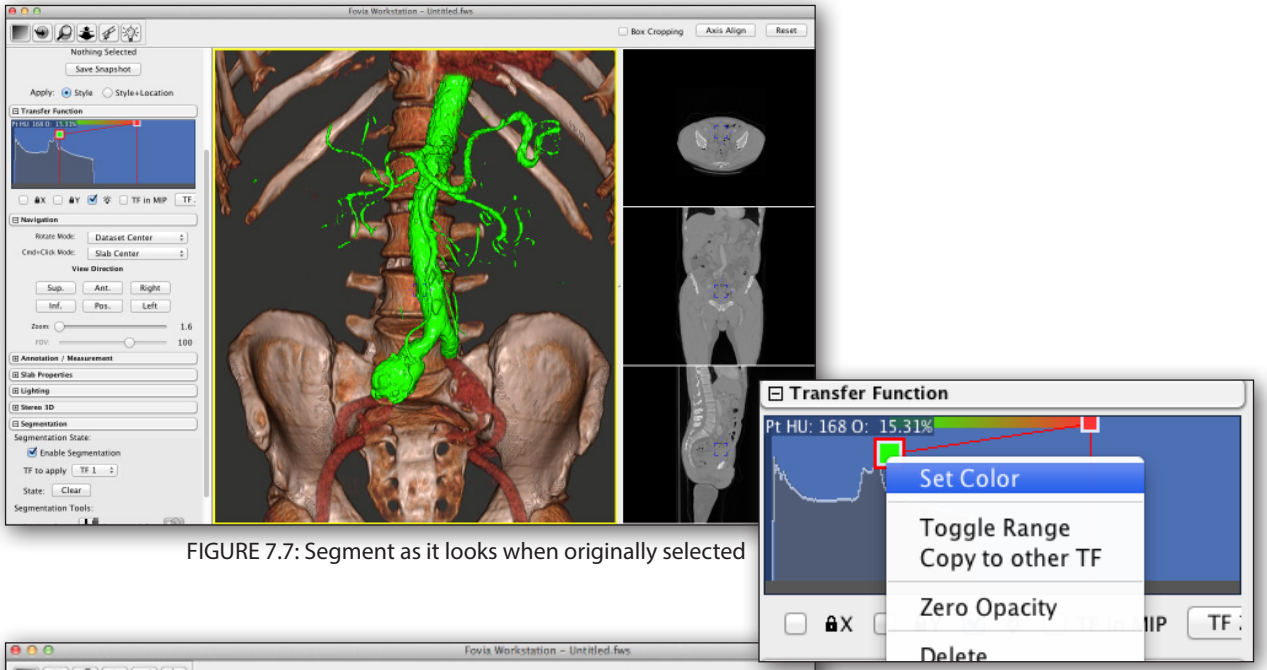

FIGURE 7.7: Segment as it looks when originally selected

FIGURE 7.8: Changing the color of the segment in the transfer function tab

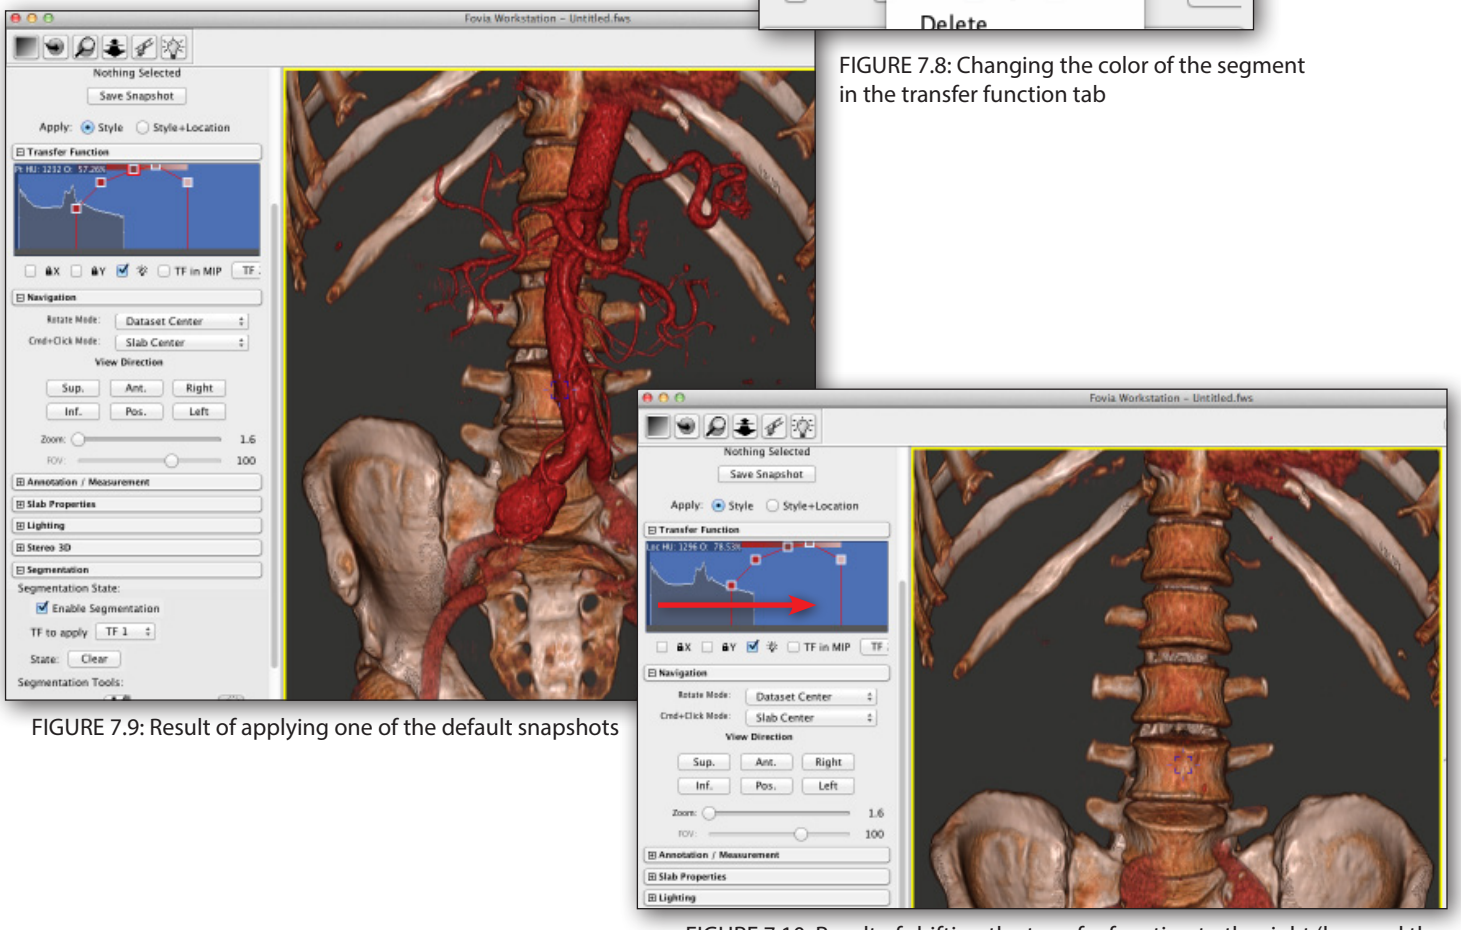

FIGURE 7.9: Result of applying one of the default snapshots

FIGURE 7.10: Result of shifting the transfer function to the right (lowered the opacity of the vessels)

A snapshot can be applied to a Transfer Function, and therefore to a segment. While the desired Transfer Function is selected, double-click on a snapshot in the Snapshots Tab. (Make sure "Style" is selected, rather than "Style + Location." Figure 7.11.) This applies the snapshot to that particular Transfer Function, as well as any segments to which that transfer function is applied.

Note that all segmentation procedures are performed on the visible scalar field range and do not affect areas where scalar field values are not visible due to the transfer function setting. For example, segmenting the skull does not affect brain tissue if the brain tissue is not visible at the moment of segmentation.

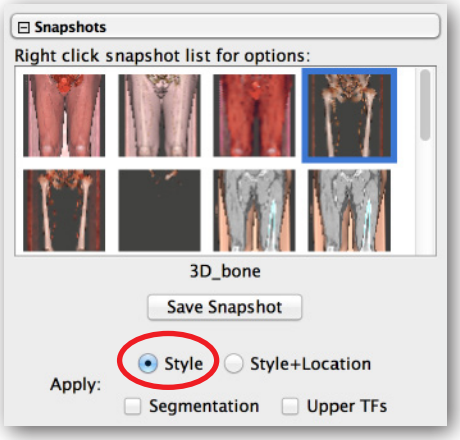

FIGURE 7.11: Settings when applying snapshot

The Segmentation Tools provide two selection options: Region Grow or Free-Hand Cut (Figure 7.12). These are defined below and elaborated in the next section.

FREE-HAND CUT (FHC): A polygonal region is drawn on the screen that includes all visible voxels contained within that region as it is projected through the volume. This segmentation method should only be used with parallel rendering modes.

REGION GROW: This segmentation method automatically segments a material or tissue type based on parameters set in the Region Grow SDK methods.

FIGURE 7.12: Region Grow and Free-Hand Cut modes

7.3 SEGMENTATION - FREE-HAND CUT

The FHC mode allows the user to have full manual control over what is segmented. To use, the hold down the left mouse button. A ray will be cast from the first click point to wherever the mouse is dragged. As the mouse is pressed and dragged, the rays combine to become a volume. All voxels within that volume will be a part of the segment when the mouse is released (Figures 7.13 and 7.14).

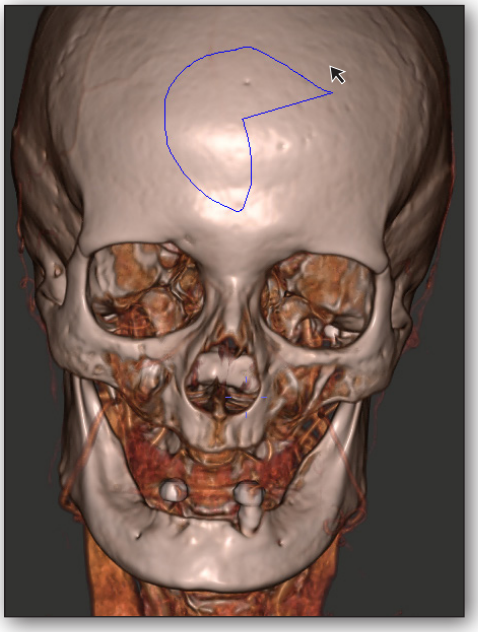

FIGURE 7.13: Drawing the segmentation rays

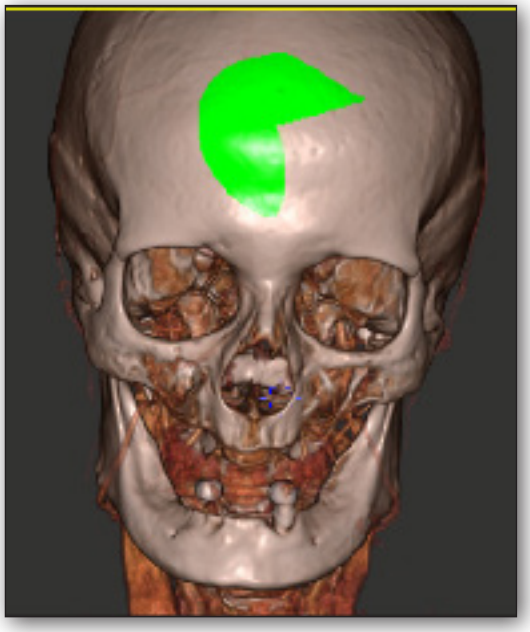

FIGURE 7.14: Filled-in segment

The user also has control over the depth of the segmentation with the "Cut Depth" dropdown menu (Figure 7.15). With the "Whole Volume" option, rays are cast through the entire volume. With "One Layer" rays are cast only through the first layer of tissue. This means that as soon as the ray exits into a totally transparent area, it will stop. This tool is useful for reclassification of films, such as skin or the colon, without affecting the tissues behind. In the "Surface" option, the ray penetrates the first 4 voxels and stops. It is useful for marking areas as if one were using a pen.

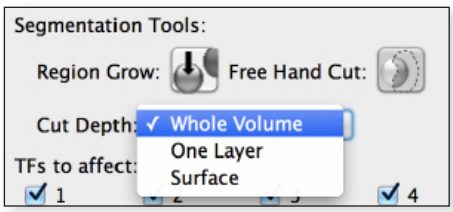

FIGURE 7.15: Segmentation depth

7.4 SEGMENTATION - REGION GROW

THRESHOLD AND RADIUS SLIDERS

The Region Grow technique supports several different methods for automatically segmenting volume material based on voxel criteria. This can highlight a structural feature such as a specific tissue or material type, or hide such structures to improve the visibility of other dataset components. In general, the region growth techniques work by first selecting an individual 'seed' voxel in the dataset, and then recursively sampling each of the voxel's 27 neighbor voxels to test if they meet the segmentation criteria. The region propagates outward from the seed voxel until adjacent voxels no longer meet the segmentation criteria.

Several options determine how the segmentation will grow, the first two being the threshold and the radius sliders. The threshold slider determines how rigorous the boundary selection will be with regard to connectivity. Moving the slider to the right will mean less connected parts can be more successfully included. (Figures 7.16 and 7.17).

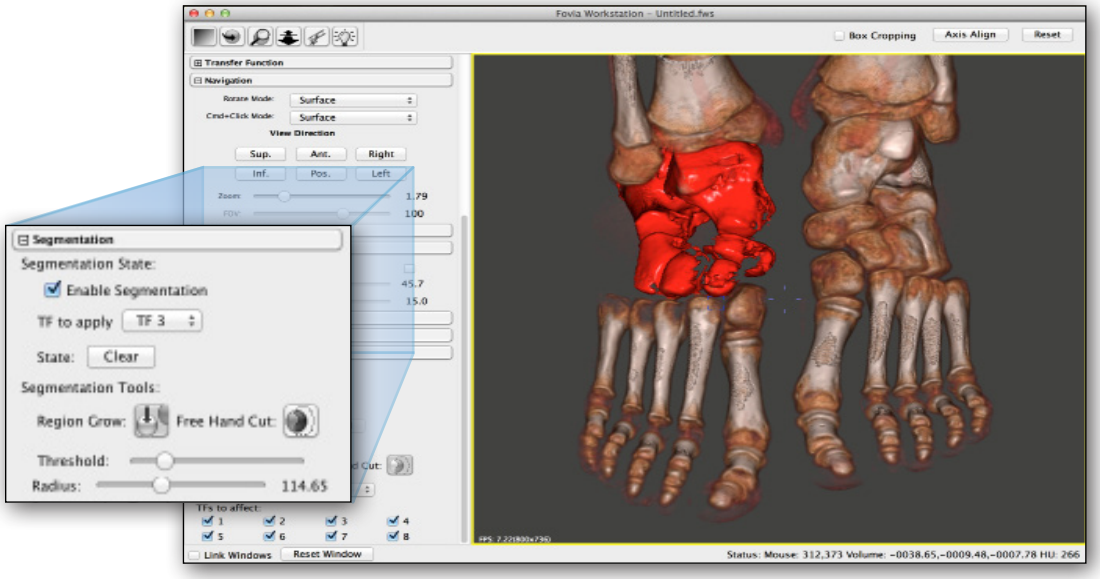

FIGURE 7.16: Low threshold

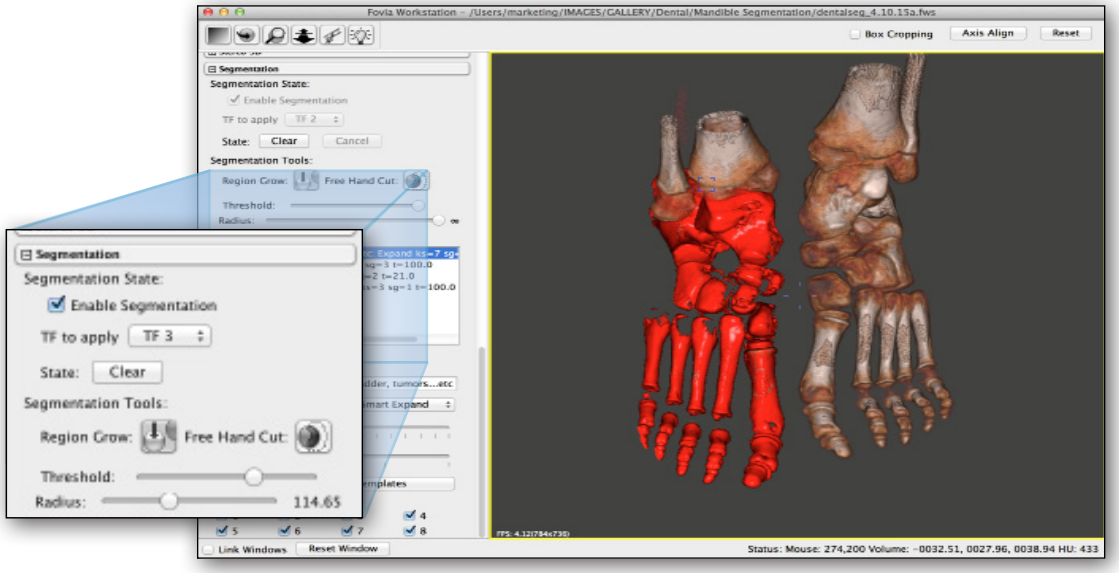

FIGURE 7.17: High threshold (same radius)

Moving the thresholding slider to the left will mean parts that are not connected will be included less. (Figures 7.18 and 7.19). This is helpful, for example, when trying to segment contrasted vessels from bone.

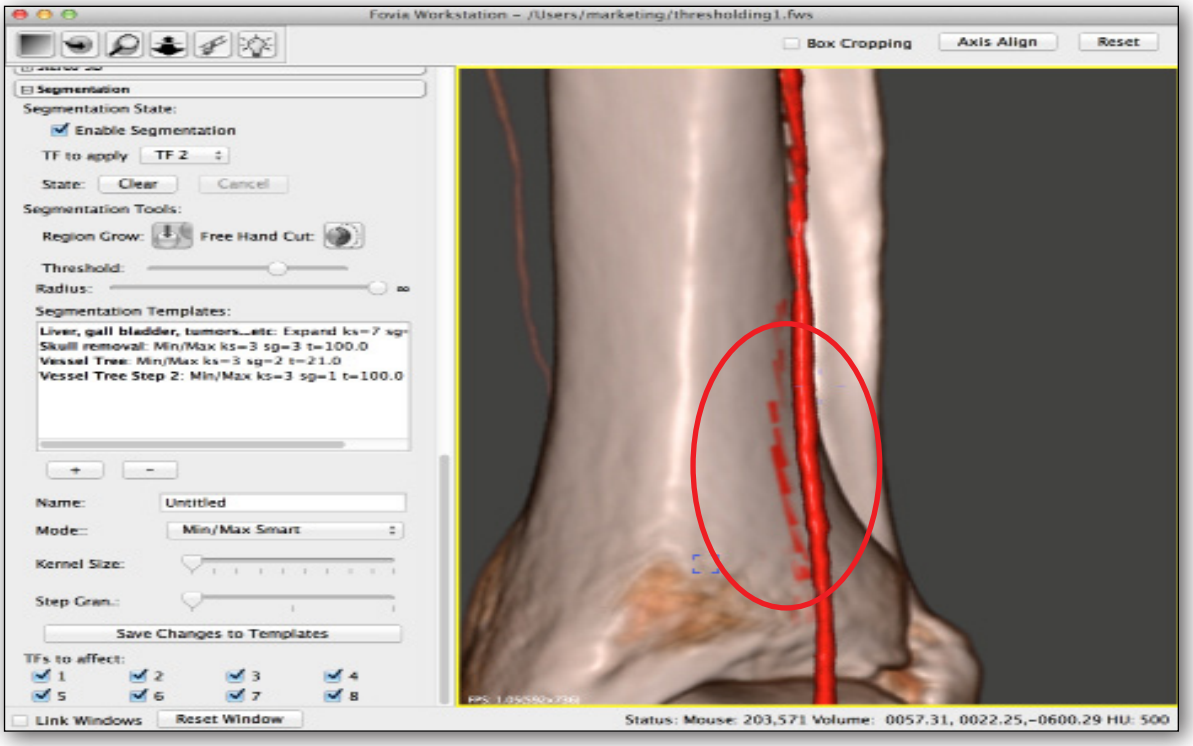

FIGURE 7.18: High threshold, includes unwanted bone

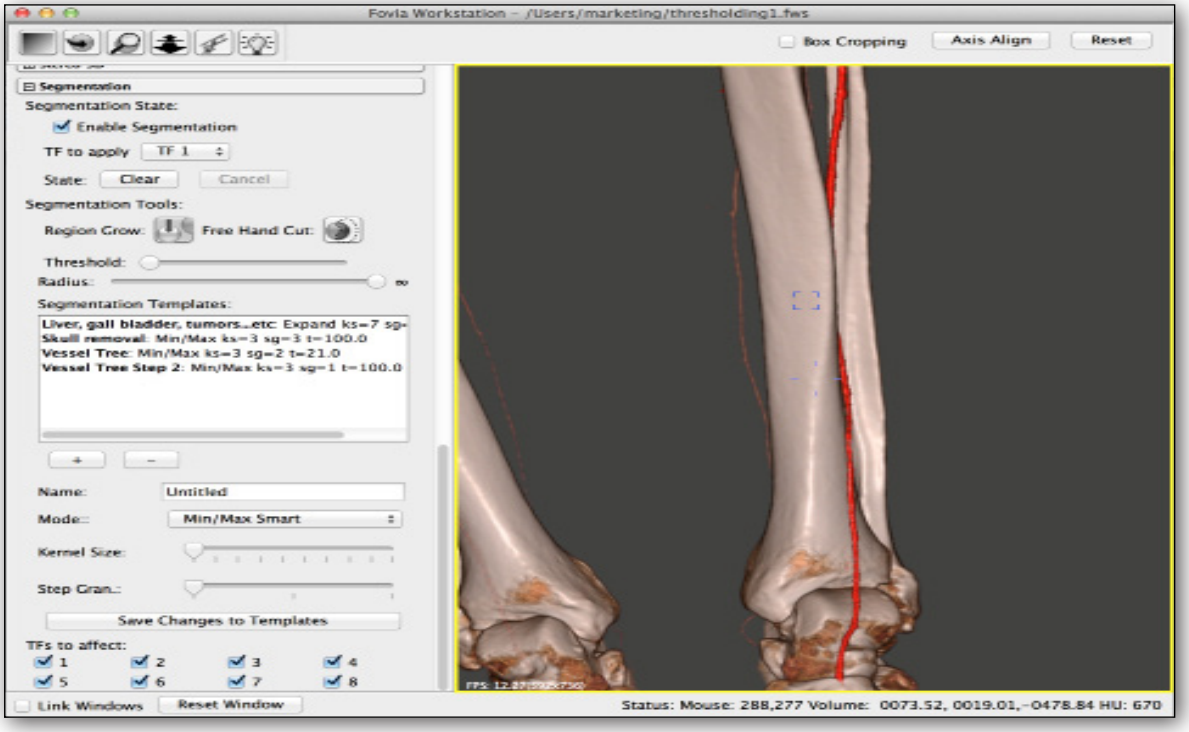

FIGURE 7.19: Low threshold excludes unwanted bone

SEGMENTATION

REGION GROW

The radius slider determines how large the radius of the growing segment will be. Moving the slider to the right includes a wider area in the growing region (Figures 7.20 and 7.21).

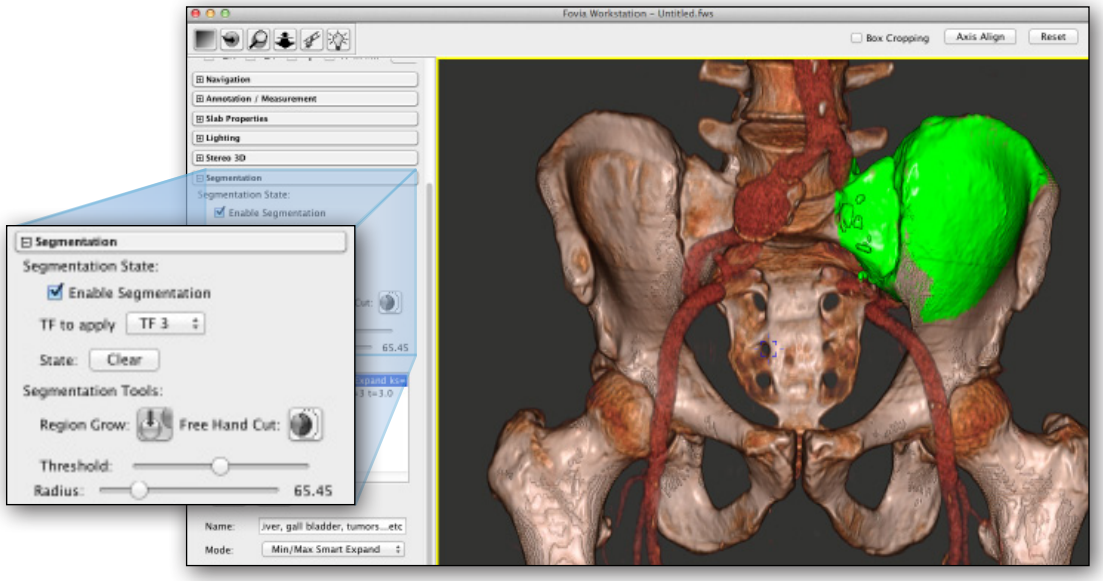

FIGURE 7.20: Small radius

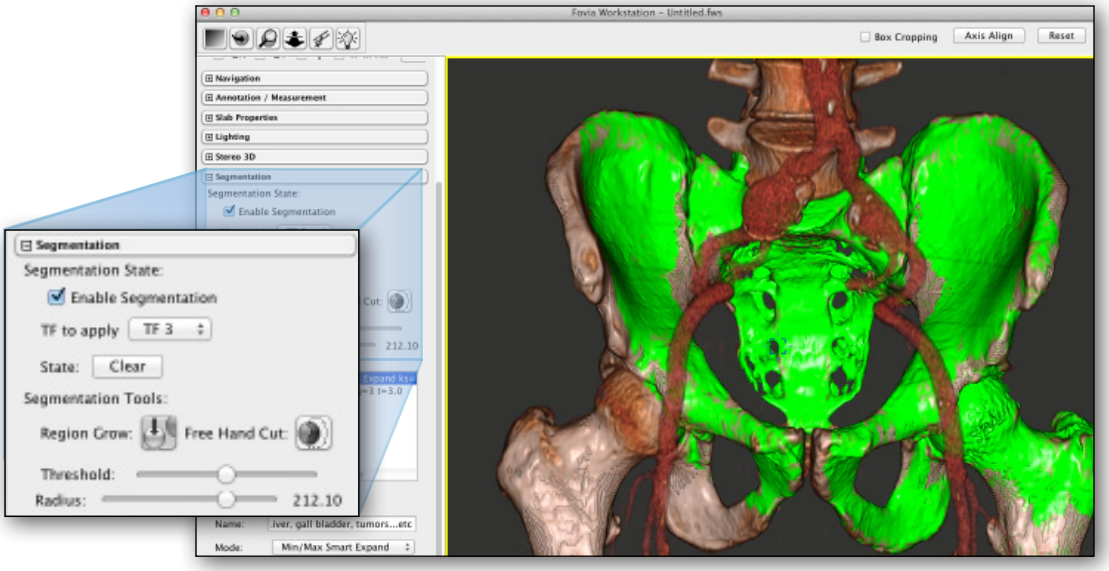

FIGURE 7.21: Large radius (same threshold)

REGION GROW ADVANCED SETTINGS

The Mode dropdown menu dictates how the software will determine the connectivity between voxels. The options include:

Min/Max: Uses a range of tissue values to determine connectivity. If the sampled point is in that range, then it's connected. If not, then it will not grow there (the threshold slider has no effect on Min/Max). This setting is best used when segmenting tissue with distinct difference in density, for example bone and muscle. When segmenting tissue of similar density Min/Max Smart or Min/Max Smart expand should be used. Note: In the workstation the min/max range is chosen as the current min/max of the range in the Transfer function of the object that you have selected, in the SDK any range can be chosen.

Min/Max Smart: Uses gradient magnitude analysis of the data to determine connectivity. This setting accounts for tissue values, and in addition, uses edgefinding to segment out data. This setting is appropriate for segmenting out tissue of simliar density.

Min/Max Smart Expand: Similar to Min/Max Smart, but utilizes special processing on the boundary of the final segmentation to get a more complete sampling of complex tissues. This setting is appropriate for segmenting out tissue of simliar density.

The last Segmentation customizations include the Kernel Size and the Step Gran Sliders:

Kernel Size: The size of the cube of data points whose gradient magnitudes are averaged for the gradient magnitude analysis. The noisier the data set, the higher the kernel size should be. However, increasing the kernel size lessens the ability to see small differences in density when looking for edges. The user needs to balance this tradeoff. A Kernel Size of 5 means a 5x5x5 voxel cube will be analyzed.

Step Granularity (Step Gran): This function is primarily adjusted for speed. It indicates how much data will be grouped together in a selection. A setting of 1 means that no grouping will occur, each voxel will be considered seprately. A setting of 2 means that every 2 voxels will be grouped and considered as one, and a setting of 3 means that every 3 voxels will be grouped and considered as one.

SEGMENTATION TEMPLATES

Segmentation templates include saved Segmentation Templates that to assist the user (Figure 7.22). The "Liver, gall bladder, tumors...etc." template is a starting point for users wanting to select soft tissue regions. The "Skull removal" template helps to select bone. Selecting a template automatically adjusts all Region Grow settings. If necessary, these settings can then be adjusted manually.

The user can also save his or her own template. Simply click on the "+" button below the list of Segmentation Templates and adjust settings as necessary. Name the template by typing in the "Name" box, also located below the Segmentation Templates.

To remove a template, select the template and click on the "-" button.

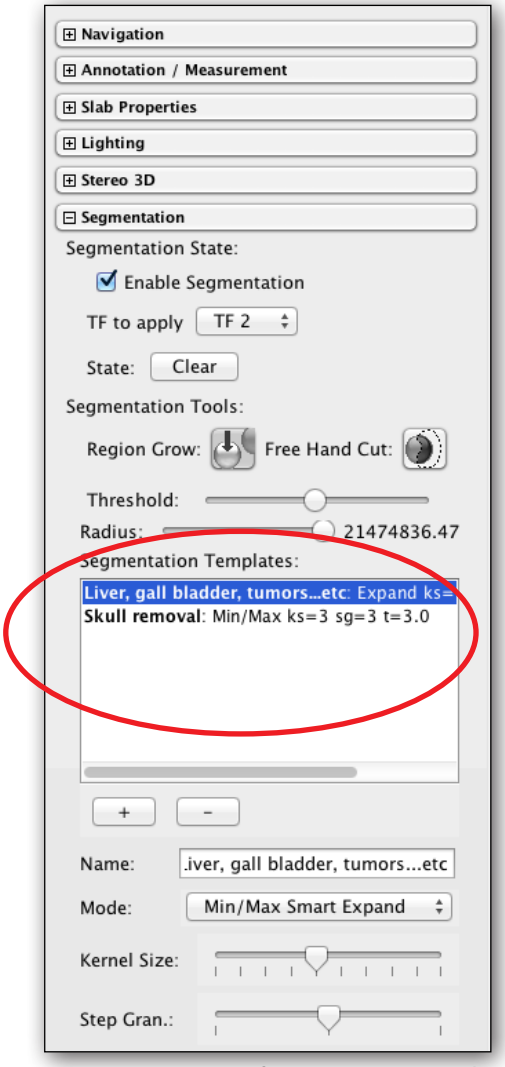

FIGURE 7.22: Location of Segmentation Templates

This section describes the *XStream HDVR* SDK features used to create, manipulate, visualize and save vertex-based mesh geometry. *XStream HDVR* provides the ability to export a watertight polygon mesh of the entire volume, or a segmented region, to a variety of data formats including STL, PLY, and OBJ. The benefit of this approach is that the editing, cleanup, and preview can happen within one application.

8.1 POLYGON FUNCTIONS - POLYGON IMPORT

8.1a IMPORTING POLYGONS

XStream HDVR supports real-time rendering and direct fusion of semi-transparent polygonal and volumetric data (without using a GPU). The Polygon Import function is located under the Annotation/Measurement Tab. Select the "Open OBJ/STL" button (Figure 8.1). A dialogue box will open to retrieve your file (Figure 8.2).

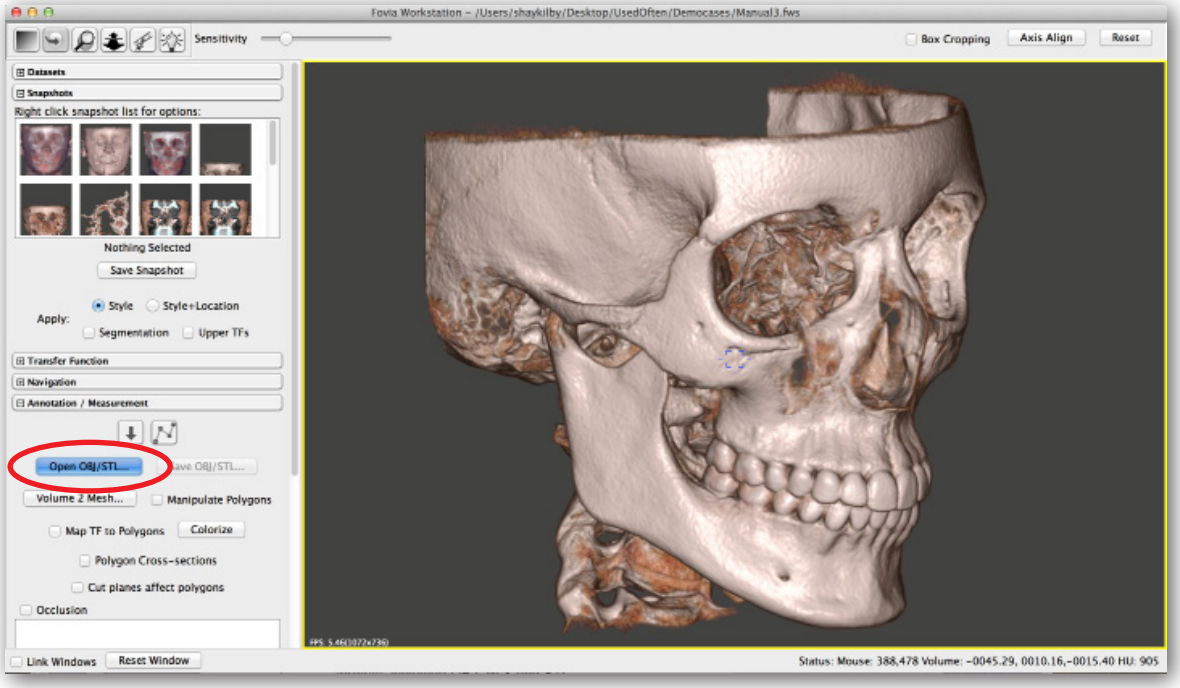

FIGURE 8.1: Polygon Import button

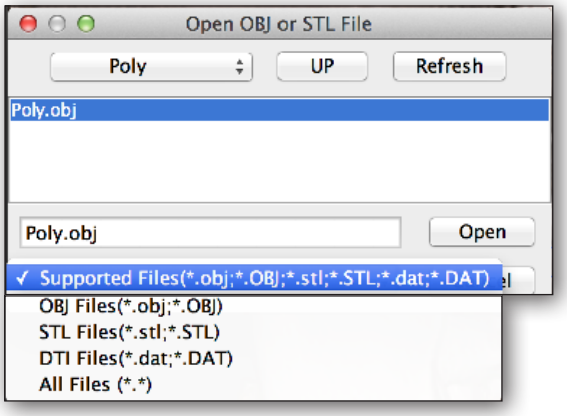

FIGURE 8.2: Configuration Dialogue-1

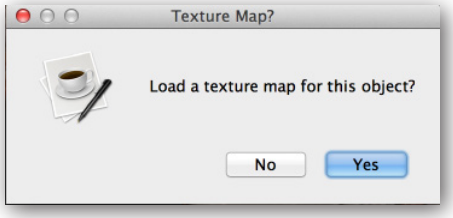

FIGURE 8.3: Configuration Dialogue-2

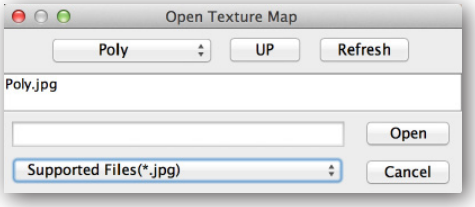

FIGURE 8.4: Configuration Dialogue-3

8.1b TEXTURE MAP

XStream HDVR supports the texture map rendering of polygon data from image files, including facial textures and live render-to-texture sources. The texture map is applied through vertex colorization. Once the texture map is imported (see previous section for instructions), select the checkbox "Manipulate Polygons" to move the polygon without moving the volume data (Figure 8.5). Uncheck this box to move the polygons and the volume together. Multiple polygons (along with their texture maps) can be imported in to the "scene."

If the texture mapped object was created with lighting added, you can turn off the lighting by unchecking the box "Enable Lighting" (Figure 8.6). You can rename objects using the "Name" dialogue box (Figure 8.6).

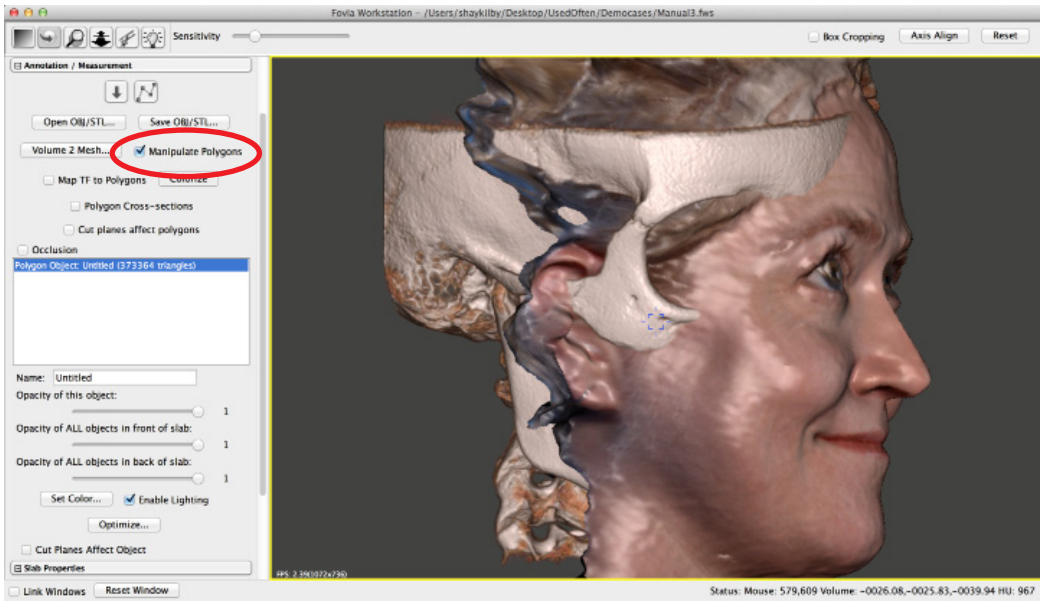

FIGURE 8.5: Manipulating Texture Map

Within the Annotation & Measurement Tab

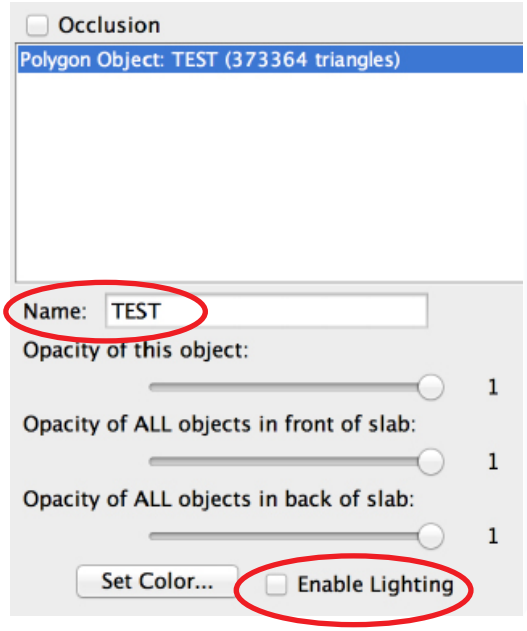

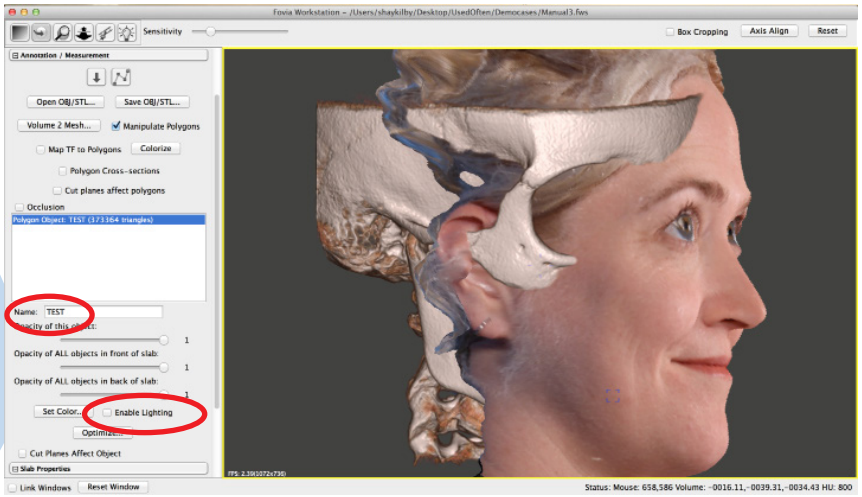

FIGURE 8.6: Polygon Naming & Lighting

POLYGON FUNCTIONS

POLYGON IMPORT

A variety of additional functions affect the poygon (Figures 8.7 and Figure 8.8). The user can adjust the opacity of an object, the opacity of all objects in front of the slab, or the opacity of all objects in back of the slab. In addition you can set the color of the polygon object. Select "Cut planes affect polygons" to adjust so that cut planes affect an object (Figure 8.7).

To get the most accurate colorization, map the Transfer Function to polygons by selecting "Map TF to Polygons." Select the "Colorize" button for the fastest colorization using the Transfer Function (Figure 8.8). The function "Polygon Cross-sections" provides a cross-sectional outline of the polygon. Optimization will be covered in section 1.3.

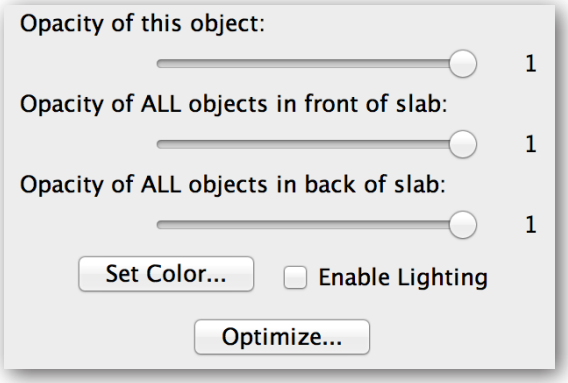

FIGURE 8.7: Additional Polygon Functions-1

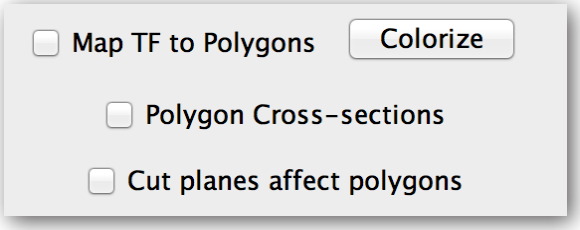

FIGURE 8.8: Additional Polygon Functions-2

8.2 POLYGON FUNCTIONS - POLYGON EXPORT

8.2a - SEGMENTATION

The entire data set can be exported to a polygon object, or segmentation can be used to export a part or parts of the dataset. As mentioned in the Segmentation Chapter (p. 39), to work with segmentation capabilities, segmentation first needs to be enabled. To do this, check the "Enable Segmentation" checkbox, located under the Segmentation Tab on the left side of the user interface (Figure 8.9).

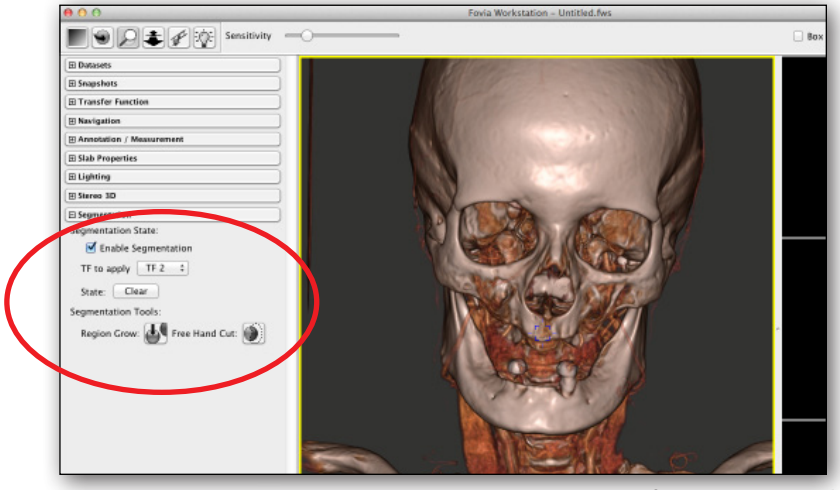

FIGURE 8.9: Location of Segmentation Tab

Datasets can be segmented to eight different transfer functions (256 in the *XStream HDVR* SDK), enabling the user to simultaneously create 8 separate polygon objects at once. Once an area of the data set is segmented out, it is ready to be exported as a polygon. (The next section will cover polygon export). For example, in Figure 8.10 the mandible can be segmented (and thus exported) separately from the maxilla and the rest of the cranium.

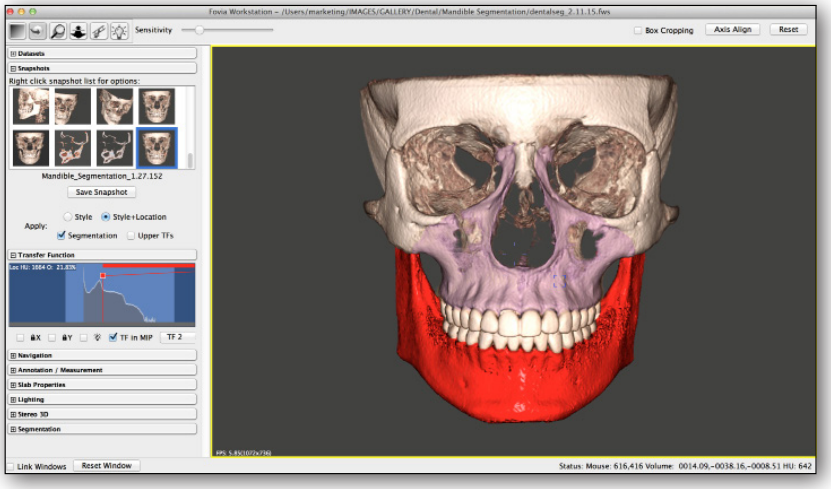

FIGURE 8.10: Mandible segment with two different transfer functions applied

8.2b POLYGON EDIT

To begin the polygon export, select the "Volume 2 Mesh" button under the Annotation/Measurement Tab (Figure 8.11). Once this button is selected, a polygon export window will open.

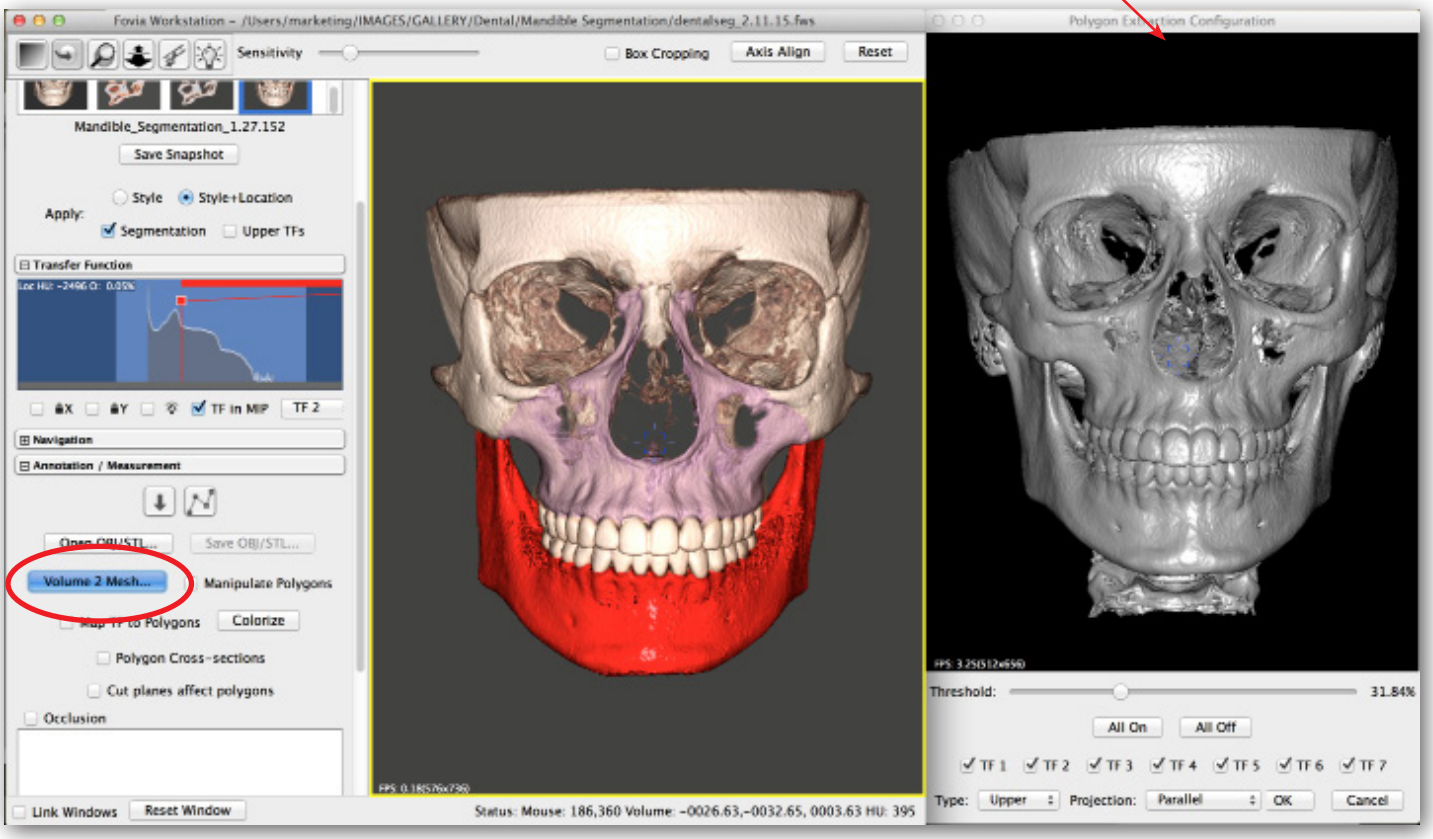

FIGURE 8.11: Polygon Export button

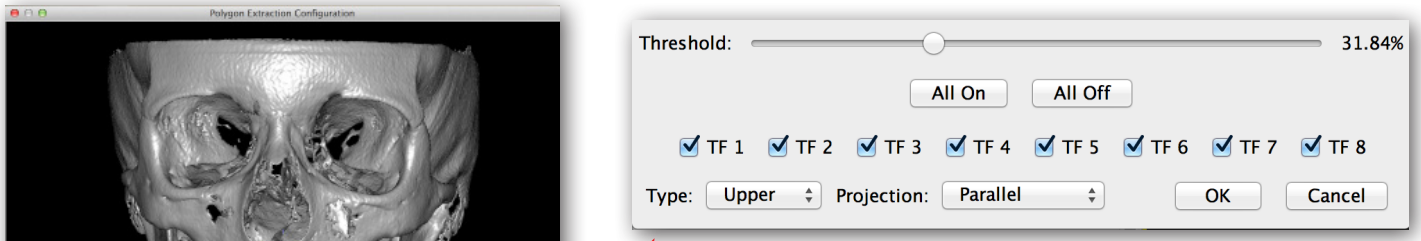

FIGURE 8.12: Polygon Export Window Functions

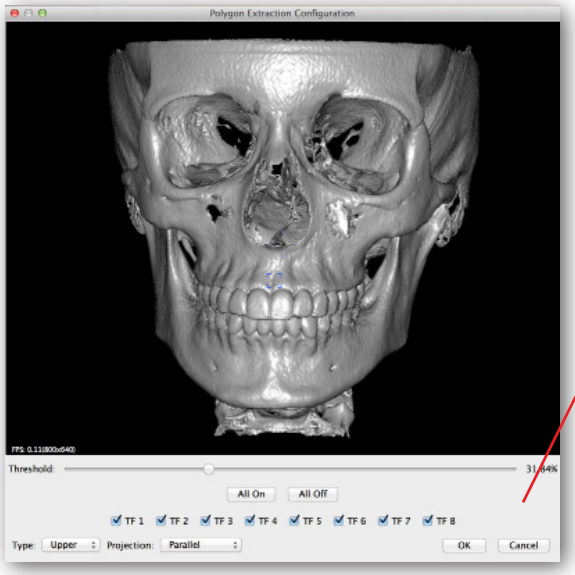

FIGURE 8.13: Polygon Export Window

The polygon export window has a variety of settings. The user can toggle individual Transfer Functions on and off, adjust thresholding of data, change the projection mode, and examine both the upper and lower limits of the thresholding.

8.3 POLYGON FUNCTIONS - POLYGON EDITING

The exported polygons can be edited in real-time. For example, if a desired portion of the mandible is missed during segmentation, the user can edit the dataset to add the omitted portion into the polygon window. See Figures 8.14-8.17 (transfer Function 1 has been turned off leaving us only the mandible, the maxilla, and the teeth in this example).

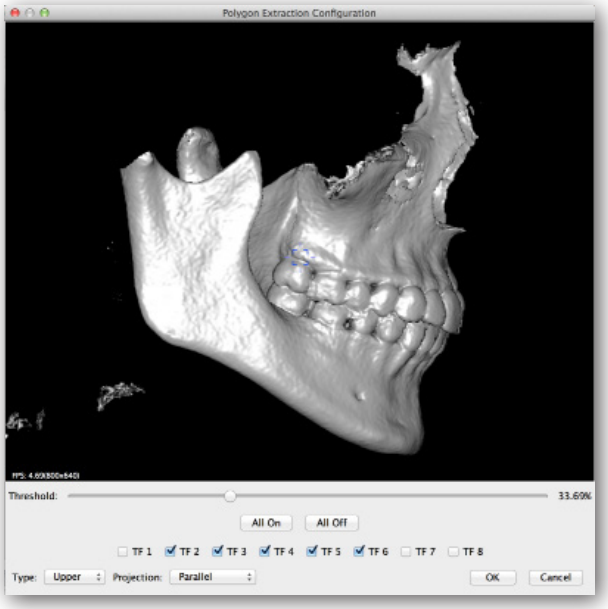

FIGURE 8.14: Stylus Process missing from Mandible Export

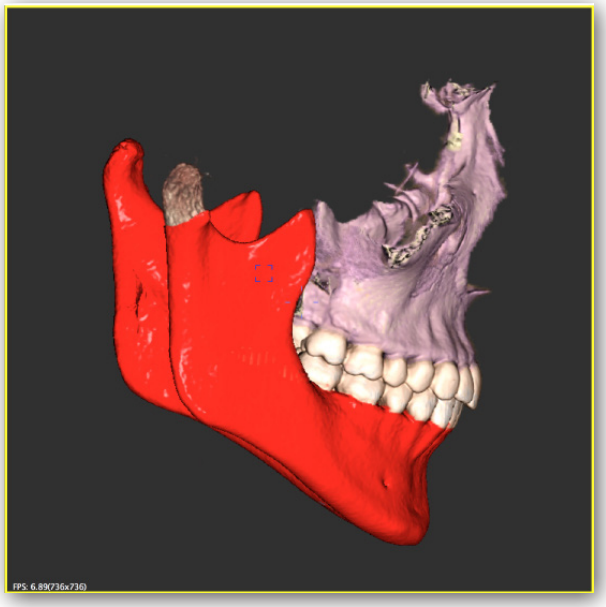

FIGURE 8.15: Stylus Process missing from Mandible Segmentation

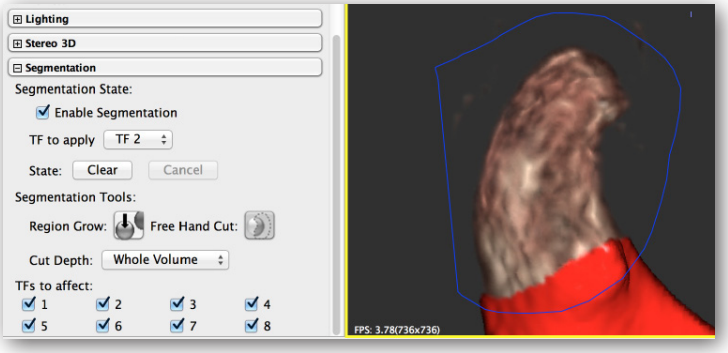

FIGURE 8.16: Select Stylus Process on Mandible Segmentation

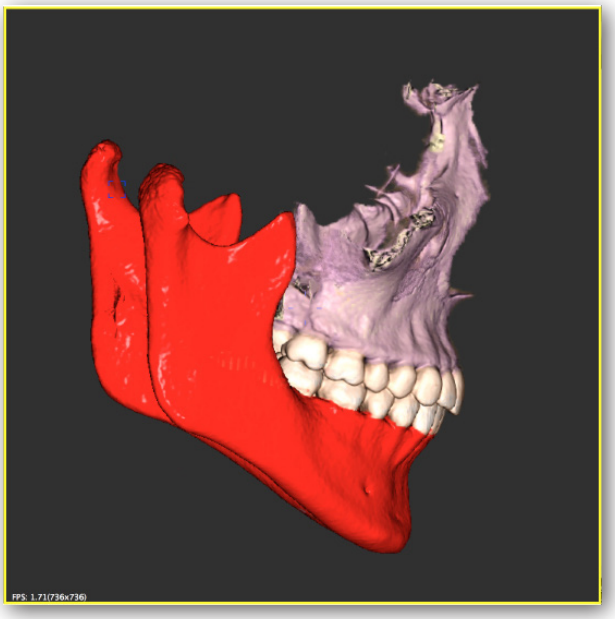

FIGURE 8.17: Corrected Mandible Segmentation

The exported mandible is now complete. It can be added to the volume rendering scene, and is ready for optimization or colorization. (Optimization is covered in the next section.)

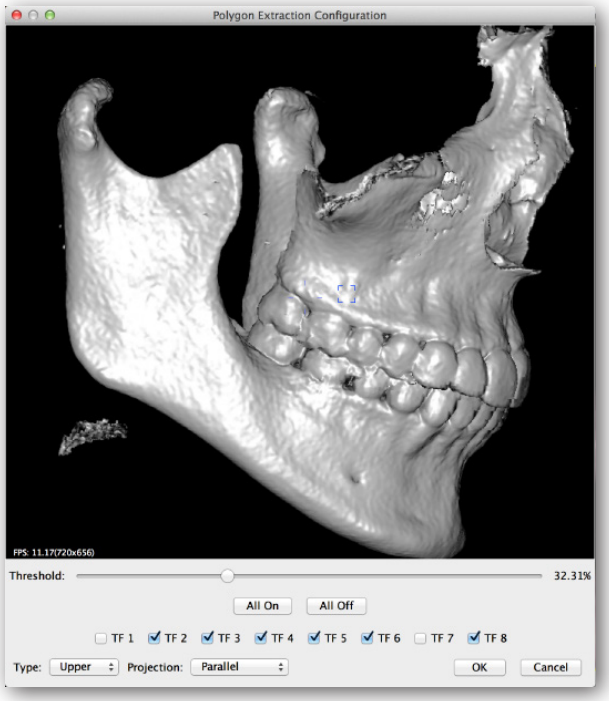

FIGURE 8.18: Corrected Mandible Export

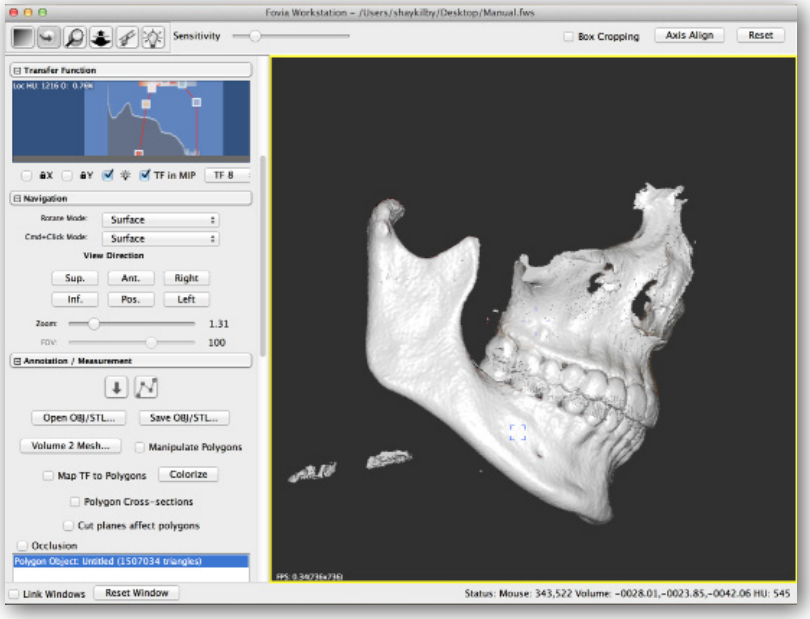

FIGURE 8.19: Polygon + Volume

8.4 POLYGON FUNCTIONS - COLORIZATION

Once the polygon has been dropped into the rendered scene it is ready for colorization or optimization. (Optimization will be covered in the next section.) The colorization button is in the Annotation/Measurement Tab (Figure 8.20). The polygons are colorized according to colors assigned to the transfer function of the underlying volume. Once colorized, the polygon object can be saved out as a .PLY file. Non-colorized polygon objects can be saved in .OBJ or .STL format.

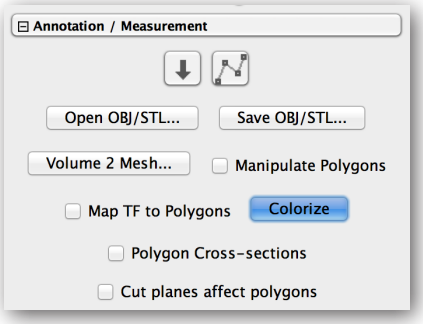

FIGURE 8.20:
Colorization Function

In general, Fovia recommends that colorization be applied before optimization in order to keep the color generation true to the initial vertex positions. A mesh optimization (discussed in the next section) changes vertex positions. However, applying decimation after colorization will also alter the appearance of the mesh as well as polygons and vertices will be moved during the optimization. The user should experiment with the order and the settings in order to determine what works best for the specific output required.

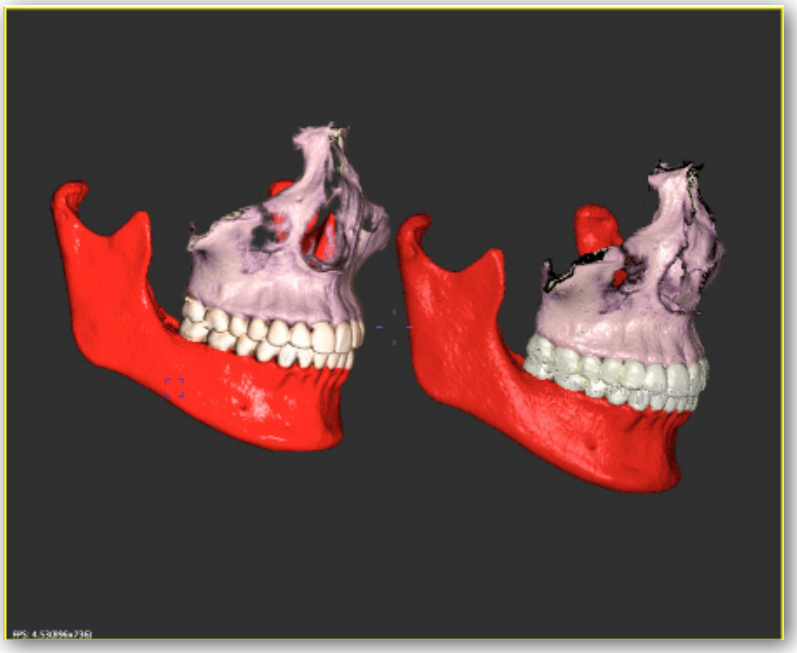

FIGURE 8.21: Volume Together with colorized polygon object

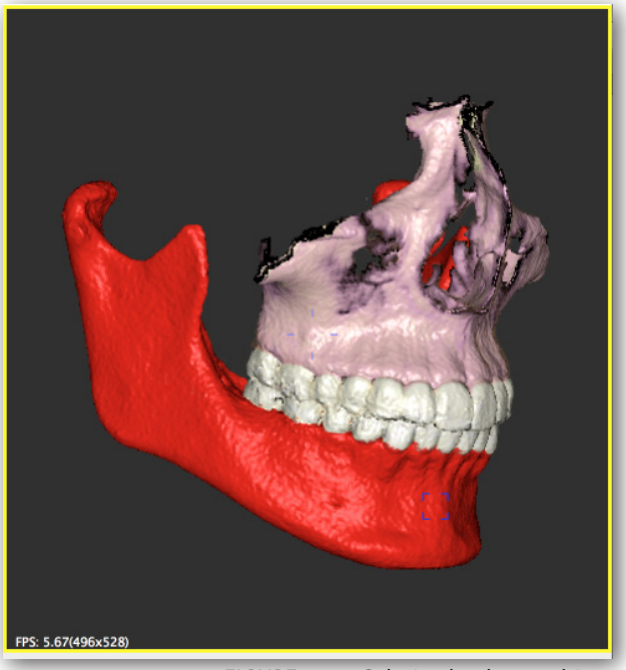

FIGURE 8.22: Colorized polygon object

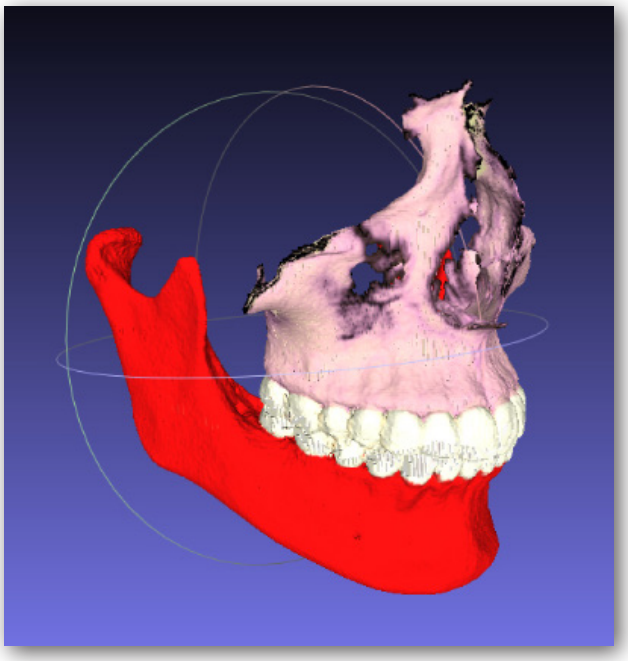

FIGURE 8.23: Exported .PLY

8.5 POLYGON FUNCTIONS - OPTIMIZATION

8.5a OPTIMIZATION FUNCTIONS

Once the polygon has been dropped into the rendered scene, it is ready for optimization. The Optimize button is located near the bottom of the Annotation/Measurement Tab (Figure 8.24). This opens the Triangle Mesh Optimization dialogue box that includes the following functions: decimation, removal of sets based on connectivity, and smoothing with a variety of sub-functions (Figure 8.25).

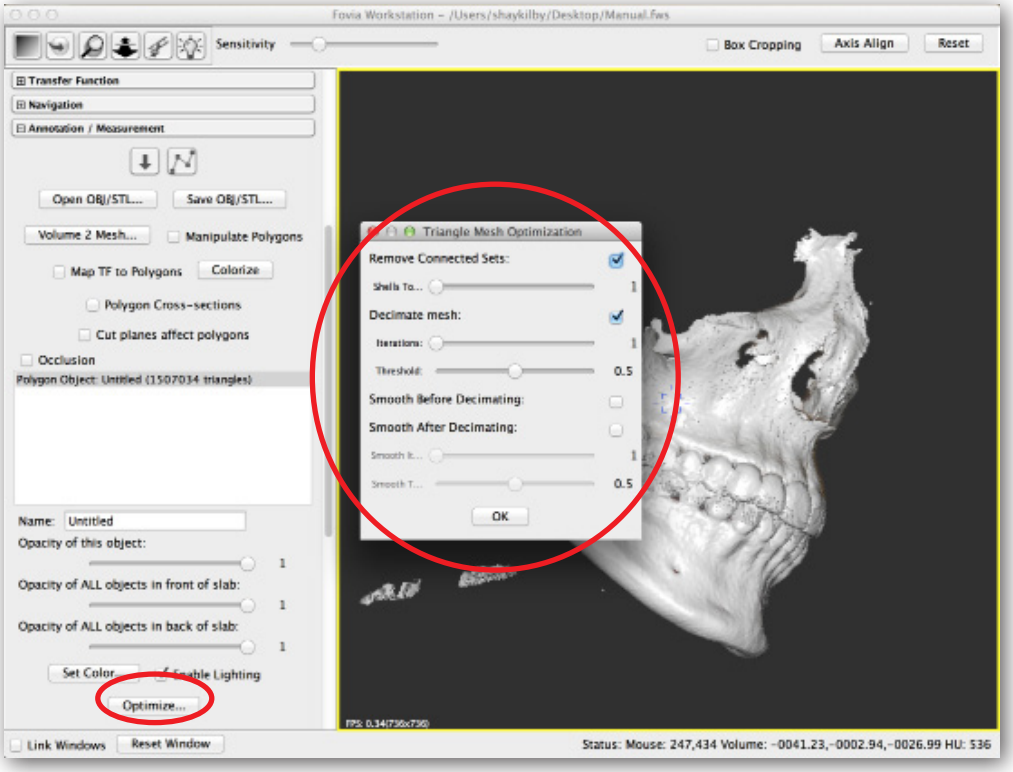

FIGURE 8.24: Polygon Optimization

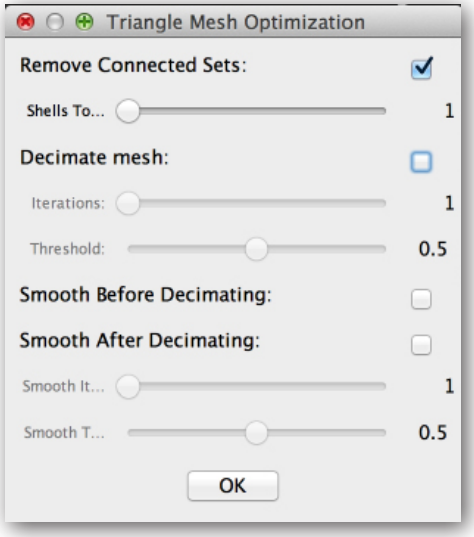

FIGURE 8.25: Polygon Optimization dialogue box

8.5b CONNECTED SETS

The first function in the optimization dialogue box, Remove Connected Sets, lets the user limit the number of connected sets that are part of polygon object (Figures 8.26-8.28).

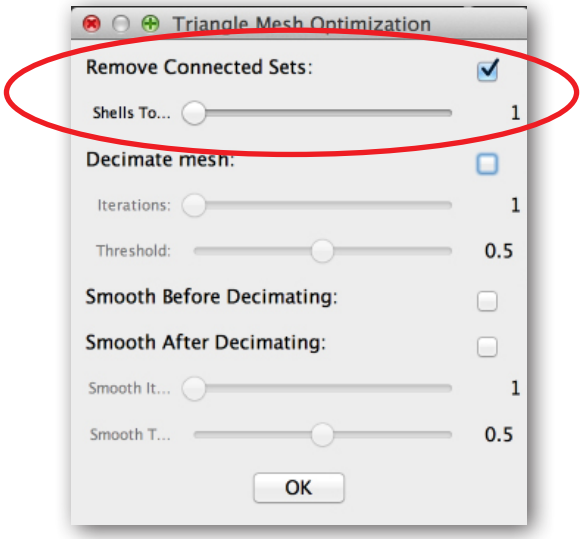

FIGURE 8.26: Remove Connected Sets

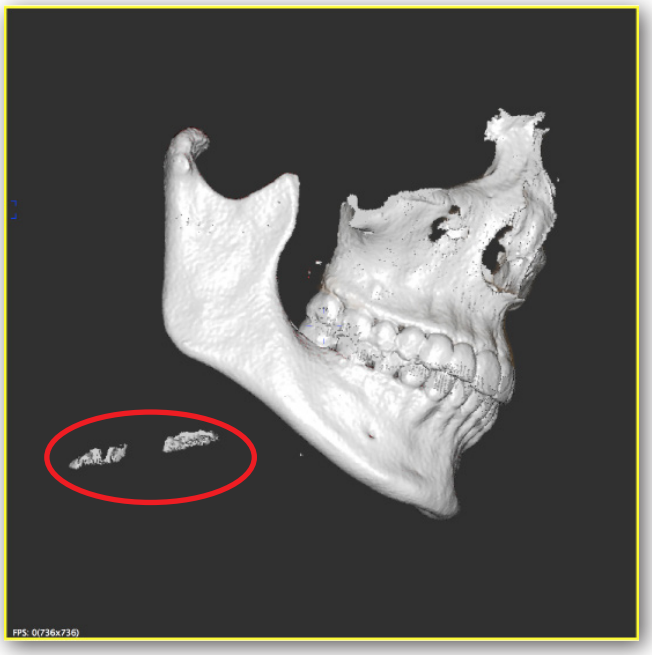

FIGURE 8.27: Removing excess polygons via the Connected Sets function

Connected Sets:

When converting a mesh to polygon geometry, there may be a significant amount of very small isolated bits of geometry scattered throughout the generated mesh data. This can result from noise in the original dataset, and/or transfer function settings that result in many small bits of matter being visualized in addition to the primary material or region of interest. Each of these disconnected bits of geometry is a single connected set. By optimizing the mesh to limit the number of connected sets that are part of the output mesh data, this extra geometric noise can be removed. Connected sets are preserved in the order of largest to smallest by polygon count. Figure 8.28 represents the polygon after the Connected Set function has been applied.

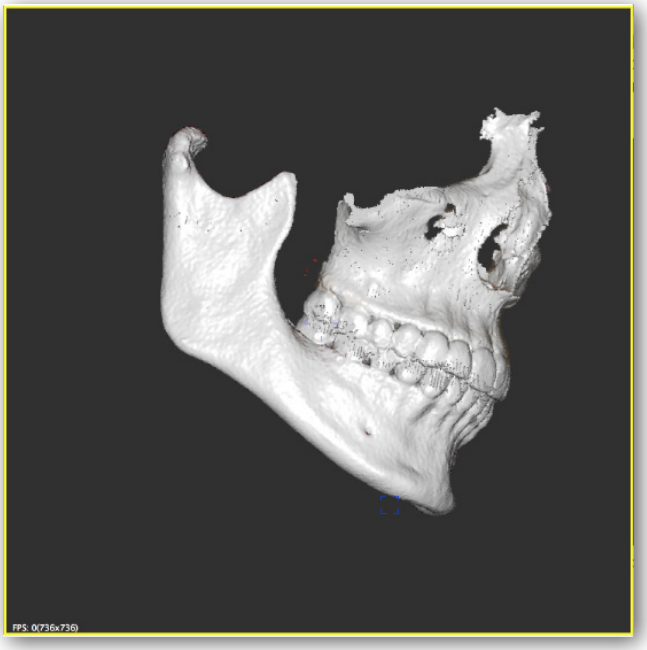

FIGURE 8.28: Excess polygons removed

8.5c DECIMATION

The second function in the dialogue box is Decimate Mesh (Figure 8.27). The algorithm that converts voxel based volume data to a polygon-based mesh often produces meshes with very high polygon counts. These meshes may contain significantly more geometric data than is necessary to produce a visually accurate polygon mesh.

The mesh optimization methods supports a decimation algorithm that reduces the number of polygons comprising the mesh, while attempting to maintain the original mesh shape. Configuration parameters are available to control how aggressive the algorithm is in simplifying the mesh and the number of simplification passes the algorithm makes. You can specify iterations and the threshold. In some cases it may be desirable to make multiple passes with a low threshold, and in others a single pass with a higher threshold, depending on the type object and how the object will be used.

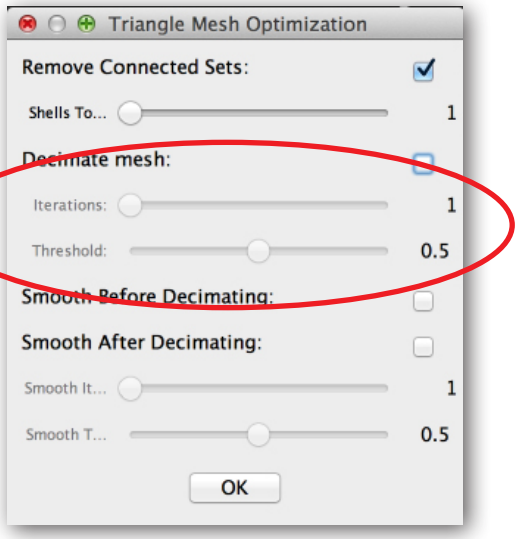

FIGURE 8.29: Decimate mesh

8.5d SMOOTHING

The next set of functions in the dialogue box is Smoothing (Figure 8.30).

The smoothing technique is Laplacian Smoothing, and this function is used to smooth out bumps and sharp edges on the mesh.. The user can smooth before decimating or after decimating. The user can set the smoothing interations as well as the smoothing thresholding. These parameters control when the smoothing is applied, how aggressive the algorithm is in applying smoothing to the mesh, and the number of smoothing passes the algorithm makes. In some cases it may be desirable to make multiple passes with a low threshold, and in others a single pass with a higher threshold.

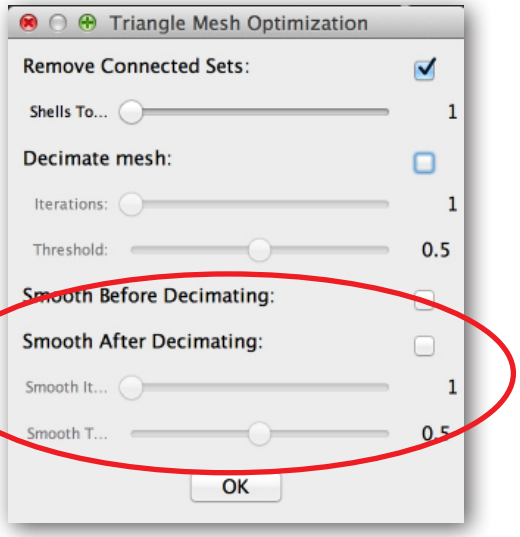

FIGURE 8.30: Smooth mesh

SLAB PROPERTIES

The Slab Properties function allows the user to manipulate a slab (thick slice) of the dataset. The volume can be clipped via a front cut plane and back cut plane, defined by the percentage into data and the slab thickness.

9.1 SLAB PROPERTIES - SLAB PROPERTIES TAB

The Slab Properties Tab, located on the left side of the user interface (Figure 9.1), has three options for the user. The first option is the "Enable Slab Navigation" button. Upon checking this button the dataset enters slab mode (Figure 9.2). The resulting dataset is defined by the front slab plane and the back slab plane. Note that when the rotate tool is used while Slab Navigation is enabled, the the data rotates through the slab boundaries.

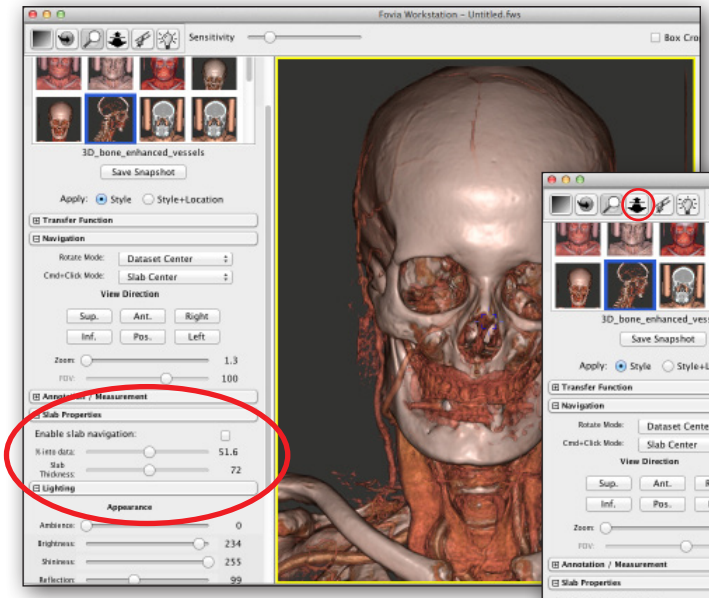

FIGURE 9.1: Location of Slab Properties Tab

FIGURE 9.2: "Enable Slab Navigation" button checked, location of Pan Slab navigation button

The second option in the Slab Properties Tab is the "% into data" slider. This moves the front slab plane and the back slab plane simultaneously through the dataset, in the same direction. The Pan Slab button in the top menu bar also provides this functionality (Figure 8.32), as does pressing and dragging the middle mouse button.

The third option is the "Slab Thickness" slider. This moves the front slab plane and the back slab plane simultaneously as well, but this time in opposite directions, thereby increasing or decreasing the thickness of the slice.

9.2 SLAB PROPERTIES - CUT PLANES

More options involving the slab cut planes are available in the right-click dropdown menu within the viewing screen (Figure 9.3), listed under "Cut Planes":

- Lock Slab to Cut Plane
  Shortcut: C
  This locks both the front and the back cut planes in place. The user can then manipulate the resulting slab as a whole with the rotate tool. Each time a cut plane is locked, it is automatically saved. The saved cut planes can be accessed under the Toggle Cut Planes option. (Note: If the following alert box appears, "Front clipping plane must be enabled to lock a cut plane," check the "Enable slab navigation" button, located in the Slab Properties tab.)

- Lock Slab to Front Cut Plane
  Shortcut: Option + C
  This locks the front cut plane in place. Once chosen, the back cut plane will disappear. It automatically saves the front cut planes, which can be accessed under the Toggle Cut Planes option.

- Toggle Cut Planes
  Each time a cut plane is locked, it is automatically saved. The user can accessed these saved planes, thereby re-cutting the dataset, by clicking the Cut Planes listed under this menu. Clicking on a Cut Plane a second time will turn it back off.

- Clear Cut Planes
  Shortcut: Shift + C
  This removes all saved Cut Planes.

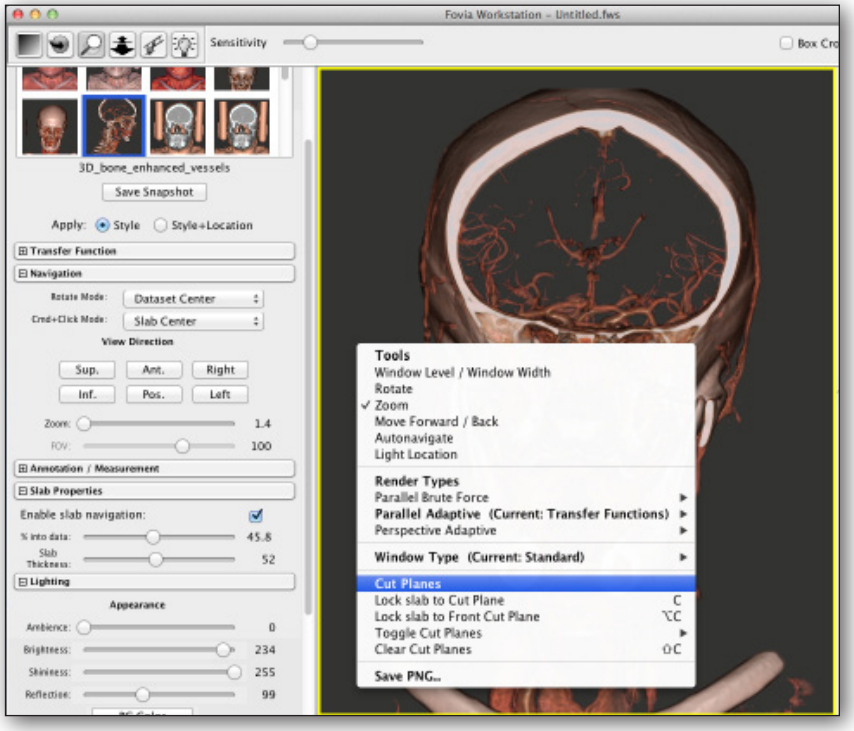

FIGURE 9.3: Cut plane options in the right-click dropdown menu

Fovia's volume rendering engine has the ability to render images and movies in resolutions up to 8192 x 8192 pixels. Options for exporting media are discussed below.

10.1 CREATING MEDIA - IMAGE EXPORT

To create an image in the workstation application, simply right-click and select "Save PNG" from the drop down menu.

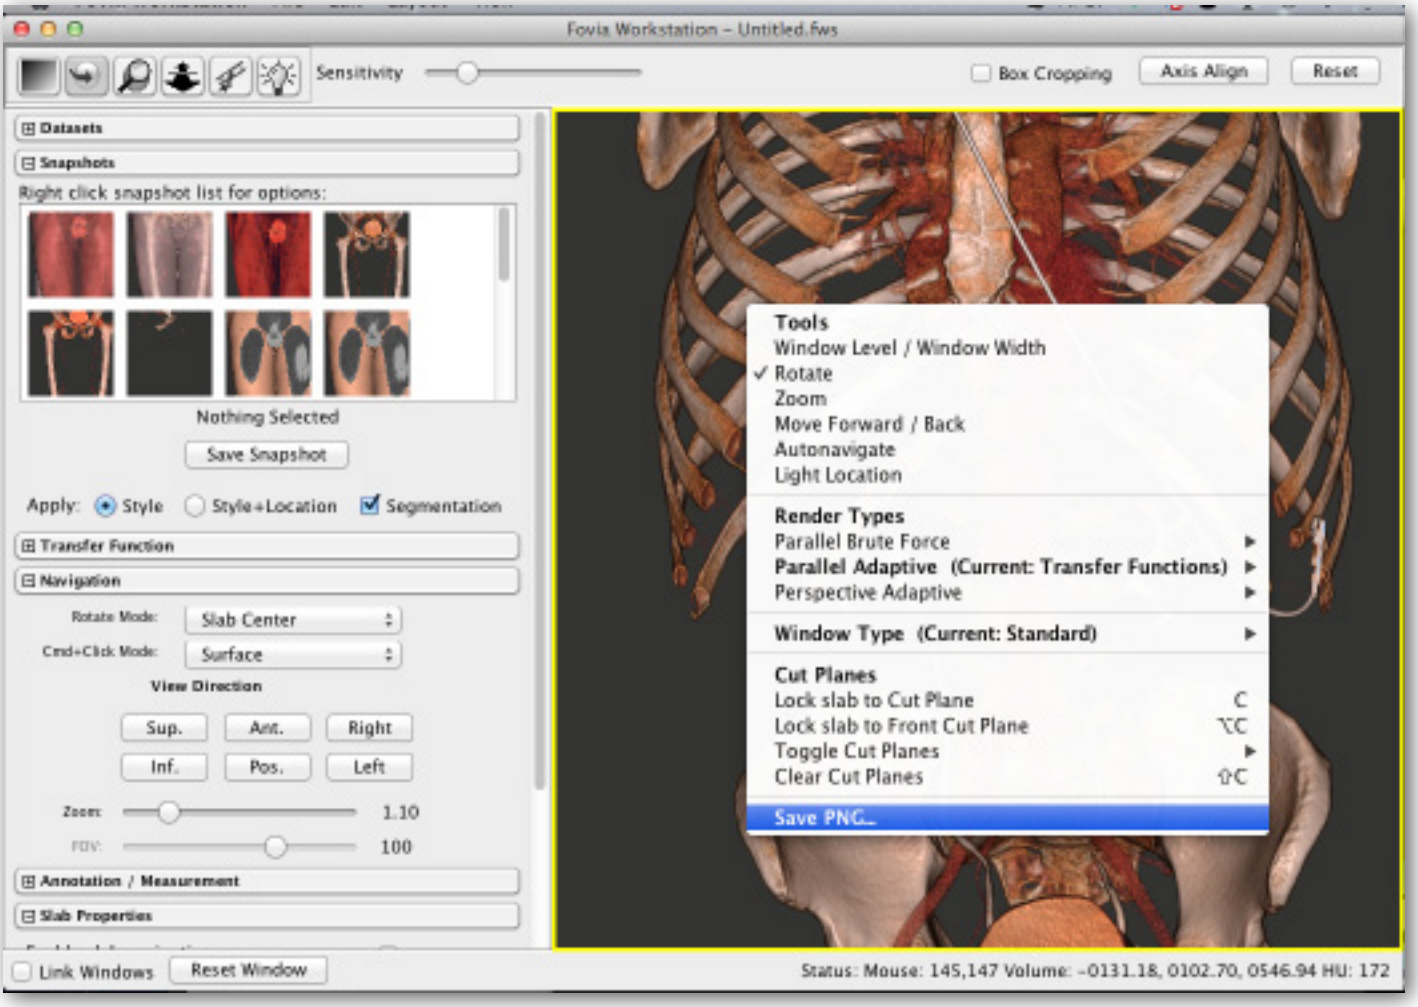

FIGURE 10.1: Save PNG

10.2 CREATING MEDIA - MOVIE EDITOR

INTRO TO MOVIE EDITOR

To open the movie editor, choose "Show Movie Editor Panel" from the "View" dropdown menu (Figure 10.2). This opens a simple Keyframe Movie Editor. This allows you to create and export QuickTime movies, as well as create editable Fovia Movie Files (.fmv).

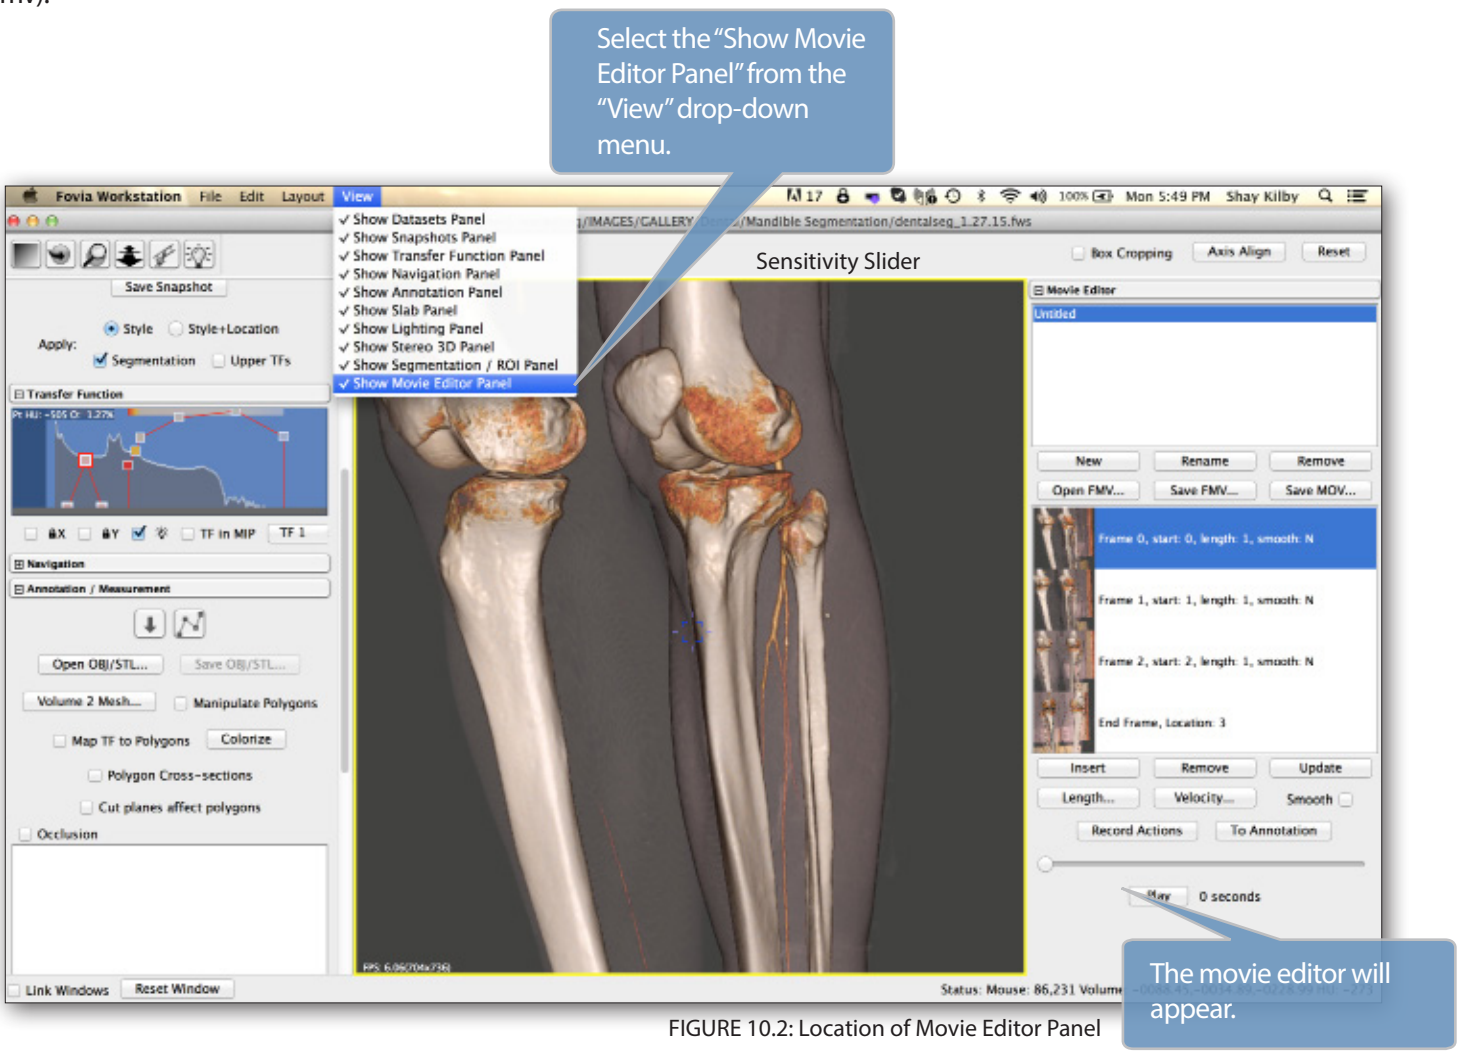

FIGURE 10.2: Location of Movie Editor Panel

MOVIE FILE MANAGER

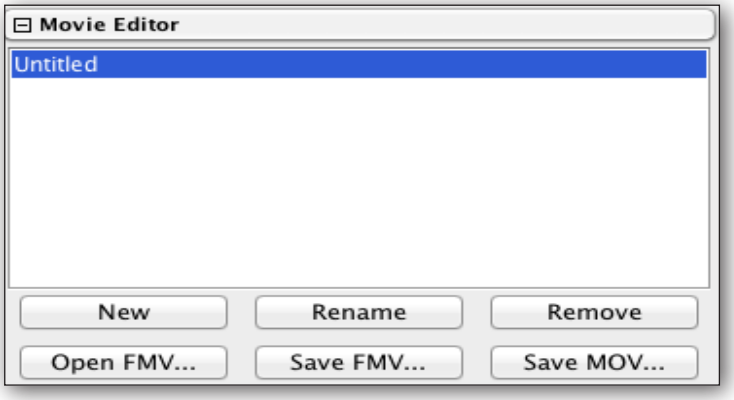

FIGURE 10.3 Naming movie projects

File Manager: To begin a new movie, create a new movie file and rename it to correspond with your current Fovia project. To continue work on an existing movie file, activate the "Open FMV" button and select the appropriate file to import. Once finished with your work session save out a .fmv file for future editing. When your movie is complete or ready to test use the Save MOVIE pop-up box to choose export options (see Section 10.3).

KEYFRAME EDITING

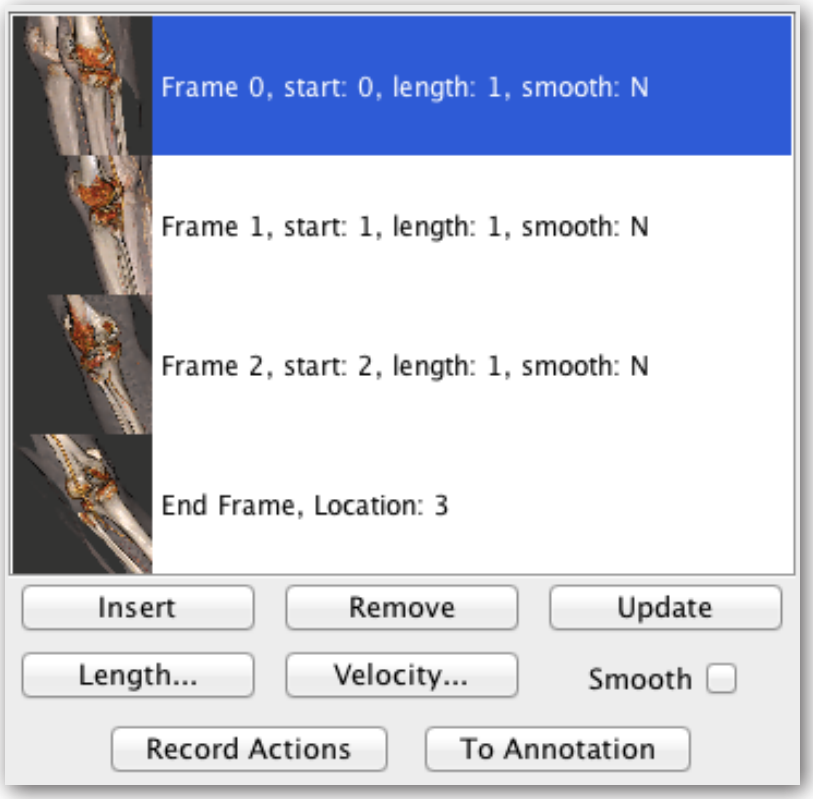

FIGURE 10.4: Keyframe Editor

The Keyframe Movie Editor is a simple editor that can help to produce a variety of effects with a few simple clicks. To begin, select a start frame in your main viewing window, then click the "insert" button and a keyframe will appear. Manipulate the data via location, segmentation, and TF edits and insert subsequent keyframes as desired.

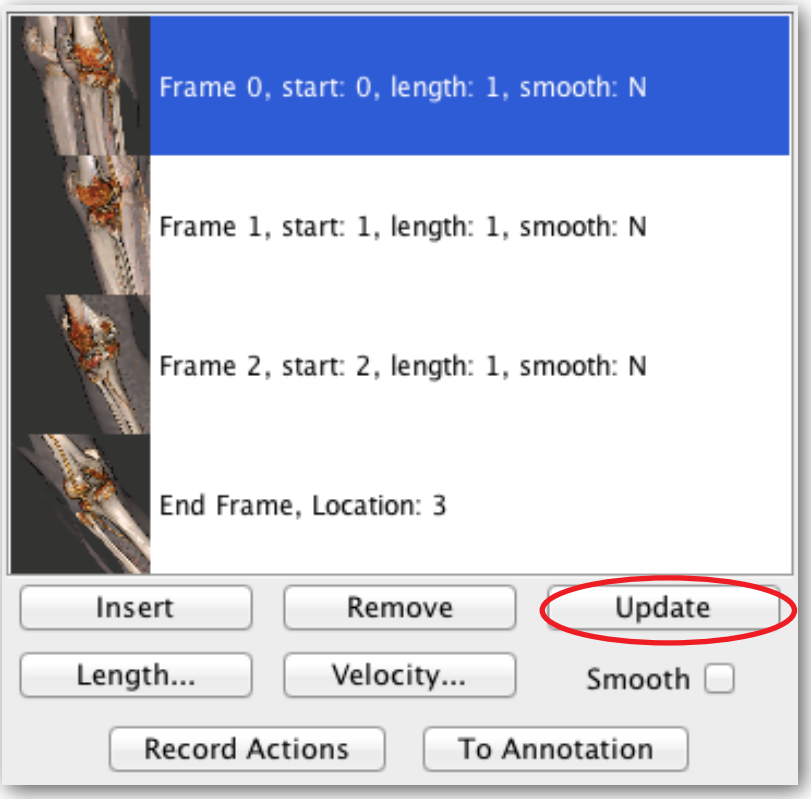

FIGURE 10.5: Updating keyframe

Keyframes can be saved to a .fmv file for editing, and are completely editable on-the-fly. The user can insert keyframes between existing keyframes, and can remove any undesireable keyframes. Additionally, existing keyframes may be relocated and updated in the transfer function (Figure 10.5). To change length and velocity, click on the corresponding buttons.

In the example below, the middle keyframe is updated to reflect a change in the TF. The transparency of the trachea has been decreased to create a more opaque look, by changing the vertical position of the Control Point corresponding to the tissue.

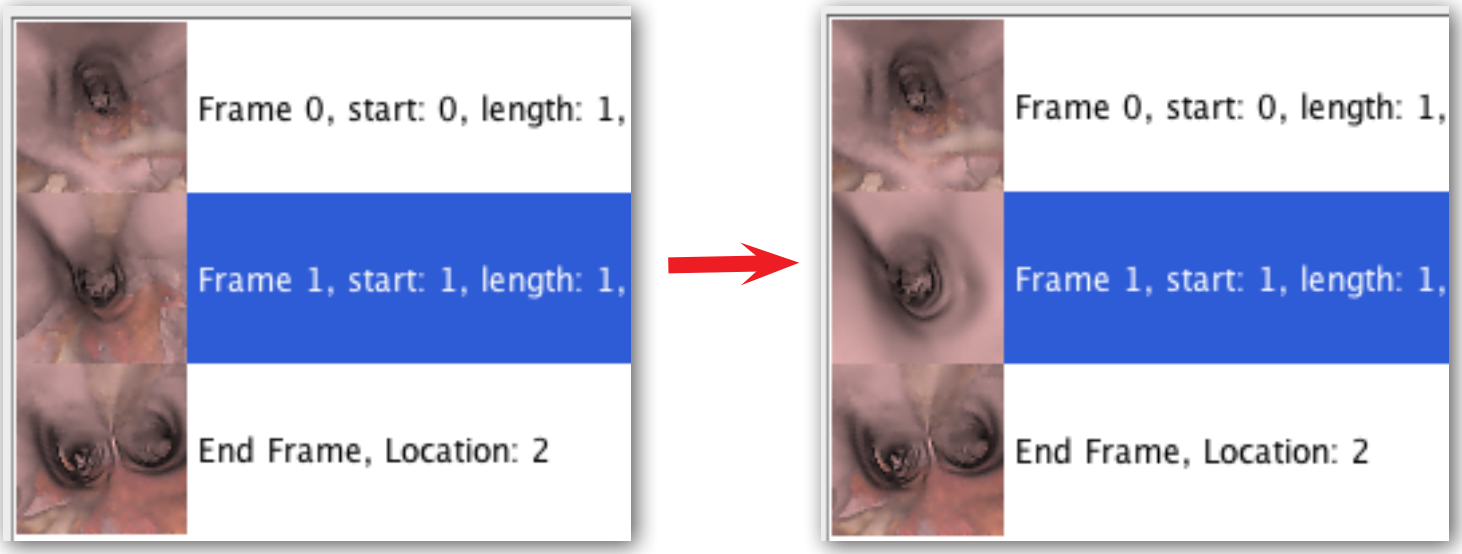

Figure 10.6: Adjusting transparency for a keyframe

As mentioned, to change the opacity of tissue for a keyframe, simply change the TF to the desired transparency and press the "Update" button in the movie editor. Below is a larger representation of what changes in the keyframe. (Figure10.7)

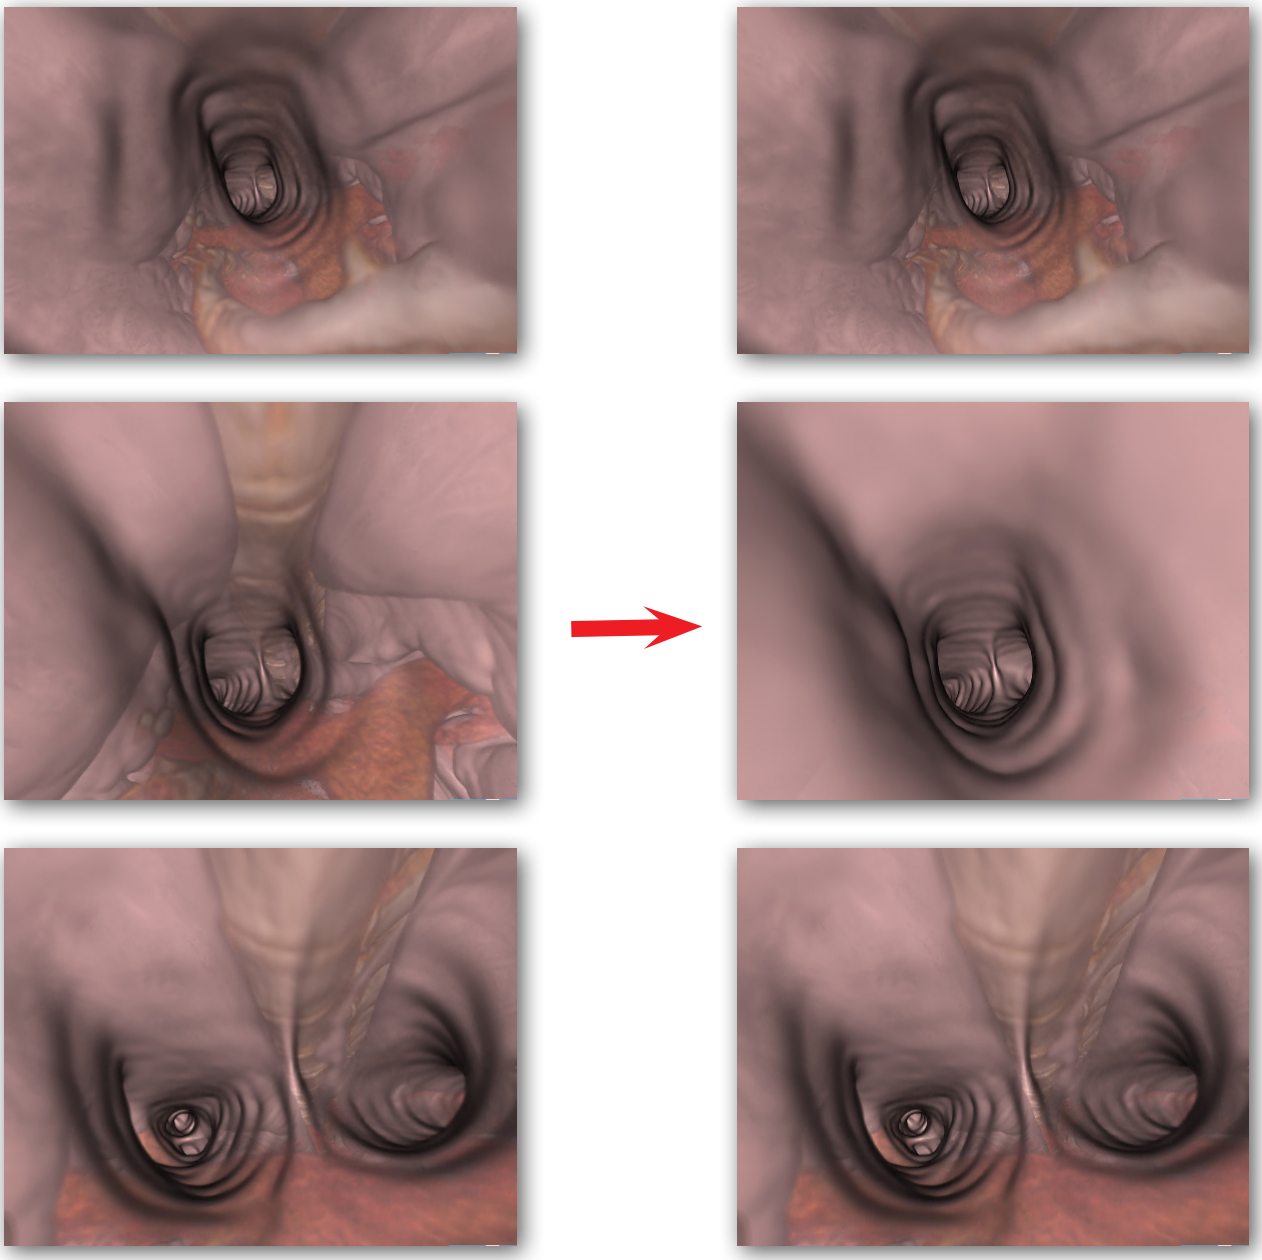

Figure 10.7: Adjusting transparency for a keyframe

The engine will automatically interpolate the change and make adjustments in themovie. In this case, the result will be a fly-through that has alternating levels of transparency with the middle part of the movie, exhibiting a higher transparency in the trachea.

10.3 CREATING MEDIA - MOVIE EXPORT

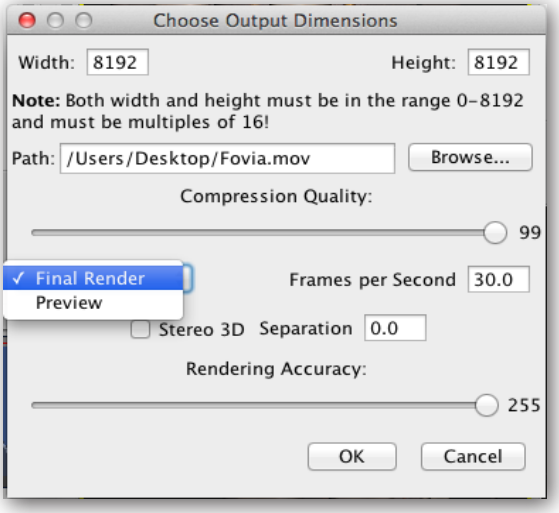

FIGURE 10.8: Saving QuickTime movies

Movie Output: The movie editor can be used to preview movies and final renders. Movies can be rendered in resolutions up to 8192 x 8192 pixels. When selecting the dimensions, keep in mind the capacity of the movie player you intend to use. Select the preview setting and render a test movie first in order to double check your movie before a final render (Figure 10.8). Once ready for a final rendering, select appropriate Compression Quality and Rendering Accuracy. (We recommend 99 and 255 respectively.)

XStream HDVR provides lighting tools with a myriad of setting options. These can be used to improve the illumination of specific surface properties of the dataset, as well as providing aesthetic appeal.

11.1 LIGHTING - INTRODUCTION

Lighting settings allow users to enhance certain surface properties. Illuminating certain tissue can accentuate noise in some images. *XStream HDVR* has methods to reduce (or alleviate) this type of lighting problem.

The software has two different types of lighting controls. The first affects the location of the lighting source, and is accessed through the navigation located at the top of the user interface (Figure 11.1). Alternatively, more options are available through the Lighting Tab, located on the left side of the user interface (Figure 11.1).

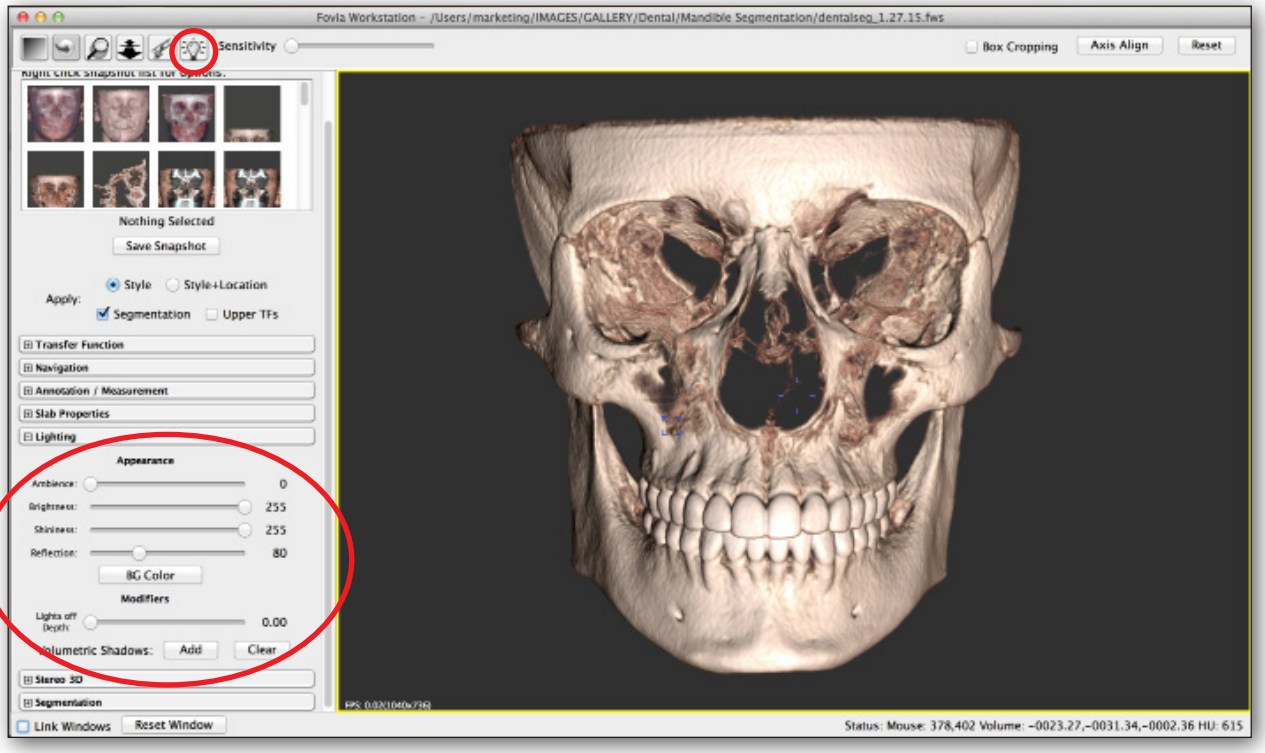

FIGURE 11.1: Location of the Lighting Navigation Tool and the Lighting Tab

11.2 LIGHTING - LIGHTING TOOL

The lighting navigation tool on the navigation bar allows the user to change the location of the light. The user can move the light source around over the volume. Additionally, the light source can be detached from the camera and left in a fixed place. This is accomplished by deslecting the "Attached to Camera" option located at the top of the screen (Figure 11.2). The volume can then be moved around.

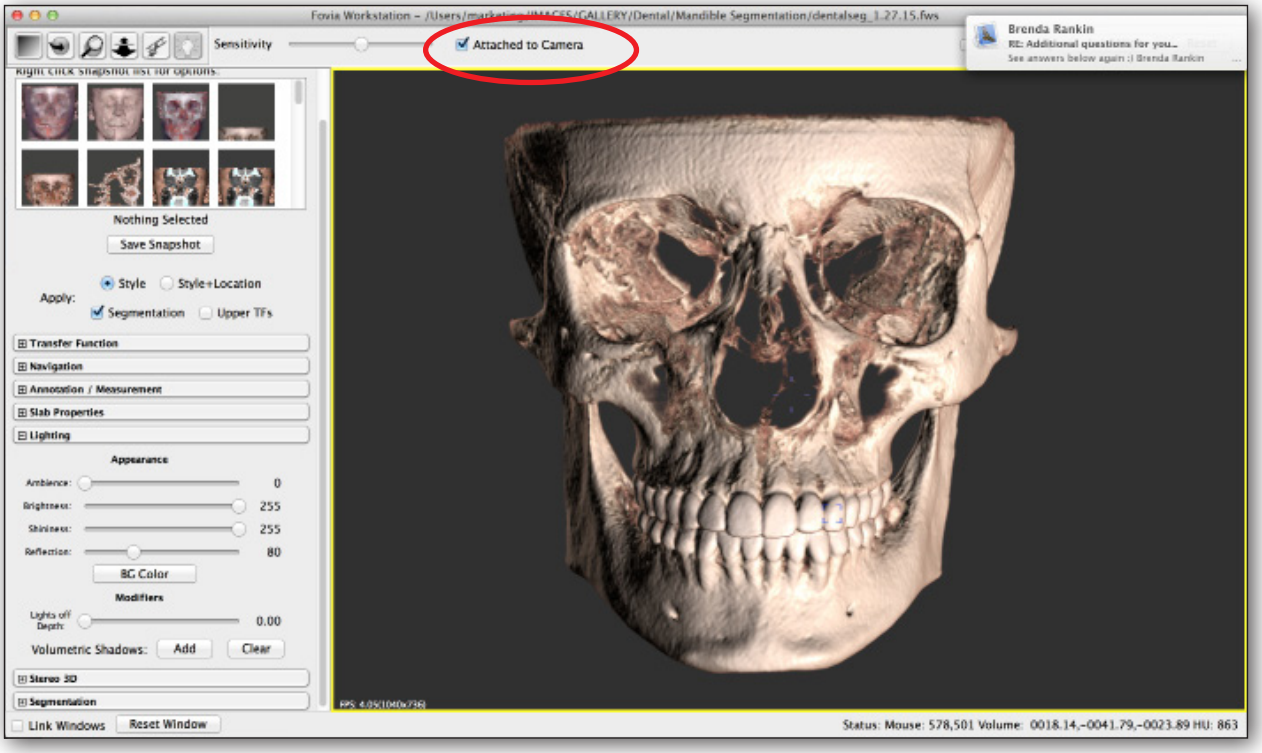

FIGURE 11.2: Deselecting "Attached to Camera"

11.3 LIGHTING - LIGHTING TAB

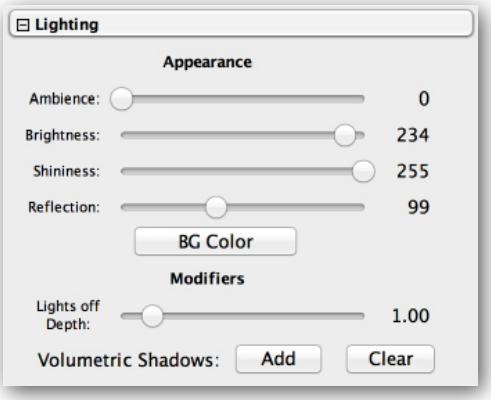

FIGURE 11.3: Lighting Tab details

The Lighting Tab provides the following additional options (Figure 11.3):

- Ambience: Adjusts the surrounding lighting of the dataset.
- Brightness: Adjusts the overall brightness of the light.
- Shininess: Determines diffuseness of specular reflection. Moving the slider to the right adds more shineness to the dataset.
- Reflection: Specular reflection. Moving the slider to the right will increase the reflection of the dataset.
- Background color change ("BG Color" button): Clicking this allows the user to change the color of the background.
- Lights off Depth: This slider controls what the light DOES NOT affect. Moving the slider to the right will increase the depth of the dataset that will be unaffected by the light.
- Volumetric Shadows: This fixes the light source onto the volume. The dataset will then casts a shadow. The shadow and light will stay static while the user rotates about the volume.

The ability to render in stereo is becoming more important due to its use in robotic-assisted surgery, as well as surgical simulation and education. *XStream HDVR* supports stereo rendering that can be displayed on a variety of displays, both passive and active. There is no restriction on its usage or type of hardware standard, since the two output images are simply generated for the left and right eyes.

12.1 STEREOSCOPY - VIEWING STEREO

The separation parameter defines how profoundly different the left and right eye images are. Think of the distance between the two eyes on the human body. If they were farther apart, one could see more radical differences between the two images, and thus perceive the 3D effect more precisely. (Figures 12.1 & 12.2).

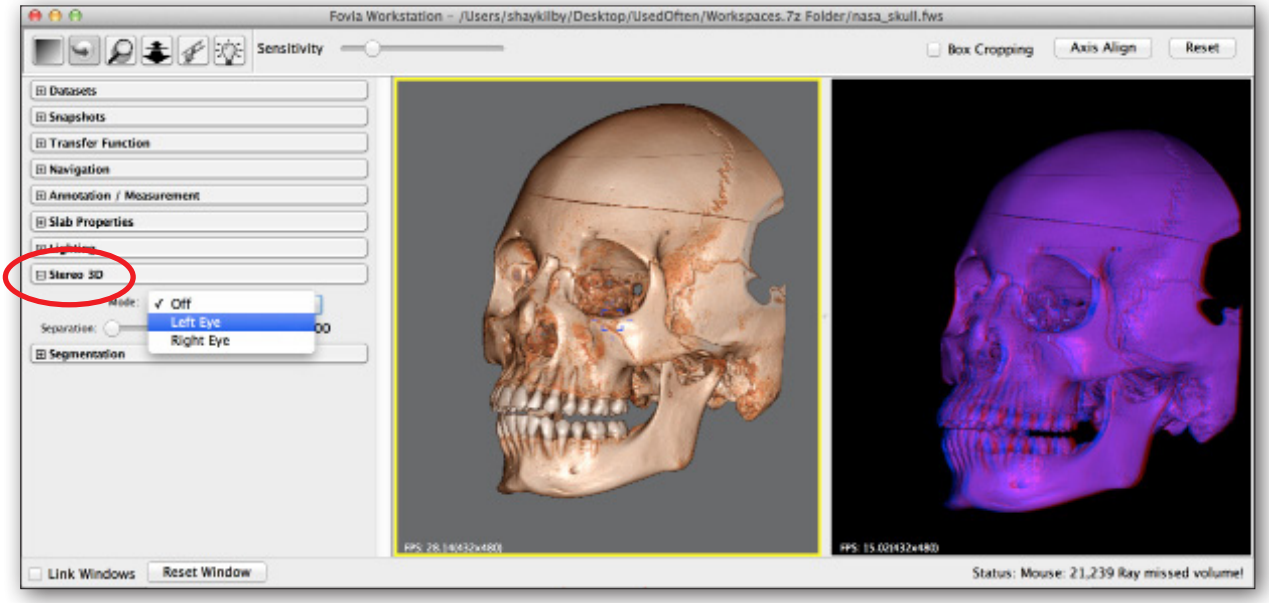

FIGURE 12.1: Stereo Window Tab

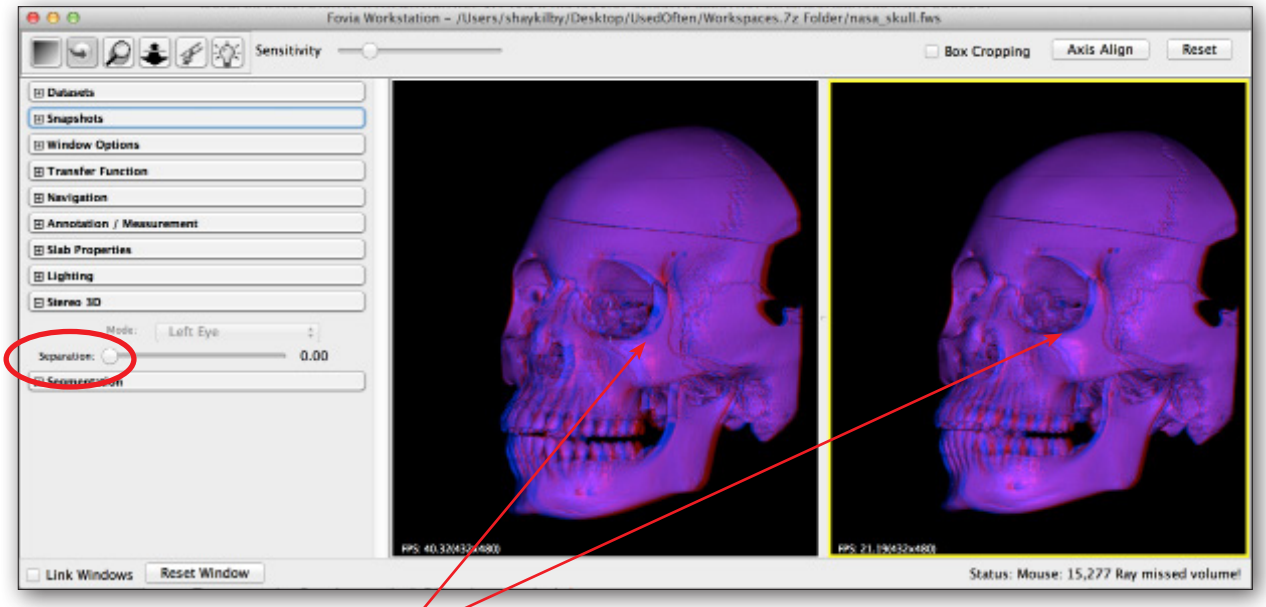

Left window has higher separation

FIGURE 12.2: Stereo Window Separation slider

To turn on the stereo window, right-click on your chosen viewing window, scroll to "Window Types" and select "Red/Blue Stereo" in the sub-drop down menu. This prepares your active workstation window for stereo viewing (Figures 12.3 and 12.4).

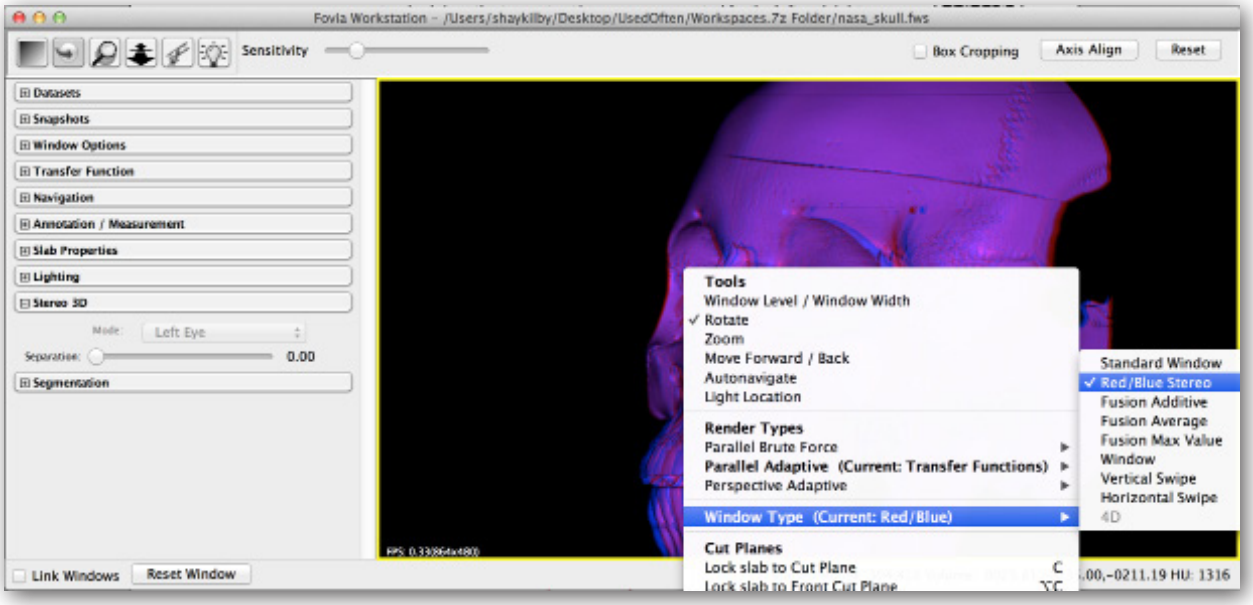

FIGURE 12.3: Stereo Window mode

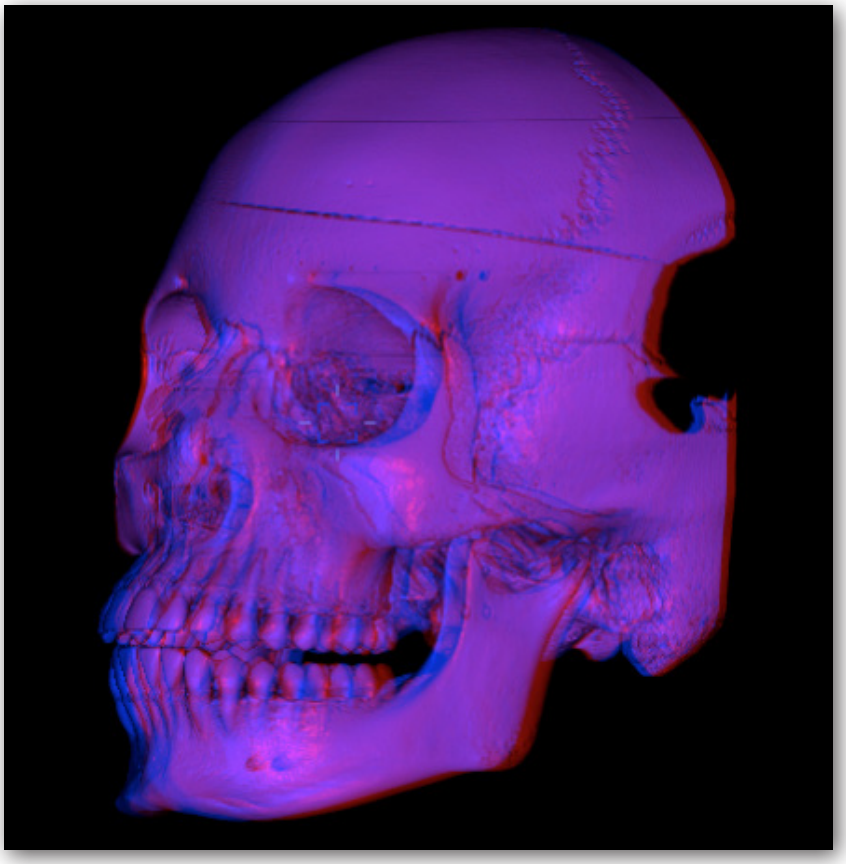

FIGURE 12.4: Stereo view

12.2 STEREOSCOPY - EXPORTING STEREO

To export stereoscopic movies in the Movie Editor window, select the checkbox "Stereo 3D" when saving a movie. Separation settings may also be selected in this dialogue box (Figure 12.5).

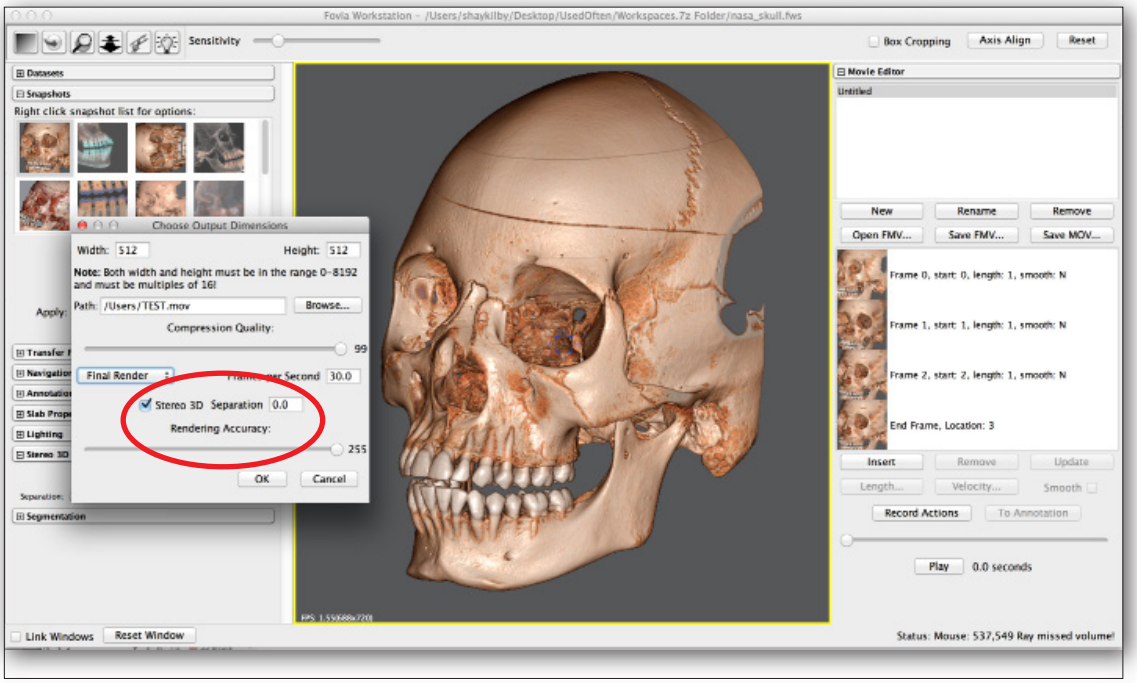

FIGURE 12.5: Stereo movies

For more about the Movie Editor please see the Creating Media Chapter (p. 62).

TABLE OF CONTENTS

MANUAL

FOViA

TABLE OF CONTENTS

MANUAL

FOViA


FOViA